Preparations for detecting and characterizing gravitational-wave signals from binary black hole coalescences

Thilina Dayanga[1, *] and Sukanta Bose[1, 2, †]

^{1}*Department of Physics & Astronomy, Washington State University, 1245 Webster, Pullman, WA 99164-2814, U.S.A.*
^{2}*Inter-University Centre for Astronomy and Astrophysics, Post Bag 4, Ganeshkhind, Pune 411 007, India*

We evaluate how well "EOBNR" waveforms, obtained from the effective one-body formalism, perform in detecting gravitational wave (GW) signals from binary black hole (BBH) coalescences modelled by numerical relativity (NR) groups participating in the second edition of the numerical injection analysis (NINJA-2). In this study, NINJA-2 NR-based signals that are available in the public domain were injected in simulated Gaussian, stationary data prepared for three LIGO-Virgo detectors with early Advanced LIGO sensitivities. Here we studied only non-spinning BBH signals. A total of 2000 such signals from 20 NR-based signal families were injected in a two-month long data set. The all-sky, all-time compact binary coalescence (CBC) search pipeline was run along with an added coherent stage to search for those signals. We find that the EOBNR templates are only slightly less efficient (by a few percent) in detecting non-spinning NR-based signals than in detecting EOBNR injections. On the other hand, the coherent stage improves the signal detectability by a few percent over a coincident search. In regards to signal parameters, such as the binary component masses and signal end-time, the magnitude and nature of the systematic errors in their measurement show some interesting but limited variations. In particular, a very small fraction of signals are systematically detected with templates of slightly more massive systems and, therefore, have a measured end-time that is earlier than the true one. The same signals have a worse match than other signals with EOBNR templates of the same parameters. We compare these observations with the results of a study where EOBNR templates were used to find EOBNR signal injections to account for any biases that might arise from the data analysis pipeline itself. We discuss how our results can be utilized by various source modelling groups and data analysis pipelines to explore ways of improving the detectability of BBH signals.

PACS numbers: 04.30.Tv,04.30.-w,04.80.Nn,97.60.Lf

I. INTRODUCTION

Inspiraling binary black holes (BBHs) are one of the most promising gravitational-wave sources that the second generation ground-based detectors, such as the Advanced Laser Interferometer Gravitational-wave Observatory (aLIGO) [1] and Advanced Virgo (AdV) [2] detectors, are likely to detect. Currently these laser interferometric detectors are being upgraded and will start collecting data in a few years' time with a sensitivity improvement of about an order of magnitude. The new detector KAGRA [3] is also expected to take data with a similar design sensitivity later this decade. These detectors will define what is being termed as the Advanced Detector Era (ADE). Second generation instruments will have sensitivity in a broader frequency band compared to the initial detectors and will observe gravitational waves (GWs) from compact binary coalescence (CBC) signals starting at a lower frequency. Observations of X-ray binaries IC10 X-1 [4] and NGC 300 X-1 [5] indicate that the masses of the stellar-mass components of a BBH can be as high as $20 - 30M_\odot$. The discovery of HLX-1 in ESO 243-49 that has a lower mass limit of approximately 500 $M_\odot$ presents strong evidence for the existence of intermediate-mass black holes [6]. As described in Refs. [7] and [8], binary black holes with component masses that high or higher will be detectable in the ADE detectors only through the merger and ringdown signals. This paper presents a search for BBH systems with total mass $25M_\odot \leq M \leq 100M_\odot$ and component masses $3M_\odot \leq m_1, m_2 \leq 97M_\odot$. Searches in real LIGO-Virgo data for BBHs in the same mass range were conducted with Inspiral-Merger-Ringdown (IMR) templates in Refs. [7, 8].

The ability to detect GW signals arising from BBH coalescences crucially depends on the accuracy of the waveform models used in designing search templates for detection pipelines. Estimation of parameters of BBH signals also demands accurate knowledge of the inspiral, merger and ringdown phases of the waveforms [14]. In 2008, the Numerical Injection Analysis (NINJA) collaboration was formed to facilitate the interaction of the Numerical Relativity and the GW Data Analysis communities with the objective of modelling CBC signals and using them to perfect GW search pipelines and parameter estimation algorithms. The main purpose of the first NINJA project (NINJA-1) [15] was to foster the exchange of numerical-relativity waveforms to evaluate the performance of a variety of data analysis pipelines in detecting them in simulated initial LIGO and Virgo detector noise. The NINJA-2 exercise, on the other hand, creates the opportunity to test and compare detection pipelines more meaningfully than in NINJA-1 owing to the strict requirements imposed on the accuracy and length of the NR-based waveforms employed by it.

*Electronic address: wdayanga@wsu.edu
†Electronic address: sukanta@wsu.edu
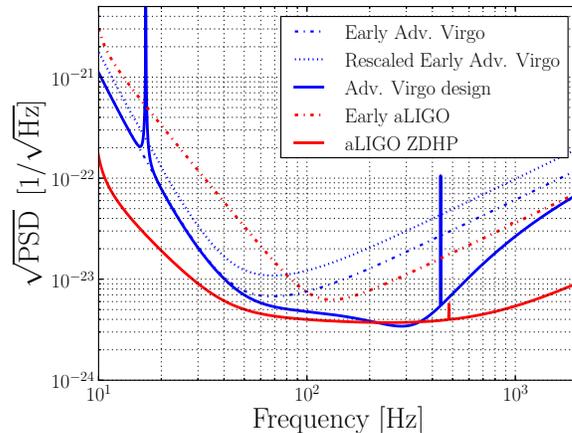
FIG. 1: Comparison of early and final Advanced LIGO (aLIGO) and Advanced Virgo (Adv. Virgo) design amplitude spectral densities (ASDs) [9]. The ASD is the square-root of the power spectral density (PSD) [10]. The red dotted and solid lines represent early and zero-detuned high-power (ZDHP) aLIGO design ASDs, respectively [11]. The solid blue curve shows the Adv. Virgo design ASD [12]. The blue dotted curve is obtained by rescaling it so that its horizon distance [13] is similar to that corresponding to the early aLIGO ASD.

The NINJA-2 project required each BBH waveform to include at least five orbits of usable data before merger, i.e. neglecting the initial burst of junk radiation. Additionally, NR waveform amplitude had to be accurate to within 5of gravitational-wave frequency) should have an accumulated uncertainty over the entire inspiral, merger and ringdown (of the numerical simulation) of no more than 0.5 rad. Following these restrictions eight NR groups have contributed 56 waveforms to NINJA-2 project [16]. Since the typical signal will have more cycles in band than what NR alone can produce at the desired accuracy with the computational resources at hand, post-Newtonian theory is used to model the remaining cycles to produce long *hybrid* signals, which we will refer to here as pN-NR hybrid signals or simply NR or NR-based signals. The NINJA-2 waveform-catalog paper [16] describes in detail the waveform requirements and method used to construct the pN-NR hybrid waveforms.

In this paper, we quantify how well "EOBNR" waveforms, obtained from the effective one-body formalism [17] by calibrating against a specific family of NR waveforms, namely, the ones obtained by the NASA-Goddard group [17], perform in detecting pN-NR BBH signals modelled for NINJA-2. In NINJA-2, a variety of NR-based signals were injected in simulated Gaussian, stationary data from three LIGO-Virgo detectors with early ADE sensitivities (see Fig. 1). These signals were constructed with contributions from various numerical relativity groups (see Table 1 in Ref. [16]), and are available in the public domain [18]. For the study reported here we focus on non-spinning BBH signals. A total of 2000 such signals from 20 pN-NR signal families were injected in a two-month long data set [18]. The coincident all-sky, all-time compact binary coalescence (CBC) search pipeline [19] was run along with an added coherent stage [20] to search for those signals. We also compare these observations with the results of a study where EOBNR templates were used to find EOBNR signal injections to account for any biases that might arise from the data analysis pipeline itself. We find that the EOBNR templates are slightly less efficient, by about a percent, in detecting non-spinning NR-based signals than in detecting EOBNR injections. Also, the magnitude and nature of the systematic error in the measurement of signal parameters show some interesting but limited variations. In particular, a very small fraction of signals are systematically detected with more massive templates and, therefore, have a measured end-time that is earlier than the true one. The same signals have a worse match than other signals with EOBNR templates of the same parameters [16]. We show how a coherent all-sky, all-time search can improve CBC detection efficiencies over a coincident analysis by improving the detection probability for any given false-alarm probability. The coherent method requires coincidence of the signal arrival times and other signal parameters, just like a coincident search method [19] does. But the former also checks for the consistency of the signal phases and amplitudes in the various detectors in a network with a physical value for the signal time-delays across the detector baselines. It also provides multi-baseline signal-based tests like the null-stream test that are effective in discriminating real signals from noise artifacts, especially, when the sensitivities of the detectors are comparable [20–23].

Owing to the high computational cost of employing a fully coherent all-sky, all-time CBC search [24], here we use the *hierarchical* coherent algorithm [20] to search for BBH signals. This algorithm and the detection statistic are described in Sec. II. Coherent data analysis methods have been formulated for searching GW signals from a variety of modelled and unmodelled sources [20, 23–27]. It is the optimal method in stationary, Gaussian noise under the Neyman-Pearson criterion [10]. In CBC searches, it has been used recently in targeted searches [23] where the time of occurrence of the signal and the sky-position of the source are known, e.g., from the observation of an electromagnetic (EM) counterpart, such as a gamma-ray burst (GRB). It was demonstrated to perform better than the targeted all-sky, all-time coincident search.

In Sec. III, we compare the performances of the coincident and the hierarchical coherent searches. We find the

latter to be somewhat better than the former and explain what factors contribute to this improvement. We argue that the level of improvement is as per expectations for a hierarchical method and that a fully coherent method should be able to yield a much better detection efficiency. We present results from multiple sanity tests that check if the signal injection recovery is consistent with our expectations of a BBH search. Moreover, the effect of signal-based discriminatory tests, such as the chi-square and the null-stream tests, on the performance of that search is also analyzed.

In Sec. IV, the accuracy with which various signal parameters are recovered is described. Those results are found to be mostly devoid of systematic errors when compared with measurements of EOBNR injection parameters by using EOBNR templates. The very few cases where a bias was found, the explanation lies in the mismatch of the pN-NR hybrid signal in those cases with the EOBNR templates, as was seen in Ref. [16].

II. COINCIDENT AND COHERENT SEARCHES FOR BINARY BLACK HOLE SYSTEMS

A. Search algorithm

The all-sky, all-time coincident search pipeline that was used to detect injected pN-NR hybrid signals is described in detail in Ref. [19]. On the other hand, the hierarchical coherent search pipeline used here was introduced in Ref. [20]. The latter pipeline consists of the former with an additional stage that computes the coherent SNR of all coincident triggers found by the former.[1] We will variously refer to these two stages as the coincident (or the first) and the coherent (or the second) stage, respectively. Both these stages comprise multiple steps. The coincident stage first splits the strain data time-series from every detector in the network into 2048 sec chunks. The noise power spectral density (PSD) [10] is estimated for each chunk and is used to construct a template bank for matched filtering. Whenever the signal-to-noise (SNR) of the filtered output crosses a preset threshold, which was chosen to be 5.5 in Refs. [7, 8] and in this work, the template parameters and the time of the trigger are saved for each detector. Next the triggers from individual detectors are compared for coincidences in mass and end-time in two or more detectors to identify multi-detector coincident GW candidate events. Triggers in a detector that do not find any coincidence with a trigger in another detector are dropped from further analysis since currently we do not have a method for assessing the noise background for single detector events. A candidate event is termed as double-coincident (triple-coincident) if the masses and end-times of the triggers in two (three) detectors are found to be very similar [28], as stipulated in advance of the searches. Since the maximum number of detectors available to this search is three, these are the only two types of coincident candidate events possible here. In a real search, a candidate event is subjected to further checks before it is announced as a GW event to ensure that it was not caused by an environmental or instrumental artifact.

The coherent stage requires coincident trigger times in order to begin the coherent analysis. This analysis includes multiple steps, which are similar to the coincident counterpart but include some important modifications. First, a template bank, termed as the coherent bank, is constructed for the coherent analysis using the parameters of the coincident triggers. The triggers identified by the coincident stage can have different mass parameters in different individual detectors due to the different noise in each detector. By definition, however, the same template must be used in every detector for computing the coherent SNR. Here, we choose the mass-template in the loudest detector for that computation. Reference [20] explains this process in detail. That common template is next used to compute the matched-filter output, which is in the form of an amplitude and phase time-series, for every detector in the coincidence. That step also computes the template normalization factors and implements signal-based vetoes.

Similar to the matched-filtering step in the coincident stage, the coherent matched-filtering step used a thresholding criterion in previous studies [20] to reduce the computational cost. This threshold value is set for the individual detectors as in the coincident stage but the value is lowered to allow more triple coincident triggers compared to the coincident counterpart. Typically, the values of previous studies used to be 5.5 for the coincident matched filtering step and 5.0 for the coherent one. One of the main goals of the coherent analysis is to convert all double-coincident triggers in the coincident stage to triples, whenever data are available from all three detectors, to improve the significance of real signals. Although a lower threshold in the coherent stage on individual detector SNRs helps to convert more double coincident events into triples many more remain as doubles when that threshold is non-zero. To improve upon earlier studies, we devised a method to handle these additional background triggers more efficiently, essentially, by reducing the maximum coincidence duration analyzed at a time. Lowering that threshold to zero allows converting all triggers that were double-coincident in the first stage to triple-coincident ones [2]. This provides additional information about every single trigger due to the phase consistency check we can impose on the signal in every participating detector. The matched-filtering step in the coherent stage is followed by the computation of the coherent SNR and the null stream for every trigger. Since we use large number of templates in our template bank,

---

[1] That stage also computes the null stream [21, 23], which will be described in Sec. III B.

[2] Results from the new pipeline show that a few signals that show up as doubles in the coincident stage are still not converted into triples in the coherent stage. This is because those (simulated) signals are very weak in the third detector, thereby, resulting in unphysical time-delays between that and the other two detectors.
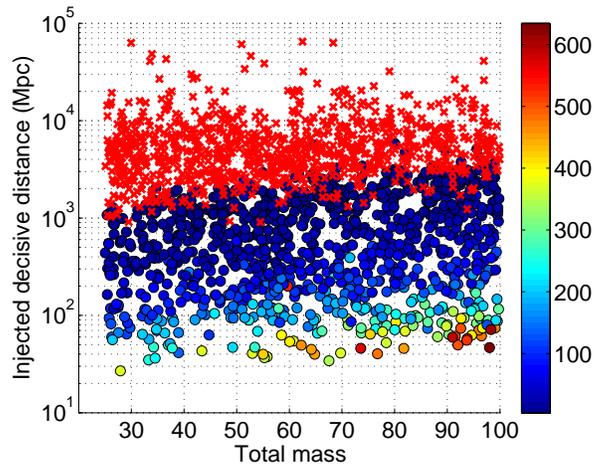
FIG. 2: The injected decisive distance of found and missed pN-NR simulated signals plotted as a function of total mass of the binary black hole system. The colorbar shows the coherent SNR. A total of 1033 injections were found out of 2000 injections by the search pipeline. Red crosses represent missed simulated signals.

same feature in detector data can be picked by different template waveforms. This leads to have multiple GW triggers at same end-time. Additionally, for coherent searches multiple sky positions can give triggers for same feature in data. These set of triggers is also know as clusters. The last step of this pipeline clusters the gravitational-wave triggers in time and sky position so that only the most significant of them is retained per cluster.

B. Detection Statistics

The BBH search algorithms exploit the knowledge of their GW signals to define the templates used for matched filtering. Nine parameters describe the GW signal from nonspinning BBH sources studied here. These are the two component masses m_1 and m_2, the luminosity distance to the source d, the right ascension and declination angles (α, δ) specifying its sky position, its orbital inclination angle ι to the line-of-sight, the angle ψ describing the orientation of its signal polarization ellipse, the signal coalescence time t_c, and the signal coalescence phase ϕ_c. Alternative mass parameters, in the form of the total mass M, the symmetrized mass-ratio $\eta \equiv m_1 m_2/M^2$, and the chirp mass $M_{\text{chirp}} \equiv \eta^{3/5} M$ are also often used to describe BBH signal parameters.

For the coincident analysis, the detection statistic used here is the one that was introduced in Ref. [8] for the *high-mass* search in LIGO-Virgo data from the sixth LIGO Science Run (S6) and the second and third Virgo Science Runs (VSR2/3). That search targeted GW signals from BBH sources with each component mass between $3 - 97 M_\odot$ and total binary mass between $25 - 100 M_\odot$. If the matched-filter output of a unit-norm mass template from a stretch of single-detector data is denoted by ρ, then the detection statistic used for long-duration templates, with a duration $t_{\text{dur}} \geq 0.2$ sec, is the *new* SNR defined as follows:

$$\rho_{\text{new}} = \begin{cases} \dfrac{\rho}{[(1 + (\chi_r^2)^3)/2]^{1/6}} & \text{for } \chi_r^2 > 1, \\ \rho & \text{for } \chi_r^2 \leq 1, \end{cases} \qquad (2.1)$$

where $\chi_r^2 \equiv \chi^2/(2p - 2)$, and χ^2 is the signal-based time-frequency discriminator studied in Ref. [29] with a chi-squared distribution of p degrees of freedom. On the other hand, for short-duration templates, with a duration smaller than 0.2 sec, the detection statistic used is

$$\rho_{\text{eff}} = \frac{\rho}{[\chi_r^2(1 + \rho^2/50)]^{1/4}}, \qquad (2.2)$$

which is termed as the *effective* SNR. These choices were arrived at after comparing their detection efficiencies with those of ρ and other alternative statistics. In a network with M detectors, the coincident detection statistic employed here is

$$\rho_{\text{coinc}} = \begin{cases} \left[\displaystyle\sum_{I=1}^{M} \left(\rho_{\text{eff}}^I\right)^2\right]^{1/2} & \text{if } t_{\text{dur}}^I < 0.2\text{sec, for any } I, \\ \left[\displaystyle\sum_{I=1}^{M} \left(\rho_{\text{new}}^I\right)^2\right]^{1/2} & \text{for all other cases}, \end{cases} \qquad (2.3)$$

where x^I denotes the value of x for the Ith detector. Employing χ^2 in the detection statistic was found to improve the performance of a search, especially, in real data, which is neither Gaussian nor stationary. Contrastingly, a network detection statistic formed from only the ρ^I is the *combined* SNR, which is just $\left[\sum_{I=1}^{M}\left(\rho^I\right)^2\right]^{1/2}$.

The coherent statistic used here is the same as the one defined in Eq. (2.30) in Ref. [20].

When describing how effective a search is in finding signals, one of the parameters used is how distant their sources are. However, from the measured strength of a signal, it is not always possible to deduce the source luminosity distance. This is because the strength of a signal in a detector is determined not only by the proximity of a source but also its location, its orbital inclination angle to the line-of-sight, and the angle ψ describing the orientation of its polarization ellipse. In fact, a source at a luminosity distance d appears to be at an *effective* distance of

$$d_{\text{eff}} = \frac{d}{\sqrt{F_+^2(\alpha,\delta,\psi)\left(1+\cos^2\iota\right)^2/4 + F_\times^2(\alpha,\delta,\psi)\cos^2\iota}}, \tag{2.4}$$

where $F_{+,\times}$ are the detector antenna-pattern functions. Due to the different orientations and, therefore, $F_{+,\times}$, the effective distance of a source can vary from one detector to another. Since for a detection we require coincidence in at least two detectors, it is essential that the larger of the two corresponding effective distances not be too large for the signal to fall below the detection threshold of the weaker detector. This is why it is useful to define the *injected decisive* distance as the injected effective distance in the second loudest detector in a coincidence. Indeed, in Fig. 2 we show how this quantity and the coherent SNR vary as a function of the total-mass of the injected sources. It is manifest that the closer sources are found with a higher coherent SNR whereas many of the very distant sources are missed.

III. DETECTING PN-NR HYBRID SIGNALS WITH EOBNR TEMPLATES

In this study we analyzed the performance of the hierarchical coherent pipeline in simulated Gaussian noise. Strain noise time-series were produced for two LIGO detectors, H1 and L1, and the Virgo detector, V1, using the design noise curves of early advanced LIGO and advanced Virgo detectors [30], respectively. Each time-series is continuous and is of two months' duration. Performance of the search pipelines was evaluated by injecting pN-NR waveforms in the simulated data and recovering them with EOBNR templates [31]. As with the LIGO-Virgo highmass search in S6-VSR2/3 data, here too the template bank used for matched-filtering spans $3 - 97M_\odot$ for each of the two component masses and $25 - 100M_\odot$ for the total mass. The noise curves used to simulate noise in this study is the same as the one used for NINJA-2 blind injection challenge [32]. Unlike in blind injection studies, it is Gaussian and stationary and has early aLIGO noise amplitude spectral density (ASD) for H1 and L1 and rescaled early advanced Virgo ASD for V1, as depicted in Fig. 1. As the name suggests, the latter is obtained by rescaling early AdV ASD to match the noise expected in V1 when H1 and L1 will be taking early aLIGO data in the future. Rescaling Virgo noise curve creates the opportunity to have three equally sensitive detectors and it benefits parameter estimation studies and multi-detector signal based vetoes in coherent analysis. It is important to note that these curves are not the 2015-predicted curves, but a best guess at them from over a year ago when NINJA-2 data sets were created. In the past, real data from the LIGO and Virgo detectors had more dissimilar ASDs than the one used in NINJA-2. This aspect of our search assumes special significance since we later study the performance of the null stream, which is more powerful for detectors in a network with similar rather than disparate sensitivities.

Since the hierarchical coherent pipeline is not a fully coherent all-sky search, it depends on some other search method to identify the GW event times. In our study we analyze GW events identified by the coincident CBC pipeline. The coherent pipeline takes GW event times found by the latter as its input and does a coherent analysis of the individual detector data from around those trigger times. Compared to a fully coherent all-sky search, which combines complete time series coherently to identify GW events, the hierarchical coherent search requires less computational power. One of the main goals of this paper is to draw lessons from the hierarchical coherent search that can benefit the development of a fully coherent search for the ADE.

One major difference between the hierarchical coherent analysis used in this paper and those employed in the past [20] is that after the coincidence stage identifies a part of the data that offers an interesting trigger, its coherent SNR is computed by involving every detector that was active at the time, even if it did not contribute to the coincidence in that stage (e.g., because the SNR in the third detector was below threshold around that event time). Previous searches used thresholds for individual matched-filter outputs to reduce the computational cost. Typical value for this SNR threshold was 5.0, which was chosen to be somewhat below the threshold in the first stage, of 5.5, chosen for the matched filtering step. By placing an SNR threshold of 5.0 the coherent pipeline allowed the conversion of some double coincident triggers (namely, H1L1, L1V1 and H1V1) to triple-coincident H1L1V1 triggers at times when all three detectors had science data, often termed as *triple time*. In this study we lower the individual detector thresholds in the coherent stage to zero. This allows all double and triple coincident triggers, in tripe time, in the coincident stage to be analyzed as triple-coincident triggers in the coherent stage. By doing so for double coincident triggers, the coherent stage uses new information available from the third detector data for constructing the coherent statistic that was not used in constructing the coincident statistic in the first stage (apart from the fact that the
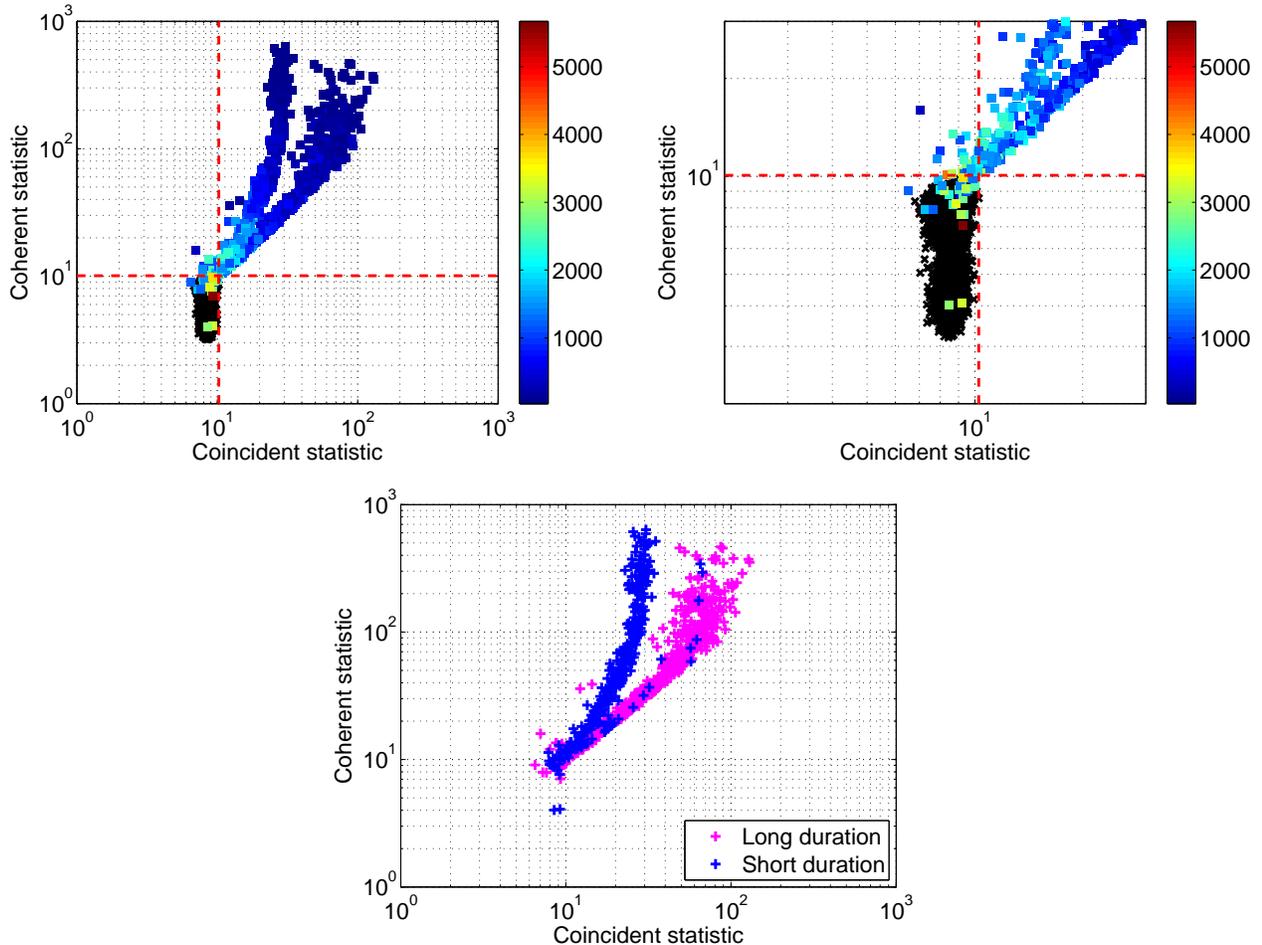
FIG. 3: **Top left:** The coherent and coincident detection statistic values of found simulated signals (colored squares) and background events (black crosses). The colorbars show the injected decisive distance in Mpc. The coherent SNR is used as the coherent detection statistic. The coincident detection statistic is ρ_{coinc}, as defined in Eq. (2.3), which is the one that was used in the high-mass search of LIGO's S6 and Virgo's VSR 2/3 data [8]. **Top Right:** Zoomed version of the left diagram shows better the distribution of the weak injections in the two statistics. **Bottom:** Distribution of pN-NR simulated injections can be mostly separated into two groups characterized by the duration, long or short, of the template they triggered in the coherent search. This plot shows that a small subset of the triggers with a long-duration template fall in the same cluster as the ones with a short-duration template, and *vice versa*. This is because a signal can trigger multiple templates; the loudest trigger in coherent statistic can correspond to a different template than that of the loudest trigger in a coincident statistic, which penalizes short-duration ones more than the long-duration ones.

third detector is weaker than the other two). This is an important change vis-a-vis the previous hierarchical coherent searches.

We analyzed the performance of the hierarchical coherent pipeline by injecting and recovering the different NR wave-forms submitted to the NINJA-2 project by various NR groups. For this study we considered only non-spinning waveforms with different mass ratios. Two sets of simulated GW signals were injected with distances between 20 Mpc to 2900 Mpc, distributed uniformly in linear or logarithmic distance. Just as for the template back, for the injections as well the total mass was chosen to be in the range $25 - 100 M_\odot$ and the component masses between $3 - 97 M_\odot$, All the injections recovered by the coincident algorithm were found by the hierarchical coherent pipeline as well, which is expected. On the other hand, although one might expect to find only triple-coincident events after the coherent stage, owing to the removal of the individual detector thresholds there, some events were found to be double-coincident. This is because these injections have very large effective distance in one of the detectors compared to the other two; the event time registered in the detector with the large effective distance (and, therefore, low SNR) typically had a large error and often did not fall within the light travel-time window corresponding to the two baselines formed with the other two detectors. This is why the hierarchical coherent pipeline did not promote them to be triple coincidences but retained them as double coincidences.

Figure 2 presents the injected decisive distance of found and missed injected BBH signals as a function of their total mass. As introduced earlier, the injected decisive distance is defined as the effective distance of a source in the detector where it was the second loudest, and is termed so because at least two detectors are required to hear the signal louder than the threshold to produce a coincident event. The closest missed injection had an injected decisive distance of 907.34 Mpc. Figure 3 compares the distribution of injected signals and background triggers using both
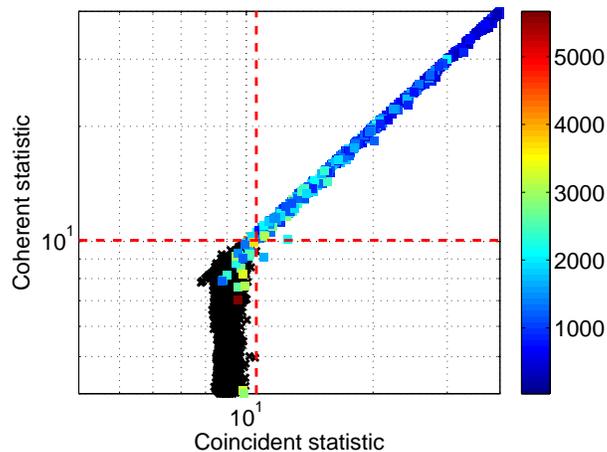
FIG. 4: The coherent and coincident detection statistic values of found simulated signals (colored squares) and background events (black crosses). The colorbar shows injected decisive distance in Mpc. The coincident detection statistic used here is the combined SNR. This plot confirms that the maximum value the coherent SNR of a trigger can assume is the combined SNR.

coincident and coherent detection statistics. The coherent analysis used the coherent network SNR as the detection statistic whereas the coincident analysis employed the same detection statistic as the one used for LIGO-Virgo BBH search of the S6-VSR2/3 data [8], namely, the one defined in Eq. (2.3). Due to the relatively poor performance of the χ^2 test for short-duration templates compared to long-duration ones, this statistic utilizes that information differently for GW triggers associated with those two categories of templates. This is reflected in its definition in Eq. (2.3), and the phenomenological case for it is explained in Ref. [8].

Figure 4 plots the coherent SNR of every found injection trigger and background trigger versus the coincident statistic, which in this case is the combined SNR. The color-filled squares denote found simulated signals and the black crosses represent background triggers. While the loudest background trigger in the coincident analysis has a coincident detection statistic value of 10.3, its counterpart in the coherent analysis has a coherent SNR of 10.0. The right plot in Fig. 4 is the zoomed version of the left plot, focusing on the most interesting region of the plot where the simulated signals and background events start mixing. The two plots show two red dashed lines, intersecting at right angles, that divide each plot into four quadrants. These lines cut across the loudest background triggers for each analysis, respectively. For instance, colored squares in the region $\rho_{\mathrm{coinc}} > 10.25$ are simulated signals found louder than the loudest background event, or simulated signals found with zero False Alarm Rate (FAR), in the coincident analysis. Similarly, the colored squares in the region $\rho_{\mathrm{coh}} > 10.0$ are simulated signals found with zero FAR in the coherent analysis. All the background triggers are confined to the bottom-left quadrant. The top-right quadrant has only injection triggers; these signals are found with a zero FAR in both coincident and coherent searches. The bottom-right quadrant has only injection triggers that have a zero FAR in the coincident search; note that these injection triggers have a non-zero FAR in the coherent search because the loudest background in the coherent search has an SNR (shown by the horizontal red line that is) greater than that of any of these inection triggers. On the other hand, the top-left quadrant has only injection triggers that have a zero FAR in the coherent search. When one counts the injection triggers in these quadrants, one finds that there are more of them in the top-left quadrant than in the bottom-right quadrant. This leads to the inference that in this study the coherent statistic performed better than the coincident one, at zero FAR.

It is also interesting to find that in Fig. 3 most of the simulated signals lie above the diagonal; this is due to the re-weighting of coincident SNR based on their χ^2 values. One also finds that the injection triggers in these plots are distributed in two branches. To investigate the reason behind this feature, we plotted the same set of simulated signals after binning them into two different sets based on the duration of the templates that detected them. Events are termed to be of a short duration if at least one of the detectors finding it has a template duration less than 0.2 sec. To be categorized as a long duration event all the participating detectors should have their template durations greater than or equal to 0.2 sec for that event. The bottom plot in that figure indeed shows that these two branches of simulated signals were formed mostly from long and short duration events. That plot also shows that a few triggers of each kind are picked by templates of the opposite kind. This is because a signal can trigger multiple templates; the loudest coherent SNR trigger in the second stage can correspond to a different template than the loudest coincident SNR trigger in the first stage. Note that the coincident statistic penalizes short-duration templates more than the long-duration ones.

The maximum value that the coherent SNR can take for any trigger, from background or injection, is the combined SNR. The combined SNR is the square root of the quadrature sum of the individual detector SNRs. Figure 4 confirms this statement for both simulated signals and background events. It is clear that most of the simulated signals fall on diagonal. A relative few, especially of the weak kind, fall below the diagonal. This is because the contribution
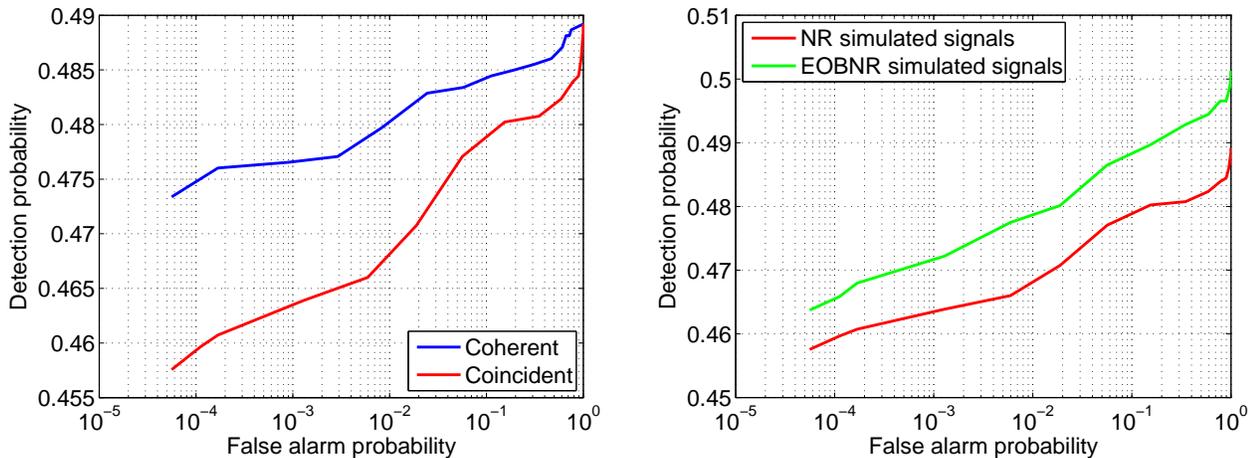
FIG. 5: **Left:** Receiver Operating Characteristic (ROC) curves from the injection and (partial) recovery of 2000 numerical relativity simulated signals. The two curves here compare the performances of the coherent and coincident searches. **Right:** The ROC curves for two different simulated signal families, namely, EOBNR and pN-NR, recovered with the same EOBNR template bank.

from their cross detector terms to the coherent SNR is less than maximal. As expected most of background events can be found below the diagonal due to the incoherence of their individual detector triggers. However, these is a small fraction of background triggers that lie on the diagonal. A majority of these are triggers were associated with sky positions that created issues for polarisation matrix inversion. Due to ill-posed polarisation matrix, detection statistic of these triggers had un-physical values larger than the combined SNR. Therefore we assigned combined SNR for detection statistic of such triggers allowing them to be on diagonal.

A. Comparing ROC curves

In this section we address the following two issues. First, we enquire if the coherent stage improves the detectability of a CBC signal, which in NINJA-2 is limited to the pN-NR hybrid kind. If an improvement is found then it will make the case for developing computationally viable fully coherent CBC searches. Second, we ask how much worse does an EOBNR template-bank perform in detecting pN-NR signals than those modelled by the EOBNR formalism itself. This exercise probes, in a limited way, if the differences in the deduction of the pN-NR waveforms, on the one hand, and the EOBNR waveforms, on the other hand, are consequential enough to affect the ability of our search pipelines to make a detection. The limitation of this exercise is that it does not address how different EOBNR and pN-NR hybrid waveforms are from a fully accurate waveform solution from General Relativity. For readers interested in that subject, we refer them to Refs. [33–37], and the references therein. We make both types of comparisons by computing the Receiver Operating Characteristic (ROC) curves [10] for each case.

The ROC curves in the left panel of Fig. 5 compare the performances of the hierarchical coherent and coincident searches for the same set of simulated pN-NR hybrid signals. The detection probability of the former is higher than the latter for all values of the false alarm probability (FAP) that could be computed in this exercise. But the region in this figure where this observation matters the most is the low FAP region around $10^{-5} - 10^{-4}$, where the first GW detections are expected to be made. Note, however, that the improvement obtained by running the additional coherent step in this coincident pipeline is limited by its hierarchical nature and, therefore, the utility of this study is that it suggests that a fully-coherent, all-sky, all-time pipeline should perform even better.

The FAP at a given value of the detection statistic, $\rho_{\mathrm{threshold}}$, is obtained as follows

$$\mathrm{FAP}(\rho_{\mathrm{threshold}}) = \frac{N_{\mathrm{background}}(\rho > \rho_{\mathrm{threshold}})}{N_{\mathrm{total\ background}}} \,, \tag{3.1}$$

where $N_{\mathrm{background}}(\rho > \rho_{\mathrm{threshold}})$ is the number of background triggers that have the value of their chosen detection statistic ρ greater than $\rho_{\mathrm{threshold}}$, and $N_{\mathrm{total\ background}}$ is the total number of background triggers, with any value of ρ, found by the search. On the other hand, the detection probability is given by

$$P_{\mathrm{detection}}(\rho > \rho_{\mathrm{threshold}}) = \frac{N_{\mathrm{recovered}}(\rho > \rho_{\mathrm{threshold}})}{N_{\mathrm{recovered}}(\rho > \rho_{\mathrm{threshold}}) + N_{\mathrm{missed}}} \,, \tag{3.2}$$

where $N_{\mathrm{recovered}}(\rho > \rho_{\mathrm{threshold}})$ is the number of recovered simulated signals with the detection statistic value greater than the value corresponding to a given false alarm probability, and N_{missed} denotes the number of simulated signals either with $\rho < \rho_{\mathrm{threshold}}$ or totally missed by the detection pipeline.

To get some insight into the relative behavior of the ROC curves in Fig. 5 note that the rate of change of FAP with respect to the detection statistic is faster in the case of the coincident statistic than in the case of the coherent statistic. However, the value of the detection statistic for an injection trigger scales inversely with its distance. Also, the detection probability decreases with increasing source distance. Thus, it follows that the rate of change of FAP with respect to the detection *probability* is faster in the case of the coincident statistic than the coherent statistic. This is exactly what is found in the left plot in Fig. 5. The fact that the two ROC curves there should meet at a high enough FAP, where the detection threshold is very small, then implies that the ROC curve of the coherent search should be above that of the coincident search. The detection probability does not go to unity in that figure because a hard cut-off of $\rho = 5.5$ was placed in each detector in the first stage, thereby, causing a fraction of the triggers to be missed. In this experiment that fraction happened to be about 51%. Figure 8 shows how the loudness of the background triggers is distributed for the two searches.

Figure 5 shows that at a FAP of 6×10^{-5} the difference between the detection probabilities of the coherent and coincident searches is approximately $(47.4 - 45.8)\% = 1.6\%$. This implies that for a total of 2000 injections the former search found 32 more than the latter, at that FAP.[3] One must bear care, however, in drawing conclusions for astrophysical searches from this simulation study because the source rate for aLIGO is estimated to be in the several tens and not thousands [38], and the aforementioned improvement will affect the detection of very few sources. Instead, the main conclusion is that while hierarchical coherent searches may not be worth investing resources into, they demonstrate the (small) improvement we expect of them and, in turn, suggest that it may be worthwhile to explore how much more gain in detection rate one can achieve with fully coherent all-sky, all-time algorithms in realistic search pipelines.

We next enquire how much the detectability of a signal suffers owing to the fact the pN-NR hybrid waveforms that are used to model the signals are not quite the same as the EOBNR waveforms used to model the templates [16]. The right panel of Fig. 5 compares the performance of two sets, with the same number of simulated signals injected at the same sky locations, but with two different signal families: The green curve corresponds to signals modelled with the EOBNR family and the red curve represents simulated signals made using the NINJA-2 pN-NR hybrid waveforms. Both sets of simulated signals were searched with EOBNR templates. As expected, simulated signals from the same family as the templates, namely EOBNR, register a better performance, at all FAP values. That figure also shows that the mismatch between the template and the signal families result in a small effect on the detection probability, which is of the order of a few percent at the most.

The variation of detection efficiency with distance also shows the expected behavior. We define it to have the same expression as $P_{\text{detection}}$ in Eq. (3.2), but now $P_{\text{detection}}$, $N_{\text{recovered}}(\rho > \rho_{\text{threshold}})$, and N_{missed} are all computed in multiple distance bins. Figure 6 shows the detection efficiencies of the hierarchical coherent and coincident searches at $\rho_{\text{threshold}}$ corresponding to FAP=0 and 10^{-3}. The detection efficiencies for all of the searches are very similar; they begin at 100% for nearby sources and fall off to zero near 2 Gpc, which is the greatest horizon distance [26] of the set of BBH sources simulated here.

B. The null-stream

Real data is neither Gaussian nor stationary and a statistic that is optimal in Gaussian and stationary noise may not remain so in real data. This makes the case for seeking a more effective detection statistic and signal-based vetoes for LIGO/Virgo science data. An advantage of multi-detector coherent searches is that a detector network with three or more baselines can (over-)determine the two waveform polarizations, in addition to identifying the source location with time-delay triangulation [26]. In such cases, one can form a linear combination of detector time-series outputs that contains no GW signal [21]. Such a combination is called the "null stream". The null stream is consistent with the noise in a detector network, and a noise artifact or a glitch in a detector that is uncorrelated with noise in the other detectors is expected to leave a residue in the null stream. Therefore, its presence can be used to veto a candidate event. The simplest example of the null-stream is the one for a network consisting two co-located, co-aligned detectors. In the simple case where both detectors in such a pair have the same sensitivity, the null stream corresponding to any search template is proportional to the difference of its matched-filter outputs from the data of the two individual detectors. When the two detectors in the pair have different sensitivities, the matched-filter outputs are inversely weighted with their respective template-norms before calculating the difference [20]. The value of the null-stream time-series at the trigger time is called the null statistic.

Figure 7 presents the null-statistic distribution for simulations and background events as a function of their coherent SNR. For very strong signals the null statistic is large since even a slightly imperfect subtraction of two large SNR values can leave a moderately large residue. This is not a major concern since the distribution of loud injections on this plot is well separated from that of the background triggers. However, the null statistic can play a critical role in the detection of weak signals. Those occur in the region where their distribution on this plot mixes with that of the background triggers. Figure 7 shows that nearly all the background triggers fall on a relatively tight band and only a

---

[3] Note that these two searches are correlated in the sense that a strong (weak) signal in one is highly likely to produce a strong (weak) signal in the other.
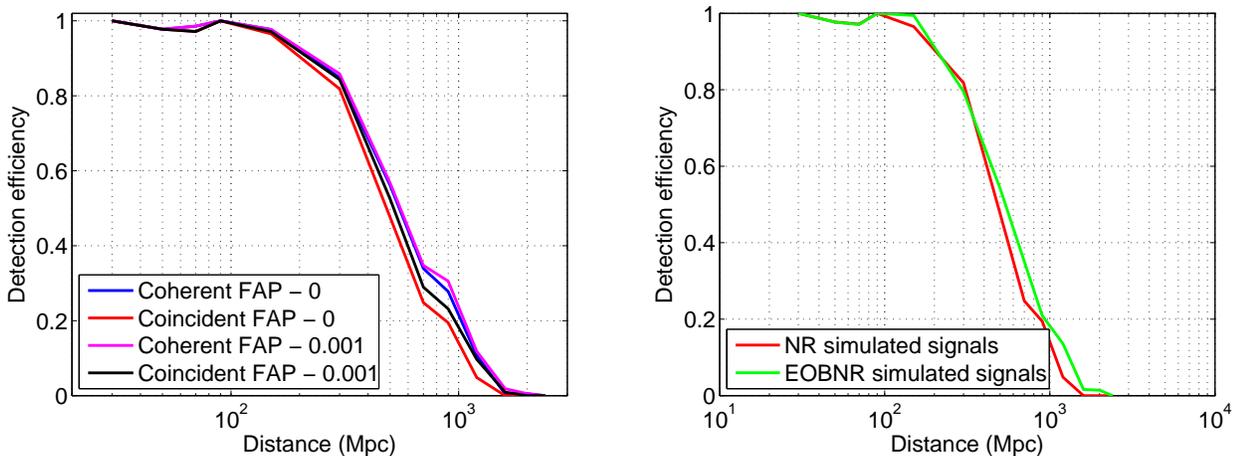
FIG. 6: **Left:** The detection efficiency is plotted as a function of the injected decisive distance (Mpc), for the hierarchical coherent and coincident searches, at false-alarm rates (FAPs) of 0.000 and 0.001. **Right:** Detection efficiency comparison of two different sets of simulated signals, namely, EOBNR and pN-NR.

couple lie outside it. Those two outliers are also the loudest background triggers in coherent SNR. Interestingly, these are triggers for which the sky positions were such that the LIGO-Virgo network could not resolve the putative signal polarizations. Just as for the injections, for these two triggers too the search pipeline used the coincident statistic as the detection statistic, the values of which were found to be well below 10.0. Therefore, Fig. 7 demonstrates that the null statistic is helpful in mitigating the impact of loud background triggers on a search pipeline's detection efficiency.

Although we will not be employing hierarchical search pipelines in ADE it is important to understand its issues in order to improve the efforts on developing fully coherent searches. As shown above, the thresholding criteria used in current coincident and hierarchical coherent pipelines can limit the detection efficiency of the search. In the current pipeline we only analyze triggers that cross the first matched-filtering stage with an SNR greater than 5.5. Additionally, we also require at least two such triggers in two different detectors within the window of light travel time between two sites to claim a gravitational-wave detection. It is possible to have gravitational-wave signals that only have threshold crossing trigger in one detector with significant coherent SNR. To investigate this in the same NR simulations, we searched for such events. Numerical relativity simulations plotted in Fig. 8 have only one threshold-crossing detector, yet their coherent SNRs are relatively strong: Out of a total of 2000 simulated pN-NR signal injections, we found that 52 had coherent SNR above 9.5. For each of these found injections only one of the three detectors in coincidence produced a single-detector SNR of above 5.5.[4] Furthermore, 35 of them had a coherent SNR higher than the loudest background trigger. Figure 9 reveals that some of these simulations had larger single-detector SNRs in Virgo compared to those in the two LIGO detectors. Due to their similar orientation and sky coverage, both LIGO detectors had very similar sensitivities in a large fraction of the sky. Additionally, their similar noise PSDs means that signals from the same source will have very similar SNRs in those detectors most of the time. But Virgo's orientation is quite different from that of the LIGO detectors. Therefore, in Fig. 9 some of the events have larger SNR contributions from Virgo than from any of the LIGO detectors.

IV. PARAMETER RECOVERY

Parameter estimation will play an important role in making astrophysical statements about the origin of the GW signals we detect (see, e.g., Ref. [39]). For instance, by measuring the masses of BBH components through GW observations we will be able to constrain stellar population synthesis models. This section presents results from our EOBNR template searches of pN-NR simulated injections to educate us on how accurately one might be able to measure the parameters of real BBH signals. In this study, parameters of those signals were estimated using the maximum likelihood method [10]. Figure 10 shows that the measured coherent SNR is within a few percent of the expected value when for signals the latter exceeds 20. Moreover, this observation holds across the whole range of total-mass values of the injected non-spinning BBH systems.

A useful construct for describing the error in the measurement of the value of a dimensional parameter is:

$$\text{fractional error} = \frac{\text{measured value - injected value}}{\text{injected value}}, \tag{4.1}$$

---

[4] For a three-detector network with single-detector thresholds of 5.5 each, the minimum combined SNR for an event is $\sqrt{3} \times 5.5 = 9.52$.
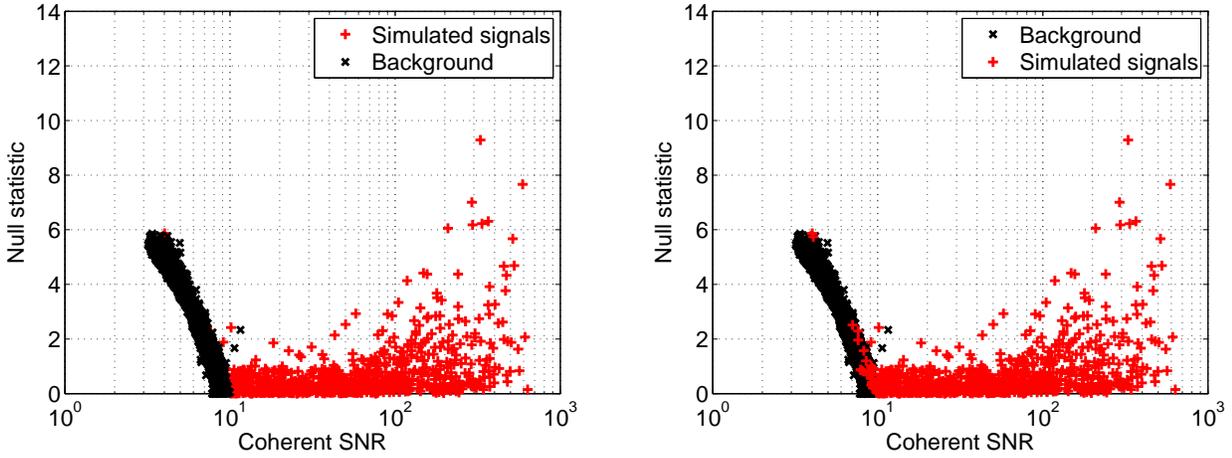
FIG. 7: **Left:** We plot the null-statistic for background and pN-NR simulated signals as a function of their coherent SNR. Most of the background triggers form a "band", which helps to separate signals from them. Notice that two of the loud background triggers prominently fall outside the band. These correspond to sky positions where the network does not resolve the GW polarizations very well [20]. In such a case, the trigger is vetoed. A better way to handle such triggers would be to not veto them but construct a more effective detection statistic by using a network's polarization resolving power at the trigger's sky position. This aspect will be explored elsewhere. **Right:** Here we show the injection triggers above the background triggers to highlight that for very weak injections, the null-statistic worsens and loses its discriminatory power.
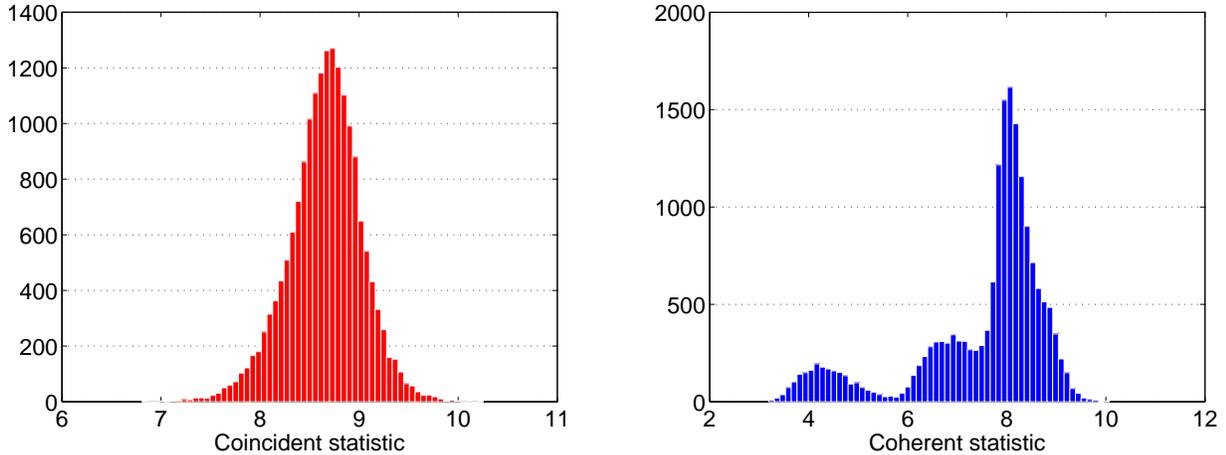
FIG. 8: **Left:** Distribution of background events in the coincident detection statistic used for the LIGO S6 - Virgo VSR2/3 high-mass search. **Right:** Distribution of background events in the coherent detection statistic.

which we will use especially for studying the error in the measured or *recovered* chirp-mass values. Comparing the injected and recovered parameter values helps us in determining if any systematics can develop in searching for pN-NR signals with a specific waveform family. Consider, e.g., the chirp-mass recovered by EOBNR templates in Fig. 11. There, in most cases, the injected chirp mass is greater than the recovered one. This can happen owing to the following possibilities: (a) For some cases of pN-NR waveforms EOBNR templates with a smaller mass have the same or similar number of cycles. If this is true, then this is information that the hybridization schemes of those pN-NR waveforms can use to correct for this bias. (b) Templates that are longer in time and have more cycles than these pN-NR injections perform better at finding them. Longer templates would accrue noise from the cycles that do not overlap with any part of the signal and would, therefore, lose SNR. However a mismatch between an injection and the template, in amplitude or phase, can cause the projection of the signal on the template manifold to select a longer template as the best fit. Indeed, the chirp-mass recovery and end-time recovery plots reveal that pN-NR injections tend to be found often with relatively higher chirp-mass EOBNR templates. Since higher chirp-mass templates will have a smaller duration than lower chirp-mass ones, all else being the same, one would expect that this bias should be accompanied by a systematic error whereby the template (or recovered) end-time occurs earlier than the injected (pN-NR) end-time. This expectation is confirmed in Fig. 12.

The end-time difference of recovered simulated signals in Fig. 12 is computed by subtracting the injected end time from the measured end time. A negative end-time difference implies that the pN-NR signal is found at a later time than its injected value. A majority of the signals were recovered within 10 msec of the injected end time. A few outliers with a greater positive (smaller negative) end-time difference arise from systems with a large (small) total
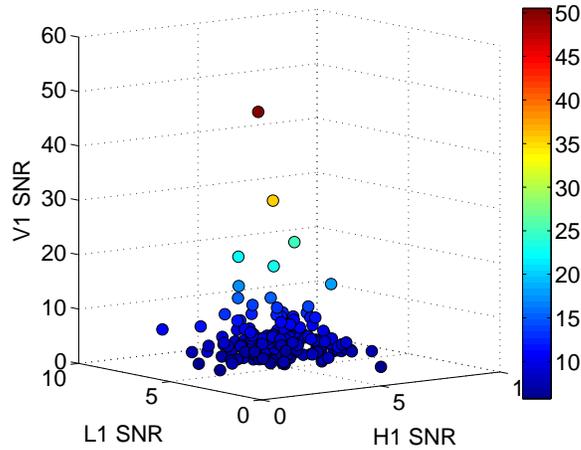
FIG. 9: The coincident pipeline uses thresholds on matched-filter outputs from all detectors in a network to keep the number of background trigger coincidences at a manageable level so that the search remains computationally viable. An event has to be coincident in at least two detectors before it can be treated as a detection candidate. While this requirement improves confidence in a detection, it also hurts detection efficiency by rejecting signals with marginally sub-threshold triggers in two of three detectors that would otherwise produce a candidate event with comparable detection confidence. This plot shows the distribution of such events, with the colorbar displaying their coherent SNR. Out of 2000 simulated pN-NR signals, 48 were found with a zero false-alarm rate that have the SNR crossing the standard 5.5 threshold in only one of the three detectors. Also 54 events were found that had SNR crossing the standard 5.5 threshold in all three detectors.
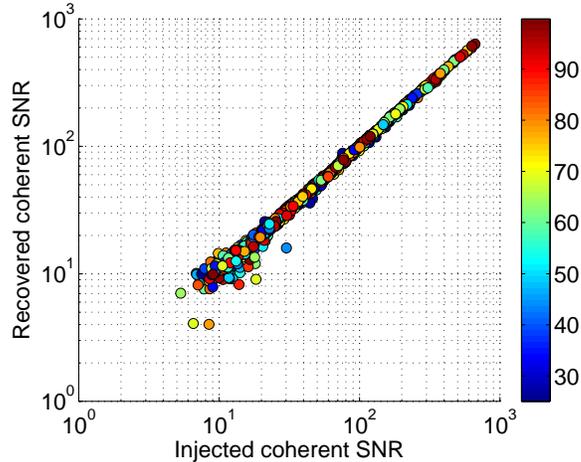
FIG. 10: Recovered coherent SNR versus its injected value for NR simulated signals. The colorbar shows the total mass of the BBH systems whose pN-NR hybrid waveforms were injected.

mass and relatively low coherent SNR.

The results for the measurement of the symmetrized mass-ratio η are presented in Fig. 13. The end-time difference is computed by subtracting the injected value from the measured one. As shown in Ref. [16], only six mass-ratio values were used to model and inject non-spinning pN-NR signals. For EOBNR signals, we produced a greater variety of mass ratios and, therefore, η, as can be seen in the top-right plot in that figure. The measured values of η are always between 0.00 and 0.25 because those are the boundaries of the template bank that was employed, and correspond to mass ratios of infinity and 1, respectively. It is interesting to note that in a large fraction of cases small (large) η, or low (high) mass ratio, templates detect large (small) η signals. This bias is, in fact, correlated with low (high) chirp-mass templates detecting high (low) chirp-mass signals, as confirmed by Fig. 14 for injections of pN-NR as well as EOBNR signals. This covariance of errors occurs because templates with low (high) η and low (high) $M_{\rm chirp}$ can have similar number of wave cycles and bandwidths as signals with high (low) η and high (low) $M_{\rm chirp}$.[5] In the same figure, also note how the degree of covariance changes somewhat as one goes from low total-mass systems (shown in blue) to high total-mass systems (shown in red). A similar behavior was observed in

---

[5] A similar correlation of errors in η and the total-mass M was found in Ref. [14] where the effect of using inspiral-only templates for searching inspiral-merger-ringdown signals on detection and parameter estimation was studied.
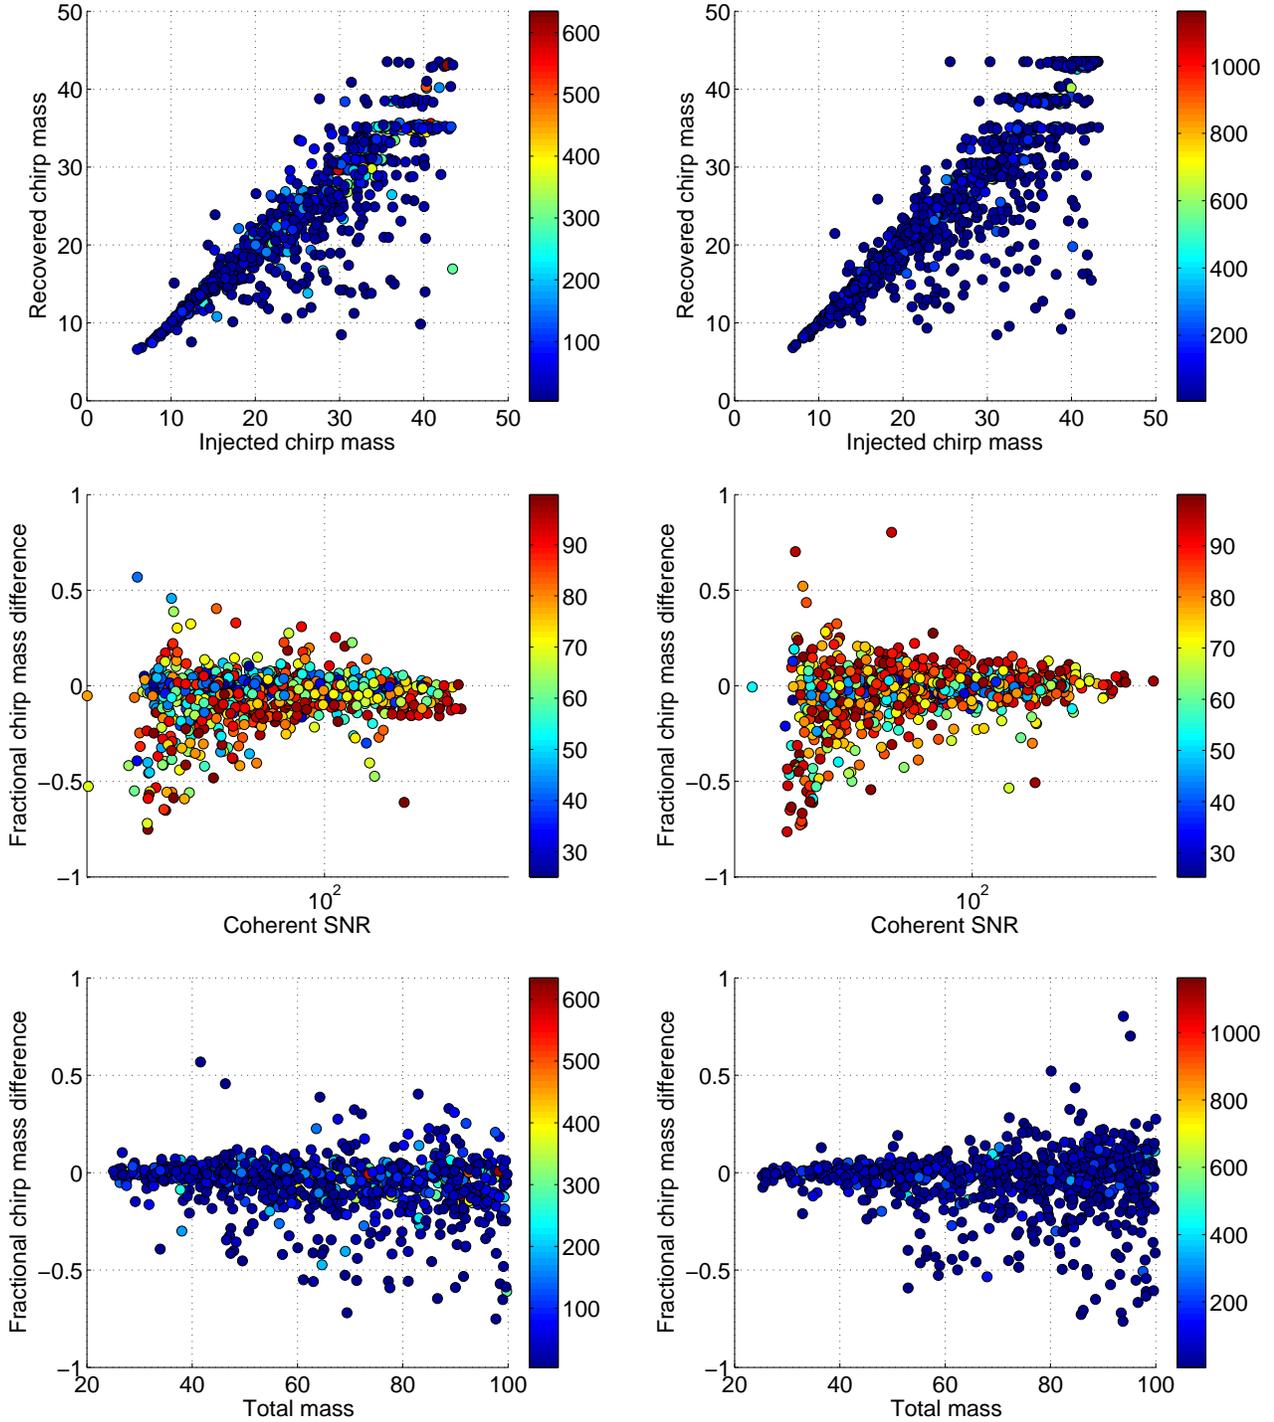
FIG. 11: **Left column:** Errors in the measurement of the chirp mass of pN-NR simulated signals. **Right column:** Same as in the left column but for EOBNR simulated signals. **Top panel:** Recovered chirp mass as a function of injected chirp mass for pN-NR simulations (left) and EOBNR simulations (right), with the colorbar showing the coherent SNR. **Middle panel:** Fractional chirp-mass difference as a function of coherent SNR with the colorbar displaying the total mass of the BBH system. **Bottom panel:** Fractional chirp-mass difference as a function of the total mass of the BBH system, with the colorbar showing the coherent SNR. According to the bottom panel the overall chirp mass recovery is slightly better for EOBNR simulations than the pN-NR ones, due to better match with the templates in the former case.

the Monte Carlo studies in Ref. [39] of statistical errors in measuring BBH mass parameters with phenomenological waveforms [40].
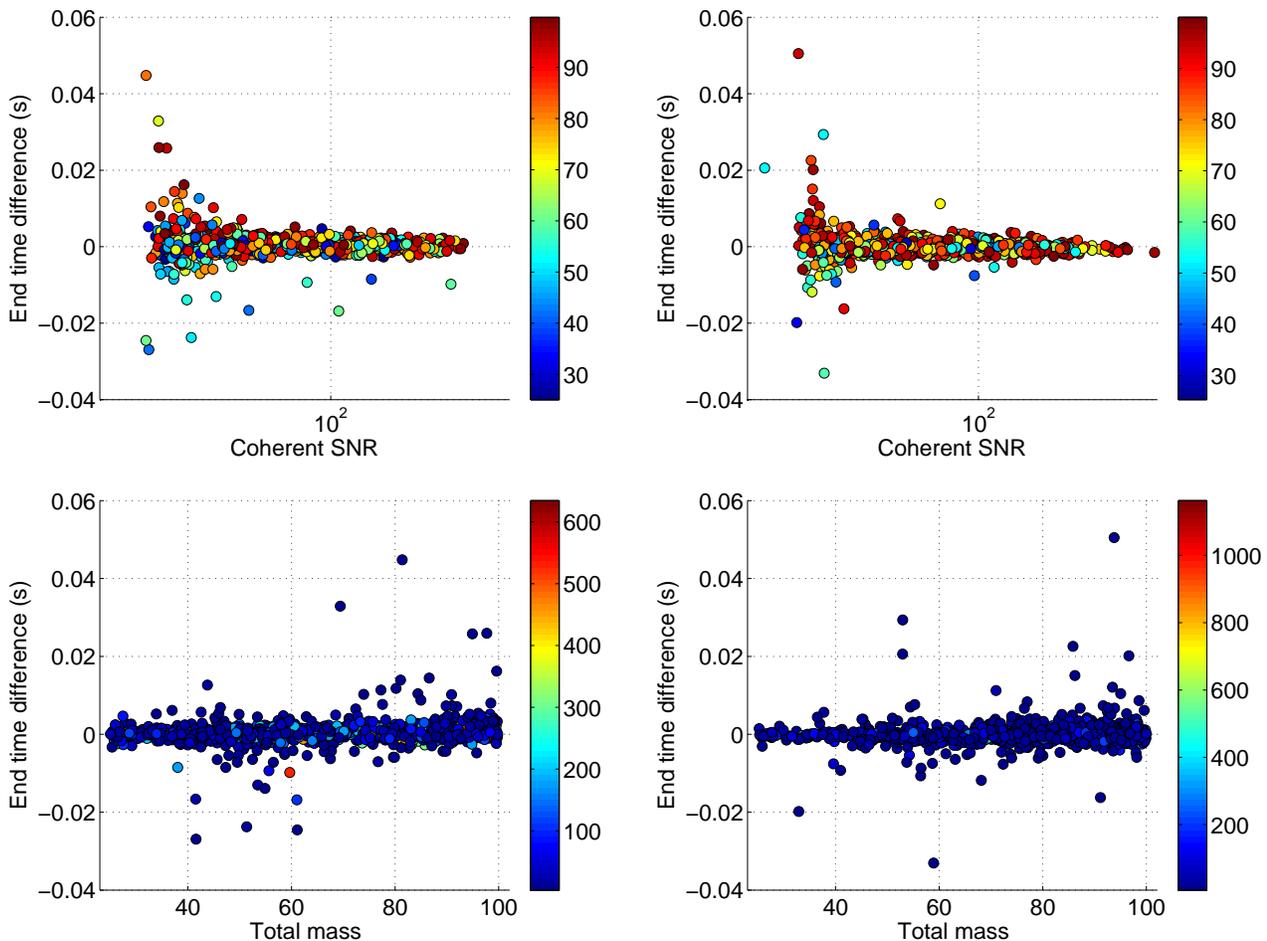
FIG. 12: Difference between the measured and injected end time of detected simulated signals (in seconds). **Top left:** The end-time difference is plotted as a function of the coherent SNR for pN-NR simulations. The colorbar represents the total mass of the binary system (in $M_\odot$). **Top right:** The same quantity is plotted for EOBNR simulations. For pN-NR simulated signals a few were recovered with relatively high negative end-time difference, i.e., ≥ 0.01 sec, in the total-mass range $40M_\odot$ to $62M_\odot$, at both high and low coherent SNRs. **Bottom panel:** The end-time difference is plotted as a function of the total mass of the system.

V. DISCUSSION

Testing the preparedness of search pipelines to detect real signals in the ADE is one of the primary goals of NINJA. This study helped to quantify the sensitivity of one of the existing CBC search pipelines to numerical-relativity based BBH waveforms in early advanced detectors *albeit* in simulated Gaussian, stationary data. As we found here, the EOBNR waveforms, which employed only NASA-Goddard NR waveforms for calibration, were able to detect pN-NR signals produced by a number of NR collaborations. This is borne out by the fact that the ROC curve of EOBNR injections tracks that of the NR injections very closely. Additionally, we studied some aspects of a coherent CBC search. Since a coherent search explores a larger dimensional parameter space than a coincident search it is more expensive, which makes the estimation of the background for the former type of search especially difficult. Therefore, we used the hierarchical coherent CBC search pipeline described in Ref. [20] on the same NR based injections. Such an exercise is also useful to teach us about a subset of the potential issues we may face in a fully coherent search in the future. Here we demonstrated that the performance of that search conforms to expectations. Specifically, the characteristics of its background are consistent with theoretical predictions. Moreover, the coherent stage provides the null-stream statistic, which is a powerful multi-baseline signal consistency test, and can be employed to improve the performance of the search. This test is especially useful for high-mass CBC searches where the chi-square test is less effective than low-mass ones, owing to fewer waveform cycles of high-mass signals in the detector band.

Finally, NINJA provided an important opportunity to test how well we might be able to measure the signal parameters. To address this question, we compared the maximum-likelihood estimates obtained by using the EOBNR family of templates. We focused our attention on only non-spinning injections. Here again the parameter accuracies of EOBNR injections are very similar to those of the pN-NR ones. The only small disagreement occurs for a few injections, mainly in the small total-mass region. For most of them, its cause was traced to the fact that compared to other injected waveforms these ones were *a priori* known to have a somewhat poorer match (by a few percent)
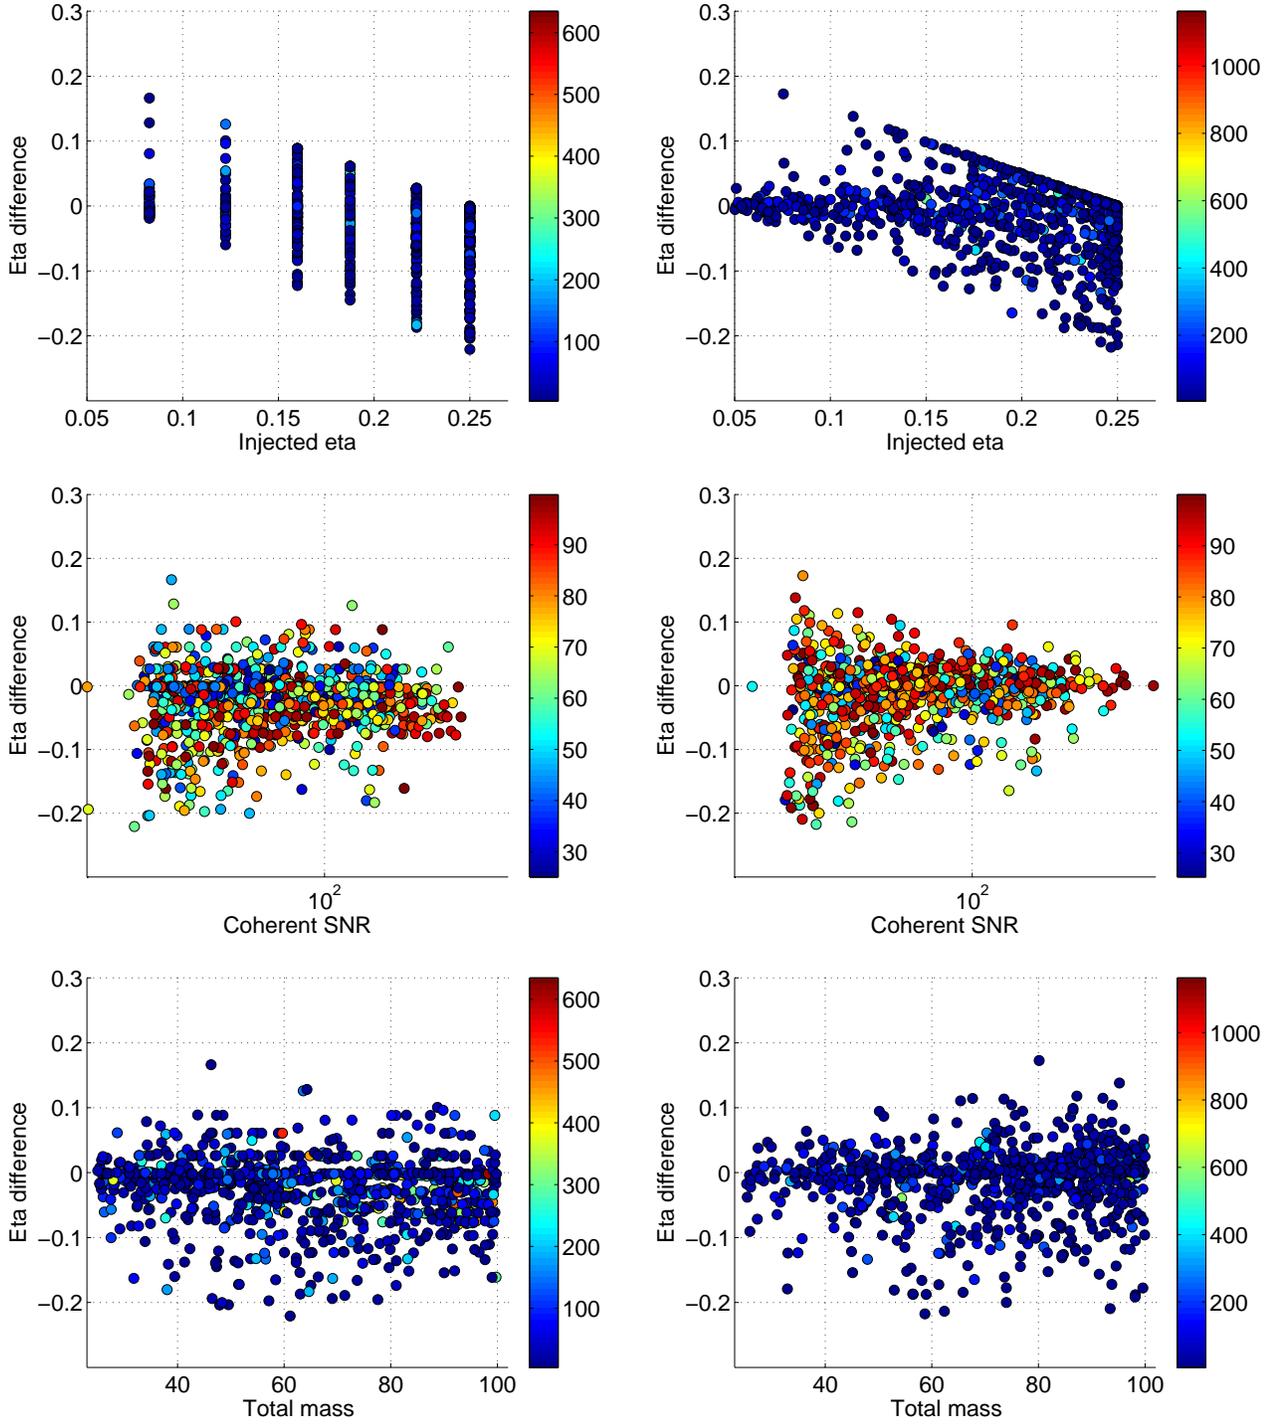
FIG. 13: **Left column:** Errors in the measurement of the symmetrized mass-ratio η or "Eta" of pN-NR simulated signals. **Right column:** Same as in the left column but for EOBNR simulated signals. **Top panel:** This panel shows the η difference as a function of injected η values. Since pN-NR simulations are only available for a discrete set of η values the left plot only has a few different values of injected η compared to EOBNR counterpart, which is plotted in the right. Since $\eta \in (0, 0.25]$, the maximum (minimum) η difference is 0.25 (-0.25). **Middle panel:** This panel shows that the η difference of pN-NR hybrids is slightly worse compared to that of EOBNR simulations due to template mismatch. This is most apparent above at large coherent SNR values. **Bottom panel:** Eta difference of two simulation categories is shown as a function of the total mass of the system.
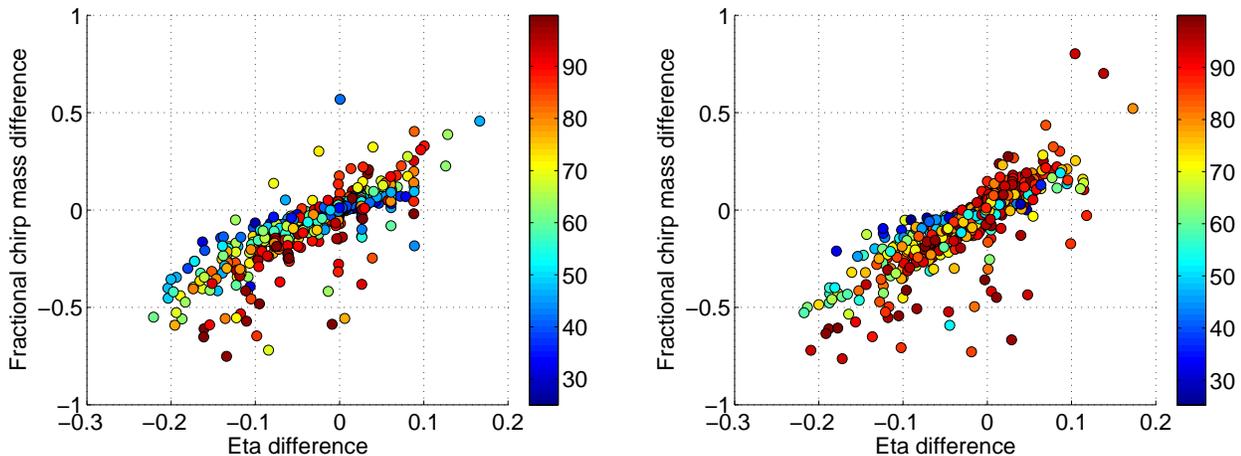
FIG. 14: Covariance of chirp-mass error with η error. The left figure shows this covariance for pN-NR injections and the right figure shows it for EOBNR injections. In both cases EOBNR templates were used to conduct the search. The colorbar shows the total mass of the BBH systems (in $M_\odot$).

[16] with the waveforms produced by the Spectral Einstein Code (SpEC). (See Ref. [41] for details about SpEC.) To study systematic errors stemming from signal-template mismatch in more detail, studies are ongoing with NR waveforms with spin and a variety of different mass ratios. Additionally, to study the impact of real data, which can be non-Gaussian and non-stationary, we plan to study the NINJA-2 waveforms injected in a recolored data set, where real data from past science runs will be scaled to have early aLIGO or aLIGO ASDs [32].

Acknowledgments

[1] Gregory M. Harry (for the LIGO Scientific Collaboration). Advanced LIGO: the next generation of gravitational wave detectors. *Class. Quant. Grav.*, 27:084006, 2010.

[2] The Virgo Collaboration. Advanced Virgo Technical Design Report. 2012.

[3] Kentaro Somiya. Detector configuration of KAGRA: The Japanese cryogenic gravitational-wave detector. *Class.Quant.Grav.*, 29:124007, 2012.

[4] Jeffrey M. Silverman and Alexei V. Filippenko. On IC 10 X-1, the Most Massive Known Stellar-Mass Black Hole. *Astrophys.J.*, 678:L17–L20, 2008.

[5] P. A. Crowther et al. NGC 300 X-1 is a Wolf-Rayet/Black-Hole binary. *MNRAS*, page L11, 2010.

[6] Sean Farrell, Natalie Webb, Didier Barret, Olivier Godet, and Joana Rodrigues. An Intermediate-mass Black Hole of Over 500 Solar Masses in the Galaxy ESO 243-49. *Nature*, 460:73–75, 2009.

[7] J. Abadie et al. Search for gravitational waves from binary black hole inspiral, merger and ringdown. *Phys. Rev.*, D83:122005, 2011.

[8] J. Aasi et al. Search for Gravitational Waves from Binary Black Hole Inspiral, Merger and Ringdown in LIGO-Virgo Data from 2009-2010. D83, 2012.

[9] Paul A. Lynn. *An introduction to the analysis and processing of signals.* Hemisphere Publishing Corporation, 79 Madison Avenue, NY, USA, 1989.

[10] Carl W. Helstrom. *Elements of signal detection and estimation.* Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1995.

[11] Advanced LIGO Anticipated Sensitivity Curves. https://dcc.ligo.org/cgi-bin/DocDB/ShowDocument?docid=2974.

[12] Advanced Virgo Baseline Design. http://pub3.ego-gw.it/codifier/includes/showTmpFile.php?doc=2219&calledFile=VIR-027A-09.pdf.

[13] Ilya Mandel, Duncan A. Brown, Jonathan R. Gair, and M. Coleman Miller. Rates and Characteristics of Intermediate-Mass-Ratio Inspirals Detectable by Advanced LIGO. *Astrophys.J.*, 681:1431–1447, 2008.

[14] Sukanta Bose, Shaon Ghosh, and P. Ajith. Systematic errors in measuring parameters of non-spinning compact binary coalescences with post-Newtonian templates. *Class.Quant.Grav.*, 27:114001, 2010.

[15] Benjamin Aylott, John G. Baker, William D. Boggs, Michael Boyle, Patrick R. Brady, et al. Testing gravitational-wave searches with numerical relativity waveforms: Results from the first Numerical INJection Analysis (NINJA) project. *Class.Quant.Grav.*, 26:165008, 2009.

[16] P. Ajith et al. The NINJA-2 catalog of hybrid post-Newtonian/numerical-relativity waveforms for non-precessing black-hole binaries . *Class. Quant. Grav.*, 29:124001, 2012.

[17] Alessandra Buonanno, Yi Pan, John G. Baker, Joan Centrella, Bernard J. Kelly, et al. Toward faithful templates for non-spinning binary black holes using the effective-one-body approach. *Phys.Rev.*, D76:104049, 2007.

[18] P. Ajith et al. Addendum to 'The NINJA-2 catalog of hybrid post-Newtonian/numerical-relativity waveforms for non-precessing black-hole binaries'. *Class. Quant. Grav.*, 30:199401, 2013.

[19] S. Babak, R. Biswas, P.R. Brady, D.A. Brown, K. Cannon, et al. Searching for gravitational waves from binary coalescence. *Phys.Rev.*, D87:024033, 2013.

[20] S. Bose, T. Dayanga, S. Ghosh, and D. Talukder. A blind hierarchical coherent search for gravitational-wave signals from coalescing compact binaries in a network of interferometric detectors. *Class. Quant. Grav.*, 28:134009, 2011.

[21] Y. Guersel and M. Tinto. Near optimal solution to the inverse problem for gravitational wave bursts. *Phys. Rev.*, D40:3884–3938, 1989.

[22] D. Talukder, S. Bose, S. Caudill, and Paul T. Baker. Improved Coincident and Coherent Detection Statistics for Searches for Gravitational Wave Ringdown Signals, in preparation.

[23] Ian Harry and Stephen Fairhurst. A targeted coherent search for gravitational waves from compact binary coalescences. *Phys. Rev. D*, 83:084002, 2011.

[24] Archana Pai, Sanjeev Dhurandhar, and Sukanta Bose. A data-analysis strategy for detecting gravitational-wave signals from inspiraling compact binaries with a network of laser-interferometric detectors. *Phys. Rev. D*, 64:042004, 2001.

[25] N. Arnaud et al. Coincidence and coherent data analysis methods for gravitational wave bursts in a network of interferometric detectors. *Phys. Rev. D*, 68:102001, 2003.

[26] Bernard F. Schutz. Networks of gravitational wave detectors and three figures of merit. *Class. Quant. Grav.*, 28:125023, 2011.

[27] S. Klimenko et al. Localization of gravitational wave sources with networks of advanced detectors . *Phys. Rev. D*, 83:102001, 2011.

[28] C. A. K. Robinson, B. S. Sathyaprakash, and Anand S. Sengupta. A geometric algorithm for efficient coincident detection of gravitational waves. *Phys. Rev.*, D78:062002, 2008.

[29] Bruce Allen. A chi^2 time-frequency discriminator for gravitational wave detection. *Phys. Rev. D*, 71:062001, 2005.

[30] J. Aasi et al. Prospects for Localization of Gravitational Wave Transients by the Advanced LIGO and Advanced Virgo Observatories. 2013.

[31] Alessandra Buonanno and Thibault Damour. Transition from inspiral to plunge in binary black hole coalescences. *Phys. Rev. D.*, 62:064015, 2000.

[32] The LIGO Scientific The NINJA-2 and The Virgo collaborations. The NINJA-2 Blind Injection Challenge, in preparation. 2013.

[33] Mark Hannam, Sascha Husa, Frank Ohme, and P. Ajith. Length requirements for numerical-relativity waveforms. *Phys.Rev.*, D82:124052, 2010.

[34] Michael Boyle. Uncertainty in hybrid gravitational waveforms: Optimizing initial orbital frequencies for binary black-hole simulations. *Phys.Rev.*, D84:064013, 2011.

[35] Sebastiano Bernuzzi, Marcus Thierfelder, and Bernd Bruegmann. Accuracy of numerical relativity waveforms from binary neutron star mergers and their comparison with post-Newtonian waveforms. *Phys.Rev.*, D85:104030, 2012.

[36] Ilana MacDonald, Samaya Nissanke, and Harald P. Pfeiffer. Suitability of post-Newtonian/numerical-relativity hybrid waveforms for gravitational wave detectors. *Class.Quant.Grav.*, 28:134002, 2011.

[37] Frank Ohme, Alex B. Nielsen, Drew Keppel, and Andrew Lundgren. Statistical and systematic errors for gravitational-wave inspiral signals: A principal component analysis. 2013.

[38] J. Abadie et al. Predictions for the Rates of Compact Binary Coalescences Observable by Ground-based Gravitational-wave Detectors. *Class.Quant.Grav.*, 27:173001, 2010.

[39] P. Ajith and Sukanta Bose. Estimating the parameters of non-spinning binary black holes using ground-based gravitational-wave detectors: Statistical errors. *Phys. Rev.*, D79:084032, 2009.

[40] P. Ajith et al. A template bank for gravitational waveforms from coalescing binary black holes: non-spinning binaries. *Phys. Rev. D*, 77:104017, 2008.

[41] Michael Boyle, Duncan A. Brown, Lawrence E. Kidder, Abdul H. Mroue, Harald P. Pfeiffer, et al. High-accuracy comparison of numerical relativity simulations with post-Newtonian expansions. *Phys.Rev.*, D76:124038, 2007.