

Long-duration transient gravitational-wave search pipeline

A. Macquet¹, M. A. Bizouard¹, N. Christensen¹ and M. Coughlin²

¹*Artemis, Université Côte d'Azur, Observatoire de la Côte d'Azur, CNRS, Nice 06300, France*

²*School of Physics and Astronomy, University of Minnesota, Minneapolis, Minnesota 55455, USA*



(Received 24 August 2021; accepted 19 October 2021; published 19 November 2021)

As the sensitivity and observing time of gravitational-wave detectors increase, a more diverse range of signals is expected to be observed from a variety of sources. Especially, long-lived gravitational-wave transients have received interest in the last decade. Because most long-duration signals are poorly modeled, detection must rely on generic search algorithms, which make few or no assumptions on the nature of the signal. However, the computational cost of those searches remains a limiting factor, which leads to suboptimal sensitivity. Several detection algorithms have been developed to cope with this issue. In this paper, we present a new data analysis pipeline to search for unmodeled long-lived transient gravitational-wave signals with duration between 10 and 10^3 s, based on an excess cross-power statistic in a network of detectors. The pipeline implements several new features that are intended to reduce computational cost and increase detection sensitivity for a wide range of signal morphologies. The method is generalized to a network of an arbitrary number of detectors and aims to provide a stable interface for further improvements. Comparisons with a previous implementation of a similar method on simulated and real gravitational-wave data show an overall increase in detection efficiency for all but one signal morphologies tested and a computing time reduced by at least a factor 10.

DOI: [10.1103/PhysRevD.104.102005](https://doi.org/10.1103/PhysRevD.104.102005)

I. INTRODUCTION

A new era in astronomy began in September 2015 with the observation of gravitational waves (GWs) from the merger of two stellar mass black holes [1]. Since then, the Advanced LIGO [2] and Advanced Virgo [3] have regularly observed a larger volume of the Universe leading, among major discoveries, to the observation of the merger of two neutron stars [4] in August 2017 associated with gamma ray burst GRB190817A [5] followed up with kilonova AT2017gfo in NGC4993 [6]. As of mid-2021, Advanced LIGO and Advanced Virgo have reported ~ 50 confirmed mergers of compact objects, black holes, and/or neutron stars [7].

Yet many GW sources have not yet been observed: core collapse supernova, isolated neutron stars, magnetars, cosmic strings, and the resulting stochastic background of GWs [8]. The diversity of the GW signal expected from these sources require different detection algorithms. When the GW signal waveform is predicted analytically, matched filter techniques can be used. In practice, this concerns mainly compact objects' binary coalescence, cosmic strings' signals [9], and GWs from pulsars [10,11]. When the GW emission is poorly modeled, detection will rely on unconstrained searches that make few assumptions about the characteristics of the signal. In the last 20 years, several search algorithms have been developed, mostly focusing on GW signals of duration less than a few seconds [12–16]. More recently,

transient GW signals of longer duration have received attention, bridging the gap between short-duration transient and continuous emission of GWs, and dedicated search algorithms have been developed [17–23].

Several astrophysical processes could be at the origin of long-duration transient GWs emission, for example, those related to core collapse supernova, compact object binary mergers, and isolated neutron stars [18]. Some of them are associated with the most energetic phenomena observed in the Universe. There is evidence [24,25] that core collapse supernovae and long gamma-ray bursts (GRBs) are connected to the death of massive stars where the iron core collapses under its own gravity, forming either a black hole or a highly magnetized neutron star, releasing an incredible amount of energy (10^{53} erg) mainly through neutrino emission, while $\sim 1\%$ goes into the kinetic energy of the explosion [26]. Once the collapse is triggered, very powerful nonspherical flows develop in the outer region of the protoneutron star that are expected to generate GWs energetically bounded to 10^{44} – 10^{47} erg [27]. The GW emission will last until the onset of the explosion or until a black hole is formed. The signal is expected to be no longer than 1–2 s.

In the collapsar model, massive stars collapse to black holes either without an initial supernova explosion or via fallback accretion after a successful but weak explosion [28]; a rotating black hole is formed when the inner layers of the star lack momentum to eject all the matter. Over a period of

minutes to hours, $0.1\text{--}5 M_{\odot}$ falls back onto the collapsed remnant, turning it into a black hole and establishing an accretion disk. GWs may be emitted by disk turbulence and disk instabilities that may lead to clumping or disk fragmentation [29,30]. The GW signal expected from accretion disk fragmentation would last $\mathcal{O}(10\text{--}100)$ s with a characteristic strain $h \sim 10^{-22}$ at 100 Hz for a source at 100 Mpc [29]. When the core collapse explosion is successful, a magnetar is formed. Convective currents and dynamical and secular nonaxisymmetric rotational instabilities in the proto-magnetar develop and may emit GWs [31]. In both scenarios, a GRB jet is launched either thanks to magnetohydrodynamical processes and neutrino pair annihilation powered by accretion or by the high Lorentz factor outflow that follows the birth of the proto-magnetar.

When a magnetar is formed, gravitational-wave emission from viscosity-driven “spin-flip” instability may last hours to days, with a detection horizon of 3–4 Mpc for Advanced LIGO/Advanced Virgo detectors and unmodeled searches [32–34].

The merger of two neutron stars will form a hot super-massive neutron star; depending on the component masses, the centrifugal forces induced by differential rotation and the stiffness of the nuclear equation of state may allow it to survive for hundreds of milliseconds before collapsing to a black hole or forming a massive neutron star [35,36]. It is very likely that the remnant is surrounded by an accretion disk that may endeavor instabilities like in the collapsar scenario. If the newly formed neutron star survives more than a few seconds, it could emit long-lived GWs through magnetic field-induced ellipticity [37,38] or r -mode instabilities [39], although the precise amplitude of such signals remains unclear. So far, no postmerger GW signals have been detected for any of the binary neutron star mergers found in LIGO and Virgo data [40–42].

Isolated neutron stars are another potential source of long-duration GW signals. Sudden speed-ups of the rotation of pulsars observed in radio and x-ray data are followed by a period of relaxation (weeks long) during which the pulsar slows down. GWs may be emitted during this period, but the amplitude is expected to be low as the rotational energy changes remain below 10^{43} erg [43–46]. Seismic phenomena in the crust of magnetars are thought to be at the origin of soft gamma repeaters and anomalous x-ray pulsars. Soft gamma repeaters’ giant flares are associated with huge emission of electromagnetic energy, up to 10^{46} erg, followed by long-duration quasiperiodic oscillations, which may be associated with GW emission over the same timescale [47–49]. The recent observation of GRB 200415a, suggesting that a magnetar giant flare may be a distinct class of short GRB, with a substantially higher volumetric rate than compact object mergers [50], is reinforcing the interest for magnetar giant flare events in nearby galaxies.

The diversity of long transient expected GW waveforms has led to the development of algorithms that do not rely on a

signal model. Coherent wave burst [13,51] and X-Pipeline [15], used for short-duration searches, have been adapted to search for transients with duration up to a few hundred seconds, while the Stochastic Transient Analysis Multi-detector Pipeline (STAMP) excess cross-power algorithm [18] has been developed to target specifically long and very long transient signals lasting up to several weeks. It has been used to search for long-duration GW transients associated with GRBs [52], for postmerger GW signals associated with GW170817 [40,42], and adapted to perform an all-sky/all-time search for long-duration GW transients in LIGO and Virgo data [53–55].

An enhanced version of the STAMP algorithm is presented in this article. It is a complete rewrite in PYTHON of the all-sky/all-time STAMP-AS pipeline that was built using the STAMP algorithm library written in MATLAB. As such, it has been optimized to search for GW signals of duration in the range $10\text{--}10^3$ s in a large dataset at a lesser computing cost than STAMP. It especially implements a hierarchical strategy, similar to the algorithm proposed in [19] to select the most interesting periods of the data without losing detection efficiency.

This paper is organized as follows. In Sec. II, we present the formalism of the analysis and the methods used to generate candidate events in the framework of a two-detector search. We describe the implementation of the pipeline and the methods used for background and efficiency estimation in Sec. III. Section IV summarizes the performance of the pipeline over simulated Gaussian noise and real data from the second LIGO-Virgo observation campaign (O2). Finally, we summarize those results in Sec. V and propose several improvements to increase the pipeline sensitivity in the future.

II. OVERVIEW OF A CROSS-CORRELATION GW TRANSIENT SEARCH ALGORITHM

A. Definitions and conventions

We are considering a network of GW detectors whose strain data time series $s(t) = n(t) + h(t)$ is a linear sum of independent detector noise $n(t)$ and the detector’s response to a GW strain amplitude given by $h(t)$. The detector noise is itself the sum of random noise and non-Gaussian noise transients, or “glitches.” The GW signal is assumed to be described by two polarization modes $h_+(t)$ and $h_{\times}(t)$ and originates from a pointlike source whose sky position is given by the right ascension and declination (α, δ) . We define $\hat{\Omega}$ as the direction to the source and $\tilde{h}(f)$ as the Fourier transform of any $h(t)$ time series. The detector’s response to a GW strain is the linear combination of the two polarizations weighted by the detector antenna factors $h(t) = F^+(t, \hat{\Omega}) \times h_+(t) + F^{\times}(t, \hat{\Omega}) \times h_{\times}(t)$. We consider an interval of duration T of GW strain data that are discrete measurements sampled at f_s . In the following, the variable t ; refers to the time segment start time.

The STAMP algorithm is an extension of the radiometer method developed to detect pointlike sources of stochastic background GWs [56]. To estimate the GW strain power spectrum of a transient signal, excess power is searched in frequency-time maps (ft maps) formed by cross-correlating the data of two spatially separated gravitational-wave detectors I and J . Following [18], an estimator of the GW power in a single ft pixel is given by

$$\hat{Y}(t; f, \hat{\Omega}) \equiv \text{Re}[Q_{IJ}(t; f, \hat{\Omega}) \tilde{s}_I^*(t; f) \tilde{s}_J(t; f)], \quad (1)$$

where

$$Q_{IJ}(t; f, \hat{\Omega}) = \frac{1}{\epsilon_{IJ}(t; \hat{\Omega})} e^{2\pi i f \hat{\Omega} \cdot \Delta \vec{x}_{IJ}/c} \quad (2)$$

is a filter function that takes into account the arrival time delay of the signal in the two detectors, whose distance is given by $\Delta \vec{x}_{IJ}$ and the pair efficiency

$$\epsilon_{IJ}(t; \hat{\Omega}) \equiv \frac{1}{2} (F_I^+(t; \hat{\Omega}) F_J^+(t; \hat{\Omega}) + F_I^\times(t; \hat{\Omega}) F_J^\times(t; \hat{\Omega})), \quad (3)$$

which weights the GW strain cross-power according to the alignment of the detectors. To normalize the cross-correlation, we compute the variance of \hat{Y} for which an estimator is

$$\hat{\sigma}_Y^2(t; f, \hat{\Omega}) = |Q_{IJ}(t; f, \hat{\Omega})|^2 P_I(t; f) P_J(t; f), \quad (4)$$

where $P_I(t; f)$ is the noise one-sided autopower spectrum. We then define the signal-to-noise ratio $\text{SNR}(t; f, \hat{\Omega})$ for a single pixel

$$\begin{aligned} \text{SNR}(t; f, \hat{\Omega}) &\equiv \frac{\hat{Y}(t; f, \hat{\Omega})}{\hat{\sigma}_Y(t; f, \hat{\Omega})} \\ &= \text{Re} \left[\frac{\tilde{s}_I^*(t; f) \tilde{s}_J(t; f)}{\sqrt{P_I(t; f) P_J(t; f)}} e^{2\pi i f \hat{\Omega} \cdot \Delta \vec{x}_{IJ}/c} \right]. \end{aligned} \quad (5)$$

$\text{SNR}(t; f, \hat{\Omega})$ depends only on the single-detector whitened statistic

$$\tilde{y}_I(t; f) \equiv \frac{\tilde{s}_I(t; f)}{\sqrt{P_I(t; f)}} \quad (6)$$

and the time delay $\tau \equiv \hat{\Omega} \cdot \Delta \vec{x}_{IJ}/c$ of the signal in the two detectors.

In the context of an all-sky search, the source direction $\hat{\Omega}$, and therefore τ , are unknown. An error in the time delay induces a dephasing in the computation of $\hat{Y}(t; f, \hat{\Omega})$ that can cause an underestimation of the SNR of coherent signals. A solution is to span all sky positions $\hat{\Omega}$ and retain the one that gives the largest SNR. That was the strategy implemented in STAMP-AS used to search for long-duration transient GW

signals in initial LIGO data [53,57] and advanced LIGO data [54,55]. However, the computational time required to process numerous sky positions was a limitation of the pipeline. Additionally, background estimation requires repeating, a large number of times, the same coherent cross-correlation of the data streams for each sky position tested using complete ft maps, while a large fraction of the pixels do not contain relevant information. As a consequence, the amount of simulated background was restricted to ~ 100 y, and the number of sky positions tested was limited to a few. All these suboptimal features resulted in a loss of sensitivity of $\sim 10\%$ – 20% [53].

The PySTAMPAS pipeline addresses these limitations by implementing the hierarchical approach proposed in [19], which consists of first identifying the most interesting clusters of pixels in single-detector autopower ft maps. In a second stage, a coherent detection statistic is computed using only the pixels that have been selected in the first stage. The computationally intensive calculations are carried out only once, allowing rapid background estimation without sacrificing the sensitivity gained by the use of coherence and spanning the whole sky positions. The gain in computational performance has also allowed the introduction of the use of several time-frequency resolutions to gain sensitivity to GW signals that may have time-varying frequency evolution. In the following sections, we describe the different computations that are performed at each stage.

B. Single-detector stage

1. Single-detector ft map

The simplest time-frequency representation of the GW time series $s_I(t)$ is a spectrogram obtained using one-sided Fourier transform of short segments of duration Δt . The short segments are first Hann windowed and overlap by 50% with each other, such that the pixel resolution is, respectively, $(\Delta t/2) \times (1/\Delta t)$ in time and frequency—the factor $\frac{1}{2}$ comes from the 50% overlap between short segments.

The spectrograms are whitened by the one-sided power spectral density of each segment $P_I(t; f)$. Two methods to compute the autopower have been implemented. The first one, inherited from STAMP, takes the average of $|\tilde{s}_I(t; f)|^2$ over time-neighboring pixels in a similar way to Welch's method. The other one considers the median over the frequency-neighboring pixels. The pros and cons of the two methods are discussed in Sec. IV A 1. For each time-frequency resolution, ft maps of the whitened statistic $\tilde{y}_I(t; f)$ are built.

The duration Δt of the Fourier transformed segments is an arbitrary choice that depends of the type of signal searched. Long-duration GW searches generally use Fourier transformed segments of duration $\simeq 1$ s, which are suited to reconstruct signals lasting $\sim 10^1$ – 10^3 s. However, when the frequency evolution of the signal is changing with time, parts of it can be better reconstructed using different resolutions.

In order to improve signal reconstruction as demonstrated by coherent wave burst [13], we opt for a multiresolution approach which consists in building several ft maps of different resolutions and combine them into a single multi-resolution ft map.

2. Clustering

The long-duration GW signal signature in ft maps appears as a cluster of pixels that a pattern recognition algorithm must be able to reconstruct without assuming a model. A year-long dataset is typically used. The unknown morphology assumption leads us to consider a seed-based clustering algorithm. The principle is to group high-energy pixels together by proximity, without imposing any preferred morphology for the cluster. For PySTAMPAS, we have adapted the BURSTEGARD algorithm, developed for STAMP [57] to multiresolution ft maps.

We consider all pixels $\tilde{y}_I(t; f)$ from every map with individual resolution $\Delta t_i \times \Delta f_i$. Pixels for which $|\tilde{y}_I(t; f)|$ exceeds a given threshold are kept to form a set of pixels for which we keep the time and frequency of the bottom left corner, Δt_i , Δf_i , and $\tilde{y}_I(t; f)$. The clustering algorithm starts with a seed pixel, the first pixel in the list, as the order does not matter. All pixels that are above threshold and within a given distance (in time and frequency) of the seed become part of the same cluster, whatever their resolution. Each new pixel added to the cluster then becomes the seed pixel and the same process is repeated recursively until no more pixels can be added. The next remaining unclustered pixel becomes the seed of the next cluster and the operations are applied again until all isolated pixels have been clustered. To eliminate clusters composed of only a few pixels, we select clusters that have a user-determined minimal number of pixels. The different parameters of the clustering (pixel threshold, radius, and minimal number of pixels per cluster) are free parameters that can be tuned considering that the number of operations is proportional to $\mathcal{O}(N \log(N))$, where N is the number of pixels above threshold. As the GW signal energy is spread over many pixels, the threshold on $|\tilde{y}_I(t; f)|$ should not be too selective, and the distance between two pixels should not be too strict as well.

C. Coherent analysis

Considering all possible detector pairs, clusters from one detector are cross-correlated with the other detector's pixels. At this stage, the clusters can be composed of pixels of different time-frequency resolution. To be able to cross-correlate pixels of different time-frequency resolution, we define virtual pixels that have resolution $\min \Delta t_i \times \min \Delta f_i$. Each of these pixels is assigned a value that is the largest $|\tilde{y}_I(t; f)|$ value of all pixels that overlap the virtual pixel. As $\Delta t_i \times \Delta f_i$ is constant over all resolution maps, the virtual pixels assigned values have the same weight. The same construction of virtual pixels is performed for the pixels of the other detector's ft map.

The cross-correlation SNR given by Eq. (5) is then computed considering the virtual pixels. As already mentioned, pixel SNR depends only on the time delay between two detectors $\tau = \hat{\Omega} \cdot \Delta \vec{x}_{IJ}/c$. In an all-sky search, the direction to the source is not known *a priori*, and an error on the time delay can cause one to underestimate the SNR of coherent signals. A solution is to span the time delay parameter space over all possible values for a given pair. The maximal SNR loss due to an error of $d\vec{\Omega}$ corresponding to $d\tau$ is

$$\text{SNR}(t; f, \hat{\Omega} + d\vec{\Omega}) = \cos(2\pi f d\tau) \text{SNR}(t; f, \hat{\Omega}). \quad (7)$$

The time delay bin size $d\tau$ is determined such that the maximal SNR loss is lower than $\epsilon \in [0, 1]$ for the maximal frequency considered in each cluster;

$$d\tau = \frac{\arccos(\epsilon)}{2\pi f_{\max}} \quad (8)$$

with f_{\max} being the maximal frequency of all pixels of the cluster, which can be much lower than the maximal frequency of the search, reducing the number of time delays to test. In the most general case, we would need to test N_τ time delay values between 0 and $\Delta x_{IJ}/c$ by steps of $d\tau$ to recover a signal accurately. However, because we are considering a phase factor, a degeneracy appears: $\epsilon = \cos(2\pi f d\tau) = \cos(2\pi f (d\tau + 1/f))$. As a consequence, for a pixel at frequency f , it is sufficient to test time delays in the interval $[0, 1/f]$ instead of $[0, \Delta x_{IJ}/c]$ to get the correct phase factor. In the case of a broadband cluster with pixel frequencies between f_{\min} and f_{\max} , this interval is the largest for $f = f_{\min}$, so we need to test time delay values between 0 and $1/f_{\min}$ by steps of $d\tau$. Finally, the number of time delays to test to recover a signal with an accuracy ϵ is

$$N_\tau = \frac{2\pi}{\arccos(\epsilon)} \frac{f_{\max}}{f_{\min}}. \quad (9)$$

In the end, N_τ remains rather small (less than a few hundred), especially compared to the thousands of sky positions that need to be tested using a regular grid in sky coordinates α , $\cos(\delta)$. Since the computations are done only over the small subset of pixels that constitute the cluster, it is possible to test hundreds of time delays in a reasonable time and therefore limit the loss of SNR to $\epsilon = 0.95$ regardless of the signal morphology, as shown in Sec. IV A 2.

The time delay τ_0 that maximizes the sum of all pixel SNRs provides a detection statistic that reflects the significance of the cluster

$$\text{SNR}_\Gamma \equiv \sum_{(t, f) \in \Gamma} \text{SNR}(t; f, \tau_0). \quad (10)$$

This detection statistic is used to test the hypothesis of a GW signal or the null hypothesis. However, hierarchical processing methods such as PySTAMPAS applied on real GW data tend to bias the selection of triggers because of the presence of noise outliers in one detector. When combined with noise fluctuation in the second detector, such triggers may have large SNR_Γ values despite being incoherent. In order to mitigate this effect, we estimate the residual noise energy that is left in one detector's data after subtracting the sum of $|\tilde{y}_I(t; f)|^2$ over all pixels belonging to the cluster. We define the quantity

$$E_I^{\text{res}} \equiv \sum_{(t,f) \in \Gamma} |\text{SNR}(t; f) - |\tilde{y}_I(t; f)|^2|. \quad (11)$$

For a coherent GW signal, recovered with the right time delay, this residual energy is expected to be much smaller than both SNR_Γ and the autopower energy E_I

$$E_I \equiv \sum_{(t,f) \in \Gamma} |\tilde{y}_I(t; f)|^2. \quad (12)$$

On the contrary, for a cluster due to a noise outlier in one of the detectors, E_I^{res} may become large in the second detector. We can then define a second discriminant variable in addition to SNR_Γ ,

$$\Sigma_{\text{res}} \equiv \sum_I E_I^{\text{res}} / E_I. \quad (13)$$

Finally, we combine these two variables into a single-detection statistic Λ defined by

$$\Lambda \equiv \frac{\text{SNR}_\Gamma}{\text{SNR}_\Gamma + \Sigma_{\text{res}}}. \quad (14)$$

Λ should tend to 1 in the presence of a coherent GW signal and take $\ll 1$ values in case of noise outliers. For convenience, we define

$$p_\Lambda \equiv -\log(|1 - \Lambda|) \quad (15)$$

such that the detection statistic increases with the significance of the trigger. Note that p_Λ is not the only possibility to combine SNR_Γ and Σ_{res} . We show in Sec. IV C that p_Λ is robust to loud noise triggers using a sample of real data from GW detectors, but other combinations may be relevant depending on the distribution of background noise.

III. DETAILS OF THE PIPELINE IMPLEMENTATION

In the following sections, we describe the implementation of PySTAMPAS in the case of a two-detector network, and we propose a generalization to the case of a network of more

than two detectors. In practice, the pipeline is implemented using PYTHON 3 and relies on the GWpy package [58].

A. Data conditioning

The GW detectors' data streams are first searched individually to reveal clusters of energy that may contain coherent GW signals. Real GW detector data are available as an ensemble of time series of different lengths. For a given pair of GW detectors, only coincident times are analyzed. This reduces the dataset to a list of coincident segments of time. For each of the coincident segments, the data are split into windows of duration T_{win} that overlap by 50%. The duration of the data window is a free parameter that can be adjusted to the typical duration of the GW signal that is being investigated. In this study, we use $T_{\text{win}} \simeq 500$ s, as it is done in previous long-duration searches [54,55]. STAMP was originally designed to search for signals with duration up to several weeks [18]. Although there is no fundamental limitation to extending PySTAMPAS to longer signals, we limit ourselves to signals in the range $10\text{--}10^3$ s in this paper. Working with very large windows leads to dropping up to $T_{\text{win}}/2$ s of data at the end of each coincident segment and increases the computing cost of clustering.

The data are first high-pass filtered to suppress energy outside the analysis frequency bandwidth whose lower boundary is adapted to the GW detectors' noise spectrum of each dataset. Real GW detector data often contain non-Gaussian, short-duration spikes (glitches) [59,60]. When the magnitude of the glitch is large, an excess of energy is present in the ft maps and generates single-detector clusters with very large energy (orders of magnitude larger than what a real GW signal would generate). The coherent step is usually not able to eliminate them completely and a better strategy consists in gating the data time series before computing the ft maps. PySTAMPAS mitigates the effect of the loud glitches by applying a Planck window on the $h_I(t)$ samples that exceeds a fixed threshold.¹ This threshold is a free parameter that should be tuned for each analysis in order to remove most of the glitches without penalizing signal recovery. After this preprocessing step, ft maps of $\tilde{y}_I(t; f)$ are built.

As shown in all GW detectors' noise spectra [61,62], real GW data contain many spectral artifacts corresponding to mechanical resonances, power lines, and pump or fanlike machines surrounding the detectors [63,64]. Most of these spectral lines are of low amplitude and relatively constant over time, while some have a time-varying frequency. These artifacts can generate false long-duration tracks in ft maps. To attenuate the impact of lines, we set to zero ("frequency notch") $\tilde{y}_I(t; f)$ pixels corresponding to a list of frequencies that are constructed following two steps:

¹These samples are found by the SciPy function `find_peaks`.

(1) For each ft map built, we compute the median value $\bar{y}(f)$ over time of $|\bar{y}(t; f)|$. Frequencies for which $\bar{y}(f)$ is higher than a given threshold are flagged.

(2) If a frequency is flagged in more than a given fraction of the total ft maps, it is added to the list.

This last condition reduces the risk that a monochromatic GW transient of duration $\lesssim T_{win}$ is mistaken for an instrumental line and notched. One should note, however, that a very long monochromatic transient signal (on the order of weeks or months) could still be flagged if it is spread over a fraction of the total ft maps higher than the threshold chosen. If a signal crosses a notched frequency, it may be divided into several parts that will be reconstructed by BURSTEGARD as separate clusters, reducing the significance of the signal. To reconnect these parts, we implement the FINDTRACK algorithm [57]. If the minimal distance between the corners of two clusters is smaller than a given radius, these clusters are connected and treated as one single cluster.

B. Coincident search

The coincident search is the proper analysis during which true GW signals are searched in the data. The individual detector's ft maps are searched for clusters of excess energy following the procedure described in Sec. II B. Two lists of clusters are extracted from a pair of detectors. They are saved along with the ft maps to be analyzed in the coherent stage following the procedure described in Sec. II C. The pipeline produces a list of coherent triggers that are ranked according to p_Λ .

C. Background estimation

In order to assess the significance of triggers found in coincidence, one has to estimate the accidental rate of noise triggers caused by instrumental and environmental effects. Like almost all GW transient search pipelines, to encompass any particular effect in the data and augment the total volume of data, we use the time-slides technique to estimate our background [65]. One data stream is time shifted with respect to the other one by an amount of time greater than the light traveling time between the detectors. Assuming the number of detectable GW signals is small, this assures that the cross-correlated data do not contain a coherent GW signal. In the meantime, non-Gaussian and nonstationary features of the detectors' noise are preserved. By repeating the analysis for many time-shift values, one simulates multiple instances of the noise.

In PySTAMPAS, time shifts are performed considering data streams split over N_{win} windows that are time ordered on a circle. Data are shifted by a multiple of windows (lags) and for each lag by a multiple of Δt_{max} the maximal time resolution (minilags). For example, considering only lags, at the n th lag, clusters from detector I that have been extracted in window i are matched with detector J data from window $(i + n)$. With this technique, the maximal number of time shifts is

$$(N_{win} - 1) \times \frac{T_{win}}{\Delta t_{max}}. \quad (16)$$

The total background lifetime simulated T_{bkg} is the number of time shifts performed times the duration of data available for a pair of detectors. The cumulative background trigger rate gives an estimation of the false-alarm rate (FAR) as a function of the detection statistic, which is used to rank the triggers.

D. Sensitivity studies

PySTAMPAS performs sensitivity studies by injecting simulated signals into the data. A simulated signal consists primarily of a “waveform,” which describes the two polarizations modes $h_+(t)$ and $h_\times(t)$ of a GW. Waveforms are stored in files in the form of two time series sampled at f_s , as well as metadata (duration, frequency range, physical model, etc.). A bank of waveforms with various properties is available to sample the rather large parameter space of long-duration transient GW signals with representative signal morphologies.

To compute the detector's response $h_I(t)$ to a given GW signal, one has to specify a waveform and the following parameters:

- (1) the time of arrival t_0 at the center of Earth;
- (2) the direction $\hat{\Omega}$ to the source;
- (3) the inclination and polarization angles (ι, ψ) that characterize the orientation of the source's reference frame with respect to Earth's equatorial frame;
- (4) a scaling amplitude factor α to modulate the strength of the signal.

Source frame GW polarizations are then rotated to be expressed in Earth's equatorial frame

$$\begin{aligned} h'_+(t) &= a_+ \cos 2\psi h_+(t) - a_\times \sin 2\psi h_\times(t), \\ h'_\times(t) &= a_+ \sin 2\psi h_+(t) + a_\times \cos 2\psi h_\times(t), \end{aligned} \quad (17)$$

where $a_+ \equiv \frac{1+\cos\iota^2}{2}$ and $a_\times \equiv \cos\iota^2$.² The polarizations are then time shifted by the delay of arrival between the detector's position \vec{r}_I and the center of Earth

$$\tau_I = \frac{\hat{\Omega} \cdot \vec{r}_I}{c} \quad (18)$$

and rescaled by the amplitude factor α such that, finally,

$$h_I(t) = \alpha [F_I^+(t; \hat{\Omega}) h'_+(t - \tau_I) + F_I^\times(t; \hat{\Omega}) h'_\times(t - \tau_I)], \quad (19)$$

where $F_I^+(t; \hat{\Omega})$ and $F_I^\times(t; \hat{\Omega})$ are the detector's sensitivity to $+$ and \times polarizations (expressed in Earth's equatorial frame) of a GW signal coming from direction $\hat{\Omega}$ at time t .

²The dependence of a_+ and a_\times on ι is correct for quadrupolar emission.

The computed response is resampled and interpolated to match with the detector's sampling, and the first and last seconds of the time series are tapered with a Hann window to avoid numerical artifacts when the signal starts or stops abruptly. The signal is injected in the data, which are then analyzed the same way as in a coincident search (restricted to the windows that overlap the injection to gain time). An injection is considered detected if the search produces a trigger within the time and frequency boundaries of the simulated signal and with a detection statistic p_Λ larger than a given threshold.

To estimate the detection sensitivity to a given waveform at a given amplitude, a statistically significant number of injections are performed with random starting time, sky position, polarization angle, and cosine of the inclination. Starting times are selected in such a way that they always fall within a coincident data segment. By computing the fraction of recovered injections for different signal amplitudes, it is possible to characterize the detection efficiency as a function of signal's strength, which is usually expressed with the root-sum-squared amplitude h_{rss} given by

$$h_{\text{rss}} \equiv \sqrt{\int (h_+^2(t) + h_\times^2(t)) dt}. \quad (20)$$

E. Generalization to a network of detectors

The search algorithm can be generalized in a straightforward manner to a network of N detectors (I, J, K, \dots) , constituting $p(N) = N(N-1)/2$ pairs. For a given time-frequency pixel $(t; f)$ and sky direction $\hat{\Omega}$, we define the total coherent SNR as the sum of cross-correlated SNRs from all detector pairs,

$$\text{SNR}(t; f, \hat{\Omega}) = \sum_{I=1}^N \sum_{J>I} \text{SNR}^{IJ}(t; f, \hat{\Omega}), \quad (21)$$

with $\text{SNR}^{IJ}(t; f, \hat{\Omega})$ the coherent SNR computed from Eq. (5) corresponding to the pair IJ . This allows us to generalize the definitions of E_I^{res} and SNR_Γ for a cluster of pixels Γ ,

$$E_I^{\text{res}} \equiv \sum_{(t,f) \in \Gamma} |\text{SNR}(t; f)/p(N) - \tilde{y}_I(t; f)|^2, \quad (22)$$

$$\text{SNR}_\Gamma \equiv \sum_{(t,f) \in \Gamma} \text{SNR}(t; f), \quad (23)$$

and finally the definition of Λ remains unchanged,

$$\Lambda \equiv \frac{\text{SNR}_\Gamma}{\text{SNR}_\Gamma + \sum_I \frac{E_I^{\text{res}}}{E_I}}. \quad (24)$$

The pipeline's implementation does not fundamentally change with $N \geq 3$ detectors. The clustering step is

performed independently over each individual detector's ft maps, following the hierarchical method of [19]. For each cluster, cross-correlation is computed for the $p(N)$ pairs to compute its ranking statistic p_Λ . However, as the degeneracy between sky direction and time delay between detectors is broken for $N \geq 3$, it is necessary in this case to test all sky positions by choosing uniformly α and $\cos(\delta)$ and select that position that maximizes SNR_Γ . Therefore, a full-scale study of the pipeline's performances over a network of three or four detectors will be necessary in the future, considering realistic detectors' sensitivity curves.

IV. PERFORMANCES AND COMPARISONS

To test the pipeline and demonstrate its performance, we consider 13 waveforms commonly used in long-duration searches [54,55] whose main characteristics are listed in Table I. Most of the waveforms are based on astrophysical models and fall into three categories: eccentric inspiral-merger-ringdown nonspinning compact binary coalescence (ECBC) [66], broadband chirps from innermost stable circular orbit waves around rotating black holes (ISCOchirp) [67,68], and accretion disk instability models (ADI) [30]. We include two *ad hoc* waveforms to better cover the parameter space; a 250-s-long sine Gaussian signal (SG-C) with a decay time of 50 s and a 20-s-long band-limited white noise burst (WNB-A). These signals of different morphology cover the time-frequency space with durations within 9–290 s and frequencies in the 10–2048 Hz range. In the following, we consider the case of a two-detector search to compare performance with STAMP-AS. If not stated differently, we are using simulated Gaussian noise following LIGO's best sensitivity during the second observing run (O2) [69] to simulate the data from the two LIGO detectors at Hanford (H1) [70] and Livingston (L1) [71].

A. Signal reconstruction

We investigate the effects of several parameters of the pipeline on the detection capability and the signal reconstruction in order to find a set of parameters that maximize the detection of a wide range of different morphology signals, while keeping the computational costs affordable.

1. Power spectral density estimation

The accuracy of the noise power spectral density (PSD) estimation plays a central role in reconstructing GW signals efficiently. Yet, this task is complicated in the case of GW detectors, as the noise contains non-Gaussian and nonstationary features such as glitches, spectral lines, and slow drifts of the noise amplitude.

Consider a detector's strain time series given by $s_I(t) = h_I(t) + n_I(t)$, where $h_I(t)$ is a deterministic GW signal and $n_I(t)$ is random noise. A good estimator of the

TABLE I. Name, parameters, duration, frequency range, and spectral morphology of waveforms used to characterize PySTAMPAS. M_i is the component compact object mass; ecc is the eccentricity of the binary orbit at 10 Hz; M_{BH} and a_{BH} are the mass and normalized spin of the black hole.

Waveform	Parameters	Duration (s)	Frequency (Hz)	Morphology
ECBC-A	$M_1 = 1.4 M_\odot, M_2 = 1.4 M_\odot, ecc = 0.2$	291	10–250	Chirp
ECBC-B	$M_1 = 1.4 M_\odot, M_2 = 1.4 M_\odot, ecc = 0.4$	178	10–275	...
ECBC-C	$M_1 = 1.4 M_\odot, M_2 = 1.4 M_\odot, ecc = 0.6$	64	10–350	...
ECBC-D	$M_1 = 3.0 M_\odot, M_2 = 3.0 M_\odot, ecc = 0.2$	81	10–180	...
ECBC-E	$M_1 = 3.0 M_\odot, M_2 = 3.0 M_\odot, ecc = 0.4$	49	10–200	...
ECBC-F	$M_1 = 3.0 M_\odot, M_2 = 3.0 M_\odot, ecc = 0.6$	15	10–200	...
ISCOchirp-A	$m_{\text{BH}} = 5.0 M_\odot$	237	1049–2048	Broadband chirp down
ISCOchirp-B	$m_{\text{BH}} = 10.0 M_\odot$	237	705–2048	...
ISCOchirp-C	$m_{\text{BH}} = 20.0 M_\odot$	236	196–1545	...
ADI-A	$m_{\text{BH}} = 5.0 M_\odot, a_{\text{BH}} = 0.3$	35	135–166	Chirp down
ADI-B	$m_{\text{BH}} = 10.0 M_\odot, a_{\text{BH}} = 0.95$	9	110–209	...
SG-C		243	402–408	Monochromatic
WNB-A		20	50–400	Band-limited white noise

one-sided PSD of the noise is given by the squared modulus of its Fourier transform,

$$P_I(t; f) \equiv \langle |\tilde{n}_I(t; f)|^2 \rangle. \quad (25)$$

This value is not directly accessible because, in case of an unknown GW waveform, it is not possible to disentangle *a priori* signal from noise. One has to rely on the observable $|\tilde{s}_I(t; f)|^2$, which may contain a GW signal. Assuming signal and noise are not correlated,

$$\langle |\tilde{s}_I(t; f)|^2 \rangle = \langle |\tilde{h}_I(t; f)|^2 \rangle + P_I(t; f). \quad (26)$$

Therefore, an assumption over the nature of the signal $\tilde{h}_I(t; f)$ must be made in order to build an unbiased estimator of $P_I(t; f)$. PySTAMPAS implements two methods to estimate the PSD that are suited for different signal morphologies.

The first method consists of taking the average of $|\tilde{s}_I(t; f)|^2$ over n_t symmetrically chosen neighboring Fourier transformed segments. The underlying assumptions are that (1) the noise is stationary over the time window considered, and (2) no signal is present in the adjacent pixels. As discussed above, (1) is often wrong because of the presence of short glitches in the data, which are therefore not factored in the PSD and appear as signal. Conversely, (2) is wrong when a monochromatic or quasimonochromatic signal is present in the data, leading these to be mistakenly included in the PSD. Degraded sensitivity to monochromatic signals is a known weakness of STAMP [53].

To address these issues, we propose to estimate the PSD by taking the median of $|\tilde{s}_I(t; f)|^2$ over n_f adjacent frequency bins. The pros and cons of this method are opposite to the first one: short glitches are well taken into account and monochromatic signals are better reconstructed.

However, signals whose frequency evolution is rapid tend to be less well reconstructed. In the case of noise only, both methods provide similar PSD estimates, except that spectral narrow features are better reconstructed with the method averaging the neighboring time segments' pixels, as shown in Fig. 1. We use the median, as it is more robust than the average to extreme values. Because of instrumental lines, it is likely that one of the neighboring frequency bins has pixels with a very high value of $s(t; f)$, which would spoil the PSD estimation.

The effect of the PSD estimation on the signal reconstruction in PySTAMPAS is illustrated in Fig. 2. Two signals with very different spectral morphologies, a broadband ISCOchirp (ISCOchirp-C) [68] and a monochromatic sine Gaussian (SG-C), are injected in Gaussian noise. By taking the median over adjacent frequency bins (hereafter referred to as “frequency-median PSD”) instead of averaging over neighboring Fourier transformed segments (“time-average




FIG. 1. Estimation of the PSD for a 100-s-long segment of LIGO Hanford data from the O2 observing run using the two different methods implemented in PySTAMPAS. The squared modulus of the Fourier transform (averaged over ten independent realizations of the noise) is shown in blue for reference.

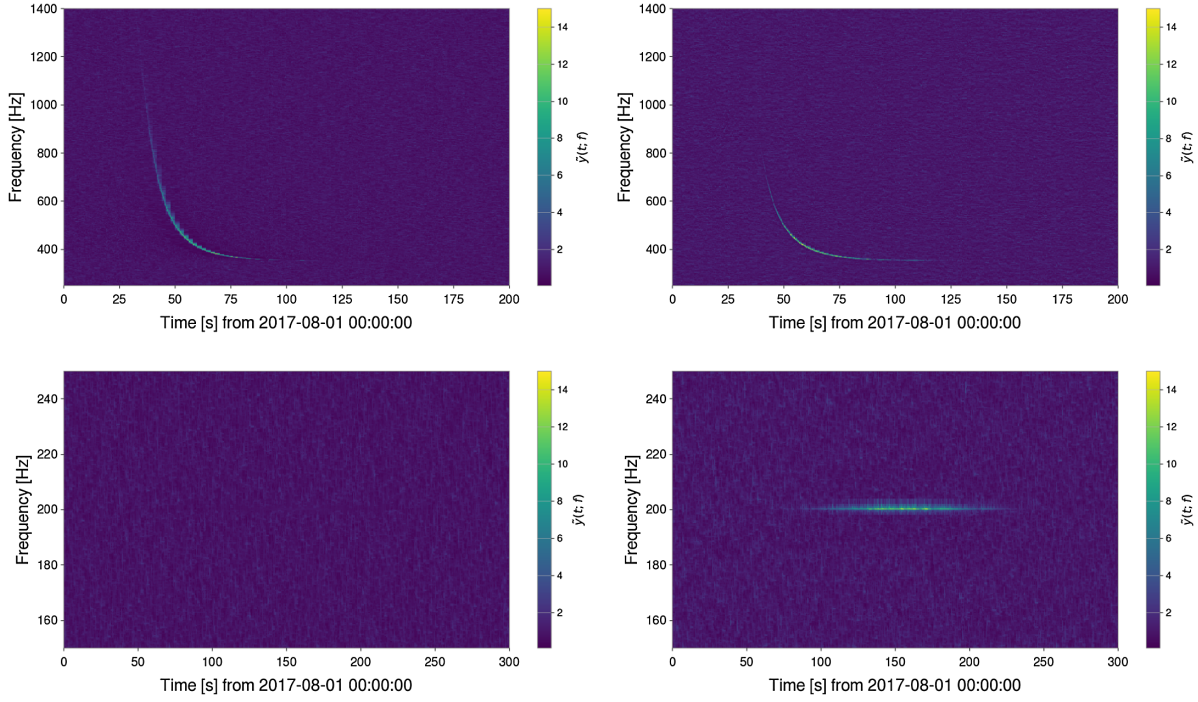


FIG. 2. Time-frequency maps of $|\tilde{y}(t; f)|$ for two injected signals (top: ISCOchirp; bottom: sine Gaussian) realized with the two different PSD methods [left: average over $n = 32$ adjacent time bins (time average); right: moving median over $n = 20$ frequency bins (frequency median)].

PSD”), the sine Gaussian signal is better reconstructed, but the fast frequency evolving part of the ISCOchirp is blurred out. The optimal choice of a PSD estimation method depends on the type of signals targeted and the characteristics of the noise, especially spectral lines and/or nonstationary features. Another way to restore the sensitivity to monochromatic triggers would be to consider a very long ($\sim 10^3$ s) time period to estimate the PSD in the case of the time adjacent pixels method. However, noise from GW detectors tends to become nonstationary over such time intervals at low frequencies (below ~ 100 Hz) [63].

2. Source sky location determination

The number of sky locations tested in the coherent step (which reduces to a single time delay parameter in the case of a two-detector network) is currently a limiting factor of all-sky searches and illustrates the necessary trade-off between detection sensitivity and computational cost [53]. The hierarchical processing implemented in PySTAMPAS allows for scanning many positions at a low cost. In Sec. II C, we have seen that N_τ , the number of time delays between detectors to be tested, depends on the ratio between the maximal and the minimal frequency of the trigger. Here, we investigate empirically the pipeline sensitivity loss as a function of the number of time delays for different waveform families.

Signal waveforms are injected coherently into Gaussian noise, simulating data from LIGO Hanford and LIGO

Livingston, from a given sky direction $\hat{\Omega}_0$, and are recovered by PySTAMPAS. We vary the number of time delays and keep the maximal SNR_Γ obtained, which is compared to $\text{SNR}_\Gamma(\hat{\Omega}_0)$, the SNR value corresponding to the true source position $\hat{\Omega}_0$.

The ratio SNR_Γ to $\text{SNR}_\Gamma(\hat{\Omega}_0)$ as a function of the number of time delays between detectors is shown in Fig. 3 for sine Gaussians of different central frequency and for a selection of waveforms of different morphology/durations. We compare the number N_τ of delays tested to get $\epsilon = 0.95$ to the theoretical prediction from Eq. (9) given in Table II. We see that the optimal value of N_τ does not depend on the signal frequency, but mainly on its frequency range $f_{\text{max}}/f_{\text{min}}$. Monochromatic sine Gaussians are recovered equally rapidly no matter their frequency and faster than signals of broader band. The empirical values are overall lower than the theoretical ones. This discrepancy comes from the fact that the clustering algorithm does not always reconstruct the entirety of the waveform, leading to a lower effective value of $f_{\text{max}}/f_{\text{min}}$. To optimize the detection efficiency while keeping the number of tested sky positions minimal, we fix ϵ such that the maximal SNR loss parameter is 5% and N_τ is determined for each cluster following Eq. (9).

3. Multiresolution and clustering

The energy of long-duration GW signals is spread over a potentially large number of pixels. This would mean it is necessary to lower the threshold on the individual pixel’s




FIG. 3. Ratio between the recovered SNR_Γ and $\text{SNR}_\Gamma(\hat{\Omega}_0)$ as a function of the number of time delays between the LIGO Hanford and LIGO Livingston detectors for different waveforms. For each value of N_τ , 50 injections have been performed. Crosses represent the theoretical number of time delays to test to reach a ratio of 0.95, computed from Eq. (9).

energy $|\tilde{y}_l(t; f)|$ and rely on the clustering algorithm to group all pixels belonging to the cluster. Clustering a large number N of pixels is computationally expensive since BURSTEGARD's time complexity is $\mathcal{O}(N \log N)$. However, because of the hierarchical implementation, that step is computed only once per ft map and is therefore no longer a bottleneck for analyzing long periods of data. Yet, the risk is to include pixels due to noise fluctuations and generate clusters that are only composed of noise pixels. By increasing the

TABLE II. Theoretical minimal number N_τ of time delays between two detectors (here LIGO Hanford and LIGO Livingston) to be considered for each waveform in order to recover the coherent signal SNR with an accuracy larger than 0.95. N_τ depends on the frequency ratio f_{\max}/f_{\min} of the signal considered as given in Eq. (9).

Waveform	f_{\max}/f_{\min}	N_τ
ADI-A	1.2	24
ADI-B	1.9	37
ISCOchirp-A	1.9	37
ISCOchirp-B	2.8	56
ISCOchirp-C	7.9	155
ECBC-A	12.5	247
ECBC-B	13.8	272
ECBC-C	17.5	346
ECBC-D	9	178
ECBC-E	10	197
ECBC-F	10	197
SG-C	1	20
WNB-A	1.2	23

TABLE III. Relative increase on the distance at 50% detection efficiency between a single-resolution approach (pixels of $1 \text{ s} \times 1 \text{ Hz}$) and the multiresolution approach implemented in PySTAMPAS (four different resolutions from $0.5 \text{ s} \times 2 \text{ Hz}$ to $4 \text{ s} \times 0.25 \text{ Hz}$), all other parameters being equal, for the astrophysical waveforms described in Table IV.

Waveform	Efficiency increase
ADI-A	+25%
ADI-B	+5%
ISCOchirp-A	+17%
ISCOchirp-B	+10%
ISCOchirp-C	+18%
ECBC-A	+23%
ECBC-B	+17%
ECBC-C	+40%
ECBC-D	+38%
ECBC-E	+41%
ECBC-F	+3%

minimal number of pixels per cluster, one can control the rate of noise clusters that are generated.

Another way to collect, as best as possible, all the energy in the ft maps is to process the data with a range of different time-frequency resolutions that match well all the different GW signal shapes. The choice of time-frequency resolutions depends on the waveform, but we have seen that, for the diversity of signals we are targeting, a limited number of time resolutions is enough to improve the detection efficiency of nonmonochromatic GW signals. Using a set of four resolutions ranging from $4 \text{ s} \times 0.25 \text{ Hz}$ to $0.5 \text{ s} \times 2 \text{ Hz}$, we report an efficiency increase by 5%–40% for the waveforms tested (at constant FAR), compared to $1 \text{ s} \times 1 \text{ Hz}$ pixels. We report the relative increase in detection efficiency for each astrophysical waveform in Table III.

It is not possible to perform a fine optimization of all PySTAMPAS parameters for a generic all-sky/all-time search because of the large parameter space, but we present in the next sections the pipeline performance for both Gaussian simulated noise and real GW data to detect long-duration GW signals using the set of parameters given in Table IV.

B. Test on simulated data

We carry out a study with simulated Gaussian noise to test the pipeline as a whole and evaluate its performance. First, we generate two sets of 14 days of stationary Gaussian GW noise following LIGO's O2 sensitivity to simulate the data from the two LIGO detectors at Hanford and Livingston. We analyze these data with PySTAMPAS, using parameters given in Table IV.

TABLE IV. PySTAMPAS parameter values used in the all-sky/all-time long-duration GW search with Advanced LIGO/Advanced Virgo data presented in this paper.

Parameters	Value
ft maps	
Window duration	512 s
Frequency range	20–2000 Hz
$\Delta t_i \times \Delta f_i$	$[4.0 \text{ s} \times 0.25 \text{ Hz} - 2.0 \text{ s} \times 0.5 \text{ Hz}]$ $[-1.0 \text{ s} \times 1.0 \text{ Hz} - 0.5 \text{ s} \times 2.0 \text{ Hz}]$
PSD estimation	
Time average	32 time bins
Frequency median	20 Hz
Clustering	
Pixel energy threshold	2.0
Clustering radius	$2 \text{ s} \times 2 \text{ Hz}$
Minimum pixels number	30
Coherent stage	
SNR loss $1 - \epsilon$	5%

Background triggers are generated following the method described in Sec. III C. We perform 128,000 time slides, simulating ~ 4900 y of background noise accounting for 34 days of CPU time on a dual-core modern processor. As a comparison, the previous version of STAMP-AS took 95 days of CPU time to perform 1000 time slides over the same data, meaning that PySTAMPAS is faster by at least one order of magnitude. In Fig. 4, showing the cumulative FAR as a function of p_Λ , the blue curves correspond to the distribution of simulated Gaussian noise triggers for the two PSD estimation methods; the shape of the two curves is similar, but the median-frequency PSD method produces ~ 60 more triggers than the time-average PSD. This has little effect on the pipeline sensitivity, as the tails of the p_Λ distribution are similar.

For each waveform described in Table I, we estimate the detection efficiency as a function of h_{rss} following the method described in Sec. III D. We fix a detection threshold corresponding to a FAR of $1/50 \text{ yr}^{-1}$ and determine the value $h_{\text{rss}}^{50\%}$ of h_{rss} for which 50% of the injections are recovered. To provide a comparison, we perform the same search with STAMP-AS over the same simulated Gaussian noise. We use the quantity $h_{\text{rss}}^{50\%}$ to estimate the detection efficiency of the search. It is inversely proportional to the typical detection range. In Fig. 5, we show the ratio of $h_{\text{rss}}^{50\%}$ between STAMP-AS and PySTAMPAS for each waveform and each PSD estimator.

For a majority of the waveforms tested, PySTAMPAS is more sensitive than STAMP-AS, up to a factor 2, with the exception of the ISCOchirp family for which detection efficiencies are worse by down a factor 0.8–1 in the best case with the time-average PSD. For this specific family, the single-detector clustering algorithm reconstructs low amplitude signals poorly because the energy is spread over too many pixels. Down to a certain amplitude, most pixels




FIG. 4. FAR obtained with data from LIGO O2 observing run versus the detection statistic p_Λ with frequency-median PSD (top) and time-average PSD (bottom). The blue curves represent the FAR obtained with Gaussian noise. FAR of triggers remaining after applying Rveto is shown by the green curve.

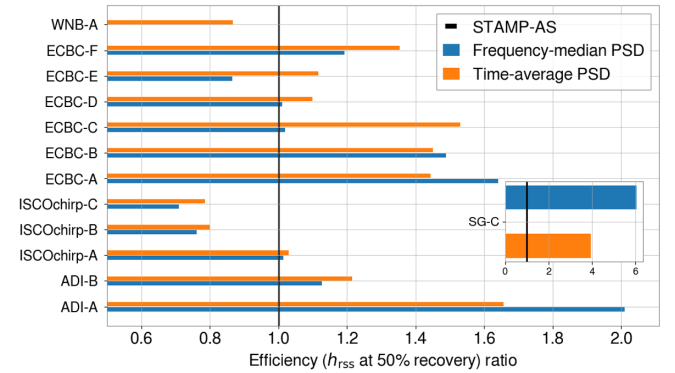


FIG. 5. Ratio between the h_{rss} at 50% detection efficiency obtained with STAMP-AS and with PySTAMPAS for a FAR = $1/50 \text{ yr}^{-1}$ for both PSD estimation methods. The white noise burst waveform WNB-A was not recovered at all using the frequency-median PSD. A ratio above 1 means that PySTAMPAS recovered the signal better than STAMP-AS.

fall below the clustering threshold and the signal is not reconstructed at all. A finer tuning of BURSTEGARD could be done to address this limitation, but this type of signal

would certainly be better reconstructed by seedless clustering algorithms. This also illustrates the difficulty of tuning the pipeline to maximize sensitivity to a wide variety of waveforms.

The *ad hoc* waveforms illustrate the most extreme cases. Detection efficiency is multiplied by ~ 6 for the monochromatic sine Gaussian signal when the PSD is computed over adjacent frequency bins as compared to STAMP-AS. On the other hand, the large band white noise burst is not recovered at all with this method and recovered almost equally well with the time-average PSD. We note that the sine Gaussian is also better recovered using the time-average PSD. This is due to the fact that we consider a wider time window to compute the PSD (32 pixels from each side instead of 8).

C. Tests on real data

Real data from GW detectors have non-Gaussian and nonstationary features that challenge pipelines. To understand the behavior of PySTAMPAS on real GW noise, we analyze LIGO data from the Advanced LIGO and Advanced Virgo O2 observing run downloaded from the Gravitational Wave Open Science Center [72,73]. The chosen period runs from August 1, 2017 00:00:00 UTC to August 15, 2017 00:00:00 UTC and contains 9.21 days of coincident data from H1 and L1. We keep the pipeline’s parameters given in Table IV, but switch on the spectral lines removal algorithm described in Sec. III A. About 5% of the total frequency bins are flagged as spectral lines and notched for each detector. As in the simulated data study, we consider both PSD estimation methods. Cumulative FAR distributions for O2 data are compared to simulated Gaussian noise FAR distributions in Fig. 4. For both PSD estimation methods, an excess of triggers is present compared to the simulated Gaussian noise distributions, meaning that the FAR of the search for a given value of p_Λ is higher than with Gaussian noise.

For the frequency-median PSD, the excess of triggers ($\sim 20\%$ more triggers in real data than in the Monte Carlo study with Gaussian noise) consist of long-duration (> 50 s), quasimonochromatic events that correspond to instrumental lines being punctually excited. These lines are too low amplitude and are not excited regularly enough to be flagged by the spectral lines removal algorithm. However, that excess becomes marginal for large value of p_Λ and thus does not affect the overall pipeline sensitivity for this set of data.

Using the time-average PSD method, the excess of triggers compared to Gaussian noise is much larger, by at least 1.5 orders of magnitude. It is dominated by short glitches with frequencies between 20 and 100 Hz that have passed the gating procedure. They generate triggers with high p_Λ that populate the tail of the distribution. To discriminate those triggers, we implement a veto, “Rveto,” based on the ratio of incoherent energy between the detectors $R = E_I/E_J$, similar to what is done for STAMP-AS in [55]. Figure 6 shows the cumulative distributions of R for background triggers and for




FIG. 6. Distribution of the incoherent energy ratio R obtained for background triggers (in blue) and GW signal triggers from the ADI-A waveform (in orange) using the time-average PSD. Rejecting triggers with $R > 4$ allows for reducing the excess of large p_Λ background triggers, while marginally affecting the pipeline efficiency to recover GW signals.

triggers recovered for a GW waveform (ADI-A). Vetoing triggers with $R > 4$ allows one to reduce by a factor 5 the number of triggers, but more interestingly, the tail of the distribution of p_Λ is drastically reduced to approach the Gaussian noise triggers estimation, while no more than 5% of GW signal triggers are vetoed. In this paper, we are just illustrating that the pipeline behavior changes considerably

TABLE V. Values of h_{rss} at 50% detection efficiency for different waveforms obtained with PySTAMPAS for the two PSD methods and STAMP-AS over O2 data from LIGO Hanford and LIGO Livingston, using a FAR threshold of $1/50 \text{ yr}^{-1}$. The last column shows the ratio between STAMP-AS and the lowest value of PySTAMPAS among the two PSD methods. White noise burst waveforms WNB-A are not recovered at all with the frequency-median PSD.

Waveform	PySTAMPAS		STAMP-AS	
	Frequency median	Time average		Ratio
ISCOchirp-A	9.18×10^{-21}	8.17×10^{-21}	6.20×10^{-21}	0.76
ISCOchirp-B	1.84×10^{-21}	2.01×10^{-21}	1.44×10^{-21}	0.78
ISCOchirp-C	8.89×10^{-22}	1.06×10^{-21}	1.01×10^{-21}	0.95
ECBC-A	9.95×10^{-22}	1.07×10^{-21}	1.55×10^{-21}	1.55
ECBC-B	8.81×10^{-22}	8.61×10^{-22}	1.34×10^{-21}	1.56
ECBC-C	8.64×10^{-22}	8.00×10^{-22}	1.35×10^{-21}	1.69
ECBC-D	1.20×10^{-21}	8.95×10^{-22}	1.48×10^{-21}	1.65
ECBC-E	1.12×10^{-21}	8.82×10^{-22}	1.89×10^{-21}	2.14
ECBC-F	9.25×10^{-22}	7.83×10^{-22}	9.64×10^{-22}	1.23
ADI-B	3.26×10^{-22}	3.26×10^{-22}	4.81×10^{-22}	1.47
SG-C	4.34×10^{-22}	6.88×10^{-22}	4.35×10^{-21}	10.0
WNB-A	...	2.0×10^{-21}	2.04×10^{-21}	1.00

in the presence of non-Gaussian and nonstationary data. We also show that simple postprocessing selection criteria can be easily developed and applied with a relatively small penalty for the overall pipeline sensitivity.

As we have done for the study with simulated Gaussian noise, we now estimate the detection efficiency of this search with the two PSD estimators and compare it to results obtained by STAMP-AS during the second Advanced LIGO observing run [55] for a FAR of $1/50 \text{ yr}^{-1}$. For the time-average PSD, signals with $R > 4$ are rejected like is done in the background study. Best results obtained for each waveform among the two PSD methods are presented in Table V and compared to STAMP-AS. The relative detection efficiency depends on the waveform, but the overall PySTAMPAS pipeline efficiency increase observed with real data is very similar to what was obtained on simulated Gaussian data.

V. CONCLUSION

In this paper, we have presented PySTAMPAS, a new data analysis pipeline designed to search for GWs of duration $\sim 10\text{--}10^3 \text{ s}$ in a network of detectors with minimal assumptions on the nature and origin of the signal. The search algorithm relies on a hierarchical method, initially designed for a seedless clustering algorithm [19], where candidate events are first identified in single-detector ft maps, and a coherent detection statistic is then computed by cross-correlating data streams from each pair of detector. This method provides a significant gain in computational efficiency compared to the initial implementation of STAMP-AS with seed-based clustering, while still benefiting from the increased sensitivity of coherent searches. This is especially critical for all-sky/all-time searches for which both the dataset and the parameter space can be very large.

The reduced computational cost allows us to implement several new features to improve the overall sensitivity of the pipeline. The use of multiresolution ft maps enables the better reconstruction of signals with fast frequency evolution. An alternative method to estimate the noise PSD is proposed that is best suited for monochromatic and quasimonochromatic signals. We also introduce a new detection statistic that compares the coherent SNR of an event to the incoherent autopower in single detectors in order to discriminate coherent GW signals from loud noise events. Additionally, it is now feasible to scan hundreds of sky positions during the coherence stage and therefore to reduce the loss of SNR due to an error in the sky position to less than 5%. The combination of these features results in a detection efficiency increased by a factor ~ 1.5 on average compared to the previous version of STAMP-AS with seed-based clustering for the different waveforms tested, which have durations between 8 and

291 s, frequencies between 10 and 2048 Hz, and various spectral morphologies. We note that the changes in detection efficiency are dependant on the type of waveform, with PySTAMPAS performing slightly less well on waveforms from the ISCOchirp family and better for the remaining waveforms. We plan to improve the tuning of the clustering algorithm to address this issue.

PySTAMPAS is able to perform all-sky or targeted searches over a full observing run and a network of detectors and provides a basis for further developments. For example, the BURSTEGARD algorithm has been used here to identify clusters of excess power pixels, but other detection algorithms could be considered, such as seedless clustering [74] or more complex pattern recognition algorithms. This will be need to be done in order for the pipeline to be fully competitive, as shown by the example of the ISCOchirp waveforms family, which are currently slightly less well recovered by PySTAMPAS. We have shown that a real GW data search requires one to develop specific trigger selection to cope with non-Gaussian and nonstationary features of GW detectors data, but another possibility of improvement could consist in implementing a better identification and subtraction of non-Gaussian features of the GW detectors noise, as well as better discriminant variables.

ACKNOWLEDGMENTS

N.C. is supported by NSF Grant No. PHY-1806990. M.C. is supported by NSF Grants No. PHY-2010970 and No. OAC-2117997. This research has made use of data, software, and/or web tools obtained from the Gravitational Wave Open Science Center (<https://www.gw-openscience.org/>), a service of LIGO Laboratory, the LIGO Scientific Collaboration, and the Virgo Collaboration. LIGO Laboratory and Advanced LIGO are funded by the U.S. National Science Foundation (NSF) as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. Virgo is funded, through the European Gravitational Observatory (EGO), by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale di Fisica Nucleare (INFN), and the Dutch Nikhef, with contributions by institutions from Belgium, Germany, Greece, Hungary, Ireland, Japan, Monaco, Poland, Portugal, and Spain [72]. The authors are grateful for computational resources provided by the LIGO Laboratory and supported by National Science Foundation Grants No. PHY-0757058 and No. PHY-0823459.

- [1] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Observation of Gravitational Waves from a Binary Black Hole Merger, *Phys. Rev. Lett.* **116**, 061102 (2016).
- [2] J. Aasi *et al.* (LIGO Scientific Collaboration), Advanced LIGO, *Classical Quantum Gravity* **32**, 074001 (2015).
- [3] F. Acernese *et al.* (Virgo Collaboration), Advanced Virgo: A second-generation interferometric gravitational wave detector, *Classical Quantum Gravity* **32**, 024001 (2015).
- [4] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral, *Phys. Rev. Lett.* **119**, 161101 (2017).
- [5] B. P. Abbott *et al.* (Virgo, Fermi-GBM, INTEGRAL, and LIGO Scientific Collaborations), Gravitational waves and gamma-rays from a binary neutron star merger: GW170817 and GRB 170817A, *Astrophys. J.* **848**, L13 (2017).
- [6] D. A. Coulter *et al.*, Swope Supernova Survey 2017a (SSS17a), the optical counterpart to a gravitational wave source, *Science* **358**, 1556 (2017).
- [7] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo during the First Half of the Third Observing Run, *Phys. Rev. X* **11**, 021053 (2021).
- [8] M. Bailes *et al.*, Gravitational-wave physics and astronomy in the 2020s and 2030s, *Nat. Rev. Phys.* **3**, 344 (2021).
- [9] T. Damour and A. Vilenkin, Gravitational wave bursts from cusps and kinks on cosmic strings, *Phys. Rev. D* **64**, 064008 (2001).
- [10] M. Sieniawska and M. Bejger, Continuous gravitational waves from neutron stars: Current status and prospects, *Universe* **5**, 217 (2019).
- [11] K. Glampedakis and L. Gualtieri, Gravitational waves from single neutron stars: An advanced detector era survey, in *The Physics and Astrophysics of Neutron Stars*, edited by L. Rezzolla, P. Pizzochero, D. I. Jones, N. Rea, and I. Vidaña (Springer International Publishing, Cham, 2018), pp. 673–736.
- [12] W. G. Anderson, P. R. Brady, J. D. E. Creighton, and E. E. Flanagan, An excess power statistic for detection of burst sources of gravitational radiation, *Phys. Rev. D* **63**, 042003 (2001).
- [13] S. Klimenko *et al.*, Method for detection and reconstruction of gravitational wave transients with networks of advanced detectors, *Phys. Rev. D* **93**, 042004 (2016).
- [14] S. Chatterji, L. Blackburn, G. Martin, and E. Katsavounidis, Multiresolution techniques for the detection of gravitational-wave bursts, *Classical Quantum Gravity* **21**, S1809 (2004).
- [15] P. J. Sutton *et al.*, X-Pipeline: An analysis package for autonomous gravitational-wave burst searches, *New J. Phys.* **12**, 053034 (2010).
- [16] N. J. Cornish and T. B. Littenberg, BayesWave: Bayesian inference for gravitational wave bursts and instrument glitches, *Classical Quantum Gravity* **32**, 135012 (2015).
- [17] R. Prix, S. Giampanis, and C. Messenger, Search method for long-duration gravitational-wave transients from neutron stars, *Phys. Rev. D* **84**, 023007 (2011).
- [18] E. Thrane *et al.*, Long gravitational-wave transients and associated detection strategies for a network of terrestrial interferometers, *Phys. Rev. D* **83**, 083004 (2011).
- [19] E. Thrane and M. Coughlin, Detecting Gravitation-Wave Transients at 5σ : A Hierarchical Approach, *Phys. Rev. Lett.* **115**, 181102 (2015).
- [20] E. Thrane, V. Mandic, and N. Christensen, Detecting very long-lived gravitational-wave transients lasting hours to weeks, *Phys. Rev. D* **91**, 104021 (2015).
- [21] D. Keitel and G. Ashton, Faster search for long gravitational-wave transients: GPU implementation of the transient \mathcal{F} -statistic, *Classical Quantum Gravity* **35**, 205003 (2018).
- [22] A. Miller *et al.*, Method to search for long duration gravitational wave transients from isolated neutron stars using the generalized frequency-Hough transform, *Phys. Rev. D* **98**, 102004 (2018).
- [23] M. Oliver, D. Keitel, and A. M. Sintes, Adaptive transient Hough method for long-duration gravitational wave transients, *Phys. Rev. D* **99**, 104067 (2019).
- [24] S. E. Woosley and J. S. Bloom, The supernova gamma-ray burst connection, *Annu. Rev. Astron. Astrophys.* **44**, 507 (2006).
- [25] S. Nagataki, Theories of central engine for long gamma-ray bursts, *Rep. Prog. Phys.* **81**, 026901 (2018).
- [26] H.-T. Janka, Explosion mechanisms of core-collapse supernovae, *Annu. Rev. Nucl. Part. Sci.* **62**, 407 (2012).
- [27] E. Abdikamalov, G. Pagliaroli, and D. Radice, Gravitational waves from core-collapse supernovae, *arXiv:2010.04356*.
- [28] A. I. MacFadyen, S. E. Woosley, and A. Heger, Supernovae, jets, and collapsars, *Astrophys. J.* **550**, 410 (2001).
- [29] A. L. Piro and E. Pfahl, Fragmentation of collapsar disks and the production of gravitational waves, *Astrophys. J.* **658**, 1173 (2007).
- [30] M. H. P. M. van Putten, Proposed Source of Gravitational Radiation from a Torus around a Black Hole, *Phys. Rev. Lett.* **87**, 091101 (2001).
- [31] N. Bucciantini, E. Quataert, B. D. Metzger, T. A. Thompson, J. Arons, and L. Del Zanna, Magnetized relativistic jets and long-duration GRBs from magnetar spindown during core-collapse supernovae, *Mon. Not. R. Astron. Soc.* **396**, 2038 (2009).
- [32] S. Dall’Osso, L. Stella, and C. Palomba, Neutron star bulk viscosity, ‘spin-flip’ and GW emission of newly born magnetars, *Mon. Not. R. Astron. Soc.* **480**, 1353 (2018).
- [33] S. Dall’Osso and L. Stella, Millisecond magnetars, *arXiv:2103.10878*.
- [34] N. Sarin, P. D. Lasky, L. Sammut, and G. Ashton, X-ray guided gravitational-wave search for binary neutron star merger remnants, *Phys. Rev. D* **98**, 043011 (2018).
- [35] V. Ravi and P. Lasky, The birth of black holes: Neutron star collapse times, gamma-ray bursts and fast radio bursts, *Mon. Not. R. Astron. Soc.* **441**, 2433 (2014).
- [36] N. Sarin and P. D. Lasky, The evolution of binary neutron star post-merger remnants: A review, *Gen. Relativ. Gravit.* **53**, 59 (2021).
- [37] S. Bonazzola and E. Gourgoulhon, Gravitational waves from pulsars: Emission by the magnetic field induced distortion, *Astron. Astrophys.* **312**, 675 (1996), <https://ui.adsabs.harvard.edu/abs/1996A%26A...312..675B/abstract>.

- [38] C. Cutler, Gravitational waves from neutron stars with large toroidal B fields, *Phys. Rev. D* **66**, 084025 (2002).
- [39] L. Lindblom, B. J. Owen, and S. M. Morsink, Gravitational Radiation Instability in Hot Young Neutron Stars, *Phys. Rev. Lett.* **80**, 4843 (1998).
- [40] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Search for post-merger gravitational waves from the remnant of the binary neutron star merger GW170817, *Astrophys. J. Lett.* **851**, L16 (2017).
- [41] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GW190425: Observation of a compact binary coalescence with total mass $\sim 3.4 M_{\odot}$, *Astrophys. J. Lett.* **892**, L3 (2020).
- [42] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso *et al.*, Search for gravitational waves from a long-lived remnant of the binary neutron star merger GW170817, *Astrophys. J.* **875**, 160 (2019).
- [43] A. Krawczyk, A. G. Lyne, J. A. Gil, and B. C. Joshi, Observations of fourteen pulsar glitches, *Mon. Not. R. Astron. Soc.* **340**, 1087 (2003).
- [44] C. A. van Eysden and A. Melatos, Gravitational radiation from pulsar glitches, *Classical Quantum Gravity* **25**, 225020 (2008).
- [45] M. F. Bennett, C. A. Van Eysden, and A. Melatos, Continuous-wave gravitational radiation from pulsar glitch recovery, *Mon. Not. R. Astron. Soc.* **409**, 1705 (2010).
- [46] A. Melatos, J. A. Douglass, and T. P. Simula, Persistent gravitational radiation from glitching pulsars, *Astrophys. J.* **807**, 132 (2015).
- [47] K. Ioka, Magnetic deformation of magnetars for the giant flares of the soft gamma-ray repeaters, *Mon. Not. R. Astron. Soc.* **327**, 639 (2001).
- [48] A. Corsi and B. J. Owen, Maximum gravitational-wave energy emissible in magnetar flares, *Phys. Rev. D* **83**, 104014 (2011).
- [49] R. Quitzow-James, J. Brau, J. A. Clark, M. W. Coughlin, S. B. Coughlin, R. Frey, P. Schale, D. Talukder, and E. Thrane, Exploring a search for long-duration transient gravitational waves associated with magnetar bursts, *Classical Quantum Gravity* **34**, 164002 (2017).
- [50] E. Burns *et al.*, Identification of a local sample of gamma-ray bursts consistent with a magnetar giant flare origin, *Astrophys. J. Lett.* **907**, L28 (2021).
- [51] M. Drago, V. Gayathri, S. Klimenko, C. Lazzaro, E. Milotti, G. Mitselmakher, V. Nacula, B. O'Brian, G. A. Prodi, F. Salemi, M. Szczepanczyk, S. Tiwari, V. Tiwari, G. Vedovato, and I. Yakushin, Coherent waveburst, a pipeline for unmodeled gravitational-wave data analysis, *arXiv:2006.12604*.
- [52] J. Aasi *et al.* (LIGO Scientific and VIRGO Collaborations), Search for long-lived gravitational-wave transients coincident with long gamma-ray bursts, *Phys. Rev. D* **88**, 122004 (2013).
- [53] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), All-sky search for long-duration gravitational wave transients with initial LIGO, *Phys. Rev. D* **93**, 042005 (2016).
- [54] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), All-sky search for long-duration gravitational wave transients in the first advanced LIGO observing run, *Classical Quantum Gravity* **35**, 065009 (2018).
- [55] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), All-sky search for long-duration gravitational-wave transients in the second advanced LIGO observing run, *Phys. Rev. D* **99**, 104033 (2019).
- [56] S. W. Ballmer, A radiometer for stochastic gravitational waves, *Classical Quantum Gravity* **23**, S179 (2006).
- [57] T. Prestegard, Unmodeled search for long-lasting gravitational-wave signals with LIGO and studies of underground seismic noise for future gravitational-wave detectors, Ph.D. thesis, University of Minnesota, 2016.
- [58] D. M. Macleod, J. S. Areeda, S. B. Coughlin, T. J. Massinger, and A. L. Urban, GWpy: A Python package for gravitational-wave astrophysics, *SoftwareX* **13**, 100657 (2021).
- [59] N. Christensen, LIGO S6 detector characterization studies, *Classical Quantum Gravity* **27**, 194010 (2010).
- [60] M. Zevin *et al.*, Gravity spy: Integrating advanced LIGO detector characterization, machine learning, and citizen science, *Classical Quantum Gravity* **34**, 064003 (2017).
- [61] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), A guide to LIGO–Virgo detector noise and extraction of transient gravitational-wave signals, *Classical Quantum Gravity* **37**, 055002 (2020).
- [62] A. Buikema *et al.* (aLIGO Collaboration), Sensitivity and performance of the Advanced LIGO detectors in the third observing run, *Phys. Rev. D* **102**, 062003 (2020).
- [63] D. Davis *et al.* (LIGO Collaboration), LIGO detector characterization in the second and third observing runs, *Classical Quantum Gravity* **38**, 135014 (2021).
- [64] P. B. Covas *et al.* (LSC Collaboration), Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO, *Phys. Rev. D* **97**, 082002 (2018).
- [65] M. Was, M.-A. Bizouard, V. Brisson, F. Cavalier, M. Davier, P. Hello, N. Leroy, F. Robinet, and M. Vavoulidis, On the background estimation by time slides in a network of gravitational wave detectors, *Classical Quantum Gravity* **27**, 015005 (2010).
- [66] E. Huerta *et al.*, ENIGMA: Eccentric, nonspinning, inspiral Gaussian-process merger approximant for the characterization of eccentric binary black hole mergers, *Phys. Rev. D* **97**, 024031 (2018).
- [67] M. H. P. M. Van Putten, Gravitational waveforms of Kerr black holes interacting with high-density matter, *Astrophys. J. Lett.* **684**, L91 (2008).
- [68] M. H. P. M. Van Putten, Directed searches for broadband extended gravitational wave emission in nearby energetic core-collapse supernovae, *Astrophys. J.* **819**, 169 (2016).
- [69] B. P. Abbott *et al.*, GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs, *Phys. Rev. X* **9**, 031040 (2019).
- [70] J. Kissel, H1 Calibrated Sensitivity Spectra Jun 10 2017, LIGO Document, LIGO, 2018, <https://dcc.ligo.org/LIGO-G1801950/public>.

- [71] J. Kissel, L1 Calibrated Sensitivity Spectra Aug 06 2017, LIGO Document No. G1801952 LIGO, 2018, <https://dcc.ligo.org/LIGO-G1801952/public>.
- [72] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Open data from the first and second observing runs of Advanced LIGO and Advanced Virgo, *SoftwareX* **13**, 100658 (2021).
- [73] Gravitational Wave Open Data Center, O2 data set, <https://doi.org/10.7935/CA75-FM95>.
- [74] E. Thrane and M. Coughlin, Searching for gravitational-wave transients with a qualitative signal model: Seedless clustering strategies, *Phys. Rev. D* **88**, 083010 (2013).