

Fast premerger detection of massive black-hole binaries in LISA based on time-frequency excess power

Francesco Nobili^{1,2,*} Malvina Bellotti^{3,4} Riccardo Buscicchio^{3,2,5} Massimo Dotti^{3,2} and Alessandro Lupi^{1,2}

¹*Como Lake Centre for Astrophysics, DiSAT, Università dell'Insubria,
via Valleggio 11, I-22100 Como, Italy*

²*INFN, Sezione di Milano-Bicocca, Piazza della Scienza 3, 20126 Milan, Italy*

³*Dipartimento di Fisica "G. Occhialini," Università degli Studi di Milano-Bicocca,
Piazza della Scienza 3, 20126 Milan, Italy*

⁴*Université Côte d'Azur, Observatoire de la Côte d'Azur, CNRS Artemis, Bd de l'Observatoire,
F-06304 Nice, France*

⁵*Institute for Gravitational Wave Astronomy and School of Physics and Astronomy,
University of Birmingham, Birmingham B15 2TT, United Kingdom*



(Received 20 February 2026; accepted 30 July 2026; published 2 September 2026)

The Laser Interferometer Space Antenna is expected to observe gravitational waves from massive black hole binaries across cosmic time. Many are anticipated to be detectable hours to weeks before coalescence. We present a fast algorithm for the premerger detection and preliminary characterization of such binaries. The method performs a search for excess power with a chirping time-frequency morphology in short-time Fourier transform spectrograms. By tiling the time-frequency plane with slices defined by the quadrupole frequency evolution, we define a signal significance relative to a fitted background distribution of instrumental noise and Galactic foreground. Individual search triggers are followed by a coherence tracker that groups triggers consistent with the same physical signal over time. Doing so, our analysis provides progressively refined estimates of the chirp mass and coalescence time. We validate our algorithm on the SangriaHM LISA Data Challenge dataset, successfully detecting all 15 injected massive black hole binaries; 14 of them hours to weeks before merger, while one is only detected after the binary coalescence. The algorithm yields chirp mass relative errors below 3% for high-SNR sources and coalescence time uncertainties of up to a few hours. With a computational cost of less than a second to process a 10-day data segment on a single core, our approach is suitable for generating real-time alerts, triggering protected observational periods, and providing informative priors for Bayesian parameter estimation.

DOI: [10.1103/PhysRevD.114.063001](https://doi.org/10.1103/PhysRevD.114.063001)

I. INTRODUCTION

The Laser Interferometer Space Antenna (LISA) will be the first space-based gravitational-wave observatory (GW) [1]. Operating in the millihertz regime, LISA will bridge the observational gap between current and future ground-based interferometers [2–4] and pulsar timing array observatories [5]. A wide variety of sources is expected to be observed [6,7]; the multitude of Galactic double white dwarfs (DWDs) producing a stochastic foreground, stellar-mass black holes in their early inspiral, compact stellar remnants orbiting massive black holes (i.e., extreme mass-ratio inspirals), and

coalescing massive black hole binaries (MBHBs). The last are characterized by loud transient signals lasting from hours to weeks in the LISA sensitivity band.

The simultaneous presence of these sources, overlapping in time and frequency, makes the LISA data analysis a complex challenge motivating the development of global fits [8–12]; pipelines capable of performing simultaneous Bayesian parameter estimation on an uncertain number of sources within the data stream. Owing to their complexity, such algorithms likely require months of accumulated data before delivering accurate results.

Nonetheless, a fast detection and preliminary characterization of transient sources is needed for time-critical scientific investigations. This is the objective of low-latency pipelines, designed to operate in real time and release rapid alerts containing preliminary source-property estimates within one hour from the reception of new data on the ground. The target alert emission timescale is valuable to trigger protected observational periods preventing scheduled gaps in the data stream, and to enable multimessenger

*Contact author: fnobili@uninsubria.it

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

observations with electromagnetic observatories. Similar low-latency strategies have been extensively developed for ground-based detectors [13,14], primarily targeting neutron star mergers, and have already enabled groundbreaking results [15–19]. Among LISA detectable sources, MBHBs are the most promising candidates for such followups. At present, massive black holes observable by LISA (i.e., in the 10^4 – $10^7 M_\odot$ mass range) are identified through electromagnetic observations, either as the bright central engines of Seyfert galaxy nuclei (see Ref. [20] for a review) or as responsible for the dynamics of nuclear stars (as for the case of Sgr A*, e.g., Ref. [21]) and of megamaser discs [22]. Mergers of MBHBs occurring in gas-rich environments may produce peculiar electromagnetic signatures both before and after coalescence [23]. While no clear electromagnetic evidence for MBHBs have been presented so far,¹ several strategies have been proposed to perform multimessenger astronomy through LISA GW detections and several EM observatories [25–28]. Notably, searches for MBHB signatures in Legacy Survey of Space and Time (LSST) photometric data [29–32], complemented with spectroscopic validations [33], and follow-up observations with the NewAthena telescope [34,35] have been put forward in the literature.

Several methods have been proposed for the rapid detection and characterization of GWs from MBHBs. One class of such methods aims at accelerating the analysis of full inspiral-merger-ringdown MBHB signals, mainly relying on approximate detection and parameter estimation schemes [36–43]. A second class focuses on the MBHB premerger signals, leveraging either approximated matched-filtering schemes [43–45] or machine-learning approaches [46–48]. In the context of LISA, several time-frequency frameworks have recently been proposed for the search and characterization of massive black hole binaries [45], extreme mass-ratio inspirals [49], and stellar-mass black holes [50]. Furthermore, novel time-delay interferometry techniques are currently being developed with the objective of pinpointing the sky position of MBHBs [51]. Finally, several time-frequency (TF) excess power methods have been presented in past literature, especially regarding agnostic burst searches [52–54], or waveform-free estimation of eccentricity [55–57] applied to data from ground-based interferometers.

In this work, we present a fast detection algorithm based on the search for excess power with a time-frequency chirping morphology through short-time Fourier transform (STFT) spectrograms. We do so inspired by the successful application of similar methods to ground-based interferometers data, where analogous detection techniques using the same data representation have been proposed [58–65]. However, we highlight two key differences with respect to

our method in the LISA context; (i) our algorithm is able to identify multiple overlapping signals in a single chunk of observation, and (ii) it tracks parameter estimates and uncertainties, progressively updating them as further data are acquired.

Datastreams of current ground-based detectors are largely noise dominated. By contrast, LISA data are signal-dominated, with a background composed of stationary instrumental noise and a cyclostationary foreground from Galactic DWDs [66–68]. We explicitly account for this through off-source noise estimation.

Overlapping signals in current ground-based detectors are extremely rare. On the contrary, long-duration LISA MBHB inspirals may overlap in time [69]. As we shall see in following sections, we make a considerable effort to retain large flexibility in the selection of spectrogram pixels used for statistical testing. The motivation is twofold; the extended duration of MBHB signals makes them susceptible to contamination from glitches and gaps. Glitches are loud and brief transients that have been observed in the data of the LISA Pathfinder mission [70,71]. Their characterization is currently a live research topic [72–74]. Gaps are downtime periods during which LISA will not collect science-quality data [75], either due to scheduled maintenance periods, or to unscheduled—temporary or permanent—data losses. Working in the TF domain is particularly advantageous to handle both artifacts, given the large flexibility in selecting spectrogram pixels unaffected by them.

The paper is organized as follows. In Sec. II, we describe our algorithm; we detail our time-frequency data representation and noise estimation, both based on STFTs, in Sec. II A; in Sec. II B, we introduce the signal model assumed and the procedure to estimate the detector-frame chirp mass $\mathcal{M}_c$ and coalescence time t_c from each excess-power detection; in Sec. II C we describe how search triggers from each segment are coherently tracked through time to identify physical sources and iteratively update their chirp mass and time-to-coalescence estimates. In Sec. III, we summarize our results, showing that our algorithm successfully detects all sources in the SangriaHM dataset with minimal computational cost. Finally, we conclude in Sec. IV with a discussion of the limitations of our approach and possible extensions for future work. In Appendixes A and B we illustrate results for each individual system, and report a detailed derivation of the background statistics, respectively.

II. METHODOLOGY

A. Time-frequency representation

Throughout our work, we assume LISA interspacecraft laser-phase measurements to be suitably delayed and combined into time-delay interferometric variables (TDI), to achieve the target GW sensitivity through synthetic laser-noise cancellation [76]. We do so coherently using the *SangriaHM* LISA Data Challenge [77,78].

¹The best candidate so far is represented by the pair of radio cores at ≈ 7 pc projected separation discussed in [24].

The simplest TDI combinations, TDI-1.0, assume a stationary array described by three constant armlengths, so that the light travel time along a given arm is the same in both propagation directions. The dataset is instead released in the form of TDI-1.5 observables, which account for the Sagnac effect induced by the rotation of the constellation—implying that the up and down optical paths are unequal—thus involving six distinct, though still time-independent, delays. Future simulated datasets will drop this approximation in favor of TDI-2.0 combinations, which further account for time-dependent armlengths [76,79]. As we shall see in Sec. II A, this would only affect the whitening procedure. The dataset is released in the form of X , Y , and Z variables. Through linear combinations of such TDI variables, we construct auxiliary, uncorrelated ones (known in literature as A , E , and T). The dataset features one year of observations with a cadence of $T_s = 5$ s, containing simulated instrumental noise [80], bright and foreground signals from a simulated Galaxy of white dwarf binaries, and 15 MBHB merger signals following the PhenomHM waveform model [81]. Henceforth, similarly to the convention adopted in Ref. [43], we will refer to the n th MBHB system, when ordered by increasing merger time, as MBHB- n (starting with $n = 0$).

Our algorithm is designed to operate in real time, processing new data as soon as they become available. To avoid ambiguity, we refer to the interval of time-domain data considered at each step as *chunk*. The choice of chunk length is adjustable depending on the typical source lifespan within the LISA band. We adopt a default time-sliding chunk of fixed 10 days length, so that incoming data points replace older ones, sliding forward every 200 minutes.

Once a data chunk is collected, we compute its TF representation via STFT, as implemented in SciPy [82]. For a discrete time series $x[k]$ sampled with cadence T_s , the STFT is obtained by taking the discrete Fourier transform (DFT) of short, overlapping segments. Each segment is selected multiplying the data by a window function $w[k - ph]$ centered at time $p\Delta t$, and of duration $\tau = NT_s$, with p denoting the segment number, and N the number of samples in each segment. The window is then shifted forward by h samples at each step, corresponding to a time shift (or *hop*) of $\Delta t = hT_s$. The STFT spectrum evaluated at time $p\Delta t$ and frequency $q\Delta f$ is therefore,

$$X[p, q] = \sum_{k=0}^{N-1} x[k]w[k - ph]e^{-i2\pi qk/N}. \quad (1)$$

In this work, we use the term *spectrogram* to denote the discrete-frequency power spectral density (PSD), defined as

$$\text{PSD}[p, q] = T_s \frac{|X[p, q]|^2}{\|w\|^2}, \quad (2)$$

which has units of Hz^{-1} . Here, $\|w\|$ is a normalization constant defined as

$$\|w\|^2 = \sum_{k=0}^{N-1} |w[k]|^2. \quad (3)$$

We employ the standard Hann window, i.e., a scaled cosine with support over $0 \leq n < N$,

$$w[n] = \frac{1}{2} - \frac{1}{2} \cos\left(\frac{2\pi n}{N-1}\right). \quad (4)$$

This window is widely used as it provides a good compromise between frequency resolution and dynamic range [83].

In our default configuration, we employ segments of duration $\tau = 10^5$ s with a 95% overlap, restricting the analysis to the frequency range between 10^{-4} and 10^{-3} Hz. While extending the search to lower frequency bands could enhance the early detection of lower-mass signals characterized by prominent inspiral phases [44], we adopt this conservative lower limit to demonstrate that our algorithm remains effective even under restrictive conditions.

We show in Fig. 1 the overall STFT scheme mapping each windowed segment to a Fourier series; to construct a spectrogram, we consider only segments whose window do not overlap with chunk borders. For a 10 days-long chunk, this amounts to consider segments centered from 0.58 up to 9.38 days. Doing so, we drop zero-padded segments as constructed in the `scipy.signal.ShortTimeFFT`




FIG. 1. Summary of STFT scheme and free parameters. In our default configuration we consider a sliding chunk of 10 days sampled with cadence $T_s = 5$ s. Each chunk is uniquely associated to a spectrogram, as defined in Eqs. (1) and (2); segments of length τ , overlapping by 95% are Hann-windowed according to Eq. (4) and Fourier transformed, yielding a frequency array associated to the time $p\Delta t$. Note that windows in the figure are overlapping by 50% for a cleaner visualization. We further mask out pixels corresponding to windows with support outside of the chunk borders (teal-bordered boxes), and pixels above 10^{-3} Hz and below 10^{-4} Hz (white, black-bordered boxes).

implementation [84], thus avoiding ringing artifacts due to the Gibbs phenomenon [85] near sharp discontinuities, which in this case happen at data chunk borders.

We further process each spectrogram through whitening. In ground-based detectors literature, methods for both on-source and off-source noise estimation have been devised [86,87]. In LISA transient searches, a major component of the background arises from Galactic sources, whose detector modulation becomes significant over weeks timescales [67]. Therefore, background PSD can be calibrated within reasonable approximation using data extending over the same time in the past, in absence of other bright transient signals.

Low-latency pipelines running on the LISA Distributed Data Processing center (DDPC) will likely be equipped with regular updates of instrumental noise, DWD confusion noise, and resolved DWDs parameter estimates, which together define the noise-hypothesis distribution, $\mathcal{H}_0$. These updates may be obtained by dedicated blocks of the ground-segment global-fit pipeline, sharing background point estimates previous to their full convergence. Such preliminary estimates of the PSD of the background process may be directly used in the whitening procedure described below, and in the computation of the distribution of the detection statistic introduced in Sec. II B, under the background assumption $\mathcal{H}_0$. However, an independent estimate of the background performed in the context of low-latency searches may prove beneficial, adding redundancy in the analysis chain. While this objective may be investigated in future work, in our analysis we estimate the background PSD directly from the SangriaHM dataset containing only instrumental noise and Galactic sources.

We consider 500 reference update times $t_u^{(k)}$ across the year of simulated data, i.e.,

$$t_u^{(k)} = 10 \text{ d} + k \times 0.71 \text{ d}. \quad (5)$$

At each update, we construct a reference PSD estimator from a 5 days-long chunk of background data X_{bkg} in $[t_u - 5 \text{ d}, t_u]$, mimicking an estimate based on previous data, as provided by the dedicated global fit blocks. We do so by computing STFTs as in Eq. (2), with windowing and overlap parameters as for the data, and taking the median value over the 76 time segments available to each t_u ,

$$\text{PSD}_0[q; t_u] = \frac{T_s}{\|w\|^2} \text{median}_{p \Delta t \in [t_u - 5 \text{ d}, t_u]} |X_{\text{bkg}}[p, q]|^2. \quad (6)$$

We therefore obtain 500 PSD estimates $\text{PSD}_0[q; t_u]$, each one a vector labeled by a reference t_u , whose frequency-content evolution over the year is shown in Fig. 2. Along with the Galactic confusion noise, this procedure naturally incorporates into the background model pixel power from faint or not-yet-resolved signals. In Fig. 2 such a scenario is represented by a bright DWD source at $f = 0.56 \text{ mHz}$.




FIG. 2. Array of PSDs used for whitening the spectrograms at each update time t_u . The Galactic confusion noise modulation is visible as a periodic brightening particularly clear above 0.5 mHz. The brightness-modulated narrow horizontal line at $f = 5.6 \times 10^{-4} \text{ Hz}$ is the signal emitted by a loud DWD source, at a distance of 0.1 kpc, far from the Galactic center, which is successfully suppressed in the whitened spectrograms.

Being located relatively nearby and far from the Galactic center (at a distance of 0.1 kpc, Ecliptic longitude $l = 139^\circ$ and latitude $\beta = 4^\circ$) its signal could potentially be missed by noise-estimation routines—typically targeting the slowly evolving power from the Galactic center—and identified only later by high-latency pipelines.

Each data spectrogram (which we identify for convenience by its last timestamp $t^\star$) is then whitened using the background PSDs,

$$\widehat{\text{PSD}}[p, q] = \frac{\text{PSD}[p, q]}{\text{PSD}_0[q; t_u]}, \quad (7)$$

where t_u is chosen to be the unique one satisfying $10 \text{ d} < t^\star - t_u < 10.71 \text{ d}$ for a given $t^\star$, i.e., the closest preceding the data chunk. At runtime, multiple algorithms will run concurrently to identify transients, including MBHB mergers and instrumental artifacts. Therefore, the chosen delay of approximately 10 days provides sufficient margin to identify and mitigate unwanted transients that could otherwise contaminate the background estimation.

In Fig. 3 we show two examples of whitened spectrograms, obtained from data segments of duration 5 and 10 days; as we shall see in Sec. II B, this is the fundamental datum for our detection strategy.

B. Detection

Our detection algorithm is based on searching for excess power in spectrograms with the morphology of a chirping signal. We adopt a simplified approach and model MBHB signals considering only the quadrupole mode. In this framework, a *chirp track* in the time-frequency plane is determined solely by the detector-frame chirp mass $\mathcal{M}_c$.




FIG. 3. The top panel shows the default setup described in Sec. II A, while the bottom panel illustrates the configuration optimized for the detection of MBHB-2, discussed in Sec. III. This source has a chirp mass of $1.2 \times 10^6 M_\odot$; it is therefore characterized by a fainter inspiral and a big portion of the total power is emitted at higher frequencies. Both spectrograms correspond to chunk identified by $t^* = 130.5$ days, but with duration of 10 (top) and 5 (bottom) days, respectively. Dashed lines denote boundaries of a chirp slice, as defined in Sec. II B; in particular, we show the slice matching the MBHB-3 signal as an outlier in our detection statistics. In the default configuration, the chirping tracks of MBHB-3 and MBHB-4 are clearly visible, whereas MBHB-2 is only marginally. By contrast, the former is clearly visible in the dedicated high-frequency configuration (bottom).

and time-to-coalescence $t_c - t$ [88],

$$f(t; \mathcal{M}_c, t_c) = \frac{1}{\pi} \left(\frac{5}{256} \frac{1}{t_c - t} \right)^{3/8} \left(\frac{G\mathcal{M}_c}{c^3} \right)^{-5/8}. \quad (8)$$

In this simple model, we neglect the time and frequency dependency of the detector response, effectively working in the *frozen low-frequency* approximation [89]. This is justified by the data chunk duration of 10 days, short enough with respect to the detector motion timescale at frequencies below 5×10^{-3} Hz. While to define a time-frequency correspondence we consider only the fundamental (2, 2)-mode, the MBHB signals in SangriaHM are generated simulating (2, 2), (2, 1), (3, 3), (3, 2), (4, 3), and

(4, 4) harmonics. The impact of higher-order modes depends on the mass ratio of the binary system and the inclination angle i_0 between the orbital angular momentum and the observer. Sources in the SangriaHM dataset are circular and nonprecessing, with mass ratios spanning the range $0.36 < q < 0.97$, and inclination angles in the range $30^\circ < i_0 < 90^\circ$. The results presented in Sec. III demonstrate that for these systems, a quadrupole-only model is sufficient to detect the signals and derive approximately correct constraints on $(\mathcal{M}_c, t_c)$.

We define a uniform grid in $\log_{10} \mathcal{M}_c$ and $t_c^* = t_c - t^*$ where, for convenience of notation, we shift the time-of-coalescence by t^* . We further define a *chirp slice* as the set of pixels $\mathfrak{s} = \{p_i, q_i, i = 1, \dots, N_{\mathfrak{s}}\}$ whose centers in the TF plane are enclosed between two chirp tracks. We label each slice by a pair of values $(\log_{10} \mathcal{M}_c, t_c^*)$ defined as the arithmetic mean over those defining the two bounding chirp tracks. To better map signals into the $(\mathcal{M}_c, t_c^*)$ space, adjacent slices are shifted so as to maintain a significant overlap. In the default configuration, we limit our search to systems with $2 \times 10^5 M_\odot < \mathcal{M}_c < 5 \times 10^6 M_\odot$ and $-1 \text{ d} < t_c^* < 20 \text{ d}$, resulting in chirp slices defined by

$$\log_{10}(\mathcal{M}_c^{(n)}/M_\odot) = \log_{10}(2 \times 10^5) + 0.012837 \times n, \quad n = 0, \dots, 104, \quad (9)$$

$$t_c^{*(m)} = -1 \text{ d} + 12000 \text{ s} \times m, \quad m = 0, \dots, 148. \quad (10)$$

Although the search ranges for $\mathcal{M}_c$ and t_c^* were tailored to the SangriaHM dataset, they can be readily extended. While the computational cost scales linearly with the number of grid points in a serial execution, the analysis of each slice is entirely independent and therefore trivial to parallelize. We only consider slices containing more than 10 TF pixels.

As shown in Fig. 3, the proposed slices naturally accommodate for pixel power from merging MBHB signals. Depending on the chosen STFT window length and the STFT hopping window [90], signals do not appear as thin lines, rather as thick bands in the TF plane; a longer window yields higher frequency resolution, making the low-frequency inspiral sharper and spreading out the rapidly chirping portion of the signal. In this case, the signal can be bracketed by two chirp tracks with relatively close $\mathcal{M}_c$ values and widely separated t_c^* values. We show the signal of MBHB-3 as an example of this behavior in the top panel of Fig. 3, where the STFT window length is ~ 1.15 days. Conversely, a shorter window reduces the frequency resolution, broadening the inspiral and sharpening the high-frequency chirp power. Therefore, the same signal is enclosed by tracks with more widely separated $\mathcal{M}_c$ and closer t_c^* . We show in the bottom panel of Fig. 3

this configuration, where we set the STFT window length to ~ 1.39 hours.

We proceed with our detection algorithm tiling the TF plane with overlapping chirp slices and test each slice for excess power. Following the procedure described in Sec. II A, we first identify the whitening PSD_0 from the closest t_u to $t^* - 10$ d. We then whiten the chunk spectrogram and compute the average power in each slice,

$$\langle \widehat{\text{PSD}} \rangle_{\mathfrak{s}} = \frac{1}{N_{\mathfrak{s}}} \sum_{i=1}^{N_{\mathfrak{s}}} \widehat{\text{PSD}}[p_i, q_i], \quad (11)$$

where the sum runs over the $N_{\mathfrak{s}}$ TF pixels selected by the slice $\mathfrak{s}$.

To quantify the significance of a given slice averaged power, we compare the value from Eq. (11) against a background distribution derived under the noise hypothesis $\mathcal{H}_0$ (i.e., in absence of MBHB signals). For this purpose, we use the SangriaHM dataset containing only instrumental noise and Galactic sources. To collect background realizations, we construct 500 whitened spectrograms equally spaced across one year, following the same procedure described in Sec. II A. Since the background PSD estimates used for whitening the spectrograms include the cyclostationary component of the Galactic DWD foreground, the collection of whitened background spectrograms can be approximately treated as independent realizations of a stationary stochastic process.

Once a set of chirp slices is defined in the $(\mathcal{M}_c, t_c^*)$ space, we collect for each slice 500 samples characterizing its $\langle \widehat{\text{PSD}} \rangle_{\mathfrak{s}}$ distribution, reflecting both the number of TF pixels and the range of frequencies spanned. We find that a Gamma distribution provides an accurate two-parameters model across all slices, parametrized by (a, b) ,

$$f(x; a, b) = \frac{(bx)^{a-1} e^{-(bx)}}{b\Gamma(a)}. \quad (12)$$

In Fig. 4 we show two examples of fitted distributions for slices with different $(\mathcal{M}_c, t_c^*)$ values. Ideally, one would need independent realizations of the background process to collect $\langle \widehat{\text{PSD}} \rangle_{\mathfrak{s}}$ samples. However, since the 500 spectrograms are obtained from 10-days-long chunks equally distributed over one year, they overlap by approximately 93%. This introduces correlations across pixel power in the definition of the background distributions, especially for chirp slices spanning longer time intervals in the TF plane (lower $\mathcal{M}_c$ values).

To justify the choice of the Gamma distribution, in Appendix B we provide an analytical derivation of the background $\langle \widehat{\text{PSD}} \rangle_{\mathfrak{s}}$ distribution under a set of simplifying assumptions. We demonstrate that the random variable $\langle \widehat{\text{PSD}} \rangle$, computed over a set of uncorrelated time-frequency pixels, follows a Gamma distribution. We then extend this derivation to account for correlated pixels (arising from




FIG. 4. Example of background distributions relative to two spectrogram slices, in orange and teal, respectively. The distribution shape is affected by number of time-frequency points selected by each slice. Supported by the derivation in Appendix B, we find a Gamma distribution to be an adequate parametric family to fit for the observed distribution, obtained as described in Sec. II B. Shaded histograms and solid curve denote the synthetic realizations and the best-fit obtained following the model in Eq. (12), respectively.

overlapping STFT windows) and obtain an explicit expression for the survival function. These results support the use of a Gamma distribution to model the background $\langle \widehat{\text{PSD}} \rangle_{\mathfrak{s}}$. However, we maintain the use of a parametric fit, as it provides a more flexible approach that can account for imperfections in the whitening process. The interested reader may find a general derivation in the context of stochastic gravitational wave backgrounds in [91].

Given the fitted background probability density $p(\langle \widehat{\text{PSD}} \rangle_{\mathfrak{s}} | \mathcal{H}_0)$ for each slice, we quantify the false alarm probability (FAP) as the survival function of $\langle \widehat{\text{PSD}} \rangle_{\mathfrak{s}}$ observed from data,

$$\begin{aligned} \text{FAP}(x) &= \text{Prob}(\langle \widehat{\text{PSD}} \rangle_{\mathfrak{s}} > x | \mathcal{H}_0) \\ &= 1 - \int_0^x p_{\langle \widehat{\text{PSD}} \rangle_{\mathfrak{s}}}(u | \mathcal{H}_0) du. \end{aligned} \quad (13)$$

For each chunk, we obtain a FAP map, containing the FAP value relative to every slice $\mathfrak{s}$, as defined by the two values $(\mathcal{M}_c, t_c^*)$. Figure 5 illustrates an example of such maps, obtained from data containing the signals of MBHB-2, 3, and 4 for different chunk timestamps t^* . A threshold on the FAP is then applied to identify significant outliers. In our default configuration, the significance threshold is set to $\text{FAP}_{\text{thr}} = 10^{-6}$. In the following section we describe the collection and processing of triggers, and justify the choice of the FAP threshold value.

C. Source identification and triggers tracking

Once the FAP maps in the $(\mathcal{M}_c, t_c^*)$ parameter space are produced by the detection algorithm, we obtain parameter




FIG. 5. FAP maps at different chunk times t^* , following the default search configuration described in Sec. II B. The white regions in the top-right corners correspond to $(\mathcal{M}_c, t_c^*)$ values associated with slices of the time-frequency plane that contain no spectrogram points because they fall outside its boundaries. We also remove the thin region corresponding to slices containing fewer than 10 points. The markers show the true chirp mass and coalescence time from the SangriaHM dataset metadata. As the chunk timestamp approaches the coalescences, two regions in the $(\mathcal{M}_c, t_c^*)$ space emerge and gradually concentrate around the true values. As discussed in Sec. III in the default configuration, MBHB-2 is not detected by our algorithm.

estimates and uncertainties from them. To this end, we implement a labeling algorithm that identifies *connected regions*; collection of adjacent slices, i.e., connected points in the FAP map domain. We progressively scan increasing FAP values, moving upward from the global minimum, and identify connected components using feature labeling [92], as implemented in `scipy.ndimage.label` [93]. Climbing up higher FAP levels, initially disconnected regions merge, as saddle points between local minima are crossed. The algorithm monitors the number of distinct regions at each threshold: a reduction in the number of regions indicates that the threshold has crossed a saddle point between two previously independent minima. The procedure continues up to a predefined maximum FAP level.

To limit the effect of noise fluctuations—particularly at the high FAP levels typical of early-detections—we apply some quality cuts when identifying regions; (i) we enforce a minimum size of 30 connected grid points, and (ii) we automatically merge distinct regions whose local minima are separated by fewer than 15 grid points. These cuts are empirically tuned in the default configuration to suppress the regions fragmentation observed for early source identification in the SangriaHM dataset. Each identified region is then assigned a best estimate of $(\mathcal{M}_c, t_c^*)$ by its point of minimum FAP. Uncertainties can then be defined in several ways, depending on the analysis requirements. In this work, we adopt for each FAP region contour levels corresponding to fixed $\log_{10}$ FAP fractions of the local minimum.

The shape of a signal-activated region in a probability map evolves over time, as illustrated in Fig. 5. We track this evolution by matching regions identified at consecutive timestamps, allowing information to accumulate and progressively refine the parameter estimates. When a new

region is detected, the algorithm checks whether it significantly overlaps (at least 10 grid points in common) with one or more region from a previous timestamp. If sufficient overlap is found, the new region is associated with the existing one. At this point, only the overlapping points are kept, and the region data are updated with the new information while preserving its original label. When two previously tracked regions merge, due to a new one overlapping with both, the algorithm keeps only the overlapping grid points, labels them as the parent region with the lowest FAP, and marks the remainder as inactive. This behavior may be observed during the first stages of detection, when multiple local minima typically appear before merging into a single region.

On the other hand, when a tracked region splits into multiple ones at a subsequent timestamp, the new region with the largest overlap inherits the previous label, and the other new region is tracked independently. This behavior is observed when two overlapping signals become separable, like the case of MBHB-3 and 4.

Thresholds on the number of points per region and on the number of overlapping points are chosen empirically and depend on the grid density. These thresholds serve as temporary countermeasures to mitigate unwanted behavior, such as consecutive region-splitting and merging. In future work, a more robust control criteria based on the significance of local minima relative to the background FAP distribution, will be implemented. In Fig. 6 we summarize diagrammatically the region-tracking scheme as just described.

In Fig. 7, we instead illustrate the temporal evolution of the MBHB-0 parameter estimates, shown as relative error in chirp mass $\Delta\mathcal{M}_c/\mathcal{M}_c$ and difference in coalescence time Δt_c with respect to the true injected values. The estimates converge toward the injected ones, reported in the




FIG. 6. Overview of region tracking scheme across data segments, as described in Sec. II C. Flowchart of the region detection and tracking algorithm applied to FAP maps over time. For each time chunk t^* , the FAP map is thresholded, connected regions are identified, refined using quality cuts, and matched to previously active (“live”) regions based on overlap. Depending on the matching outcome, regions are newly created, updated (many-to-one), or split (one-to-many), with labels propagated accordingly. Regions that are not updated for more than two days are deactivated and saved.

SangriaHM dataset metadata; however, a small residual bias in the chirp mass ($\gtrsim 2\%$) and coalescence time (a few hours) may persist up to merger.

The current implementation of our algorithm enforces the retention of only overlapping points in the $\mathcal{M}_c$ - t_c space during region tracking across successive timestamps. This choice ensures that constraints on $\mathcal{M}_c$ established at the onset of tracking are preserved over time. However, an alternative version of the algorithm can be implemented without this requirement. In this case, the shape of the

regions tracked at each timestamp in the $\mathcal{M}_c$ - t_c space is no longer influenced by their past evolution. This approach can help prevent the loss of region tracking when early estimates are biased, but it may also increase confusion in the presence of overlapping sources. An example of this behavior is discussed in Appendix A and illustrated in Fig. 12.

The choice of the FAP detection threshold determines the algorithm’s sensitivity to noise fluctuations. To quantify this effect, we evaluate the number of spurious regions

detected when the algorithm is applied to data containing only instrumental noise and Galactic DWDs. Using a threshold of $\text{FAP}_{\text{thr}} = 10^{-6}$, which we adopt as our default value, the algorithm reports 10 false detections over one year of simulated data running on channel A, and 9 on channel E. False detections arise when background fluctuations cause the $\langle \text{PSD} \rangle_s$ in certain time-frequency slices to exceed the chosen threshold, leading the tracking algorithm to identify spurious regions. The A and E TDI channels are noise-orthogonal, and we verify that all occurrences of false positives are not concurrent across them. The absence of a concurrent trigger in both channels could be used, in principle, as a smoking gun to identify noise-driven spurious detections. Particular attention should be paid, however, as even genuine GW signals can exhibit unequal power in the A and E channels, depending on the source extrinsic parameters (as reported in Sec. III). We leave a characterization of multi-channel analysis to strengthen detections and remove false positives to a future version of the algorithm. Setting the value of FAP_{thr} entails a trade-off between the number of false detections and the earliest time at which sources can be identified. The optimal choice depends on whether we prioritize a conservative search with fewer false positives or aim to maximize our sensitivity to faint signals.

We apply this detection strategy in parallel to the TDI channels A and E. The response of the two channels to a given source differs depending on its position relative to the detector, as encoded in their antenna pattern functions [76]. In current analysis, searches in the two channels are conducted independently. The activated regions in the resulting FAP maps are then compared, and a flag is raised when overlapping regions are identified in both channels, following the same set of rules described in Sec. II C.

III. RESULTS

The algorithm successfully identifies all 15 massive black hole binaries included in the SangriaHM dataset. Figure 7 illustrate the evolution of the t_c and $\mathcal{M}_c$ estimates for the MBHB-0 signal, while in Appendix A we systematically illustrate results for all sources in the SangriaHM dataset.

Detection and tracking performance vary across the 15 sources. All binaries but MBHB-2² are detected before merger, with MBHB-0, 3, 4, 10, 11, and 13 identified between 2 and 17 days in advance. For these sources, the best estimates for $\mathcal{M}_c$ and t_c overlap with the true values with an accuracy up to two grid steps a few days before merger. In the case of premerger detections, FAP minima typically approach the true values as more data are accumulated. Signals with a more prominent inspiral phase

²MBHB-2 requires an adaptation of the algorithm configuration, discussed further in the text.




FIG. 7. Evolution of $\mathcal{M}_c$ and t_c estimates for the signal MBHB-0, shown in top and bottom panel, respectively. Results are shown relative to the true injection, hence the red horizontal lines denote them. The black solid lines denote the best estimate, corresponding to the minimum FAP value, at each t^* . Shaded regions denote the point estimates uncertainties, as determined by the region tracking scheme described in Sec. II C; decreasingly lighter shades correspond to regions where $\log_{10}$ FAP is within 95%, 90%, 70%, and 50% of the minimum value, respectively. For reference, a few red markers label the minimum FAP values achieved at corresponding t^* .

yield the highest precision on $\mathcal{M}_c$ estimates at merger time,³ with relative errors below 3% (MBHB-0, 3, 4, 10, 13). MBHB-7 and 11 exhibit fractional uncertainties between 4% and 7%, while signals with weaker inspiral—such as MBHB-2, 5, 6, 8, and 14—show uncertainties ranging from 20% to 40%. Finally, MBHB-1 is recovered with the highest uncertainty, reaching a relative error of 138%. These discrepancies may arise from the strong approximations we make in signal modeling. However, we do not consider this a critical issue, as the goal of our method is to provide a rapid and approximate source parameters estimate, to serve as prior constraints to downstream algorithms performing more accurate parameter estimation.

Depending on their parameters—most notably the detector frame masses—MBHBs enter the LISA band and accumulate signal-to-noise ratio (SNR) at different frequencies. In Fig. 8, we show the evolution of the cumulative SNR for all sources and highlight the SNR values at the time of first detection. Signals such as MBHB-1, MBHB-2,

³The estimates of $\mathcal{M}_c$ and t_c^* continue to evolve after the merger time t_c as additional data become available, since the spectrograms do not extend up to the current timestamp t^* , as described in Sec. II A.




FIG. 8. Evolution of SNR for every source in the SangriaHM dataset. (Left panel) Cumulative SNR for each source, as the mission progresses through the simulated year of data. The computation is carried out adapting the code base in Spadaro *et al.* [94]. We employ the IMRPhenomXHM approximant [95] to produce waveforms analogous to the signals in SangriaHM, along with an analytical model of the instrumental and DWD confusion noise. Individual signals, numbered by gray boxes, are simulated up to their time-of-coalescence (dashed gray lines), and a star marker denotes the time of first detection by our algorithm. (Right panel) SNR value at the time of first detection, for a threshold $\text{FAP}_{\text{thr}} = 10^{-6}$. The first-detection timestamps shown in the figure are consistent with the results presented in Figs. 10 and 11; the default algorithm configuration is used for the majority of signals, while MBHB-1 and MBHB-2 are analyzed using a time–frequency configuration with a shorter STFT window, as described in Sec. II B. As discussed in Sec. III, MBHB-2 is identified 1.75 hours after coalescence, hence is not shown in the right panel.

and MBHB-14 have large chirp masses ($\mathcal{M}_c > 10^6 M_\odot$). For these systems, the accumulated SNR rises steeply, and detection occurs only close to, or slightly after, the merger. Adapting the configuration of the time–frequency representation can be beneficial for the detection of such signals. In the case of MBHB-1 and MBHB-2, we decrease the frequency resolution from 10^{-5} Hz to 2×10^{-4} Hz, consequently increasing the time resolution and reducing spectral leakage at higher frequencies, thereby enhancing the late-inspiral part of the signal. This strategy is described in Sec. II B and illustrated in the bottom panel of Fig. 3. The results shown in Figs. 10 and 11 relative to these signals are obtained using such dedicated configuration.

With our default configuration, MBHB-2 is not detected, as its signal overlaps significantly with the stronger emissions of MBHB-3 and MBHB-4, particularly at lower frequencies. We identified a strategy to mitigate such overlap in time–frequency by adjusting the spectrogram frequencies. Our default setup spans frequencies between 10^{-4} Hz and 10^{-3} Hz. By shifting to a higher-frequency band (5.0×10^{-4} Hz $< f < 5.0 \times 10^{-3}$ Hz), we effectively suppress the loud inspiral portions of MBHB-3 and MBHB-4. Moreover, as discussed above, we reduce the frequency resolution at the advantage of an increased time resolution. In Fig. 9 we show the FAP maps obtained for these different configurations, demonstrating that MBHB-2 becomes clearly separable from other signals and is successfully detected. However, the signal is correctly detected 1.75 hours after coalescence. Although a region of the parameter space exceeds the detection threshold before merger, the corresponding parameter estimates are significantly biased, and the region is not tracked further due to the lack of overlapping points at

subsequent timestamps. This demonstrates that running multiple searches in parallel, each optimized for different frequency ranges, may be necessary during operations to address source overlap. The minimal numerical burden of our search allows the adoption of such strategy remaining well below a minute, i.e., the envisioned low-latency analysis timescales.

For each source in the dataset, we additionally compare the evolution of $\min \log_{10} \text{FAP}_{\text{A,E}}$, i.e., the FAP minima as recorded by the region-tracking algorithm in the two TDI channels. Depending on the source parameters, their ratio




FIG. 9. Comparison of FAP maps for the same data chunk under different time–frequency configurations. The left panel shows the default settings, while the right panel uses a higher-frequency band with lower frequency resolution. In the left panel, the regions activated by MBHB-3 and MBHB-4 are clearly visible, but no distinct peak appears at the location of MBHB-2. In contrast, the FAP map on the right displays a discernible low-FAP region corresponding to MBHB-2.

typically spans up to two orders of magnitude, with $0.1 < \min \log_{10} \text{FAP}_A / \min \log_{10} \text{FAP}_E < 10$. For MBHB-1, MBHB-2, and MBHB-13, the measured power in channel E is consistently higher than in channel A throughout the tracking. The opposite behavior is observed for MBHB-6, MBHB-7, and MBHB-10, while MBHB-5 is not detected in channel E. The observed differences in measured power across channels lead to different first-detection times, with offsets of up to ± 4 days. For all other signals, we measure a comparable amount of power in both channels.

We assessed the algorithm computational performance by measuring the runtime of its main components on a single core. The workflow consists of two main stages; the computation of probability maps and the region-tracking algorithm. Based on 200 repetitions, we recorded a mean computing time for a single spectrogram of 0.039 ± 0.004 s. Generating probability maps is more demanding, with an average runtime of 0.43 ± 0.03 s per map, i.e., the dominant computational cost of the algorithm. The total processing time for the region-tracking algorithm varies depending on the number of regions being simultaneously tracked. The average runtime per chunk is 0.14 ± 0.06 s.

An issue identified during our analysis is the spurious activation of regions in the FAP maps caused by the loud signals of already merged sources. When the merger phase of an MBHB enters the spectrogram, a large amount of power is concentrated in the high-frequency chirp portion. Slices of the $(t_c, \mathcal{M}_c)$ grid intersecting such signal may exceed the detection threshold in $\langle \text{PSD} \rangle_g$, even though they merely cross the signal rather than following its chirping evolution. These merger-induced activations appear in the FAP maps as new regions characterized by low-chirp masses and may be mistakenly interpreted as distinct sources. For instance, following the merger of MBHB-0, approximately 20 spurious regions are activated. This issue can be mitigated through a masking strategy. Once a source is confidently identified (e.g., $\text{FAP} \approx 10^{-10}$), we remove the time-frequency pixels associated with slices in a small neighborhood around the minimum-FAP point of the tracked region. This procedure effectively removes the tracked signal from the time-frequency plane, while leaving the computation of $\langle \text{PSD} \rangle_g$ over other slices unaffected. We tested this approach and verified that such spurious regions are efficiently suppressed.

Spurious region activation may also occur in other situations. In rare cases, a tracked region may fragment due to statistical fluctuations, temporarily producing two nearby regions; however, the spurious region is typically short-lived and disappears within a few hours. In the presence of multiple overlapping signals, their interaction may also generate spurious regions in the FAP map. This is illustrated by the case of MBHB-3 and MBHB-4, which have similar chirp masses; their corresponding regions become clearly separable only about 5 days before the

merger of MBHB-3, during which approximately 15 spurious regions are activated. We verified that these can again be mitigated using the masking strategy described above. Since MBHB-4 is the loudest signal in the dataset, it is detected before MBHB-3 despite merging later. Applying the masking to MBHB-4 safely removes both its signal and the spurious regions arising from its confusion with MBHB-3, enabling the detection of MBHB-3 up to 8 days before merger.

Finally, we offer a few considerations regarding glitches. In the time-frequency domain, known glitches appear as short-lived (minutes to hours), high-amplitude transients spanning a broad frequency range, potentially mimicking the merger phase of an MBHB coalescence [73]. Although a systematic study of this effect is deferred to future work, we expect that, in the current configuration, unflagged glitches could be misidentified as loud, high-mass signals, which are characterized by the rapid inspiral described above. Aside from this potential mislabeling, we anticipate the algorithm to be fairly robust: a glitch occurring during the early inspiral of an MBHB produces a clearly distinguishable time-frequency signature, and we have shown that the algorithm can handle overlapping sources in that regime.

IV. CONCLUSIONS

In this work, we presented a fast prototype algorithm to detect MBHBs with LISA in low-latency. The main objective is to issue early alerts with approximate chirp mass and coalescence time estimates, potentially triggering protected observational periods and providing informative priors for more detailed parameter-estimation algorithms.

Our detection strategy is based on the search for excess power with the morphology of a chirping signal in time-frequency data. We compute the STFT of 10-day data segments and accumulate the power over pixels defined by the quadrupole frequency evolution. The algorithm outputs a significance map which measures the FAP of the average power inside each slice, relative to a background distribution. The background distribution is constructed from data segments containing only instrumental noise and Galactic sources, and is phenomenologically fitted with a Gamma function. At each timestep, the significance maps are reprocessed by a region-tracking algorithm: it targets points in the $(\mathcal{M}_c, t_c)$ space exceeding a significance threshold, labels them as potential sources, and follows their temporal evolution while progressively refining the estimates of $\mathcal{M}_c$ and t_c . Multiple potential sources can therefore be tracked simultaneously from a single chunk of data. A complete analysis for a single 10-days-long chunk of data is typically complete in 0.4 s using a single core.

For sources characterized by long inspirals, detected a few days before merger, parameter estimates yield typical fractional errors of 1–40% in chirp mass, and a few hours in coalescence time. The majority of MBHB sources in the

SangriaHM dataset are detected premerger, typically hours to days before coalescence. The signal of a particularly massive and high-redshift source is detected only after merger, primarily due to overlap with two other signals exhibiting louder inspiral phases. Some signals required tuning of the time-frequency configuration for identification, as discussed in Sec. III. In a realistic scenario, such tuning would be replaced by global optimization over multiple simulated datasets, producing a set of search configurations to be run in parallel for a complete exploration of data.

In Davies *et al.* [44], the authors adapt several well-established methods from ground-based detectors to the LISA context for both the search and parameter estimation of MBHB merger signals. The search algorithm is based on matched filtering against an optimally populated template bank of premerger waveforms, allowing them to confidently detect signals with different morphologies once they reach $\text{SNR} \sim 8$. The algorithm proposed was tested on a set of five representative MBHBs with $7.3 \times 10^5 M_\odot < \mathcal{M}_c < 8.7 \times 10^6 M_\odot$ (detector frame), and $0.125 < q < 1$. In our approach, as shown in Fig. 8, SangriaHM sources first detection does not occur at a fixed SNR threshold. Instead they take place over a fairly broad range of values, approximately $4 \lesssim \text{SNR} \lesssim 30$.

The results obtained for the sources in the SangriaHM dataset are consistent with those reported in Deng *et al.* [43], which presented the first application of a premerger detection algorithm on the Sangria dataset (which differs from SangriaHM by having only the (2, 2) harmonic present in the signals of MBHBs). The two approaches are fundamentally different, as the method presented in Deng *et al.* [43] is based on the maximization of the $\mathcal{F}$ -statistic on the intrinsic parameters of the source. The algorithm in Deng *et al.* [43] is able to detect every source in the Sangria dataset previous to its merger. However, our algorithm provides useful preliminary estimates of the chirp mass and coalescence time with a computational cost roughly two orders of magnitude lower. In fact, the computational cost of the end-to-end analysis of single 10-day data chunk is currently dominated by FAP map generation (~ 0.4 s) and region tracking (~ 0.14 s). By contrast, the algorithm presented in Deng *et al.* [43] currently requires ~ 15 min to analyze 15 days of data, where the main computational cost resides in the waveform evaluation.

In the current implementation, our algorithm is applied independently to *A* and *E* TDI channels. In future work we will investigate a detection statistic coherently combining information from multiple channels.

As discussed in Sec. II B, the background distributions are estimated under the idealized assumption that all transient signals occurring before the last 10 days of data have been detected and subtracted. It is foreseeable that an online estimate of the PSD will be provided by other blocks

of the LISA ground-segment (e.g., the noise-estimation global-fit blocks), possibly including an accurate characterization of the time-evolving confusion noise [67].

Alternative time–frequency representations could replace the STFT without altering the algorithm logic and structure. Notably, the constant-Q transform [59,96] have been shown to provide a multiresolution time-frequency representation that can reduce the impact of spectral leakage effects. Preliminary analyses suggest that using the constant-Q transform may allow better constraints of t_c and $\mathcal{M}_c$. A systematic comparison between the two methods is left for future work.

The time-frequency domain offers a straightforward way to handle data gaps, as the flexible selection of time–frequency bins allows the isolation of those affected by such artifacts. We plan to implement an automated procedure to mitigate their impact on detection and preliminary parameter estimation.

Currently, the algorithm discards information from the amplitude and phase evolution of the signals, which may encode crucial information about source extrinsic parameters. Future developments incorporating the analysis of such features are under investigation for an improved source characterization including critical, additional source parameters like their sky localization.

ACKNOWLEDGMENTS

The authors are grateful to Q. Baghi, F. Pozzoli, A. Spadaro, J. B. Bayle, C. Cossou, R. Srinivasan, L. Copparoni, F. Rigamonti, M. Le Jeune, S. Deng, S. Babak, and D. Bandopadhyay for discussions and valuable inputs. R. B. is supported by the Italian Space Agency grant “Phase B2/C activity for LISA mission,” Agreement No. 2024-NAZ-0102/PE, MUR Grant “Progetto Dipartimenti di Eccellenza 2023–2027” (BiCoQ), the ICSC National Research Centre funded by NextGenerationEU. Computational work was performed at CINECA with allocations through INFN and Bicocca.

DATA AVAILABILITY

There are no publicly available research data or software supporting this manuscript. Requests for further information or data should be sent to the authors.

APPENDIX A: $\mathcal{M}_c$ AND t_c ESTIMATES ACROSS ALL SANGRIA SOURCES

In Fig. 10 and 11, we show the evolution of the $\mathcal{M}_c$ and t_c estimates for all Sangria sources.

The results shown for MBHB-1, MBHB-13, and MBHB-14 are obtained from the search performed using channel E, as the measured power is consistently higher than in channel A, as discussed in Sec. II C. In channel A, sources MBHB-1 and MBHB-14 are detected only after merger.




FIG. 10. Evolution of $\mathcal{M}_c$ and t_c estimates across the first nine Sangria sources, as a function of the data processed up to t^* . Each subplot corresponds to one source. Red horizontal lines in each subplot correspond to the true injected values, i.e., zero relative differences with respect to our estimates. Blue vertical lines marks the merger time for each system, and black solid curves shows the best estimate at each t^* . Shaded regions denote the point estimates uncertainties; decreasingly lighter shades correspond to regions where $\log_{10}$ FAP is within 95%, 90%, 70%, and 50% of the minimum value, respectively. Parameter estimates for MBHB-2 are tracked from 1.75 hours after coalescence, while all other signals are detected before merger. Signals with a more prominent inspiral phase yield the highest precision on $\mathcal{M}_c$ estimates at merger time ($\Delta \mathcal{M}_c / \mathcal{M}_c \lesssim 3\%$), while signals with weaker inspiral show uncertainties ranging from 20% to 40%. The highest relative error is registered for MBHB-1 ($\Delta \mathcal{M}_c / \mathcal{M}_c = 138\%$). Preliminary constraints of the coalescence time within a few hours with respect to the true value are generally obtained days before merger (hours in the case of signals characterized by shorter inspirals).




FIG. 11. Evolution of $\mathcal{M}_c$ and t_c estimates across the last six Sangria sources, following the same conventions in Fig. 10.

The panel corresponding to MBHB-4 shows estimates extending to approximately ~ 8 days before merger. In fact, this signal is already detected about ~ 17 days before merger; however, a mislabeling issue arises during the tracking of the corresponding region in the FAP map. In Fig. 12, we present results obtained with the same analysis without enforcing the retention of only overlapping points in the region-tracking described in Sec. II C. In this case, the algorithm continuously tracks the signal from the first detection through merger. Nevertheless, during the first ~ 7 days of evolution the tracking is affected by significant biases due to the interaction with the overlapping signal MBHB-3.

Furthermore, we observe a qualitatively different behavior of the FAP contour levels compared to the case in which the retention of only overlapping points is enforced. In the present configuration, the activated region is no longer bounded by the initial detection estimates, making the contour levels more susceptible to oscillations. Focusing on the final ~ 5 days before merger, we find that all contour levels associated with the estimate of t_c become

progressively more constrained with time. A different behavior is observed for $\mathcal{M}_c$; although the 90%–95% contours around the FAP minimum remain relatively well confined, the 50%–70% contours broaden toward larger values. This is due to the higher-frequency portion of the signal becomes dominant over the inspiral, close to merger. As a result, the activated regions in the FAP maps become increasingly elongated in the $\mathcal{M}_c$ direction while simultaneously tightening in the t_c direction, as illustrated in Fig. 5.

Both configurations of the tracking algorithm represent valid alternatives. Enforcing the retention of only overlapping points across successive timesteps yields a more constrained evolution in the $\mathcal{M}_c$ – t_c space, at the cost of potentially discarding information when the initial estimates are biased. Conversely, allowing regions to be tracked without enforcing overlap produces evolving estimates that more faithfully reflect the instantaneous structure of the FAP maps. While the t_c estimates naturally become more constrained as the merger is approached, a sensible strategy for $\mathcal{M}_c$ is to interpret the 90%–95%




FIG. 12. Evolution of $\mathcal{M}_c$ and t_c estimates for MBHB-4, following conventions in Figs. 10 and 11. Results are obtained without retaining overlapping points, only. The signal is seen first approximately ~ 17 days before merger, whereas enforcing the overlap criterion restricts tracking to about ~ 8 days before merger. Nevertheless, during the first ~ 7 days of evolution the tracking is affected by biases induced by the overlapping signal MBHB-3.

contours around the FAP minimum as meaningful measures of uncertainty.

APPENDIX B: SLICE-AVERAGED PIXEL POWER BACKGROUND DISTRIBUTION

We provide here a short derivation of results relevant to justify the choice of parametric fit for the background distribution of slice-averaged power, i.e., $\langle \widehat{\text{PSD}} \rangle_{\mathcal{H}_0}$. As detailed in Sec. II, spectrograms are defined by Hann-windowed short Fourier transforms,

$$S[p, q] = \sum_{k=0}^{N-1} w[k - ph] x[k] e^{-i2\pi qk/N}, \quad (\text{B1})$$

suitably whitened by the estimator of the quasi

stationary $\text{PSD}_0[p; t_u]$ in Eq. (6). For completeness we report here the value of normalization constant for the Hann window; using the trigonometric identity $\cos^2 \theta = \frac{1}{2}(1 + \cos 2\theta)$ we obtain from Eq. (4),

$$w[n]^2 = \frac{1}{4}(1 - 2\cos \theta_n + \cos^2 \theta_n) = \frac{3}{8} - \frac{1}{2}\cos \theta_n + \frac{1}{8}\cos(2\theta_n), \quad (\text{B2})$$

$$\theta_n = 2\pi n/(N-1), \quad (\text{B3})$$

which yields

$$\|w\|^2 = \sum_{n=0}^{N-1} w[n]^2 = \frac{3}{8}(N-1). \quad (\text{B4})$$

Consequently, $S[p, q]$ is a circular complex Gaussian random variable with zero mean,

$$\mathbb{E}[S[p, q]] = \sum_{k=0}^{N-1} w[k - ph] \mathbb{E}[x[k]] e^{-i2\pi qk/N} = 0, \quad (\text{B5})$$

and second moment

$$\mathbb{E}[|S[p, q]|^2] = \mathbb{E}\left[\left(\sum_{k=0}^{N-1} x[k] w[k - ph] e^{-i2\pi qk/N}\right) \times \left(\sum_{\ell=0}^{N-1} x[\ell] w[\ell - ph] e^{-i2\pi q\ell/N}\right)^*\right] \quad (\text{B6})$$

$$= \sum_{k=0}^{N-1} \sum_{\ell=0}^{N-1} w[k - ph] w[\ell - ph] e^{-i2\pi q(k-\ell)/N} \mathbb{E}[x[k] x[\ell]] \quad (\text{B7})$$

$$= \sigma_x^2 \sum_{k=0}^{N-1} w^2[k - ph] = \sigma_x^2 \|w\|^2 \quad (\text{B8})$$

$$= \frac{3\sigma_x^2}{8}(N-1), \quad (\text{B9})$$

where $\mathbb{E}[\cdot]$ denotes the expectation value operator. The amplitude $R \equiv |S|$ is therefore Rayleigh-distributed, with probability density specified by the parameter σ_r ,

$$p_R(r; \sigma_r) = \frac{r}{\sigma_r^2} \exp\left(-\frac{r^2}{2\sigma_r^2}\right), \quad (\text{B10})$$

$$\sigma_r^2 = \mathbb{E}[|S[p, q]|^2]/2. \quad (\text{B11})$$

For reproducibility, we note that Eq. (B10) is consistent with the Rayleigh distribution implemented in `scipy.stats`, with the scale parameter chosen as $\sigma_r \|w\|^2 / \sqrt{f_s} = \sigma_x / \sqrt{2f_s}$.

Equation (B10) represents the probability distribution of a single pixel amplitude spectral density, white Gaussian noise realisations. We verify the correctness of our result by generating white Gaussian noise STFTs from nonoverlapping segments and collecting samples of a given pixel in each spectrogram. A comparison between the analytical model and the numerically simulated one is shown in Fig. 13. We obtain the distribution of STFT pixel power, including the normalization factor $T_s/\|w\|^2$ from Eq. (2), i.e., an exponential distribution with rate parameter λ ,

$$p(s; \lambda) = \lambda e^{-\lambda s} s > 0, \quad (\text{B12})$$




FIG. 13. Comparisons between analytical models and empirical distributions for STFT pixel amplitude spectral density (left), power spectral density (middle) and slice averaged power (right). The teal curves in left and middle panel denote the Rayleigh and exponential distributions predicted in Eqs. (B10) and (B12), compared against empirical distributions of simulated spectrogram pixels (gray histograms). (Right panel) Empirical and analytical distributions of slice STFT power averaged over 6 (26, 77) pixels in orange (teal, purple), shown as histograms and solid curves, respectively. The prediction of Gamma distributed values in Eq. (B18) justifies our choice of parametric fit shown in Fig. 4.

$$\lambda = \frac{8\sigma_x^2}{3N(N-1)}, \quad (\text{B13})$$

or equivalently by the scale parameter $\theta = 1/\lambda$ for the `scipy.stats` implementation. To determine how the average slice power in Eq. (11) is distributed, we define it as

$$p_{\langle \text{PSD} \rangle_s}(x) = \int_0^\infty \cdots \int_0^\infty \delta\left(x - \frac{1}{N_s} \sum_{i=1}^{N_s} x_i\right) \prod_{i=1}^{N_s} p(x_i; \lambda) dx_i, \quad (\text{B14})$$

and obtain its Laplace transform,

$$\mathcal{L}\{p_{\langle \text{PSD} \rangle_s}\}(t) = \int_0^\infty e^{-ts} p_{\langle \text{PSD} \rangle_s}(s) ds \quad (\text{B15})$$

$$= \prod_{i=1}^{N_s} \int_0^\infty e^{-tx_i/N_s} p(x_i; \lambda) dx_i \quad (\text{B16})$$

$$= \left(\frac{\lambda N_s}{\lambda N_s + t} \right)^{N_s}. \quad (\text{B17})$$

Through inverse Laplace transform we recover the explicit expression for the probability distribution,

$$p_{\langle \text{PSD} \rangle_s}(s; \lambda, N_s) = \frac{(\lambda N_s)^{N_s}}{\Gamma(N_s)} s^{N_s-1} e^{-s\lambda N_s}, \quad s \geq 0, \quad (\text{B18})$$

which is a Gamma distribution with shape parameter N_s and scale parameter $1/\lambda N_s$. This result supports the choice of a Gamma distribution as a parametric fit in Sec. II B. Nonetheless, we point out that the chosen STFT therein uses overlapping window segments with fixed hop size $h < N$, thereby violating the independence of pixels

assumed in Eq. (B14), while preserving individual ones Gaussianity. We account for it computing the covariance between two pixels specified by p, q and p', q' , respectively, i.e.,

$$\begin{aligned} \mathbb{E}[S[p, q]S[p', q']^*] &= \mathbb{E}\left[\left(\sum_{k=0}^{N-1} w[k-ph]x[k]e^{-i2\pi qk/N}\right) \right. \\ &\quad \times \left.\left(\sum_{\ell=0}^{N-1} w[\ell-p'h]x[\ell]e^{-i2\pi q'\ell/N}\right)^*\right] \end{aligned} \quad (\text{B19})$$

$$= \sum_{k=0}^{N-1} \sum_{\ell=0}^{N-1} w[k-ph]w[\ell-p'h]e^{-i2\pi(qk-q'\ell)/N} \mathbb{E}[x[k]x[\ell]^*] \quad (\text{B20})$$

$$= \sigma_x^2 \sum_{k=0}^{N-1} w[k-ph]w[k-p'h]e^{-i2\pi(q-q')k/N}. \quad (\text{B21})$$

A complete expression for the covariance can be obtained numerically for specific windows and segment overlaps. To take such correlation into account, we note that our aim is to construct the survival function for the slice-averaged pixel power. As just shown, pixel STFT values of a slice $\mathfrak{s}$ will be jointly-distributed according to a multivariate complex Gaussian, with covariance $\Sigma_{\mathfrak{s}}$, i.e.,

$$\mathbf{s} = \{S[p_i, q_i], (p_i, q_i) \in \mathfrak{s}\} \sim \mathcal{CN}(0, \Sigma_{\mathfrak{s}}). \quad (\text{B22})$$

Once a slice is chosen, the covariance can be diagonalized through eigenvalue decomposition $\Sigma_{\mathfrak{s}} = U\Lambda U^\dagger$, where U is unitary and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_{N_s})$ is the matrix of eigenvalues. The detection statistic $\langle \text{PSD} \rangle_s$ in Eq. (11) is defined as the squared norm of $\|\mathbf{s}\|^2 = \mathbf{s}^\dagger \mathbf{s}$, scaled by a conventional normalization constant. Knowing that the

auxiliary variables $\mathbf{y} = U^\dagger \mathbf{s}$ are uncorrelated complex Gaussian random variables,

$$y_i \sim \mathcal{CN}(0, \sigma_i^2), \quad \sigma_i^2 = \lambda_i, \quad (\text{B23})$$

and that unitary transformations preserve the squared norm,

$$\langle \widehat{\text{PSD}} \rangle_{\mathbf{s}} \propto \|\mathbf{s}\|^2 = \|\mathbf{y}\|^2 = \sum_{i=1}^{N_{\mathbf{s}}} y_i^2, \quad (\text{B24})$$

we finally obtain for the survival function,

$$\text{Prob}(\langle \widehat{\text{PSD}} \rangle_{\mathbf{s}} > s | \mathcal{H}_0) = \sum_{k=1}^{N_{\mathbf{s}}} \left(\exp \left(-\frac{s}{\sigma_k^2} \right) \prod_{\substack{j=1 \\ j \neq k}}^{N_{\mathbf{s}}} \frac{\sigma_j^2}{\sigma_j^2 - \sigma_k^2} \right), \quad s \geq 0. \quad (\text{B25})$$

-
- [1] P. Amaro-Seoane, H. Audley, S. Babak, J. Baker *et al.*, [arXiv:1702.00786](#).
 - [2] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham *et al.*, *Living Rev. Relativity* **23**, 3 (2020).
 - [3] A. Abac, R. Abramo, S. Albanesi, A. Albertini *et al.*, *J. Cosmol. Astropart. Phys.* **03** (2026) 081.
 - [4] D. Reitze, R. X. Adhikari, S. Ballmer, B. Barish, L. Barsotti *et al.*, *Bull. Am. Astron. Soc.* **51**, 35 (2019).
 - [5] J. P. W. Verbiest, L. Lentati, G. Hobbs, R. van Haasteren, and P. B. Demorest, *Mon. Not. R. Astron. Soc.* **458**, 1267 (2016).
 - [6] M. Colpi, K. Danzmann, M. Hewitson, K. Holley-Bockelmann, P. Jetzer *et al.*, [arXiv:2402.07571](#).
 - [7] P. Amaro-Seoane, J. Andrews, M. Arca Sedda, A. Askar *et al.*, *Living Rev. Relativity* **26**, 2 (2023).
 - [8] S. Deng, S. Babak, M. Le Jeune, S. Marsat, É. Plagnol, and A. Sartirana, *Phys. Rev. D* **111**, 103014 (2025).
 - [9] M. L. Katz, N. Karnesis, N. Korsakova, J. R. Gair, and N. Stergioulas, *Phys. Rev. D* **111**, 024060 (2025).
 - [10] T. B. Littenberg and N. J. Cornish, *Phys. Rev. D* **107**, 063004 (2023).
 - [11] S. H. Strub, L. Ferraioli, C. Schmelzbach, S. C. Stähler, and D. Giardini, *Phys. Rev. D* **110**, 024005 (2024).
 - [12] G. Astorino, L. V. Dall'Armi, R. Buscicchio, J. Pomper, A. Ricciardone, and W. D. Pozzo, *Eur. Phys. J. C* **86**, 395 (2026).
 - [13] L. P. Singer and L. R. Price, *Phys. Rev. D* **93**, 024013 (2016).
 - [14] E. Marx, W. Benoit, A. Gunny, R. Omer, D. Chatterjee *et al.*, *Phys. Rev. D* **111**, 042010 (2025).
 - [15] B. P. Abbott *et al.* (MASTER Collaborations), *Nature (London)* **551**, 85 (2017).
 - [16] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. Lett.* **119**, 161101 (2017).
 - [17] B. P. Abbott *et al.*, *Astrophys. J. Lett.* **848**, L12 (2017).
 - [18] G. Hallinan, A. Corsi, K. P. Mooley, K. Hotokezaka, E. Nakar *et al.*, *Science* **358**, 1579 (2017).
 - [19] E. Troja, L. Piro, H. van Eerten, R. T. Wollaeger, M. Im *et al.*, *Nature (London)* **551**, 71 (2017).
 - [20] D. W. Weedman, *Annu. Rev. Astron. Astrophys.* **15**, 69 (1977).
 - [21] R. Abuter *et al.* (GRAVITY Collaboration), *Astron. Astrophys.* **615**, L15 (2018).
 - [22] R. C. E. van den Bosch, J. E. Greene, J. A. Braatz, A. Constantin, and C.-Y. Kuo, *Astrophys. J.* **819**, 11 (2016).
 - [23] T. Bogdanović, M. C. Miller, and L. Blecha, *Living Rev. Relativity* **25**, 3 (2022).
 - [24] C. Rodriguez, G. B. Taylor, R. T. Zavala, A. B. Peck, L. K. Pollack, and R. W. Romani, *Astrophys. J.* **646**, 49 (2006).
 - [25] A. Mangiagli, C. Caprini, M. Volonteri, S. Marsat, S. Vergani, N. Tamanini, and H. Inchauspé, *Phys. Rev. D* **106**, 103017 (2022).
 - [26] P. Saini, S. A. Bhat, and K. G. Arun, *Phys. Rev. D* **106**, 104015 (2022).
 - [27] G. Lops, D. Izquierdo-Villalba, M. Colpi, S. Bonoli, A. Sesana, and A. Mangiagli, *Mon. Not. R. Astron. Soc.* **519**, 5962 (2023).
 - [28] M. Dotti, F. Mannucci, R. Buscicchio, M. Colpi, Q. D'Amato *et al.*, [arXiv:2512.21359](#).
 - [29] Ž. Ivezić *et al.* (LSST Collaboration), *Astrophys. J.* **873**, 111 (2019).
 - [30] C. Xin and Z. Haiman, *Mon. Not. R. Astron. Soc.* **506**, 2408 (2021).
 - [31] C. Xin and Z. Haiman, *Mon. Not. R. Astron. Soc.* **533**, 3164 (2024).
 - [32] C. Xin, M. Isi, W. M. Farr, and Z. Haiman, *Astrophys. J.* **999**, 149 (2026).
 - [33] L. Bertassi, E. Sottocorno, F. Rigamonti, D. D'Orazio, M. Eracleous, Z. Haiman, and M. Dotti, *Astron. Astrophys.* **702**, A165 (2025).
 - [34] L. Piro, M. Colpi, J. Aird, A. Mangiagli, A. C. Fabian *et al.*, *Mon. Not. R. Astron. Soc.* **521**, 2577 (2023).
 - [35] L. Piro, M. Ahlers, A. Coleiro, M. Colpi, E. d. O. Wilhelm *et al.*, *Exp. Astron.* **54**, 23 (2022).
 - [36] N. J. Cornish, *Phys. Rev. D* **105**, 044007 (2022).
 - [37] M. L. Katz, *Phys. Rev. D* **105**, 044055 (2022).
 - [38] C. R. Weaving, L. K. Nuttall, I. W. Harry, S. Wu, and A. Nitz, *Classical Quantum Gravity* **41**, 025006 (2023).
 - [39] C. Hoy, C. R. Weaving, L. K. Nuttall, and I. Harry, *Classical Quantum Gravity* **41**, 245012 (2024).
 - [40] A. Sharma, A. S. Sengupta, and S. Mukherjee, *Phys. Rev. D* **111**, 042009 (2025).
 - [41] A. Jan, R. O'Shaughnessy, D. Shoemaker, and J. Lange, *Phys. Rev. D* **111**, 064079 (2025).
 - [42] J. E. Gammal, R. Buscicchio, G. Nardini, and J. Torrado, *Phys. Rev. D* **112**, 063010 (2025).

- [43] S. Deng, S. Babak, and S. Marsat, *Phys. Rev. D* **112**, 043010 (2025).
- [44] G. Cabourn Davies, I. Harry, M. J. Williams, D. Bandopadhyay, L. Barack *et al.*, *Phys. Rev. D* **111**, 043045 (2025).
- [45] R. Tenorio and D. Gerosa, *Phys. Rev. D* **111**, 104044 (2025).
- [46] W.-H. Ruan and Z.-K. Guo, *Phys. Rev. D* **109**, 123031 (2024).
- [47] N. Houba, S. H. Strub, L. Ferraioli, and D. Giardini, *Phys. Rev. D* **110**, 062003 (2024).
- [48] M.-C. Isfan, L.-I. Caramete, A. Caramete, D. Tonoiu, and A. Nicolin-Żaczek, *Classical Quantum Gravity* **42**, 225001 (2025).
- [49] L. Speri, R. Tenorio, C. Chapman-Bird, and D. Gerosa, *Phys. Rev. D* **113**, 024061 (2026).
- [50] D. Bandopadhyay, C. E. A. Chapman-Bird, and A. Vecchio, *arXiv:2510.19047*.
- [51] R. C. Barroso, Y. Lemi re, F. Mauger, and Q. Baghi, *Classical Quantum Gravity* **42**, 115016 (2025).
- [52] S. Chatterji, L. Blackburn, G. Martin, and E. Katsavounidis, *Classical Quantum Gravity* **21**, S1809 (2004).
- [53] S. Klimenko *et al.*, *Phys. Rev. D* **93**, 042004 (2016).
- [54] A. Martini *et al.*, *Classical Quantum Gravity* **43**, 055016 (2026).
- [55] K. S. Tai, S. T. McWilliams, and F. Pretorius, *Phys. Rev. D* **90**, 103001 (2014).
- [56] N. Bose and A. Pai, *Phys. Rev. D* **104**, 124021 (2021).
- [57] J. Fernandes, P. Tiwari, and A. Pai, *arXiv:2603.09150*.
- [58] R. Tenorio, D. Keitel, and A. M. Sintes, *Universe* **7**, 474 (2021).
- [59] S. Chatterji, L. Blackburn, G. Martin, and E. Katsavounidis, *Classical Quantum Gravity* **21**, S1809 (2004).
- [60] S. Klimenko, I. Yakushin, A. Mercer, and G. Mitselmakher, *Classical Quantum Gravity* **25**, 114029 (2008).
- [61] V. Necula, S. Klimenko, and G. Mitselmakher, *J. Phys. Conf. Ser.* **363**, 012032 (2012).
- [62] N. J. Cornish and T. B. Littenberg, *Classical Quantum Gravity* **32**, 135012 (2015).
- [63] S. Klimenko, G. Vedovato, V. Necula, F. Salemi, M. Drago *et al.*, *cwb pipeline library: 6.4.1*, Zenodo (2021).
- [64] M. J. Szczepa czyk, F. Salemi, S. Bini, T. Mishra, G. Vedovato *et al.*, *Phys. Rev. D* **107**, 062002 (2023).
- [65] A. Licciardi, D. Carbone, L. Rondoni, and A. Nagar, *Phys. Rev. D* **111**, 084044 (2025).
- [66] G. Nelemans, *Classical Quantum Gravity* **26**, 094030 (2009).
- [67] R. Buscicchio, F. Pozzoli, D. Chirico, and A. Sesana, *arXiv:2511.03604*.
- [68] F. Pozzoli, R. Buscicchio, A. Klein, V. Korol, A. Sesana, and F. Haardt, *Phys. Rev. D* **111**, 063005 (2025).
- [69] A. Antonelli, O. Burke, and J. R. Gair, *Mon. Not. R. Astron. Soc.* **507**, 5069 (2021).
- [70] M. Armano *et al.*, *Phys. Rev. Lett.* **116**, 231101 (2016).
- [71] M. Armano *et al.*, *Phys. Rev. Lett.* **120**, 061101 (2018).
- [72] E. Castelli, Q. Baghi, J. G. Baker, J. Slutsky, J. Bobin, N. Karnesis, A. Petiteau, O. Sauter, P. Wass, and W. J. Weber, *Classical Quantum Gravity* **42**, 065018 (2025).
- [73] A. Spadaro, R. Buscicchio, D. Vetrugno, A. Klein, S. Vitale, R. Dolesi, W. J. Weber, and M. Colpi, *Phys. Rev. D* **108**, 123029 (2023).
- [74] M. Armano, H. Audley, J. Baird, P. Binetruy, M. Born *et al.*, *Phys. Rev. D* **106**, 062001 (2022).
- [75] O. Burke, S. Marsat, J. R. Gair, and M. L. Katz, *Phys. Rev. D* **111**, 124053 (2025).
- [76] M. Tinto and S. V. Dhurandhar, *Living Rev. Relativity* **24**, 1 (2020).
- [77] M. Le Jeune and S. Babak, LISA Data Challenge Sangria (LDC2a) (2022).
- [78] Q. Baghi, *arXiv:2204.12142*.
- [79] M. Tinto, S. Dhurandhar, and D. Malakar, *Phys. Rev. D* **107**, 082001 (2023).
- [80] LISA Science Study Team, LISA science requirements document (2018).
- [81] S. Khan, S. Husa, M. Hannam, F. Ohme, M. P rrer, X. J. Forteza, and A. Boh , *Phys. Rev. D* **93**, 044007 (2016).
- [82] SciPy Developers, *SciPy user guide: Short-time Fourier transform*, <https://docs.scipy.org/doc/scipy-1.16.2/tutorial/signal.html#tutorial-stft> (2025).
- [83] A. Nuttall, *IEEE Trans. Acoust. Speech Signal Process.* **29**, 84 (1981).
- [84] SciPy Developers, *SciPy user guide: Sliding windows*, <https://docs.scipy.org/doc/scipy-1.16.2/tutorial/signal.html#sliding-windows> (2025).
- [85] S. Gottlieb, J.-H. Jung, and S. Kim, *Commun. Comput. Phys.* **9**, 497 (2011).
- [86] K. Chatziioannou, C.-J. Haster, T. B. Littenberg, W. M. Farr, S. Ghonge, M. Millhouse, J. A. Clark, and N. Cornish, *Phys. Rev. D* **100**, 104004 (2019).
- [87] M. Isi, S. Mastroianni, M. Pitkin, and O. J. Piccinni, *Phys. Rev. D* **102**, 123027 (2020).
- [88] M. Maggiore, *Gravitational Waves. Vol. 1: Theory and Experiments* (Oxford University Press, New York, 2007).
- [89] S. Marsat, J. G. Baker, and T. D. Canton, *Phys. Rev. D* **103**, 083011 (2021).
- [90] L. F. Richardson and W. F. Eddy, *arXiv:1807.07797*.
- [91] Q. Baghi, N. Karnesis, and J.-B. Bayle, *arXiv:2602.12781*.
- [92] J. R. Weaver, *Am. Math. Mon.* **92**, 711 (1985).
- [93] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy *et al.*, *Nat. Methods* **17**, 261 (2020).
- [94] A. Spadaro, R. Buscicchio, D. Izquierdo-Villalba, D. Gerosa, A. Klein, and G. Pratten, *Phys. Rev. D* **111**, 023004 (2025).
- [95] C. Garc a-Quir s, M. Colleoni, S. Husa, H. Estell s, G. Pratten, A. Ramos-Buades, M. Mateu-Lucena, and R. Jaume, *Phys. Rev. D* **102**, 064002 (2020).
- [96] A. Virtuoso and E. Milotti, *Phys. Rev. D* **109**, 102010 (2024).