

A hierarchical approach to matched filtering using a reduced basis

Rahul Dhurkunde, Henning Fehrmann, and Alexander H. Nitz

Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut), D-30167 Hannover, Germany and
Leibniz Universität Hannover, D-30167 Hannover, Germany

(Dated: October 26, 2021)

Searching for gravitational waves from compact binary coalescences (CBC) is performed by matched filtering the observed strain data from gravitational-wave observatories against a discrete set of waveform *templates* designed to accurately approximate the expected gravitational-wave signal, and are chosen to efficiently cover a target search region. The computational cost of matched filtering scales with both the number of templates required to cover a parameter space and the in-band duration of the waveform. Both of these factors increase in difficulty as the current observatories improve in sensitivity, especially at low frequencies, and may pose challenges for third-generation observatories. Reducing the cost of matched filtering would make searches of future detector's data more tractable. In addition, it would be easier to conduct searches that incorporate the effects of eccentricity, precession or target light sources (e.g. subsolar). We present a hierarchical scheme based on a reduced bases method to decrease the computational cost of conducting a matched-filter based search. Compared to the current methods, we estimate without any loss in sensitivity, a speedup by a factor of ~ 18 for sources with signal-to-noise ratio (SNR) of at least $= 6.0$, and a factor of 8 for SNR of at least 5. Our method is dominated by linear operations which are highly parallelizable. Therefore, we implement our algorithm using graphical processing units (GPUs) and evaluate commercially motivated metrics to demonstrate the efficiency of GPUs in CBC searches. Our scheme can be extended to generic CBC searches and allows for efficient matched filtering using GPUs.

I. INTRODUCTION

The first gravitational wave (GW) detection in 2015 marked the dawn of GW astronomy [1]. The first two observation runs of LIGO [2] and VIRGO [3] detectors (O1 and O2) reported over a dozen confident detections [4, 5]. The number of detections has rapidly increased to over 50 with the most recent O3 observing run [6, 7]. To date, all gravitational-wave observations have come from compact binary coalescences (CBC); the vast majority of sources were from binary black holes (BBH) [6, 7], but notably two binary neutron star mergers (NS) [8, 9], and recently two neutron star – black hole NSBH mergers [10] have been observed. These observations have helped us to understand the physics of compact objects [11, 12] and their dynamical evolution [11]. As the gravitational-wave observatories become more sensitive, the increased number of CBC sources will allow us to determine merger rate [13] and population distribution [14]. Upcoming third-generation observatories such as the Einstein telescope [15], cosmic explorer [16], and LISA [17] are expected to detect new kinds of astrophysical sources [15, 18–20].

Matched filtering is the most widely used technique to detect CBC signals [5, 21–23]. The method is optimal for stationary Gaussian noise [24]. While the detector data contains non-Gaussian noise transients [25, 26], which require the use of vetoing techniques [27, 28], matched filtering remains the dominant computational cost of a search algorithm [29]. In this work, we focus only on the implementation of matched filtering. Matched filtering requires accurate models of the expected gravitational waveform; CBCs can be modeled

using different techniques [30–32]. The parameters of a binary merger are categorized into intrinsic (e.g. masses and spins) and extrinsic (e.g. binary orientation and location). To search for sources with unknown intrinsic parameters, we must select a discrete bank of templates which span the parameter space. These templates are chosen such that the minimum match (MM) between the data and at least one template from the bank is sufficiently large [33, 34]. For example, a minimum match value of ~ 0.97 would imply that at least 97% of the SNR of any signal with parameters within the search area could be recovered. To identify a potential signal, gravitational-wave strain data is convolved with every template in a bank to calculate a signal-to-noise time series. Extrinsic parameter are often analytically maximized over. A possible candidate is identified if the SNR rises above a pre-determined threshold, passes various tests of signal consistency [27, 28] and data quality [25, 35, 36], and is statistically significant [4, 37]. The computational cost of matched filtering, and so also the entire search, scales linearly with the number of templates in a bank and also increases with the duration of the observable signal, though generally sub-linearly.

With advancements in the current and future detectors, it is expected that observation of signals at an increasingly lower frequency will become possible [15, 18, 19, 38]. As the low frequency cutoff of a search decreases, both the size of the template bank and the signal duration grows rapidly, leading to increased computational costs [38–40]. While current template-based searches have confident detections only from quasi-circular aligned-spin binaries [5, 6, 21, 22], sources that exhibit measurable eccentricity or precession of the orbital plane could provide unique astrophysical

insights [41, 42]. While a few searches have included the effect of eccentricity [39, 43, 44] for parts of parameter space, many searches neglect the effects from eccentricity of the orbital plane [45] and precession of the orbit [46] in part due to the increased computational cost relative to normal searches [7, 39, 40, 47]. For example, it has been shown that the template bank including precession is at least 10 times bigger than one without precession [47]. Furthermore, the lower mass boundary of subsolar primordial black hole searches is limited by computational cost considerations [7, 38–40]. Development of cost-efficient filtering algorithm would allow searches to be conducted more easily, with higher sensitivity, and in uncharted regions of parameter space.

The computational costs for matched filtering data with templates consist of redundant computations due to significant overlap of templates with each other in the neighborhood. This redundancy is eliminated by using an orthonormal bases to filter data instead of templates [48]. The costs of filtering scale linearly with the number of bases and can be reduced by rejecting the basis vectors of lower importance [48]. Disregarding contributions from a few basis vectors leads to a loss in SNR, but this loss is kept under a tolerance by tuning the number of relevant bases p involved for filtering purpose [48]. When considering large number of templates, the value of p is much smaller, and therefore, it is possible to filter data with a fewer number of bases. Current online (low-latency) searches [23, 49] employ this technique and are in good agreement with searches that do not use this approximation [50, 51].

Since the basis vectors do not correspond to any physical source, match filtering outputs from each basis are weighted and linearly combined to give an SNR time-series for a unique template, in a *reconstruction* process. The reconstruction is performed for each template and at every time sample which occurs additional costs to matched filtering. The naive costs for reconstructing SNR time-series for T templates can be estimated in terms of a matrix multiplication which scales as $\mathcal{O}(NpT)$, where p is the number of bases and N the number of samples in the data. On the other hand, the direct template-based filtering is widely done using a Fast Fourier Transform (FFT) - based algorithm which requires $\mathcal{O}(NT \log N)$ operations. Usually $\log N \sim \mathcal{O}(10)$ and $p \sim \mathcal{O}(10^2 - 10^3)$ which suggests that naive reduced bases filtering is more expensive than the template-based filtering.

In this work, we demonstrate a cost-efficient matched filtering method by employing a new hierarchical scheme using a reduced set of bases. The reduced bases are obtained by applying principal component analysis (PCA) on a template bank. A two-stage hierarchical scheme is then invoked to compute the SNR time series for each template. In the first stage, an intermediate time-series is computed that corresponds to a binned average of the complete SNR time-series. In the second step we, do a full time resolution (non-averaged)

reconstruction using the reduced bases inside the bins where the average SNR exceeds a threshold. We demonstrate our method on simulated Gaussian noise and a population of CBC signals. To estimate the improvement from our method, we compare it against the flat template based filtering scheme used in current searches. We observe that our method attains a speed up by a factor of ~ 8 for a threshold of $\text{SNR} = 5$. Furthermore, we expect the performance of our method to increase at higher SNR thresholds, and similarly observe a performance gain of ~ 18 times for $\text{SNR} = 6$.

Matched filtering is dominated by mathematical operations which are easily parallelizable across different threads or compute cores. Since a GPU is designed to accommodate a large number of threads, we employ GPUs in this work for an efficient implementation of matched filtering in parallel. To investigate the relative performance of different hardware, we compare the matched filtering implementations on GPUs with the central processing units (CPUs) that is currently used in the PyCBC search pipelines. We use two metrics to quantify the performance - cost, and power efficiency while filtering data. We observe that GPUs are much more efficient than CPUs in performing matched filtering. In this analysis, we have restricted ourselves to aligned-spins, dominant mode, and a single detector analysis. However, our method can be easily extended to multiple detectors, and search scenarios including eccentricity, precession, or higher-order modes.

The rest of the paper is organized as follows. In section II, we give a brief overview of the matched filtering process and motivate why we need a new approach different from previous works. In section III we describe our search method and the implementation. In section IV we present our results and compare them with the existing search methods. Finally we conclude in section V.

II. COMPACT BINARY COALESCING SEARCHES

The core of any modeled searches for CBC signals is the matched filtering technique which involves searching the interferometric data for a modeled waveform of the expected GW signal [5, 21–23]. In this section we briefly describe the process of matched filtering and introduce some ideas for efficient filtering algorithms. We also discuss previous efforts to improve the performance of matched filtering in section II B.

GW signals from a non-eccentric CBC sources are characterised by 15 parameters [30–32]. These parameters are divided in two categories, 1) Intrinsic parameters - (m_1, m_2) component masses, and three dimensional spin vectors $(\vec{\chi}_1, \vec{\chi}_2)$. 2) Extrinsic parameters (in the observer frame of reference) - standard spherical coordinates (D, i, ψ) , sky-location (θ, ϕ) , and

lastly (t_c, ϕ_c) the time and the phase at the coalescence. The anticipated signal is then accurately modeled in terms of these parameters with the help of various analytical and numerical techniques [30–32].

To search for the modeled signal $\tilde{h}(f)$ also referred as template, matched filtering is performed in the Fourier domain to quantify the likelihood of data containing the particular template. The matched filter statistic is a correlation between the Fourier transform of the data ($\tilde{s}(f)$) and the template ($\tilde{h}(f)$) weighted by the noise power spectral density (PSD) $S_n(f)$ [24]. It can be shown that matched filter is an optimal detection statistic for distinguishing signals in the presence of stationary Gaussian noise [24]. The mathematical form of the complex matched filter statistic is

$$\langle s|h \rangle = 4 \int_0^\infty \frac{\tilde{s}(f)\tilde{h}^*(f)}{S_n(f)} df. \quad (1)$$

The output of the matched filter after normalizing with the correlation of the template with itself $\langle h|h \rangle^{1/2}$ is the signal-to-noise ratio SNR

$$\rho^2 = \frac{(\text{Re}[\langle s|h \rangle])^2}{\langle h|h \rangle}. \quad (2)$$

Apriori the parameters of gravitational-wave signals are unknown and to search for the intrinsic parameters, a discrete template bank is used to *cover* the intrinsic parameter space. The notion of cover is to sample enough points in the parameter space such that the match between data and at least one template is above a minimum match value. In current searches typically a minimum match of 0.97 is used [29, 52], and lattice-based [52–54], stochastic methods [55, 56] or hybrid methods [33, 34] are applied to sample the points in the parameter space. Since we are considering aligned-spins with the orbital angular momentum in the $+z$ direction, the intrinsic parameter space consists only of $\zeta = (m_1, m_2, \chi_{1z}, \chi_{2z})$ parameters. Search over the two categories of binary parameters are handled differently - the intrinsic parameters are searched by repeatedly matched filtering for every template. Whereas, the extrinsic parameters - sky location, the orientation of the binary, and distance to the source are accounted as an overall phase ϕ_0 and an amplitude A [29]

$$\tilde{h} = A e^{i\phi_0} \tilde{h}_0(\zeta) e^{2\pi i f t_c}. \quad (3)$$

In the Eq. (3), A and ϕ_0 are unknown functions of $(D, i, \theta, \phi, \psi)$, and h_0 depends on the intrinsic parameters. The unknown amplitude A is a nuisance parameter that is eliminated by normalizing the SNR with the norm of the template as seen in Eq. (2). The unknown phase ϕ_0 is maximised using a quadrature, which is equivalent to maximizing the norm of the complex SNR [29]. Finally, the position of the signal

is determined by searching for the time of coalescence of binary, represented by the t_c parameter. Variation in t_c is expressed as time-translations, and is separated using $e^{2\pi i f t_0}$. Substituting Eq. (3) in Eq. (1), the matched filter output (SNR) at $t = t_0$ is given by Eq. (4). The SNR as a function of time can be obtained efficiently by performing an inverse FFT (IFFT) of the Eq. (4) [29, 53]

$$\langle s|h_0 \rangle(t_0) = 4 \int_0^\infty \frac{\tilde{s}(f)\tilde{h}_0^*(f; \zeta)}{S_n(f)} e^{2\pi i f t_0} df. \quad (4)$$

Data segment with SNR above a predetermined threshold is referred as a *trigger* which may contain a true GW signal. The ambiguity is due to the assumption of stationary Gaussian noise for the matched filter statistic. However, triggers due to non-stationary glitches can give rise to false alarms which lower our confidence of identifying true GW signals [25, 26]. Additional signal consistency test introduced in [27, 28, 57] is performed to down-rank triggers due to glitches. Furthermore, it is ensured that only coincident triggers from multiple detectors are considered – triggers corresponding to the same template and observed within the light travel time window between the detectors [58]. Amongst the various steps mentioned, matched filtering requires the dominant computational costs of a search. Hence, our focus is to optimize the matched filtering process.

To summarize the matched filtering procedure, the intrinsic and extrinsic parameters are searched separately using a template bank, analytical techniques respectively. First, the PSD weighted correlation of data with a single template in ζ is computed, and then an inverse FFT is executed to obtain a complex time-series. Taking the modulus of the complex times-series and normalizing it by the norm of the template gives the SNR times-series. The above steps are repeated for all the templates in the template bank to search over the intrinsic parameters. Throughout this paper, we refer to the method of matched filtering data with the templates as the *template method* for simplicity.

It is clear from above that matched filtering operation scales linearly with the number of templates. In the case when the size of the template bank is large, a search can be limited by the computational costs required for filtering the data with templates. It is possible to numerically reduce the size of the template banks to a fewer number of basis vectors and filter data directly with the bases. We now give a brief introduction to performing matched filtering with a reduced bases.

A. Matched Filtering using a Reduced Bases

Consider a region in the parameter space described by $\zeta = (m_1, m_2, \chi_{1z}, \chi_{2z})$. Discrete templates are used to cover this region, and as a result of the mismatch criterion, templates are strongly correlated in the vicinity of each other. The correlation between

the templates incurs a redundancy in matched filtering computations, instead, an orthonormal basis can be used to eliminate these correlated computations [48]. Commonly used methods for computing an orthonormal bases is the principal component analysis (PCA) and the singular value decomposition (SVD). The PCA approach to obtain bases is by performing an eigenvalue decomposition (EVD) of the covariance matrix $\mathbf{C} = \mathbf{T}^\top \mathbf{T}$ constructed using the templates (see Eq. (5a)). Whereas, SVD is applied directly to a matrix containing the templates \mathbf{T} (see Eq.(5b)).

$$\mathbf{C} = \mathbf{P}\mathbf{L}\mathbf{P}^\top, \quad (5a)$$

$$\mathbf{T} = \mathbf{U}\mathbf{S}\mathbf{P}^\top. \quad (5b)$$

In the case when templates are centered – the column means of \mathbf{T} are zero, then both SVD and PCA yield the same orthonormal bases. The bases are represented as columns of \mathbf{P} which are ranked by their corresponding eigenvalues in \mathbf{L} or singular values in \mathbf{S} . It can be easily shown that $\mathbf{L} = \mathbf{S}^2$ and that the two methods are similar. Therefore, either of them can be used to obtain the bases.

Consider a set of p_t basis vectors for the parameter region ζ denoted by $\tilde{p}(f)$ in the Fourier domain. Every template in this region h_ζ is expressed in terms of a unique linear combination of the bases. Since the matched filtering operation is also linear, we can filter the data \tilde{s} only using the bases and rewrite Eq. (4) in terms of $\tilde{p}(f)$

$$\rho_r(t) = \sum_{k=0}^{p_t-1} 4 \int_0^\infty \frac{\tilde{s}(f)c_{k,\zeta}^* \tilde{p}_k^*(f)}{S_n(f)} e^{2\pi i f t} df, \quad (6)$$

where the template $\tilde{h}_\zeta(f) = \sum_{k=0}^{p_t-1} c_{k,\zeta} \tilde{p}_k$, is a linear combination of the bases $\tilde{p}(f)$ and the unique decomposition coefficients $c_{k,\zeta}$. The coefficient $c_{k,\zeta}$ is obtained by computing the scalar product of \tilde{h}_ζ and the k^{th} basis vector \tilde{p}_k .

Matched filtering with the basis is first done by performing an IFFT of the correlation between \tilde{s} and \tilde{p}_k , which results in a complex time-series defined as $\beta_k(t) = \text{IFFT}(\langle \tilde{s} | p_k \rangle)$. Afterwards, each β_k is weighed accordingly using the respective decomposition coefficients and then combined to give SNR time-series corresponding to the template $\tilde{h}_\zeta(f)$. The process of multiplying the coefficients $c_{k,\zeta}$ with β_k is the reconstruction step, as it reconstructs the SNR time-series using the contribution from the bases

$$\rho_r(t) = \sum_{k=0}^{p_t-1} c_{k,\zeta}^* \beta_k. \quad (7)$$

Using the complete orthonormal bases $p_t = T$ where T is the number of templates, Eq. (7) reproduces exact results as Eq. (4).

Instead of using the full bases, it is possible to approximately reconstruct the SNR with fewer bases. Eigenvalues σ are arranged in decreasing order and only the first p basis vectors are chosen and the rest are discarded. It can be shown that the first p basis vectors span an approximate lower rank subspace of the original parameter space. Neglecting contribution from some basis vectors leads to an average loss in SNR, which is shown to be a function proportional to the eigenvalues [48]

$$\left\langle \frac{\delta \rho}{\rho} \right\rangle = 1 - \frac{\left| \sum_{k=0}^{p-1} \sigma_k^2 \right|}{\left| \sum_{k=0}^{p_t-1} \sigma_k^2 \right|}. \quad (8)$$

The equation above indicates that the number of relevant bases p can be fine-tuned based on the choice of tolerance in loss of SNR. For detection purposes, we want to keep this loss under the mismatch value (0.03) due to the discreteness of the template bank.

Reconstruction of SNR time-series is performed for every template and thus, the reduced bases approach requires additional costs to matched filtering. To estimate the reconstruction costs in brief (exact costs are estimated later in this paper), consider a data segment having N samples, T number of templates, and reduced p basis vectors. The number of operations required for the reconstruction step is $O(NpT)$. Meanwhile, comparing the costs for template based filtering; which requires $O(NT \log N)$, and thus, the actual comparison boils down to $\log N$ and p . The exact values for $\log N$ and p vary over the parameter region, but considering a ballpark it is observed that p is typically 1-2 order in magnitude bigger than $\log N$. Hence, the reduced bases approach loses all the computational advantage of filtering against fewer bases.

B. Comparison with Current Methods

Different methods in the past have been implemented to speed up the process of filtering either by reducing the latency of the search [22, 59, 60] or by decreasing the required computational costs [61–63]. For the former case, a reduced bases filtering technique along with multi-rate sampling is used with an intent to decrease the latency of matched filtering. And in the latter case, the aim is to reduce the filtering costs by using a multi-stage hierarchical filtering method. Some of these techniques have been already implemented in the current search pipelines [23, 49, 51], and are deployed in different search scenarios. We briefly discuss these various strategies and contrast our methodology next.

1. Reduced bases filtering with or without multi-rate sampling

To discretize a continuous signal, the Nyquist-Shannon criterion [64] determines the sampling rate to be at least twice the highest resolvable frequency ($1/dt \geq 2f_{max}$) of the anticipated signal. This gives the relation for the number of samples in the data $N = (df/dt)^{-1}$ where $1/dt, 1/df$ are sampling rate and sampling frequency respectively, suggesting the filtering costs increase when searching for higher frequencies. Because the frequency evolution of these signals is *chirp* like; rapidly increasing towards the merger, this allows a low sampling rate at the earlier times and can be increased subsequently as the signal evolves. Using multiple sampling rates the matched filtering costs are reduced significantly [22, 59, 60].

Multi-rate sampling has been adopted by the MBTA [22] LLOID [60] and SPIIR [59] schemes that are implemented in current online pipelines for the prompt detection of signals [22, 23, 49]. The MBTA method performs matched filtering in the Fourier domain using the standard FFT approach to obtain SNR time-series. Whereas the LLOID and SPIIR methods perform time domain filtering by employing FIR or IIR filters respectively to compute an equivalent form of the matched filter Eq. (4). These filters are specially designed for whitening the data, a process which causes most latency in matched filtering [22, 23, 49]. The overall latency is further improved by using a reduced bases obtained by performing SVD of the IIR/FIR filters. Results from the online pipelines are very well in agreement with the rigorous offline searches [21, 50], justifying the viability of multi-rate sampling and reduced bases filtering in CBC searches.

The number of templates T drastically increases when searching in sub-solar regions [7, 38–40] or with additional parameters [39, 47] e.g. eccentricity. In such search scenarios obtaining a reduced bases for the complete template bank is computationally limited. This is because SVD is performed on a template matrix whose size scales linearly with T and might require infeasible amounts of memory for the template matrix. To address this issue we choose the PCA approach of computing the bases, which is performed on a covariance matrix whose size is independent of T , and therefore, making it feasible to obtain orthonormal bases even for large template banks.

The major drawback of a reduced basis filtering method is the large reconstruction costs, because of which this scheme is not implemented in extensive offline searches where the computational costs play an important role. Another approach to reducing the costs is to split the filtering step into multiple stages in a hierarchical fashion. Next, we discuss established hierarchical methods and compare them with our new hierarchical scheme.

2. Hierarchical methods

The crux of any hierarchical search is to perform a coarse and a fine search over the parameters involved in the hierarchy. In the past, the work in [62] proposed a two-stage hierarchical filtering on just a single chirp mass parameter, and the same work was extended in [63] for three parameters - the component masses and the time of coalescence. The most recent work [61] in hierarchical approach to matched filtering extended the scheme to multiple detectors analysis. For the first stage data is down-sampled at 512 Hz, and filtered using a coarse bank of $MM = 0.9$. In the second stage, data is sampled at the full rate of 2048 Hz, and filtered with a fine bank having $MM = 0.97$. Their method achieved $\sim 20x$ speed up in comparison to the one-step search on simulated data containing only Gaussian noise.

To assign significance of any event, it is important to estimate the noise *background* which is the trigger distribution due to noise only events [58]. In the work [61], the performance gain comes at the expense of a poor estimation of the background. Since they do not follow the noise triggers until the second stage, the true background is mimicked by scaling the first stage background. This leads to an improper estimation of the significance for an event.

In this method, we present a new hierarchical method aimed at reducing the reconstruction costs. We perform a two-stage hierarchical reconstruction of the SNR time-series for the complete bank. We follow up all the triggers for $SNR \geq 5$ till the second stage, and therefore, are able to accurately determine the original noise background. Furthermore, our method incurs no loss in the search sensitivity. We discuss the methodology in detail in the next section.

III. REDUCED BASES HIERARCHICAL MATCHED FILTER

In this section, we describe our new hierarchical approach to reduced bases matched filtering in detail. We first briefly review the PCA method in general and discuss a technique to reduce the costs of performing PCA. Then we explain how PCA is applied on a template bank and the hierarchical method of filtering data along with its implementation on GPUs. Finally, we estimate the matched filtering costs in detail for the template based method and the reduced bases hierarchical scheme to compare the relative gain in performance.

A. PCA using Non-uniform Sampling

We first briefly explain the PCA procedure applied on n vectors denoted by \mathbf{v} . Every vector is centered by subtracting the mean vector $\mathbf{v}_s = \mathbf{v} - \mathbf{b}$, where $\mathbf{b} = 1/n \sum_{i=0}^n v_i$ is the mean vector. To ensure each

vectors get equal weight, they are normalized w.r.t. the inner product defined on the vector space. Using the normalized and centered vectors $\hat{\mathbf{v}}_s$, a covariance matrix is created $\mathbf{C} = \hat{\mathbf{v}}_s^\top \hat{\mathbf{v}}_s$. An EVD is performed to get the orthonormal bases vectors \mathbf{p}_t of the \mathbf{C} matrix, which are ranked by the corresponding eigenvalues σ . The basis vectors corresponding to small eigenvalues are discarded, and the resulting set of reduced bases is denoted by \mathbf{p} . Projecting $\hat{\mathbf{v}}_s$ onto \mathbf{p} gives the decomposition coefficients $\mathbf{D} = \mathbf{p} \hat{\mathbf{v}}_s$. The reduced bases $\hat{\mathbf{v}}_s$ and the decomposition coefficients \mathbf{D} are used to retrieve an approximate version of $\hat{\mathbf{v}}_s$, given by $\hat{\mathbf{v}}_s^{approx} = \mathbf{D}^\top \mathbf{p}$.

Even though the costs of PCA are amortized, PCA can be time-intensive and difficult to perform on a large collection of longer duration templates e.g. low-mass or lower-frequency. Amongst the various steps involved in PCA, the dominant computational and memory costs are for the covariance matrix - both scale quadratically with the required number of samples N_t by the templates. To put the scaling relations into perspective, consider the complete O2 bank sampled at a constant sampling frequency of 128s, the rough estimate for the size of \mathbf{C} is ~ 550 GBs. The EVD of \mathbf{C} is difficult to distribute across several machines, and thus, the size of the covariance matrix is constrained due to the memory of a single machine. Therefore, it is crucial to reduce the size of the templates to make the PCA faster and feasible in the low-mass regime. We use a much smaller set of frequencies to efficiently sample the templates and the bases. Once the bases are obtained, we interpolate them back to the original sampling frequency.

In this work, we consider the frequency range from [15, 1024] Hz. The principle idea behind efficient sampling is to adjust the sampling rate according to the number of oscillations of a function within a given frequency bin. To account for the complete frequency range in our sampling analysis, we consider the template with the longest bandwidth and identify all the frequencies corresponding to the zero crossings of this template. The identified frequencies are used to define edges of the non-overlapping bins of different sizes. We further sample every bin by using five uniformly spaced frequencies within, and together they make up the complete set of non-uniform sampling frequencies. We also ensure that our sampling criteria is never less than $df = 1/128$, which helps us avoid oversampling the dense bins at very low frequencies. The number of frequencies per bin is chosen empirically based on the relative error induced in the templates for not using the full sampling rate.

To test the accuracy of our sampling method, we check the overlap between the templates generated using non-uniform and uniformly sampled frequencies. For this purpose, templates evaluated at non-uniform frequencies are linearly interpolated to a constant sampling frequency of 128s. We obtain a mismatch of $< 10^{-4}$ in the overlap due to interpolation of the templates, which is much smaller than the error due to discreteness of the template bank and, therefore, can

be safely neglected. This justifies that our non-uniform frequencies are viable for sampling the templates. Using our sampling method, we obtain $N_t = 11074$ which significantly reduces the memory required for \mathbf{C} from ~ 550 GBs to only ~ 0.9 GBs. Therefore, saving a lot of computational resources and simultaneously speeding up the PCA process.

B. Implementing PCA on a Template Bank

In this work, we use the template bank as described in [65] which was also used for the PyCBC analysis of the O2 observing run [5]. The parameters ranges used in this bank are - total mass $M \in [2, 100]$ and mass ratio $q \in [1, 98]$. We restrict ourselves to the aligned-spin case where the spins for NS are up to 0.05 and up to 0.998 for BHs. The minimum match criterion used in this bank is 0.97, and the bank contains $\sim 400,000$ templates.

We divide the parameter space into smaller regions to reduce the number of local bases p as they contribute linearly to the dominant reconstruction costs. The complete bank is split into smaller sub-banks, and then PCA is performed on each of them individually. Splitting of template bank is performed in the (τ_0, τ_3) coordinates along iso- τ_0 lines because the metric is roughly Euclidean in these coordinates [53]. The parameter τ_0 roughly corresponds to the duration of a template in seconds, that scales as $\tau_0 \propto \mathcal{M}^{-5/3} f_0^{-8/3}$, where \mathcal{M} is the chirp mass and f_0 is the lower frequency used in the analysis. We choose to split the τ_0 range into 64 equal parts, each of them containing 6250 templates. While optimizing the splitting is not in the scope of this work, we observe no significant improvement by briefly testing the splitting with a few smaller or bigger equal parts than 64. In Fig. 1 we show the complete parameter space along with an example sub-region which is used as a case study for further analysis.




FIG. 1. O2 template bank [65] used in this work, and the small orange region corresponds to the sub-bank used as a case study containing 6250 templates.

The PCA operation begins by sampling templates at the previously obtained non-uniform frequencies using the IMRPhenomPv2 waveform model [30]. Templates are then whitened using the aLIGO PSD and normalized to unity. The template matrix for the m^{th} sub-bank \mathbf{T}^m is constructed by storing templates row-wise such that \mathbf{T}^m has the dimensions of $(T/64 \times N)$. In the case when T is not a multiple of 64, we can simply choose another divisor close to 64. We observed the mean vector of \mathbf{T} to be almost zero, and hence, skip the mean subtraction step. Covariance matrices for each sub-bank \mathbf{C}^m are evaluated by multiplying the template matrix with its transpose $\mathbf{T}^m \times (\mathbf{T}^m)^\top$. In the next step, we perform the EVD of \mathbf{C}^m to obtain the basis vectors and their corresponding eigenvalues. For this purpose we employ the Lanczos algorithm [66] - an efficient algorithm to obtain the p largest eigenvalues. Invoking Eq. (8) for a tolerance of 10^{-5} , the number of relevant eigenvalues p obtained for a few different sub-banks are shown in the Table I.

sub-bank index	parameter τ_0	p
1	[0.1, 5.1]	64
.	.	.
34(case study)	[98.0, 103.4]	254
.	.	.
64	[442.5, 595.7]	200

TABLE I. Example of a few sub-banks corresponding to their respective τ_0 ranges in the second column. The third column is the number of relevant eigenvalues p for the respective sub-region.

We then compute the decomposition coefficients

essential to reconstruct the whitened templates in \mathbf{T}^m . These coefficients are unique for each template and are obtained by multiplying the two matrices - bases matrix \mathbf{P}^m and the template matrix \mathbf{T}^m . The resulting matrix is the decomposition matrix \mathbf{D}^m containing the p unique coefficients for every template in the m^{th} sub-bank and has the dimensions $(p \times T/64)$. Finally, the original whitened templates can be approximately reconstructed by multiplying \mathbf{D} and \mathbf{P} . In the next section, we discuss the two-stage hierarchical reconstruction of the matched filter output using \mathbf{D} and \mathbf{P} .

C. Hierarchical Reconstruction of the SNR Time-series

Matched filtering with the basis vectors as per Eq. (7) is performed using the following steps:

- compute FFT of the data $s(t)$ with N sample points at a uniform sampling rate $1/dt$ to obtain $\tilde{s}(f)$.
- linearly interpolate the basis vectors $\tilde{p}_k(f)$ at the uniform frequencies (multiples of f_{max}/N).
- filter data with every basis vector to obtain β_k - compute inverse FFT of the product $\tilde{s}(f)\tilde{p}_k^*(f)/\sqrt{S_n(f)}$.
- average β_k in bins of w samples to obtain β_k^{avg} .
- perform first stage reconstruction to obtain averaged SNR time-series.
- perform second stage reconstruction around the triggers from the first stage.

The first step in the reduced basis matched filtering process is to compute forward FFT of the data $s(t)$. Since $s(t)$ is sampled uniformly with a rate of $1/dt = Ndf$, the Fourier transformed data $\tilde{s}(f)$ is obtained at frequencies given by integer multiples of f_{max}/N , where $f_{max} = 2/dt$. Now to filter the data with the bases, the correlation of data and bases are computed at the uniform frequencies, and for this reason, the bases are linearly interpolated from f_{min} to f_{max} in df steps. Since the bases are already whitened and the denominator in the matched filtering Eq. (6) requires $S_n(f)$, we multiply the correlation product with $S_n(f)^{-1/2}$ to get the appropriate denominator. For a basis vector \tilde{p}_k , the filtered output β_k time-series is obtained by computing the inverse FFT of the weighted correlation. Finally, for every data p different time-series (bases output) are stored in a separate β matrix of size $T \times p$.

The reconstruction of SNR time-series for every sample and each template requires large computational costs. Since we are only interested in triggers exceeding a certain threshold, it is better not to reconstruct the complete SNR times-series, rather only in the vicinity of the triggers. We propose a two-step hierarchical scheme for reconstruction, which performs a coarse

reconstruction, and then a finer reconstruction around the triggers obtained in the first stage. In the first stage, we consider non-overlapping bins of w samples. Then the outputs from each basis vector β_k are averaged in the bins to give the averaged time-series

$$\beta_k^{avg}[t_i] = \sum_{j=0}^{w-1} \beta_k[t_{i \times w + j}] / w, \quad (9)$$

where $i = \{0, \dots, N/w - 1\}$ corresponds to different bins, and $j = \{0, \dots, w - 1\}$ goes over the samples inside each bin. The bin index i can be considered as a sample point in the shortened averaged time-series. By linearly combining the β_k^{avg} with the decomposition coefficients $c_{k,\zeta}$ similar to Eq. (7), results into an averaged SNR times-series for the template h_ζ . We identify the first stage triggers and their corresponding bins having the average SNR above a first stage threshold ρ_I . In the next step we perform a finer reconstruction for every sample inside all the triggering bins. Triggers from the second stage which are above a second stage threshold ρ_{II} are referred as the final triggers.

Our aim is to minimize the total costs for the hierarchical method without losing any sensitivity. The sensitivity is determined by computing the fraction of triggers recovered at a given SNR, or conversely, the SNR threshold at which all the triggers are recovered. We compute the sensitivity using the latter approach – by comparing the hierarchical distribution of the triggers with the original distribution. The parameters w, ρ_I are fine-tuned under the constraint of reaching a fixed target SNR (ρ_{target}) to optimize the total costs. We give a detailed description to compute ρ_{target} in the later subsection III F.

D. Fast First Stage Filtering using Templates

To reduce the first stage costs the average SNR can be obtained by the faster template method which is mathematically equivalent to reduced bases filtering. Using the template method, the average SNR time-series can be obtained by binned averaging of the integrand in the Eq. (4) and then performing an IFFT of shorter length. To demonstrate the averaging process mathematically, we consider a single bin b with w samples, and compute the averaged SNR for the same, which is denoted by $\langle \rho_\theta(t) \rangle_b$. The w time samples in the bin are represented by $t = wb + r$, where $r \in [0, w-1]$. We then write the discretized version of the Eq. (4) averaging the SNR time-series over w samples

$$\langle \rho_\theta(t) \rangle_b = \frac{1}{Nw} \sum_{r=0}^{w-1} 4\Delta f \sum_{f=0}^{N-1} \frac{\tilde{s}[f] \tilde{h}^*[f]}{S_n[f]} e^{2\pi i f (wb+r)/N}. \quad (10)$$




FIG. 2. Demonstration of the hierarchical scheme, the first panel shows the reconstructed SNR time-series without any averaging and a single trigger ($\text{SNR} \geq 4$). In the middle panel is the averaged SNR series obtained from binned averaging of $w = 16$ samples (shown in green in the first panel), and the red region corresponds to the bin containing the original trigger. The last panel shows a finer reconstruction performed only for the bin having average SNR greater than the first stage threshold.

Now, breaking the summation over $f \in [0, N - 1]$ into a double sum by expressing $f = lN/w + f'$, we can rewrite the single summation of any function $\tilde{g}(f)$ in the Fourier domain as

$$\sum_{f=0}^{N-1} \tilde{g}[f] = \sum_{f'=0}^{N/w-1} \sum_{l=0}^{w-1} \tilde{g}[l\frac{N}{w} + f'], \quad (11)$$

where $l \in [0, w - 1]$ and $f' \in [0, N/w - 1]$. This modifies the Eq. (10) to

$$\begin{aligned} \langle \rho_\theta(t) \rangle_b &= \frac{4w\Delta f}{N} \sum_{f'=0}^{N/w-1} e^{2\pi i f' \frac{wb}{N}} \\ &\times \underbrace{\frac{1}{w^2} \sum_{l=0}^{w-1} \frac{\tilde{s}[l\frac{N}{w} + f'] \tilde{h}^*[l\frac{N}{w} + f']}{S_n[l\frac{N}{w} + f']} \sum_{r=0}^{w-1} e^{2\pi i (l\frac{N}{w} + f') \frac{r}{N}}}_{=\Omega(f')} \\ &= \frac{4w\Delta f}{N} \sum_{f'=0}^{N/w-1} e^{2\pi i f' \frac{wb}{N}} \Omega(f'). \end{aligned} \quad (12)$$

To simplify Eq. (12), we introduce a new frequency series $\Omega(f')$ which represents the binned average of the PSD weighted correlation of the data and template. The second summation term in $\Omega(f')$ has a closed form

solution and can be computed analytically. We notice that the last line of Eq. (12) has a similar form to Eq. (4) but, the integrand replaced with $\Omega(f')$ having only N/w number of samples. Suggesting that the average SNR time-series can be obtained by an IFFT of $\Omega(f')$, where one can think of b as the new time equivalent variable which ranges between $[0, N/w - 1]$. Therefore, it is possible to perform the first stage using bases or templates, however, the second stage filtering needs to be done with the bases.

E. Implementation

We now discuss the implementation of our method which is divided into two parts - 1) preparation and 2) matched filtering. The preparation stage is implemented partially on GPUs, whereas matched filtering is performed entirely on the GPUs. We use several Nvidia GV100 GPUs, each having memory of 32 GB as well as several RTX 2070 Super each with 8 GB of memory. Our code is written in C language and uses various optimized libraries for different purposes. Operations on the GPU are performed using CUDA [67] an application programming interface by Nvidia.

In the preparation stage, we perform PCA on the template bank to obtain the reduced bases and the respective decomposition coefficients. Matrix multiplications on the GPUs in this stage are performed using the *cuBLAS* library from CUDA. We begin by computing the covariance matrix in several parts in parallel using cuBLAS. Afterward we combine all the parts to obtain the final matrix \mathbf{C} . The Lanczos algorithm for EVD of \mathbf{C} is implemented on CPUs using the SLEPc [68] and PETSc [69] libraries. We obtain the decomposition coefficients matrix \mathbf{D} by multiplying the matrices \mathbf{T} and \mathbf{P} . The preparation stage is computed in advance and performed only once. Results from this stage - the bases and the decomposition coefficients, are stored on hard drives for the later matched filtering stage. To reduce the input/output (IO) bandwidth, we compress the PCA results before writing them on the hard drive.

In the next stage, we read the output from the previous stage to match filter data using our hierarchical scheme. To reduce the time-intensive memory transfers between CPU and GPU, we load the matrices \mathbf{P} and \mathbf{D} at once onto the GPUs. We divide the data into several smaller segments such that we can optimally utilize the memory of the GPUs while filtering each segment in parallel. We use data segments of 128s and a sampling rate of 2048 Hz. Each data segment has $N = 128 \times 2048$ samples which overlap with $N/2$ samples from the previous segment. We employ the cuFFT library from CUDA to perform the FFTs in this stage. Using cuFFT we perform FFT for a batch of data segments in parallel. In the next step, we interpolate the basis vectors and multiply them with \tilde{s} along with $S_n(f)^{-1/2}$. Afterward we perform in-place

batched IFFTs to obtain the filtering output from the bases. The in-place technique saves GPU memory by recycling the allocated input memory to write the output.

To perform the hierarchical reconstruction we first average the output from bases to obtain the β^{avg} matrix. The first stage reconstruction is done using cuBLAS by multiplying β^{avg} and \mathbf{D} , which outputs the average SNR time-series. Next, we use a dedicated function on the GPU to find triggers with average SNR above ρ_T . Once the triggers are identified, we store their bin indices along with their corresponding average SNRs. Since these triggers are not contiguous in memory, further reconstruction of the triggering bins cannot be performed by simple matrix multiplication. Therefore, instead of using the optimized cuBLAS library, we use a custom-built function on the GPU to perform the second reconstruction.

F. Cost Estimation

In this subsection, we estimate the floating-point operations required by the two different matched filtering schemes. The purpose of estimating the number of operations is to get a rough idea of the scaling relations involved for the total filtering costs. Moreover, it will allow us to estimate the improvement in performance due to the proposed hierarchical scheme. In both methods, we split the data into blocks of N samples, having an overlap of $N/2$ samples with the previous block. This is generally done to avoid corrupt SNR samples at the start and end of a data segment [29]. Hence, filtering a single block results in $N/2$ unique SNR time-samples. Most operations involve complex numbers unless otherwise specified. Throughout the cost estimation, we consider 6 operations for multiplication and 2 operations for the addition of two complex numbers. To estimate the costs for the FFTs we consider a split-radix method.

We first estimate the costs for the template method which will be our baseline comparison. The first step is to compute the forward real-to-half-complex FFT of the data with N samples, which requires $3/2N \log N$ operations per block [70]. Computing the integrand of the matched filtering Eq.(1) for T templates requires $6NT$ operations. Finally, an inverse complex-to-complex FFT is required to obtain the SNR time-series for each template, and this attributes to $5TN \log N$ operations, where each IFFT requires $5N \log N$ operations. In total the template method for a single block requires $N \log N (3/2 + 5T) + 6TN$ operations. Usually, the number of templates is huge ($T \gg 1$), so we can neglect the cost for the forward FFT of data. Therefore, the total floating-point operations z_{basic} for filtering $N/2$ data samples with the template method can be approximated to

$$z_{basic} = NT(5 \log N + 6) \quad (13)$$

Now, we estimate the costs for the two stage

hierarchical filtering. As shown in the subsections III C and III D, the first stage can be performed either by using the bases or the templates. We evaluate the first stage costs by considering the faster template method described in the section III D. Starting with the forward FFT of the data which needs $3/2N\log N$ operations. Afterward the weighted correlation of the matched filter in Eq. (12) is then obtained for every template in $6NT$ multiplicative operations. Then, we perform binned averaging of the correlations to get reduced frequency series $\Omega(f')$ of size N/w for each template, and this requires $2NT/w$ operations. For every template, we obtain the average time series by computing the IFFT of $\Omega(f')$ in $5NT/w\log(N/w)$ operations. Hence, the number of floating-point operations required for the first stage is

$$z_{first} = NT\left(\frac{5}{w}\log\left(\frac{N}{w}\right) + 6 + \frac{2}{w}\right) \quad (14)$$

In the next stage, we compute a finer reconstruction of w points around each first stage trigger. The costs for the second stage are calculated in terms of the number of first stage triggers. We denote the number of first stage triggers for a single template by $f(w, \rho_I)$. We assume that the number of triggers does not vary over templates, and thus, can be obtained from a single test template. This assumption is justified for the templates in the vicinity of the test template as the number of triggers would be roughly the same. In addition, since $f(w, \rho_I)$ decreases rapidly at higher SNRs, the error in the total costs due to our assumption is negligible. Using the above assumption, the second stage requires $z_{second} = pwf(w, \rho_I)T$ operations. Summing up the costs from both the stages, the total floating-point operations required for the hierarchical method are

$$z_{total} = NT\left(\frac{5}{w}\log\left(\frac{N}{w}\right) + 6 + \frac{2}{w}\right) + 4pwf(w, \rho_I)T \quad (15)$$

From Eq. (15) we see that the final costs are obtained in terms of two nuisance parameters w and ρ_I . We must choose an appropriate first stage threshold ρ_I because a potential trigger not identified in the first stage is missed completely. Only triggers that are followed till the second stage are considered while estimating the sensitivity of the filtering method. To ensure that we recover all the triggers using a reliable first stage threshold in the two stage filtering scheme, we compare the hierarchical distribution of second stage triggers against the trigger distribution from the flat scheme. The idea is to identify a target SNR (ρ_{target}) as a function of (w, ρ_I) , such that the hierarchical scheme recovers 99% of the total triggers from the flat scheme. Considering a certain first stage configuration given by specific values of (w, ρ_I) , we denote the number of final triggers above ρ_{II} as $n_{final}(\rho_{II})$. Similarly, using the same threshold ρ_{II} , we denote triggers from the flat scheme as $n_{flat}(\rho_{II})$. The

target SNR ρ_{target} is then defined as

$$\rho_{target}(w, \rho_I) = \left(\min(\rho_{II}) \mid n_{final}(\rho_{II}) \geq 0.99n_{flat}(\rho_{II})\right). \quad (16)$$

IV. ACCURACY AND PERFORMANCE OF ANALYSIS WITH SIMULATED DATA

In the previous section, we demonstrated our method and estimated the required costs for filtering. In this section, we present the accuracy of the hierarchical filtering results and assess the reduction in the required number of operations w.r.t. the baseline – template method. Furthermore, we measure the gain in performance by implementing matched filtering on GPUs relative to the established CPU implementations.

We use the sub-region (shown in Fig. 1 in orange) to demonstrate our method. This sub-region covers the parameter ranges $M \in [5.72, 12.05]$ and $q \in [1.0, 11.05]$. Using a tolerance of 10^{-5} as per Eq. (8), we obtain $p = 254$ for this sub-region. We want to estimate a conservative reduction in the total costs that scale linearly with p . Following this reason, we choose the mentioned sub-region as it corresponds closely to the average $\langle p \rangle$.

Current offline and online searches [50, 71] generally use SNR threshold of $\sim 4 - 5$ for the single detector triggers. However, searches involving a large number of templates are affected by increased background due to noise triggers [39, 47], and thus, higher SNR thresholds are used to detect events at a constant false alarm rate. These kinds of search scenarios also happen to be the case where cost-efficient algorithms are necessary. Therefore, in this work, we target SNR thresholds of 5 and above.

A. Accuracy of the SNR

We expect two primary contributions to the SNR loss in our method - *truncation of the number of eigenvalues* and *interpolation of the bases*. Error due to truncating the eigenvalues is translated as the SNR loss via the Eq. (8), and is regulated by choosing an appropriate tolerance. The loss due to linear interpolation of the bases is quantified in terms of the mismatch between the interpolated and fully sampled templates. We evaluate the total loss by computing the relative error in the SNR time-series obtained using our method and the template method. We filter simulated colored Gaussian noise from the PSD to acquire the SNR time-series. Comparing triggers from every template, we note the maximum relative error in the SNR values and plot it against the SNR thresholds as shown in Fig. 3.

1. Number of required operations




FIG. 3. Maximum relative error between the original and reconstructed SNR values as a function of SNR threshold. The incurred error scales inversely with the SNR thresholds and is smaller than the error due to the discreteness of the templates for relevant SNR thresholds.

Based on our sensitivity requirements we are interested only in $\text{SNR} = 5$ and above, and from Fig. 3 we observe the relative error to be $\sim 0.4\%$ for $\text{SNR} = 5.0$. This amount of error can be tolerated because the observed loss is less than the error due to mismatch of the templates, which is up to 3%. We also notice that the relative error decreases even further with higher SNR thresholds. Therefore, it is justified that our method successfully recovers the SNR values for search scenarios requiring SNR thresholds ≥ 5 .

B. Comparing Performance with Template Based Matched Filtering

Our hierarchical method is characterised by two parameters, the averaging bin size w , and the first stage threshold ρ_I . For testing the hierarchical scheme, we use four different averaging bin size $w \times dt$ with $w \in [2, 4, 8, 16]$, and various different values of ρ_I . We drop the dt factor when referring w , to suggest the reader that we simply average over w samples. We test our scheme using noise generated from the aLIGO PSD sampled at 2048 Hz. We also check our method for a population of BBH signals within the case study sub-region. Finally, we compare the estimated costs required by the hierarchical method against the template method. We use PyCBC software library [72] to generate simulated data containing Gaussian noise and to perform injections.

Using the estimates in section III F, we compare the number of operations required by the hierarchical scheme with the template method. We begin by testing only Gaussian noise generated using the aLIGO PSD. For this purpose, we generate a total of ~ 7.4 days of data using different seeds sampled at 2048 Hz. For filtering purposes, the simulated data is then divided in smaller segments of 128s with an overlap of 64s from the previous segment.

To estimate the total costs for the hierarchical method Eq. (15), we first determine the number of first stage triggers $f(w, \rho_I)$ by varying the hierarchical parameters. We neglect the variation of $f(w, \rho_I)$ for different templates for reasons discussed previously in section III F. $f(w, \rho_I)$ is obtained simply by iterating over different values of the first stage threshold ρ_I for different (fixed) w and by counting the total first stage triggers above the same threshold. The observed $f(w, \rho_I)$ w.r.t. ρ_I for different values of w is shown in Fig. 4.




FIG. 4. The number of first stage triggers $f(w, \rho_I)$ per second as a function of ρ_I for different averaging bin sizes w (shown in the legend). As w is increased, ρ_I must be decreased to recover the same number of triggers f .

It is seen from Fig. 4 that the number of first stage triggers decreases with increasing ρ_I , as expected. But more interestingly, we notice that less number of triggers are recovered when the averaging is done over more samples for the same ρ_I , i.e. bigger w . Therefore, to maintain the same sensitivity or the number of final triggers while w increases, the first stage threshold must be lowered.

A particular combination of the parameters (w, ρ_I) for no loss in sensitivity determines a specific target SNR ρ_{target} , as discussed in section III F. The target SNR is computed for a fixed w and different values of ρ_I , by

iterating over different values of second stage SNR ρ_{II} using Eq. (16). In Fig. 5 we plot the relationship between ρ_I and ρ_{target} for different w . Using the results obtained in Fig. 5, we can choose various combinations of w and ρ_I to reach a desired ρ_{target} . We notice from the plot that to reach the same ρ_{target} , bigger w requires a lower value of ρ_I . This information combined with the previous plot suggests that for a specific ρ_{target} , choosing a lower w will result in more number of first stage triggers.




FIG. 5. The plot shows different combinations of w and ρ_I to achieve a specific target SNR ρ_{target} without losing any sensitivity. The choice of w and ρ_I determines the total cost of the hierarchical method for the corresponding ρ_{target} .

Now, we combine the results from Fig. 4 and Fig. 5 to estimate the final costs Eq. (15) in terms of ρ_{target} as shown in the Fig. 6. These costs are normalized by the costs required by the template method Eq. (13) (shown by the horizontal orange line). To demonstrate the contributions in the total costs from the individual stages separately, we plot the first stage and the total costs together. The first stage costs are constant w.r.t. ρ_{target} and scale inversely with w . On the other hand, due to the rapid increase of the first stage triggers $f(w, \rho_I)$ at low SNRs, the second stage costs become dominant. It is also inferred from the figure that the second stage costs are more for larger w at a constant ρ_{target} . We notice from the Fig. 6, the total costs converge to the first stage costs at higher SNRs, and hence, infer that the first stage leads to the dominant costs for the hierarchical matched filtering.

From Fig. 6 we observe a reduction in total costs compared to the baseline for all choices of w complying to our desired target of $\rho_{target} \geq 4.0$. We notice for $\rho_{target} = 5.0$, the setting with $w = 8$ achieves a relative speed up factor of 8, which corresponds to a reduction of $\sim 88\%$ in the total costs. It is observed that the

hierarchical method performs better with increasing SNR thresholds. Our method achieves the best computational gain of ~ 18 times which is equivalent to a reduction of $\sim 94\%$ in the total costs, for $\rho_{target} = 6.0$ using $w = 16$. We have not tested the scheme for higher values of w , but by extrapolating the obtained results we may infer that the hierarchical method might perform better with $w > 16$ for even higher SNR thresholds.

To ensure that higher SNR triggers are recovered using only ρ_I , and also to sanity check our method for recovering CBC signals. We test our scheme for a population of 1000 CBC injections randomly generated within the case study region. For every injection, we create separate strains of 128 sec and place signals randomly into simulated data containing only colored Gaussian noise. The signals correspond to SNRs ranging between $[4.5, 30]$. We use two test cases of $\rho_{target} = 4.0$ and 6.0 using $w = 16$. Using the appropriate first stage cutoff ρ_I from Fig. 5 we recover all the injections that are above the respective test cases of SNR threshold. Therefore, we verify the reliability of our method to recover CBC signals using only ρ_I even at higher SNRs.

We now make a few key remarks. Our aim is to optimize the costs by fine-tuning the hierarchical parameters by recovering all the triggers with $\text{SNR} \geq 5$. The total costs of the hierarchical method are dominated by the first stage at higher SNR thresholds, which can be reduced by choosing a bigger w . On the other hand, the second stage costs are dominant at lower SNRs which grow as w increases. As seen in Fig. 6, for different SNRs we obtain different optimal window lengths, e.g. for $\rho_{target} = 5.0$ the most optimal setting is with $w = 8$. Depending upon the threshold criterion required for the search the optimal choice of w can be chosen based on Fig. 6. In case, the search demands higher SNR thresholds than shown in the plot, a larger w may be preferred to reduce the costs further.

2. Observed Performance

In this section, we measure the performance of matched filtering implementations on GPUs to estimate a realistic improvement compared to the established search pipelines. For this purpose, we use a widely quoted performance metric – the *throughput* of a search method, used for determining the number of templates analyzed in real-time. Consider a data segment of N secs filtered against T templates and the filtering process takes t seconds, then the throughput would be NT/t templates processed in real-time. Once again, we filter simulated colored Gaussian noise sampled at 2048 Hz for 64 seconds to evaluate this metric.

We benchmark the template method implemented using the optimized *cuFFT* library from CUDA. The template method in-situ took roughly 100ms to filter 64 secs of data against 6250 templates per GPU. Therefore, achieving an in-situ performance of $4000 \times$




FIG. 6. The computational cost (in terms of FLOP) for hierarchical filtering ~ 7.4 days of simulated data containing only Gaussian noise against 6250 templates. We show the costs for different (fixed) w which are normalized by the direct template filtering cost (orange line) versus the target SNR. The first stage costs are indicated by dot-dash lines, and the lines with markers correspond to the total costs. On the right, the vertical axis shows the relative speed-up factor (log scale) compared to the template method. The best speed up for a desired ρ_{target} is achieved by choosing w accordingly from this plot.

10^3 templates processed in real-time on a single Nvidia GV100. We want to remark that the second stage reconstruction is not optimized and thus, we could not benchmark the hierarchical scheme to its full potential. Considering the costs from Fig. 6, we estimate the expected peak performance of a completely optimized hierarchical method, which suggests that the hierarchical implementation may require only 12ms to perform the cuFFT equivalent filtering. Hence, we expect an increase of roughly an order of magnitude in the throughput (second row in Table II) if the second stage is fully optimized. Work is in progress for optimizing the second stage.

We now compare performances using previously quoted numbers from currently used search pipelines – PyCBC live [71] and PyCBC offline [50], as shown in Table II. The PyCBC search methods are implemented on multiple CPU cores, whereas ours are on multiple

GPUs. Depending upon the search method the throughput is standardized by templates processed in real-time per core (GPU) for PyCBC (hierarchical) implementation. We notice from the first column in Table II that using the latest GPUs gives an enormous improvement in the throughput. However, the PyCBC numbers are not quoted from an up-to-date hardware implementation and hence, a fair comparison might require the latest hardware.

Furthermore, we also present commercially motivated metrics to benchmark the performance, measuring the cost and energy efficiency of hardware while filtering. These metrics are computed by normalizing the throughput by the total cost or the energy consumption of the hardware respectively. The two metrics for the different search schemes are listed in Table II.

It is apparent from Table II, that due to GPU’s ability to perform tasks in a highly parallelized fashion,

Method	Throughput	Throughput/ Euro	Throughput/ W
cuFFT(in-situ)	4000×10^3	400	14×10^3
Hierarchical scheme (expected)	3300×10^4	3300	116×10^3
PyCBC live	6300	17	31
PyCBC offline	12,000	32	60

TABLE II. Comparing different implementations of the matched filtering schemes on GPUs (first and second row) with the established PyCBC schemes on CPUs (third and fourth row). In the second column we show the throughput for the respective methods, and in the third and fourth columns are the throughput per euro and per watt of power consumption respectively. The expected peak performance of the hierarchical method is estimated for $\text{SNR} = 5.0$ in the second row.

GPUs can analyze $\sim 10\times$ templates more than CPUs for the same costs. We also notice that considering the same power consumption a single GPU is equivalent to $\sim 10^3$ CPU cores for analyzing templates at a given instant. These metrics suggest that GPUs are highly energy and cost-efficient in performing matched filtering, which motivates their application in CBC searches. In addition, proper implementation of our hierarchical method will allow further improvement in the efficiency of the hardware and will help reduce the time required to perform extensive offline searches.

V. DISCUSSION AND FUTURE PROSPECTS

In this work, we have demonstrated using simulated data containing Gaussian noise an efficient way of matched filtering. We filter using a reduced bases and employ a new hierarchical method to reduce the reconstructions costs. Compared to the template based filtering, our method is $\sim 18\times$ faster than the template based filtering methods without losing sensitivity at a threshold $\text{SNR} = 6$, and $8\times$ for $\text{SNR} = 5$. The gain in performance increases with higher SNR thresholds and is currently estimated for a specific region of the parameter space. Our method is successful in recovering the original flat search background, and thus, does not compromise the significance of detected candidates with SNR above the SNR threshold.

We demonstrate the advantages of implementing matched filtering methods on the latest GPUs. We compare the throughput of GPU implementation of matched filtering with the CPU implementation of current methods. Benchmarking the in-situ performance of template method implementation on GPUs, we observe a performance gain of 2-3 orders in magnitude compared to the PyCBC search pipelines. Our results

indicate a significant improvement in performance, which may motivate the development of a fully optimized second stage reconstruction. In addition, we present two new metrics to compare the performance of the matched filtering implementation on different hardware. Analyzing these metrics suggests that GPUs are more cost and energy-efficient in performing matched filtering than CPUs. Hence, the utilization of GPUs is encouraged for current or future searches.

A possible avenue to improve the described method would be to find better ways of performing the first stage and a faster implementation of the second stage. In this work, we use a constant sampling rate for matched filtering. Multi-rate sampling can be implemented to further improve the performance of the hierarchical method for cases where latency is a strong requirement or the duration of signals is significantly longer than those tested here.

In the near future, detectors will become more sensitive and thus the cost-effective hierarchical method proposed here can be useful for exploring sub-solar regimes or searching for low frequency long duration signals. Our method might play a role in reducing the computational costs for the future 3G detectors where the template bank size can be at least an order of magnitude larger [47, 73] than the current CBC banks. Furthermore, this method can also be employed in new regions of the parameter space to perform computationally intensive searches for sources exhibiting precession or eccentricity.

ACKNOWLEDGMENTS

We thank Badri Krishnan and Kipp Cannon for providing useful comments. We would like to acknowledge Marlin Schäfer, Yifan Wang for reading the manuscript and suggestions. We thank the Atlas

cluster computing team at Albert-Einstein Institute (AEI) Hannover for support.

-
- [1] B. P. et al. Abbott (LIGO Scientific Collaboration and Virgo Collaboration), “Observation of gravitational waves from a binary black hole merger,” *Phys. Rev. Lett.* **116**, 061102 (2016).
- [2] B. P. Abbott *et al.* (LIGO Scientific, Virgo), “GW150914: The Advanced LIGO Detectors in the Era of First Discoveries,” *Phys. Rev. Lett.* **116**, 131103 (2016), arXiv:1602.03838 [gr-qc].
- [3] F. Acernese *et al.* (VIRGO), “Advanced Virgo: a second-generation interferometric gravitational wave detector,” *Class. Quant. Grav.* **32**, 024001 (2015), arXiv:1408.3978 [gr-qc].
- [4] B. P. Abbott *et al.* (LIGO Scientific, Virgo), “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), arXiv:1811.12907 [astro-ph.HE].
- [5] Alexander H. Nitz, Thomas Dent, Gareth S. Davies, Sumit Kumar, Collin D. Capano, Ian Harry, Simone Mozzon, Laura Nuttall, Andrew Lundgren, and Márton Tápai, “2-OGC: Open Gravitational-wave Catalog of binary mergers from analysis of public Advanced LIGO and Virgo data,” *Astrophys. J.* **891**, 123 (2020), arXiv:1910.05331 [astro-ph.HE].
- [6] R. Abbott *et al.* (LIGO Scientific, Virgo), “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), arXiv:2010.14527 [gr-qc].
- [7] Alexander H. Nitz, Collin D. Capano, Sumit Kumar, Yi-Fan Wang, Shilpa Kastha, Marlin Schäfer, Rahul Dhurkunde, and Miriam Cabero, “3-OGC: Catalog of gravitational waves from compact-binary mergers,” (2021), arXiv:2105.09151 [astro-ph.HE].
- [8] B. P. Abbott *et al.* (LIGO Scientific, Virgo), “GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral,” *Phys. Rev. Lett.* **119**, 161101 (2017), arXiv:1710.05832 [gr-qc].
- [9] B. P. Abbott *et al.* (LIGO Scientific, Virgo), “GW190425: Observation of a Compact Binary Coalescence with Total Mass $\sim 3.4M_{\odot}$,” *Astrophys. J. Lett.* **892**, L3 (2020), arXiv:2001.01761 [astro-ph.HE].
- [10] R. Abbott *et al.* (LIGO Scientific, KAGRA, VIRGO), “Observation of Gravitational Waves from Two Neutron Star–Black Hole Coalescences,” *Astrophys. J. Lett.* **915**, L5 (2021), arXiv:2106.15163 [astro-ph.HE].
- [11] B. S. Sathyaprakash and B. F. Schutz, “Physics, Astrophysics and Cosmology with Gravitational Waves,” *Living Rev. Rel.* **12**, 2 (2009), arXiv:0903.0338 [gr-qc].
- [12] Charalampos Markakis, Jocelyn S. Read, Masaru Shibata, Koji Uryu, Jolien D. E. Creighton, John L. Friedman, and Benjamin D. Lackey, “Neutron star equation of state via gravitational wave observations,” *J. Phys. Conf. Ser.* **189**, 012024 (2009), arXiv:1110.3759 [gr-qc].
- [13] B. P. Abbott *et al.* (LIGO Scientific, Virgo), “The Rate of Binary Black Hole Mergers Inferred from Advanced LIGO Observations Surrounding GW150914,” *Astrophys. J. Lett.* **833**, L1 (2016), arXiv:1602.03842 [astro-ph.HE].
- [14] B. P. Abbott *et al.* (LIGO Scientific, Virgo), “Binary Black Hole Population Properties Inferred from the First and Second Observing Runs of Advanced LIGO and Advanced Virgo,” *Astrophys. J. Lett.* **882**, L24 (2019), arXiv:1811.12940 [astro-ph.HE].
- [15] Michele Maggiore *et al.*, “Science Case for the Einstein Telescope,” *JCAP* **03**, 050 (2020), arXiv:1912.02622 [astro-ph.CO].
- [16] David *et al.* Reitze, “Cosmic explorer: The u.s. contribution to gravitational-wave astronomy beyond ligo,” *Bulletin of the AAS* **51** (2019), <https://baas.aas.org/pub/2020n7i035>.
- [17] The eLISA Consortium, “The gravitational universe,” (2013), arXiv:1305.5720 [astro-ph.CO].
- [18] Ken K. Y. Ng, Shiqi Chen, Boris Goncharov, Ulyana Dupletsa, Ssohrab Borhanian, Marica Branchesi, Jan Harms, Michele Maggiore, B. S. Sathyaprakash, and Salvatore Vitale, “On the single-event-based identification of primordial black hole mergers at cosmological distances,” (2021), arXiv:2108.07276 [astro-ph.CO].
- [19] Stanislav *et al.* Babak, “Science with the space-based interferometer LISA. V: Extreme mass-ratio inspirals,” *Phys. Rev. D* **95**, 103012 (2017), arXiv:1703.09722 [gr-qc].
- [20] Matthew Evans *et al.*, “A horizon study for cosmic explorer: Science, observatories, and community,” (2021), arXiv:2109.09882 [astro-ph.IM].
- [21] Surabhi Sachdev *et al.*, “The GstLAL Search Analysis Methods for Compact Binary Mergers in Advanced LIGO’s Second and Advanced Virgo’s First Observing Runs,” (2019), arXiv:1901.08580 [gr-qc].
- [22] F. Aubin *et al.*, “The MBTA pipeline for detecting compact binary coalescences in the third LIGO–Virgo observing run,” *Class. Quant. Grav.* **38**, 095004 (2021), arXiv:2012.11512 [gr-qc].
- [23] Xiaoyang Guo, Qi Chu, Zhihui Du, and Linqing Went, “GPU-Optimised Low-Latency Online Search for Gravitational Waves from Binary Coalescences,” (2018).
- [24] “Gravitational-wave data analysis,” in *Gravitational-Wave Physics and Astronomy* (John Wiley & Sons, Ltd, 2011).
- [25] B. P. Abbott *et al.* (LIGO Scientific, Virgo), “Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914,” *Class. Quant. Grav.* **33**, 134001 (2016), arXiv:1602.03844 [gr-qc].
- [26] Miriam Cabero *et al.*, “Blip glitches in Advanced LIGO data,” *Class. Quant. Grav.* **36**, 15 (2019), arXiv:1901.05093 [physics.ins-det].
- [27] Bruce Allen, “ χ^2 time-frequency discriminator for gravitational wave detection,” *Phys. Rev. D* **71**, 062001 (2005), arXiv:gr-qc/0405045.

- [28] B P Abbott *et al.* (LIGO Scientific, Virgo), “Effects of data quality vetoes on a search for compact binary coalescences in Advanced LIGO’s first observing run,” *Class. Quant. Grav.* **35**, 065010 (2018), arXiv:1710.02185 [gr-qc].
- [29] Bruce Allen, Warren G. Anderson, Patrick R. Brady, Duncan A. Brown, and Jolien D. E. Creighton, “FINDCHIRP: An Algorithm for detection of gravitational waves from inspiraling compact binaries,” *Phys. Rev. D* **85**, 122006 (2012), arXiv:gr-qc/0509116.
- [30] Sebastian Khan, Katerina Chatziioannou, Mark Hannam, and Frank Ohme, “Phenomenological model for the gravitational-wave signal from precessing binary black holes with two-spin effects,” *Phys. Rev. D* **100**, 024059 (2019), arXiv:1809.10113 [gr-qc].
- [31] Alejandro Bohé *et al.*, “Improved effective-one-body model of spinning, nonprecessing binary black holes for the era of gravitational-wave astrophysics with advanced detectors,” *Phys. Rev. D* **95**, 044028 (2017), arXiv:1611.03703 [gr-qc].
- [32] Jonathan Blackman, Scott E. Field, Mark A. Scheel, Chad R. Galley, Christian D. Ott, Michael Boyle, Lawrence E. Kidder, Harald P. Pfeiffer, and Béla Szilágyi, “Numerical relativity waveform surrogate model for generically precessing binary black hole mergers,” *Phys. Rev. D* **96**, 024058 (2017), arXiv:1705.07089 [gr-qc].
- [33] Soumen Roy, Anand S. Sengupta, and Nilay Thakor, “Hybrid geometric-random template-placement algorithm for gravitational wave searches from compact binary coalescences,” *Phys. Rev. D* **95**, 104045 (2017), arXiv:1702.06771 [gr-qc].
- [34] Soumen Roy, Anand S. Sengupta, and Parameswaran Ajith, “Effectual template banks for upcoming compact binary searches in Advanced-LIGO and Virgo data,” *Phys. Rev. D* **99**, 024048 (2019), arXiv:1711.08743 [gr-qc].
- [35] B P Abbott *et al.* (LIGO Scientific, Virgo), “Effects of data quality vetoes on a search for compact binary coalescences in Advanced LIGO’s first observing run,” *Class. Quant. Grav.* **35**, 065010 (2018), arXiv:1710.02185 [gr-qc].
- [36] Derek Davis *et al.* (LIGO), “LIGO detector characterization in the second and third observing runs,” *Class. Quant. Grav.* **38**, 135014 (2021), arXiv:2101.11673 [astro-ph.IM].
- [37] Alexander H. Nitz, Collin Capano, Alex B. Nielsen, Steven Reyes, Rebecca White, Duncan A. Brown, and Badri Krishnan, “1-OGC: The first open gravitational-wave catalog of binary mergers from analysis of public Advanced LIGO data,” *Astrophys. J.* **872**, 195 (2019), arXiv:1811.01921 [gr-qc].
- [38] Leone Bosi and Edward K. Porter, “Data Analysis Challenges for the Einstein Telescope,” *Gen. Rel. Grav.* **43**, 519–535 (2011), arXiv:0910.0380 [gr-qc].
- [39] Alexander H. Nitz and Yi-Fan Wang, “Search for gravitational waves from the coalescence of sub-solar mass and eccentric compact binaries,” (2021), 10.3847/1538-4357/ac01d9, arXiv:2102.00868 [astro-ph.HE].
- [40] B. P. Abbott *et al.* (LIGO Scientific, Virgo), “Search for Subsolar Mass Ultracompact Binaries in Advanced LIGO’s Second Observing Run,” *Phys. Rev. Lett.* **123**, 161102 (2019), arXiv:1904.08976 [astro-ph.CO].
- [41] Yi-Fan Wang and Alexander H. Nitz, “Prospects for detecting gravitational waves from eccentric subsolar mass compact binaries,” *Astrophys. J.* **912**, 53 (2021), arXiv:2101.12269 [astro-ph.HE].
- [42] Mark Hannam, “Modelling gravitational waves from precessing black-hole binaries: Progress, challenges and prospects,” *Gen. Rel. Grav.* **46**, 1767 (2014), arXiv:1312.3641 [gr-qc].
- [43] B. P. Abbott *et al.* (LIGO Scientific, Virgo), “Search for Eccentric Binary Black Hole Mergers with Advanced LIGO and Advanced Virgo during their First and Second Observing Runs,” *Astrophys. J.* **883**, 149 (2019), arXiv:1907.09384 [astro-ph.HE].
- [44] Alexander H. Nitz, Amber Lenon, and Duncan A. Brown, “Search for Eccentric Binary Neutron Star Mergers in the first and second observing runs of Advanced LIGO,” *Astrophys. J.* **890**, 1 (2019), arXiv:1912.05464 [astro-ph.HE].
- [45] Thibault Damour, Achamveedu Gopakumar, and Bala R. Iyer, “Phasing of gravitational waves from inspiralling eccentric binaries,” *Phys. Rev. D* **70**, 064028 (2004), arXiv:gr-qc/0404128.
- [46] Theocharis A. Apostolatos, Curt Cutler, Gerald J. Sussman, and Kip S. Thorne, “Spin induced orbital precession and its modulation of the gravitational wave forms from merging binaries,” *Phys. Rev. D* **49**, 6274–6297 (1994).
- [47] Ian Harry, Stephen Privitera, Alejandro Bohé, and Alessandra Buonanno, “Searching for Gravitational Waves from Compact Binaries with Precessing Spins,” *Phys. Rev. D* **94**, 024012 (2016), arXiv:1603.02444 [gr-qc].
- [48] Kipp Cannon, Adrian Chapman, Chad Hanna, Drew Keppel, Antony C. Searle, and Alan J. Weinstein, “Singular value decomposition applied to compact binary coalescence gravitational-wave signals,” *Phys. Rev. D* **82**, 044025 (2010), arXiv:1005.0012 [gr-qc].
- [49] B. P. Abbott *et al.* (LIGO Scientific, Virgo), “Low-latency Gravitational-wave Alerts for Multimessenger Astronomy during the Second Advanced LIGO and Virgo Observing Run,” *Astrophys. J.* **875**, 161 (2019), arXiv:1901.03310 [astro-ph.HE].
- [50] Samantha A. Usman *et al.*, “The PyCBC search for gravitational waves from compact binary coalescence,” *Class. Quant. Grav.* **33**, 215004 (2016), arXiv:1508.02357 [gr-qc].
- [51] T. Adams, D. Buskulic, V. Germain, G. M. Guidi, F. Marion, M. Montani, B. Mours, F. Piergiovanni, and G. Wang, “Low-latency analysis pipeline for compact binary coalescences in the advanced gravitational wave detector era,” *Class. Quant. Grav.* **33**, 175012 (2016), arXiv:1512.02864 [gr-qc].
- [52] S. Babak *et al.*, “Searching for gravitational waves from binary coalescence,” *Phys. Rev. D* **87**, 024033 (2013), arXiv:1208.3491 [gr-qc].
- [53] B. S. Sathyaprakash and S. V. Dhurandhar, “Choice of filters for the detection of gravitational waves from coalescing binaries,” *Phys. Rev. D* **44**, 3819–3834 (1991).
- [54] S. Babak, R. Balasubramanian, D. Churches, T. Cokelaer, and B. S. Sathyaprakash, “A Template bank to search for gravitational waves from inspiralling compact binaries. I. Physical models,” *Class. Quant. Grav.* **23**, 5477–5504 (2006), arXiv:gr-qc/0604037.

- [55] P. Ajith, N. Fotopoulos, S. Privitera, A. Neunzert, and A. J. Weinstein, “Effectual template bank for the detection of gravitational waves from inspiralling compact binaries with generic spins,” *Phys. Rev. D* **89**, 084041 (2014), arXiv:1210.6666 [gr-qc].
- [56] Ian W. Harry, Bruce Allen, and B. S. Sathyaprakash, “A Stochastic template placement algorithm for gravitational wave data analysis,” *Phys. Rev. D* **80**, 104014 (2009), arXiv:0908.2090 [gr-qc].
- [57] Prasanna Joshi, Rahul Dhurkunde, Sanjeev Dhurandhar, and Sukanta Bose, “Optimal χ^2 discriminator against modeled noise transients in interferometric data in searches for binary black-hole mergers,” *Phys. Rev. D* **103**, 044035 (2021), arXiv:2006.12901 [gr-qc].
- [58] Samantha A. Usman *et al.*, “The PyCBC search for gravitational waves from compact binary coalescence,” *Class. Quant. Grav.* **33**, 215004 (2016), arXiv:1508.02357 [gr-qc].
- [59] Shaun Hooper, Shin Kee Chung, Jing Luan, David Blair, Yanbei Chen, and Linqing Wen, “Summed Parallel Infinite Impulse Response (SPIIR) Filters For Low-Latency Gravitational Wave Detection,” *Phys. Rev. D* **86**, 024012 (2012), arXiv:1108.3186 [gr-qc].
- [60] Kipp Cannon, Romain Cariou, Adrian Chapman, Mireia Crispin-Ortuzar, Nickolas Fotopoulos, Melissa Frei, Chad Hanna, Erin Kara, Drew Keppel, Laura Liao, Stephen Privitera, Antony Searle, Leo Singer, and Alan Weinstein, “TOWARD EARLY-WARNING DETECTION OF GRAVITATIONAL WAVES FROM COMPACT BINARY COALESCENCE,” *The Astrophysical Journal* **748**, 136 (2012).
- [61] Bhooshan Gadre, Sanjit Mitra, and Sanjeev Dhurandhar, “Hierarchical search strategy for the efficient detection of gravitational waves from nonprecessing coalescing compact binaries with aligned-spins,” *Phys. Rev. D* **99**, 124035 (2019), arXiv:1807.06803 [astro-ph.IM].
- [62] S. D. Mohanty and S. V. Dhurandhar, “A hierarchical search strategy for the detection of gravitational waves from coalescing binaries,” *Phys. Rev. D* **54**, 7108–7128 (1996).
- [63] Anand S. Sengupta, Sanjeev V. Dhurandhar, Albert Lazzarini, and Tom Prince, “Extended hierarchical search (EHS) algorithm for detection of gravitational waves from inspiraling compact binaries,” *Class. Quant. Grav.* **19**, 1507–1512 (2002), arXiv:gr-qc/0109088.
- [64] C.E. Shannon, “Communication in the presence of noise,” *Proceedings of the IRE* **37**, 10–21 (1949).
- [65] Tito Dal Canton and Ian W. Harry, “Designing a template bank to observe compact binary coalescences in Advanced LIGO’s second observing run,” (2017), arXiv:1705.01845 [gr-qc].
- [66] Xuansheng Wang, Beidun Chen, Jianqiang Sheng, Hongying Zheng, Tangren Dan, and Xianfeng Wu, “An improved lanczos algorithm for principal component analysis,” in *Proceedings of 2020 the 6th International Conference on Computing and Data Engineering, ICCDE 2020* (Association for Computing Machinery, New York, NY, USA, 2020) p. 70–74.
- [67] NVIDIA, Péter Vingelmann, and Frank H.P. Fitzek, “Cuda, release: 10.2.89,” (2020).
- [68] V. Hernandez, J. E. Roman, A. Tomas, and V. Vidal, *A survey of software for sparse eigenvalue problems*, Tech. Rep. STR-6 (Universitat Politècnica de València, 2009) available at <https://slep.c.upv.es>.
- [69] Satish Balay, Shrirang Abhyankar, Mark F. Adams, Jed Brown, Peter Brune, Kris Buschelman, Lisandro Dalcin, Alp Dener, Victor Eijkhout, William D. Gropp, Dinesh Kaushik, Matthew G. Knepley, Dave A. May, Lois Curfman McInnes, Richard Tran Mills, Todd Munson, Karl Rupp, Patrick Sanan, Barry F. Smith, Stefano Zampini, Hong Zhang, and Hong Zhang, “PETSc Web page,” (2021), <https://petsc.org/>.
- [70] Steven G. Johnson and Matteo Frigo, “A modified split-radix fft with fewer arithmetic operations,” *IEEE Transactions on Signal Processing* **55**, 111–119 (2007).
- [71] Alexander H. Nitz, Tito Dal Canton, Derek Davis, and Steven Reyes, “Rapid detection of gravitational waves from compact binary mergers with PyCBC Live,” *Phys. Rev. D* **98**, 024050 (2018), arXiv:1805.11174 [gr-qc].
- [72] Alex Nitz *et al.*, “gwastro/pycbc,” (2021).
- [73] Leone Bosi and Edward K. Porter, “Data Analysis Challenges for the Einstein Telescope,” *Gen. Rel. Grav.* **43**, 519–535 (2011), arXiv:0910.0380 [gr-qc].