

Fully automated end-to-end pipeline for massive black hole binary signal extraction from LISA data

Michael L. Katz¹*

*Max-Planck-Institut für Gravitationsphysik, Albert-Einstein-Institut,
Am Mühlenberg 1, 14476 Potsdam-Golm, Germany*

 (Received 1 November 2021; accepted 27 January 2022; published 25 February 2022)

The LISA Data Challenges Working Group within the LISA Consortium has started publishing datasets to benchmark, compare, and build LISA data analysis infrastructure as the Consortium prepares for the launch of the mission. We present our solution to the dataset from LISA Data Challenge (LDC) 1A containing a single massive black hole binary signal. This solution is built from a fully automated and graphics-processing-unit-accelerated pipeline consisting of three segments: a brute-force initial search, a refining search that uses the efficient likelihood computation technique of heterodyning (also called relative binning) to locate the maximum likelihood point, and a parameter estimation portion that also takes advantage of the speed of the heterodyning method. This pipeline takes tens of minutes to evolve from randomized initial parameters throughout the prior volume to a converged final posterior distribution. Final posteriors are shown for both datasets from LDC 1A: one noiseless data stream and one containing additive noise. A posterior distribution including higher harmonics is also shown for a self-injected waveform with the same source parameters as is used in the original LDC 1A dataset. This higher-mode posterior is shown in order to provide a more realistic distribution on the parameters of the source.

DOI: [10.1103/PhysRevD.105.044055](https://doi.org/10.1103/PhysRevD.105.044055)

I. INTRODUCTION

The era of gravitational-wave physics and astronomy is here in full force. Detections of ~ 10 s of binaries by ground-based observing groups [1,2] have shown the power of this new messenger for understanding the Universe. The future space-based gravitational wave detector, the Laser Interferometer Space Antenna (LISA) [3], will open a new window into the milliHertz regime of the gravitational-wave spectrum. This frequency range contains many sources of interest including massive black hole binaries (MBHB) stemming from the mergers of their host galaxies (e.g., [4]), stellar-origin black hole binaries earlier in their evolution that can act as progenitors to ground-based sources (e.g., [5]), extreme-mass-ratio-inspirals involving a compact object bound to an MBH in a close and energetic orbit (e.g., [6]), and Galactic binaries typically consisting of two white dwarf stars in a bound system at large separations exhibiting a quasimonochromatic signal (e.g., [7]).

MBHBs are a primary focus of the LISA mission due to their unique capabilities in helping to answer many scientific questions in astrophysics, cosmology, and fundamental physics [3,8–15]. This includes probing the mass spectrum of MBHBs and their formation channels, as well as better understanding their interplay within the formation and evolution of galaxies. These systems with total detector-frame masses of $\sim 10^5$ – $10^7 M_\odot$ are expected to have high signal-to-noise ratios (SNRs) of order ~ 100 – 1000 s, allowing for in-depth extraction of parameters and detailed tests of general relativity. Recent event rate estimates for MBHBs detectable by LISA range from ~ 1 – 20 per year, depending on the underlying assumptions [15–19].

While there is a large amount of science that can be done by observing these MBHB systems, properly analyzing the LISA data is not an easy task. All of the sources (including all other source classes) will be emitting signals that will overlap throughout the lifetime of the LISA mission, both requiring the use of global fitting of source parameters and preventing direct access to pure signal-free noise information as is available for current ground-based observing runs. Due to the long-lived signals, the time dependence of the detector orientation will need to be included in the analysis to ensure unbiased extraction of source parameters. There is expected to be nonstationary noise effects in the form of data gaps and glitches, as well as expected drift of the overall instrument sensitivity over its observing lifetime. The new LISA Data Challenges (LDC) Working Group is

*michael.katz@aei.mpg.de

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Open access publication funded by the Max Planck Society.

tasked with slowly incorporating all of these effects and building the initial data analysis pipelines that will, over time, prove the capability of the LISA community to properly extract source parameters and scientific information from the future realistic LISA data stream. The LDC provides datasets that help researchers to compare and contrast methods and results for the same underlying data, allowing the community to more fully understand its range of capabilities.

Earlier work over the past few decades to address the analysis of MBHBs with LISA can be found in [20–43]. More recent analyses in [44,45] were performed with a more modern LISA sensitivity [3,46], full inspiral-merger-ring-down waveform models that included higher harmonic modes, and a more accurate representation of the LISA response from [47]. Waveforms in [45] also contained aligned spins. The effect of data gaps on MBHB analysis was recently looked at without noise in [48]. There have also been papers addressing the simultaneous analysis of MBHBs with LISA and other space-based detectors (e.g., [49–51]).

The first round of the new LDC datasets (LDC-1A) focuses on single sources injected into stationary instrumental LISA noise. *This work will present a solution for the LDC-1A dataset containing a single MBHB source. Specifically, we will describe a fully automated end-to-end pipeline capable of extracting accurate posterior distributions on source parameters in under one hour of runtime with no human intervention or prior knowledge of the signal in the data.* Authors in [52] presented a full pipeline for the search and parameter estimation of MBHBs in response to the LDC-1A dataset. This was recently extended to the MBHB portion of the new LDC-2A challenge [53]. This work is similar in its origins and goals to our work presented here. The duration of the pipelines between the two works are similar. The parameter estimation stage of the pipelines are both run with fast, residual-based likelihood methods, but the search stages are different. In [52], the search portion relies on a sequence of marginalizations over certain parameters and a fixed LISA detector. This setup improves the speed of the search, but it requires updating when more physics are included (e.g., precession). In our work, we achieve a fast search by using computational acceleration from graphics processing units (GPUs), rather than marginalization. Our search uses a brute-force approach in the initial stages, which should extend to more complicated waveform models, as well as searching for sources in their early inspiral where movement of the LISA detector will be required for the earliest possible detection. We follow this initial stage with a maximum likelihood refining step that makes use of the very efficient likelihood method of heterodyning [54,55]. We use the specific heterodyning formulation given in [56]. In [56], they refer to their heterodyning implementation as “relative binning.”

In addition to providing a solution to the LDC datasets, we also provide a posterior distribution with a waveform

containing higher harmonic modes for the same source parameters as the signal in the LDC data. The LDC data only includes the dominant $l = m = 2$ harmonic. The higher-harmonic posterior is provided to give the reader a better understanding of the constraints on parameters and the posterior distribution in general when higher modes are included in the waveform. MBHBs detectable with LISA are expected to have multiple detectable harmonics given the high SNR of these sources.

We expect the methods presented here to extend to datasets with multiple MBHBs. This assumes the predicted rates for MBHB detections are roughly correct, meaning the merger-ringdown portion of coalescence, where most of the SNR is accumulated for MBHBs, is unlikely to overlap with another signal’s merger ringdown since this portion of the signal has a short duration of \sim hour or less. MBHBs that have overlapping merger-ringdown signals will require more sophisticated methods such as global fitting.

In Sec. II we briefly describe the waveform, LISA response function, and likelihood computations used in this work. The pipeline that uses these methods and the final posteriors it produces are provided in Sec. III. We then discuss the pipeline in the larger context of LISA data analysis in Sec. IV and conclude our remarks in Sec. V.

II. BAYESIAN METHODS FOR MASSIVE BLACK HOLE BINARY PARAMETER EXTRACTION WITH LISA

The gravitational-wave data stream, $d(t)$, is usually represented as a sum of noise and signal contributions: $d(t) = s(t) + n(t)$, where $s(t)$ is the true underlying gravitational-wave signal and $n(t)$ is the noise contribution. In our case, $d(t)$ is provided by the LDC dataset. The data from the LDC are shown in both the time and frequency domain in Fig. 1. We will analyze two datasets with the same $s(t)$ injection but with different $n(t)$. In one case, $n(t) = 0$ for all values of t , which we will refer to as the “noiseless” case. In the “noisy” or “noise-infused” dataset, $n(t)$, which is modeled as a stochastic process, is generated by the LDC Working Group. The noise curve used is the “SciRDv1” noise curve from the LDC code base [46]. The Fourier transform of the noise is represented as $\tilde{n}(f) = \mathcal{F}\{n(t)\}$. Because the noise in the time domain is assumed to be Gaussian and stationary, the noise in the frequency domain is Gaussian with mean zero and a covariance matrix that is diagonal with values given by the power spectral density in the noise: $\langle \tilde{n}(f)\tilde{n}(f') \rangle = \frac{1}{2}S_n(f)\delta(f - f')$ (with $f \geq 0$), where $S_n(f)$ represents the power in the noise at a given frequency.

The goal of parameter estimation is to determine the posterior distribution on the parameters that describe the underlying signal given the specific set of noisy data and the modeling of the data generation process. This conditional probability can be determined using Bayes’ theorem:

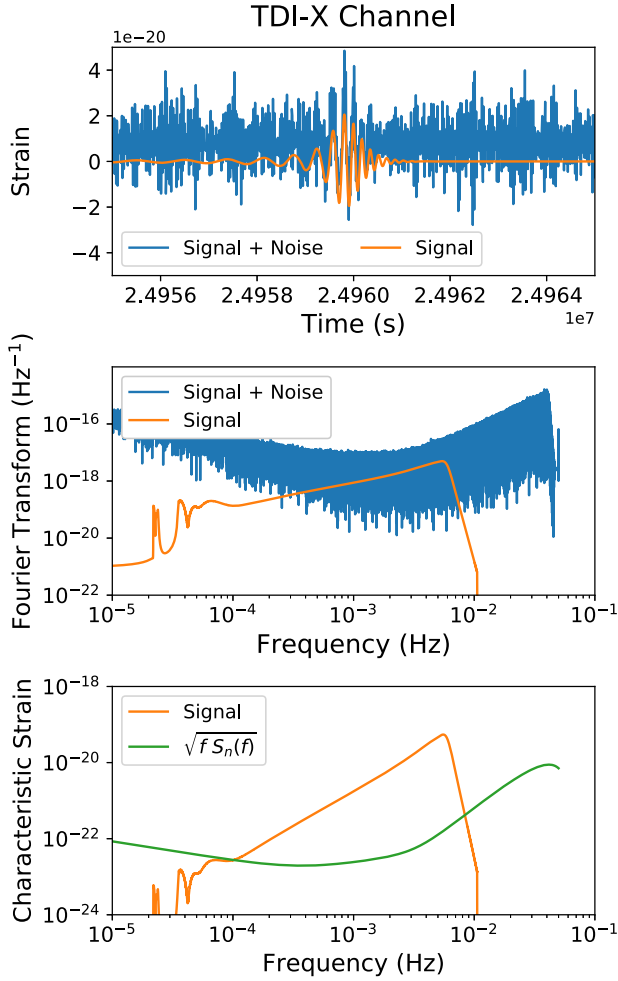


FIG. 1. The data from the LDC-1A dataset containing one MBHB signal. The top panel shows the time domain representation of the TDI-X channel (see Sec. II A), with an enlargement around the merger ringdown. The middle panel is the corresponding frequency domain representation shown as the Fourier transform of the data. The dataset without noise is shown in orange, while the dataset containing noise is shown in blue. The bottom panel shows the noiseless signal (orange) in the characteristic strain representation: $h_c^2 = 4|f^2 \tilde{h}(f)|^2$. The sensitivity curve (green) is also shown as the characteristic strain in the noise: $h_n^2 = f S_n(f)$ [57,58].

$$p(\vec{\Theta}|d, \Lambda) = \frac{p(d|\vec{\Theta}, \Lambda)p(\vec{\Theta}|\Lambda)}{p(d|\Lambda)}, \quad (1)$$

where $\vec{\Theta}$ is the vector of parameters that describe the signal $s(t)$ and Λ is the chosen underlying model. In reality, Λ will be an approximation to the true signal given by nature. However, in this work, we assume that $s(t)$ is exactly given by Λ and avoid assessing any biases due to incorrect modeling. We will, therefore, drop Λ from the notation for the remainder of this paper. See [59–61] for information on waveform modeling errors related to the waveforms used in this paper.

On the right-hand side of Eq. (1), $p(\vec{\Theta})$ represents the prior probability distribution of the parameters that describe the signal. This helps to weight the posterior distribution based on prior knowledge. In the denominator, there is the evidence, which represents the marginalization of the numerator over all of parameter space: $p(d) = \int_{\vec{\Theta}} p(d|\vec{\Theta})p(\vec{\Theta})d\vec{\Theta}$. In gravitational-wave applications, the evidence is intractable to calculate directly. In Markov chain Monte Carlo (MCMC) simulations, the method that will be used to determine the posterior, the evidence enters only as a multiplicative factor. Therefore, we can neglect its direct computation here. The evidence could be computed in MCMC using stepping-stone sampling [62] or thermodynamic integration [63,64].

The key component of Eq. (1), where actual signal modeling and gravitational-wave specific computations are included, is the likelihood $\mathcal{L} = p(d|\vec{\Theta})$. The gravitational-wave likelihood is determined by matching a template, $h(t)$, generated using parameters $\vec{\Theta}$, to the data stream. In a noiseless situation, if $h(t)$ is the correct signal model, then the maximum likelihood of zero would occur where $\vec{\Theta} = \vec{\Theta}_{\text{true}}$, giving $h(t) = s(t)$. The likelihood for a specific template is given by

$$\begin{aligned} \log \mathcal{L} &\propto -\frac{1}{2} \langle d - h | d - h \rangle \\ &= -\frac{1}{2} (\langle d | d \rangle + \langle h | h \rangle - 2 \langle d | h \rangle), \end{aligned} \quad (2)$$

where we define $\langle a | b \rangle$ to be the noise-weighted inner product between time-domain data streams $a(t)$ and $b(t)$. This inner product is general and can include nonstationary noise effects by including a proper noise-covariance matrix. Under the assumption of Gaussianity and stationarity, we can write down the inner product necessary for the likelihood:

$$\langle a | b \rangle = 4 \text{Re} \sum_{i=AET} \int_0^\infty \frac{\tilde{a}^i(f) \tilde{b}^i(f)}{S_n^i(f)} df, \quad (3)$$

where the sum is over the three time-delay interferometry observables A , E , T (see below in Sec. II A). Notice the frequency range is over positive frequencies because the signals are real valued. In this work, we assume $S_n^i(f)$ are known analytic functions, and, therefore, neglect assessing or including their uncertainty. The noise curves for time-delay interferometry (TDI) observables A and E are identical, but T is different. Their formulas are given in [44].

The square root of the template-only term in the likelihood, $\sqrt{\langle h | h \rangle}$, represents the optimal SNR achievable for a given template. In general, the parameters we build the template with will not be exactly equal to the true parameters (where $\mathcal{L} = 0$ in the noiseless case). The extracted SNR of a specific template against the data is

given by $\langle d|h \rangle / \sqrt{\langle h|h \rangle}$. The extracted SNR will be an important component of the automated extraction pipeline.

A. Template generation

The main component necessary for computing the likelihood is the template. The LDC dataset is created by generating the waveform in the frequency domain and then transforming it to the time domain. For this dataset, the PHENOMD waveform model [59,60] is used. This model includes only the $l = m = 2$ harmonic with aligned spins. In reality, multiple harmonics for MBHBs will be observable. To get a sense of the posterior that would be produced with the inclusion of the higher harmonic modes, we will use the PHENOMHM waveform [61] by injecting a signal with parameters equivalent to the LDC injection parameters. This will be discussed further in Sec. III C.

The time-domain transform of the LDC injection waveform is then put through the time-domain LISA response function that assumes equal-armlength orbits for the LISA spacecraft. This step projects the signal on to the arms of the LISA constellation and then combines these projections using TDI to obtain the three TDI observables: X , Y , Z [65–69]. The LDC provides these three observables. These observables are correlated in their noise properties, so they are transformed into uncorrelated observables: A , E , and T . This transform is given by [70]

$$A = \frac{1}{\sqrt{2}}(Z - X), \quad (4)$$

$$E = \frac{1}{\sqrt{6}}(X - 2Y + Z), \quad (5)$$

$$T = \frac{1}{\sqrt{3}}(X + Y + Z). \quad (6)$$

When generating the template, it is these uncorrelated observables that are determined for each likelihood computation.

For our current pipeline, we will work in the frequency domain and generate the Fourier transform of $h(t)$: $\tilde{h}(f)$. Note that, prior to the analysis, we transform the data to get $\tilde{d}(f)$. Quasicircular aligned-spin waveforms, like PHENOMD and PHENOMHM, require 11 parameters: $\vec{\Theta} = \{M_T, q, a_1, a_2, D_L, \phi_{\text{ref}}, \iota, \lambda, \beta, \psi, t_{\text{ref}}\}$. M_T is the binary total mass: $M_T = m_1 + m_2$. The mass ratio is $q = m_2/m_1$ with $m_2 < m_1$. The dimensionless spins, with range $-0.99 \leq a_i \leq 0.99$, for m_1 and m_2 are a_1 and a_2 , respectively. The spins are aligned with the orbital angular momentum of the system: a negative spin indicates anti-parallel alignment. The luminosity distance is given as D_L . The reference phase and polarization angles are ϕ_{ref} and ψ , respectively. The orbital inclination is ι . The ecliptic longitude is λ and β is the ecliptic latitude. The time of

coalescence is t_{ref} . In this work, t_{ref} and ϕ_{ref} are both set at $f_{\text{ref}} = f_{\text{max}}$, where f_{max} is a value determined internally in the PHENOMD code where the value of $f^2 A_{22}(f)$ is maximized [A_{22} is the amplitude of the (2,2) mode]. All of the extrinsic parameters given here are defined with respect to the LISA constellation reference frame.

With the intrinsic parameters and the luminosity distance, we can generate the source frame waveform scaled for the distance as the amplitude, $A(f)$, and phase, $\phi(f)$:

$$\tilde{h}_{lm}(M_T, q, a_1, a_2, D_L) = A_{lm}(f) e^{-i\phi_{lm}(f)}, \quad (7)$$

where (l, m) are the harmonic mode indices.¹

The scaled source-frame waveform must then be transformed into the TDI observables through the frequency-domain transfer function, $\mathcal{T}(f, t_{lm}(f))$. This function is described in detail in [44,47]. It is determined by the extrinsic parameters $\{\phi_{\text{ref}}, \iota, \lambda, \beta, \psi, t_{\text{ref}}\}$ and is time and frequency dependent due to the evolution of the LISA constellation throughout its orbit. The time-frequency dependence for each harmonic, $t_{lm}(f)$, is determined from the stationary-phase approximation used to produce the frequency-domain waveforms by taking the derivative of the phase with respect to frequency:

$$t_{lm}(f) = t_{\text{ref}} - \frac{1}{2\pi} \frac{d\phi_{lm}(f)}{df}. \quad (8)$$

The LISA response depends on these extrinsic parameters given in the Solar System Barycenter (SSB) reference frame. The sampler, discussed in Sec. III, generates posterior samples in the LISA constellation reference frame. These parameters are converted to the SSB frame prior to the likelihood computation. This conversion is discussed in [44].

By combining $h_{lm}(f)$ and $\mathcal{T}(f, t_{lm}(f))$ within each harmonic, we produce the templates for each TDI channel:

$$\tilde{h}^{A,E,T}(f) = \sum_{lm} \mathcal{T}^{A,E,T}(f, t_{lm}(f)) \tilde{h}_{lm}(f). \quad (9)$$

B. Heterodyned likelihoods

For two out of the three sections of the pipeline, we take advantage of the heterodyned likelihood technique [54,55]. The specific linear formulation implemented for this work can be found in [56]. We will refer to this fast likelihood technique as “heterodyning” for the remainder of the paper to avoid confusion with other LISA-related literature, but please note the specific maths and formulation presented below are often referred to as “relative binning,” the name

¹Even though the LDC waveform only contains the $l = m = 2$ mode, we include the generic description that includes higher harmonics in preparation for the higher harmonic posterior presented in Sec. III C.

provided in [56]. We will briefly summarize the heterodyning technique and its specific implementation used here. We refer the interested reader to [55,56] for more information.

Heterodyning simplifies the frequency-domain likelihood computation by operating on residuals between templates, rather than direct comparisons of templates to data. This involves precomputing a handful of quantities representing the contribution to the likelihood from $\langle d|h \rangle$ and $\langle h|h \rangle$ terms at the full frequency resolution of the Fourier transform of the data stream.

We assume that a ‘‘reference template,’’ $\tilde{h}_0(f)$, has been found that is similar in its characteristic shape and properties to the true signal. This reference template is compared to the data at the Fourier frequencies. Online computations will involve computing new templates, $\tilde{h}(f)$, that are close in the waveform domain to the reference template. Discretizing the inner product and using the residual ratio between the two templates, $\tilde{r}(f) = \tilde{h}(f)/\tilde{h}_0(f)$, gives

$$\begin{aligned} \langle d|h \rangle &\approx Z[\tilde{d}(f), \tilde{h}(f)] \\ &= 4\text{Re} \sum_f \frac{\tilde{d}(f)\tilde{h}_0^*(f)}{S_n(f)} \tilde{r}^*(f) \Delta f. \end{aligned} \quad (10)$$

Z here represents the heterodyned approximation to the inner product value. A similar equation is given for the $\langle h|h \rangle$ term in the likelihood. In practice, $\tilde{r}(f)$ can be smoothly approximated over a large (relative to the Fourier bin width) frequency bin, b , as a complex linear segment:

$$\tilde{r}(f) = \frac{\tilde{h}(f)}{\tilde{h}_0(f)} = \tilde{r}_0(f) + \tilde{r}_1(f)(f - f_b) + \mathcal{O}(f^2), \quad (11)$$

where f_b is the frequency at the center of bin b . Equation (10) can then be recast over large frequency bins while including $\tilde{r}(f)$ up to linear order:

$$Z[\tilde{d}(f), \tilde{h}(f)] = 4\text{Re} \sum_b \sum_{f \in b} \frac{\tilde{d}(f)\tilde{h}_0^*(f)}{S_n(f)} \quad (12)$$

$$\times (\tilde{r}_0^*(f) + \tilde{r}_1^*(f)(f - f_b)) \Delta f. \quad (13)$$

Since the terms in $\tilde{r}(f)$ are constant over the bin, they can be separated from the within-bin summation:

$$Z[\tilde{d}(f), \tilde{h}(f)] = 4\text{Re} \sum_b A_0(b)\tilde{r}_{0,b}^* + A_1(b)\tilde{r}_{1,b}^*, \quad \text{with} \quad (14)$$

$$A_0(b) = \sum_{f \in b} \frac{\tilde{d}(f)\tilde{h}_0^*(f)}{S_n(f)} \Delta f \quad \text{and} \quad (15)$$

$$A_1(b) = \sum_{f \in b} \frac{\tilde{d}(f)\tilde{h}_0^*(f)}{S_n(f)} (f - f_b) \Delta f. \quad (16)$$

The complimentary equation for $\langle h|h \rangle$ is given by

$$\begin{aligned} Z[\tilde{h}(f), \tilde{h}(f)] &= 4\text{Re} \sum_b \{B_0(b)|\tilde{r}_{0,b}|^2 \\ &\quad + 2B_1(b)(\text{Re}[\tilde{r}_{0,b}\tilde{r}_{1,b}^*])\}, \quad \text{with} \end{aligned} \quad (17)$$

$$B_0(b) = \sum_{f \in b} \frac{|\tilde{h}_0(f)|^2}{S_n(f)} \Delta f \quad \text{and} \quad (18)$$

$$B_1(b) = \sum_{f \in b} \frac{|\tilde{h}_0(f)|^2}{S_n(f)} (f - f_b) \Delta f. \quad (19)$$

The tangible reason for using heterodyning is the strong reduction in the number of individual frequencies required for an accurate likelihood computation. Typical LISA data streams are of order $\sim 10^6$ – 10^7 points. Heterodyned likelihoods use ~ 100 s of points. During online computations, $\tilde{h}(f)$ is determined at the edges of each large frequency bin. The terms in $\tilde{r}(f)$ to linear order are then computed by dividing the new waveform by the reference waveform, and then determining the slope and intercept of the complex linear stretching across each frequency bin.

The main difference between the likelihood computation in this work compared to [56] is the inclusion of the LISA response. The response function is sufficiently smooth that it does not harm the ‘‘closeness’’ assumption that the residual waveform can be smoothly represented as a complex linear function over a larger frequency bin.

C. A note on the use of graphics processing units

All codes used in this work for likelihood computations, including heterodyning, waveform generation, and the LISA response function, are performed on GPUs. GPUs strongly accelerate these computations. A general description for the GPU coding architecture can be found in [45]. There have been some crucial improvements from the codes described in [45]. The main difference is that the waveform and response codes, which were formerly run in parallel on CPU cores, are now redesigned for large batches of computations on the GPUs. This is not a *major* change as the core of the code remains constant. Rather, the outer layers of the waveform and response codes were adjusted to facilitate proper usage of GPU parallelization across GPU blocks and threads. This crucial change creates the immense speed of the batched heterodyning technique. The timing of the various likelihood computations will be discussed in the next section. The MBHB waveform and likelihood code is publicly available² [71]. All other codes used in this work are available upon request to the author. They will eventually be made public.

III. AUTOMATED EXTRACTION PIPELINE

The pipeline consists of three distinct modules: an initial search phase starting across the prior volume, a refining search phase to locate the maximum likelihood point and

²You can find it here: github.com/mikekatz04/BBHx.

TABLE I. The lower and upper bounds of the prior distributions used in all modules of the pipeline presented in this work. All priors are uniform distributions, except for the prior on M_T , ι , and β . The prior for M_T is log-uniform. The prior on ι and β are uniform distributions on $\cos \iota$ and $\sin \beta$. The sky and orientation parameters are given in the LISA reference frame. T_{obs} is the observation time. For the LDC datasets, $T_{\text{obs}} \sim 1.33$ years.

Parameter	Lower bound	Upper bound
$\ln M_T$	$\ln(10^5)$	$\ln(10^8)$
q	5×10^{-2}	1
a_1	-1	1
a_2	-1	1
D_L	10^{-2}	10^3
ϕ_{ref}	0	2π
$\cos \iota$	-1	1
λ	0	2π
$\sin \beta$	-1	1
ψ	0	π
t_{ref}	0	T_{obs}

build up the high likelihood population, and a parameter estimation portion to produce the posterior. Each module passes information to the next flowing in a fully automated way from module to module.

All modules use the same priors. They are given in Table I. Uniform distributions are used on parameters $\{q, a_1, a_2, D_L, \phi_{\text{ref}}, \lambda, \psi, t_{\text{ref}}\}$. The total mass prior is log-uniform. The ι prior is uniform in $\cos \iota$, while the prior on β is uniform in $\sin \beta$.

All pipeline modules also use the same MCMC sampler. This sampler is based on the overall structure of `emcee` [72] with the parallel tempering scheme from `PTMCMC` [73]. The specific implementation used here for the parallel tempering method is given in [45]. There are two proposals used. The first is the typical “stretch” proposal described in [74], which is used for $\sim 85\%$ of proposals. The second is a sky-mode hopping proposal. As explained in [44], when sampling in the LISA reference frame, as is done in this work, there exists eight distinct sky modes in terms of the extrinsic parameters. There are four longitudinal modes at $\{\lambda + (0, 1, 2, 3)\pi/2, \psi + (0, 1, 2, 3)\pi/2\}$ and two latitudinal modes: $\{\pm\beta, \pm \cos \iota, \pm \cos \psi\}$. Three varieties of this proposal are used. The first is to move the walker to a mode drawn randomly with replacement from all eight sky modes. This is used for $\sim 4\%$ of proposals. Proposals to change just the longitudinal or latitudinal mode are used $\sim 3\%$ and $\sim 8\%$ of the time, respectively. The relative percentages of the sky-mode hopping proposals reflect that at higher frequencies, the eight modes will reduce to just the two latitudinal modes at the proper ecliptic longitude. In this situation, longitudinal mode-hopping proposals are not efficient.

A. Module 1: direct full-template search

Harnessing the power of GPUs allows for a fast and direct likelihood computation against the data as described

in [45]. This computation takes \sim few milliseconds. We use 80 walkers per temperature with four temperatures ranging from one (target distribution) to infinity (prior). Only four temperatures are used (compared to 10 for the later segments) due to the longer likelihood computation time as compared to the heterodyning methods used below. This setup allows the chains to take more steps in the same amount of wall time, which was found to be advantageous for the initial search. The starting positions of all walkers are generated randomly from the prior distribution. These positions are then evolved with the MCMC techniques given above. This procedure is used to obtain samples from the posterior, which will be preferentially located towards higher likelihood values. The goal of this portion of the pipeline is to locate the proper frequency band where the signal lies. In other words, this means roughly locating the true value on the total mass. This is crucial to the success of the heterodyned likelihood since it can give erroneous results at frequencies that are not mutually included in both the reference template (h_0) and tested template (h).

This portion of the pipeline takes \sim minutes or less to complete. The sampler is outfit with a “stopping” function that performs a test every 150 iterations. The test is based on the extracted SNR (see Sec. II). When it reaches a user-specified value, the sampling ends and passes the data file information to the next pipeline module. The evolution of the sampler in this module is shown in Fig. 2 in terms of the likelihood and the extracted SNR over time. In both the noisy and noiseless cases, the sampler climbs the likelihood surface quickly and consistently. This figure also includes a rough visualization of the different durations of this module for two different extraction SNR convergence limits (20 and 200). In practice, a lower extracted SNR threshold (~ 10 – 20) will be used to terminate the first module in the pipeline given that the SNR of the signal in real data will be unknown. However, a higher threshold could be used to further refine parameters with this direct method if desired, as long as the empirically determined extracted SNR continues to rise. This may be desirable if the heterodyned likelihood has issues with the determined frequency band over which the calculation is performed. Further refining the parameters may give a more accurate frequency band depending on the specific source. In the case examined here, the lower extracted SNR threshold (20) is successful in finding a proper frequency band over which the heterodyning computation is performed.

An advantage of this method is that it is completely general (given noise and detector assumptions) to frequency domain waveforms. Therefore, it can be easily extended to include precession, eccentricity, etc. This module could be run all the way through posterior production if desired, but this is not as efficient (by a factor of ~ 20 – 100) as the chosen fast likelihood methods described in the following sections.

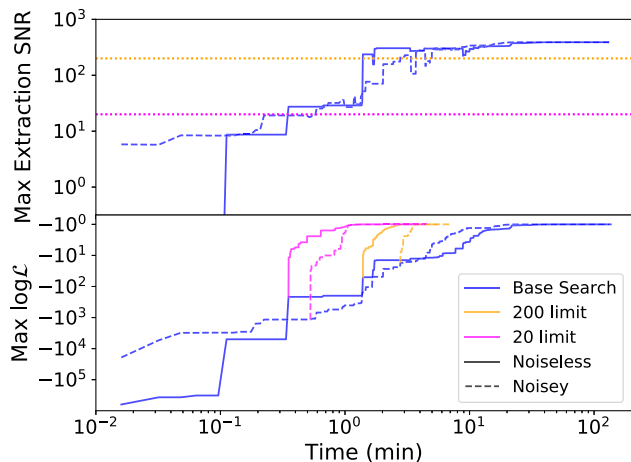


FIG. 2. The evolution of the maximum log-likelihood and extraction SNR over time as the base search module proceeds. The base search runs for the noiseless and noisy cases are shown in blue solid and dashed lines, respectively. The noisy log-likelihood is normalized so that the maximum log-likelihood found by the sampler is set to zero, equivalent to the noiseless maximum log-likelihood. The base search module culminates when a user-specified extraction SNR is found by the sampler. Two extraction SNR limits are shown here: 20 (magenta) and 200 (orange). In the top panel, these limits are marked with dotted lines to indicate the turnover point: when the blue line reaches these horizontal lines for the first time, it will end sampling in the base search module and move on to the second module. In the bottom panel, the evolution of the log-likelihood during the second module is also shown over time for both extraction SNR limits. These lines begin at the moment when the initial search reaches the desired extraction SNR and turns over into the second module; they end at the time when the pipeline has converged on a maximum likelihood point.

B. Module 2: heterodyned search

After reaching the extracted SNR threshold, the highest likelihood point is used as a reference template for heterodyning. The frequency content contained in the reference template over the observation time, T_{obs} , is set as the frequency band over which the heterodyning method is applied [$\tilde{h}_0(f) \neq 0$]. The frequency grid is chosen to have 128 log-spaced frequencies. With these settings, each likelihood computation takes $\sim 5 \mu\text{s}$. The reference template is updated every 5000 iterations of the sampler by replacing it with the highest likelihood point in the population.

This segment of the pipeline uses 400 walkers in 10 temperatures (4000 total), taking advantage of the immense power of the low-memory, batched heterodyned likelihood computations.

Each walker is assigned an initial position from the top 4000 likelihood points resulting from module 1 and is evolved forward according to the same MCMC method as module 1. The stopping function for this section of the pipeline is determined by the convergence of the maximum

likelihood located by the sampler. Every 150 iterations of the sampler, the stopping function checks the maximum likelihood point. If this point does not change its value by 1 during multiple consecutive checks, then the sampling is stopped and the data is saved and passed on to the parameter estimation portion in the final module. The number of consecutive checks necessary to stop is a user-defined parameter that is set in this work to 30. This is conservative allowing for 30×150 iterations = 4.5×10^3 iterations of the sampler after finding the initial convergent likelihood value. Typically, this segment is \sim minutes, depending on the exact stopping criterion settings. The evolution of the log-likelihood for module 2 with different extraction SNR limits from module 1 is shown in Fig. 2 as the magenta (20 limit) and orange (200 limit) lines in the bottom panel. As the tracks flatten out near $\log \mathcal{L} \approx 0$, the maximum likelihood begins to converge. The length of this flat portion gives the amount of time required for the convergence criterion to be met (\sim minutes). Ultimately, the use of heterodyning in module 2 improves the speed of the search by a factor of ~ 20 – 100 . The higher end of that range indicates how long it takes the base search in module 1 to satisfy the log-likelihood convergence criterion of module 2.

C. Module 3: heterodyned parameter estimation

The final module generates the posterior starting from information around the maximum likelihood point obtained during module 2. The exact same procedure is used from module 2 in terms of the heterodyned likelihood settings, timing, and updates. A burn-in of 1000 iterations is imposed. The sampler then collects posterior samples for 5×10^4 iterations. With 4000 walkers, this produces 4×10^8 posterior evaluations with 4×10^7 evaluations along the cold chain probing the target distribution.

The posterior samples for the noise-infused and noiseless cases are shown in Figs. 3 and 4, respectively, for the intrinsic parameters. Figures 5 and 6 display the posterior distribution on the extrinsic parameters for the noiseless and noisy cases, respectively. The mean and standard deviation for unimodal parameters in all tests are given in Table II. The parameters of the maximum likelihood point attained from the sampler during each MCMC run are shown in Table III. The posterior information for the waveform with higher-order modes is also added to the tables and figures. Extrinsic parameters in these figures and tables are given in the SSB frame.

The injected signal has a strong SNR of ~ 390 . Therefore, the constraints on intrinsic parameters are quite strong, even for the case where only the (2,2) mode is included. The sky location is fairly constrained and exhibits the character expected. This signal reaches high enough frequencies that the base eightfold sky degeneracy in the LISA reference frame is reduced to just the latitudinal modes at the true longitude. The two-dimensional sky




FIG. 3. The posterior distribution on the intrinsic parameters of the source in the noiseless LDC dataset. The blue contours represent the main posterior distribution. The red distribution is shown to give an idea of more realistic posteriors that include higher-order harmonic modes. The true injection values are marked in black. The posterior distributions contain the 1σ , 2σ , and 3σ contours.

posterior shows the transformation of this characteristic feature to the SSB frame. The two sky modes also lead to multimodal distributions in the inclination and reference time (in the SSB frame). The most likely value for t_{ref} found by the sampler is approximately 30 seconds before the value given in the LDC dataset. This is due to a differing convention in the reference frequency at which t_{ref} is assigned. In the LDC, t_{ref} is set at the highest frequency in the waveform. As described in Sec. II A, we set t_{ref} at f_{max} , where $f^2 A_{22}(f)$ reaches a maximum. The difference in t_{ref} between these two conventions for this specific source is ~ 30 seconds, which confirms we have recovered the correct value. The distance posterior in both the noise and noiseless cases are typical for sources with the

(2,2)-mode-only waveforms. The true value is closer to the observer than the mean and mode of the distribution.

The effect of the noise is also visible. This particular noise realization produces only a small shift in the noiseless distribution. The parameter posterior from the noisy data is qualitatively and quantitatively similar to the noiseless distribution in all parameters.

It is clear, from examining the distributions over time as more samples are collected, that intrinsic parameters converge rapidly while the sky and orientation parameters converge more slowly. A posterior distribution in the extrinsic parameters roughly equivalent to the final distribution is found as fast as the intrinsic parameters are located; however, it does take significantly more time for




FIG. 4. The posterior distribution on the intrinsic parameters from the noisy LDC dataset represented with blue contours. The true injection values are marked in black. The posterior distribution contains the 1σ , 2σ , and 3σ contours.

the sampler to fully converge the tails in the extrinsic parameter distributions.

In order to better inform the reader of expected and more realistic parameter constraints and posterior distributions, we provide a posterior that shows the effect of analyzing higher-order harmonic modes for the same source that is used for the LDC-1A dataset. To obtain the higher-mode distributions, we injected a PHENOMHM signal with the same parameters as the LDC source and then used the brute-force likelihood approach from Sec. III A subbing in PHENOMHM for PHENOMD. The sampling begins with spreading the walkers around the true injection point.

The higher-mode distribution is better in constraining parameters by about an order of magnitude. The higher harmonics also strongly inform the marginalized distance posterior to the point where it is now Gaussian and

spanning the true value as its mean. This is due to the broken degeneracy between inclination and distance that exists in the (2,2)-only case. This broken degeneracy also improves sampling efficiency. See [47] for more detail on comparing sampling effects and posteriors from (2,2)-mode-only waveforms versus waveforms with higher modes.

IV. DISCUSSION

This pipeline is designed for generic, quasicircular frequency domain waveforms. It will not require changes to handle precessing and/or eccentric binaries because it does not rely on any type of maximization over parameters or F-statistic-like computations. However, these types of maximization techniques could be added for further




FIG. 5. The posterior distribution, shown as blue contours, on the extrinsic parameters of the source in the noiseless LDC dataset. The sky and orientation parameters given here are in the SSB frame. The red distribution represents the posterior from the PHENOMHM injection and template fitting. The true injection values are marked in black. The posterior distributions contain the 1σ , 2σ , and 3σ contours.

efficiency during the initial base search. This type of change should not affect the heterodyned likelihood sections of the pipeline because the maximization techniques will maintain a strong determination of the total mass. This is essential to the success of the heterodyned likelihood modules because heterodyning requires an accurate estimation of the frequency band containing the waveform.

While the pipeline focuses on finding signals that have already merged, it can, in principle, work for signals prior to merger. It does not assume the detector is fixed in time, so it can handle searching over longer stretches of waveform. Specific studies of detection of these sources prior to merger will be needed to more fully understand its capabilities in the premerger realm. For premerger




FIG. 6. The posterior distribution on the extrinsic parameters from the noisy LDC dataset represented with blue contours. The sky and orientation parameters are shown in the SSB reference frame. The true injection values are marked in black. The posterior distribution contains the 1σ , 2σ , and 3σ contours.

detection, we expect the full base template search to work. However, its relation to and success of the heterodyned likelihood modules is uncertain due to the smaller frequency range over which premerger sources will evolve. In principle, this should be okay, but it requires verification.

When searching for signals after the merger-ringdown, higher-mode waveforms may ensure faster convergence to the true values due to the higher degree of Gaussianity in

the higher-mode posteriors. With that said, the peaks in the posterior are thinner, so locating this initial peak may be more difficult than with (2,2)-mode-only waveforms.

Fast searches are necessary for early detection in low-latency pipelines in order to maximize the potential for simultaneous observation through electromagnetic counterparts. The earlier an MBHB source is detected, the higher the likelihood for coincident detection. The pipeline

TABLE II. The mean and standard deviation values associated with the various tests performed. These values are only shown for the mass, spin, and distance parameters because these marginalized distributions are unimodal. The sky and orientation parameters are all multimodal indicating the observed mean value is not representative of the true source. From left to right, the values stated apply to the LDC dataset without noise, the LDC dataset with noise, and the self-injected, noiseless higher modes dataset analyzed with PHENOMHM.

Parameter	Injection	(2,2) Noiseless	(2,2) Noise	HM Noiseless
m_1	2.599137×10^6	$2.60 \times 10^{6+3 \times 10^4}$ -3×10^4	$2.60 \times 10^{6+3 \times 10^4}$ -3×10^4	$2.599 \times 10^{6+5 \times 10^3}$ -5×10^3
m_2	1.242861×10^6	$1.24 \times 10^{6+1 \times 10^4}$ -1×10^4	$1.24 \times 10^{6+1 \times 10^4}$ -1×10^4	$1.243 \times 10^{6+3 \times 10^3}$ -3×10^3
a_1	7.534822×10^{-1}	$7.6 \times 10^{-1+3 \times 10^{-2}}$ -2×10^{-2}	$7.5 \times 10^{-1+3 \times 10^{-2}}$ -2×10^{-2}	$7.53 \times 10^{-1+3 \times 10^{-3}}$ -2×10^{-3}
a_2	6.215875×10^{-1}	$6.0 \times 10^{-1+7 \times 10^{-2}}$ -7×10^{-2}	$6.2 \times 10^{-1+7 \times 10^{-2}}$ -7×10^{-2}	$6.22 \times 10^{-1+3 \times 10^{-3}}$ -3×10^{-3}
D_L	5.600578×10^1	$1.1 \times 10^{2+3 \times 10^1}$ -4×10^1	$8.76 \times 10^{1+3 \times 10^1}$ -3×10^1	$5.6 \times 10^{1+3 \times 10^0}$ -3×10^0

TABLE III. Parameters at the maximum likelihood values found by the sampler for the three different tests are shown. These values are given for the LDC dataset without noise, the LDC dataset with noise, and the self-injected higher modes dataset, from left to right, respectively. Both runs without noise located a maximum log-likelihood ≥ -0.2 . If the $\langle d|d \rangle$ term is removed from the likelihood (i.e., $\langle d|d \rangle = 0$), then the maximum log-likelihood located by the sampler for the run containing noise was $\sim +76005$.

Parameter	Injection	(2,2) Noiseless	(2,2) Noise	HM Noiseless
m_1	2.599137×10^6	2.615348×10^6	2.597597×10^6	2.599091×10^6
m_2	1.242861×10^6	1.235689×10^6	1.243229×10^6	1.243054×10^6
a_1	7.534822×10^{-1}	7.721731×10^{-1}	7.512426×10^{-1}	7.539068×10^{-1}
a_2	6.215875×10^{-1}	5.716351×10^{-1}	6.256857×10^{-1}	6.213842×10^{-1}
D_L	5.600578×10^1	2.859041×10^1	6.923802×10^1	5.620339×10^1
ι	1.224532×10^0	1.403617×10^0	1.169769×10^0	1.223004×10^0
λ	3.509097×10^0	3.519460×10^0	3.487416×10^0	3.509118×10^0
β	2.927012×10^{-1}	2.655787×10^{-1}	2.870121×10^{-1}	2.931597×10^{-1}
t_{ref}	2.495997×10^7	2.495996×10^7	2.495997×10^7	2.495997×10^7

presented in this work takes \sim minutes to locate the maximum likelihood point associated with the source. This gives an initial point estimate of the sky location. The parameter estimation pipeline can then be run for as long as is necessary to refine the sky-location distribution. As discussed in Sec. III C, a roughly accurate map of the extrinsic parameters is attained very quickly during sampling. Therefore, in ~ 10 minutes, the pipeline could generate posteriors on the sky parameters that are roughly equivalent to the fully converged distributions. A different and optimized method for generating fast sky maps is discussed in [52].

V. CONCLUSION

We presented a solution to the LDC-1A dataset containing a single MBHB signal. The solution involves a fully automated pipeline that starts with randomized initial points throughout the prior volume. A brute-force likelihood approach takes advantage of GPU acceleration to make the initial detection of a signal within the data. This portion of the pipeline is flexible and should expand well to signals prior to merger. The brute-force likelihood can also be directly used for the rest of the pipeline to verify the

heterodyning results or if the heterodyning technique proves unreliable in any specific situation. The second stage of the pipeline uses a GPU-accelerated batched computation that employs the heterodyned likelihood technique. This method is extremely efficient at locating the maximum likelihood point. Once converged, the final step takes over, running again with the heterodyning technique to quickly build a final posterior distribution. We provided final posterior distributions for both the noiseless and noise-infused datasets provided by the LDC. Additionally, a posterior was presented for the inclusion of higher-order modes to give the reader a better sense of the true constraints we will be able to place on the signal parameters. For this higher-modes posterior, we injected our own signal with the same parameters as the source from the LDC dataset.

The noisy LDC dataset analyzed in this work was provided with stationary and Gaussian noise properties. The LISA mission will exhibit nonstationary and non-Gaussian noise effects such as from glitches, data gaps, the confusion foreground from Galactic sources, and the generally expected drift of the instrument sensitivity over the observation window. Future datasets produced and provided by the LDC will address these difficult problems. The pipeline designed in this work heavily leverages

frequency-domain information. When dealing with the nonstationary information in the future, this pipeline will most likely require a switch to the time-frequency or wavelet domain [75].

Another realistic complication to be added is the occurrence of multiple simultaneous signals. The LDC has provided a new dataset (LDC-2A) that combines the millions of (mostly unobservable) Galactic binaries with ~ 15 MBHB sources while maintaining Gaussian and stationary noise assumptions. This dataset is the next logical test for the pipeline designed here. We expect this pipeline to succeed in the search for sources in this new dataset. With the merger ringdown, where most of the SNR is accumulated, occurring on a timescale of minutes to hours, we expect the ~ 15 MBHB sources in the data to be sufficiently separated so as to easily distinguish individual sources. The posterior distributions from individually performed parameter estimation runs may be biased in this more-complicated dataset by the presence of other MBHB signals and/or Galactic binary sources. To avoid these biases, a global fit will have to be performed

(e.g., [7]). The posterior samples attained from our individual source pipeline may be useful for providing the global fit with direct information on the MBHBs, potentially vastly decreasing the time to convergence of the MBHB portion of the global fit algorithm.

While these more realistic settings require further development of this efficient and automated pipeline, the initial implementation shows promising results indicating we have taken a strong first step towards achieving these long-term LISA analysis objectives.

ACKNOWLEDGMENTS

M. L. K. would like to thank Sylvain Marsat, Jonathan Gair, Alvin Chua, Stanislas Babak, Lorenzo Speri, and Ollie Burke for helpful discussions towards the success of this work. This research was supported in part through the computational resources and staff contributions provided for the Quest/Grail high performance computing facility at Northwestern University. This paper also employed use of SCIPY [76], NUMPY [77], and MATPLOTLIB [78].

-
- [1] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. X* **9**, 031040 (2019).
 - [2] LVK Collaboration, *Living Rev. Relativity* **21**, 3 (2018).
 - [3] P. Amaro-Seoane *et al.*, [arXiv:1702.00786](https://arxiv.org/abs/1702.00786).
 - [4] M. C. Begelman, R. D. Blandford, and M. J. Rees, *Nature (London)* **287**, 307 (1980).
 - [5] A. Sesana, *Phys. Rev. Lett.* **116**, 231102 (2016).
 - [6] S. Babak, J. Gair, A. Sesana, E. Barausse, C. F. Sopuerta, C. P. L. Berry, E. Berti, P. Amaro-Seoane, A. Petiteau, and A. Klein, *Phys. Rev. D* **95**, 103012 (2017).
 - [7] T. Littenberg, N. Cornish, K. Lackeos, and T. Robson, *Phys. Rev. D* **101**, 123021 (2020).
 - [8] B. F. Schutz, *Nature (London)* **323**, 310 (1986).
 - [9] D. E. Holz and S. A. Hughes, *Astrophys. J.* **629**, 15 (2005).
 - [10] A. Petiteau, S. Babak, and A. Sesana, *Astrophys. J.* **732**, 82 (2011).
 - [11] S. Burke-Spolaor, *Classical Quant. Grav.* **30**, 224013 (2013).
 - [12] J. R. Gair, M. Vallisneri, S. L. Larson, and J. G. Baker, *Living Rev. Relativity* **16**, 7 (2013).
 - [13] E. Barausse, J. Bellovary, E. Berti, B. Holley-Bockelmann, K. Farris, B. Sathyaprakash, and A. Sesana, *J. Phys. Conf. Ser.* **610**, 012001 (2015).
 - [14] T. Bogdanović, *Astrophys. Space Sci. Proc.* **40**, 103 (2015).
 - [15] A. Klein *et al.*, *Phys. Rev. D* **93**, 024003 (2016).
 - [16] E. Berti, A. Sesana, E. Barausse, V. Cardoso, and K. Belczynski, *Phys. Rev. Lett.* **117**, 101102 (2016).
 - [17] J. Salcido, R. Bower, T. Theuns, S. McAlpine, M. Schaller, R. Crain, J. Schaye, and J. Regan, *Mon. Not. R. Astron. Soc.* **463**, 870 (2016).
 - [18] M. L. Katz, L. Z. Kelley, F. Dosopoulou, S. Berry, L. Blecha, and S. L. Larson, *Mon. Not. R. Astron. Soc.* **491**, 2301 (2020).
 - [19] M. Bonetti, A. Sesana, F. Haardt, E. Barausse, and M. Colpi, *Mon. Not. R. Astron. Soc.* **486**, 4044 (2019).
 - [20] C. Cutler and É. Flanagan, *Phys. Rev. D* **49**, 2658 (1994).
 - [21] A. Vecchio, *Phys. Rev. D* **70**, 042001 (2004).
 - [22] E. Berti, A. Buonanno, and C. M. Will, *Phys. Rev. D* **71**, 084025 (2005).
 - [23] K. Arun, *Phys. Rev. D* **74**, 024025 (2006).
 - [24] R. N. Lang and S. A. Hughes, *Phys. Rev. D* **74**, 122001 (2006).
 - [25] J. Thorpe, S. McWilliams, B. Kelly, R. Fahey, K. Arnaud, and J. Baker, *Classical Quant. Grav.* **26**, 094026 (2009).
 - [26] S. T. McWilliams, J. I. Thorpe, J. G. Baker, and B. J. Kelly, *Phys. Rev. D* **81**, 064014 (2010).
 - [27] S. T. McWilliams, B. J. Kelly, and J. G. Baker, *Phys. Rev. D* **82**, 024014 (2010).
 - [28] S. McWilliams, R. Lang, J. Baker, and J. Thorpe, *Phys. Rev. D* **84**, 064003 (2011).
 - [29] K. Arun, B. R. Iyer, B. Sathyaprakash, S. Sinha, and C. van den Broeck, *Phys. Rev. D* **76**, 104016 (2007).
 - [30] M. Trias and A. M. Sintes, *Phys. Rev. D* **77**, 024030 (2008).
 - [31] E. K. Porter and N. J. Cornish, *Phys. Rev. D* **78**, 064005 (2008).
 - [32] V. Baibhav, E. Berti, and V. Cardoso, *Phys. Rev. D* **101**, 084053 (2020).
 - [33] D. A. Brown, J. Crowder, C. Cutler, I. Mandel, and M. Vallisneri, *Classical Quant. Grav.* **24**, S595 (2007).
 - [34] N. J. Cornish and E. K. Porter, *Classical Quant. Grav.* **23**, S761 (2006).

- [35] J. Crowder, N. J. Cornish, and J. L. Reddinger, *Phys. Rev. D* **73**, 063011 (2006).
- [36] E. Wickham, A. Stroeer, and A. Vecchio, *Classical Quant. Grav.* **23**, S819 (2006).
- [37] C. Röver, A. Stroeer, E. Bloomer, N. Christensen, J. Clark, M. Hendry, C. Messenger, R. Meyer, M. Pitkin, J. Toher, R. Umstätter, A. Vecchio, J. Veitch, and G. Woan, *Classical Quant. Grav.* **24**, S521 (2007).
- [38] F. Feroz, J. R. Gair, M. P. Hobson, and E. K. Porter, *Classical Quant. Grav.* **26**, 215003 (2009).
- [39] J. R. Gair and E. K. Porter, *Classical Quant. Grav.* **26**, 225004 (2009).
- [40] A. Petiteau, Y. Shang, and S. Babak, *Classical Quant. Grav.* **26**, 204011 (2009).
- [41] S. Babak *et al.*, *Classical Quant. Grav.* **27**, 084009 (2010).
- [42] E. K. Porter and J. Carré, *Classical Quant. Grav.* **31**, 145004 (2014).
- [43] E. K. Porter and N. J. Cornish, *Phys. Rev. D* **91**, 104001 (2015).
- [44] S. Marsat, J. G. Baker, and T. Dal Canton, *Phys. Rev. D* **103**, 083011 (2021).
- [45] M. L. Katz, S. Marsat, A. J. K. Chua, S. Babak, and S. L. Larson, *Phys. Rev. D* **102**, 023033 (2020).
- [46] LISA Science Study Team, LISA Science Requirements Document (2018).
- [47] S. Marsat and J. G. Baker, [arXiv:1806.10734](https://arxiv.org/abs/1806.10734).
- [48] K. Dey, N. Karnesis, A. Toubiana, E. Barausse, N. Korsakova, Q. Baghi, and S. Basak, *Phys. Rev. D* **104**, 044035 (2021).
- [49] W.-H. Ruan, C. Liu, Z.-K. Guo, Y.-L. Wu, and R.-G. Cai, *Research* **2021**, 6014164 (2021).
- [50] K. J. Shuman and N. J. Cornish, [arXiv:2105.02943](https://arxiv.org/abs/2105.02943).
- [51] C. Zhang, Y. Gong, and C. Zhang, *Phys. Rev. D* **104**, 083038 (2021).
- [52] N. J. Cornish and K. Shuman, *Phys. Rev. D* **101**, 124008 (2020).
- [53] N. J. Cornish, [arXiv:2110.06238](https://arxiv.org/abs/2110.06238) [*Phys. Rev. D* (to be published)].
- [54] N. J. Cornish, [arXiv:1007.4820](https://arxiv.org/abs/1007.4820).
- [55] N. J. Cornish, *Phys. Rev. D* **104**, 104054 (2021).
- [56] B. Zackay, L. Dai, and T. Venumadhav, [arXiv:1806.08792](https://arxiv.org/abs/1806.08792).
- [57] L. Finn and K. Thorne, *Phys. Rev. D* **62**, 124021 (2000).
- [58] C. J. Moore, R. H. Cole, and C. P. L. Berry, *Classical Quant. Grav.* **32**, 015014 (2015).
- [59] S. Husa, S. Khan, M. Hannam, M. Pürrer, F. Ohme, X. Forteza, and A. Bohé, *Phys. Rev. D* **93**, 044006 (2016).
- [60] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. Forteza, and A. Bohé, *Phys. Rev. D* **93**, 044007 (2016).
- [61] L. London, S. Khan, E. Fauchon-Jones, C. García, M. Hannam, S. Husa, X. Jiménez-Forteza, C. Kalaghatgi, F. Ohme, and F. Pannarale, *Phys. Rev. Lett.* **120**, 161102 (2018).
- [62] P. Maturana-Russel, R. Meyer, J. Veitch, and N. Christensen, *Phys. Rev. D* **99**, 084006 (2019).
- [63] P. M. Goggans and Y. Chi, *AIP Conf. Proc.* **707**, 59 (2004).
- [64] N. Lartillot and H. Philippe, *Syst. Biol.* **55**, 195 (2006).
- [65] M. Tinto and J. W. Armstrong, *Phys. Rev. D* **59**, 102003 (1999).
- [66] J. Armstrong, F. Estabrook, and M. Tinto, *Astrophys. J.* **527**, 814 (1999).
- [67] F. B. Estabrook, M. Tinto, and J. W. Armstrong, *Phys. Rev. D* **62**, 042002 (2000).
- [68] S. V. Dhurandhar, K. R. Nayak, and J.-Y. Vinet, *Phys. Rev. D* **65**, 102002 (2002).
- [69] M. Tinto and S. V. Dhurandhar, *Living Rev. Relativity* **8**, 4 (2005).
- [70] M. Vallisneri, *Phys. Rev. D* **71**, 022001 (2005).
- [71] M. Katz, [mikekatz04/bbhx](https://arxiv.org/abs/2105.02943): Second release (2021).
- [72] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, *Publ. Astron. Soc. Pac.* **125**, 306 (2013).
- [73] W. Vousden, W. Farr, and I. Mandel, *Mon. Not. R. Astron. Soc.* **455**, 1919 (2016).
- [74] J. Goodman and J. Weare, *Commun. Appl. Math. Comput. Sci.* **5**, 65 (2010).
- [75] N. J. Cornish, *Phys. Rev. D* **102**, 124038 (2020).
- [76] P. Virtanen *et al.*, *Nat. Methods* **17**, 261 (2020).
- [77] S. van der Walt, S. C. Colbert, and G. Varoquaux, *Comput. Sci. Eng.* **13**, 22 (2011).
- [78] J. D. Hunter, *Comput. Sci. Eng.* **9**, 90 (2007).