

Efficient GPU-accelerated multisource global fit pipeline for LISA data analysis

Michael L. Katz^{1,2,*}, Nikolaos Karnesis,³ Natalia Korsakova⁴, Jonathan R. Gair², and Nikolaos Stergioulas³

¹*NASA Marshall Space Flight Center, Huntsville, Alabama 35811, USA*

²*Max-Planck-Institut für Gravitationsphysik, Albert-Einstein-Institut,
Am Mühlenberg 1, 14476 Potsdam-Golm, Germany*

³*Department of Physics, Aristotle University of Thessaloniki, Thessaloniki 54124, Greece*

⁴*Astroparticule et Cosmologie, CNRS, Université Paris Cité, F-75013 Paris, France*



(Received 7 May 2024; accepted 10 December 2024; published 24 January 2025)

The large-scale analysis task of deciphering gravitational-wave signals in the LISA data stream will be difficult, requiring a large amount of computational resources and extensive development of computational methods. Its high dimensionality, multiple model types, and complicated noise profile require a global fit to all parameters and input models simultaneously. In this work, we detail our global fit algorithm, called “Erebor,” designed to accomplish this challenging task. It is capable of analyzing current state-of-the-art datasets and then growing into the future as more pieces of the pipeline are completed and added. We describe our pipeline strategy, the algorithmic setup, and the results from our analysis of the LDC2A Sangria dataset, which contains massive black hole binaries, compact galactic binaries, and a parametrized noise spectrum whose parameters are unknown to the user. The Erebor algorithm includes three unique and very useful contributions: GPU acceleration for enhanced computational efficiency; ensemble Markov Chain Monte Carlo (MCMC) sampling with multiple MCMC walkers per temperature for better mixing and parallelized sample creation; and special online updates to reversible-jump (or transdimensional) sampling distributions to ensure sampler mixing and accurate initial estimates for detectable sources in the data. We recover posterior distributions for all 15 (6) of the injected massive black hole binaries (MBHB) in the LDC2A training (hidden) dataset. We catalog ~ 12000 galactic binaries (~ 8000 as high confidence detections) for both the training and hidden datasets. All of the sources and their posterior distributions are provided in publicly available catalogs.

DOI: [10.1103/PhysRevD.111.024060](https://doi.org/10.1103/PhysRevD.111.024060)

I. INTRODUCTION

In the mid-2030s, the Laser Interferometer Space Antenna (LISA) will launch into space to measure gravitational waves emanating from a variety of astrophysical and cosmological sources in the millihertz frequency band [1,2]. LISA will add important information to the gravitational-wave spectrum, following on the discoveries of ground-based observing networks at higher frequencies [3] and pulsar timing arrays at lower frequencies [4–7]. LISA will detect different astrophysical sources: compact object binaries, also referred to as galactic binaries (GB), inside and close to the Milky Way Galaxy in slowly evolving orbits [e.g. [1,8–11]] massive black hole binaries

(MBHB) in the centers of galaxies [e.g. [8,12–15]]; and extreme mass ratio inspirals where a smaller compact object orbits around a massive black hole [e.g. [8,16]]. LISA will also be able to detect astrophysical and cosmological stochastic signals of sufficient amplitude [e.g. [1]]. All of the signals will exist simultaneously in the LISA data stream with MBHBs providing quasitransient signals over days to weeks; GB signals existing for the duration of the mission with almost monochromatic signals; and extreme-mass-ratio inspirals (EMRIs) somewhere in between with long-duration (\sim year) signals that will merge after a long time inspiralling.

In addition to this astrophysical variety, LISA will have a complicated noise profile with nonstationary, non-Gaussian noise effects such as sensitivity drift, gaps, and glitches [17–21]. Additionally, the orbit and performance of the LISA experiment will need to be included in any global analysis [22–25]. There will also be an initial data reduction pipeline that must be run properly prior to performing the scientific data analysis. This step includes instrument calibration [26] and reduction of primary noises such as laser noise, clock noise, and tilt-to-length noise [27–29]. The techniques of time delay interferometry (TDI)

*Contact author: michael.l.katz@nasa.gov

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.

[e.g. [30,31] and sources within], global time synchronization, and ranging [32] will be used to reduce these crucial noise contributions.

This complicated analysis requires a global fit over all parameters characterizing the models for the signals and the noise. The design of this type of pipeline has been, and will continue to be, a large developmental project requiring many participating groups and expertise. LISA global fit pipelines are likely to be built on the concept of the global fit “wheel”: separate modules specifically designed for different source or instrumental analyses will run in parallel and communicate with each other throughout the global fit run [21]. An example “wheel” illustration is shown in Fig. 1. This takes advantage of the small correlations expected between different source classes and instrumental models to perform the complex analysis in a much more computationally friendly manner.

In order to benchmark and test the implementation of these algorithms across the LISA community, the LISA data challenges working group has released a series of datasets of simulated LISA data. The LDC2A dataset, also

called “Sangria” [33], was formulated to test the current state of the art in the LISA global fitting community: simultaneous fitting of MBHBs, GBs throughout the Galaxy, and a stationary, Gaussian and overall simplified noise profile. It contains one year of data with the aforementioned signals and noise combined into a single dataset. It also comes with information related to the exact input populations to be used for debugging and developing experimental pipelines [34].

There are a few groups working on this large-scale problem. In [21,35,36], the authors describe a pipeline that can successfully analyze the LDC2A dataset and produce output posterior distributions for the detected sources and noise properties. They also provide details of their methods and findings and discuss future projects related to their global fit pipeline.

Recently, [37] demonstrated a highly efficient maximum likelihood solution for GBs and MBHBs, which also included a fit to the overall noise profile including the confusion foreground signal generated by the many unresolved GBs. They performed this analysis for one year of data, updating week by week, thus showing its utility and capabilities to handle time evolution. This useful pipeline will provide a low-cost initial estimate of source parameters, populations and noise properties. The information contained in this estimate can also be used to further improve the efficiency of full posterior estimation-based global fit samplers.

In this work, we present our implementation of a LISA global fit pipeline for analyzing the sources and noise profile in the LDC2A dataset. We perform this full computation on the “training” data and, for the first time, on the “hidden” dataset provided by the LDC. Our goal was to learn valuable lessons from [21], while trying to build an independent algorithm with important differences in order to further expand the LISA global fit knowledge space.

Our pipeline follows a similar Bayesian approach as presented in [21]. We use reversible-jump Markov chain Monte Carlo (RJMCMC; also known as transdimensional Markov chain Monte Carlo) techniques [38] to fit for and characterize the uncertain number of GB sources in the data. We also adapt the concept of the global fit “wheel” (Fig. 1) running large modules of the sampler in a blocked Gibbs fashion [39–41], sending out and receiving updates to and from all other modules periodically [21].

There are three important large-scale differences between our pipeline and [21]. The pipeline in [21] used a large number of CPUs run in parallel. In contrast, we designed our algorithm specifically to maximize the efficiency and usage of graphics processing units (GPU). GPUs are specialized hardware that are designed for a high level of parallel processing. GPUs are generally more cost and energy efficient than CPU clusters with equivalent computational power and are expected to continue to improve at a more rapid rate than CPUs [42]. For these

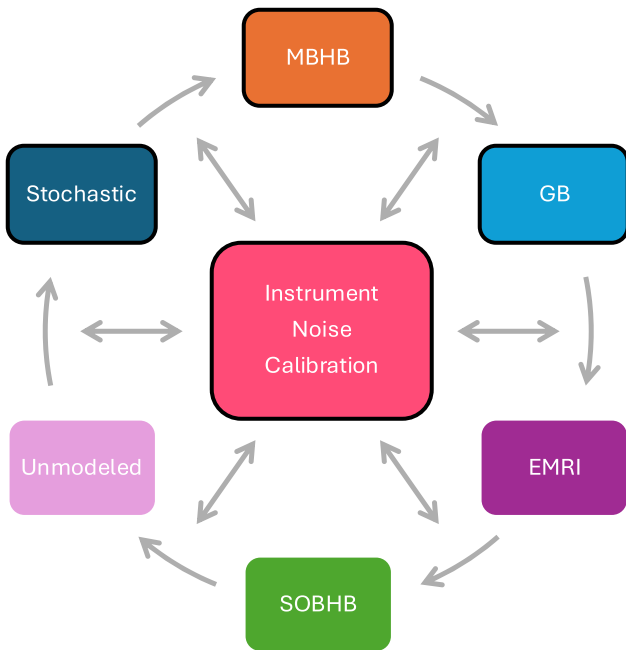


FIG. 1. The “Global Fit Wheel” representing an example flow for a global fit algorithm across different source and instrument subanalyses for LISA [35]. In clockwise order: massive black hole binaries (MBHB), galactic binaries (GB), extreme-mass-ratio inspirals (EMRI), stellar-origin black hole binaries (SOBHB), unmodeled gravitational-wave sources (e.g. bursts), and stochastic signals like astrophysical and cosmological backgrounds. The instrument, noise, and calibration will enter the process throughout to ensure updated experimental information. Boxes outlined in black show which components are currently included in their *initial* implementation state in Erebor. However, as of now we do not include any instrument or calibration analyses. That is left to future work.

reasons, we designed our algorithm for GPUs in order to take advantage of the computational and energy efficiency, which we expect to further improve over the coming years.

The second main difference is we use an ensemble sampling setup for our Markov chain Monte Carlo runs [43]. Rather than having one walker per temperature as used in [21], we use many walkers per temperature by leveraging the MCMC software package Eryn [43–45]. A “walker” is one point in parameter space representing a current state of the sampler. As walkers evolve due to probabilistic distributions, each state they evolve through forms the chain of samples that comes together to make the posterior distribution. In addition to the intrinsic benefits of ensemble sampling, within the LISA global fit context it allows us to marginalize over the various model types more efficiently: walkers can contain different residuals or noise estimates that come from the random differences during sampling of other parts of the pipeline. Ensemble sampling is also better for parallelization of MCMC sample creation: there is more than one walker in the cold chain; each walker can be evaluated independently in parallel across the GPU. We also employ the “group stretch” MCMC proposal, which requires an ensemble setup to use and minimal tuning compared to typical covariance proposals [44].

Finally, using one or more extra GPUs, we perform online search and refitting techniques to continue to improve and update our RJMCMC proposals as the sampling run proceeds. This allows our algorithm to adapt and ensure all detectable sources are found and characterized, as well as to maintain strong mixing in the sampler. This also makes it much easier for us to restart our global fit at later times rather than requiring us to build it up over time.

Our pipeline finds all 15 (6) MBHB sources in the training (hidden) data and provides posterior estimates for all of them. We determine the posterior distributions on parametrized models of the detector noise and galactic foreground noise components. We produce a catalog of ~ 12000 GB sources (~ 8000 with high confidence) and characterize the uncertainty in the template (or source) count found in the sampler at various times.

In the next section (Sec. II), we will detail and explain the gravitational-wave signals and likelihood computations for MBHB and GB signals. In Sec. III, we explain our parametrized model for the fit to the noise and foreground confusion. An overview of the MCMC methods employed is provided in Sec. IV. In Sec. V, we detail our full global fit algorithm. We follow that with our results for our LDC2A analysis in Sec. VI and a broader discussion of those results in Sec. VII. We end with some concluding remarks in Sec. VIII.

II. GRAVITATIONAL-WAVE LIKELIHOODS

For the current implementation of the pipeline, there are two gravitational-wave sources: MBHBs and GBs.

These are the two types of sources included in the LDC2A dataset (also called “Sangria”). This dataset is the first LISA Data Challenges (LDC) dataset designed to test multiple-source pipelines.

The two sources have vastly different waveform properties and generation methods. Here, we will detail the overall properties that apply to both and describe the specific models used to generate the signal templates.

Gravitational-wave signals are real-valued signals in the time domain denoted by $s(t)$. These signals are measured by the detector by observing their projection on the sensitive axes of the detector. In LISA, multiple signals will be detected simultaneously: $S(t) = \sum_j \sum_{i \in j} s_{ij}$ with j indicating the source type (MBHBs or GBs) and i indexing the sources within type j . The output of the detector measurements, from a data analysis perspective, are these signals combined with the noise in the detector: $d(t) = S(t) + n(t)$, where $n(t)$ is the noise component. To match these signals, we construct template waveforms, $H(t) = \sum_j \sum_{i \in j} h_{ij}$, built from theory and use these to detect and statistically characterize the signals present in the data stream. We assume that the noise is distributed according to a Gaussian distribution; therefore, we use a Gaussian likelihood based on the residual between the data stream and the specified combination of template waveforms. We will discuss more about this choice of likelihood and the noise properties in Sec. III. The natural log of the likelihood is given by

$$\ln \mathcal{L} = -\frac{1}{2} \langle d - H | d - H \rangle - \sum_f \ln S_n(f), \quad (1)$$

where $\sum_f \ln S_n(f)$ needs to be included when there is noise uncertainty, with $S_n(f)$ representing the power spectral density of the noise. We will also refer to this likelihood as the “base likelihood” or \mathcal{L}_b . The first term on the right-hand side is an inner product:

$$\langle b(t) | c(t) \rangle = 4\text{Re} \sum_f \frac{\tilde{b}(f)^* \tilde{c}(f)}{S_n(f)} \Delta f, \quad (2)$$

for real-valued time series $b(t)$ and $c(t)$, where $\tilde{b}(f)$ denotes the Fourier transform of $b(t)$ and Δf is the frequency spacing of the discrete Fourier transform. This term quantifies the quality of the template fitting by looking at the size of the residual. Both summations in Eqs. (1) and (2) are over all *positive* frequencies contained in the discrete Fourier transform of the observed discretely sampled time-domain signals.

There are a few useful quantities that we will also define here that will be helpful later on. When there is *only* one source in the data $H(t) \equiv h_1(t)$. In these circumstances, it can be useful to separate the source term in Eq. (1)

into constituent pieces using the bilinearity of the inner product:

$$\ln \mathcal{L}_1 \propto -\frac{1}{2}(\langle d|d \rangle + \langle h_1|h_1 \rangle - 2\langle d|h_1 \rangle). \quad (3)$$

The optimal signal-to-noise ratio (SNR) of the template waveform is then given by $\sqrt{\langle h_1|h_1 \rangle}$. The observed SNR of the source (when it is the only source in the data) is $\langle d|h_1 \rangle / \sqrt{\langle h_1|h_1 \rangle}$. When there are multiple sources in the data, the observed SNR of a single source is more complicated to define because of contributions from overlaps with the other sources in the data.

For the purposes of this work, we use two “special” likelihood formulations that are related to adjusting the data stream by one source. The first is the change in log likelihood associated with adding a single source, $a(t)$, to the current set of templates ($H + a$):

$$\begin{aligned} \ln \mathcal{L}_{+1} &\propto -\frac{1}{2}\langle d - (H + a)|d - (H + a) \rangle, \\ &= \ln \mathcal{L}_b - \frac{1}{2}[\langle a|a \rangle - 2\langle d - H|a \rangle], \\ \therefore \Delta \ln \mathcal{L}_{+1} &= -\frac{1}{2}[\langle a|a \rangle - 2\langle d - H|a \rangle]. \end{aligned} \quad (4)$$

In other words, if we have already calculated the base likelihood [Eq. (1)], we can determine the likelihood after adding one source by only computing terms related to the addition of the single template. This helps in the frequency domain to restrict the summation in Eq. (2) to frequencies at which the new signal is present. This is especially useful for individual GBs as they have a very small range of frequencies over which they emit gravitational radiation. Similarly to Eq. (4), we can also write down the change in likelihood due to the swapping of one source in the current template, $r(t)$, for another, $a(t)$:

$$\begin{aligned} \ln \mathcal{L}_{+1/-1} &\propto -\frac{1}{2}\langle d - (H + a - r)|d - (H + a - r) \rangle, \\ &= \ln \mathcal{L}_b - \frac{1}{2}[\langle a|a \rangle + \langle r|r \rangle - 2\langle a|r \rangle \\ &\quad - 2\langle d - H|a \rangle + 2\langle d - H|r \rangle], \\ \therefore \Delta \ln \mathcal{L}_{+1/-1} &= -\frac{1}{2}[\langle a|a \rangle + \langle r|r \rangle - 2\langle a|r \rangle \\ &\quad - 2\langle d - H|a \rangle + 2\langle d - H|r \rangle]. \end{aligned} \quad (5)$$

This once again allows for single-source-width computations. However, in Eq. (5), the number of frequencies needed to evaluate the $\langle a|r \rangle$ term will be between the frequency width of one and two sources depending on how much $a(t)$ and $r(t)$ overlap in frequency. If they perfectly overlap, then the effective width is the width of one source. If they do not overlap at all, the effective width is the

combination of the frequency widths of both sources. If they overlap imperfectly, the width will be between these two extremes.

As previously mentioned, observations from the triangular LISA detector will be preprocessed onground to reduce laser frequency noise using TDI [30]. The outputs of this pipeline are called “TDI observables.” The most basic set of these are the observables constructed from each set of two arms: X_1, Y_1, Z_1 . The “1” denotes the 1.5 Generation TDI observables [46,47]. The Sangria dataset uses the 1.5 Generation TDI observables, and assumes equal and constant detector arm lengths (TDI Generation 1.5 can handle a rigid and rotating LISA constellation; the arm lengths do not have to be equal). During the actual mission, we will be using second generation TDI observables [48], which are slightly more complicated. These observables will account for the rotation and flexing of the LISA constellation, meaning the arm lengths can vary with time. This is an ongoing area of study and will be added into the pipeline presented here as a part of future work.

The X_1, Y_1, Z_1 variables are correlated in their noise properties. Therefore, we make a linear transformation to A_1, E_1, T_1 variables [49,50], which, under idealized assumptions about the LISA instrument, would have uncorrelated noise properties:

$$A_1 = \frac{1}{\sqrt{2}}(Z_1 - X_1), \quad (6a)$$

$$E_1 = \frac{1}{\sqrt{6}}(X_1 - 2Y_1 + Z_1), \quad (6b)$$

$$T_1 = \frac{1}{\sqrt{3}}(X_1 + Y_1 + Z_1). \quad (6c)$$

For the sake of simplicity, in this work we do not consider the T_1 channel as it generally suppresses gravitational-wave signals. These uncorrelated variables simplify the analysis, but, as in the case of first versus second generation TDI, they will not be used in the final analysis. A more complicated likelihood matrix will be needed to properly take into account the existence of correlations between the channels.

A. Input template models

All waveform models used in this pipeline as templates are mathematically equivalent to those used in previous works. For that reason, we are not going to describe all of the details. Rather, we point the reader to the necessary citations to understand the specifics of the template generation methods used.

The MBHB waveforms are those used in [51]. One example waveform for an MBHB is shown in the frequency domain in Fig. 2. We use the aligned-spin IMRPhenomD waveform model [52,53]. A few




FIG. 2. An example of an MBHB waveform, GB waveform, LISA instrumental noise, and the galactic foreground (GF) confusion contribution. The MBHB waveform is shown in blue. The GB waveform is shown in orange. The GB waveform is very narrow band in frequency. For this reason, we also show an enlargement on the GB signal.

computational improvements have been made since that paper. The most important of these is a combination of the waveform generation with the likelihood computation. Previously the waveform was constructed and then stored in memory. It would then be passed to a likelihood function where the likelihood would be computed. Now, as each value of the waveform is generated at each frequency, within the same GPU kernel the likelihood at that frequency is immediately computed and entered into the full likelihood summation. This avoids the need to store full waveforms in memory, allowing us to run more likelihood computations simultaneously across a single GPU.

The GB waveforms are based on the FastGB [54] algorithm. FastGB is an efficient GB waveform generation method that works by leveraging the small \dot{f} observed in GB sources to separate the waveform generation into a fast-time scale piece and a slow-time scale piece. They were updated and used for circumbinary exoplanet analysis with LISA in [55]. For this work, we reimplemented FastGB in a special form for efficient usage on the GPU. This is discussed further in Appendix F. An example GB waveform is shown in Fig. 2, as well as an enlargement of its very localized frequency domain structure.

For more specific information on each waveform model, we refer the reader to the publications cited above.

III. NOISE TREATMENT

In the Sangria dataset, as well as our analysis, the detector noise is modeled as stationary and Gaussian. In reality we expect a variety of nonstationary effects including data gaps, glitches, and sensitivity drift over time [2]. The effect of removing these assumptions is a topic of

ongoing research and other LDC datasets will be produced for this purpose. In the future, these advances will be combined with the multiple-source pipeline presented here as those techniques are further developed.

Under the assumption of stationarity and Gaussianity, the noise covariance matrix in the frequency domain reduces to a vector (the matrix diagonal) due to the absence of correlations between different frequency components. Hence, the summations shown in Eqs. (1) and (2) are one-dimensional over frequency rather than two-dimensional, which would be needed if there were off-diagonal elements of the noise covariance matrix.

The exact noise in the detector is uncertain and needs to be fitted as an additional component of the analysis. There are a few different suggested methods for this process. We consider a spectral model for the noise as the sum of two major noise components: the acceleration and optical metrology system (OMS) noises. The overall PSD, S_n , in the A and E channels is given by [56]

$$S_n(x) = 8 \sin^2 x [2S_{\text{pm}}(3 + 2 \cos x + \cos 2x) + S_{\text{op}}(2 + \cos x)], \quad (7)$$

where $x = 2\pi(f/f^*)$ with $f^* = c/L$ (the speed of light divided by the LISA arm length). S_{pm} is related to the acceleration noise, S_{acc} , by

$$S_{\text{pm}}(f) = S_{\text{acc}}(2\pi cf)^{-2} \left[1 + \left(\frac{0.4 \times 10^{-3} \text{ Hz}}{f} \right)^2 \right] \times \left[1 + \left(\frac{f}{8 \times 10^{-3} \text{ Hz}} \right)^4 \right]. \quad (8)$$

Similarly, S_{op} is related to the OMS noise, S_{oms} :

$$S_{\text{op}}(f) = S_{\text{oms}} \left(\frac{2\pi f}{c} \right)^2 \times \left[1 + \left(\frac{2 \times 10^{-3} \text{ Hz}}{f} \right)^4 \right]. \quad (9)$$

In our global fit pipeline, we fit for $\sqrt{S_{\text{oms}}}$ and $\sqrt{S_{\text{acc}}}$ in both channels (i.e., four parameters in total). The injection noise for the LDC2A was generated using $\sqrt{S_{\text{oms}}} = 7.9 \times 10^{-12} \text{ m Hz}^{-0.5}$ and $\sqrt{S_{\text{acc}}} = 2.4 \times 10^{-15} \text{ ms}^{-2} \text{ Hz}^{-0.5}$ [47]. Figure 2 shows the injection noise PSD as the gray dashed line.

Please note the above model is used for fitting the data. Due to a previous omission from the LDC documentation [34], we did not include a third noise, backlink noise, that was included in the LDC2A dataset noise injection. We discuss this further in the Sec. VI.

A. Galactic foreground noise parametrization

To characterize the galactic foreground noise we also choose a simple parametrized model that describes the shape of the expected foreground spectrum. In future work,

we will be exploring other potential models to parametrize the noise and foreground.

Our foreground noise model contains five parameters: A_{fg} represents the overall amplitude of the foreground; α is a power law index; f_k is the “knee” frequency where the underlying slopes change; and s_1 and s_2 are the slopes before and after the knee frequency, respectively. The analytic formula for the foreground noise *sensitivity*, h_{fg} , is given by [57]

$$h_{fg}(f) = \frac{1}{2} A_{fg} f^{-7/3} e^{-s_1(f^\alpha)} [1 + \tanh(-s_2(f - f_k))]. \quad (10)$$

The contribution to the PSD in the A and E channels from the galactic foreground is given by

$$S_{fg} = 6x^2(\sin^2 x) h_{fg}(f), \quad (11)$$

where $x = 2\pi(f/f^*)$ as in the previous section. This contribution is directly added to S_n given in Eq. (7) to give the total PSD used in the inner product [Eq. (2)]. Please note that, while we parametrize the instrument noise separately for the A and E channels, we parametrize the foreground contribution as one foreground of five parameters that is then shared across both channels.

An example foreground is shown in Fig. 2 for parameters $(A_{fg}, \alpha, f_k, s_1, s_2) \approx (3.3 \times 10^{-44}, 1.2, 2.8 \text{ mHz}, 1763.3, 1686.3)$.

IV. MARKOV CHAIN MONTE CARLO

Throughout the entire pipeline, both in search and parameter estimation, Markov chain Monte Carlo (MCMC) techniques are heavily used. Here we will give a quick primer and references for these MCMC methods. As we describe further details of the full pipeline later, we will indicate where each of these methods is used. We begin with Bayes’ rule:

$$p(\vec{\theta}|D, \Lambda) = \frac{p(D|\vec{\theta}, \Lambda)p(\vec{\theta})}{p(D|\Lambda)} = \frac{\mathcal{L}(\vec{\theta})p(\vec{\theta}|\Lambda)}{Z(\Lambda)}, \quad (12)$$

where $p(\vec{\theta}|\Lambda)$ is the prior probability for model Λ ; $p(\vec{\theta}|D, \Lambda) \equiv \pi(\vec{\theta})$ is the posterior probability of the parameters of model Λ given data D ; $p(D|\vec{\theta}, \Lambda) \equiv \mathcal{L}(\vec{\theta})$ is the likelihood function; and $p(D|\Lambda) \equiv Z(\Lambda) = \int_{\vec{\theta}} \mathcal{L}(\vec{\theta})p(\vec{\theta})d\vec{\theta}$ is the evidence for the model Λ .

We build our entire pipeline around the Eryn sampler [43–45], which is an ensemble-based sampler (ensemble: multiple walkers per temperature rather than just one) designed for use in both fixed-dimensional and transdimensional MCMC settings.

As discussed in previous global fit work [21], the number of templates needed to fit the data is uncertain.

Therefore, we have to use sampling methods that allow for the addition and removal of templates over the course of an MCMC run. For this, we employ the use of reversible-jump (or transdimensional) MCMC [38]. “Between-model” proposals are so called because they propose to change the underlying model or model dimension (Λ) rather than just the model parameters ($\vec{\theta}$).

In our situation, where all templates are of the same kind, the addition and subtraction of templates is considered “nested,” which vastly simplifies our problem. In a nested setting, this means changing the source count by +1 or −1. The acceptance fraction for nested models when adding one component is given by [58]

$$\alpha_{RJ} = \min \left[1, \frac{\mathcal{L}(\vec{\theta}_{k+1})p(\vec{\theta}_{+1})}{\mathcal{L}(\vec{\theta}_k)q(\vec{\theta}_{+1})} \right], \quad (13)$$

where k is the current number of templates in the data, “+1” indicates the proposed model addition, and q is the proposal distribution. In other words, the acceptance fraction becomes the ratio of the likelihood values of the new and previous point multiplied by the ratio of the prior density to the proposal density of the newly added source.

The prior terms for the first k binaries, $p(\vec{\theta}_k)$, cancel out because they are found in both the numerator and the denominator. When proposing to remove a source, the acceptance fraction is inverted and the “+1” is replaced with a “−1” now indicating the source that is to be removed. At the moment we only use RJMCMC in the GB sampler. All RJ moves within that sampler are based on the acceptance ratio given in Eq. (13), but they have different proposal distributions ($q(\vec{\theta}_{+1})$). These distributions will be described below.

MCMC moves that conserve the model type and/or the model dimensionality (or source count) are referred to as “in-model” or “within-model” proposals. These proposals only change the parameters contained in the current model. In *fixed-dimensionality* or *fixed-model* settings where reversible jump is not needed, we use proposals built around the affine-invariant or “stretch” proposal [59]. The stretch proposal relies on the ensemble setup where multiple MCMC walkers are run simultaneously. As a reminder, a walker represents an instance of a moving sampler state in parameter space that will build up the MCMC chain samples over time. The stretch proposal moves a given walker using the parameter space locations of the other walkers in the ensemble:

$$Y_k = X_j + z(X_k - X_j), \quad (14)$$

where X_k and Y_k are the current and proposed positions of the walker of interest, respectively; X_j is another walker drawn uniformly from the remaining members of the ensemble (excluding the k th walker of interest); and z is

a randomly distributed variable drawn from a distribution with specific requirements [59]. A main benefit of this type of proposal is that it does not require specific tuning-based adjustments. For more information on the stretch proposal see [59]. In the PSD and MBHB portions of the global fit sampler we use the regular stretch proposal as described in [43].

In our GB sampler where we employ reversible jump, we use the “group stretch” proposal discussed in [44] because the regular stretch proposal does not scale properly to a reversible jump setting where there is, in general, a different number of model instances for each walker. The key difference between the base stretch and group stretch proposals is the determination of X_j in Eq. (14). Rather than drawing from the current walkers of the ensemble, X_j is drawn from a “stationary group” [44] that stays fixed through a large number of proposals (required to satisfy detailed balance). For this group, we use a flattened set of all templates contained in the cold chain in all current walkers within the sampler. Individual groups are formed by sorting this large set of templates by frequency and taking the n_{friends} closest sources in frequency to the GB of interest. We update this group every ~ 50 proposals and are careful to satisfy detailed balance during the changeover. The purpose of this setup is to facilitate efficient in-model sampling in RJMCMC without needing to calculate or tune a covariance matrix or similar such proposal.

We employ the technique of parallel tempering in all modules of our pipeline. We discuss the specifics of our tempering implementation in Appendix D. There, we include details about how a typical tempering scheme must be adjusted to handle a residual-based analysis as is used here.

V. THE GLOBAL FIT

The global fit algorithm presented in this work is very similar in nature to the methods presented in [21], with some important differences. The main idea is that the sampler uses a blocked Gibbs sampling technique [39–41] to analyze the MBHBs, PSD (and foreground noise), and GBs as separate sampler entities. Then, periodically, each individual sampler will read out its current cold-chain information to a controlling process. The controlling process will then transmit this to all participating individual samplers so that they update their current residuals and PSD information. It is this operation that marginalizes over other source and PSD posterior distributions contained in the larger global fit. This is effectively the “wheel” method shown in Fig. 1.

Our global fit algorithm has four simultaneously running modules: PSD and foreground fitting; MBHB fitting; GB fitting; and a fourth module that contains a search and updating proposal setup for the GBs. A schematic diagram of this process is shown in Fig. 3.

A. Noise and stochastic foreground signal PSD fitting

The PSD and foreground noise fitting module is the most straightforward of the four. As described above, each channel (A and E) will have a set of two parameters (acceleration noise and OMS noise) to describe the base LISA sensitivity. This results in four total parameters representing the PSD, with the two sets of parameters acting only on their associated channel.

The foreground noise, as detailed in Sec. III, is parametrized by five parameters (amplitude, power-law index, knee frequency, and slopes before and after the knee frequency). In contrast to the four PSD parameters, the foreground noise represents an astrophysical model meaning it is computed independently of each channel and is then transformed and added across both A and E channels.

MCMC sampling of these parameters directly uses Eryn and its constituent resources. A fixed-dimensional MCMC is run with ten temperatures ranging from the cold chain sampling the posterior ($T = 1$) to an infinite temperature chain sampling the prior distribution, with eight geometrically spaced temperature rungs in between that are allowed to adapt over time [60]. The regular stretch proposal is used in this fixed-dimensional application.

Before beginning to search for astrophysical signals, we run the PSD and foreground sampling operation until it converges to an initial estimate. This estimate is usually okay for the instrumental noise parameters; however, the foreground fit strongly compensates for the lack of subtraction of GW signals. Over the course of the global fit run, as GW signals are found and removed from the residual, the foreground and instrumental noise parameters converge to a their proper state.

During each update from the controlling process, residuals from the GB and MBHB modules are drawn uniformly from the batch of cold-chain residuals coming in. Each randomly drawn residual is assigned to one walker in the PSD cold chain. Then, in a vertical fashion, the residual assigned to walker w in the cold chain is used for all walkers of index w in all of the higher temperature chains. With this in mind, and for computational efficiency, the ensemble permutations during temperature swapping, which is the default in Eryn, is turned off. Temperature swaps are, therefore, only proposed for walkers with the same walker index w located in neighboring temperature chains. With this setup, it is not required to calculate new likelihoods as described in Appendix D because the underlying residual and PSD against which the likelihoods are computed are the same.

B. Searching for MBHBs

With the initial noise PSD fitting complete, we extract the maximum likelihood PSD and foreground noise components. As previously mentioned, the foreground noise at this point is much higher than it should be, since no astrophysical signals have been subtracted. The MBHBs in

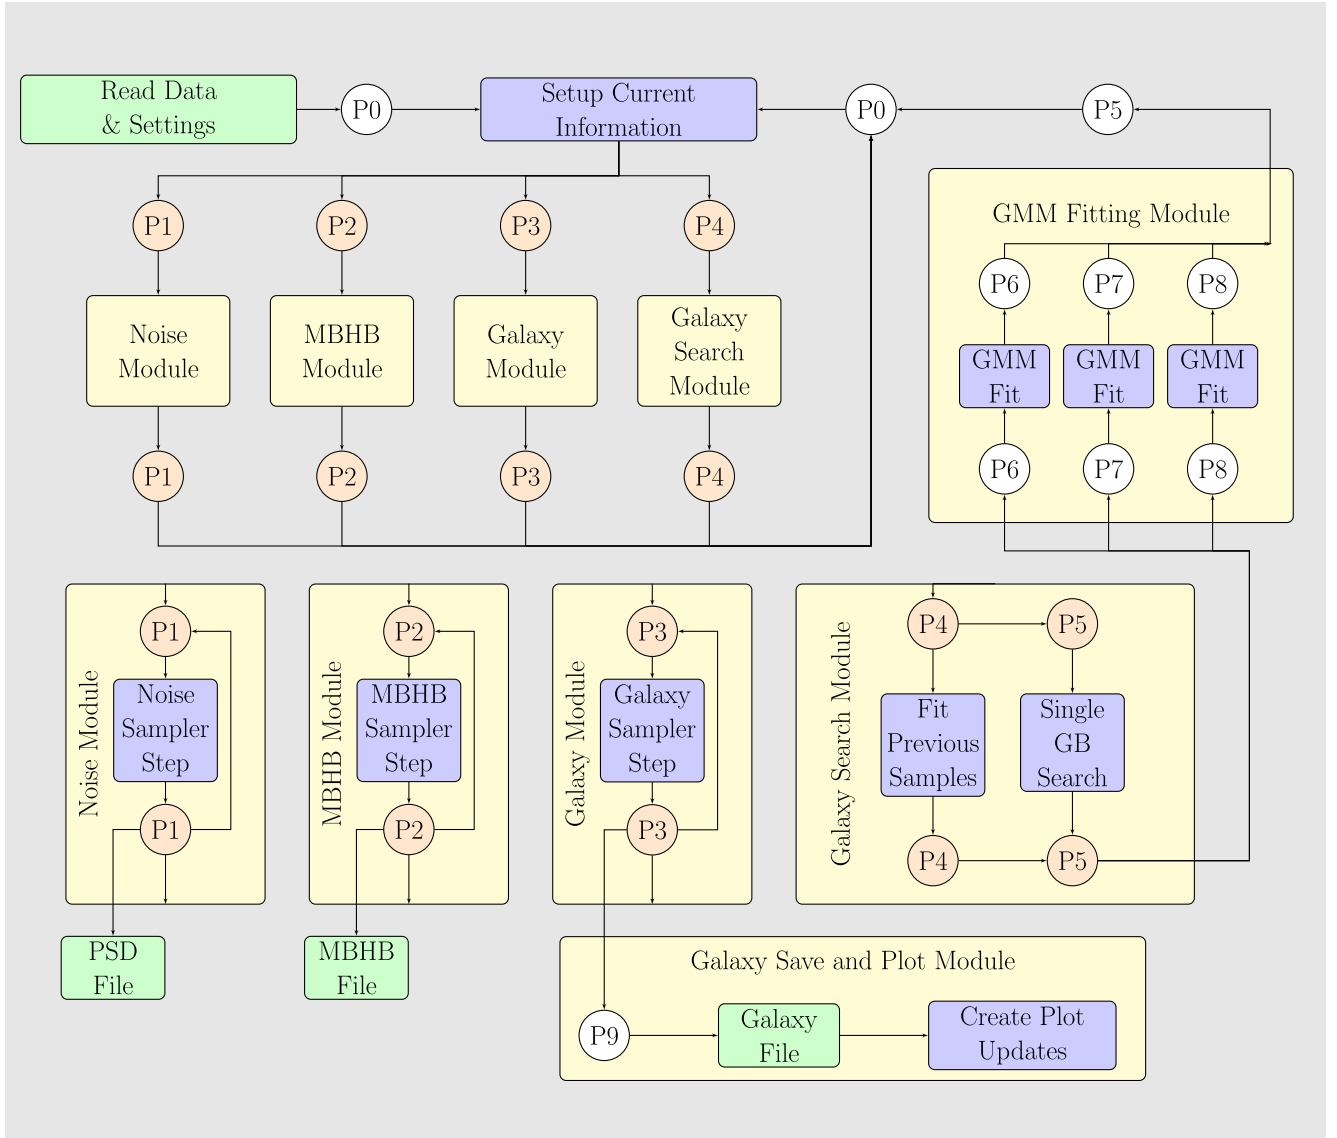


FIG. 3. Process diagram for the Erebor algorithm. Each circle with “P” indicates a separate process. Orange circles indicate processes that communicate with and use GPUs. P1, P2, and P3 communicate with one GPU each. P4 and P5 communicate with one GPU total between the two. File-related operations are in green. Separate but communicating modules are shown in yellow. Major algorithmic partitions are given in blue.

the LDC2A dataset are loud enough that all sources can still be detected in this setting. In the future, when lower SNR MBHB signals are included, this method will need to be extended as the confusion noise component is successively reduced. We will discuss this further in Sec. VII.

The one-year dataset is split into multiple segments, usually determined by the number of GPUs available. In our case, running four GPUs, we split the dataset into eight segments, running two of them per GPU so that we are analyzing the full year simultaneously.

Within each segment, we further split the larger pieces into chunks of $3/2$ days, overlapping the first $1/2$ day and last $1/2$ day with neighboring chunks. Therefore, each

chunk focuses on a specific $1/2$ day, but it contains data from the $1/2$ day before and $1/2$ day after. The purpose of this is to remove, as best as possible, any confusion between MBHB mergers, which can inhibit the search.

These chunks are then transformed to the frequency domain and stacked as “separate” data streams within the MCMC sampler. Within each of the eight larger segments, we allow the merger time to vary only across that segment (i.e., the edges of the prior on the merger time are the edges of the $1/8$ -year time window). When it is time to evaluate the likelihood for a sample in the MCMC run, we determine which $1/2$ -day window its merger time falls into. We then calculate the likelihood using the associated $3/2$ -day

window, centered on the $1/2$ day in which the merger occurs. This, once again, allows us to focus on merger detections specifically. There are advantages and disadvantages to this which will be further discussed in Sec. VII.

This strategy will produce discontinuities in the likelihood when the merger time crosses from one window to another. However, in practice these discontinuities create only small disturbances in the likelihood. They are further minimized by including the $1/2$ day before and after the window of interest. Therefore, this does not cause any noticeable issue in the sampling allowing this search process to proceed properly.

Each of the eight larger pieces of data are analyzed independently. The MCMC in each chunk is run until the likelihood converges to a maximum value (remaining fixed for 500 sampler iterations) and every sample in the cold chain is within a few tens of that maximum value. The detected and optimal SNRs are then calculated. Right now, we apply a threshold of 20 to the smallest of these two SNRs. This choice is loosely based on what was found empirically in previous works [e.g. [51]], but in the future we plan to further investigate this choice of threshold and its effect on the MBHB search results. Under an SNR of 20, we begin to confuse potential sources with a combination of the noise and the GB signals. However, as stated earlier, the loud MBHBs in the Sangria datasets allowed us to leave the resolution of this issue to later work.

If the source found has detected and optimal SNRs greater than 20, the source is kept and added to a running list of found sources. These sources are then subtracted from subsequent runs. This process is repeated until each of the eight larger-time pieces are unable to locate a source above the SNR threshold.

This process usually results in finding more sources than are actually there. We prune our catalog for duplicate entries. The probability of an extreme match between two astrophysical MBHB signals is infinitesimally small; therefore, we assume that a high match between the signals indicates imperfect extraction of the first signal found, with a second or third signal completing the extraction process. The first source is usually very close to the true point with subsequent found sources attaching to a much smaller fraction of the signal. We keep, as is, only the highest SNR source within each of these multiple-source groups. After performing this pruning operation, we end up with the 15 (6) MBHBs that are truly injected in the LDC2A training (hidden) dataset, indicating its success. This will obviously need to be tested on more datasets in the future, but our initial tests give us confidence it will perform correctly. In addition to the sources of interest, this search also provides a set of samples that act as a starting point for each source in the larger global fitting run to come.

Following the culmination of the MBHB search pipeline, we run the global fit sampling operation without the GBs: the PSD, foreground, and MBHB parameters are allowed to

vary until the PSD and foreground estimation properly accounts for the MBHBs removed from the residual. This process culminates when it converges to a maximum likelihood value. Following this occurrence, GBs are added back in and the full global fit run commences.

C. MBHB parameter estimation

The MBHB parameter estimation operation is more demanding than the PSD module, but not as difficult as the GB module. It runs in a very similar way to the PSD module: it uses ten temperatures spaced the same way as the PSD run; a similar update function to read in and out information; and a vertically arranged set of residuals and PSDs for marginalization, removing the permutation property of the ensemble tempering scheme (for computational efficiency).

The main difference between the MBHB part of the sampler and the PSD module is that there are multiple MBHBs in each of the datasets. In this iteration of our global fit algorithm, we operate the MBHB portion as a fixed-dimensional sampler that assumes the search has found all sources of interest. In future work, this will be adapted into a transdimensional setup that is similar in nature to how we currently run the GB module.

We implement a special proposal move designed specifically for sampling across the MBHB sources. It is described in Algorithm 3 in Appendix H. The main idea, like the global sampler as a whole, is to sample each source in a parallel tempered fashion while marginalizing over only the cold chain information from the other sources.

We leverage the fast heterodyned likelihood functions available in `bbhx` [61,62] for improved efficiency. For more information on heterodyning MBHB waveforms, see [51,62]. The MBHB proposal algorithm is given in the Appendix.

The “inner proposal” that operates each individual step of the sampling operation is chosen each step at random from the four proposals detailed in [51]. The main proposal chosen with $\sim 90\%$ probability is just the regular stock stretch proposal from Eryn [44,45,59]. The other three split the last 10%. These are “sky moves” that leverage sky-mode symmetry by changing the sky-mode in which the point lies, while keeping all other parameters fixed. The three different moves include an only-latitudinal move, an only-longitudinal move, and a move that is both latitudinal and longitudinal. See [51] for more information on these moves.

D. Galactic binary search and parameter estimation

By far the most complicated operation of our global fit pipeline is the detection and characterization of GBs. The complication is due to many aspects of the problem, including the large number of expected signals, a high degree of confusion between signals, and the fact that there is an unknown number of signals.

To accomplish this task, following the lead of [21], we use reversible jump MCMC [38] to characterize the number and type of signals detectable in the LISA data. In this section, we will describe the overall properties of the GB algorithm, detail the various proposals used in the GB sampler, and discuss how we proceed naturally through search and parameter estimation stages.

With GBs in the global fit, there are three scales of concern. The smallest scale represents the individual sources (dimensionality, $D \sim 8$). The largest scale involves all sources across the whole frequency band ($D \sim 10^5$). The middle scale is the scale over which individual source posteriors interact with each other. While this varies across the LISA frequency band, this scale involves the interaction of ~ 2 –50 sources via the overlap in their frequency domain spectra ($D \sim 16$ –400).

The GB algorithm is designed to primarily handle the middle scale (source interaction) while successfully meeting both the small- and large-scale requirements. We break up the computation across the full band into a sequence of subbands whose widths are primarily determined by the widths of the GB frequency spectra. The subband to which a source belongs is determined based on its initial frequency parameter. With one year of observation, the subband widths employed are 128 frequency bins (each of width $1/T_{\text{obs}}$) at low frequencies (< 1 mHz), 256 at middle frequencies (< 10 mHz), and 1024 at high frequencies (> 10 mHz). For more information on this construction, please see [54].

Within each subband is really where the MCMC-related operations are run. This means proposals, tempering, and accounting for source changes are all performed at the subband level. Within the ensemble setup, we leverage multiple walkers and multiple temperatures for each subband. However, there are a few important limitations. First, while sources are grouped into the subbands by their initial frequency parameter, their spectrum may extend beyond what we call the “band edge” of the subband. Figure 4 shows the leakage of the GB template models into neighboring subbands: the purple template waveforms cross over band edges represented with black dashed lines. This is natural and represents our inability to specifically localize each source to its own set of frequencies. For this reason, subbands are run in an “odds and evens” way, where the odd-numbered bands are run first and even numbered bands are run second. This holds fixed sources with spectra that may leak across the band edge while sources within the band of interest are moved with MCMC operations.

We allow samples to go halfway into neighboring bands while the sampling within a band of interest is occurring. This prevents the effect of “hard” edges barring samples from moving across them in their initial frequency parameter. After a sequence of sampling operations within the band of interest, if a cold chain sample falls into a

neighboring subband, the subband ownership over that sample is shifted to the neighboring band it now falls under. This treatment of these “edge effects” is similar to [21].

Second, similar to the PSD and MBHBs, we want all residuals to be cold-chain information. In other words, we want all computations to be performed on residuals representative of the posterior distribution ($T = 1$) which ensures the marginalization efficiency is maximized. For this reason, all subbands at higher temperatures are marginalized against neighboring subbands from their associated walker index (w) in the cold chain. In the GB sampler, because the tempering is not as much of a computational bottleneck, we choose to do a full ensemble-permuted tempering setup. This helps to filter the various subband states more efficiently to the cold chain while properly marginalizing over the various residuals and PSD information coming from the other parts of the global sampler.

We run our sampler with 18 walkers in 24 temperatures, allowing for a large number of running subbands to better sift through all the signals and output the desired information. Our temperature ladder at the start ranges from one to infinity with the temperatures in the middle spaced geometrically by a factor of 1.2. These temperatures do adjust over time. Each subband is given an independent temperature ladder. This means that the subbands will generally have different temperature ladders that will ideally adjust to their differing characteristics. Only their cold chain information will be able to directly translate between subbands because the cold chain is the only temperature rung guaranteed to be at the same temperature as its neighbors at all times.

As the sampling proceeds, a user-defined number of in-model proposals are performed consecutively (we use ~ 30) followed by one reversible jump (RJ) proposal and then temperature swapping. For in-model proposals, we use the group stretch proposal discussed in Sec. IV.

For between-model moves, we use a mix of various proposals. The most basic RJ proposal is a draw from the prior distribution. We refer to this proposal as the “prior RJ” proposal. It is the least informative proposal, but is helpful to ensure the entire space is explored.

E. Online updating of RJ sampling distributions

In addition to the prior RJ proposal, we developed two proposals based on fitting Gaussian mixture models (GMM) [63] to empirical sample distributions for resampling within the larger posterior estimation run. The first of the GMM-based proposals, the “search RJ proposal,” searches every subband independently and simultaneously for the maximum likelihood source that still remains to be subtracted from the best-fit residual. To do this, it runs a sampler that is fixed dimensional, parallel tempered, and stretch-proposal based. Once the highest likelihood point still remaining in each subband is found, samples are

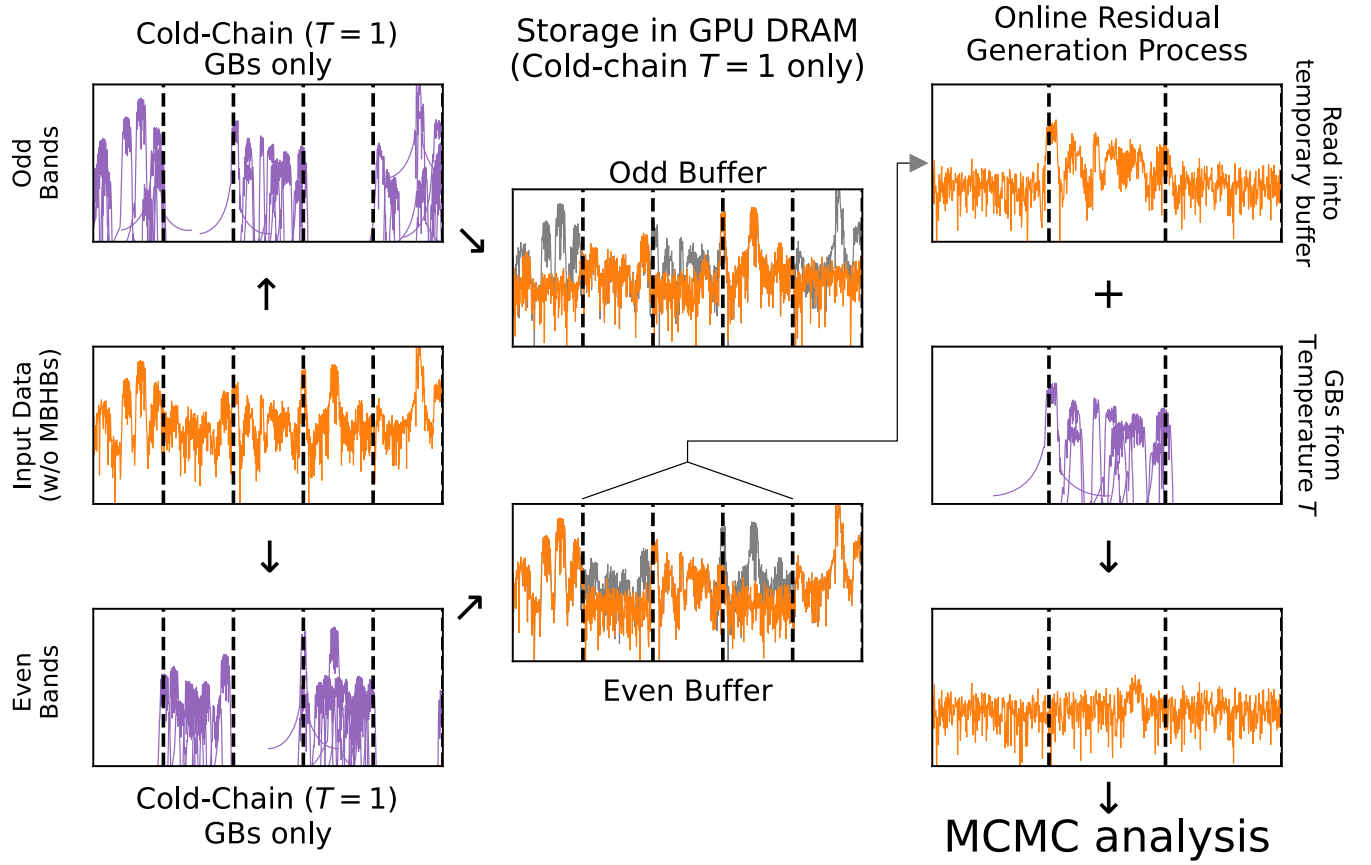


FIG. 4. Illustration of the subband structure used to analyze the GB population, as well as the special storage of residual information for maximal memory efficiency. The left column, in the middle, shows the input data with the MBHB signals subtracted from it. The dashed black lines divide the window into subbands. These subbands represent the analysis unit of the GB algorithm. The subbands are run in parallel rotating between odd and even groups. During each specific subband analysis, we need information from three subbands: the subband of interest, the subband before, and the subband after the subband of interest. In the cold chain ($T = 1$), this is straightforward. At higher temperatures ($T > 1$), the subbands before and after the band of interest need to come from the cold chain ($T = 1$) because neighboring subbands will have generally different temperature ladders with only the cold chain able to transfer information across the subband edge. If we stored three subbands for each subband of interest, it takes (per TDI channel) $\sim 3 \times n_t \times n_w \times n_b \times (\text{length of each subband}) \approx 3 \times n_t \times n_w \times (\text{data length})$ 16-byte numbers. This can quickly fill available GPU DRAM (device RAM). Instead, we store *cold chain* odd-band and even-band buffers that are read in at runtime and adjusted to the above three subband setup needed for analysis. This uses (per TDI channel) $2 \times n_w \times (\text{data length})$ 16-byte numbers, reducing the storage needed by a factor of $\sim n_t/2$. To form these odd and even buffers, we duplicate the input data (with MBHBs removed) and subtract waveforms belonging to odd or even bands from the two instances of input data (purple waveforms in the top and bottom plots in the left column). The resulting buffers are illustrated above in the center column. The orange shows the stored residual and the gray shows the original input data to use as reference to see which subbands have been updated. This storage method separates into odd and even bands while maintaining any signal leakage across subband edges. The online residual formation process is shown in the right column. When running online computations on an odd (even) band at any temperature, the three subbands needed for the analysis are read in from the even (odd) cold-chain buffer to a temporary GPU buffer. Then, the proper residual for analysis is formed by removing the waveforms belonging to the temperature, walker, and subband of interest.

generated around each point to build a distribution to sample from within the larger sampler. We fit each source separately to its own GMM and then combine them all into one larger GMM by assigning each source-specific GMM as having equal probability. More information about the algorithmic side of this process can be found in Appendix G.

A similar process is performed for the other GMM-based proposal. However, rather than the inputs being points

found from a search operation, the inputs are samples that repeatedly appear in the sampler. Once the sampler reaches a user-defined number of iterations, we go back roughly ten (thinned) samples within each walker and perform a cataloging operation (see Sec. V G). The samples that appear regularly are grouped into distinct distributions that are then fit individually with a GMM. Similar to the search RJ proposal, we combine all these individual distributions

into one larger GMM by assigning equal probability to each individual GMM. We refer to this proposal as the “refit RJ proposal.” The main point of using this proposal is to ensure satisfactory mixing at higher temperatures by proposing points that have already been found, but may not appear in all higher temperature chains depending on their SNR.

F. Transition from search to parameter estimation

At the start of the full global fit run, after the initial search for the MBHBs, overall instrumental noise and foreground confusion level, we begin the GB sampling run with special settings. The purpose of the initial search settings is to prevent a specific problem: when starting with zero GBs in the model, it is common for the sampler to fill itself with an excess of low-SNR sources while it takes time to find the louder sources. This is mainly caused by proposals like the prior proposal, since it uses an uninformed distribution. This can clog up the sampler with many unnecessary computations, significantly slowing its convergence. Additionally, even in a more refined state that is close to the ground truth, the sampler efficiency is highly dependent on the total number of templates currently being used. Therefore, many unnecessary and false detections cause the sampler to slow down.

This is more of an issue when starting directly from one year of data. In [21] the authors start with 1.5 months of data and gradually build up to use the whole year. They use the information from the previous duration to help the sampler more quickly converge for longer durations. This is a different but effective way to deal with the issue of small SNR source buildup when beginning sampling runs at longer observation times.

The approach we take to solve this issue, when using one year of data from the beginning, is to set up our RJ proposals in a specific way. We use only the search RJ proposal to both add and remove sources. This proposal is selected on average 80% of the time. We do not use the refit RJ proposal until later in the search run. We use the prior proposal $\sim 20\%$ of the time, but, importantly, *we only allow it to remove sources*. Preventing the prior RJ proposal from adding sources prevents the addition of a large chunk of low SNR sources. It also helps to “prune” the ongoing search by easily removing any sources that evolve through RJ or in-model moves to an SNR too low for them to remain in the sampler consistently.

When the GB sampler begins, the search RJ proposal is finding the loudest sources in each subband. This distribution is then used to propose new points, which ensures the loudest sources enter the model first. As the search RJ proposal updates, it will use the best-fit data residual, which allows it to extract sources and move to lower and lower SNRs over the course of the run. To keep the large number of extra low SNR sources out of the data residual, we use an optimal SNR limit of 7 in the search RJ proposal at the beginning.

The search operation finds ~ 600 sources in its first iteration across all subbands. As the search proceeds, the number found in each iteration decreases. The end of the search phase comes when the maximum number of templates in the sampler converges. At this point, the settings are switched to those appropriate for full parameter estimation.

During parameter estimation an SNR limit of 5 is used for the search RJ proposal. The search proposal and refit RJ proposal are each used $\sim 20\%$ of the time. The prior RJ proposal is used for $\sim 60\%$ of proposals (both for addition and removal). These settings allow the sampler to focus on sources at the limit of detection, that may not have been found properly during the search operation, while still ensuring good mixing of chains at higher temperatures where the larger SNR sources are added and removed.

G. Cataloging operation

Our cataloging operation begins with an optimal SNR cut on all collected individual samples from the parameter estimation run: optimal $\text{SNR} > 7$. We begin with one sample (meaning $\sim 10^4$ templates from across the band) as the control group against which we make our initial overlap computations. We group all other templates in every other sample into a test group. For each binary in the test set, we find the maximum overlap with the sources of the control group. If this maximal overlap is greater than 0.9, the test set source is added to the group of the control source it matches with. We repeat this step with the next sample as the control and all other nongrouped sources as the test set. This continues until all samples have been treated as the control. Then, to remove any imperfections, we merge any groups that have an overlap greater than 0.9 with each other.

At the culmination of the cataloging operation, each of these groups is assigned a “confidence” value, C . The confidence is the number of samples contained in a cataloged group divided by the total MCMC sample count. For louder sources with any appreciable SNR, this confidence fraction will be close to or equal to 1 because they will appear in every sample extracted from the global fit. We note the cataloging operation is not perfect. It requires making some set of assumptions and/or cuts. For this reason, a few samples may be missing from clearly detected sources, making their confidence values close to but not equal to 1.

For sources near the SNR detection limit ~ 7 , the confidence value may decrease because the RJMCMC algorithm may locate a model configuration that removes the marginally detectable source from the current sample. Confidence values for these sources can range from 0 to 1. Since full global fit runs and results have only recently been produced, the cataloging step—its choices and ramifications—are not well understood and are expected to improve in the future. Similarly, the relation of our

confidence readings to these cataloging choices is a topic for future investigations. We also do not account for two or more cataloged binaries fitting one real source. Our initial checks indicate this does not happen often; a more detailed analysis of where and why it happens is left for future work. With these important points in mind, we believe our cataloging operation is good enough to prove our global fit algorithm's capabilities.

After the full operation is done and confidences are accounted for, the samples that make up the catalog are saved to HDF5 files to be read in and analyzed using LISACATTOOLS [64]. Many of the plots shown below were produced using this LISA catalog analysis package.

H. Large-scale algorithmic setup

The global fit is an extremely compute intensive procedure. For this reason, it is absolutely essential to have efficient sampling and likelihood computations. We focus here on the large-scale setup. More specific information about the proposals, tempering, and the distribution fitting pipelines can be found in the Appendix.

Our overall goal was to use GPUs to handle as much of the global fit analysis as possible. The most important reasons for this are the efficiency of GPUs when it comes to handling highly parallelizable problems and the fact that GPUs are expected to continue to improve at a faster rate than CPUs in the near future leading up to the LISA mission [42].

Every separate module (MBHBs, GBs, PSD, etc.) in our algorithm is run simultaneously using `mpi4py` [65–68] for launching and communicating between all current computing processes. Each separate module uses a single GPU of its own (except for the MBHB search operation that uses as many GPUs as you give it). All of the main sampling modules run directly out of `Eryn` [44,45] for sampling. The only customized parts of this are the special MBHB and GB proposals, as well as the GB tempering piece. They are, however, directly worked into the highly modular framework of `Eryn` for adding customized proposals and tempering setups.

Besides the specific proposals mentioned previously for the PSD, foreground, and MBHB analysis, their GPU-accelerated, vectorized likelihood functions were important to the efficiency of this pipeline. The fast GPU-based PSD likelihood was created specifically for this pipeline to ensure its computational and memory efficiency. As mentioned previously, the MBHB sampler uses heterodyning [61,62] to improve its likelihood evaluation time.

For rapid evaluation of GB sampling operations while maintaining memory efficiency, we use a special setup for the storage of the GB residuals in the GB sampler. We store two sets of residuals: one for cold-chain residuals from odd-numbered subbands and one set for cold-chain residuals from even-numbered subbands. However, both of these sets of residuals contain the original data with all

MBHBs subtracted. Therefore, to recover the full residuals across the entire band, it is necessary to add the two sets of residuals together and then subtract out the original data (with MBHBs subtracted). This process is illustrated in Fig. 4.

The main idea here is that all subbands within the analysis need information only from the neighboring *cold-chain* subbands. To illustrate this usage, imagine we want to operate at temperature t for walker index w in subband index b which we take to be odd. We read in subbands $b - 1$, b , and $b + 1$ from the even-numbered cold-chain residual of walker w into a GPU buffer. The only template waveforms subtracted from band b (which is odd) in the even-numbered cold-chain residual are sources in the neighboring bands that leak over into the band of interest. We subtract from the buffer residual the waveforms contained in $(\text{walker}, \text{temperature}, \text{subband}) = (w, t, b)$. This produces, in the buffer, the full residual at the start of sampling with all current sources in the sampler removed from the subband of interest. This process is illustrated in Fig. 4. This simulates the subband of interest at each temperature by reading in the proper cold-chain residuals of its neighbors and filling information in the subband of interest with its own sources at its specific temperature. This means we do not need to store any of the higher-temperature residuals in GPU global memory. Memory optimization on GPUs is very important. If we stored every higher temperature residual, the GPU memory burden would be $n_w \times n_t \times (2 \text{ channels}) \times (\text{memory size of each residual})$. We reduce this by a factor of $n_t/2$ since we have twice the number of cold-chain residuals.

The alternative easier choice is to store three subbands worth of information ($b - 1, b, b + 1$) for every (w, t, b) combination (with $b - 1$ and $b + 1$ coming from the cold chain). This would cost $n_w \times n_t \times n_b \times (\text{two channels}) \times (\text{memory size of three subbands})$, which is three times the cost of storing all higher-temperature residuals.

When the proposal operation (see Appendix A) on the GPU is complete, the odd-numbered residual will be updated with the new sources in $(\text{walker}, \text{temperature}, \text{subband}) = (w, 0, b)$.

VI. LDC2A RESULTS

Our analysis of the LDC2A data is presented below. Overall, our pipeline performed well. It found, extracted, and characterized resolved signals of the expected types; characterized the detector noise and confusion noise foreground; and accomplished all of this without any human intervention. We ran our pipeline for ~ 7 days on 4 NVIDIA A100 GPUs and ~ 12 CPUs.

Figure 5 shows the high level picture of the source extraction and noise analysis. It shows the residual compared to the injection in the A channel (left) and E channel (center), as well as the template waveforms for the resolved sources subtracted to form those residuals (right).

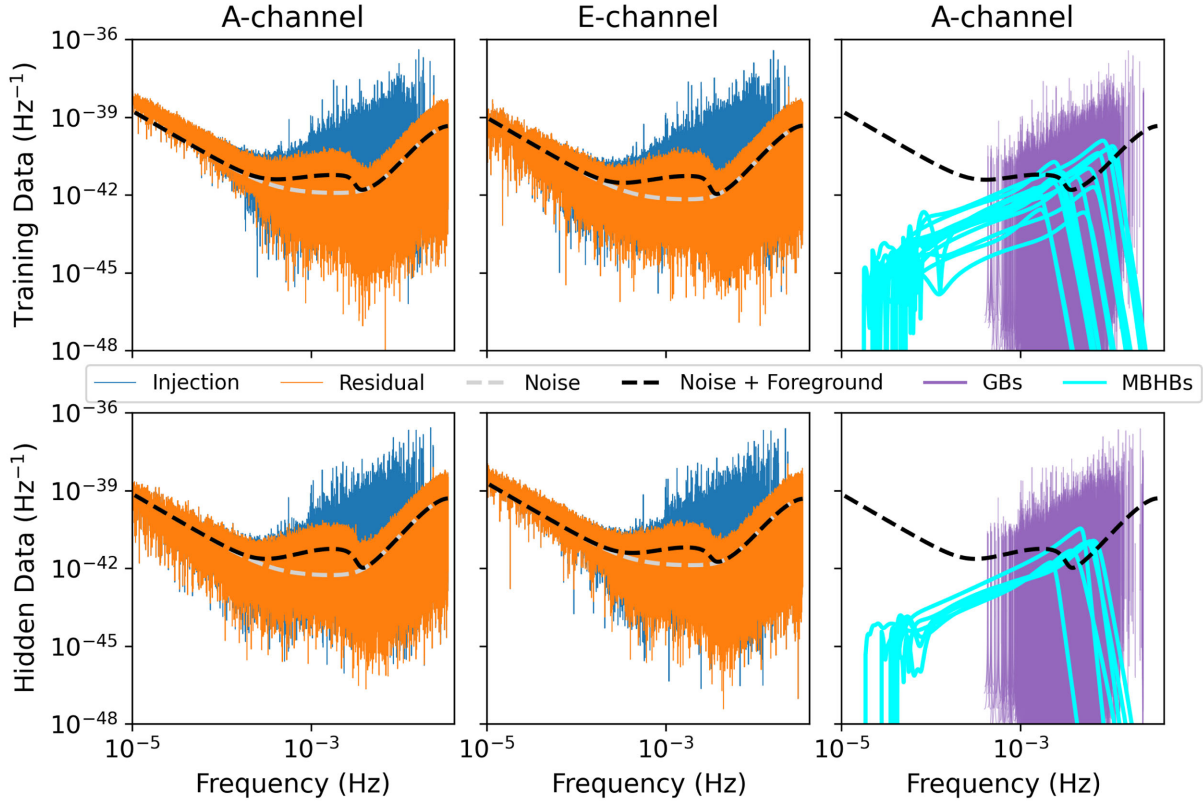


FIG. 5. Comparison of injected (blue) versus residual (orange) data streams. The best fit noise profile is also shown with a black dashed line. The gray dashed line represents only the instrumental noise contribution. The difference between the black and gray dashed lines is the galactic confusion noise. The left and center columns are for the A and E channels, respectively. The column to the right shows all template waveforms subtracted to make the given residual, with GBs shown in purple and MBHBs shown in cyan. The upper row is the training dataset. The bottom row is the hidden dataset.

The training and hidden datasets are shown in the top and bottom rows, respectively.

The residual plot shows that we were able to extract all appreciable signals (note, the line in the residual spectrum at the right of the plot is an artifact of data generation and not a missed GB signal). This is true for both the training and hidden datasets. We also found and extracted all 15 MBHB signals from the training dataset, with no false or missed detections. For the hidden dataset, we extracted all six MBHB signals with no false or missed detections.

The low-frequency noise fit for the A channel (E channel) in the training (hidden) data shows a small inconsistency with the residual data. We believe this is due to two factors. First, the injected noise model contains acceleration noise, optical metrology system noise, and backlink noise (see [23] for more information on these noises). Our fitted model given in Eq. (7) only contains two of those noises: OMS noise and acceleration noise. This causes a mismatch in the noise model at low frequencies. This was due to a previous oversight related to the LDC documentation [34]. Second, our preprocessing operation did not perform the necessary detrending of the time-domain data prior to transforming the data. We have verified that

detrending by removing a linear fit to the data over time removes much of this issue. This does lead to a small increase (~ 100 – 300) in the number of detected binaries and will slightly affect the individual posteriors of various sources. However, we did not rerun our entire pipeline because the observed changes were small enough that the necessary computational expense was not worthwhile.

A. Massive black hole binary posteriors

The posteriors on the MBHBs are shown in Fig. 6 for both the training and hidden datasets. In this figure, we focus on the observed masses and detector-frame merger time information to concisely represent our results. The single-source posterior windows are labeled in order of coalescence time. Full posteriors can be created and analyzed for all sources by anyone using the LISACATTOOLS software package [64] and our final catalogs [69,70].

Both black hole masses lie within the 2σ credible interval for all injected sources, with the exception of two sources (one in each of the training and hidden datasets), for which the true masses were within 3σ credibility. Most source coalescence times show unimodal Gaussian distributions with widths of the order of a few tens of seconds with the

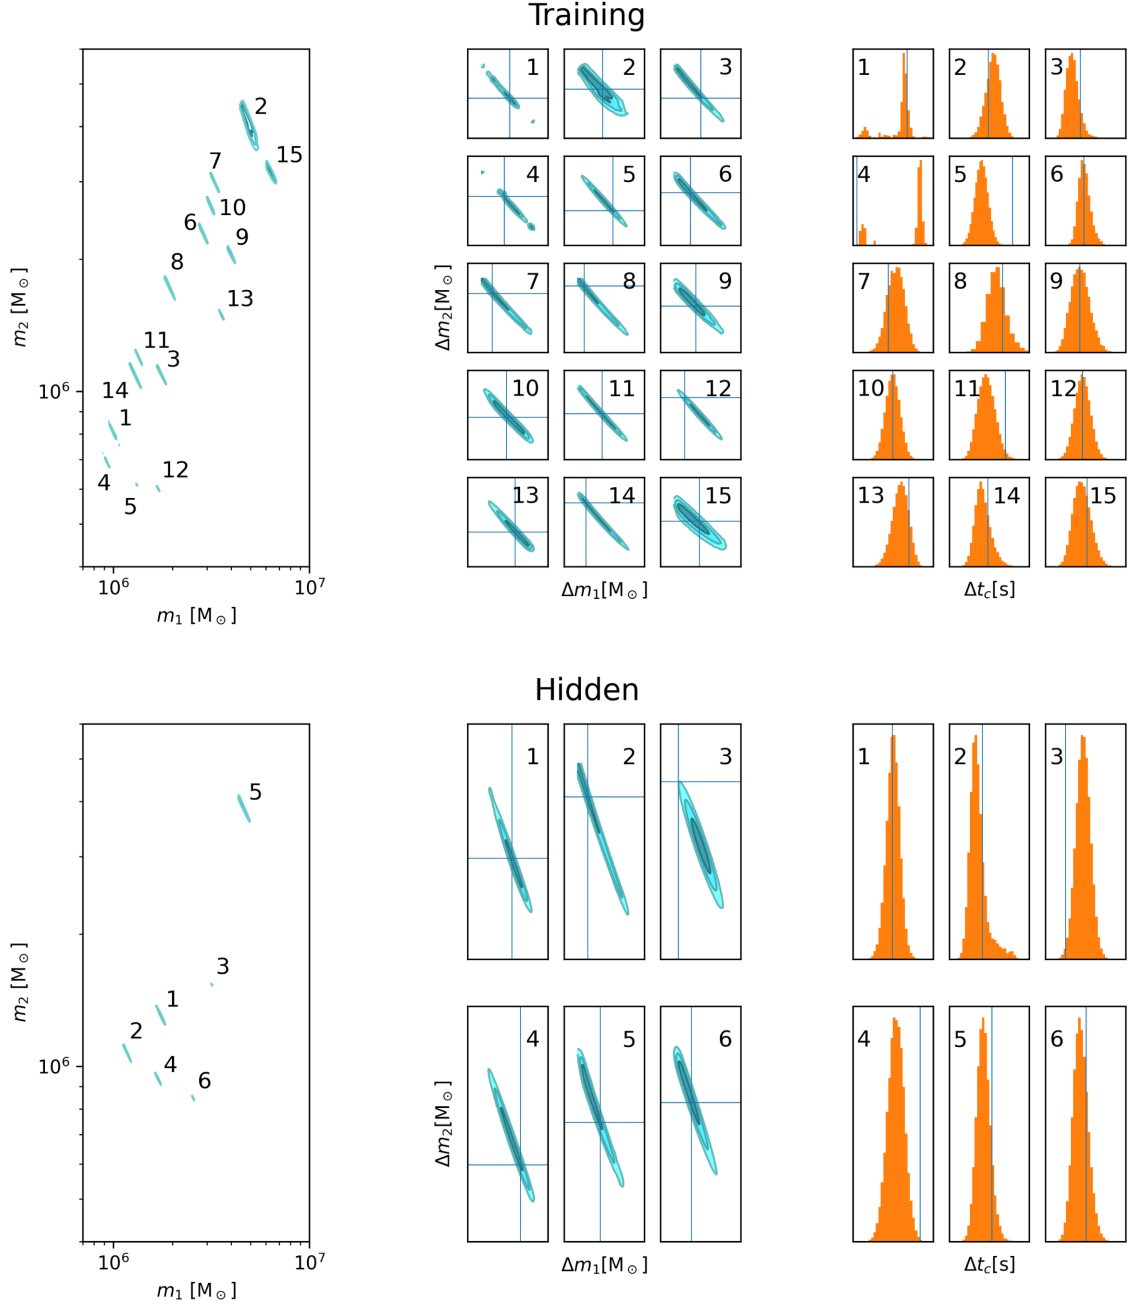


FIG. 6. Posterior distributions for the MBHB sources. The left-hand column shows posteriors on the masses of the binary components, including contours at 1, 2, and 3σ credibility. Sources are numbered in order of their merger time. The middle column shows the same results, but now enlarged and expressed as a difference to the injected value. The injected values are marked by blue lines. The right-hand column shows 1D marginalized posterior distributions for the coalescence time. Injected values are again marked with blue lines. The top/bottom row of plots shows results for the training/hidden dataset.

injected value lying within the posterior bounds. The fifth source in the training dataset contains the true coalescence time far into the tail of its distribution. The first and fourth sources in the training dataset show multimodal posteriors. This is likely caused by the same effect in both cases, but for the first source the larger of the posterior modes contains the true time of coalescence while for the fourth source, it is the smaller of the posterior modes. Similar

results were obtained by other groups analyzing this dataset [e.g. [21]]. It is currently uncertain what causes these additional modes (separated by ~ 10 s) and will be a topic for further analysis in the future.

B. Recovery of galactic binaries

Our algorithm collected samples for and cataloged ~ 12000 galactic binary sources including known

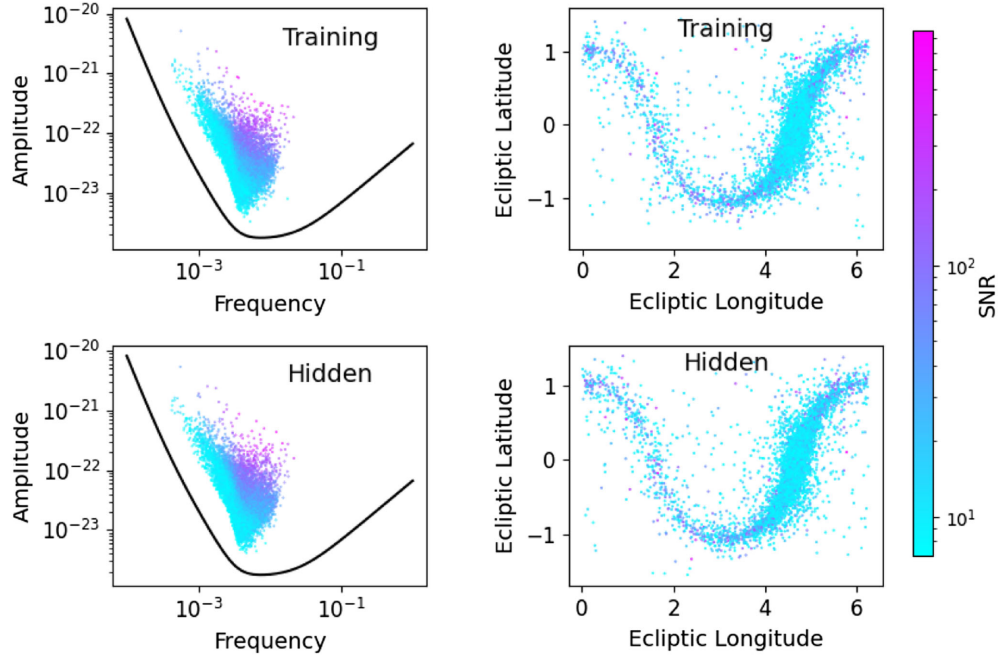


FIG. 7. Point estimates for the cataloged population of GB signals observed in our global fit. The left and right plots show point estimates colored according to their SNR for amplitude versus frequency and ecliptic latitude versus ecliptic longitude, respectively. These point estimates are the median frequency sample for each cataloged source. The training (hidden) dataset is given in the top (bottom) row.

(or verification), detached, and interacting binaries. Approximately 8000 of these were detected with high confidence ($\mathcal{C} > 0.9$). All of these sources are included in the same galaxy sampling module. Known GBs are found and characterized in the same way as unknown GBs. We do not specifically search for them, but we do use them as markers of the performance of the algorithm.

Figure 7 shows information about the detected population of GBs from the training dataset after cataloging. The $\mathcal{A} - f$ plot is color coded according to the SNR of the detection. This plot also shows the expected dearth of sources between the sensitivity curve and the detected population. This difference is the characteristic confusion noise. The sky location point estimates are also given in the same figure. The sources congregate as expected along the galactic plane.

Our comparison to the injection catalog for each dataset is shown in Fig. 8. This plot is similar to Fig. 13 in [21] and the right side of Fig. 5 in [71]. Since the injected catalog technically has $\sim 10^7$ binaries, an *optimal* SNR cut of 7 is applied. This narrows the tested injection set to $\sim 10^4$ injection sources to compare against our cataloged source samples. We take every sample contained for all sources in our catalog and compare them all against each source in the injection set. The maximum overlap achieved per injection source is plotted on the horizontal axis.

The cumulative distribution function of this maximal overlap against each injection source is shown in the top

plots (training on the left and hidden on the right). Their associated survival function is shown in the bottom plots. The curves are also labeled by their confidence, \mathcal{C} , i.e., the fraction of MCMC samples that contained a given source (in the range $[0, 1]$). As a reminder, the choice of a detection statistic in RJMCMC is nuanced and depends a lot on choices made during cataloging, especially on the margins. Understanding and optimizing this is a topic for future discussion.

Our algorithm performed similarly for hidden and training sets with a $>90\%$ match rate for both runs. The confidence lines are quite close together indicating that most of the samples have confidences greater than 0.9.

With greater than 8000 detected sources of high confidence, it is difficult to analyze a large number on an individual basis. Therefore, we try to pick sets of sources that can help to indicate how well the sampler is performing. The three groups we have focused on are high-frequency detached GBs, high-frequency interacting GBs, and the known (or verification) GBs. Groups of eight $f - \dot{f}$ 2D posteriors for each of these three categories are shown in Fig. 9 (for both datasets).

The first group in the top row are the known Verification Binaries (VB). We have chosen the eight highest-frequency VBs with $SNR \gtrsim 15$ arranged in descending order according to their frequency. These sources all exist in regions with some level of confusion. HMCnc is above the main confusion noise frequencies, yet it still suffers from

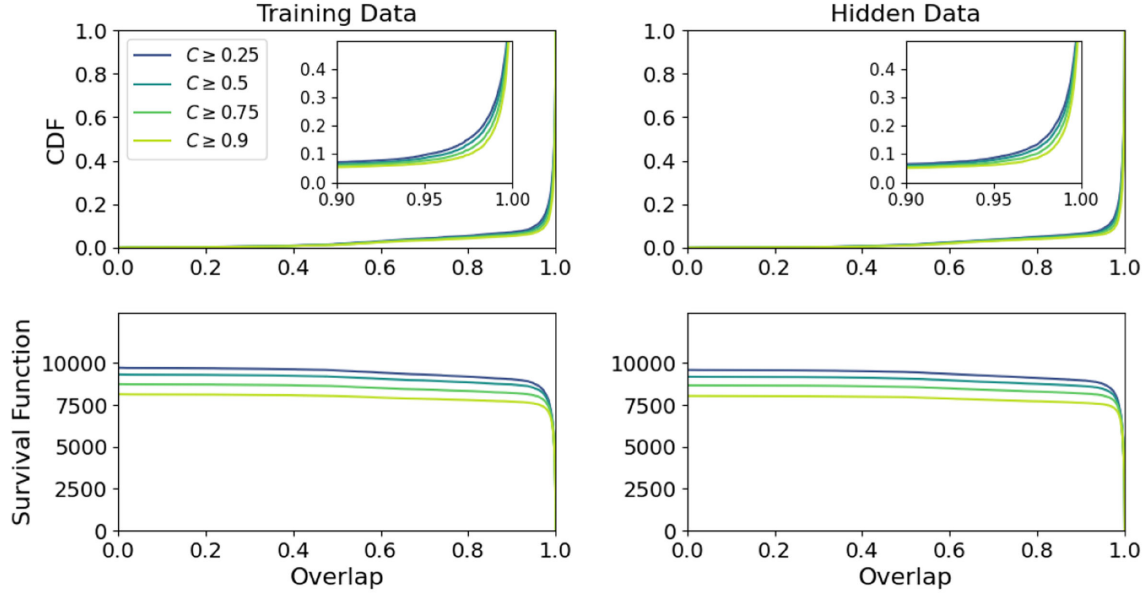


FIG. 8. The cumulative distribution function and survival function for the best match of a source within the GB output catalogs against the injected population. The injected population consists of all sources with optimal SNR > 7 . The lines are separated by their minimum confidence value, C . The confidence is the number of samples in a catalog group divided by the total number of samples. As the confidence value decreases, more sources are included. The training dataset is shown on the left and the hidden dataset is shown on the right.

confusion with nearby detectable sources. The rest of the VGBs are all at frequencies where the confusion noise is expected to play an important role.

In [21], they perform a targeted search for these sources by including direct information from EM observations. While this is a good idea in reality and will most likely be the method used during the final LISA data analysis, we chose to analyze these sources as a part of the larger population without any targeted assumptions. This allows us to test the performance of our algorithm against a specific subset of known sources and compare our posterior distributions to [21] to understand how deliberately targeting the VGBs affects their constraints.

In Fig. 9 we include a horizontal dashed line to indicate $\dot{f} = 0$. For HMCnc and ZTFJ1539, the injected value is well separated from $\dot{f} = 0$ and the posterior excludes this value. This indicates a confident detection, which will provide a stronger constraint on the source chirp mass. As we move to lower frequencies, shown on the right of the figure, the $\dot{f} = 0$ line and the injected value converge until $\dot{f} = 0$ is well contained within the posterior uncertainty. This is an indication that there are weak constraints on \dot{f} .

The middle row of Fig. 9 shows posteriors for the eight highest-frequency detached GBs. As a reminder, “detached” means that the only driving force of the gravitational-wave chirp is general relativity. These sources are very useful to understand how well the sampler’s within-dimensional or in-model moves are working. These sources are found at frequencies where the density of sources is low enough that we do not expect them to overlap. Therefore, we can check

our global fit posteriors against single-source, fixed-dimensional MCMC runs. We have verified that this procedure produces identical posteriors.

GBs at high frequency will also have strong constraints on their \dot{f} parameter. For this reason, none of the plots in the middle row of the figure show a line at $\dot{f} = 0$ because it is beyond the range of the axes.

The final row shows posteriors for the interacting GBs. While these sources have a strong astrophysical basis [e.g. [72]], the actual \dot{f} values injected into this population were assigned *ad hoc* rather than with a specific parametrized model, largely due to the lack of a robust astrophysical prediction for this population. Nonetheless, we were able to constrain \dot{f} values well enough to show at least a few of these sources have negative chirps.

As we move towards the right of the bottom row in Fig. 9, to sources at lower frequency, the constraints weaken, but still show potential detections of a negative chirp, and hence a difference to the predictions of vacuum GR. We are unable to translate this into a meaningful bound on the astrophysical parameters or chirp mass because of the *ad hoc* nature of the injected population. However, this should be pursued with future datasets.

C. Recovery of instrumental noise and confusion foreground

The fitting of the instrumental and confusion foreground noises is unique within the global fit in that the algorithm is fast and efficient, but it is slow to converge because it is

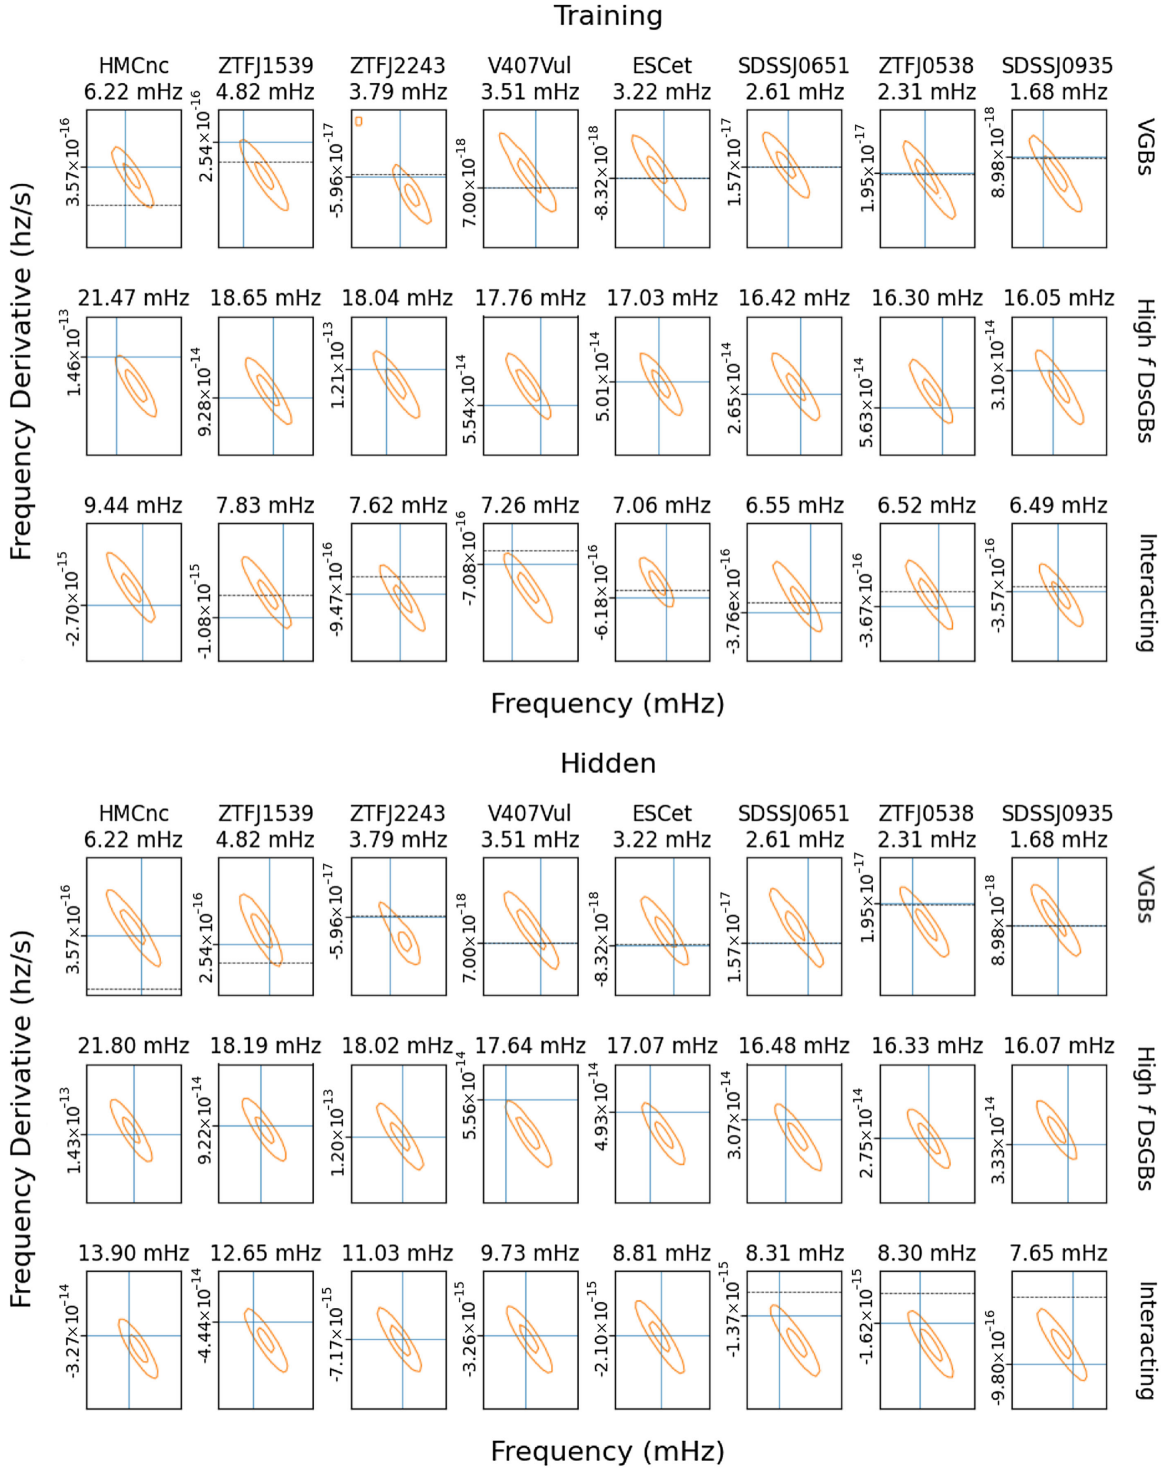


FIG. 9. Marginalized 2D posterior distributions in the $f - \dot{f}$ plane. The top row represents the highest-frequency VGBs arranged in order of decreasing frequency. These VGBs interact with other surrounding sources. The VGB parameters are not inferred using a targeted search as in [21], but simultaneously fitted with the remainder of the population. The middle row shows posteriors for the highest-frequency detached binaries in the randomly generated galactic injection population. These detached sources all occur at frequencies high enough to avoid confusion with other detectable sources. The bottom row shows posteriors for the highest frequency interacting GBs with $\dot{f} \neq \dot{f}_{\text{GW}}$. The orange contours show the 1σ and 2σ contours with the blue solid vertical and horizontal lines representing the injection values from the input population. The injected \dot{f} value is given in units of Hz/s along the vertical axis. The horizontal dashed black line is $\dot{f} = 0$, which gives a sense of the scale and constraint on \dot{f} . The training dataset is the top grid of plots and the hidden dataset is shown in the bottom grid of plots.

highly dependent on the extraction of all detectable GBs. In other words, the noise PSD and confusion foreground information only converge once the galaxy fit has converged. As the GBs are found and removed from the residual, the foreground noise estimate decreases gradually down towards its final state.

As discussed at the beginning of this section, there was a mismatch in the low-frequency noise fitting due to incorrect detrending of the data and using a fitted noise model that did not match the injected model (missing the backlink noise component). This most likely affected the foreground estimation as well, but it is hard to determine by how much. We leave this to future work.

The posteriors on all of the instrumental noise and foreground signal parameters were found to be fairly Gaussian in both runs. The uncertainty in the total PSD was very small, such that when plotting the PSD as a function of frequency, as in Fig. 5, the confidence intervals are smaller than the width of the plotted line. However, individual parameters in the PSD model are less well constrained. We were able to constrain the confusion foreground amplitude to $\sim 5\%$ – 10% . We believe this is larger than expected [e.g. [73]] due to parameter correlations built into our model parametrization clearly seen in the foreground posterior distribution. Samples for the instrumental noise and foreground parameters over the course of the search and parameter estimation runs are included in our output catalogs [69,70].

D. Algorithmic performance

In this section, we will discuss some initial findings about the algorithmic performance of each module of the pipeline. The first three segments of the pipeline—initial instrumental and confusion noise fit, MBHB search, and MBHB plus PSD information mixing—all occur fairly rapidly.

The initial overall noise PSD fitting takes under 30 min. There are some potential ways to improve this timing, but given it is a small fraction of the whole pipeline, only a small amount of initial effort was put into optimizing this segment. The MBHB search, which results in a reasonable posterior estimate of each source, takes \sim hours with roughly eight processes running simultaneously over four GPUs. This could also be improved in various ways, but it is still a small percentage of the overall pipeline run time. The mixing of the MBHBs and PSD fits also takes less than one hour.

The galaxy analysis is the clear driver of the global fit computational cost. The PSD information and MBHBs need time to converge to a posterior, but in the absence of GBs, this would take $\lesssim 1$ day. The ~ 7 days of total runtime is to allow for convergence of the galaxy analysis. The most challenging part about the GB analysis is the scale: with $\sim 10^4$ sources, the overall runtime can be slow, even when the evaluation cost per source is extremely efficient.

Our approximate MCMC generation efficiency is one *proposal* per binary source per $\sim 1 \mu\text{s}$, with actual waveform generation occurring at a rate of one waveform computation per $\sim 0.3 \mu\text{s}$. These timing estimates occur when performing a large number of these computations in parallel as we do in our algorithm.

In different frequency regimes, the algorithm faces different difficulties. At low ($\lesssim 1$ mHz) or high ($\gtrsim 10$ mHz) frequencies, there are fewer resolvable sources which means these regimes generally do not greatly contribute to the overall runtime of the algorithm. At low frequencies, the primary complication is the confusion foreground. At high frequencies, the sources are well separated, which is generally easier to deal with than the confused regimes.

The middle frequency band ($1 \text{ mHz} \lesssim f \lesssim 10 \text{ mHz}$) is the most computationally challenging. The foreground confusion noise is above the instrument PSD at the start of this middle frequency window. It then falls below the PSD around ~ 4 mHz. While the confusion foreground decreases over this range, the confusion between *detectable* sources increases, reaching its peak around ~ 3 – 5 mHz and then decreasing again before sources become more isolated at $\gtrsim 10$ mHz.

VII. DISCUSSION

The LISA global fit is an extremely complex task that is difficult to accomplish. Here we have presented a new global fit pipeline capable of extracting multiple sources of different types, and a fit to a stationary noise profile. Our pipeline is not the only global fit algorithm that has been developed, but it is extremely useful to now have multiple global fit algorithms within the community. Having multiple algorithms will help with research into many open questions pertinent to the planning of the LISA mission. One example is how all the output catalogs produced by the multiple pipelines that will be run as part of the ground segment will be conglomerated into a single output catalog to be released by ESA. This will require making a variety of mixing assumptions and choices between the outputs of the individual global fit algorithms.

As well as a number of minor differences, our algorithm adds three main novelties in comparison to [21]: (i) GPU acceleration, (ii) ensemble sampling, and (iii) the online fitting of RJMCMC proposal distributions.

A. GPU acceleration

The GPU setup has improved computational efficiency relative to [21]. It is difficult to directly compare the timing between the two implementations. Our implementation starts at one whole year of data analyzing it all at once, while [21] builds up the analysis over time, using information from early times for later analyses. This time buildup is a more natural way to design a global fit algorithm given that the mission will be producing data

for several years, and we do not want to wait until the end of data taking to start an analysis. In the future, we will adapt our algorithm to build up over time as well. It is hard to predict how this will affect the timing of the algorithm, so that comparison is left for future work.

If we compare with what we have, $\mathcal{O}(7)$ days on 4 GPUs, to $\mathcal{O}(5)$ days on $\mathcal{O}(1000)$ CPUs, our algorithm is about $\sim 5\times$ more efficient in terms of power usage and cost (multiplying power usage by a flat price per kWh). This is encouraging and we hope it will continue to improve in the future.

When building algorithms for GPUs or CPUs, there is a trade-off between performance, availability, and ease of use. The goal of this code is to have building blocks that work for both machine types (most likely optimized for the GPU), with wrappers that run them in the desired fashion. Being CPU/GPU agnostic will also aid in development.

The GPU development is necessary in our opinion because it is clear that algorithms designed for the GPU will be more energy efficient. This is already true at the present time and, as we move into the future, it is expected that GPUs will advance at a much more rapid pace than CPUs, meaning the efficiency difference between the algorithms will most likely increase further. Our galaxy sampler generally saturates the GPU with computations. Therefore, newer generations of GPUs like the NVIDIA H100 should improve the evaluation time of our algorithm because of the greater number of compute cores available.

Currently, computing facilities with CPU hardware are more widely available than facilities with GPUs. However, the GPU demands of our algorithm are not great, and most new computing clusters that have $\mathcal{O}(1000)$ CPUs also have multiple modern GPUs available.

B. Ensemble sampling

Ensemble sampling is a very important aspect of our algorithm. It allows for two main advantages: multiple cold chains generating samples in a more parallelizable fashion and better mixing between the temperatures due to the random permutation of the walkers when swapping between temperatures. The full ensemble, the total number of walkers in all temperatures, provides a large number of computations for building up subband searches and proposals that feed down to the cold chains.

It is hard to assess or quantify the impact of ensemble sampling and tempering on the efficiency of this sampler. Qualitatively, it seems it can only help. At a minimum, more parallelization is possible since all walkers in all temperatures have likelihood computations that are all independent of one another. The improved parallelization will also strongly benefit as GPU technology and capacity improves.

In the MBHB and PSD sampling modules, we currently do not employ ensemble tempering. This is because it

requires the computation of extra likelihoods that we have currently deemed to be unnecessary. However, in the future this may be useful to improve convergence. We could perform an ensemble tempering move periodically after a specified number of iterations, while maintaining the “vertical” tempering structure during the majority of proposals. This would surely help the mixing. Future work will determine its actual impact on efficiency and sampling quality, and optimize the frequency of vertical and ensemble tempering proposals.

C. Single source search-to-parameter estimation information transfer

Another important piece of our algorithm is how we build up the galaxy population. We start with a single-source MCMC search that gives a posterior estimate. This estimate is then transformed into a proposal with a GMM and sampled directly as an RJMCMC proposal distribution.

This search-to-proposal pipeline was essential to the success of our galaxy search due to the fact that we analyze the entire year of data at once. If we relied on basic RJMCMC proposals, it would take prohibitively long to converge because all the loud sources would be fit with multiple templates at the beginning. In [21], they avoid this issue by building up over time. This time buildup method finds louder sources first when they are the only detectable sources and then adds quieter sources as they become detectable over time.

In addition to the careful addition of sources to build up the galaxy solution, this search-to-proposal method ensures that we will never leave a source of appreciable SNR in the data. It will always locate the best-fit template that remains in the data residual (if there are any left).

This method can be extended beyond the galaxy fit. It is at its core a general method for adding sources into a larger global fit running on an RJMCMC backend. This would allow for much greater flexibility compared to fixed source searches like the MBHB search we currently use. It would be more capable of ensuring sources are removed down to the lowest detectable SNRs that are only found as the full pipeline converges.

Other types of astrophysical sources, such as EMRIs, stellar-origin black hole binaries, and MBHBs, could also be treated using this framework. We can sample them with a single-source search and parameter estimation algorithm to get an initial estimate for the source posterior. We will then fit these posterior estimates with GMMs or normalizing flows [e.g. [74]]. Once they are transferred into the global fit and allowed to fit simultaneously with all other sources, the initial estimated posteriors will be refined. The usefulness of this method cannot be overstated as we add more sources and source types in the future. We expect it to play a key role in ensuring our global fit has the full range of required capabilities.

D. MBHB recovery

We recovered all 15 (6) of the MBHB sources in the LDC2A training (hidden) dataset. This search was done in a rather *ad hoc* manner and could be improved. The search currently uses an SNR cut of 20. However, this SNR is computed against a PSD that is far too high since this computation comes before fitting any GBs. Therefore, this SNR 20 cut corresponds to a much higher optimal SNR.

Having this higher-than-expected effective PSD at the beginning of the MBHB search helps to eliminate spurious sources at the cost of sacrificing our sensitivity to low-SNR MBHBs. In the future, we will have to improve this. One potential solution is to use the same system of adding GBs for MBHBs (as mentioned in the previous section). This will allow for the search and addition of MBHBs with RJMCMC rather than the fixed search we currently employ. This would mean that as the overall noise PSD converges and GBs are subtracted, we will have the ability to add low-SNR MBHBs that cannot be found without near complete GB subtraction.

In our current implementation that uses the entire year of data in one analysis, we split up the MBHB search computation into separate frequency windows and analyze them simultaneously. Within this setup, we have not considered the problem of premerger detection of MBHBs. The inclusion of this feature of the global fitting process is left for future work.

In the future, there are also obvious improvements to be made in terms of the quality and type of MBHB waveforms used in the fit. The LDC2A has just the $l = m = 2$ mode of the PhenomD waveform [52,53]. In the near future, there are plans to re-release the LDC2A dataset with higher-order modes included using the PhenomHM waveform [75]. While still not the most up-to-date or complete waveform available (it lacks features such as precession and eccentricity), it will be slightly close to what will be required for the final global fit algorithm. Although these waveforms are more expensive to generate, higher modes are expected to make the search and parameter estimation of these important sources easier while producing posteriors that are much more constrained in most cases [e.g. [51]]. Therefore, we expect the algorithm will perform as well or better when this waveform improvement is included. The inclusion of precession is also not expected to change this conclusion much, other than adding further cost to the waveform evaluation. Adding other physical effects like eccentricity are likely to have a bigger impact, but this is quite uncertain, in part because eccentric waveforms for MBHB systems have only recently started to become available [e.g. [76]].

E. Galaxy extraction

The extraction of the GB population is the largest and most difficult part of the LISA global fit (at least for currently simulated datasets). Overall, our global fit

algorithm performed very well. It was able to confidently find, characterize, and remove the detectable GB sources. It is more energy efficient than other pipelines, such as [21], and we expect the difference in the efficiency between CPU and GPU algorithms to increase over time as discussed in previous sections.

One aspect of the sampler setup that should be further investigated in the future is the dynamic selection of dimensions within the temperature ladder. Currently, we use dynamic temperature adaptation to adapt the temperatures over the course of the run, but the number of walkers and temperatures for each subband is kept fixed and equal across the full frequency range. Ideally, we would place more resources where they are needed. For example, we should remove temperatures/walkers from edge samplers that do not need as large of a dynamic range due to the low number of templates needed to fit the data within those subbands. We would then use these resources in the center of the band, where the highest source counts and confusion exist, by adding temperatures/walkers. There are a few small planned changes to our algorithm to accomplish this goal.

We also need to examine a larger variety of astrophysical models for the GB population, for example, eccentric, triple, or accreting systems. The detached GBs are expected to make up the majority of the resolved catalog, so the fact that our algorithm performed so well means the largest piece of the fit is in place. Adding other astrophysical GB populations will complicate the fit, but probably only within a subset of the overall galaxy fitting problem. The two most difficult aspects of adding additional astrophysics to the GB population (in terms of data analysis) will be to create accurate, efficient template models and to understand how well those models can be matched by a basic detached GB template. If the match is high, it complicates the RJMCMC process by requiring efficient interchange between the model types. This is a topic for future work.

F. Noise and instrument analysis

For the current assumptions used to construct the LDC2A dataset, our instrumental noise and confusion signal PSD fitting algorithm performed well. It was able to fit both features and produce good posterior information.

The noise model assumed for the LDC2A analysis was purely stationary and Gaussian, including for the foreground. In reality, the foreground noise is cyclostationary. As the LISA constellation orbits the sun, the foreground confusion gains a directionality as the constellation moves toward and away from the galactic center. The fit that we give is effectively an average of this time dependence over the course of the observation. This not only affects our information about the Galaxy, it also affects the sensitivity used for each individual MBHB merger. The foreground noise locally in time to the MBHB merger may be above or

below the average estimated from the full-year fit. This will affect the characterization of these signals [20].

In the future, we will also have to include realistic orbits and instrumental effects in our global fit. This means using numerically generated orbits to construct the LISA response and include fitting for nonstationary noise from sensitivity drift, gaps, glitches, etc. in the global fit. Additionally, we will move to a more instrument-centric PSD parametrization to both input useful prior information into the global fit and to extract information about the instrument performance from it.

VIII. CONCLUSION

We have presented a new global fit algorithm capable of successfully analyzing the LDC2A training and hidden datasets from start to finish with no human intervention. We extracted and characterized all MBHB signals in both datasets without any false alarms. Our galaxy sampler extracted $\sim 10^4$ GBs with $>90\%$ match rate to the input population. Along with these astrophysical sources, we were able to fit the PSD and foreground confusion noise.

The most important advancements in this pipeline are in the galaxy sampler because of its large computational burden: (i) GPU-accelerated sampling operations; (ii) ensemble sampling and tempering; and (iii) online fitting of RJMCMC distributions via refitting of samples and single-source MCMC runs against the current residual. The GPU-accelerated pipeline is more efficient in energy usage and cost than comparable CPU-based pipelines, and is expected to continue to improve into the future. The ensemble sampling allows for better mixing and better marginalization between source and noise modules across the global fit. The refitting and single-source MCMC pipelines used to generate the RJMCMC proposal distributions are very helpful to building up the galaxy solution. We expect these methods to extend to other sources as more advanced pipelines are constructed. We have also published our catalogs for both training and hidden LDC2A datasets [69,70]. Our new open-source algorithm and its catalog outputs now available to the public will be critical tools for encouraging access to and advancement of LISA data and astrophysical analysis across the global LISA community.

ACKNOWLEDGMENTS

The authors would like to thank Tyson Littenberg for helpful discussions in finishing this initial implementation. We would also like to thank Christian Chapman-Bird, Lorenzo Speri, Ollie Burke, Stanislav Babak, Sylvain Marsat, Robert Rosati, Alvin Chua, amongst others for helpful discussions along the way. N.K. acknowledges support from the CNES for the exploration of LISA science. N.K. acknowledges the funding from the European Union’s Horizon 2020 research and innovation

program under the Marie Skłodowska-Curie Grant Agreement No. 101065596. This code made use of NumPy [77], Matplotlib [78], SciPy [79], and CuPy [80].

DATA AVAILABILITY

The analysis codes (and the specific commits) used were LISA analysis tools [81,82], Eryn [45,83], BBHx [61,84], and GBGPU [85,86]. We plan to clean up the code and add documentation and tutorials in the coming months. At this point, the full code will be released and made readily available to all. Our output catalogs are available for download on Zenodo. The training dataset catalogs [69] can be found here [87]. The hidden dataset catalogs [70] can be found here [88]. These data can be read in and analyzed with LISACATTOOLS [64] and LISA analysis tools [81].

APPENDIX A: GB PROPOSAL ALGORITHM

The GB proposal algorithm builds on the special GPU-based waveforms to handle the sampling and likelihood computations in an efficient and light-weight manner. Algorithm 1 describes the GB sampling implementation.

Due to complications with changing the dimensionality of the sorted and grouped GBs used in the proposal, the RJ proposals are currently run a bit differently from in-model Proposals. They are implemented in exactly the same way as in-model proposals, but the inputs are handled differently. For in-model proposals, the initial points and their associated information are organized and entered into the GPU kernel. It then runs multiple proposals for each binary and reads out the final information.

As discussed in Appendix C, RJ proposals are generated in a batched way across the whole band by proposing to remove all sources that are currently there and proposing to add a generated source for each array entry currently absent of a source. From an algorithmic perspective, we propose all points for the unused array entries upfront. This is true for all walkers, all temperatures, all unused array entries. We organize and group them with the sources that are in the sampler currently. This specialized grouping information is then input into the GPU kernel with two extra pieces of information: whether each binary is already in the sampler or not, as well as the PDF value for all sources (both existing and proposed) as determined from the generated distribution.

The GPU kernel then runs through all of these existing and proposed sources in randomized order within the same subband. We treat this like an in-model proposal simulating an RJ step by making the source of interest almost identical before and after the proposal, adjusting only its amplitude. If the source exists in the sampler, meaning it is a proposed removal, we set the initial amplitude to the true value and its proposed amplitude to a negligibly small number. If the source of interest is a proposed addition, we set the final

amplitude to the true value and the initial amplitude to a negligibly small number. We accept or reject the proposed change after each individual proposed binary addition or removal. All sources and source parameters, other than for the specific source for which we propose a change, remain fixed throughout this sampling process.

APPENDIX B: SETTING UP IN-MODEL GB PROPOSALS

In-model proposals are MCMC proposals that do not change the dimensionality. In our case here, this is the underlying model count. As discussed in [44], the regular stretch proposal does not scale well to reversible jump situations where the dimensionality varies. For this reason, we implement a “group stretch proposal” as detailed in [44]. The group stretch proposal is a generalization of the stretch proposal to reversible-jump situations with varying dimensionality and multiple intertwined maxima. The main idea of the group proposal is to set a stationary group of points from all leaves and walkers in the sampler from which to draw the auxiliary point needed in the stretch proposal. This point is drawn from a distribution of points in the stationary group that are “close” to the current point of interest from which the proposal originates. Here, “close” means as determined by some user-defined distance metric.

For the GB sampler, the initial frequency parameter is the most obvious parameter on which to base a distance metric to use to cluster points into nearby distributions. The GB group proposal begins by selecting the n_{friends} points out of the previously stored stationary distribution that are closest

in initial frequency parameter. It then draws uniformly from that group of nearby points to use in the stretch proposal.

In the limit where every walker contains every template and no template’s posterior distribution interacts closely with another source, the group proposal will converge to the performance of the base stretch proposal. However, in realistic situations, the efficiency of the group proposal is highly situation dependent.

We update the stationary distribution every ~ 30 iterations. One downside to this setup is it does not work well when moving a template that appears in a low fraction of walkers (e.g., a source that is not real or a source that is marginally detectable). In this case, the moves are rarely accepted. While recognizing this shortcoming, based on empirical runs, we decided to use only this group proposal rather than dealing with the tuning and complications of implementing a covariance-based proposal. In the future, we plan to add a covariance proposal to help with these cases.

APPENDIX C: RJ GB PROPOSAL BASIS

When running the RJ proposals, we do not use the typical “birth-death” model where one source is either added or subtracted for each RJ proposal. We set an array of $n_{\text{bin,max}}$ boolean values for each walker. True values in the array indicate a binary exists. False values indicate there is no binary at that array entry. The value of $n_{\text{bin,max}}$ is chosen to be sufficiently high so as not to affect the outcome of the experiment. During *each* RJ proposal, we propose to switch every value of this array. We propose to remove all binaries that exist and add a binary for each value in our array that

ALGORITHM 1. Galaxy sampling algorithm.

$j \leftarrow (t, w, b)$ $\mathbf{r}' \leftarrow \mathbf{r}_{(0,w,b-1)}^{\text{even}} + \mathbf{r}_{(0,w,b)}^{\text{even}} + \mathbf{r}_{(0,w,b+1)}^{\text{even}}$ $\mathbf{r}' \leftarrow \mathbf{r}' - \sum_{\theta \in \theta_j} \mathbf{h}(\theta)$ $\theta_{j,0} \leftarrow \theta_j$ $p \leftarrow 0$ while $p < n_{\text{prop}}$ do $k \leftarrow \text{randint}(0, g_j)$ $\theta_k \leftarrow \theta_j[k]$ $\theta'_k \leftarrow \text{Make Proposal on } \theta$ if accepted then $\theta_j[k] \leftarrow \theta'_k$ $\mathbf{r}' \leftarrow \mathbf{r}' + \mathbf{h}(\theta_k)$ $\mathbf{r}' \leftarrow \mathbf{r}' - \mathbf{h}(\theta'_k)$ end if $p \leftarrow p + 1$ end while if $t = 0$ then $\mathbf{r}_{(0,w,b)} \leftarrow \mathbf{r}_{(0,w,b)} + \sum_{\theta \in \theta_{j,0}} \mathbf{h}(\theta)$ $\mathbf{r}_{(0,w,b)} \leftarrow \mathbf{r}_{(0,w,b)} - \sum_{\theta \in \theta_j} \mathbf{h}(\theta)$ end if	\triangleright All information contained for band (t, w, b) . Assume b is odd. \triangleright Applies to frequencies contained in bands $b - 1$, b , and $b + 1$. \triangleright Subtract current sources in band b , walker w , and temperature t . $\triangleright g_j$ is the number of GBs in band $j = (t, w, b)$. \triangleright Includes accept/reject and likelihood computations described in Eq. (5). \triangleright Remove $\mathbf{h}(\theta_k)$ from \mathbf{r}' . \triangleright Add $\mathbf{h}(\theta'_k)$ to \mathbf{r}' . \triangleright Remove starting waveforms from residual in global memory. \triangleright Add new group waveforms to residual in global memory.
---	--

ALGORITHM 2. Galaxy tempering algorithm.

```

j  $\leftarrow$  (t, p, b)
Current likelihood array:  $\text{len}(\mathcal{L}_c) = n_t$ 
Swapped likelihood array:  $\text{len}(\mathcal{L}_s) = n_t$ 
t  $\leftarrow$  0
while t <  $n_t$  do
  w  $\leftarrow$  p[t]
  j  $\leftarrow$  (t, w, b)
  r'  $\leftarrow$  r(0,w,b-1) + r(0,w,b+1) -  $\sum_{\theta \in \theta_j} \mathbf{h}(\theta)$ 
   $\mathcal{L}_c[\mathbf{t}] = \mathcal{L}(\theta_j)$ 
  t  $\leftarrow$  t + 1
end while
t  $\leftarrow$   $n_t - 1$ 
while t > 0 do
  wt  $\leftarrow$  p[t]
  wt-1  $\leftarrow$  p[t - 1]
  r'  $\leftarrow$  r(0,wt,b-1) + r(0,wt,b+1) -  $\sum_{\theta \in \theta_{(t-1,w_{t-1},b)}} \mathbf{h}(\theta)$ 
   $\mathcal{L}_p[\mathbf{t}] \leftarrow \mathcal{L}(\theta_{(t-1,w_{t-1},b)})$ 
  r'  $\leftarrow$  r(0,wt-1,b-1) + r(0,wt-1,b+1) -  $\sum_{\theta \in \theta_{(t,w_t,b)}} \mathbf{h}(\theta)$ 
   $\mathcal{L}_p[\mathbf{t} - 1] \leftarrow \mathcal{L}(\theta_{(t,w_t,b)})$ 
  Accept or reject swap with Eq. (D2).
  if swap accepted then
     $\theta' \leftarrow \theta_{(t,w_t,b)}$ 
     $\theta_{(t,w_t,b)} \leftarrow \theta_{(t-1,w_{t-1},b)}$ 
     $\theta_{(t-1,w_{t-1},b)} \leftarrow \theta'$ 
     $\mathcal{L}_c[\mathbf{t}] \leftarrow \mathcal{L}_p[\mathbf{t}]$ 
     $\mathcal{L}_c[\mathbf{t} - 1] \leftarrow \mathcal{L}_p[\mathbf{t} - 1]$ 
    if t = 1 then
       $\mathbf{r}_{(0,w_0,b)} \leftarrow \mathbf{r}_{(0,w_0,b)} + \sum_{\theta \in \theta_{(1,w_1,b)}} \mathbf{h}(\theta)$ 
       $\mathbf{r}_{(0,w_0,b)} \leftarrow \mathbf{r}_{(0,w_0,b)} - \sum_{\theta \in \theta_{(0,w_0,b)}} \mathbf{h}(\theta)$ 
    end if
  end if
  t  $\leftarrow$  t - 1
end while

```

\triangleright **t** is the array over the temperature indices (range (n_t)).

\triangleright **p** is the array of permuted walkers with $\text{len}(\mathbf{p}) = n_t$.

\triangleright **w** is the walker index associated with temperature **t**.

\triangleright Store all current likelihood values.

\triangleright Residual in w_t with sources from band $(t - 1, w_{t-1}, b)$.

\triangleright Use **r'** in likelihood computation.

\triangleright Residual in w_{t-1} with sources from band (t, w_t, b) .

\triangleright Use updated **r'** in likelihood computation.

\triangleright New cold-chain parameters ($t = 1, t - 1 = 0$).

\triangleright Remove old cold-chain sources from global residual.

\triangleright Add new cold-chain sources to global residual.

begins as “false.” We randomize the order of these additions and subtractions and run through them one-by-one using Algorithm 1. Theoretically speaking, this method is equally valid compared to the single birth-death setup usually used for nested RJ problems. Additionally, our RJ proposal distributions span the entire space, which we believe creates more flexibility and power in proposal generation.

APPENDIX D: PARALLEL TEMPERING IN A RESIDUAL-BASED ANALYSIS

We leverage parallel tempering in every module in our global fit sampler. Parallel tempering “heats up” the likelihood surface by raising it to the power of the inverse temperature: $\mathcal{L}_\beta = \mathcal{L}^\beta$, with $\beta = 1/T$. Samples are then probabilistically swapped between temperature chains (see below). This allows for the algorithm to explore the full likelihood surface using the higher temperature chains to move between likelihood maxima. As the sampling proceeds, the temperatures across the chains adapt to

optimize swapping percentages according to the prescription given in [60]. You can find more information about parallel tempering and its implementation in Eryn in [44].

In all of the sampling discussed below, we leverage the ensemble nature of Eryn which allows for many walkers per temperature. The ensemble setup creates better tempered mixing because the order of walkers within a temperature is randomly permuted before proposing temperature swaps. This means that, from one step to the next, sets of walkers for proposed swaps will always be newly randomly generated.

The typical expression for the acceptance, $\alpha_{T,1}$, of temperature swaps between the parameters, $\vec{\theta}_n$ and $\vec{\theta}_m$, of walkers at temperatures $1/T_i = \beta_i$ and $1/T_j = \beta_j$ is given by [60],

$$\alpha_{T,1} = \min \left[1, \left(\frac{\mathcal{L}_n}{\mathcal{L}_m} \right)^{\beta_j - \beta_i} \right], \quad (\text{D1})$$

where \mathcal{L}_k denotes the cold-chain likelihood evaluated at the parameters $\vec{\theta}_k$.

In our global fit algorithm, each source type and the overall noise has its own temperature ladder setup, only sharing cold-chain information with each other. We use the ensemble setup within each source or noise module to marginalize over different instances of residuals (GBs plus MBHBs) and total noise PSDs (total means including the stochastic confusion noise) that come from the posterior distributions of other modules. For example, if we have n_w walkers in each temperature, we may choose to have n_w different realizations of the PSD or n_w different residuals from the MBHB analysis in the GB part of the sampler (or vice versa).

Having residuals or PSDs attached to a specific walker complicates the swapping of samples between rungs on the temperature ladder. Equation (D1) assumes that the likelihoods for the two states being swapped are equivalent. This is generally not the case if the two walkers have different residuals or PSDs. Suppose that we are proposing a swap between a walker, labeled i , in the chain at inverse temperature β_I , currently at parameter values $\vec{\theta}_n$, with a walker, j , at temperature β_{II} , at parameter values $\vec{\theta}_m$. We denote the cold-chain likelihood for walker i , evaluated at $\vec{\theta}_n$ by $\mathcal{L}_{i,n}$. When deciding whether to swap the parameter values, $\vec{\theta}_n$ and $\vec{\theta}_m$, of the two walkers, we must account for the difference between the likelihood functions associated with the walkers. The correct acceptance fraction, $\alpha_{T,2}$, is given by

$$\alpha_{T,2} = \min \left[1, \frac{(\mathcal{L}_{i,m})^{\beta_I} (\mathcal{L}_{j,n})^{\beta_{II}}}{(\mathcal{L}_{i,n})^{\beta_I} (\mathcal{L}_{j,m})^{\beta_{II}}} \right]. \quad (\text{D2})$$

Notice that in the case where the two walkers are sampling the same likelihood we have $(\mathcal{L}_{i,n} = \mathcal{L}_{j,n} = \mathcal{L}_n$ and $\mathcal{L}_{j,m} = \mathcal{L}_{i,m} = \mathcal{L}_m)$, Eq. (D2) reduces to Eq. (D1) as expected.

We use independent temperature ladders that adapt on their own. This type of issue has important ramifications for the efficiency of the sampler. It effectively means that for every swap proposed, two new likelihoods need to be computed in addition to the original likelihoods that enter the swapping operation. We will discuss how this applies differently to each module of the global fit in more detail below and in Algorithms 1–3. We use independent temperature ladders that adapt on their own for PSD, each different MBHB, and each subband in our GB sampler (see below). We communicate information between these ladders only through the cold chain ($T = 1$).

APPENDIX E: GB TEMPERING ALGORITHM

The tempering algorithm is very similar to the proposal algorithm with a few important changes to the information

passed and how it is passed. This includes the ordering of the subbands to ensure permutation of walkers during temperature swaps. The tempering scheme is shown in Algorithm 2.

APPENDIX F: GB WAVEFORM GENERATION

Besides the source count uncertainty, the most difficult aspect of the GB algorithm is the large scale needed. With $\sim 10^4$ binaries expected from ~ 1 year of observation, even when using the most efficient waveform generators, the sheer number of waveforms needed throughout the MCMC run can be computationally difficult or prohibitive. To address this, we have developed an extremely efficient waveform generation pipeline.

Ultimately, we use functions based on GBGPU [85] for generating GB waveforms efficiently. However, we use a special implementation that heavily leverages the usage of GPU blocks and shared memory to make the waveform generation as efficient and light weight as possible. For this task, we implement the waveforms in C/CUDA. We use a special FFT CUDA package [89] to perform FFTs of the slow part of the waveform within GPU shared memory. Therefore, we never read out the waveforms to global memory, which ensures maximal efficiency for our computations.

This waveform generation method can build the most common waveform in our analysis (256 frequency bins) at a rate of one waveform per ~ 200 ns when properly batched. This is roughly ~ 1000 times faster than those generated on a single CPU. From initial estimates, this translates to a power usage efficiency improvement of $\sim 70\times$.

APPENDIX G: ALGORITHM FOR FITTING PROPOSAL DISTRIBUTIONS

While the other modules are running on their own GPUs, one extra process and one or two extra GPUs are used to search for and continuously improve upon the proposal distributions for the GBs. The cost of adding this GPU operation needs to be weighed against its helpful increase in efficiency of the overall process. We believe this step is incredibly helpful for getting the sampler to converge more rapidly.

There are two types of distribution fitting performed in the GB proposal fitting module. The first is the fitting of the distributions generated from searching for new binaries in each subband. The algorithm for this is a customized, batched, parallelized fixed-dimensional MCMC code based on Eryn. It runs with ~ 10 temperatures and ~ 100 walkers per temperature using the stretch proposal to find the maximum likelihood source that still exists in a given subband (meaning it is yet to be found and subtracted from the residual data).

This algorithm runs simultaneous, independent MCMC runs for each subband (~ 900). The priors are modified so that the initial frequency spans its assigned subband only.

Therefore, in this setup, the sources cannot move across subband edges.

The starting points for each independent MCMC are generated from these modified priors. A user-defined threshold, $n_{\text{it,max}}$, is set to determine the number of iterations required for maximum likelihood convergence. This means the sampler must not find (within each specific subband) a higher likelihood value for $n_{\text{it,max}}$ consecutive iterations. For the results presented here we use $n_{\text{it,max}} = 500$.

Once this convergence criterion is reached, the best likelihood point is taken for each subband. A new “ball” of points around this max likelihood point is generated. The MCMCs are then relaunched with the same settings and stopping criterion. After this converges, we can be sure we have found the maximum likelihood point available in each subband. We then do a test of each source against an optimal SNR limit of 7 to start. We eventually lower this to 5 over the course of the run (see Sec. V F for more information). The MCMC is then run again for each source in each subband with an SNR above the limit. We generate 30 samples (after thinning by a factor of 25) for each of the 100 walkers giving us 3000 samples per source.

When all of these sample groups have been generated, we use mpi4py [65–68] to spread out this process to multiple processors which use Scikit-learn [63] to fit the GMM distributions. Ideally, this operation would be performed in batches on a GPU. This is a topic for future work.

We also fit the points that have been in the large-scale sampler recently. This includes all sources in the cold chain that are appreciably detectable (optimal SNR $\gtrsim 7$). We group these sources based on our cataloging operation described in Sec. V G. After they have been grouped, we perform the same type of GMM fitting described above for the search proposal distributions. As mentioned previously, the purpose of this proposal is to help the mixing of the higher temperature chains by providing them with more direct proposals for binaries that are at higher SNR.

In the future, we also plan to include GB population distributions fit with normalizing flows [74]. We will investigate using these fitted distributions as our prior and/or proposal distribution functions.

APPENDIX H: MBHB PROPOSAL ALGORITHM

The MBHB proposal algorithm is given in Algorithm 3.

ALGORITHM 3. MBHB sampling algorithm within the global fit.

```

j ← 0
while j < n_steps do
   $\vec{n} \leftarrow \text{permutation}(\text{range}(0, n_{\text{mbhb}}))$  ▷ Randomly order the MBHBs each time.
  i ← 0
  while i < len(n) do
    l ← n[i]
     $\mathbf{h}_{(0,\mathbf{w},l)} \leftarrow f(\theta_{(0,\mathbf{w},l)})$  ▷ w represents all walker indexes here. The “0” is the temperature index.
     $\mathbf{r}'_{\mathbf{w}} \leftarrow \mathbf{r}_{\mathbf{w}} + \mathbf{h}_{(0,\mathbf{w},l)}$  ▷ Remove current cold-chain waveforms from residuals ( $\mathbf{d}_{\mathbf{w}} - \mathbf{H}_{\mathbf{w}}$ ).
     $w_{\text{ref}} \leftarrow \text{argmax}_{\theta_{(0,\mathbf{w},l)}} [\mathcal{L}(\theta_{(0,\mathbf{w},l)})]$  ▷ Use maximum likelihood point as reference waveform for heterodyning.
     $\mathbf{h}_{\text{ref}} \leftarrow \mathbf{h}_{(0,w_{\text{ref}},l)}$ 
    p ← 0
    while p < n_prop do
      Make proposal step for all (t, w, l) ▷ Use heterodyned likelihood with residual  $\mathbf{r}'_{\mathbf{w}}$  and reference waveform  $\mathbf{h}_{\text{ref}}$ .
      p ← p + 1
    end while
     $\mathbf{h}'_{(0,\mathbf{w},l)} \leftarrow f(\theta_{(0,\mathbf{w},l)})$  ▷ Add the new cold-chain waveforms back into the residual ( $\mathbf{d}_{\mathbf{w}} - \mathbf{H}_{\mathbf{w}}$ ).
     $\mathbf{r}_{\mathbf{w}} \leftarrow \mathbf{r}'_{\mathbf{w}} - \mathbf{h}'_{(0,\mathbf{w},l)}$ 
    i ← i + 1
  end while
  FILE ←  $\theta_{(\mathbf{t},\mathbf{w},1)}$ 
  HEAD PROCESS ←  $\theta_{(0,\mathbf{w},1)}$ 
   $\mathbf{r}_{\mathbf{w}} \leftarrow \mathbf{r}_{\mathbf{w},\text{new}}$  ▷ Residual update from HEAD PROCESS.
  j ← j + 1
end while

```

- [1] P. Amaro-Seoane *et al.* (LISA Collaboration), Laser interferometer space antenna, [arXiv:1702.00786](#).
- [2] M. Colpi *et al.*, LISA definition study report, [arXiv:2402.07571](#).
- [3] R. Abbott *et al.* (KAGRA, Virgo, and LIGO Scientific Collaborations), GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run, *Phys. Rev. X* **13**, 041039 (2023).
- [4] G. Agazie *et al.* (NANOGrav Collaboration), The NANOGrav 15 yr data set: Evidence for a gravitational-wave background, *Astrophys. J. Lett.* **951**, L8 (2023).
- [5] J. Antoniadis *et al.* (EPTA, InPTA: Collaborations), The second data release from the European Pulsar Timing Array—III. Search for gravitational wave signals, *Astron. Astrophys.* **678**, A50 (2023).
- [6] H. Xu *et al.*, Searching for the nano-hertz stochastic gravitational wave background with the Chinese Pulsar Timing Array Data Release I, *Res. Astron. Astrophys.* **23**, 075024 (2023).
- [7] A. Zic *et al.*, The Parkes Pulsar Timing Array third data release, *Pub. Astron. Soc. Aust.* **40**, e049 (2023).
- [8] P. A. Seoane *et al.* (LISA Collaboration), Astrophysics with the laser interferometer space antenna, *Living Rev. Relativity* **26**, 2 (2023).
- [9] T. Kupfer *et al.*, LISA Galactic binaries with astrometry from Gaia DR3, *Astrophys. J.* **963**, 100 (2024).
- [10] M. A. Keim, V. Korol, and E. M. Rossi, The large magellanic cloud revealed in gravitational waves with LISA, *Mon. Not. R. Astron. Soc.* **521**, 1088 (2023).
- [11] S. Rieck, A. W. Criswell, V. Korol, M. A. Keim, M. Bloom, and V. Mandic, LISA detectability of a stochastic gravitational wave background from unresolved white dwarf binaries in the LMC, *Mon. Not. R. Astron. Soc.* **531**, 2642 (2024).
- [12] M. C. Begelman, R. D. Blandford, and M. J. Rees, Massive black hole binaries in active galactic nuclei, *Nature (London)* **287**, 307 (1980).
- [13] M. Volonteri, M. Habouzit, and M. Colpi, The origins of massive black holes, *Nat. Rev. Phys.* **3**, 732 (2021).
- [14] A. Askar, M. B. Davies, and R. P. Church, Formation of supermassive black holes in galactic nuclei—I. Delivering seed intermediate-mass black holes in massive stellar clusters, *Mon. Not. R. Astron. Soc.* **502**, 2682 (2021).
- [15] M. Bonetti, A. Sesana, F. Haardt, E. Barausse, and M. Colpi, Post-Newtonian evolution of massive black hole triplets in galactic nuclei—IV. Implications for LISA, *Mon. Not. R. Astron. Soc.* **486**, 4044 (2019).
- [16] 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, Science with the space-based interferometer LISA. V: Extreme mass-ratio inspirals, *Phys. Rev. D* **95**, 103012 (2017).
- [17] Q. Baghi, I. Thorpe, J. Slutsky, J. Baker, T. Dal Canton, N. Korsakova, and N. Karnesis, Gravitational-wave parameter estimation with gaps in LISA: A Bayesian data augmentation method, *Phys. Rev. D* **100**, 022003 (2019).
- [18] Q. Baghi, N. Korsakova, J. Slutsky, E. Castelli, N. Karnesis, and J.-B. Bayle, Detection and characterization of instrumental transients in LISA Pathfinder and their projection to LISA, *Phys. Rev. D* **105**, 042002 (2022).
- [19] A. Spadaro, R. Buscicchio, D. Vetrugno, A. Klein, D. Gerosa, S. Vitale, R. Dolesi, W. J. Weber, and M. Colpi, Glitch systematics on the observation of massive black-hole binaries with LISA, *Phys. Rev. D* **108**, 123029 (2023).
- [20] M. C. Digman and N. J. Cornish, LISA gravitational wave sources in a time-varying galactic stochastic background, *Astrophys. J.* **940**, 10 (2022).
- [21] T. B. Littenberg and N. J. Cornish, Prototype global analysis of LISA data with multiple source types, *Phys. Rev. D* **107**, 063004 (2023).
- [22] M. L. Katz, J.-B. Bayle, A. J. K. Chua, and M. Vallisneri, Assessing the data-analysis impact of LISA orbit approximations using a GPU-accelerated response model, *Phys. Rev. D* **106**, 103001 (2022).
- [23] J.-B. Bayle, O. Hartwig, M. Lilley, A. Hees, C. Chapman-Bird, G. Woan, and P. Wolf, End-to-end simulation and analysis pipeline for LISA, in *57th Rencontres de Moriond on Gravitation* (2023); [arXiv:2305.09702](#).
- [24] J. Page and T. B. Littenberg, Bayesian time delay interferometry, *Phys. Rev. D* **104**, 084037 (2021).
- [25] J. Page and T. B. Littenberg, Bayesian time delay interferometry for orbiting LISA: Accounting for the time dependence of spacecraft separations, *Phys. Rev. D* **108**, 044065 (2023).
- [26] E. Savalle, J. Gair, L. Speri, and S. Babak, Assessing the impact of instrumental calibration uncertainty on LISA science, *Phys. Rev. D* **106**, 022003 (2022).
- [27] S. Paczkowski, R. Giusteri, M. Hewitson, N. Karnesis, E. D. Fitzsimons, G. Wanner, and G. Heinzel, Postprocessing subtraction of tilt-to-length noise in LISA, *Phys. Rev. D* **106**, 042005 (2022).
- [28] M.-S. Hartig and G. Wanner, Tilt-to-length coupling in LISA Pathfinder: Analytical modeling, *Phys. Rev. D* **108**, 022008 (2023).
- [29] G. Wanner, S. Shah, M. Staab, H. Wegener, and S. Paczkowski, In-depth modeling of tilt-to-length coupling in LISA's interferometers and TDI Michelson observables, *Phys. Rev. D* **110**, 022003 (2024).
- [30] M. Tinto and S. V. Dhurandhar, Time-delay interferometry, *Living Rev. Relativity* **24**, 1 (2021).
- [31] O. Hartwig, J.-B. Bayle, M. Staab, A. Hees, M. Lilley, and P. Wolf, Time-delay interferometry without clock synchronization, *Phys. Rev. D* **105**, 122008 (2022).
- [32] J. N. Reinhardt, M. Staab, K. Yamamoto, J.-B. Bayle, A. Hees, O. Hartwig, K. Wiesner, S. Shah, and G. Heinzel, Ranging sensor fusion in LISA data processing: Treatment of ambiguities, noise, and onboard delays in LISA ranging observables, *Phys. Rev. D* **109**, 022004 (2024).
- [33] LISA Data Challenges Working Group, Sangria dataset, [10.5281/zenodo.7132177](#) (2022).
- [34] LISA Data Challenges Working Group, LISA Data Challenges, <https://lisa-ldc.lal.in2p3.fr/> (2022).
- [35] T. Littenberg, N. Cornish, K. Lackeos, and T. Robson, Global analysis of the gravitational wave signal from galactic binaries, *Phys. Rev. D* **101**, 123021 (2020).
- [36] K. Lackeos, T. B. Littenberg, N. J. Cornish, and J. I. Thorpe, The LISA data challenge Radler analysis and time-dependent ultra-compact binary catalogues, *Astron. Astrophys.* **678**, A123 (2023).

- [37] S. H. Strub, L. Ferraioli, C. Schmeltzbach, S. C. Stähler, and D. Giardini, Global analysis of LISA data with galactic binaries and massive black hole binaries, *Phys. Rev. D* **110**, 024005 (2024).
- [38] P. J. Green, Reversible jump Markov chain Monte Carlo computation and Bayesian model determination, *Biometrika* **82**, 711 (1995).
- [39] C. Ritter and M. Tanner, Facilitating the Gibbs sampler: The Gibbs stopper and the Griddy–Gibbs sampler, *J. Am. Stat. Assoc.* **87**, 861 (1992).
- [40] P. Muller, A generic approach to posterior integration and Gibbs sampling, Technical Report, 91-09, Purdue University, 1991.
- [41] W. R. Gilks and P. Wild, Adaptive rejection sampling for Gibbs sampling, *J. R. Stat. Soc. C (Appl. Stat.)* **41**, 337 (1992).
- [42] D. Harris, Applications on one of the world’s largest supercomputers demonstrate nvidia gpus lead in energy efficiency, <https://blogs.nvidia.com/blog/gpu-energy-efficiency-nersc/>.
- [43] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, EMCEE: The MCMC Hammer, *Publ. Astron. Soc. Pac.* **125**, 306 (2013).
- [44] N. Karnesis, M. L. Katz, N. Korsakova, J. R. Gair, and N. Stergioulas, Eryn: A multi-purpose sampler for Bayesian inference, *Mon. Not. R. Astron. Soc.* **526**, 4814 (2023).
- [45] M. Katz, N. Karnesis, and N. Korsakova, mikekatz04/Eryn: First full release, [10.5281/zenodo.7705496](https://doi.org/10.5281/zenodo.7705496) (2023).
- [46] M. Tinto and J. W. Armstrong, Cancellation of laser noise in an unequal-arm interferometer detector of gravitational radiation, *Phys. Rev. D* **59**, 102003 (1999).
- [47] LISA Data Challenges Working Group, Sangria dataset description <https://lisa-ldc.lal.in2p3.fr/challenge2a>.
- [48] D. A. Shaddock, M. Tinto, F. B. Estabrook, and J. W. Armstrong, Data combinations accounting for LISA spacecraft motion, *Phys. Rev. D* **68**, 061303 (2003).
- [49] T. A. Prince, M. Tinto, S. L. Larson, and J. W. Armstrong, The LISA optimal sensitivity, *Phys. Rev. D* **66**, 122002 (2002).
- [50] M. Vallisneri, Synthetic LISA: Simulating time delay interferometry in a model LISA, *Phys. Rev. D* **71**, 022001 (2005).
- [51] M. L. Katz, Fully automated end-to-end pipeline for massive black hole binary signal extraction from LISA data, *Phys. Rev. D* **105**, 044055 (2022).
- [52] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. Forteza, and A. Bohé, Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era, *Phys. Rev. D* **93**, 044007 (2016).
- [53] S. Husa, S. Khan, M. Hannam, M. Pürrer, F. Ohme, X. Forteza, and A. Bohé, Frequency-domain gravitational waves from nonprecessing black-hole binaries. I. New numerical waveforms and anatomy of the signal, *Phys. Rev. D* **93**, 044006 (2016).
- [54] N. J. Cornish and T. B. Littenberg, Tests of Bayesian model selection techniques for gravitational wave astronomy, *Phys. Rev. D* **76**, 083006 (2007).
- [55] M. L. Katz, C. Danielski, N. Karnesis, V. Korol, N. Tamanini, N. J. Cornish, and T. B. Littenberg, Bayesian characterization of circumbinary sub-stellar objects with LISA, *Mon. Not. R. Astron. Soc.* **517**, 697 (2022).
- [56] D. Quang Nam, J. Martino, Y. Lemi  re, A. Petiteau, J.-B. Bayle, O. Hartwig, and M. Staab, Time-delay interferometry noise transfer functions for LISA, *Phys. Rev. D* **108**, 082004 (2023).
- [57] N. Karnesis, S. Babak, M. Pieroni, N. Cornish, and T. Littenberg, Characterization of the stochastic signal originating from compact binary populations as measured by LISA, *Phys. Rev. D* **104**, 043019 (2021).
- [58] M. Sambridge, K. Gallagher, A. Jackson, and P. Rickwood, Trans-dimensional inverse problems, model comparison and the evidence, *Geophys. J. Int.* **167**, 528 (2006).
- [59] J. Goodman and J. Weare, Ensemble samplers with affine invariance, *Commun. Appl. Math. Comput. Sci.* **5**, 65 (2010).
- [60] W. Voudsen, W. Farr, and I. Mandel, Dynamic temperature selection for parallel tempering in Markov chain Monte Carlo simulations, *Mon. Not. R. Astron. Soc.* **455**, 1919 (2016).
- [61] M. Katz, mikekatz04/bb  x: First official public release, [10.5281/zenodo.5730688](https://doi.org/10.5281/zenodo.5730688) (2021).
- [62] N. J. Cornish, Heterodyned likelihood for rapid gravitational wave parameter inference, *Phys. Rev. D* **104**, 104054 (2021).
- [63] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, Scikit-learn: Machine learning in Python, *J. Mach. Learn. Res.* **12**, 2825 (2011).
- [64] LISACATTOOLS, <https://tlittenberg.github.io/lisacattools/>.
- [65] L. Dalc  n and Y.-L. L. Fang, mpi4py: Status update after 12 years of development, *Comput. Sci. Eng.* **23**, 47 (2021).
- [66] L. D. Dalc  n, R. R. Paz, P. A. Kler, and A. Cosimo, Parallel distributed computing using Python, *Adv. Water Resour.* **34**, 1124 (2011).
- [67] L. Dalc  n, R. Paz, M. Storti, and J. D’El  a, MPI for Python: Performance improvements and MPI-2 extensions, *J. Parallel Distrib. Comput.* **68**, 655 (2008).
- [68] L. Dalc  n, R. Paz, and M. Storti, MPI for Python, *J. Parallel Distrib. Comput.* **65**, 1108 (2005).
- [69] M. L. Katz, N. Karnesis, N. Korsakova, J. R. Gair, and S. Nikolaos, Erebor LDC2A training dataset output catalogs, [10.5281/zenodo.11118992](https://doi.org/10.5281/zenodo.11118992) (2024).
- [70] M. L. Katz, N. Karnesis, N. Korsakova, J. R. Gair, and N. Stergioulas, Erebor LDC2A hidden dataset output catalogs, [10.5281/zenodo.11131522](https://doi.org/10.5281/zenodo.11131522) (2024).
- [71] S. H. Strub, L. Ferraioli, C. Schmeltzbach, S. C. Stähler, and D. Giardini, Accelerating global parameter estimation of gravitational waves from galactic binaries using a genetic algorithm and GPUs, *Phys. Rev. D* **108**, 103018 (2023).
- [72] S. Nissanke, M. Vallisneri, G. Nelemans, and T. A. Prince, Gravitational-wave emission from compact galactic binaries, *Astrophys. J.* **758**, 131 (2012).
- [73] M. Muratore, J. Gair, and L. Speri, Impact of the noise knowledge uncertainty for the science exploitation of cosmological and astrophysical stochastic gravitational wave background with LISA, *Phys. Rev. D* **109**, 042001 (2024).

- [74] N. Korsakova, S. Babak, M.L. Katz, N. Karnesis, S. Khukhlaev, and J.R. Gair, Neural density estimation for galactic binaries in LISA data analysis, *Phys. Rev. D* **110**, 104069 (2024).
- [75] L. London, S. Khan, E. Fauchon-Jones, C. García, M. Hannam, S. Husa, X. Jiménez-Forteza, C. Kalaghatgi, F. Ohme, and F. Pannarale, First higher-multipole model of gravitational waves from spinning and coalescing black-hole binaries, *Phys. Rev. Lett.* **120**, 161102 (2018).
- [76] A. Ramos-Buades, A. Buonanno, M. Khalil, and S. Ossokine, Effective-one-body multipolar waveforms for eccentric binary black holes with nonprecessing spins, *Phys. Rev. D* **105**, 044035 (2022).
- [77] C.R. Harris *et al.*, Array programming with NumPy, *Nature (London)* **585**, 357 (2020).
- [78] J.D. Hunter, Matplotlib: A 2D graphics environment, *Comput. Sci. Eng.* **9**, 90 (2007).
- [79] P. Virtanen *et al.* (SciPy 1.0 Contributors), SciPy 1.0: Fundamental algorithms for scientific computing in Python, *Nat. Methods* **17**, 261 (2020).
- [80] R. Okuta, Y. Unno, D. Nishino, S. Hido, and Crissman, CuPy: A NumPy-compatible library for nvidia gpu calculations (2017).
- [81] M. Katz, C. Chapmanbird, L. Speri, N. Karnesis, and N. Korsakova, mikekatz04/lisaanalysis tools: First main release, 10.5281/zenodo.10930980 (2024).
- [82] <https://github.com/mikekatz04/LISAanalysis tools/tree/9305e9963b1920d32fcb74fb341dd136063577c5>
- [83] <https://github.com/mikekatz04/Eryn/tree/58347f5433ed9670b3c0791f5e395c5d00bea865>
- [84] <https://github.com/mikekatz04/BBHx/tree/646a4e268e37e4e1d5c87d0d94bdd4e8ce9b69e9>
- [85] M. L. Katz, mikekatz04/gbgpu: First official public release!, 10.5281/zenodo.6500434 (2022).
- [86] <https://github.com/mikekatz04/GBGPU/tree/0a9e8429f64d4823c00cfe572c93c2f1da8181c0>
- [87] <https://zenodo.org/records/11000534>
- [88] <https://zenodo.org/records/10989860>
- [89] Nvidia cufftx, <https://docs.nvidia.com/cuda/cufftx/>.