Assessing the data-analysis impact of LISA orbit approximations using a GPU-accelerated response model

Michael L. Katz,[1,*] Jean-Baptiste Bayle,[2] Alvin J. K. Chua,[3] and Michele Vallisneri[2]

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

[2] *Jet Propulsion Laboratory, California Institute of Technology, Pasadena CA 91109, USA*

[3] *Theoretical Astrophysics Group, California Institute of Technology, Pasadena CA 91125, USA*

(Dated: April 18, 2022)

The analysis of gravitational-wave (GW) datasets is based on the comparison of measured time series with theoretical templates of the detector's response to a variety of source parameters. For the Laser Interferometer Space Antenna (LISA), the main scientific observables will be the so-called time-delay interferometry (TDI) combinations, which suppress the otherwise overwhelming laser noise. Computing the TDI response to GWs involves projecting the GW polarizations onto the LISA constellation arms, and then combining projections delayed by a multiple of the light propagation time along the arms. Both computations are difficult to perform efficiently for generic LISA orbits and GW signals. Various approximations are currently used in practice, e.g., assuming constant and equal armlengths, which yields analytical TDI expressions. In this article, we present `fastlisaresponse`, a new efficient GPU-accelerated code that implements the generic TDI response to GWs in the time domain. We use it to characterize the parameter-estimation bias incurred by analyzing loud Galactic-binary signals using the equal-armlength approximation. We conclude that equal-armlength parameter-estimation codes should be upgraded to the generic response if they are to achieve optimal accuracy for high (but reasonable) SNR sources within the actual LISA data.

I. INTRODUCTION

The launch of the Laser Interferometer Space Antenna (LISA) in the early 2030s will extend the reach of gravitational-wave (GW) astronomy to the millihertz band, complementing the plentiful observations of compact binaries obtained by ground-based GW detectors [1, 2]. GW sources in the millihertz band include Galactic binaries, typically consisting of two white dwarfs emitting quasi-monochromatic GWs; massive–black-hole binaries, producing loud signals that sweep across the band as the black holes inspiral and merge; stellar-origin black-hole binaries, which are detected by LISA at large orbital separations and will eventually merge in the band of ground-based detectors; and extreme mass-ratio inspirals, consisting of a stellar-mass compact object orbiting a massive black hole. The combination of observations from all these sources will provide immense scientific return [3].

Analyzing the LISA data stream is a tall task. Detector noise is expected to be non-stationary due to the presence of glitches, data gaps, and drifting noise levels in the LISA components. Furthermore, a large number of overlapping GW signals will be present for the duration of the observation. Therefore, estimating detector noise will be more complicated than in the ground-based case, where signal-free stretches of data can be used to characterize noise directly. An additional complication is the motion of the LISA spacecraft: the compact-binary coalescence signals observed from the ground are short enough that detectors can be assumed be to be static, allowing for a detector response function that is a function of frequency only. By contrast, the movement of the LISA constellation in its Solar orbit leads to a response function that depends on both frequency and time.

It is useful to decompose the LISA GW response function in two stages: in the first, we project the GW polarizations onto the evolving LISA arms, computing projections along all six interferometric links (forward and backward along each arm); in the second, the projections are combined with appropriate time shifts to form time-delay interferometry (TDI) observables, which suppress the laser noise that would otherwise drown the GW signals. For this reason, the LISA sensitivity is typically quoted for TDI observables [3]. Both stages require knowledge of the LISA orbits: the projections involve the spacecraft positions and the light-propagation times, or delays, along all six links, while TDI requires sufficiently accurate delays for proper laser-noise suppression.

Many LISA analysis codes make the simplifying assumption of *equal-armlength* orbits. This unphysical model neglects the breathing of the armlengths in actual orbits, as well as the orbital corrections required to realign the spacecraft to counter accumulating drifts. The equal-armlength model has a period of one year and is completely determined by two parameters. The computation of TDI expressions is especially convenient because fixed armlength delays allow the resummation of single-arm responses using trigonometric identities over GW phasing terms [see, e.g., 4]. A more realistic but still unphysical model places spacecraft on eccentric Keplerian orbits. It also has a period of one year and is completely determined by three parameters.

The actual LISA orbits are influenced by the grav-

* michael.katz@aei.mpg.de

itational attraction of the planets, as well as non-gravitational effects such as solar wind; thus they display drifts leading to unequal and evolving armlengths[1], and non-periodic trajectories. To account for these effects fully, TDI observables must be constructed by evaluating single-arm GW projections at delayed times that are changing continually along the orbits. In this work, we present the first efficient implementation on a graphics processing unit (GPU) of this *general* GW response for arbitrary LISA orbits. This new code can be used directly in GW searches and parameter estimation, avoiding any orbit or delay approximation that may affect the accuracy of results.

Conversely, the code enables an exploration of the effects of orbit approximations in LISA data analysis. Using Galactic-binary waveforms as an example and our GPU code as *truth*, we evaluate the loss of parameter-estimation accuracy incurred by computing TDI observables using the equal-armlength approximation: for moderately loud signals, we find significant bias in the recovery of important astrophysical parameters such as frequency, frequency derivative, and sky position. We also test a hybrid approximation that uses realistic orbits for projections, and equal armlengths for TDI. While such hybrid templates fit the truth much better, we still find consequential bias for louder (but realistic) sources. Therefore, it is our recommendation that LISA analysis codes used for actual data should fully model the effects of the LISA orbits. Tools such as our GPU code will ensure that the additional detail does not pose an undue burden on computation.

This paper is organized as follows: in section II, we describe the statistical methods used in GW analysis with LISA; in section III, we introduce the waveform templates used in our study; in section IV, we discuss the formulation and implementation of the response function; in section V, we describe the results of our analysis; in section VI, we discuss the implications of our findings; last, in section VII we present our conclusions.

II. BAYESIAN ANALYSIS ON GRAVITATIONAL-WAVE SIGNALS

A. Likelihood function

The analysis of GW signals is generally performed within a Bayesian framework, driven by Bayes' rule,

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

where $d = d(t)$ is the measured GW data; $\vec{\Theta}$ is the vector of parameters representing a GW source, and is associated with the assumed model of the signal, Λ; $p(\vec{\Theta}|d,\Lambda)$ is the posterior probability distribution on the parameters of the source; $p(d|\vec{\Theta},\Lambda) = \mathcal{L}$ is the probability that the data is represented by the chosen model and model parameters, also called the likelihood; $p(\vec{\Theta}|\Lambda)$ is the prior probability on the parameters; and $p(d|\Lambda)$ is the integral of the numerator over all parameter space, also referred to as the evidence. In this study, we will produce the posterior distribution by drawing samples from it with Markov chain Monte Carlo (MCMC) techniques. Typically, the evidence is intractable in the GW setting. With MCMC methods, the evidence enters as a constant factor and can, therefore, be neglected. In order to determine the posterior, we must compute the likelihood. Under the assumption of stationary and Gaussian noise, we write down the noise-weighted inner product between two time domain signals $a(t)$ and $b(t)$,

$$\langle a|b\rangle = 4\,\mathrm{Re}\int_0^\infty \frac{\tilde{a}(f)^*\tilde{b}(f)}{S_n(f)}\,\mathrm{d}f\,, \tag{2}$$

where $\tilde{a}(f)$ is the Fourier Transform of $a(t)$ and $S_n(f)$ is the one-sided power spectral density (PSD) of the noise. We use the "SciRDv1" curve from the `tdi` package from the LISA Data Challenge (LDC) [5] for the PSD. The log-likelihood is a linear combination of inner products,

$$\begin{aligned}\log\mathcal{L} &\propto -\frac{1}{2}\langle d-h|d-h\rangle \\ &= -\frac{1}{2}(\langle d|d\rangle + \langle h|h\rangle - 2\langle d|h\rangle)\,,\end{aligned} \tag{3}$$

where $h(t)$ is the template that models the true signal, $s(t)$. The true signal together with the noise contribution, $n(t)$, makes the data stream: $d(t) = s(t) + n(t)$. For all studies in this work, we use no additive noise ($n(t) = 0$). This allows for the direct analysis of the likelihood surface without any shift due to the random noise process.

The optimal signal-to-noise ratio (SNR) ρ_{opt} achievable for a template is $\sqrt{\langle h|h\rangle}$. When fitting a template against the data (and assuming the true signal is perfectly modeled by the template), the SNR will be normally distributed about the optimal SNR with variance equal to 1. The extraction SNR ρ_{ex} is given by

$$\rho_{\mathrm{ex}} = \frac{\langle d|h\rangle}{\sqrt{\langle h|h\rangle}}\,. \tag{4}$$

The extraction SNR is fundamental to signal processing in terms of understanding the detection of a signal in a time series. We use it in this work as a diagnostic indicator to help understand detection capabilities of different template models. Assuming no additive noise, the overlap (normalized cross-correlation) between the template and data is given by $\rho_{\mathrm{ex}}/\sqrt{\langle s|s\rangle}$. The mismatch is then defined as 1 minus the overlap.

[1] Note that analytic Keplerian orbits also yield unequal and evolving armlengths. Variation of armlengths in the case of realistic orbits is an order of magnitude larger.

An additional statistic that is of use in this work is the *fitting factor*. The fitting factor is the overlap between the data and the template whose parameters are located at the global maximum in the likelihood surface, which is found using stochastic sampling methods (discussed below).

B. Markov chain Monte Carlo sampler

We use the MCMC technique to draw samples from the posterior distribution. The sampler is a combination[2] of the overall architecture of `emcee` [6] with the parallel tempering methods from `ptemcee` [7]. In all of the sampling runs used in this work, four temperatures are used ranging from 1 (target distribution) to ∞ (prior). Each temperature contains 50 walkers. The only proposal used was the "Stretch" proposal [8].

III. TEMPLATE GENERATION

As the Markov chain evolves, templates are generated at each step to evaluate the likelihood for all walkers. The first step in this process is to generate $h_+^{\rm SSB}(t)$ and $h_\times^{\rm SSB}(t)$, the two polarization time series at the Solar system's barycenter (SSB). For all tests in this paper, we focus on single Galactic white dwarf binary, whose time domain waveform in its source frame (scaled by the distance and denoted with an "S" superscript) is given by

$$h_+^{\rm S} = A\cos\left(2\pi\left(f_0 t + \frac{1}{2}\dot{f}_0 t^2 + \frac{1}{6}\ddot{f}_0 t^3\right) + \phi_0\right), \quad (5a)$$

$$h_\times^{\rm S} = A\sin\left(2\pi\left(f_0 t + \frac{1}{2}\dot{f}_0 t^2 + \frac{1}{6}\ddot{f}_0 t^3\right) + \phi_0\right), \quad (5b)$$

where A is the amplitude of the GW at the SSB; $\{f_0, \dot{f}_0, \ddot{f}_0\}$ are the initial GW frequency and its first two time derivatives; and ϕ_0 is the initial phase of the GW. The amplitude A is given by

$$A(\mathcal{M}_c, f_0, d_L) = 2\frac{(G\mathcal{M}_c)^{5/3}}{c^4 d_L}(\pi f)^{2/3}, \quad (6)$$

where $\mathcal{M}_c$ is the chirp mass and d_L is the luminosity distance. Assuming that GW emission solely drives the evolution of the binary, the frequency derivative $\dot{f}_{\rm GW}$ is given by

$$\dot{f}(\mathcal{M}_c, f) = \frac{96}{5}\pi^{8/3}\left(\frac{G\mathcal{M}_c}{c^3}\right)^{5/3} f^{11/3}. \quad (7)$$

For a given value of f_0, choosing the initial frequency derivative $\dot{f}_0$ effectively fixes the chirp mass of the binary.

[2] The sampler code is available upon request to the authors. It will soon be made publicly available.

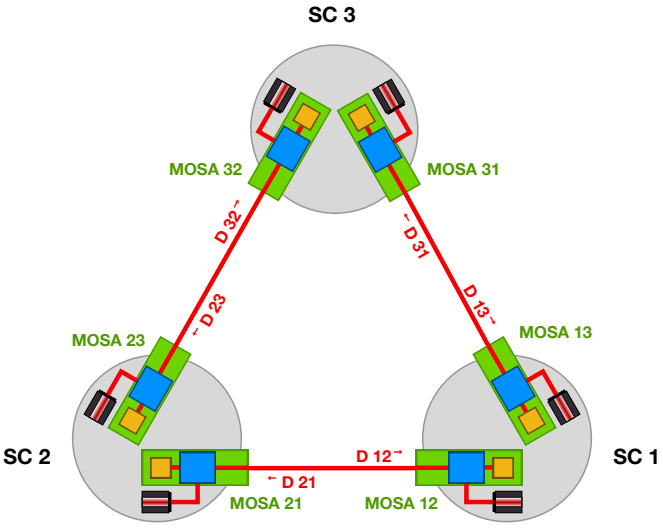

Figure 1. Indexing conventions.

With f_0 and $\dot{f}_0$ set, we can deduce $\ddot{f}_0$ assuming once again that gravitational radiation is the sole driver of the binary evolution,

$$\ddot{f}_0 = \frac{11}{3}\frac{\dot{f}_0^2}{f_0}. \quad (8)$$

Further derivatives have very small effects and are ignored here.

These source-frame GW orientations are then transformed to the SSB frame using the orbital inclination ι and polarization angle ψ,

$$\begin{bmatrix} h_+^{\rm SSB} \\ h_\times^{\rm SSB} \end{bmatrix} = \begin{bmatrix} \cos 2\psi & -\sin 2\psi \\ \sin 2\psi & \cos 2\psi \end{bmatrix} \begin{bmatrix} -(1+\cos^2\iota)h_+^{\rm S} \\ -(2\cos\iota)h_\times^{\rm S} \end{bmatrix}. \quad (9)$$

The parameter set used above, $\{A, f_0, \dot{f}_0, \phi_0, \iota, \psi\}$, combined with the two sky-localization angles, give the eight parameters needed to generate our Galactic binary waveform. The sky-localization angles represent the ecliptic longitude λ and latitude β. These angles will appear in the detector response described below.

IV. INSTRUMENT RESPONSE

We follow the standard conventions given in fig. 1. Spacecraft are indexed from 1 to 3 clockwise when looking down on the z-axis. Moveable optical sub-assemblies (MOSAs) are indexed with two numbers ij, where i is the index of the spacecraft the system is mounted on (local spacecraft), and j is the index of the spacecraft the light is received from (distant spacecraft).

Measurements are indexed according to the MOSA they are performed on. Light travel times are indexed according to the MOSA they are measured on, i.e., the receiving spacecraft. All equations in this document possess the symmetries of the triangular constellation, i.e.,

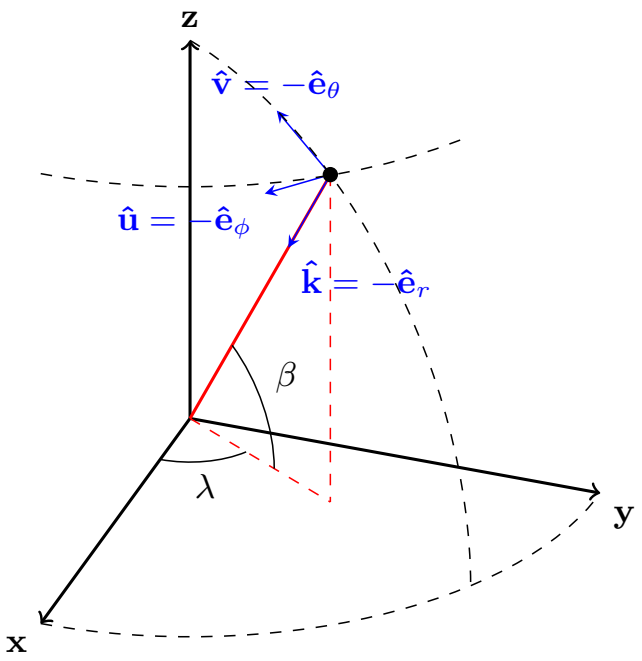

Figure 2. Source localization in the SSB frame. Adapted from the LDC Manual. The propagation vector is $\hat{\mathbf{k}}$. The polarization vectors are $\hat{\mathbf{u}}$ and $\hat{\mathbf{v}}$.

3 reflections and 3 rotations. All associated transformations on the indices can be generated from one circular permutation of indices and one reflection. Therefore, we will only give one expression and leave the reader to form all remaining expressions using this set of index transformations.

A. Reference frames

We present here the formulation of the LISA response to gravitational waves, following the conventions proposed by the LDC Manual[3].

The SSB Cartesian coordinate system is defined by $(\mathbf{x}, \mathbf{y}, \mathbf{z})$, such that $(\mathbf{x}, \mathbf{y})$ is the plane of the ecliptic. The SSB frame is used to express the LISA spacecraft coordinates, as well as the binary sky-localization. We introduce the SSB spherical coordinates (θ, ϕ) as illustrated in fig. 2, based on the orthonormal basis vectors $(\hat{\mathbf{e}}_r, \hat{\mathbf{e}}_\theta, \hat{\mathbf{e}}_\phi)$. The source localization is parametrized by the *ecliptic latitude* $\beta = \pi/2 - \theta$ and the *ecliptic longitude* $\lambda = \phi$. The basis vectors read

$$\hat{\mathbf{e}}_r = (\cos\beta\cos\lambda, \cos\beta\sin\lambda, \sin\beta)\,, \tag{10a}$$
$$\hat{\mathbf{e}}_\theta = (\sin\beta\cos\lambda, \sin\beta\sin\lambda, -\cos\beta)\,, \tag{10b}$$
$$\hat{\mathbf{e}}_\phi = (-\sin\lambda, \cos\lambda, 0)\,. \tag{10c}$$

The propagation vector is $\hat{\mathbf{k}} = -\hat{\mathbf{e}}_r$. We define the *polarization vectors* as $\hat{\mathbf{u}} = -\hat{\mathbf{e}}_\phi$ and $\hat{\mathbf{v}} = -\hat{\mathbf{e}}_\theta$. This produces a direct orthonormal basis in $(\hat{\mathbf{u}}, \hat{\mathbf{v}}, \hat{\mathbf{k}})$.

B. Projection on the constellation arms

The deformation induced on link 12, measured on MOSA 12, is denoted as $H_{12}(t)$. It is computed by projecting the SSB GW strain on the link unit vector $\hat{\mathbf{n}}_{12}$ (computed from the spacecraft positions),

$$\begin{aligned} H_{12}(t) = {} & h_+^{\mathrm{SSB}}(t) \times \xi_+(\hat{\mathbf{u}}, \hat{\mathbf{v}}, \hat{\mathbf{n}}_{12}) \\ & + h_\times^{\mathrm{SSB}}(t) \times \xi_\times(\hat{\mathbf{u}}, \hat{\mathbf{v}}, \hat{\mathbf{n}}_{12})\,, \end{aligned} \tag{11}$$

where we assume that the link unit vector $\hat{\mathbf{n}}_{12}$ is constant during the light travel time. The *antenna pattern functions* are given by

$$\xi_+(\hat{\mathbf{u}}, \hat{\mathbf{v}}, \hat{\mathbf{n}}_{12}) = (\hat{\mathbf{u}}\cdot\hat{\mathbf{n}}_{12})^2 - (\hat{\mathbf{v}}\cdot\hat{\mathbf{n}}_{12})^2\,, \tag{12a}$$
$$\xi_\times(\hat{\mathbf{u}}, \hat{\mathbf{v}}, \hat{\mathbf{n}}_{12}) = 2(\hat{\mathbf{u}}\cdot\hat{\mathbf{n}}_{12})(\hat{\mathbf{v}}\cdot\hat{\mathbf{n}}_{12})\,. \tag{12b}$$

Light emitted by spacecraft 2 at t_2 reaches spacecraft 1 at t_1. These two times t_1 and t_2 are related by $H_{12}(\mathbf{x}, t)$,

$$t_1 \approx t_2 + \frac{L_{12}}{c} - \frac{1}{2c}\int_0^{L_{12}} H_{12}(\mathbf{x}(\lambda), t(\lambda))\,\mathrm{d}\lambda\,. \tag{13}$$

We approximate the wave propagation time to first order as $t(\lambda) \approx t_2 + \lambda/c$. Also, $\mathbf{x}(\lambda) = \mathbf{x}_2(t_2) + \lambda\hat{\mathbf{n}}_{12}(t_2)$, where $\mathbf{x}_2(t_2)$ represents the position of the emitter spacecraft at emission time. Using these two expressions, we can further refine H_{12} as

$$\begin{aligned} H_{12}(\mathbf{x}(\lambda), t(\lambda)) &= H_{12}\left(t(\lambda) - \frac{\hat{\mathbf{k}}\cdot\mathbf{x}(\lambda)}{c}\right) \\ &= H_{12}\left(t_2 - \frac{\hat{\mathbf{k}}\cdot\mathbf{x}_2(t_2)}{c} + \frac{1 - \hat{\mathbf{k}}\cdot\hat{\mathbf{n}}_{12}(t_2)}{c}\lambda\right). \end{aligned} \tag{14}$$

Combining eqs. (13) and (14) and differentiating the resulting expression with respect to t_2 yields the relative frequency shift, y_{12}, experienced by light as it travels along link 12,

[3] Available at `https://lisa-ldc.lal.in2p3.fr`.

$$y_{12}(t_2) \approx \frac{1}{2\Big(1-\hat{\mathbf{k}}\cdot\hat{\mathbf{n}}_{12}(t_2)\Big)}\left[H_{12}\left(t_2-\frac{\hat{\mathbf{k}}\cdot\mathbf{x}_2(t_2)}{c}\right)-H_{12}\left(t_2-\frac{\hat{\mathbf{k}}\cdot\mathbf{x}_1(t_1)}{c}+\frac{L_{12}}{c}\right)\right]. \tag{15}$$

Here, we have introduced the receiver spacecraft position at reception time $\mathbf{x}_1(t_1) = \mathbf{x}_2(t_2) + L_{12}\hat{\mathbf{n}}_{12}(t_2)$. Using $t_1 \approx t_2 + L_{12}/c$ and the fact that the spacecraft moves slowly compared to the propagation timescale, we obtain $\mathbf{x}_2(t_2) \approx \mathbf{x}_2(t_1)$ and $\hat{\mathbf{n}}_{12}(t_1) \approx \hat{\mathbf{n}}_{12}(t_2)$. Finally,

$$y_{12}(t_1) \approx \frac{1}{2\Big(1-\hat{\mathbf{k}}\cdot\hat{\mathbf{n}}_{12}(t_1)\Big)}\left[H_{12}\left(t_1-\frac{L_{12}(t_1)}{c}-\frac{\hat{\mathbf{k}}\cdot\mathbf{x}_2(t_1)}{c}\right)-H_{12}\left(t_1-\frac{\hat{\mathbf{k}}\cdot\mathbf{x}_1(t_1)}{c}\right)\right], \tag{16}$$

where the equation for y_{12} is now solely a function of reception time t_1. Combining eqs. (11) and (16) gives y_{12} as a function of t_1 in terms of the GW strain.

The y_{ij} time series along each of the six links are the final quantities output of the projection step. They are then combined in various ways to compute the TDI observables.

C. Time-delay interferometry

TDI combinations are defined as linear combinations of time-shifted measurements. The first and second-generation Michelson combinations, X_1 and X_2, are given by [9],

$$\begin{aligned} X_1 = {} & y_{13} + \mathbf{D}_{13}y_{31} + \mathbf{D}_{131}y_{12} + \mathbf{D}_{1312}y_{21} \\ & - [y_{12} + \mathbf{D}_{12}y_{21} + \mathbf{D}_{121}y_{13} + \mathbf{D}_{1213}y_{31}]\,, \end{aligned} \tag{17}$$

$$\begin{aligned} X_2 = {} & X_1 + \mathbf{D}_{13121}y_{12} + \mathbf{D}_{131212}y_{21} + \mathbf{D}_{1312121}y_{13} \\ & + \mathbf{D}_{13121213}y_{31} - [\mathbf{D}_{12131}y_{13} + \mathbf{D}_{121313}y_{31} \\ & + \mathbf{D}_{1213131}y_{12} + \mathbf{D}_{12131312}y_{21}]\,. \end{aligned} \tag{18}$$

Delay operators are defined by

$$\mathbf{D}_{ij}x(t) = x(t - L_{ij}(t))\,, \tag{19}$$

where $L_{ij}(t)$ is the delay time along link ij at reception time t. Because light travel times evolve slowly with time, we compute chained delays as simple sums of delays rather than nested delays, i.e.,

$$\mathbf{D}_{i_1,i_2,\ldots,i_n}x(t) = x\left(t - \sum_{k=1}^{n-1} L_{i_k i_{k+1}}(t)\right). \tag{20}$$

While this approximation cannot be used to study laser-noise suppression upstream of the LISA data analysis, it is sufficient when computing the GW response function. Note that these equations are left unchanged (up to a sign) by reflection symmetries. However, applying the three rotations generates the three Michelson combinations, X, Y, Z, for both generations.

These Michelson combinations have correlated noise properties. An uncorrelated set of TDI variables, A, E, T, can be obtained from linear combinations of X, Y, Z given by [10]

$$A = \frac{1}{\sqrt{2}}(Z - X)\,, \tag{21a}$$

$$E = \frac{1}{\sqrt{6}}(X - 2Y + Z)\,, \tag{21b}$$

$$T = \frac{1}{\sqrt{3}}(X + Y + Z)\,. \tag{21c}$$

The inner products from eq. (3) are really sums over the three channels,

$$\langle a|b\rangle = \sum_{i=A,E,T} \langle a^i|b^i\rangle\,. \tag{22}$$

Note that A, E, T are only exactly orthogonal (or uncorrelated in noise properties) in the equal-armlength limit. In this exploratory work, we maintain the use of A, E, T even in the limit of breathing constellation arms, because we are not considering additive noise and believe this approximation is good enough to inform us of the effect of orbital assumptions on the analysis. In future work, and for the actual LISA analysis, we will have to determine the noise information and properly compute the likelihood with off-diagonal terms representing the correlation between the various TDI observables [11].

D. Orbital trajectories

The orbital trajectory of the LISA constellation affects the projections of the GWs onto the constellation arms, as well as the computation of the TDI observables through the armlengths (or, equivalently, the light travel times along each link).

In this study, we focus on two classes of orbits. The standard *equal-armlength* orbital configuration is used in most data analysis codes. The three spacecraft follow heliocentric Keplerian orbits that leave armlengths constant to leading order in the orbit eccentricity [12]. This is realized when the constellation plane has an angle of 60° with the ecliptic. In our study, the semi-major axis is set to 1 au and the mean inter-spacecraft distance to the nominal value of 2.5×10^9 m. In the frequency domain,

which is used for the vast majority of current LISA analysis techniques, the armlength is strictly fixed to a single constant value. This is necessary to derive the simpler analytic TDI expressions used in frequency-domain template generation [see, e.g., 4]. In the time domain, the "equal armlengths" are only nearly constant, but their variations remain much smaller than the typical values seen in realistic orbits (see below).

The other class of orbits we examine are numerically generated orbits provided by the European Space Agency (ESA) [13]. These orbits take into account all relevant bodies in the Solar system, and are optimized to minimize, amongst others, constellation breathing and fuel consumption necessary to insert the constellation in the correct orbit.

For both equal-armlength and ESA orbital configurations, the light propagation delay time series along the six constellation links are computed by `LISA Orbits` [14], using a second-order post-Minkowskian expansion of the light travel time equation.

One issue in trying to compare equal-armlength orbits to accurate orbits is finding the proper parameterization of the equal-armlength configuration that minimizes any intrinsic bias between the two models. We used the Nelder-Mead optimization algorithm to determine the two free parameters for equal-armlength orbits, chosen as λ_1, the longitude of the periastron for spacecraft 1, and m_1, the mean anomaly at initial time for spacecraft 1. The cost function $f(\lambda_1, m_1)$ was written as the sum, for each spacecraft i, of the absolute distances between the spacecraft positions according to equal-armlength orbits $\mathbf{x}_i^{\mathrm{EQ}}$ and their true positions given by accurate ESA orbits $\mathbf{x}_i^{\mathrm{ESA}}$, integrated over four years of mission,

$$f(\lambda_1, m_1) = \int_{t=0}^{4\,\mathrm{yr}} \sum_{i=1,2,3} \left\| \mathbf{x}_i^{\mathrm{EQ}}(\lambda_1, m_1, t) - \mathbf{x}_i^{\mathrm{ESA}}(t) \right\| . \tag{23}$$

To check the result, we run the estimator 10 times, with different initial guesses for λ_1 and m_1 randomly chosen between 0 and 2π. In all cases, the algorithm converges quickly to the same set of parameters in less than 100 iterations.

Examples of the difference between the two orbital configurations in satellite 1's orbital parameters and light travel time for one link are shown in fig. 3. All three physical coordinates (x, y, z) are given in the ecliptic coordinate frame. The maximum and mean difference in x between the two orbits are $\sim 0.04\,\mathrm{au}$ and $\sim 0.008\,\mathrm{au}$, respectively. The y coordinate differences have roughly the same properties. The z-coordinate maximum difference is $\sim 4 \times 10^{-4}\,\mathrm{au}$ and the mean difference is $\sim 10^{-4}\,\mathrm{au}$. These spatial coordinate properties are consistent across all three spacecraft. The delay times for link 12 show an expected non-periodic difference between orbits. The delay times for the equal-armlength orbits are roughly constant with a variation range of $\sim 1 \times 10^{-3}\,\mathrm{s}$. The ESA orbit delay time for link 12 varies from $\sim 8.154\,\mathrm{s}$ to $\sim 8.387\,\mathrm{s}$ (range of $\sim 0.232\,\mathrm{s}$), clearly displaying a non-periodic behavior.

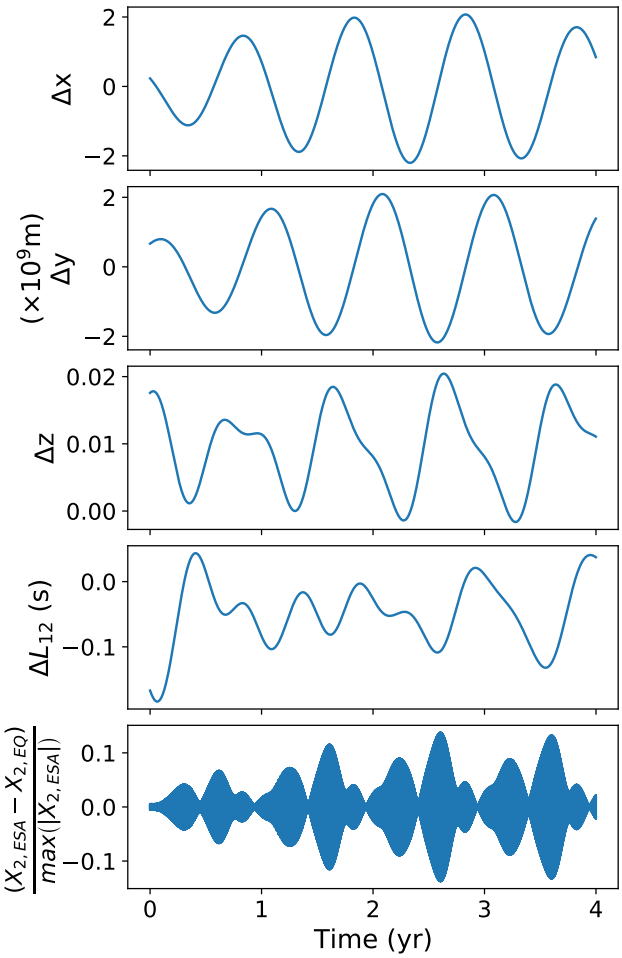

Figure 3. Orbital trajectory difference as determined from the two orbital configurations tested: equal-armlength orbits ("EQ") and orbits from ESA ("ESA"). The top three plots show, from top to bottom, the difference in the x, y, and z coordinates of spacecraft 1's position over four years, in the ecliptic reference frame. The second plot from the bottom shows the difference in the light travel time along link 12. The bottom plot shows an example of relative differences in the TDI X_2 channel between the two orbital models, over four years of observation, for a binary with parameters $\{A, f_0, \dot{f}_0, \iota, \phi_0, \psi, \lambda, \beta\} = \{10^{-22}, 10\,\mathrm{mHz}, 1.4 \times 10^{-14}\,\mathrm{Hz\,s^{-1}}, 1.11, 4.58, 0.45, 5.23, 1.23\}$.

The orbit of the LISA constellation is included in the Λ model parameter from eq. (1). In the case of equal-armlength orbits, Λ will be an incorrect model that will incur some mismatch against the true waveforms. A visual comparison of the TDI X_2 channel between the two models, over the course of observation, is shown in fig. 3. We will address this quantitatively in section V.

E. Implementation

Codes performing these generic time-domain calculations have been available for some time [15]. Our implementation of this process takes advantage of strong computational acceleration using GPUs. Our code, `fastlisaresponse` can be found on GitHub [16]. All computations of the waveform, response function, and likelihood are performed on GPUs, only transferring the final value of the likelihood back to the CPU processing stream. The simple Galactic binary waveform is produced using `CuPy`, an effective drop-in replacement for `NumPy` [17] built for NVIDIA GPUs.

The response function is coded in `C++/CUDA` and wrapped into `Python` with a special `Cython` [18] setup script from [19]. The input orbital information is interpolated and then evaluated at each evenly-spaced time step in the observed data stream. This information is pre-computed and stored for the entirety of the sampling run. All time points in all six projections are evaluated in parallel on a separate thread of the GPU. The projections require interpolation of $h_+^{\rm SSB}$ and $h_\times^{\rm SSB}$ at desired time points. This is performed using centered Lagrange interpolating polynomials of user-defined order. Typically, an order of 25 is used to be conservative. Lower orders, even as low as less than 10, are generally accurate enough and can improve the speed of the computation.

To perform this interpolation, a set of points from the original waveform is required. These points span before and after the desired interpolation time. Typically, this is not a concern for central processing unit (CPU) programming because efficient memory access is not a limiting factor. For GPU programming, with slower memory access from global memory, gathering separate sub-arrays from the original waveform for each point in time where interpolation is needed can be the bottleneck. Therefore, the GPU shared memory is leveraged by storing sub-arrays spanning enough length for use by multiple time points during the computation. The use of shared memory led to a crucial improvement in the overall performance of this code.

The same process is performed to compute the TDI variables, by running all time points for all TDI delay combinations (eqs. (17) and (18)) in parallel along separate threads. The same Lagrange polynomials are used to interpolate the six projections to form the proper TDI combinations.

To check the correctness of `fastlisaresponse`, we compared, for a small set of Galactic binary strain as input, the resulting TDI time series to those obtained on CPU with the time-domain generic response code `LISA GW Response` [20] and `PyTDI` [21]. In all cases we tested, time series matched down to numerical precision after correcting for the different conventions on initial times implemented in these codes.

The speed of this new code is illustrated in fig. 4 comparing CPU to GPU performance as a function of the duration of a template waveform. The speed comparison is shown for the projection piece ("Proj."), TDI 1, TDI 2, and the full response function combining the projection with TDI 2. These tests were performed on a single AMD EPYC 7713 CPU processor and an NVIDIA Tesla A100 GPU. As the duration of the data stream increases, the GPU performance increase levels out at $\sim 600\times$ faster than the CPU. The response function for one year of observation takes $\sim 10\,$ms ($\sim 10^4\,$ms for the CPU). The Lagrangian interpolation order does affect this time by a factor of order unity between the minimum and maximum orders tested of 1 and 25, respectively. In current frequency-domain response methods, the response function is usually evaluated in $\sim$ ms on a CPU and $\sim$ µs when evaluated in batches on a GPU [22]. The main reason for this efficient speed is the ability to operate with a much sparser or shorter array of frequencies compared to the number of time points in the time-domain signal. Frequency-domain response functions are usually evaluated with $\sim 2^{10}$ points or less compared to the $\sim 2^{22}$ points in the equivalent one-year time-domain evaluation.

V. POSTERIOR ANALYSIS

We analyze the effect of using different orbits within MCMC by taking advantage of the computational efficiency of the GPU implementation. We tested binaries across a multidimensional grid in SNR, ecliptic latitude, frequency, and frequency derivative. Two SNR values were used: an SNR of 30 was chosen to examine if orbital effects are visible in common, quieter Galactic binary signals. An SNR of 500 was chosen to represent louder sources, and allow us to ensure that the response function remains accurate enough for our "best" sources [23]. Two ecliptic latitudes were examined; one configuration is near-planar to the ecliptic with $\beta = 0.09\,$rad; the other is near-polar with $\beta = 1.2\,$rad.

In order to span the range of interest of Galactic binary frequencies and frequency derivatives, we tested three independent parameter configurations with $(f, \dot{f}) \in \{(7\times10^{-4}, 1\times10^{-19}),\ (2\times10^{-3}, 1.2\times10^{-17}),\ (1\times10^{-2}, 1.4\times10^{-14})\}$ (frequency in Hz and frequency derivative in $\mathrm{Hz\,s^{-1}}$). The three configurations were chosen to represent a binary at low-frequency and low-frequency derivative (chirp mass of $\sim\ 0.08\,\mathrm{M}_\odot$); one binary at middle frequency and middle frequency derivative (chirp mass of $\sim\ \ 0.2\,\mathrm{M}_\odot$); and one binary at high frequency and frequency derivative (chirp mass of $\sim\ 0.75\,\mathrm{M}_\odot$). The amplitude for each source tested was adjusted to achieve the desired SNR value.

The other parameters were randomly chosen from their respective prior distributions (see below). Our tests were computed with $\{\iota, \phi_0, \psi, \lambda\} = \{1.11, 4.58, 0.45, 5.23\}$ (in rad). Prior to performing MCMC, we tested a larger grid of values with basic mismatch calculations. It was clear from these computations that the main factors affecting the mismatch were the frequency, frequency derivative, and the ecliptic latitude. Therefore, we believe our grid

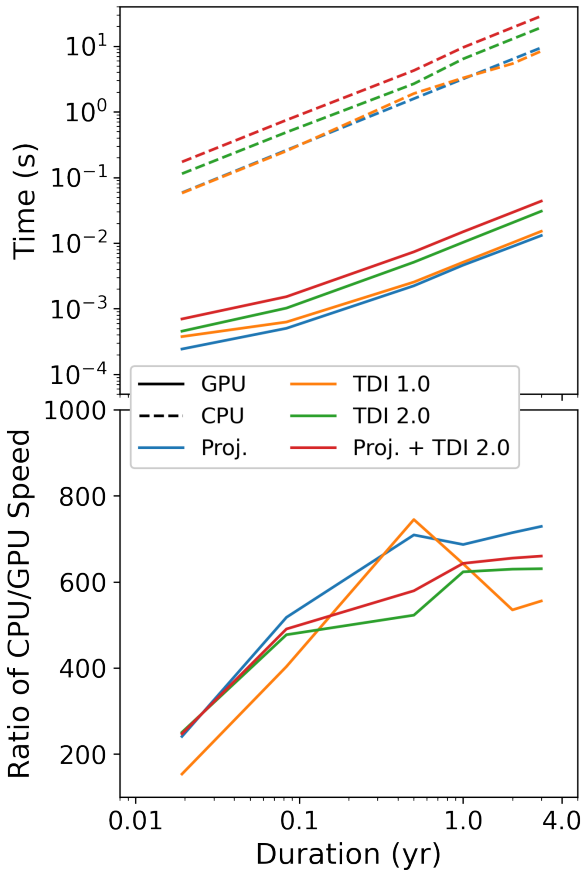

Figure 4. Performance of `fastlisaresponse`. The horizontal axis represents the simulation duration. In the top panel, GPU and CPU performance are shown in solid and dashed lines, respectively. The timing of the projection portion is shown in blue. TDI generations 1 and 2 are shown in orange and green, respectively. The total evaluation time of the response, combining the projection portion and TDI 2, is shown in red. The bottom panel shows the runtime ratio. These timing tests were performed on an NVIDIA Tesla A100 GPU and a single AMD EPYC 7713 CPU.

is representative of the parameter space for Galactic binaries.

For each point in the grid, we performed three MCMC runs. For all three runs, a signal built from ESA orbits was injected. Then, each run can test one of three different orbital configurations in the template waveforms. The first template was generated with ESA orbits. This run acts as a *control* with a template-signal fitting factor of 1, indicating no inherent bias in the template production. The second run involved using the equal-armlength orbits ("EQ") to generate the template. In this case, the fitting factor is strictly less than 1, allowing us to investigate the parameter-estimation bias incurred from using incorrect LISA constellation orbital properties. A third test that we performed is a middle ground between the first two, where the projection operation (section IV B) is performed with ESA orbits, while the TDI operation (section IV C) is computed with equal-armlength orbits. This model tests whether switching to accurate orbits in the projection portion of the response is sufficient for proper data analysis (this would allow for accurate projections while maintaining the efficiency of analytic TDI methods that require an equal-armlength configuration). We will refer to this *hybrid* orbital arrangement as "ESAEQ".

Parameter	Lower Bound	Upper Bound
$\ln A$	$\ln(A^* \times 10^{-2})$	$\ln(A^* \times 10^{2})$
f_0 (mHz)	0.5	12
$\dot{f}_0$ (Hz^2)	$\dot{f}(f_0^*, \mathcal{M}_c = 10^{-3} M_\odot)$	$\dot{f}(f_0^*, \mathcal{M}_c = 1 M_\odot)$
$\cos\iota$	-1	1
ϕ_0	0	2π
ψ	0	π
λ	0	2π
$\sin\beta$	-1	1

Table I. Prior distributions used in our analysis. The priors on f_0, ϕ_0, ψ, and λ are uniform distributions. The prior on the amplitude, A, is log-uniform, spanning two orders of magnitude above and below the injected value of A^*. The inclination prior (ι) and the ecliptic latitude prior (β) are uniform in $\cos\iota$ and $\sin\beta$, respectively. The prior on $\dot{f}_0$ is adjusted to ensure reasonable values for the chirp of a Galactic binary, given an initial frequency f_0.

The prior distributions used on the parameters are given in table I. Uniform distributions were used for parameters f_0, ϕ_0, ψ, and λ. The inclination (ι) prior is uniformly distributed in the cosine of the inclination angle. The ecliptic latitude (β) prior is uniform in $\sin\beta$. The prior on A is log-uniform between $A^* \times 10^{-2}$ and $A^* \times 10^{2}$, where A^* is the injected amplitude. Like the prior on A, the prior on the frequency derivative ($\dot{f}_0$) is adjusted based on the injection to ensure physically reasonable, but also encapsulating, limits on the chirp mass of the system. The frequency derivative is given a uniform prior, with boundaries computed as $\dot{f}_0^{\mathrm{min,max}} = \dot{f}_{\mathrm{GW}}(f_0, \mathcal{M}_c^{\mathrm{min,max}})$, see eq. (7), with physically representative $\mathcal{M}_c^{\mathrm{min}} = 10^{-3} M_\odot$ and $\mathcal{M}_c^{\mathrm{max}} = 1 M_\odot$.

The assessment of bias from MCMC results is visualized in fig. 5. In this figure, we examine a Galactic binary with parameters $\{\mathrm{SNR}, \beta, (f, \dot{f})\} = \{500, 1.2, (7 \times 10^{-4}\,\mathrm{Hz}, 1 \times 10^{-19}\,\mathrm{Hz\,s^{-1}})\}$. We focus the estimation of ecliptic longitude and latitude, shown on the horizontal and vertical axes, respectively. The mean estimates from parameter estimation runs with four different orbital configurations are shown with scatter points. The four configurations are ESA (orange star), EQ (blue plus), ESAEQ (green dot), and EQESA (purple dot). As the name implies, EQESA involves using equal-armlength orbits for the projections and ESA orbits for

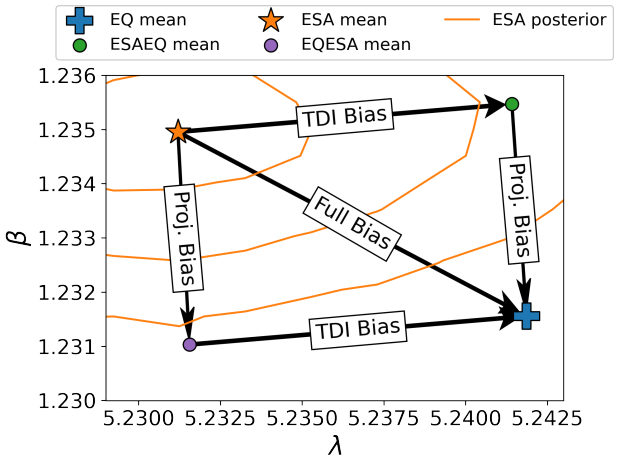

Figure 5. An example of bias created using inaccurate orbital information. The true sky localization is represented by the orange star, coincidental with the mean of the ESA template run. The orange lines represent the 1σ, 2σ, and 3σ contours for the two-dimensional marginalized distribution using accurate ESA orbits. The EQ template mean is shown with a blue plus. Connecting the ESA mean to the EQ mean yields the overall bias when using the EQ orbital configuration. ESAEQ and EQESA configuration means are shown with green and purple dots, respectively. Connecting ESA to ESAEQ represents the bias incurred by using equal-armlength orbits in the projection portion of the response function. Alternatively, connecting the ESA mean to the ESAEQ mean reveals the bias when using equal-armlength orbits in the TDI portion of the response function. The two bias components seem, visually, to add linearly. This is because the overall bias is small enough. In general, the overall bias from using incorrect orbits is inherently non-linear, as the TDI portion operates directly on the projections.

the TDI computation.

Connecting the ESA mean to the EQ mean in fig. 5 gives the overall bias on the sky location when using equal-armlength orbits. The contributions of the projection and TDI portions of the response to the overall bias can also be determined: the projection bias connects the mean of the ESA configuration to the mean of the EQESA configuration. The arrow from the ESA mean to the ESAEQ mean gives the bias from TDI. While the visualization appears here to indicate that the biases add linearly, this process is inherently non-linear as any projection bias will feed into the TDI computation. In addition, the two-dimensional marginalized posterior distribution for the ESA orbital configuration is shown as its 1σ, 2σ, and 3σ contours (orange lines). For the source studied here, the bias in the ESAEQ configuration is beyond 1σ, indicating the TDI bias from equal-armlength orbits is significant for higher SNR sources. We will discuss these effects further below.

This bias examination over the full posterior distributions was performed for each point on our multidimensional parameter grid. Here, we highlight important two-dimensional marginalized posterior distributions to help illustrate the effect of using incorrect orbits. Figures 6 and 7 show the two-dimensional posteriors in $f_0 - \dot{f}_0$ (initial frequency and frequency derivative) and $\beta - \lambda$ (sky localization angles), respectively. These parameter are of high scientific interest, as they characterize the intrinsic Galactic binary systems and provide opportunities for electromagnetic counterpart observations.

These figures are grouped by SNR (30 on the left and 500 on the right). Within each SNR group, $\beta = 1.23$ is shown in the first row and $\beta = 0.09$ is shown in the bottom row. Then, from left to right, within each row, are the low $f_0 - \dot{f}_0$, middle $f_0 - \dot{f}_0$, and high $f_0 - \dot{f}_0$ cases. For each MCMC run shown in these plots, the log-likelihood and mismatch marginalized over phase for the misaligned orbits at the injection parameters are given in table II. Also in table II, the log-likelihood and mismatch are given at the best-fit location (maximum log-likelihood) found by the sampler.

The resulting observations are fairly consistent across all three figures. Low-SNR sources generally show minimal relative bias from orbital misalignment. This does not strictly hold for the sky localization distributions of the high and middle $f_0 - \dot{f}_0$ sources. For the high $f_0 - \dot{f}_0$ source at higher polar latitude, this effect barely reaches 1σ deviation for the EQ comparison, and is almost entirely in the ecliptic longitude. For the nearly-planar source localization, the relative bias reaches beyond 2σ. This is expected, as a larger projection of the GW signal along the plane of LISA motion helps constraining more tightly the ecliptic longitude. In the middle $f_0 - \dot{f}_0$ configurations, the near-planar latitude system displays a bias that is close to, but slightly less than, 1σ. The nearly-polar case shows minimal bias. The frequency dependence appears clearly: tending towards higher frequencies at lower SNR will produce observable biases in the source localization. In addition to these findings on the EQ configuration, we observe that the ESAEQ configuration is perfectly acceptable for lower SNR, as it does not produce any observable bias.

In all high-SNR cases, the EQ templates produce strong biases on the sky localization of the sources with higher frequencies. The ESAEQ recovery of the signal is biased by 1σ or greater for all frequency configurations at high ecliptic latitudes; this reveals that the TDI portion alone contributes an observable bias on these parameters. At lower ecliptic latitudes, in the ESAEQ configuration, smaller the bias (less than 1σ) are observed for middle and low $f_0 - \dot{f}_0$. For the high $f_0 - \dot{f}_0$ case at near-planar ecliptic latitude, the bias in the ESAEQ case does reach 1σ.

These high-SNR cases also produce observable biases on the other parameters (there are minimal biases on the other parameters at low SNR). As an example, fig. 7 shows posterior distributions for f_0 and $\dot{f}_0$. In the low $f_0 - \dot{f}_0$ cases, $\dot{f}_0$ is not well-constrained such that we only

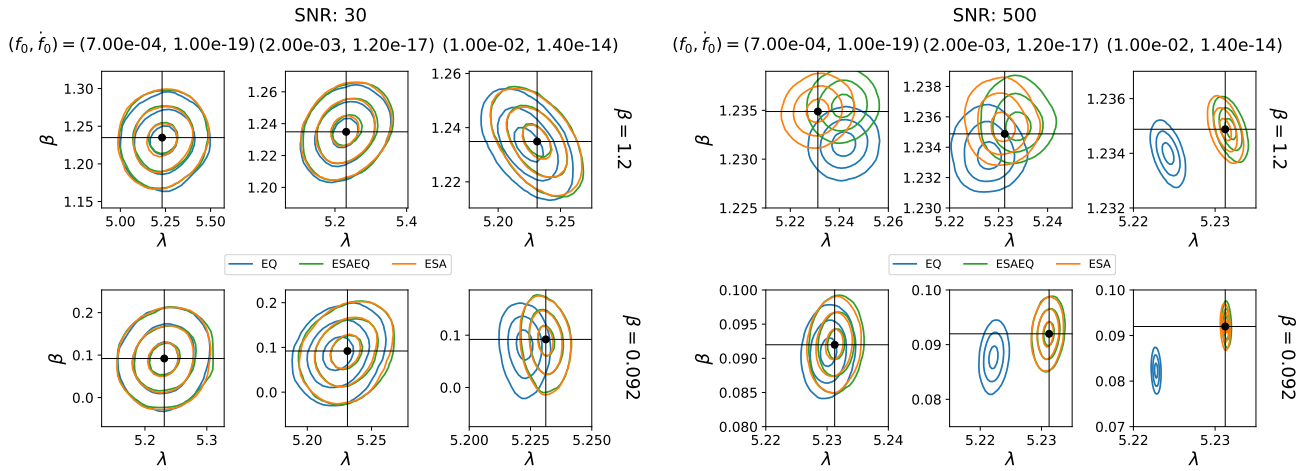

Figure 6. Two-dimensional marginalized posterior distributions on the sky location, across a variety of injected parameters and SNRs. The left set of plots shows an SNR of 30 and the right set an SNR of 500. The lower SNR is chosen to resemble a high number of detectable and characterizable Galactic binaries. The higher SNR represents the best sources we expect to detect. Within each SNR set, the columns represent $(f_0, \dot{f}_0)$ combinations of $(7 \times 10^{-4}, 1 \times 10^{-19})$, $(2 \times 10^{-3}, 1.2 \times 10^{-17})$, $(1 \times 10^{-2}, 1.4 \times 10^{-14})$ from left to right, respectively. The top (bottom) row in each SNR set is for the off-planar (on-planar) sources with $\beta = 1.2$ ($\beta = 0.092$). Within each subplot, the posterior distributions are shown for three orbital models: ESA in orange, EQ in blue, and ESAEQ in green. The deviation of the mean of the EQ or ESAEQ models compared to the mean of the ESA model gives the bias associated with the two inaccurate orbital configurations.

SNR	β	$f_0 - \dot{f}_0$	Orbit	Inj. $\ln \mathcal{L}$	Inj. MM	Best $\ln \mathcal{L}$	Best MM
30	1.2	low	EQ	-2.9×10^{-1}	1.7×10^{-2}	-2.0×10^{-1}	-1.3×10^{-3}
30	1.2	low	ESAEQ	-8.0×10^{-2}	1.1×10^{-2}	-1.9×10^{-1}	-1.3×10^{-2}
30	1.2	mid	EQ	-3.2×10^{-1}	1.7×10^{-2}	-2.3×10^{-1}	5.2×10^{-3}
30	1.2	mid	ESAEQ	-7.9×10^{-2}	1.1×10^{-2}	-1.9×10^{-1}	1.8×10^{-3}
30	1.2	high	EQ	-1.1	1.4×10^{-2}	-2.7×10^{-1}	-7.0×10^{-3}
30	1.2	high	ESAEQ	-7.0×10^{-2}	7.8×10^{-3}	-1.6×10^{-1}	3.2×10^{-3}
30	0.092	low	EQ	-3.9×10^{-1}	1.4×10^{-2}	-2.8×10^{-1}	-2.3×10^{-3}
30	0.092	low	ESAEQ	-9.8×10^{-2}	1.0×10^{-2}	-1.9×10^{-1}	-7.1×10^{-4}
30	0.092	mid	EQ	-6.8×10^{-1}	1.4×10^{-2}	-3.1×10^{-1}	7.4×10^{-3}
30	0.092	mid	ESAEQ	-9.7×10^{-2}	1.0×10^{-2}	-1.5×10^{-1}	-1.2×10^{-2}
30	0.092	high	EQ	-8.2	2.0×10^{-2}	-3.8×10^{-1}	-1.2×10^{-2}
30	0.092	high	ESAEQ	-8.6×10^{-2}	7.7×10^{-3}	-2.1×10^{-1}	5.9×10^{-3}
500	1.2	low	EQ	-8.0×10^{1}	1.7×10^{-2}	-2.8×10^{1}	2.5×10^{-4}
500	1.2	low	ESAEQ	-2.2×10^{1}	1.1×10^{-2}	-4.5	2.9×10^{-4}
500	1.2	mid	EQ	-8.8×10^{1}	1.7×10^{-2}	-3.3×10^{1}	-9.8×10^{-5}
500	1.2	mid	ESAEQ	-2.2×10^{1}	1.1×10^{-2}	-6.3	-3.3×10^{-4}
500	1.2	high	EQ	-3.1×10^{2}	1.4×10^{-2}	-3.0×10^{1}	4.8×10^{-4}
500	1.2	high	ESAEQ	-1.9×10^{1}	7.8×10^{-3}	-4.3	9.4×10^{-5}
500	0.092	low	EQ	-1.1×10^{2}	1.4×10^{-2}	-5.9×10^{1}	-6.3×10^{-4}
500	0.092	low	ESAEQ	-2.7×10^{1}	1.0×10^{-2}	-1.3×10^{1}	5.2×10^{-4}
500	0.092	mid	EQ	-1.9×10^{2}	1.4×10^{-2}	-5.3×10^{1}	3.2×10^{-4}
500	0.092	mid	ESAEQ	-2.7×10^{1}	1.0×10^{-2}	-1.3×10^{1}	7.7×10^{-4}
500	0.092	high	EQ	-2.3×10^{3}	2.0×10^{-2}	-6.5×10^{1}	2.2×10^{-4}
500	0.092	high	ESAEQ	-2.4×10^{1}	7.7×10^{-3}	-8.9	-4.2×10^{-4}

Table II. Log-likelihood and mismatch information for all parameter-estimation runs tested. The first four columns label each run with the SNR, ecliptic latitude (β), $f_0 - \dot{f}_0$ regime, and orbital configuration from left to right, respectively. The next two columns labelled "Inj. $\ln \mathcal{L}$" and "Inj. MM" give the log-likelihood and mismatch at the injection parameters marginalized over the initial phase. The final two columns represent the log-likelihood and mismatch at the best-fit location found by the sampler (not marginalized over phase). A mismatch of less than zero (overlap > 1) can occur when the fitting-factor of the model is not 1 as in the cases displayed here.

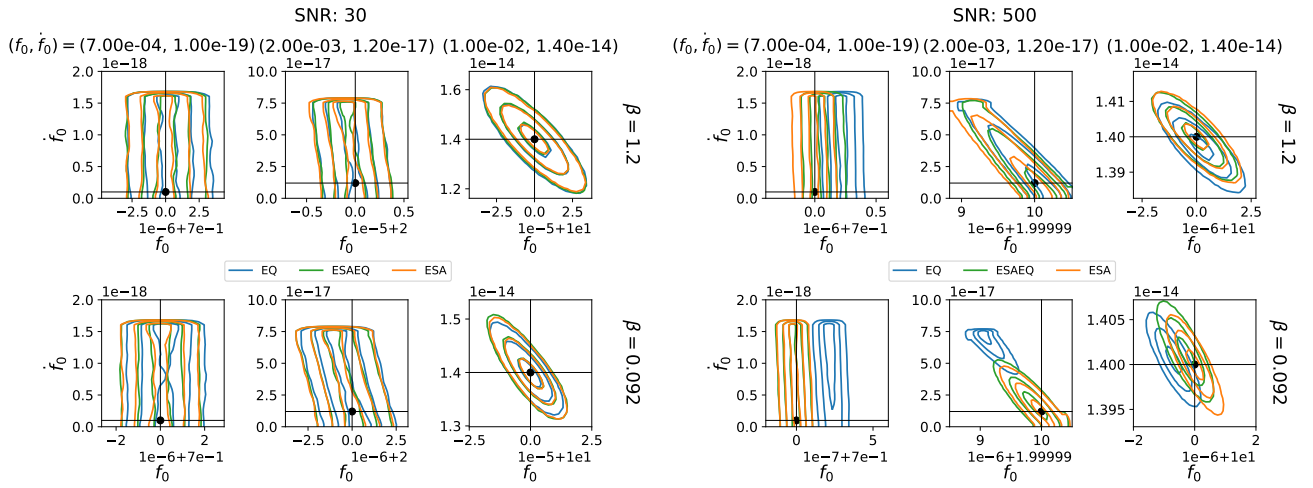

Figure 7. Two-dimensional marginalized posterior distributions, with f_0 along the horizontal axis and $\dot{f}_0$ along the vertical axis. These posteriors are shown across a variety of injected parameters and SNR combinations. C.f. fig. 6 for details on the arrangement of the subplots.

study the effect of orbital misalignment on f_0 distributions. At both latitudes, frequency is biased by more than 3σ for the EQ configuration. The ESAEQ configuration does not exhibit large biases at low latitudes; at high ecliptic latitude, however, f_0 is biased by almost 2σ. The high $f_0 - \dot{f}_0$ posteriors are highly Gaussian, an indication that $\dot{f}_0$ is well constrained. The EQ biases at high latitude are mainly in $\dot{f}_0$ (around 1σ) with only a small shift in f_0 (less than 1σ). The ESAEQ case does not show any observable bias, indicating that the overall bias is mainly due to the projection portion of the response. For the low-latitude case, we observe an EQ bias almost entirely in frequency ($\sim 3\sigma$), with a noticeable contribution from the ESAEQ setup ($\sim 1\sigma$ bias). As expected, the frequency derivatives of high-SNR, middle $f_0 - \dot{f}_0$ sources are only slightly constrained, and highly correlated with f_0 (unlike in the low $f_0 - \dot{f}_0$ case). At high latitude, the EQ and ESAEQ setups are only biased by about 1σ or less. The low-latitude case shows that using EQ templates yields very strong biases in frequency and frequency derivative (towards the maximum values allowed by the chosen prior). The ESAEQ template only shows a small bias, indicating that the overall bias originates in the projection portion of the response.

In addition to the parameters discussed above, we also observe a bias on the amplitude A of the source. This is expected, because a slight change in orbital configurations (even while maintaining other parameters fixed) creates slight deviations in the template SNR; the amplitude space is then explored by the sampler so that the template SNR matches that of the injected signal, which ultimately leads to bias on the amplitude. The chirp mass can be determined using the posterior distributions for f_0 and $\dot{f}_0$, c.f., eq. (7). Then, the chirp mass, f_0, and A can be combined to determine the luminosity distance. As a consequence, slight errors on the amplitude (due to variations of the SNR because of wrong orbital assumptions) will lead to errors on luminosity distance of the source.

VI. DISCUSSION

The creation of a fast generic time-domain response function for LISA will finally allow the LISA Community to perform full Bayesian inference with accurate LISA constellation orbital information. Prior to the creation of this fast response function, all LISA-based analyses were either too slow to scale to the level necessary for MCMC-style techniques or were performed with templates built from approximated methods. Now, with the construction of any time-domain template that is fast enough for MCMCs, there is an immediately implementable response function ready for the task.

This not only produces a more accurate analysis, but it also allows us to perform a wider range of tests to understand the impact of various response approximations. This is an important question, because, as previously mentioned, all current analysis codes use such approximations, such as the equal-armlength setup examined in this paper. The main reason for this is that these orbits are analytic, which allows for a variety of flexible implementations.

The advantage of analytic orbits manifests differently in the link projection and TDI portions of the LISA response. In the projections, due to LISA's time-dependent motion, the positions of the spacecraft have to be computed for a given array of times. In fast models, where the response can be computed on a *sparse grid* in the time-domain (such as `FastGB` [24]), the projections remain tractable with more accurate orbits, because they only involves computing spacecraft locations at given times.

When evaluating frequency-domain waveforms and response functions (such as `bbhx` [22]), the time-frequency correspondence of each harmonic mode is used to determine the position of the LISA constellation at a given frequency. Therefore, if the orbital information is input into the response code as an interpolant, the projections in the frequency domain waveforms will also be quite flexible to orbital adjustments.

During TDI, the frequency domain is greatly affected by the equal-armlength approximation. Currently, all MCMC-capable response codes evaluate the TDI combinations in the frequency domain. Working in the frequency domain and assuming equal-armlengths vastly simplifies the TDI computation. In the time-domain response function used in this work, each term in eq. (17) or eq. (18) is interpolated and computed separately and then added together. The frequency-domain equal-armlength response assumes the delays are all constant and equivalent to L/c, where L is the chosen rigid armlength of the LISA constellation. This reduces the necessary interpolation in the time-domain to basic phase shifts in the frequency-domain, greatly simplifying and accelerating the computation.

It is important to understand how well the equal-armlength approximation performs in order to understand how useful the current frequency-domain response functions will be to LISA analysis. As was presented in the previous section, it is clear equal-armlength orbits will not be accurate enough for parameter estimation purposes. They are mostly fine at lower SNRs, but at any higher SNRs, the biases, relative to the increased precision in our parameter estimation, are not acceptable. However, all response codes in use could easily be updated for accurate orbits in the projection portion of the computation. Therefore, if the ESAEQ templates exhibited an acceptably small bias, the frequency-domain response codes would probably remain sufficient for parameter estimation.

However, it is clear that, at higher (but reasonable) SNRs, even this type of half-approximation does not hold for all sources tested. *This incursion of bias from equal-armlength orbits indicates frequency-domain response codes will have to be updated to a more accurate response formulation for use in parameter estimation.* One potential solution is to perform the analysis in the time-frequency or wavelet domain [25], but this requires further investigation.

Due to the ability of the EQ templates to produce a high empirical SNR and reasonable, albeit biased, posteriors, the equal-armlength approximation will probably be useful in search, especially if it remains a much faster computation. This will require verification in future work, especially for sources that require more than one year of observation for detection (our tests in this work are over one year). From fig. 3, we see that the difference in the orbits, as well as the difference in the waveforms that manifest from those orbits, gradually increases over longer durations.

In this work, we have only tested Galactic binaries, the simplest of LISA sources in terms of their waveform complexity. The inaccuracy of orbital information plays a large role for these sources, since their GW signals are in the data stream for the entire mission duration. Extreme mass-ratio inspirals (EMRIs) are similarly long-lived sources that might be observable over the mission lifetime, or a significant fraction. The signals from such sources are expected to have a much higher complexity than that of Galactic-binary signals, or indeed any other source in the LISA catalog. Nevertheless, a Galactic-binary signal essentially resembles that from a particular (albeit unrealistic) type of EMRI: an early-stage quasi-circular EMRI, where a single, slowly evolving, harmonic mode dominates the signal. Therefore, we conjecture that the findings presented here on orbital information accuracy will, at best, similarly apply to EMRIs—rendering equal-armlength approximations to response functions unusable in parameter estimation for high-SNR sources. At worst, the bias incurred from even full treatments of the orbit and response function might be compounded by the higher waveform complexity, and might end up as the dominant source of error for these precisely modeled sources.

Massive black-hole binaries are going to be observed for much shorter stretches of time than Galactic binaries (a few days or weeks) [e.g., 26]. They will also be observed at higher SNRs. When analyzing massive black-hole binaries, due to their short-duration signal, the equal-armlength orbital parameters can be optimized (see section IV D) to a smaller number of time points, potentially making this approximation more accurate. With that said, in a full parameter estimation or global fit setting, it would be unwise to assume different underlying orbits for different sources. Therefore, we categorize this technique as another potential option when *searching* for massive black-hole binaries. Given fixed underlying orbits, we draw the conclusion that equal-armlength orbits for massive black-hole binaries are also not good enough for inference, viewing our Galactic binary example once again as conservative. The higher SNRs of massive black-hole binaries will risk larger bias, and may be more affected by instantaneous differences between accurate and inaccurate orbits if their mergers lie during a small time segment where the orbits are less similar.

In a more realistic analysis setup, the parameters of various overlapping sources must be estimated simultaneously by a *global fit* technique. One can reasonably think that orbital parameters will be be included in this global fit, and therefore argue that the findings presented do not apply. However, if one uses a simple orbital model, such as the equal-armlength configuration, the fitted orbital parameters for the simple model will be, at best, identical to the parameters fitted to the true orbits, i.e., the exact same parameters used in this conservative study. Therefore, we believe that the results presented here remain valid in a global fit analysis.

Using a more accurate response function for parame-

ter estimation, and in particular a more realistic orbital model, as we advocate here, requires some knowledge of the positions of the LISA spacecraft, as well as the light travel times along the 6 constellation links. The light travel times can be computed from the spacecraft position, velocity and acceleration vectors [e.g., 27]. The LISA spacecraft follow the free-falling test masses they host, such that their trajectories are, to great precision, simple geodesics in the Solar System. These geodesics are easily computed using standard ephemerides and some initial conditions (this is how ESA orbits used in this paper have been calculated). The initial conditions are provided by on-ground estimation of absolute positions and velocities of the LISA satellites by ESA's ESTRACK system. Ground-tracking position estimates have $\sim 10\,\mathrm{km}$ precision [28]; this is several orders of magnitude below the difference between the fitted equal-armlength orbits and ESA orbits used in this study (c.f., fig. 3). Therefore, we are confident that spacecraft trajectories will be known with sufficient accuracy not to be a limiting factor in source parameter recovery.

VII. CONCLUSION

In this article, we presented a new open-source tool to computing the general LISA TDI response to GWs in the time domain, for arbitrary orbits and waveforms. The code is flexible and can be run on both CPU and GPU; however, a specialized implementation of the response function for GPUs makes fast enough ($\sim 10\,\mathrm{ms}$ for 4 years worth of data) to use in stochastic sampling algorithms, such as MCMCs.

We took advantage of this newly-found speed to test the parameter-estimation bias incurred by current analysis codes that adopt equal-armlength LISA orbits for the sake of computational efficiency. We found that the resulting waveform templates produce unacceptable bias in source frequency and position for loud Galactic-binary waveforms; however the corresponding matched-filtering SNR would still be appropriate to identify sources in a searches. We also tested hybrid templates that use accurate orbits for GW projection, but equal armlengths for TDI. Such templates could be implemented with straightforward modifications to equal-armlength codes, and they enable accurate parameter recovery except for the loudest sources (SNR ~ 500), where they produce significant bias.

These results indicate the importance of including accurate orbits in the LISA signal model, and they highlight the need to upgrade current approximate methods in preparation for real data. We expect that these conclusions would be strengthened for GW signals more complex than the Galactic binaries studied here, such as EMRIs or massive black-hole binaries. However, more studies must be carried out to obtain a definitive answer on the impact of orbital approximations for these types of sources. Other features of real data, such as nonstationary noise, gaps, glitches, etc., would also impact data-analysis performance, but we do not expect them to affect the general findings presented here.

ACKNOWLEDGMENTS

J.-B.B. was supported by a NASA postdoctoral fellowship administered by USRA. A.J.K.C. acknowledges support from the NASA LISA Preparatory Science grant 20-LPS20-0005. M.V. was supported by the NASA LISA study office. This research was supported in part through the computational resources and staff contributions provided for the Quest/Grail high performance computing facility at Northwestern University. This paper also employed use of `SciPy` [29] and `Matplotlib` [30]. Part of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004).

[1] B. P. Abbott *et al.* (LIGO and Virgo Scientific Collaborations), Phys. Rev. X **9**, 031040 (2019), arXiv:1811.12907 [astro-ph.HE].
[2] LIGO–Virgo–Kagra Collaboration, Living Reviews in Relativity **21**, 3 (2018), arXiv:1304.0670 [gr-qc].
[3] P. Amaro-Seoane and et al, (2017), arXiv:1702.00786 [astro-ph.IM].
[4] S. Marsat and J. G. Baker, "Fourier-domain modulations and delays of gravitational-wave signals," (2018), unpublished, arXiv:1806.10734 [gr-qc].
[5] LISA Science Study Team, "LISA Science Requirements Document," (2018), atrium.in2p3.fr.
[6] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, Pub. Astron. Soc. Pacific **125**, 306 (2013), arXiv:1202.3665 [astro-ph.IM].
[7] W. Vousden, W. Farr, and I. Mandel, MNRAS **455**, 1919 (2016), arXiv:1501.05823 [astro-ph.IM].
[8] J. Goodman and J. Weare, Comm. App. Math. Comp. Sci. **5**, 65 (2010).
[9] M. Tinto, F. B. Estabrook, and J. W. Armstrong, Phys. Rev. D **69**, 082001 (2004), arXiv:gr-qc/0310017.
[10] M. Vallisneri, Phys. Rev. D **71**, 22001 (2005), arXiv:gr-qc/0407102 [gr-qc].
[11] M. Vallisneri and C. R. Galley, Class. Quant. Grav. **29**, 124015 (2012), arXiv:1201.3684 [gr-qc].
[12] S. V. Dhurandhar, K. R. Nayak, S. Koshti, and J. Y. Vinet, Classical and Quantum Gravity **22**, 481 (2005), arXiv:0410093 [arXiv:gr-qc].
[13] W. Martens and E. Joffre, J. Astronaut. Sci. **68**, 402– (2021), arXiv:2101.03040 [gr-qc].

[14] J.-B. Bayle, A. Hees, M. Lilley, and C. Le Poncin-Lafitte, "LISA Orbits," (2022).
[15] M. Vallisneri, Phys. Rev. D **71**, 022001 (2005), arXiv:gr-qc/0407102.
[16] M. Katz and J.-B. Bayle, "lisa-on-gpu," (2022).
[17] S. van der Walt, S. C. Colbert, and G. Varoquaux, Computing in Science & Engineering **13**, 22 (2011).
[18] S. Behnel, R. Bradshaw, C. Citro, L. Dalcin, D. S. Seljebotn, and K. Smith, Computing in Science & Engineering **13**, 31 (2011).
[19] R. McGibbon and Y. Zhao, "npcuda-example," (2012).
[20] J.-B. Bayle, Q. Baghi, and A. Renzini, "LISA GW Response," (2022).
[21] M. Staab, J.-B. Bayle, and O. Hartwig, "PyTDI," (2022).
[22] M. L. Katz, Phys. Rev. D **105**, 044055 (2022), arXiv:2111.01064 [gr-qc].
[23] N. Cornish and T. Robson, J. Phys. Conf. Ser. **840**, 012024 (2017), arXiv:1703.09858 [astro-ph.IM].
[24] N. J. Cornish and T. B. Littenberg, Phys. Rev. D **76**, 083006 (2007), arXiv:0704.1808 [gr-qc].
[25] N. J. Cornish, (2020), 10.1103/PhysRevD.102.124038, arXiv:2009.00043 [gr-qc].
[26] M. L. Katz, S. Marsat, A. J. K. Chua, S. Babak, and S. L. Larson, Phys. Rev. D **102**, 023033 (2020), arXiv:2005.01827 [gr-qc].
[27] A. Hees, S. Bertone, and C. Le Poncin-Lafitte, Phys. Rev. D **89**, 064045 (2014), arXiv:1401.7622 [gr-qc].
[28] W. Martens and E. Joffre, (2021), personal communication.
[29] P. Virtanen, R. Gommers, T. E. Oliphant, and SciPy 1.0 Contributors, Nature Methods **17**, 261 (2020).
[30] J. D. Hunter, Computing In Science & Engineering **9**, 90 (2007).