

Fast and Fourier: Extreme Mass Ratio Inspiral Waveforms in the Frequency Domain

Lorenzo Speri¹, Michael L. Katz^{1,2}, Alvin J. K. Chua^{3,4,*}, Scott A. Hughes⁵, Niels Warburton⁶, Jonathan E. Thompson⁷, Christian E. A. Chapman-Bird⁸, Jonathan R. Gair¹

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

²NASA Marshall Space Flight Center, Huntsville, Alabama 35811, USA

³Department of Physics, National University of Singapore, Singapore 117551

⁴Department of Mathematics, National University of Singapore, Singapore 119076

⁵Department of Physics and MIT Kavli Institute, Cambridge, MA 02139, United States

⁶School of Mathematics and Statistics, University College Dublin, Belfield, Dublin 4, Ireland

⁷Theoretical Astrophysics Group, California Institute of Technology, Pasadena, CA 91125, U.S.A.

⁸SUPA, University of Glasgow, Glasgow G12 8QQ, United Kingdom

Correspondence*:

Alvin J. K. Chua
alvincjk@nus.edu.sg

ABSTRACT

Extreme Mass Ratio Inspirals (EMRIs) are one of the key sources for future space-based gravitational wave interferometers. Measurements of EMRI gravitational waves are expected to determine the characteristics of their sources with sub-percent precision. However, their waveform generation is challenging due to the long duration of the signal and the high harmonic content. Here, we present the first ready-to-use Schwarzschild eccentric EMRI waveform implementation in the frequency domain for use with either graphics processing units (GPUs) or central processing units (CPUs). We present the overall waveform implementation and test the accuracy and performance of the frequency domain waveforms against the time domain implementation. On GPUs, the frequency domain waveform takes in median 0.044 seconds to generate and is twice as fast to compute as its time domain counterpart when considering massive black hole masses $\geq 2 \times 10^6 M_{\odot}$ and initial eccentricities $e_0 > 0.2$. On CPUs, the median waveform evaluation time is 5 seconds, and it is five times faster in the frequency domain than in the time domain. Using a sparser frequency array can further speed up the waveform generation, reaching up to 0.3 seconds. This enables us to perform, for the first time, EMRI parameter inference with fully relativistic waveforms on CPUs. Future EMRI models which encompass wider source characteristics (particularly black hole spin and generic orbit geometries) will require significantly more harmonics. Frequency-domain models will be essential analysis tools for these astrophysically realistic and important signals.

Keywords: gravitational wave, LISA, EMRI, data analysis, frequency domain

1 INTRODUCTION

The future space-based Laser Interferometer Space Antenna (LISA) will observe gravitational waves in the millihertz regime (Amaro-Seoane et al., 2017; Gair et al., 2022). This region of the gravitational wave spectrum is rich in sources, including Galactic double white dwarf (WD) binaries (Seto, 2001), massive black hole (MBH) binaries (Klein et al., 2016), and stellar origin black hole (SOBH) binaries early in their evolution (Klein et al., 2016; Moore et al., 2019). Another key source class is extreme mass-ratio inspirals (EMRIs) (Babak et al., 2007). These are compact binaries with a mass ratio $\mu/M \simeq 10^{-4} - 10^{-6}$, where $\mu \sim 1 - 100M_{\odot}$ is the mass of the orbiting secondary and $M \sim 10^5 - 10^7M_{\odot}$ is the mass of the MBH. EMRIs are expected to form in dense stellar clusters of galactic nuclei (Amaro-Seoane, 2018a; Pan and Yang, 2021), where their detection rate ranges from $\sim 1 - 10^4$ per year with observable signal-to-noise ratios (SNR) expected to be $\sim 20 - 1000$ over the duration of the signal (Babak et al., 2017; Amaro-Seoane et al., 2007; Porter, 2009). The details depend on the precise formation mechanism (Amaro-Seoane, 2018b; Berry et al., 2019a; Arca-Sedda and Capuzzo-Dolcetta, 2019; Gair et al., 2004), but it is anticipated that the majority of EMRIs are expected to have eccentricities at plunge in the range $0 < e_p \lesssim 0.2$, with a long tail extending to larger e_p (Babak et al., 2017). This high eccentricity is why EMRIs have some of the richest and most complicated gravitational waveforms of any compact binary system. The small mass ratio of EMRIs also means that they evolve slowly, typically completing $\sim 10^4 - 10^5$ orbits over years whilst in the LISA band. The long-lasting, complex waveforms of EMRIs present a substantial challenge for both the modeling of these binaries and the LISA data analysis task (Babak et al., 2009; Cornish, 2011; Babak et al., 2010, 2008; Arnaud et al., 2006; Barack and Pound, 2018; Chua and Cutler, 2022; Speri et al., 2022).

Modeling and extracting EMRI signals from the LISA data stream will provide binary system measurements with sub-percent level precision (Babak et al., 2017; Berry et al., 2019b). This precision will enable precise tests of general relativity (Barack and Cutler, 2007; Gair et al., 2013; Maselli et al., 2020, 2022; Barsanti et al., 2022). Additionally, studying EMRIs will enhance our understanding of the mass function of MBHs (Gair et al., 2010), the dense stellar environments in galactic cores (Amaro-Seoane et al., 2007), and the gas disks surrounding MBHs (Barausse et al., 2007; Barausse and Rezzolla, 2008; Gair et al., 2011; Yunes et al., 2011; Barausse et al., 2014, 2015; Speri et al., 2023). Furthermore, gravitational wave signals from EMRIs could be used to constrain cosmological parameters (MacLeod and Hogan, 2008; Laghi et al., 2021) or measure phase calibration errors (Savalle et al., 2022).

Extracting this wealth of information from EMRIs in the LISA data stream poses two significant challenges: accuracy and speed. The former is limited by the waveform modeling accuracy, whereas the latter is limited by the combined computational cost of the waveform generation, the discrete Fourier transform (DFT) and the inner product necessary to compute the likelihood function or any detection statistic. A first attempt to address these problems was presented in the release of the FASTEMRIWAVEFORMS (*few*) computational framework (Chua et al., 2021; Katz et al., 2021), which can compute fully relativistic waveforms rapidly. This package combines a set of standalone modules to create EMRI waveforms on both graphics processing unit (GPU) and central processing unit (CPU) hardware. It was shown that the GPU version is ≈ 2500 times faster than its CPU counterpart for Schwarzschild spacetimes. GPUs are designed with a large number of cores and specialized hardware that allows them to perform multiple tasks simultaneously, making them highly efficient in handling parallel processing tasks. Compared to CPUs, which typically have fewer cores, GPUs can process a much larger number of tasks concurrently, resulting in significantly improved performance for certain applications. GPUs are

particularly well-suited to efficient EMRI waveform generation because these signals consist of a large number of harmonic modes that are summed in parallel.

The waveform generation in `few` is currently performed in the time domain (TD), with typical speeds using GPUs and CPUs on the order of tens of milliseconds and tens of seconds respectively. For GPUs, the computational costs of the likelihood are negligible (less than a millisecond), therefore making the waveform generation the bottleneck of the computations. However, for CPUs all of the aforementioned operations contribute significantly to the overall data analysis cost. We expect that a frequency domain implementation of EMRI waveforms would reduce the likelihood evaluation cost. However, it is not immediately clear how the frequency domain implementation of the EMRI waveforms would comparatively perform on GPUs and CPUs.

In this work, we expand upon the `few` package by implementing a frequency domain (FD) EMRI waveform based upon Hughes et al. (2021). We provide a module that can compute EMRI waveforms for eccentric inspirals into a non-rotating black hole at adiabatic order (Hinderer and Flanagan, 2008; Hughes et al., 2021). The implementation supports use with both CPUs and GPUs and its accuracy and performance are explored. This work is fully reproducible using the software provided at https://github.com/lorenzsp/EMRI_FrequencyDomainWaveforms

The paper is organised as follows. We review the general construction of EMRI waveforms in Section 2.1 and we discuss the frequency domain implementation in Section 2.2. In Section 2.3 we present the data analysis tools that are used to test the accuracy of the waveform generation. Finally, in section 3 we present the results for the performance and accuracy tests of the frequency domain implementation.

2 METHODS

2.1 Extreme Mass Ratio Inspiral Waveforms

The gravitational waves emitted by an EMRI are among the most complicated compact binary signals due to their long duration and rich harmonic content. This complexity stems from the dynamics of the secondary object inspiralling into the central MBH. Such orbits are characterized by three degrees of freedom associated with the dimensionless fundamental frequencies of Kerr geodesic orbits $\Omega_{r,\theta,\phi}$ (Fujita and Hikida, 2009; Schmidt, 2002). These quantities are determined by the dimensionless spin magnitude of the MBH a , and the quasi-Keplerian orbital parameters of p (semi-latus rectum or separation in units of M), e (eccentricity), and $x_I \equiv \cos I$ (cosine of the angle I which describes the orbit's inclination from the equatorial plane). At adiabatic order, for a given system with primary mass M , secondary mass μ , initial orbital parameters (p_0, e_0, x_{I0}) , and initial phases $\Phi_{\varphi 0}, \Phi_{\theta 0}, \Phi_{r 0}$, the orbital evolution can be determined by solving the following system of ordinary differential equations:

$$\begin{aligned}\frac{d}{dt}p &= \frac{\mu}{M} f_p(a, p, e, x_I) \\ \frac{d}{dt}e &= \frac{\mu}{M} f_e(a, p, e, x_I) \\ \frac{d}{dt}x_I &= \frac{\mu}{M} f_{x_I}(a, p, e, x_I) \\ \frac{d}{dt}\Phi_{\varphi,\theta,r} &= \Omega_{\varphi,\theta,r}(a, p, e, x_I)/M.\end{aligned}$$

The orbital-element fluxes f_{p,e,x_I} account for the gravitational wave emission and are interpolated from pre-computed grids. In the `few` package, the system of ordinary differential equations is solved using an 8th-order Runge-Kutta integrator with adaptive step size, resulting in sparse orbital trajectories of ~ 100 data points in length.

Once the trajectory is fully specified by $\{p(t), e(t), x_I(t), \Phi_\varphi(t), \Phi_\theta(t), \Phi_r(t)\}$, we compute the complex gravitational-wave amplitudes $A_{lmkn}(p(t), e(t), x_I(t))$. Each amplitude mode is associated with the multipole l of the radiation and the harmonic $\{m, n, k\}$ of the fundamental orbital frequencies, where m , k , and n label the azimuthal, polar, and radial modes, respectively. In the `few` package, the amplitudes and phases are interpolated using a “not-a-knot” cubic spline to handle the large number of modes using a sparse number of points. In the interest of computational efficiency, we control the number of harmonic modes with a further mode-selection step. We cumulatively sum modes in order of decreasing power, and truncate this summation at a threshold value of $(1 - \epsilon)$ times the total power. Only the modes that pass this threshold are included in the waveform computation (see Katz et al. (2021) for further details).

The gravitational wave signal at a large distance from the source can be represented in terms of the complex time-domain dimensionless strain (Hughes et al., 2021):

$$h = h_+ - ih_\times = \frac{\mu}{d_L} \sum_{lmkn} A_{lmkn}(t) S_{lmkn}(t, \theta) e^{im\phi} \exp\{-i[m\Phi_\varphi(t) + k\Phi_\theta(t) + n\Phi_r(t)]\}, \quad (1)$$

where t is the time of arrival of the gravitational wave at the Solar System barycenter, θ is the source-frame polar viewing angle, ϕ is the source-frame azimuthal viewing angle, and d_L is the luminosity distance of the source. The functions $S_{lmkn}(t, \theta)$ are spin-weighted spheroidal harmonics. For inspiral into a Schwarzschild black hole (the focus of this analysis), these functions are identical to the spin-weighted spherical harmonics. For the general Kerr case, they differ from the spherical harmonics (though they can be very usefully expanded into spherical harmonics (Hughes, 2000)), and in particular they depend on the orbital frequencies. Because the orbital frequencies evolve with time, the spheroidal harmonics evolve as well (Hughes et al., 2021), complicating the description of the waveform. Future work will describe how we handle this complication. In any case, the usual transverse-traceless gravitational wave polarizations can be obtained from the real and imaginary parts of the complex conjugate of h .

A key element to be noted in Eq. (1) is that the EMRI waveform can be viewed as a sum of simple “voices”, where each voice corresponds to a mode (l, m, k, n) . The time-domain waveform can thus be rewritten in terms of a sum over voices V :

$$h = \sum_V H_V(t) e^{-i\Phi_V(t)}, \quad (2)$$

where each voice is characterized by a complex amplitude $H_V(t)$ and a phase $\Phi_V(t)$. The voice-by-voice decomposition was suggested to one of the present authors by L. S. Finn, and first presented in Hughes (2001)¹. In this work, we provide the first ready-to-use implementation of this decomposition for frequency-domain EMRI waveforms.

¹ Note that a frequency domain waveform implementation was also presented in Piovano et al. (2020) with the limitation to circular Kerr orbits.

2.2 The frequency domain

EMRI waveforms can be represented in the frequency domain by using the stationary phase approximation. Because amplitude, phase, and frequency in an EMRI evolve slowly (as long as the two-timescale approximation is valid (Miller and Pound, 2021)), we expect the stationary phase approximation (SPA) to provide a high-quality approximation to the Fourier transform of the signal (Droz et al., 1999). Here, we review the computation of the Fourier transform of an EMRI waveform in the frequency domain presented in Hughes et al. (2021). We begin by taking the Fourier transform of each voice of Eq. (2) as follows:

$$\tilde{h}(f) \equiv \int_{-\infty}^{\infty} h(t)e^{2\pi ift} dt = \int_{-\infty}^{\infty} H(t)e^{i[2\pi ft - \Phi(t)]} dt \quad (3)$$

where we have not written the subscript for ease of notation. To compute the stationary phase approximation to the signal, it is necessary to expand the phase evolution, written as

$$\Phi(t) = \Phi(t_S) + 2\pi F(t - t_S) + \pi \dot{F}(t - t_S)^2 + \frac{\pi}{3} \ddot{F}(t - t_S)^3 + \dots \quad (4)$$

where

$$F \equiv \frac{1}{2\pi} \frac{d\Phi}{dt} \Big|_{t_S}, \quad \dot{F} = \frac{dF}{dt} \equiv \frac{1}{2\pi} \frac{d^2\Phi}{dt^2} \Big|_{t_S}, \quad \ddot{F} \equiv \frac{1}{2\pi} \frac{d^3\Phi}{dt^3} \Big|_{t_S}. \quad (5)$$

We have introduced the signal's instantaneous frequency and the instantaneous first and second derivatives of the frequency at the stationary time $t = t_S$ for which $F = f$. We expanded up to third order in the phase evolution because for some voices the frequency evolution is not monotonic. If the frequency derivative vanishes ($\dot{F} = 0$), the standard SPA Fourier transform is singular and cannot be used (Klein et al., 2014). This is highly relevant to EMRI signals because the frequency associated with many voices rises to a maximum and then decreases. In particular, this occurs for EMRI voices which involve harmonics of the radial frequency. This is because Ω_r becomes very small (approaching $|\Omega_\phi|$) in the weak-field limit, and goes to zero as systems approach the last stable orbit; a maximum of Ω_r exists between these two “small” limits.

We now use Eq. (4) to obtain the time-frequency correspondence for each voice, i.e. we find the time t such that $F(t) = f$. The frequency evolution of each voice $F(t) = (m\Omega_\varphi(t) + k\Omega_\theta(t) + n\Omega_r(t))/(2\pi M)$ is directly obtained from the trajectory evolution of each of the fundamental frequencies. During the waveform production process, the fundamental frequency evolution is interpolated with a cubic spline allowing for a good approximation using a third-order polynomial. Doing so, the condition $f = F(t)$ between two knots can be rewritten as:

$$f = A + B(t - t_i) + C(t - t_i)^2 + D(t - t_i)^3, \quad (6)$$

where A, B, C, D are real numbers given by the sum of the spline coefficients of three frequencies multiplied by their mode number, and t_i is the time at the beginning of the spline segment.

The fundamental frequency, fundamental phase, and amplitude spline coefficients are determined along the sparse trajectory prior to the final waveform summation phase (see Katz et al. (2021) for more information). Additionally, using all values of f and each set of frequencies found along the sparse trajectory for each voice, the spline segment in which each value of f falls is predetermined before the summation. Once in the summation, the goal is to determine $t(f)$ as efficiently as possible. This is difficult in cases where the voice is not monotonic because $t(f)$ is not a bijective mapping for this case: two values

of t can be represented by the same value of f . To address this issue, we solve for the cubic roots of Eq. 6 within each segment where f is found (one or two segments per voice depending on f). We take the real-valued roots and determine which one or two actually exist between the bounds of the segment we are currently examining (it can be two only for the segment in which the turnover occurs). Since the frequency evolution is slowly varying within each segment, the condition $\dot{F} = 0$ cannot occur more than one time per segment. These time values occurring within our segment are then used to determine all necessary quantities for waveform building because all informative splines are fit as a function of t .

We define $t_j(f)$ to be the j th time (first or second) at which $F(t) = f$, and write $\dot{F}_j \equiv \dot{F}[t_j(f)]$, $\ddot{F}_j \equiv \ddot{F}[t_j(f)]$. These higher-order derivatives are also computed using the fundamental frequency spline information. With that, for each value of f for each voice, the Fourier transform can be represented as:

$$\tilde{h}(f) \simeq \sum_{j=1}^N H[t_j(f)] e^{i[2\pi f t_j(f) - \Phi(t_j(f))]} \times \int_{-\infty}^{\infty} e^{-i\pi[\dot{F}_j(t')^2 + \ddot{F}_j(t')^3/3]} dt', \tag{7}$$

where $N \in \{1, 2\}$ is the number of time values associated with each value of f . To perform this integral, we set $\alpha = \gamma + 2\pi i \dot{F}$, with γ real and positive, define $\beta = 2\pi \ddot{F}$, and use

$$\int_{-\infty}^{\infty} e^{-\alpha t^2/2 - i\beta t^3/6} dt = \frac{2}{\sqrt{3}} \frac{\alpha}{|\beta|} e^{\alpha^3/3\beta^2} K_{1/3}(\alpha^3/3\beta^2), \tag{8}$$

where $K_n(z)$ is the modified Bessel function of the second kind. Taking the limit $\gamma \rightarrow 0$, we find

$$\tilde{h}(f) \simeq \frac{2}{\sqrt{3}} \sum_{j=1}^N H[t_j(f)] e^{i[2\pi f t_j(f) - \Phi(t_j(f))]} \frac{i\dot{F}_j}{|\dot{F}_j|} e^{-2\pi i \dot{F}_j^3/3\ddot{F}_j^2} K_{1/3}(-2\pi i \dot{F}_j^3/3\ddot{F}_j^2). \tag{9}$$

This result defines our frequency domain waveform. Expanding about “small” \ddot{F}_j , it can be shown that this yields the standard SPA when $\ddot{F}_j \rightarrow 0$ and $N = 1$ (see Hughes et al. (2021)).

As previously mentioned, this waveform is built into the framework of \mathbb{f}_{EW} . All parts of the waveform generation prior to the “waveform summation module” are identical to the original time-domain implementation (Katz et al., 2021). The waveform summation described above is implemented for both CPUs and GPUs. For the purposes of GPU parallelization, the summation is grouped by harmonic voice and segment of the splines defined in time. This creates a rectangular computation grid because all spline information occurs on the same sparse time evolution returned from the trajectory integrator yielding a constant number of segments across all spline constituents. On the GPU, we place each segment within each voice on a separate GPU block (larger parallelization unit). Within each block, we use the parallel threads (smaller parallelization unit) to evaluate all values of interest in frequency for that specific voice and segment. The waveform is then constructed by “atomically” adding from each of these blocks and threads to an overall waveform vector stored in GPU global memory.

2.3 Data analysis setup

An EMRI waveform for non-rotating black holes is uniquely determined in the reference frame of the Solar System barycenter by setting the following parameters:

$$\{M, \mu, p_0, e_0, d_L, \theta_S, \phi_S, \theta_K, \phi_K, \Phi_{\varphi 0}, \Phi_{r0}\}, \tag{10}$$

where θ_S and ϕ_S are the polar and azimuthal sky location angles given in the Solar System barycenter reference frame, and θ_K and ϕ_K are the azimuthal and polar angles describing the orientation of the orbital angular momentum. The relation to the source frame angles θ, ϕ can be found in Katz et al. (2021). We do not include the effect of the detector response function, but we leave this to future works (Marsat and Baker, 2018).

We check the performance and accuracy of the FD waveform over the EMRI parameter space confined to $\log(M/M_\odot) \in [\log 10^5, \log 10^7)$, $\log(\mu/M) \in [\log 10^{-6}, \log 10^{-4})$, $e_0 \in [0.001, 0.7)$ and $\Phi_{\varphi 0}, \Phi_{r0} \in [0, 2\pi)$ by uniformly drawing 5000 parameter realizations. The initial semi-latus rectum is fixed to obtain an inspiral with a time duration of $0.99T$, where T is the duration of the observation. We consider h_+ and h_\times in the Solar System barycenter as our two channels, and we fix $\{d_L = 1 \text{ Gpc}, \theta_S = \pi/3, \phi_S = \pi/3, \theta_K = \pi/3, \phi_K = \pi/3\}$ throughout this work. To check our implementation, we quantify how similar two waveforms are using the mismatch²

$$\mathcal{M}(a, b) = 1 - \sum_{\alpha} \frac{\langle a|b \rangle_{\alpha}}{\sqrt{\langle a|a \rangle_{\alpha} \langle b|b \rangle_{\alpha}}} \quad (11)$$

where we have introduced the inner product $\langle \cdot | \cdot \rangle$ of each channel α ,

$$\langle a(t)|b(t) \rangle = 4 \text{Re} \int_0^{\infty} \frac{\tilde{a}^*(f)\tilde{b}(f)}{S_n(f)} df \approx 4 \text{Re} \sum_i \frac{\tilde{a}^*(f_i)\tilde{b}(f_i)}{S_n(f_i)} \frac{1}{T}, \quad (12)$$

where $f_i = i/T$, with i from 1 up to the total number of frequencies $N_f = N/2$ with N the number of time points. Note that we will also consider frequency arrays with spacing different from the one previously mentioned. The tilde indicates the Fourier transform, and $S_n(f)$ is the one-sided noise spectral density for the LISA data streams, which we take from Babak et al. (2021). We use the mismatch $\mathcal{M}(h_{\text{FD}}, h_{\text{TD}})$ to compare the frequency domain implementation with the DFT of the time-domain implementation sampled with a time interval $\Delta t = T/N$. When taking the DFT of a time series, we expect some level of leakage. We mitigate this effect by applying the Hann window to the time domain waveform (Virtanen et al., 2020). To account for the effect of windowing in the FD waveform, we perform a convolution of the FD signal with the DFT of the window function. We do not maximize over the phase since the initial conditions of the frequency and phase evolution of the TD and FD waveforms are the same by construction. Global relative time and phase shifts are not observable parameters, so by not optimizing over them we are being conservative with the model and folding data processing into the model comparison.

We show the impact of windowing and spectral leakage in Figure 1 for a system with $\mu = 10M_\odot$, $M = 10^6M_\odot$, $e_0 = 0.6$, $p_0 = 9.56$ for an observation of $T = 1$ year, a mode content threshold $\epsilon = 10^{-2}$, and a sampling interval of $\Delta t = 10$ seconds. We show the amplitude squared, $|\tilde{h}_+(f)|^2$, of the frequency domain waveform (dashed orange line) and of the DFT of the time domain waveform (solid blue line) for the plus polarization in the upper panel of Figure 1. For reference, we also show the LISA sensitivity curve used to calculate the inner products. In the upper panel of Figure 1 we do not apply any windowing and therefore the FD waveform is zero only where there is no mode content. Conversely, the DFT of the TD waveform is affected by spectral leakage, especially for frequencies larger than 10^{-2} Hz. If we apply the Hann window consistently to both the FD and TD signals as previously described, we obtain the spectrum

² In the literature this is also denoted unfaithfulness.

shown in the lower panel of Figure 1. The effect of windowing reduces the total power but improves the agreement between the waveforms such that the mismatch decreases from 8×10^{-4} to 4×10^{-6} .

It is often quoted that the average mismatch for a waveform needs to be $1 - 0.9^{1/3} \approx 1 - 0.965$ to recover 90% of the signals in an experiment (Owen and Sathyaprakash, 1999; Apostolatos, 1995). However, the requirements on the mismatch for parameter estimation (PE) studies are not as straightforward to characterize, primarily as there is no one-to-one relationship between the level of mismatch and the level of bias that is incurred during inference. A sensible requirement for a PE study is to ask that the systematic (or mismodeling) error coming from an approximate waveform is lower than the statistical error coming from the noise in the data. While the former is independent of the signal-to-noise ratio ($\text{SNR} = \sqrt{\langle h_+ | h_+ \rangle + \langle h_\times | h_\times \rangle}$), the latter does depend on the SNR, and therefore the mismatch requirement on a waveform for a PE study is SNR-dependent. Therefore, we perform PE on the system that yields the largest value of $\mathcal{M} \times \text{SNR}^2$.

In particular, we check that the posterior distribution obtained with an FD waveform is equivalent to the one obtained with a TD waveform. To do this, we use the likelihood given by (Whittle, 1953):

$$p(s|h) \propto \exp \left[-\frac{1}{2} \sum_{\alpha=+, \times} \langle s - h | s - h \rangle_{\alpha} \right], \quad (13)$$

where we ignored the normalization factor that needs to be included when fitting for the shape of the noise spectral density S_n . We do not inject any noise in the datastream since we want to check whether any waveform systematic arises due to the difference between the FD and TD implementations. Note that if we apply a window to our data, we should take into account its effects in the likelihood. In fact, the expectation value of the inner product of the noise $n(t) = s(t) - h(t)$ is different if a window is applied


$$\mathbb{E}[\langle w(t) n(t) | w(t) n(t) \rangle] \neq \mathbb{E}[\langle n(t) | n(t) \rangle] = N \quad (14)$$

We do not take into account the effect of windowing in the likelihood because this would lead to a non-diagonal noise covariance matrix and an increase in the computational cost of the inner product. This is not a problem for the purpose of our study because we are consistently neglecting the effect of windowing in the likelihood when analyzing the data with different waveform models (FD or TD) (Burke, 2021). We remind the reader that we apply the windowing solely because the two waveforms (FD and TD) are in two different domains and we perform the parameter estimation only to check the presence of waveform systematics. However, we remark that the effect of windowing (and gaps) in LISA data is an important issue that remains to be addressed, but it is beyond the scope of this work.

3 RESULTS

The frequency and time domain waveforms are built using the same trajectory and amplitude modules. However, the final signal output is obtained using different mode summations and prescriptions. Therefore, any waveform difference might be attributed to the stationary phase approximation and/or spectral leakage. However, any difference in the speed of the waveform generation is due to the mode summation algorithm of the two domains. In the following, we investigate the waveform mismatch and speed by comparing the time and frequency domain implementations over the parameter space for different observation time-spans T , sampling intervals Δt and mode content thresholds ϵ . We consider as our fiducial values an observation

Figure 1. Spectrum of the frequency domain (FD) waveform compared to the spectrum of the discrete Fourier transform (DFT) of the time domain (TD) waveform. In the upper plot, we show the amplitude squared $|\tilde{h}_+(f)|^2$ when we do not include any windowing when computing the spectrum of the waveforms. In the lower panel, we apply the Hann window to both the TD and FD waveforms. Note that there are harmonics contributing to the low-frequency region of the spectrum. We consider an EMRI system composed of a secondary object with mass $\mu = 10M_\odot$ orbiting around a non-rotating black hole of mass $M = 10^6M_\odot$ with initial eccentricity $e_0 = 0.6$ and semi-latus rectum $p_0 = 9.56$ fixed such that the inspiral plunges after 1 year of observation.



time-span of $T = 4$ years (yrs), a sampling interval of $\Delta t = 5$ seconds, and a threshold $\epsilon = 10^{-5}$ and we vary each of these one at a time.

3.1 Accuracy analysis

We calculate the mismatch between the frequency and time domain waveforms as described in Sec. 2.3. The mismatches obtained over the parameter space are shown in the upper part of Figure 2. The median of the mismatch distributions are between 4×10^{-6} and 10^{-5} , whereas the largest 95% quantile of the mismatch is 9.7×10^{-5} .

The distributions obtained with a different mode content $\epsilon = 10^{-5}$ (solid blue histogram) and $\epsilon = 10^{-2}$ (dashed green histogram) are almost identical. This is expected since these systems are affected in the same way by spectral leakage. When comparing the mismatches obtained with different sampling intervals $\Delta t = 5$ s (solid blue histogram) and $\Delta t = 10$ s (dotted orange histogram) we obtain lower mismatches for larger sampling intervals. We attribute this difference to the fact that the total number of frequencies $T/\Delta t$ contributing to the mismatch is different. For the systems with lower $T/\Delta t$, the spectral leakage is weaker since the total power is lower. To confirm this we take the ratio of the mismatches $\mathcal{M}_{\Delta t=10\text{s}}/\mathcal{M}_{\Delta t=5\text{s}}$ and we find that it to be approximately constant as a function of the mass. For the same reason the mismatches obtained for $T = 2$ years (dashed-dotted red histogram) are slightly lower than the ones obtain for the fiducial value $T = 4$ years (solid blue histogram). However, the mismatches obtained for $(T, \Delta t) = (2 \text{ yrs}, 5 \text{ s})$ (dashed-dotted red histogram) are slightly larger than the ones obtained for $(T, \Delta t) = (4 \text{ yrs}, 10 \text{ s})$ (dotted orange histogram). This is probably due to the fact that we fix p_0 for a given inspiral duration, and the maximum p_0 allowed by the current implementation is $p_0 \approx 16$.

In the lower panel of Figure 2 we show the behavior of the mismatch as a function of the central black hole mass M . The correlation coefficients between the mass M and the mismatch \mathcal{M} is approximately ~ 0.6 . This is caused by the shape of the LISA sensitivity curve and the fact that for larger M systems the spectrum shifts to lower frequencies and the spectral leakage present at the high frequencies becomes more dominant. For large M the fiducial distribution of mismatches (blue crosses) is wider than the one obtained with $\Delta t = 10$ seconds, but both distributions reach approximately the same mismatches of $\sim 10^{-5}$ for $M \sim 10^7$. This is no longer the case when M becomes smaller. For $M \sim 10^5$, the mismatches of $\Delta t = 10$ seconds are generally smaller than the ones obtained from $\Delta t = 5$ seconds³. This confirms that for larger masses the spectral leakage dominates, whereas for smaller masses the mismatch is more affected by the total power in the signal. This behavior is also confirmed for systems with $T = 2$ years (red crosses). In fact, if we compare the systems with $(T, \Delta t) = (4 \text{ yrs}, 10 \text{ s})$ and $(T, \Delta t) = (2 \text{ yrs}, 5 \text{ s})$ that have the same power, the distribution for the latter system reaches larger values than the distribution for the former system at all masses because the maximum resolvable frequency, $1/(2\Delta t)$, is higher. We note that for lower masses, the distribution for $(T, \Delta t) = (2 \text{ yrs}, 5 \text{ s})$ has lower mismatches than the distribution for $(T, \Delta t) = (4 \text{ yrs}, 5 \text{ s})$ because the total power $T/\Delta t$ is smaller.

We use Bayesian parameter estimation to assess whether the FD waveforms are affected by systematic errors. We use the Eryn package (Karnesis et al., 2023) to perform a Markov chain Monte Carlo analysis to estimate the posterior distribution of the EMRI intrinsic parameters. Since this analysis is computationally expensive and cannot be performed over the 5000 realizations considered in the mismatch analysis, we decided to focus on the EMRI parameters that yield the largest $\mathcal{M} \times \text{SNR}^2$. For the values of $T = 4$ yrs, $\Delta t = 5$ s or $\Delta t = 10$ s, $\epsilon = 10^{-5}$ (blue solid and dotted orange histograms), we find that the largest value

³ The points that abruptly rise to larger mismatches for the orange distribution are caused by the fact that, for some low mass systems, the sampling interval $\Delta t = 10$ seconds is not small enough to resolve the largest frequency.

Figure 2. *Upper panel:* Mismatch between the frequency domain (FD) and time domain (TD) waveforms over the parameter space of EMRI systems and for different observation time-spans, T , sampling intervals, Δt and mode content thresholds, ϵ . The different parameter realizations are drawn uniformly from the following ranges: $\log(M/M_\odot) \in [\log 10^5, \log 10^7)$, $\log(\mu/M) \in [\log 10^{-6}, \log 10^{-4})$, $e_0 \in [0.001, 0.7)$, $\Phi_{\varphi 0}, \Phi_{r0} \in [0, 2\pi)$. *Lower panel:* Dependence of the mismatch on the central black hole mass M . The points that abruptly rise to larger mismatches for the orange dots are caused by the fact that, for some low-mass systems, the sampling interval $\Delta t = 10$ seconds is not small enough to resolve the largest frequency.

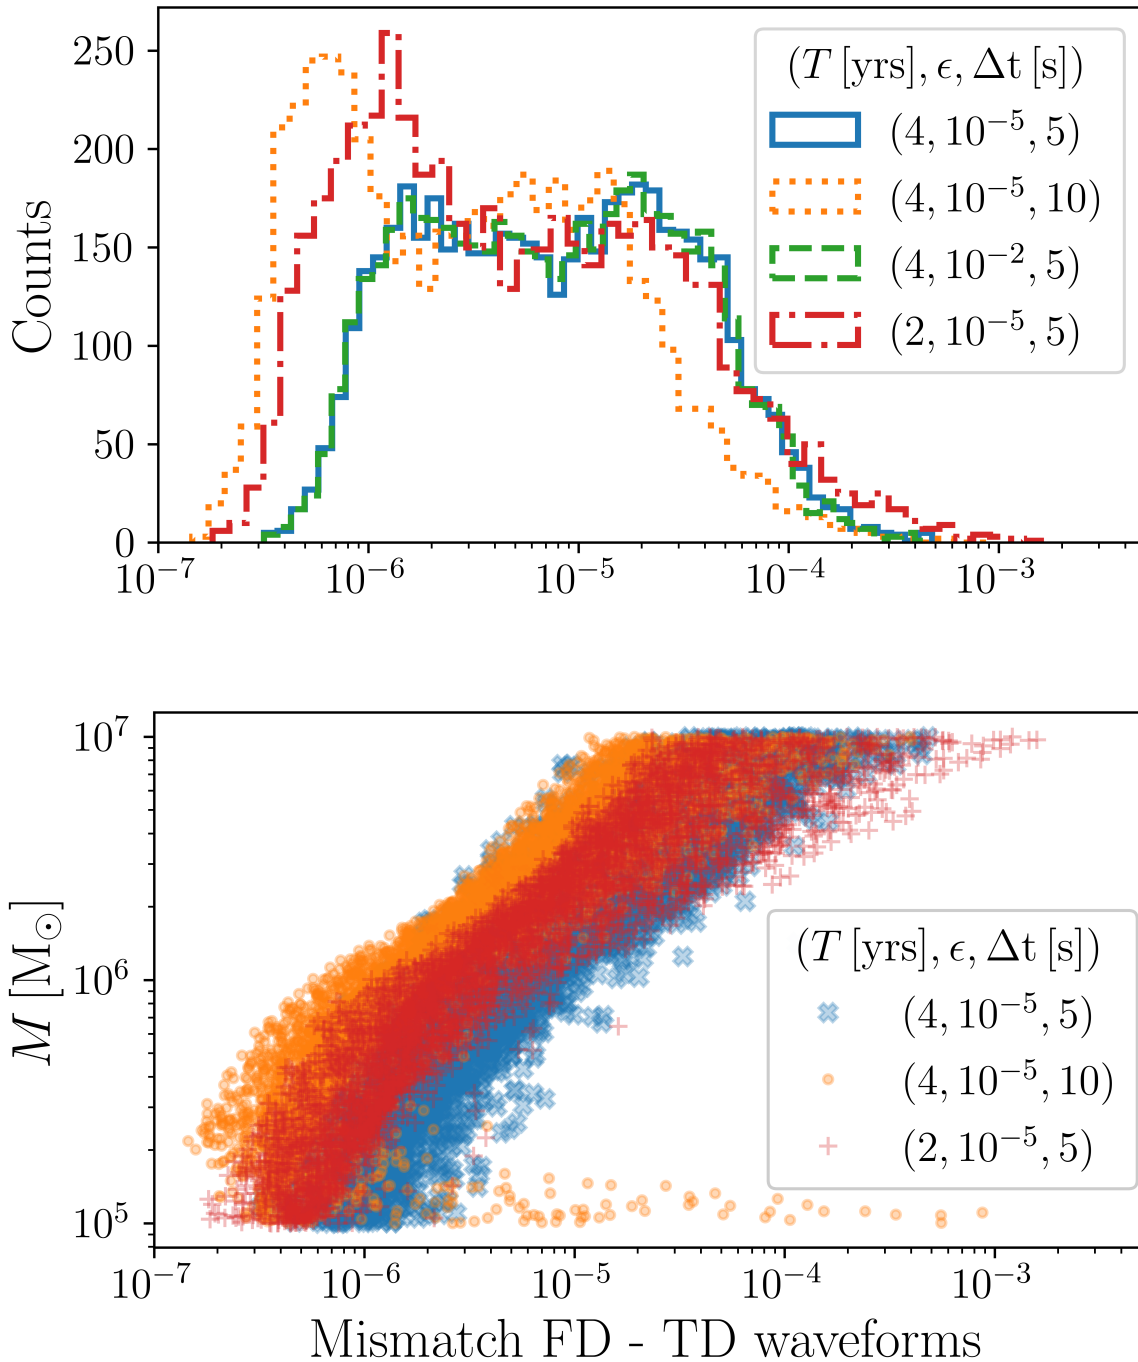


Table 1. EMRI parameters of the source with largest $\mathcal{M} \times \text{SNR}^2$, where \mathcal{M} is the mismatch between the frequency and time domain waveforms. The posteriors distributions obtained with different waveform implementations are shown in Fig. 3 for a mode content of $\epsilon = 10^{-5}$.

$M [M_\odot]$	$\mu [M_\odot]$	$p_0 [M]$	e_0	$T[\text{yrs}]$	SNR
3.67004174×10^6	292.058317	13.7091019	0.579413083	4	78

of $\mathcal{M} \times \text{SNR}^2$ is reached by the system with parameters reported in Table 1. The posterior distribution of this system is shown in Figure 3 for the three scenarios in which the FD and TD are used as injections or templates. The Kullback-Leibler divergences (Kullback and Leibler, 1951) between each pair of posteriors are of order 10^{-4} for all the considered parameters. Therefore, we conclude that we do not find any significant loss of accuracy or precision for the worst point in parameter space.

3.2 Computational Cost of the Waveform generation

The computational cost of the time and frequency domain implementations differ only in the final waveform summation and output domain. Here, we compare the waveform speeds without including the DFT cost. We make this choice to highlight the difference in the waveform generation cost and we provide the speed of the DFT for a reference system. The speed of the waveform is evaluated using the NVIDIA A100 GPU and the computer cluster Hypatia⁴, for GPU and CPU timings reported below. We warn the reader that the computational cost of the waveform can vary depending on the computing resources used.


By drawing 5000 realizations from the parameter space as discussed in the previous section, we show the speed-up factor given by the ratio of the TD and FD waveform evaluation speeds per each realization. The results for the GPU speed-up factor are shown in Figure 4. The upper panel shows the speed-up as a function of sampling interval, observation time and mode content, while the lower panel shows the dependence on eccentricity and central mass for the fiducial system. Note that the median speeds of all considered configurations are 0.044 and 0.055 seconds for the FD and TD waveforms, respectively. The distributions shown in the upper panel of Figure 4 depend on the priors chosen for the parameter space.

Since the mode threshold ϵ controls the mode content, we can deduce that the FD waveform generation is faster than the TD one when the harmonic content is larger, i.e., ϵ is smaller (see solid blue and dashed green histograms). This is a key result because future expansion of the `few` package will implement EMRI systems with a central rotating black hole where the number of harmonics is expected to increase by an order of magnitude.

As shown by the solid blue and red dashed-dotted histograms in the upper panel of Figure 4, the FD waveform creation is faster than the TD one for longer signals. This is expected since the TD generation increases with the number of time points. Instead, the FD generation slows down with an increasing number of frequencies. This is confirmed by comparing the orange and blue histograms, which differ only by the time sampling interval Δt . The FD generation for our fiducial case is two times faster for systems with $M > 2 \times 10^6$ and $e_0 > 0.2$. The performance of the FD waveform compared to the TD one improves with increasing eccentricity and central mass M . However, the speed-up trend shown in the lower panel of Figure 4 moves to higher eccentricities and higher masses when $T = 2$ years, resulting in a smaller speed-up region. In the limit of short signal durations we expect the TD generation to be faster than the FD one, and contrary in the long duration regime.

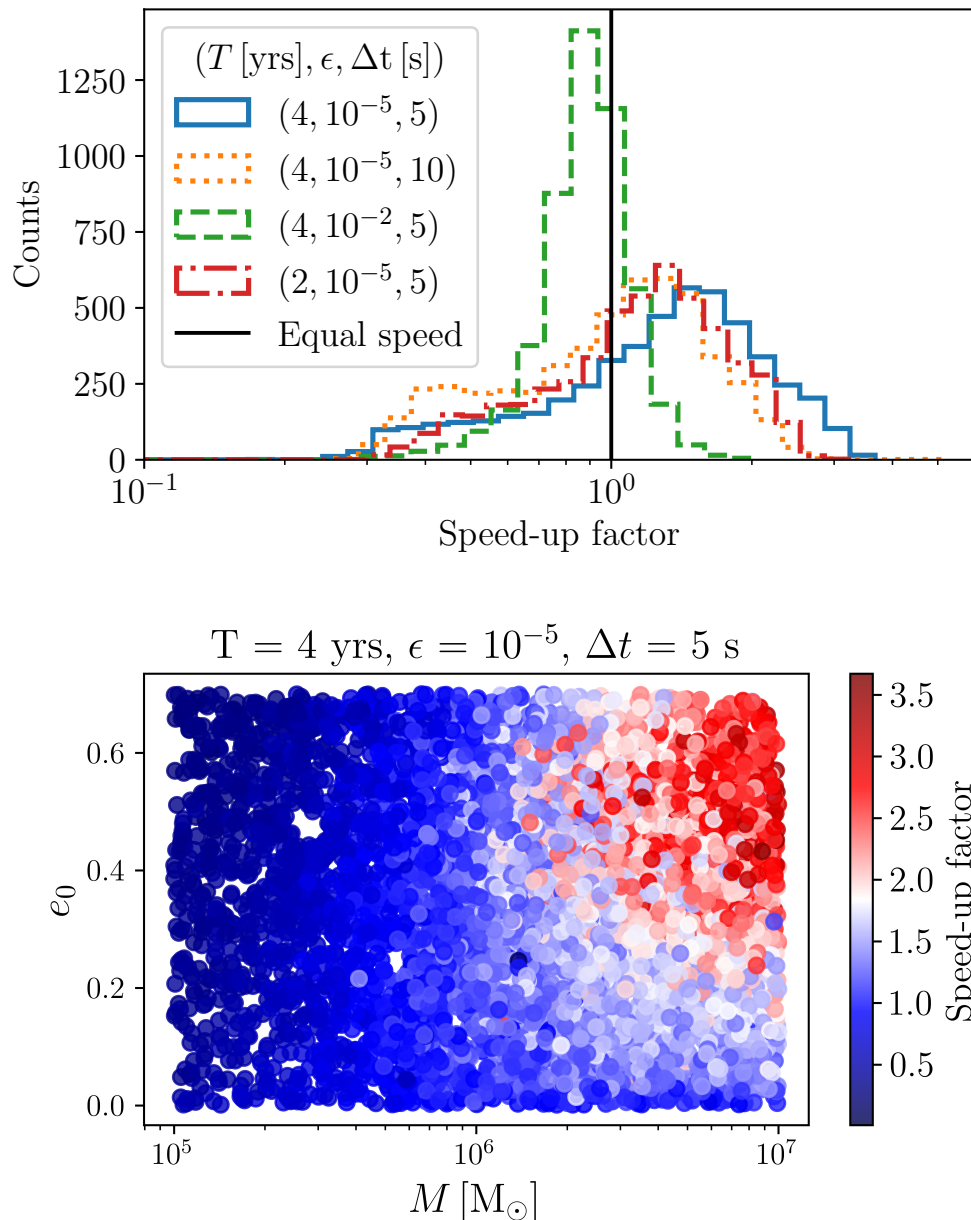
⁴ Hypatia consists of 262 GIGABYTE compute nodes (dual-socket, sixteen-core, SMT-enabled AMD EPYC (Naples) 7351 (2.40 GHz), of which 14 provide 8 GB RAM per core (4 GB per core otherwise).

Figure 3. Posterior distributions obtained with different waveform implementations for an EMRI system with parameters given in Table 1. In red we show the posterior distribution obtained using a frequency domain (FD) template to recover an injected waveform generated in the time domain (TD). The red distribution matches the posteriors obtained when the injected and template waveforms are in the same domain (blue and green posteriors). The observation time, the sampling interval and mode content are $(T, \Delta t, \epsilon) = (4 \text{ yrs}, 10 \text{ s}, 10^{-5})$. The Kullback-Leibler divergences between each pair of marginal posteriors are of order 10^{-4} for all the considered parameters.



The CPU waveform generation can take up to hundreds of seconds depending on $(T, \epsilon, \Delta t)$. Due to the significantly higher computational cost, we do not present the full scan of the parameter space, but from 10 parameter realizations we conclude that for $(T, \epsilon, \Delta t) = (4 \text{ yrs}, 10^{-2}, 10 \text{ s})$ the FD waveform generation is on average 5 times faster than the TD one on CPUs and takes 5 seconds in median. For reference, we also report in Table 2 the main findings after timing the FD and TD speeds for the system of Table 1 with $(T, \epsilon, \Delta t) = (4 \text{ yrs}, 10^{-5}, 10 \text{ s})$. The FD waveform is seven times faster than the TD waveform on CPUs,

Figure 4. *Upper panel:* Speed-up factor of the GPU waveform generation in the time domain (TD) versus frequency domain (FD) over the parameter space of EMRI systems and for different observation time-spans T , sampling interval Δt and mode content ϵ . The speed-up factor is given by the ratio of the TD generation speed divided by the FD generation speed per each parameter realization. The different parameter realizations are drawn uniformly from the following ranges: $\log(M/M_\odot) \in [\log 10^5, \log 10^7]$, $\log(\mu/M) \in [\log 10^{-6}, \log 10^{-4}]$, $e_0 \in [0.001, 0.7]$, $\Phi_{\varphi 0}, \Phi_{r 0} \in [0, 2\pi)$. *Lower panel:* Speed-up factor as a function of central black hole mass M and initial eccentricities e_0 for $T = 4$ years, $\Delta t = 5$ seconds and mode content $\epsilon = 10^{-5}$.



but only 1.3 times faster on GPUs. The DFT operation is particularly expensive on CPUs taking 7 seconds compared to the 0.4 milliseconds on a GPU.

Table 2. Computational cost of the waveform generation in frequency and time domain on GPU and CPU obtained from an EMRI system with parameters given in Table 1 and $(T, \epsilon, \Delta t) = (4 \text{ yrs}, 10^{-5}, 10 \text{ s})$. The speed of the DFT and of the downsampled frequency domain generation is also reported. The last row indicates the size of the frequency array.

Operation =	FD waveform (Downsampled)	TD waveform	DFT
Speed CPU [s] =	13.7 (0.4)	90.6	7.1
Speed GPU [s] =	0.048 (0.035)	0.064	0.0004
Number of frequencies N_f =	6311631 (7864)	6311631	6311631


3.3 Frequency downsampling

One of the main advantages of the FD generation is that we can evaluate the waveform on a sparser frequency array than the one imposed by the frequency resolution $1/T$. We use an evenly-spaced frequency array that spans only the region where the injected waveform is non-zero $f \in [0, f_{\max}]$. We change the frequency spacing to downsample the array and we define the total number of frequencies to be $N_f = f_{\max}/\Delta f$. We show in Table 2 the computational cost of the waveform generation in frequency and time domain on GPU and CPU obtained for an EMRI system with the parameters specified in Table 1 and $(T, \epsilon, \Delta t) = (4 \text{ yrs}, 10^{-5}, 10 \text{ s})$. As shown in Table 2 the frequency downsampling strongly affects the CPU speed-up factors, reducing the cost by a factor of 34. For GPUs, instead, the frequency downsampling mildly affects the speed-up factor between FD and TD generation since the GPU timing is already of the order of tens of milliseconds.

One of the downsides of reducing the number of frequencies is the loss of accuracy in the inner product calculation, with a possible impact on parameter inference. Here, we study the loss of accuracy as a function of the number of frequencies. As a proof of concept, we consider the source with the parameters reported in Table 1 and obtain the posterior distributions using the FD waveform with different input frequencies and with $(T, \epsilon, \Delta t) = (4 \text{ yrs}, 10^{-2}, 10 \text{ s})$. In this analysis, we did not include windowing since the injected and template waveforms are in the same domain. The absence of windowing also affects the total SNR, so we rescaled the distance to keep the SNR reported in Table 1 for the finest frequency resolution. The input array is given by evenly-spaced frequencies between zero and a maximum frequency f_{\max} . The spacing is set by choosing the total number of frequencies N_f .


In Figure 5, we show the posteriors obtained with different levels of downsampling and report the CPU and GPU likelihood speeds. The solid blue posterior is obtained using a frequency array defined by the DFT resolution $f_{\max} = 1/\Delta t$ and $N_f = 6311631$, whereas the maximum frequency of the dashed orange and dash-dotted green posteriors is set by the maximum harmonic frequency present in the signal $f_{\max} = 2.8 \text{ mHz}$. The only difference between the three cases is due to the approximation of the inner product in calculating the likelihood. Even using only 0.05% of the frequencies, the downsampled analysis with $N_f = 3554$ (dashed orange histograms) provides a posterior indistinguishable from the one of the full frequency array (solid blue). When further decreasing N_f to 35 the downsampling starts to affect the obtained posteriors. The frequency downsampling leads to a faster likelihood evaluation which drops from 4.26 seconds for the full array to 0.34 for CPUs. Further reducing the number of frequencies does not significantly improve the speed, highlighting that the bottleneck of the likelihood computation might be due to other waveform computation steps, such as mode selection (Katz et al., 2021). The GPU likelihood timing is reduced by a factor of two when using downsampling.

Figure 5. Posterior distribution for an EMRI system with parameters given in Table 1. The posteriors are obtained using the frequency domain waveform with a specified frequency array $f \in [0, f_{\max}]$ with N_f number of frequencies. The solid blue posterior considers a maximum frequency set by $f_{\max} = 1/(2\Delta t)$, whereas the maximum frequency of the dashed orange and dash-dotted green posteriors is set by the maximum frequency present in the signal. The computational cost of the likelihood evaluation is reported in seconds for GPU and CPU. The mode content is fixed to $\epsilon = 10^{-2}$, and the injected waveforms are shown in Figure 6.



The reason for the deviation in the posteriors can be explained by looking at the absolute square of the characteristic strain $|f\tilde{h}(f)|^2$ in Figure 6. The FD waveform spectrum obtained with $N_f = 3554$ (dashed orange line) well represents the spectrum obtained with the full array $N_f = 6311631$ (solid blue line) in the frequency range where the signal is above the LISA sensitivity. Instead, the spectrum obtained with $N_f = 35$ (dotted green line) has a very sparse frequency array that cannot describe all the features present in the EMRI signal. We conclude that the usage of frequency downsampling can be an important tool for

Figure 6. Characteristic strain spectrum $|f\tilde{h}(f)|^2$ obtained with the frequency domain waveform evaluated on different frequency arrays with N_f frequencies. The solid blue posterior considers a maximum frequency set by $f_{\max} = 1/(2\Delta t) = 50$ mHz, whereas the maximum frequency of the dashed orange and dash-dotted green posteriors is set by the maximum frequency present in the signal $f_{\max} = 2.8$ mHz. The parameters of the EMRI system are reported in Table 1 and the posteriors for these different levels of frequency downsampling are shown in Figure 5.



exploratory studies of EMRI parameter inference on CPUs. This motivates the FD implementation as an alternative to the TD one for CPU facilities. However, we stress that it is essential to check the validity of the downsampling before using it.

4 DISCUSSION AND CONCLUSIONS

In this work, we presented a ready-to-use implementation of an EMRI frequency domain waveform, and we compared its accuracy and performance against its time domain counterpart. We found the largest mismatch to be at 3×10^{-3} , and 95% of the generated waveforms over the EMRI parameter space have a mismatch below 10^{-4} . We further investigated the accuracy with a Bayesian analysis of the worst point in the parameter space and found no visible bias in the recovered parameters.

The frequency domain's computational cost on GPUs is lower than the time domain cost for central MBH masses of $> 10^6$ and for initial eccentricities > 0.2 and high numbers of harmonic modes. This suggests that the frequency domain model might be more suitable when including a large number of harmonics and we expect the FD implementation to be an important alternative to the TD one when extending the waveform model to Kerr spacetimes, where the number of modes is expected to increase by an order of magnitude. Thanks to the modularity of the FD implementation, the current model can be extended to Kerr spacetimes once amplitude and trajectories are available. These results also have implications for the data

analysis applications of the FD waveform model. If we expect the search phase of EMRIs to be conducted using the fastest model, for instance, with a lower mode content, then the TD generation may be preferable. However, including all the modes might be important when performing tests of general relativity, and therefore the FD waveform could be used in this scenario.

The speed-up factor of the waveform model on CPUs is ≈ 7 times faster than the time domain version for the considered system, and, in median, 5 times faster for different configurations. However, the waveform evaluation still takes order of tens of seconds, making the FD waveform not fast enough to perform EMRI parameter inference on CPUs. The frequency domain formulation allows us to choose the input frequency array giving the possibility of downsampling. This was explored for an EMRI system in which we found that the CPU likelihood evaluation can be reduced to 0.3 seconds when reducing the number of frequencies to be evaluated. This allows users without GPU resources to run EMRI parameter inference with a fully relativistic waveform for the first time. Frequency downsampling can be used for exploratory studies aimed at assessing the uncertainties that might be realized in parameter estimation. However, this must be done carefully especially when injecting noise. The noise must be scaled appropriately to avoid obtaining misleading results, and the possible downsampling will depend on the EMRI parameters. Downsampling will in general not be possible when analyzing real data, but it can be used as a tool for understanding EMRI parameter estimation.

In the current setup, the CPU evaluation time converges to values determined by the mode selection. In future works, we plan to improve this module, therefore enhancing the performance of the CPU FD implementation. Other approaches that use the sparsity of the frequency array to speed up have already been presented in (Cornish, 2021; Zackay et al., 2018). However, the non-monotonic modes and the large number of harmonics pose challenges in applying such techniques to EMRI waveforms. It is not clear yet whether relative binning or heterodyning will significantly speed up the parameter estimation of EMRI sources.

In this work, we did not include the effect of the LISA response function. This is crucial in using the FD waveforms for realistic data analysis studies. A frequency domain response was already presented in Marsat and Baker (2018). However, it still needs to be assessed whether the approximations used in Marsat and Baker (2018) apply to EMRIs. This should be investigated in future work.

Looking to the future, we note that post-adiabatic corrections to the waveform phase are important for precision science with EMRIs. These corrections are known for quasi-circular inspirals (Wardell et al., 2023), and the path to eccentric orbit calculations was laid out in (Leather and Warburton, 2023). Although the necessary offline second-order self-force calculations are computationally demanding, the online generation of the inspiral is structurally the same, and as fast, as the adiabatic model (Miller and Pound, 2021). With the modular FD implementation presented in this work the FD post-adiabatic waveform will be available as soon as the post-adiabatic corrections to the inspiral phases are known.

Finally, a full realistic data analysis pipeline has not been developed for EMRIs yet (see, however, Babak et al. (2009); Cornish (2011)). We stress the importance of developing waveform models in the time, frequency and time-frequency domains in order to explore the advantages and disadvantages of each domain (Cornish, 2020).

CONFLICT OF INTEREST STATEMENT

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

AUTHOR CONTRIBUTIONS

LS: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Project administration, Resources, Software, Supervision, Validation, Visualization, Writing – original draft.

MLK: Conceptualization, Data curation, Formal Analysis, Investigation, Methodology, Resources, Software, Supervision, Validation, Writing and Editing.

AJJC: Conceptualization, Methodology, Project administration, Supervision, Writing - review and editing.

SAH: Conceptualization, Formal analysis, Funding acquisition (at least of the portion relevant to developing the waveform data / fluxes), Writing - review and editing.

NW: Conceptualization, Data curation, Formal analysis, Software, Writing - review and editing.

JET: Writing - review and editing.

CEAC-B: Writing - review and editing.

JRG: Writing - review and editing.

FUNDING

AJJC and JET acknowledge support from the NASA LISA Preparatory Science grant 20-LPS20-0005. NW acknowledges support from a Royal Society - Science Foundation Ireland University Research Fellowship. SAH gratefully acknowledges support from US National Science Foundation Grant PHY-2110384, and past support from NASA Astrophysics Theory Program Grant 80NSSC18K109. This publication has emanated from research conducted with the financial support of Science Foundation Ireland under Grant numbers 16/RS-URF/3428, 17/RS-URF-RG/3490 and URF/R/221054. CEAC-B is supported by STFC studentship 2446638.

ACKNOWLEDGMENTS

L. S. acknowledges Ollie Burke for the very helpful discussions on the impact of windowing in LISA data analysis. This work makes use of the Black Hole Perturbation Toolkit (Black Hole Perturbation Toolkit, 2017).

DATA AVAILABILITY STATEMENT

This study is fully reproducible using the publicly available scripts at https://github.com/lorenzsp/EMRI_FrequencyDomainWaveforms.

REFERENCES

- Amaro-Seoane, P. (2018a). Relativistic dynamics and extreme mass ratio inspirals. *Living Rev. Rel.* 21, 4. doi:10.1007/s41114-018-0013-8
- Amaro-Seoane, P. (2018b). Relativistic dynamics and extreme mass ratio inspirals. *Living Rev. Rel.* 21, 4. doi:10.1007/s41114-018-0013-8
- Amaro-Seoane, P., Audley, H., Babak, S., Baker, J., Barausse, E., Bender, P., et al. (2017). Laser Interferometer Space Antenna. *ArXiv e-prints*

- Amaro-Seoane, P., Gair, J. R., Freitag, M., Miller, M. C., Mandel, I., Cutler, C. J., et al. (2007). TOPICAL REVIEW: Intermediate and extreme mass-ratio inspirals—astrophysics, science applications and detection using LISA. *Classical and Quantum Gravity* 24, R113–R169. doi:10.1088/0264-9381/24/17/R01
- Apostolatos, T. A. (1995). Search templates for gravitational waves from precessing, inspiraling binaries. *Phys. Rev. D* 52, 605–620. doi:10.1103/PhysRevD.52.605
- Arca-Sedda, M. and Capuzzo-Dolcetta, R. (2019). The MEGaN project II. Gravitational waves from intermediate-mass and binary black holes around a supermassive black hole. *MNRAS* 483, 152–171. doi:10.1093/mnras/sty3096
- Arnaud, K. A. et al. (2006). A How-To for the Mock LISA Data Challenges. *AIP Conf. Proc.* 873, 625–632. doi:10.1063/1.2405109
- Babak, S., Fang, H., Gair, J. R., Glampedakis, K., and Hughes, S. A. (2007). “Kludge” gravitational waveforms for a test-body orbiting a Kerr black hole. *Phys. Rev. D* 75, 4005
- Babak, S., Gair, J., Sesana, A., Barausse, E., Sopuerta, C. F., Berry, C. P. L., et al. (2017). Science with the space-based interferometer LISA. V. Extreme mass-ratio inspirals. *Phys. Rev. D* 95, 3012
- Babak, S., Gair, J. R., and Porter, E. K. (2009). An algorithm for the detection of extreme mass ratio inspirals in LISA data. *Classical and Quantum Gravity* 26, 135004. doi:10.1088/0264-9381/26/13/135004
- Babak, S., Petiteau, A., and Hewitson, M. (2021). LISA Sensitivity and SNR Calculations
- Babak, S. et al. (2008). Report on the second Mock LISA Data Challenge. *Class. Quant. Grav.* 25, 114037. doi:10.1088/0264-9381/25/11/114037
- Babak, S. et al. (2010). The Mock LISA Data Challenges: From Challenge 3 to Challenge 4. *Class. Quant. Grav.* 27, 084009. doi:10.1088/0264-9381/27/8/084009
- Barack, L. and Cutler, C. (2007). Using LISA extreme-mass-ratio inspiral sources to test off-Kerr deviations in the geometry of massive black holes. *PRD* 75, 42003. doi:10.1103/PhysRevD.75.042003
- Barack, L. and Pound, A. (2018). Self-force and radiation reaction in general relativity. *Reports on Progress in Physics* 82, 16904. doi:10.1088/1361-6633/aae552
- Barausse, E., Cardoso, V., and Pani, P. (2014). Can environmental effects spoil precision gravitational-wave astrophysics? *PRD* 89, 104059. doi:10.1103/PhysRevD.89.104059
- Barausse, E., Cardoso, V., and Pani, P. (2015). Environmental Effects for Gravitational-wave Astrophysics. In *Journal of Physics Conference Series*. vol. 610 of *Journal of Physics Conference Series*, 12044. doi:10.1088/1742-6596/610/1/012044
- Barausse, E. and Rezzolla, L. (2008). Influence of the hydrodynamic drag from an accretion torus on extreme mass-ratio inspirals. *PRD* 77, 104027. doi:10.1103/PhysRevD.77.104027
- Barausse, E., Rezzolla, L., Petroff, D., and Ansorg, M. (2007). Gravitational waves from extreme mass ratio inspirals in nonpure Kerr spacetimes. *PRD* 75, 64026. doi:10.1103/PhysRevD.75.064026
- [Dataset] Barsanti, S., Franchini, N., Gualtieri, L., Maselli, A., and Sotiriou, T. P. (2022). Extreme mass-ratio inspirals as probes of scalar fields: eccentric equatorial orbits around Kerr black holes
- [Dataset] Berry, C. P. L., Hughes, S. A., Sopuerta, C. F., Chua, A. J. K., Heffernan, A., Holley-Bockelmann, K., et al. (2019a). The unique potential of extreme mass-ratio inspirals for gravitational-wave astronomy
- Berry, C. P. L., Hughes, S. A., Sopuerta, C. F., Chua, A. J. K., Heffernan, A., Holley-Bockelmann, K., et al. (2019b). The unique potential of extreme mass-ratio inspirals for gravitational-wave astronomy
- [Dataset] Black Hole Perturbation Toolkit (2017). Black Hole Perturbation Toolkit. (bhptoolkit.org)
- Burke, O. (2021). *Extreme precision and extreme complexity: source modelling and data analysis development for the laser interferometer space antenna*. Ph.D. thesis, The University of Edinburgh

- Chua, A. J. K. and Cutler, C. J. (2022). Nonlocal parameter degeneracy in the intrinsic space of gravitational-wave signals from extreme-mass-ratio inspirals. *Phys. Rev. D* 106, 124046. doi:10.1103/PhysRevD.106.124046
- Chua, A. J. K., Katz, M. L., Warburton, N., and Hughes, S. A. (2021). Rapid generation of fully relativistic extreme-mass-ratio-inspiral waveform templates for lisa data analysis. *Phys. Rev. Lett.* 126, 051102. doi:10.1103/PhysRevLett.126.051102
- Cornish, N. J. (2011). Detection strategies for extreme mass ratio inspirals. *Classical and Quantum Gravity* 28, 094016. doi:10.1088/0264-9381/28/9/094016
- Cornish, N. J. (2020). Time-frequency analysis of gravitational wave data. *PRD* 102, 124038. doi:10.1103/PhysRevD.102.124038
- Cornish, N. J. (2021). Heterodyned likelihood for rapid gravitational wave parameter inference. *Phys. Rev. D* 104, 104054. doi:10.1103/PhysRevD.104.104054
- Droz, S., Knapp, D. J., Poisson, E., and Owen, B. J. (1999). Gravitational waves from inspiraling compact binaries: Validity of the stationary phase approximation to the Fourier transform. *Phys. Rev. D* 59, 124016. doi:10.1103/PhysRevD.59.124016
- Fujita, R. and Hikida, W. (2009). Analytical solutions of bound timelike geodesic orbits in Kerr spacetime. *Classical and Quantum Gravity* 26, 135002. doi:10.1088/0264-9381/26/13/135002
- Gair, J. R., Barack, L., Creighton, T., Cutler, C., Larson, S. L., Phinney, E. S., et al. (2004). Event rate estimates for LISA extreme mass ratio capture sources. *Classical and Quantum Gravity* 21, S1595–S1606. doi:10.1088/0264-9381/21/20/003
- Gair, J. R., Flanagan, E. E., Drasco, S., Hinderer, T., and Babak, S. (2011). Forced motion near black holes. *PRD* 83, 44037. doi:10.1103/PhysRevD.83.044037
- Gair, J. R., Hewitson, M., Petiteau, A., and Mueller, G. (2022). *Space-based Gravitational Wave Observatories*. doi:10.1007/978-981-15-4702-7_3-1
- Gair, J. R., Tang, C., and Volonteri, M. (2010). LISA extreme-mass-ratio inspiral events as probes of the black hole mass function. *PRD* 81, 104014. doi:10.1103/PhysRevD.81.104014
- Gair, J. R., Vallisneri, M., Larson, S. L., and Baker, J. G. (2013). Testing General Relativity with Low-Frequency, Space-Based Gravitational-Wave Detectors. *Living Reviews in Relativity* 16, 7. doi:10.12942/lrr-2013-7
- Hinderer, T. and Flanagan, E. E. (2008). Two-timescale analysis of extreme mass ratio inspirals in Kerr spacetime: Orbital motion. *PRD* 78, 64028. doi:10.1103/PhysRevD.78.064028
- Hughes, S. A. (2000). Evolution of circular, nonequatorial orbits of kerr black holes due to gravitational-wave emission. *Phys. Rev. D* 61, 084004. doi:10.1103/PhysRevD.61.084004
- Hughes, S. A. (2001). Evolution of circular, nonequatorial orbits of Kerr black holes due to gravitational wave emission. II. Inspiral trajectories and gravitational wave forms. *Phys. Rev. D* 64, 064004. doi:10.1103/PhysRevD.64.064004. [Erratum: *Phys.Rev.D* 88, 109902 (2013)]
- Hughes, S. A., Warburton, N., Khanna, G., Chua, A. J. K., and Katz, M. L. (2021). Adiabatic waveforms for extreme mass-ratio inspirals via multivoice decomposition in time and frequency. *Phys. Rev. D* 103, 104014. doi:10.1103/PhysRevD.103.104014. [Erratum: *Phys.Rev.D* 107, 089901 (2023)]
- Hughes, S. A., Warburton, N., Khanna, G., Chua, A. J. K., and Katz, M. L. (2021). Adiabatic waveforms for extreme mass-ratio inspirals via multivoice decomposition in time and frequency. *arXiv e-prints*, arXiv:2102.02713
- Karnesis, N., Katz, M. L., Korsakova, N., Gair, J. R., and Stergioulas, N. (2023). Eryn: A multi-purpose sampler for Bayesian inference

- Katz, M. L., Chua, A. J. K., Speri, L., Warburton, N., and Hughes, S. A. (2021). Fast extreme-mass-ratio-inspiral waveforms: New tools for millihertz gravitational-wave data analysis. *Phys. Rev. D* 104, 064047. doi:10.1103/PhysRevD.104.064047
- Klein, A., Barausse, E., Sesana, A., Petiteau, A., Berti, E., Babak, S., et al. (2016). Science with the space-based interferometer elisa: Supermassive black hole binaries. *Phys. Rev. D* 93, 024003. doi:10.1103/PhysRevD.93.024003
- Klein, A., Cornish, N., and Yunes, N. (2014). Fast Frequency-domain Waveforms for Spin-Precessing Binary Inspirals. *Phys. Rev. D* 90, 124029. doi:10.1103/PhysRevD.90.124029
- Kullback, S. and Leibler, R. A. (1951). On Information and Sufficiency. *The Annals of Mathematical Statistics* 22, 79–86. doi:10.1214/aoms/1177729694
- Laghi, D., Tamanini, N., Del Pozzo, W., Sesana, A., Gair, J., and Babak, S. (2021). Gravitational wave cosmology with extreme mass-ratio inspirals. *arXiv e-prints*, arXiv:2102.01708
- Leather, B. and Warburton, N. (2023). Applying the effective-source approach to frequency-domain self-force calculations for eccentric orbits
- MacLeod, C. L. and Hogan, C. J. (2008). Precision of Hubble constant derived using black hole binary absolute distances and statistical redshift information. *Phys. Rev. D* 77, 043512. doi:10.1103/PhysRevD.77.043512
- Marsat, S. and Baker, J. G. (2018). Fourier-domain modulations and delays of gravitational-wave signals
- Maselli, A., Franchini, N., Gualtieri, L., and Sotiriou, T. P. (2020). Detecting scalar fields with Extreme Mass Ratio Inspirals. *Phys. Rev. Lett.* 125, 141101. doi:10.1103/PhysRevLett.125.141101
- Maselli, A., Franchini, N., Gualtieri, L., Sotiriou, T. P., Barsanti, S., and Pani, P. (2022). Detecting fundamental fields with LISA observations of gravitational waves from extreme mass-ratio inspirals. *Nature Astron.* 6, 464–470. doi:10.1038/s41550-021-01589-5
- Miller, J. and Pound, A. (2021). Two-timescale evolution of extreme-mass-ratio inspirals: waveform generation scheme for quasicircular orbits in Schwarzschild spacetime. *Phys. Rev. D* 103, 064048. doi:10.1103/PhysRevD.103.064048
- Moore, C. J., Gerosa, D., and Klein, A. (2019). Are stellar-mass black-hole binaries too quiet for LISA? *Mon. Not. Roy. Astron. Soc.* 488, L94–L98. doi:10.1093/mnras/slz104
- Owen, B. J. and Sathyaprakash, B. S. (1999). Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement. *Phys. Rev. D* 60, 022002. doi:10.1103/PhysRevD.60.022002
- Pan, Z. and Yang, H. (2021). Formation Rate of Extreme Mass Ratio Inspirals in Active Galactic Nucleus. *arXiv e-prints*, arXiv:2101.09146
- Piovano, G. A., Maselli, A., and Pani, P. (2020). Extreme mass ratio inspirals with spinning secondary: a detailed study of equatorial circular motion. *Phys. Rev. D* 102, 024041. doi:10.1103/PhysRevD.102.024041
- Porter, E. K. (2009). An Overview of LISA Data Analysis Algorithms. *GW Notes* 1, 4
- Savalle, E., Gair, J., Speri, L., and Babak, S. (2022). Assessing the impact of instrumental calibration uncertainty on LISA science. *Phys. Rev. D* 106, 022003. doi:10.1103/PhysRevD.106.022003
- Schmidt, W. (2002). Celestial mechanics in Kerr space-time. *Class. Quant. Grav.* 19, 2743. doi:10.1088/0264-9381/19/10/314
- Seto, N. (2001). Proposal for determining the total masses of eccentric binaries using signature of periastron advance in gravitational waves. *Phys. Rev. Lett.* 87, 251101. doi:10.1103/PhysRevLett.87.251101. [Erratum: *Phys. Rev. Lett.* 101, 209901 (2008)]

- Speri, L., Antonelli, A., Sberna, L., Babak, S., Barausse, E., Gair, J. R., et al. (2023). Probing Accretion Physics with Gravitational Waves. *Phys. Rev. X* 13, 021035. doi:10.1103/PhysRevX.13.021035
- Speri, L., Karnesis, N., Renzini, A. I., and Gair, J. R. (2022). A roadmap of gravitational wave data analysis. *Nature Astron.* 6, 1356–1363. doi:10.1038/s41550-022-01849-y
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., et al. (2020). SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17, 261–272. doi:10.1038/s41592-019-0686-2
- Wardell, B., Pound, A., Warburton, N., Miller, J., Durkan, L., and Le Tiec, A. (2023). Gravitational Waveforms for Compact Binaries from Second-Order Self-Force Theory. *Phys. Rev. Lett.* 130, 241402. doi:10.1103/PhysRevLett.130.241402
- Whittle, P. (1953). The Analysis of Multiple Stationary Time Series. *Journal of the Royal Statistical Society. Series B (Methodological)* 15, 125–139
- Yunes, N., Kocsis, B., Loeb, A., and Haiman, Z. (2011). Imprint of Accretion Disk-Induced Migration on Gravitational Waves from Extreme Mass Ratio Inspirals. *PRL* 107, 171103. doi:10.1103/PhysRevLett.107.171103
- Zackay, B., Dai, L., and Venumadhav, T. (2018). Relative Binning and Fast Likelihood Evaluation for Gravitational Wave Parameter Estimation