

Computation of stochastic background from extreme-mass-ratio inspiral populations for LISA

Federico Pozzoli 

Dipartimento di Scienza e Alta Tecnologia, Università degli Studi dell'Insubria, I-22100 Como, Italy

Stanislav Babak

APC, AstroParticule et Cosmologie, CNRS, Université de Paris, F-75013 Paris, France

Alberto Sesana and Matteo Bonetti 

Dipartimento di Fisica "G. Occhialini", Università degli Studi di Milano-Bicocca, 20126 Milano, Italy

Nikolaos Karnesis 

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



(Received 14 February 2023; accepted 17 October 2023; published 27 November 2023)

Extreme mass ratio inspirals (EMRIs) are among the primary targets for the Laser Interferometer Space Antenna (LISA). The extreme mass ratios of these systems result in relatively weak gravitational wave signals, that can be individually resolved only for cosmologically nearby sources (up to $z \approx 2$). The incoherent piling up of the signal emitted by unresolved EMRIs generate a confusion noise, that can be formally treated as a stochastic GW background (GWB). In this paper, we estimate the level of this background considering a collection of astrophysically motivated EMRI models, spanning the range of uncertainties affecting EMRI formation. To this end, we employed the innovative *augmented analytic kludge* waveforms and used the full LISA response function. For each model, we compute the GWB SNR and the number of resolvable sources. Compared to simplified computations of the EMRI signals from the literature, we find that for a given model the GWB SNR is lower by a factor of ≈ 2 whereas the number of resolvable sources drops by a factor of 3 to 5. Nonetheless, the vast majority of the models result in potentially detectable GWB which can also significantly contribute to the overall LISA noise budget in the 1–10 mHz frequency range.

DOI: [10.1103/PhysRevD.108.103039](https://doi.org/10.1103/PhysRevD.108.103039)

I. INTRODUCTION

Galactic nuclei are among the densest structures in the Universe, and it has been shown that they host a massive black hole (MBH) in their center. At such high densities, generally exceeding $10^6 M_\odot \text{pc}^{-3}$, the high rate of strong and weak gravitational encounters among stars and compact objects (COs) efficiently redistribute their energy and angular momentum, occasionally resulting in close encounters with the central MBH. From this stems a wide variety of spectacular phenomena driven by extreme dynamics such as stellar tidal disruptions [1], hypervelocity stars ejection [2], relativistic captures of COs [3], and quasi-periodic eruptions [4].

Specifically, when scattered on very low angular momentum orbits, COs decouple from the influence of the stellar environment and, together with the central MBH, evolve as a relativistic binary approximately in isolation. The energy of the binary is gradually released through the emission of gravitational waves (GWs), causing the inspirals of the CO onto the MBH.

Considering the enormous difference between the two masses of the binary (typically $1\text{--}50 M_\odot$ for the CO and $10^5\text{--}10^9 M_\odot$ for the MBH), such events are called extreme mass ratio inspirals (EMRIs), with mass ratio in the range $q = \mu/M = 10^{-9}\text{--}10^{-4}$, with M the mass of the MBH and μ the mass of the CO. Several channels of EMRI formation have been proposed, which modify the process described above either by adding further physical effects such as resonant relaxation and BH-BH scattering events [3,5], or by invoking different formation processes, like binary tidal separation [6], captures of cores of giants [7], massive star capture or production in accretion discs [8].

The GW emitted by EMRIs is a primary target for the forthcoming space mission LISA (Laser Interferometer Space Antenna) [9]. The observatory consists of three spacecrafts forming an equilateral triangle with $\sim 2.5 \times 10^6$ km armlength. Lasers will be sent both ways between each pair of spacecrafts, and they will measure the phase changes between the transmitted and the received beam, induced by

the deformation of the spacetime due to the passing GW. The LISA sensitivity will span the 0.1–100 mHz frequency range, enabling the detection of GWs produced by different kinds of events such as massive black hole binary (MBHB) coalescences [10], galactic compact binaries (GCB) [11], inspiraling stellar origin black hole (SOBH) binaries [12], and EMRIs.

Due to their tiny mass ratio, EMRIs evolve slowly, completing $\sim 10^4$ – 10^6 cycles in LISA’s sensitive frequency range before eventually plunging onto the central MBH [13,14]. A large number of cycles allow measuring the parameters of the binary with exquisitely high precision. Therefore, EMRIs are ideal sources to map MBH spacetime [15,16], perform tests of general relativity [17], and possibly detect the presence of gas around the central MBH [18,19]. Measuring the properties of a population of EMRI signals could additionally provide information on the mass distribution of MBHs [20] and their host stellar environment [5]. Concerning waveform modeling, the large number of orbital cycles requires an extremely faithful waveform, since even a tiny dephasing over one cycle can jeopardize the correct signal recovery through matched filtering [21]. Additionally, the vast majority of EMRIs are not expected to be individually detectable with LISA because they are too far away or too far from the final plunge onto the MBH. Therefore, thousands of EMRIs will be present in the LISA data without reaching the “individual detection threshold”, and their GW emission can pile up incoherently, forming a confusion noise [22,23]. In the worst case scenario, this stochastic gravitational wave background (GWB) could even overcome the level of the anticipated instrumental noise, possibly jeopardizing the detection of other sources. This is, for example, the case with the collective signal from unresolvable GCBs, which constitute the primary astrophysical limitation for the LISA mission in the frequency range [0.2, 3] mHz. The GCB confusion noise is expected to be highly anisotropic, unlike the EMRIs one, and its amplitude in the LISA detector fluctuates due to the satellite constellation motion. This modulation makes this signal at least partially subtractable from the LISA error budget when searching for an underlying stochastic background.

Despite the characterization of GWB from EMRI being essential for the fulfillment of the LISA goals, it has so far raised little attention in the vast literature involving LISA. The first estimation has been done by the pioneering work of [24], which used basic piecewise approximations for the inclination and eccentricity averaged GW signal from unresolved EMRIs and considered early estimates of the EMRI rates, in terms of a redshift independent MBH mass function. Subsequently, [23] improved the estimate by considering the EMRI populations of [25], but still employing a straightforward waveform for the GW signal, namely, a simplified version of the analytic kludge (AK) used by [24].

In this study, we want to contribute by adding further elements to make the estimate more credible. To provide a realistic estimate, we couple the astrophysical motivated model developed in [25] to the augmented analytic kludge (AAK) [26] waveform, calibrated against the numerical kludge (NK) model [27]. Most importantly, we perform the first realistic injection of the collective distribution of EMRI signals in LISA by directly computing the GW response of the instrument in the time-delay interferometry channels. We use the recursive algorithm developed in [28] to subtract resolvable sources and estimate the residual GWB level. Our approach not only ensures a better GWB characterization, but it also contributes to the development of tools for realistic signals injection into the LISA data stream, which is critical for future investigations and for the development of appropriate analysis pipelines in the context of the LISA Consortium.

The plan of the paper is as follows. In Sec. II, we provide all the technical aspects concerning the LISA sensitivity curve and the TDI technique. In Sec. III, we briefly review kludge waveforms used to model the GW signal from EMRIs, highlighting their differences and justifying the use of AAK. Section IV describes the astrophysical models chosen to construct the EMRI populations and defines the corresponding catalogs. Finally, Secs. V and VI are dedicated respectively to presenting the results and to drawing our final conclusions. Throughout the paper, we adopt geometrical units $G = c = 1$.

II. LISA

A. Sensitivity

The definition of the sensitivity is closely related to the SNR which for a deterministic source can be defined as

$$\text{SNR}^2 = 4\text{Re} \left(\int_0^\infty df \frac{\tilde{h}(f)\tilde{h}^*(f)}{S_n(f)} \right), \quad (1)$$

where we have introduced the one-sided noise power spectral density (PSD) $S_n(f)$ and the signal $\tilde{h}(f)$ in the frequency domain. The detector noise is described as a stochastic variable $n(t)$, and the noise PSD is defined as

$$E[\tilde{n}(f)\tilde{n}(f')] = \frac{1}{2}\delta(f - f')S_n(f) \quad (2)$$

where $E[\dots]$ is the expectation value, $\tilde{n}(f)$ is the Fourier transform of $n(t)$, and δ is the Dirac delta distribution.

In the analysis, we adopt the noise model of ESA’s science requirement document [29]. That noise model has been approximated into two effective functions. The high-frequency components are represented by the readout noise S_{OMS} , whereas the low-frequency components by the single mass-acceleration noise S_{acc} . Each LISA spacecraft hosts a test mass designed to follow geodesic motion only.

However, LISA Pathfinder showed that they are not perfectly free falling. External spurious forces push them out of their geodesics, resulting in noninertial residual accelerations, which are collected in S_{acc} . Instead, S_{OMS} collects all the optical metrology system noise entering via power measurements at the photodetectors.

We highlight that [23], besides the sky-averaged sensitivity for LISA, took also into account the confusion noise generated by unresolved GCB (mostly white dwarf, WD, binaries). That component has been added according to the fitted formula obtained in [28]. Generally, the GCB population produces a stochastic GWB that effectively degrades the instrumental sensitivity at frequencies below 1 mHz. Here, we ignore the CGB GWB to produce robust estimates that depend on the LISA noise only. Although including the CGB GWB might affect the EMRI GWB observability, the main contribution to the SNR of the latter comes from $f > 2$ mHz, where the former should have a minor impact.

B. Time-delay interferometry

The main source of noise for LISA will be laser frequency noise. It is due to the frequency instability of the onboard lasers. Ideally, a laser operates at a single frequency with zero linewidth. In the real world, however, a laser has a finite linewidth because of phase fluctuations, which cause instantaneous frequency shifts away from the central frequency. Equal-arm interferometer detectors can observe GW by canceling the laser frequency fluctuations affecting the light injected into their arms. This is done by comparing phases of split beams propagated along the equal arms of the detector. The laser frequency fluctuations affecting the two beams undergo the same delay and cancel out at the photodetector. However, for a not-equal arms interferometer (as LISA) the exact cancellation of the laser frequency fluctuations does not take place at the photodetector. For this reason, a postprocess technique has been developed. The algorithm is called time-delay interferometry (TDI), and it was first proposed by [30]. To understand how TDI works, we have to introduce some notation. Figure 1 represents the constellation of LISA schematically. The spacecrafts are labeled 1, 2, 3, and they are oriented clockwise. The separating distances are denoted as L_1 , L_2 , L_3 , with L_i being the opposite spacecraft, and we refer with L'_i to the reverse path made by a photon. Primed or unprimed indices refer to beams traveling clockwise or counterclockwise, respectively. In the literature [31], we can find the derivation of the link response $y_{sr}(t, \hat{k})$ from a generic GW traveling along \hat{k} , employing common approximations such as the first-order expansion of wave propagation time and assuming the spacecraft remains stationary during this propagation period. y_{sr} is indexed with two numbers sr , where s is the index of the sender spacecraft, and j is the index of the receiver spacecraft.

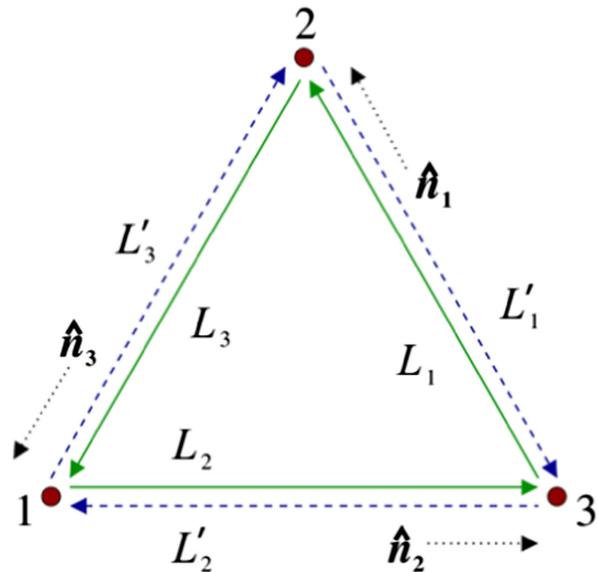


FIG. 1. Schematic LISA configuration. The image has been taken from [32].

To study the response of each link (single arm) to the GW signal, it is useful to introduce the delay operator \mathcal{D}_l ,

$$\mathcal{D}_l y_{sr}(t) = y_{sr}(t - L_l), \quad (3)$$

that we later express through its shorthand notation $y_{sr,l}$.

The main idea of TDI is to add the delayed signals of LISA together so that the laser frequency noise terms add up to zero. This amounts to seeking data combinations that cancel the laser frequency noise. There exist different versions of TDI, corresponding to different levels of accuracy of the LISA geometry taken into account. In this work, we adopt the first generation of TDI, that is valid for a static LISA constellation and does not account for rotation and flexing. The first condition emphasizes that all armlengths are constant in time, while the second condition implies that the light travel time in one arm is independent of the light propagation direction, meaning that $L_i(t) = L_i$ and $L_i = L'_i$. A direct consequence is that the delayed operators commute.

A possible realization of TDI defines three variables X , Y , and Z , corresponding to pairwise Michelson-like interferometers. As an example, in Eq. (4) we report the expression of the X channel in the first generation of TDI, while Y and Z are obtained by cyclic permutation of indices $1 \rightarrow 2 \rightarrow 3 \rightarrow 1$. $X(t)$ can be visualized as the difference of two sums of measurements, each corresponding to a specific light path from the laser on board spacecraft 1,

$$X(t) = [(y_{31} + y_{13,2}) + (y_{21,22} + y_{12,322})] - [(y_{21} + y_{12,3}) + (y_{31,33} + y_{13,233})]. \quad (4)$$

redshift z . Thus, to summarize, in the kludges family, 14 parameters are needed to describe the EMRI waveform: $\{M, \mu, a, p, e, \phi, \alpha, \tilde{\gamma}, \iota, D_L, \theta_S, \phi_S, \theta_K, \phi_K\}$. Actually, they would be 17, also considering the spin of the secondary object, but its effect on the dynamics and GW emission is small, and it is irrelevant to our analysis.

A. Analytic kludge

The AK has been developed by Barack and Cutler [24]. It approximates the EMRI system as being, at any instant, a Keplerian binary emitting a quadrupolar wave [37]. In this waveform, the orbital evolution is given by five first-order ordinary differential equations of $(\phi, \nu, \tilde{\gamma}, \alpha, e)$. The equations combine several leading order contributions from PN (post-Newtonian) dynamics. For ν and e , the expressions are accurate through 3.5PN $\tilde{\gamma}$, and α equations are accurate through 2PN order. The PN expressions are expressed by Eqs. (28)–(31) in [24]. The EMRI orbital elements are integrated between $t = 0$ and the observed time T_{obs} . If the EMRI plunges during the observation, the inspiral evolution has to be truncated, and this can be done in two different ways. The first consists of the ‘‘Schwarzschild’’ analytic kludge which uses the last stable orbit for nonspinning MBH as a proxy for the plunge. Alternatively, the ‘‘Kerr’’ analytic kludge takes into account the MBH spin information, identifying the last stable orbit defined by a Kerr separatrix. The choice of different plunging conditions defines how close to the MBH horizon we can approach and, therefore, the bandwidth of the signal and its SNR. The Schwarzschild version of AK underestimates signal’s strength, while the Kerr version strongly overestimates it as PN equations approximate poorly dynamics near the horizon.

The AK waveform is very efficient from a computational point of view, and for that reason, it is widely used to study EMRIs. However, its poor accuracy does not allow us to represent the expected signal.

B. Numerical kludge

The NK waveform has been developed by [27], and the idea behind it is to combine the accurate particle trajectory to an approximate expression for the GW emission. The first step in constructing a NK waveform is to compute the trajectory that the inspiraling body follows in the Boyer-Lindquist coordinates of the Kerr spacetime of the central black hole. Bound black hole orbits (in the absence of GW radiation) admit four constants of the motion, allowing us to rewrite the geodesic equations as a system of first-order differential equations. These four constants are the rest mass, the energy E , the axial angular momentum L_z , and the Carter constant Q [38]. The first-order equations can be written in the following form:

$$\begin{aligned} \frac{dr}{d\tau} &= \pm\sqrt{V_r}, \\ \frac{d\theta}{d\tau} &= \pm\sqrt{V_\theta}, \\ \frac{d\phi}{d\tau} &= V_\phi, \\ \frac{dt}{d\tau} &= V_t. \end{aligned} \quad (6)$$

Solutions of the geodesic equations [Eq. (6)] are uniquely determined if we specify E , L_z , and Q and the initial orbital position. Often these quantities are expressed in terms of $(r_a, r_p, \theta_{\text{min}})$: from the roots of V_r we can determine the periastron r_p and the apoastron r_a , and following their definition, we arrive to compute the semilatus rectum and the eccentricity. The angle θ_{min} is instead determined by the root with the smallest value of V_θ (turning point). Exploiting Q and L_z , we can obtain the inclination angle ι , as $\cos \iota = L_z / \sqrt{Q + L_z^2}$. The geodesic motion has to be augmented with a dissipation due to GWs, which (in adiabatic approximation) is described by orbit average changes on the orbital constants. Then, Eq. (6) are integrated to obtain the trajectory of the CO. Finally, the waveform can be obtained from the inspiral trajectory using quadrupole (or quadrupole-octupole) approximation as it were in the flat spacetime.

The NK waveform provides a good agreement with the numerical Teukolsky-based waveform, thus superseding the AK. The model, however, is computationally more expensive since it requires the integration of the trajectory both in phase and coordinate space.

C. Augmented analytic kludge

The AK model can be up to 15 times faster [39] than the NK model at generating waveforms. This efficiency could increase for longer integrations, but not at higher eccentricity because more harmonics are needed in the AK approximation. However, AK waveforms suffer, already at the early-inspiral stage, from dephasing with respect to NK waveforms due to the mismatched frequencies in the two models. The AAK model was first introduced in [40] to cure this issue. The AAK waveform uses information from the NK model to improve the faithfulness of AK waveforms without significantly increasing their computational cost. The main idea is to extend parameters of the AK model beyond their physical meaning to match the frequencies of NK waveforms. AAK generates a small section of trajectory with NK, and then maps the AK trajectory to the NK result and finds out the best-fit parameters. Let us give a more detailed description of this procedure.

A bound geodesic orbit in Kerr spacetime is characterized by three fundamental frequencies $\Omega_{r,\theta,\phi}$ for the radial, polar, and azimuthal components of the motion. These frequencies take a simple form, choosing a timelike

parameter $\lambda = \int dt/\Sigma$ (the so-called *Mino time*). Indeed, the frequencies are given by

$$\begin{aligned}\Omega_r &= \frac{2\pi}{\Lambda_r \Gamma}, \\ \Omega_\theta &= \frac{2\pi}{\Lambda_\theta \Gamma}, \\ \Omega_\phi &= \lim_{N \rightarrow \infty} \frac{1}{N^2 \Lambda_r \Lambda_\theta \Gamma} \int_0^{N\Lambda_r} d\lambda_r \int_0^{N\Lambda_\theta} V_\phi d\lambda_\theta, \quad (7)\end{aligned}$$

where $\Lambda_{r,\theta}$ represent the radial and polar period, and $\Gamma = \langle dt/d\lambda \rangle$ the average rate (in analogy with the Lorentz factor),

$$\begin{aligned}\Lambda_r &= 2 \int_{r_p}^{r_a} \frac{dr}{\sqrt{V_r}}, \\ \Lambda_\theta &= 4 \int_{\theta_{\min}}^{\pi/2} \frac{d\theta}{V_\theta}, \\ \Gamma &= \lim_{N \rightarrow \infty} \frac{1}{N^2 \Lambda_r \Lambda_\theta} \int_0^{N\Lambda_r} d\lambda_r \int_0^{N\Lambda_\theta} V_t d\lambda_\theta. \quad (8)\end{aligned}$$

In terms of the fundamental frequencies, the periapsis and Lense-Thirring precession rates are given by $\Omega_{\text{pre}} = \Omega_\phi - \Omega_r$ and $\Omega_{\text{LT}} = \Omega_\phi - \Omega_\theta$, respectively. We can introduce for a BH with mass M the dimensional frequencies $\omega_{r,\theta,\phi} = \frac{\Omega_{r,\theta,\phi}}{2\pi M}$, which in the Newtonian limit tend to the orbital frequency, i.e., $f_{\text{orb}} = \omega_r = \omega_\theta = \omega_\phi$. In this limit, the periapsis and Lense-Thirring precession are zero, but in AK they are introduced manually through the PN equation. In particular, the precession frequency would be $f_{\text{pre}} = \dot{\gamma} + \dot{\alpha}$ and Lense-Thirring frequency $f_{\text{LT}} = \dot{\alpha}$. Thus, we can match the three frequencies in the AK and in the Kerr formalism, producing an endomorphism over the AK parameter space. Typically, the frequencies are expressed in terms of (M, a, p) , but also the set (e, ι, p) can be used. For example, given (M, a, p) of a BH, the match between frequencies is done by introducing the nonphysical values $(\tilde{M}, \tilde{a}, \tilde{p})$ as

$$\begin{aligned}f_{\text{orb}} &= \dot{\phi}(\tilde{M}, \tilde{a}, \tilde{p}) = \omega_r(M, a, p), \\ f_{\text{LT}} &= \dot{\alpha}(\tilde{M}, \tilde{a}, \tilde{p}) = \omega_\theta(M, a, p) - \omega_\phi(M, a, p), \\ f_{\text{pre}} &= \dot{\gamma}(\tilde{M}, \tilde{a}, \tilde{p}) + \dot{\alpha}(\tilde{M}, \tilde{a}, \tilde{p}) \\ &= \omega_\phi(M, a, p) - \omega_r(M, a, p).\end{aligned} \quad (9)$$

Substituting the parameters $\tilde{M}, \tilde{a}, \tilde{p}$ with (M, a, p) in the AK model provides a correction of its frequencies along the entire inspiral trajectory. Next, the waveform can be generated using the AK framework with the improved orbital motion.

An improved AAK model is implemented by [26] as a part of the package *FastEmriWaveform* (FEW). The

“classic” AAK described above, builds a trajectory by using the frequency evolution from the numerical kludge and mapping it onto the frequency basis. The new improved AAK, instead, computes directly the fundamental frequencies exploiting a 5PN trajectory developed in [41]. Finally, a fast method of plunge handling has been added to the AAK implementation. In general, the CO plunges when its orbit along the phase-space trajectory becomes unstable (separatrix).

The separatrix depends upon three parameters (e, p, ι) . In the literature, orbits along the separatrix with $e = 0$ are referred to as the innermost stable spherical orbit (ISSO). If $\iota = 0$, this orbit is usually called the innermost stable circular orbit (ISCO) instead. In the FEW framework, the trajectory has been computed until it reaches 0.1 of the separatrix, computed numerically as shown in [42]. Noteworthy, another EMRI waveform has been embedded in FEW; it is a fully relativistic waveform based on the interpolated self-force calculations. However, this model covers only nonspinning BHs.

IV. FROM EMRI POPULATION TO GWB

In order to provide astrophysical motivated estimates of the GWB generated by a cosmic population of EMRIs, we consider several models presented in [25], reported in Table I for completeness.

The main idea behind the 12 different models was to bracket the expected range of EMRI rates spanning the wide range of uncertainties affecting EMRI formation, including the cosmic evolution of the MBH mass function, the relation between MBH mass and density of the surrounding stellar environment, the impact of core scouring following galaxy mergers, and so on. For each model, the distribution of cosmic EMRIs was built, and mock catalogs were generated by Monte Carlo sampling from the distribution. In particular, 10 Monte Carlo realizations of the expected population of EMRIs plunging within 1 year were generated. Put together, they consist of a library which includes all EMRI events occurring in the Universe during 10 years for each of the 12 models. Every catalog provides the primary mass M , the redshift of the event z , and MBH spin a of each event. For the EMRIs, it has been considered a secondary mass of $10M_\odot$ or $30M_\odot$ as shown in the table. The expected event rate (number of sources per year) varies across models from tens to tens of thousands. This large range reflects the astrophysical uncertainties on the model parameters, e.g., the choice of the number of direct plunges or the MBH mass function. Full details about the models can be found in [25]. The GW signal modeling requires 14 parameters listed in Sec. III. For each EMRI inside the catalogs, we associate the cosine of inclination angle ι by drawing it randomly from a uniform distribution between $[-1, 1]$ since we consider that EMRIs form in a spherical bulge and not in a disklike structure. Note that ι

TABLE I. List of EMRI models taken from [25] and considered in this work. Column 1 defines the label of each model. For each model the following quantities are specified: the MBH mass function (column 2), the MBH spin model (column 3), whether or not the effect of cusp erosion is included (column 4), the $M - \sigma$ relation (column 5), the ratio of plunges to EMRIs (column 6), the mass of the COs (column 7), and the total EMRI merger rate (yr^{-1}) up to $z = 4.5$ (column 8). In columns 9 and 10, the detected EMRI rate per year is reported for two different variants of AK waveforms (AKS and AKK truncate the waveform, respectively, at Schwarzschild and Kerr ISCO). The full definition and implementation of each ingredient entering the models can be found in [25].

Model	Mass function	MBH spin	Cusp erosion	$M - \sigma$ relation	N_p	CO mass [M_\odot]	Total	EMRI rate [yr^{-1}] Detected (AKK)	Detected (AKS)
M1	Barausse12	a98	yes	Gultekin09	10	10	1600	294	189
M2	Barausse12	a98	yes	KormendyHo13	10	10	1400	220	146
M3	Barausse12	a98	yes	GrahamScott13	10	10	2770	809	440
M4	Barausse12	a98	yes	Gultekin09	10	30	520	260	221
M5	Gair10	a98	no	Gultekin09	10	10	140	47	15
M6	Barausse12	a98	no	Gultekin09	10	10	2080	479	261
M7	Barausse12	a98	yes	Gultekin09	0	10	15800	2712	1765
M8	Barausse12	a98	yes	Gultekin09	100	10	180	35	24
M9	Barausse12	aflat	yes	Gultekin09	10	10	1530	217	177
M10	Barausse12	a0	yes	Gultekin09	10	10	1520	188	188
M11	Gair10	a0	no	Gultekin09	100	10	13	1	1
M12	Barausse12	a98	no	Gultekin09	0	10	20000	4219	2279

spans in the interval $[0, \pi]$ and that prograde orbits are associated for $0 \leq \iota \leq \pi/2$, while a retrograde orbit for $\pi/2 \leq \iota \leq \pi$. To determine the initial eccentricity and the semilatus, we exploit [25], which has shown that, evolving large samples of COs, the eccentricity distribution at plunge is nearly flat in the range $0 < e_{\text{pl}} < 0.2$. Then, we use the same strategy adopted in [23]:

- (i) For each event in the catalogs, we draw the eccentricity at the last stable orbit e_p in the range $[0, 0.2]$. Since we already know M and the spin a , we can also estimate the radius at the last stable orbit.
- (ii) We integrate the orbital elements of the event backwards in time for T_{back} years (i.e., using the orbit-averaged equations in [14]).
- (iii) We draw $N_{\text{back}} = \text{int}(T_{\text{back}}/10)$ points randomly in the range $[0; T_{\text{back}}]$ in order to select different evolutionary points of each EMRI. Here, the division by 10 is made because, as we mentioned, we collect 10 catalogs of EMRIs for every model, each representing one year of observation. T_{back} is computed as $T_{\text{back}} = 20 \left(\frac{M}{10^4 M_\odot} \right)$ yr, because the time taken to cover the same amount of gravitational radii scales linearly with MBH mass. Thus, EMRIs with low mass primarily emit in the LISA band only over the last years of their inspiraling, while EMRIs with larger MBH masses emit in the LISA band for much longer prior to the plunge.
- (iv) For each N_{back} , we collect the semimajor axis $a_{\text{sm},0}$ and the eccentricity e_0 at the corresponding backward time. Following the definition of semilatus, we compute it as $p_0 = a_{\text{sm},0}(1 - e_0^2)$.

With this procedure, we are creating N_{back} copies of each EMRI from the catalogs of [25] with same redshift, masses, spin, and inclination angles. Nevertheless, the procedure is not expected to introduce any bias in the computation of the background because the catalogs provide a smooth coverage of the relevant MBH mass redshift and spin range.

Concerning the sky position (θ_S, ϕ_S) and spin orientation (θ_K, ϕ_K) , we assume that the corresponding unit vectors are isotropically distributed on a sphere. Finally, the three initial phases corresponding to the orbital phase, the precession phase of the periapsis, and the precession phase of the orbital plane (Lense-Thirring) are uniformly distributed between 0 and 2π . To compute the waveform with AAK, we convert them in terms of radial, polar, and azimuthal phases $(\Phi_{r,0}, \Phi_{\theta,0}, \Phi_{\phi,0})$ using the system analogues to Eq. (9).

A. Definition of the catalog

After completing the list of parameters, we compute the signal of each EMRI. Since we want to evaluate the power spectrum of the unresolved GW background in the LISA detector, we have to work in frequency domain using the optimal combination of TDI described in Sec. II B. However, the AAK code returns the signal in time domain. To inject the signal in the detector, we therefore perform the following steps:

- (1) Computation of $h_+(t), h_\times(t)$ with AAK + 5PN trajectory, considering an observation time of 4 years and a measuring cadence of 15 s.
- (2) Computation of variables $A(t), E(t), T(t)$.

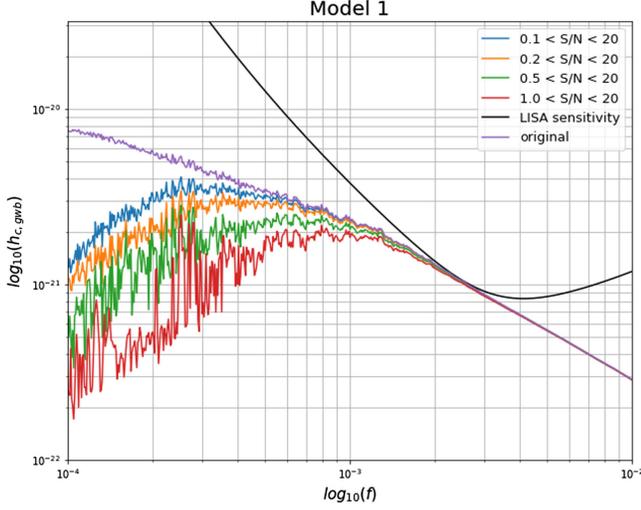


FIG. 3. Characteristic strain of the GWB generated for model M1 listed in Table I. The colors refer to different SNR thresholds for the EMRI included in the computation, as labeled in figure. The gravitational signal has been computed using a simplified, inclination-polarization averaged version of the AK waveform.

(3) Computation of Fourier transforms $\tilde{A}(f)$, $\tilde{E}(f)$, $\tilde{T}(f)$.

This procedure has a variable computational cost depending on the initial parameters of the source. For example, a source with high eccentricity e_0 takes more time because the number of harmonics needed in the computation is higher.

The whole procedure takes on average ~ 120 seconds per source, which is a non-negligible computational time, given that the populations we built are composed of a number of sources which spans from tens of thousands up to several millions. For instance, model M12 counts almost 16 million EMRIs, meaning the background evaluation would take about 60 years on a single CPU. We, therefore, need to find a strategy to speed up the process. A possibility is to select only the sources that produce a significant contribution to the GWB. For this purpose, we evaluate a preliminary SNR for each EMRI exploiting the inclination-polarization averaged version of the AK waveform, following the same approach as in [23] and truncating the EMRI evolution at the Schwarzschild last stable orbit considering 4 years of observation time. Then, we introduce a cut in the population, removing all the sources below a chosen SNR threshold. Finally, we compute what fraction of the total SNR is lost due to the cut. We tried different cuts, and we decided to fix the threshold to $\rho = 1$, which allows removing the greatest number of sources without losing too much total SNR. After removing the weak sources, we observe an SNR loss of $\sim 5\%$ only (cf. columns 4 and 5 of Table II). Figure 3 shows an example of this simplified GWB computation for model M1 for different SNR cuts. The selection affects the GWB mostly at $f < 1$ mHz, where LISA starts to lose sensitivity. Conversely, where

TABLE II. The effect of source selection. Columns 2 and 3 are the initial and final number in the EMRIs catalogs before and after the source selection process. Column 4 is the original SNR of GWB calculated using the simplified version of the AK waveform assuming 4 years of observation time. Column 5 is the SNR after removing sources with $\text{SNR} < 1$. Column 6 is the final SNR computed after removing also sources with $e_0 > 0.9$ or $p_0 < 10$. The last column corresponds to the number of detected EMRIs, defined as those with $\text{SNR} > 20$.

Model	N_{tot}	N_f	SNR_{tot}	$\text{SNR}_{\rho > 1}$	SNR_f	Detections
M1	1225158	31764	571	534	460	366
M2	580149	17030	434	418	352	386
M3	2030059	68100	1174	1121	839	1123
M4	43607	11166	960	926	750	1713
M5	1772409	44011	727	695	535	622
M6	396800	2486	48	44	27	35
M7	8218425	303348	4940	4748	3960	4619
M8	89010	3620	63	61	51	53
M9	872231	31356	501	480	419	426
M10	823589	30597	478	459	408	406
M11	34724	287	4.53	4.11	2.32	0
M12	16547658	394583	6475	6200	4553	5580

the instrumental curve is closer to the EMRI GWB level, the backgrounds are similar to each other.

The choice of the AAK waveform imposes two further selection criteria: (i) a cut in eccentricity and (ii) a cut in semilatus rectum. The number of harmonics needed to construct the waveform, and consequently the computational cost, is a steep function of eccentricity. We verified that for sources with eccentricity bigger than 0.9, the waveform code does not terminate the computation, and therefore, we cut from the catalog all sources with $e_0 > 0.9$. Moreover, we also remove the sources with initial semilatus smaller than 10 since the global fit for the inspiral turns out to be difficult for such tight EMRIs. The cut in semilatus removes only $\sim 1\%$ of the sources from the original catalogs, as reported in the third column of Table III. To summarize, the final EMRIs catalogs are composed by sources with $e_0 < 0.9$, $p_0 > 10$ and individual SNR greater than 1, calculated with a simplified AK waveform. It is important to notice that cutting in eccentricity and the semilatus rectum induced the loss of a significant part of the original background information. As reported in column 6 of Table II, the final SNR, after taking into account this further selection, is about 20%–30% lower compared to the full population. Despite the number of removed EMRIs at high eccentricity being larger, most of the signal loss is due to the selection in the semilatus rectum. Indeed, we found that the percentages of SNR lost due to semilatus selections are greater than the other two, as seen in Table IV. Figure 4 represents the distribution of the sources in the plane eccentricity/semilatus for model M1, even if we can find a similar behavior in the other model. First, we observe that $e > 0.9 \cap p < 10 = \emptyset$, and the two cuts are independent.

TABLE III. Percentage of EMRIs removed due to the selection procedure aiming at speeding up the GWB evaluation. Column 2 refers to the sources subtracted with $e_0 > 0.9$, while Column 3 is the fraction of sources with $p_0 < 10$.

Model	$\%_{e>0.9}$	$\%_{p<10}$
M1	26.5	1.11
M2	30.3	0.92
M3	28.0	0.84
M4	26.8	1.08
M5	22.2	1.17
M6	10.3	1.56
M7	26.5	1.11
M8	28.1	1.12
M9	26.3	1.14
M10	26.8	1.07
M11	11.1	1.35
M12	21.6	1.15

TABLE IV. Percentage of SNR after different sources selections. Column 2 is the total percentage lost; column 3 is referred to after removing the dimmest source with $\text{SNR} < 1$, while columns 4 and 5 are, respectively, the SNR after deleting sources with $p < 10$ and $e > 0.9$.

Model	$\% \text{SNR}_{\text{tot}}$	$\% \text{SNR}_p$	$\% \text{SNR}_e$	$\% \text{SNR}_e$
M1	20	7	11	2
M2	18	4	11	3
M3	28	5	19	4
M4	22	4	13	6
M5	26	4	18	4
M6	43	8	30	5
M7	20	4	11	5
M8	16	3	12	1
M9	16	4	9	3
M10	14	4	8	2
M11	49	9	39	1
M12	29	4	18	7

The gradient of color shows that the $p < 10$ region hosts multiple sources with single SNR between 10 and 20, which, in principle, could contribute to increasing the final background, whereas only low SNR sources are found at $e > 0.9$.

Despite the above-mentioned caveats, we stress once again the efficiency of this selection: we are able to maintain almost 80% of the original signal just considering on average 3% of the sources in the catalogs, making the computation cost manageable.²

²Notice that the computational cost is still non-negligible. Taking, for instance, M12, the computation of all its sources after the selections requires almost one week exploiting 100 CPU.

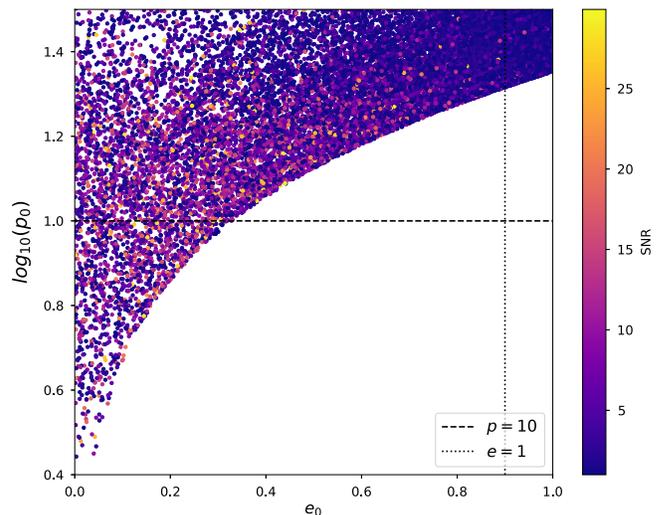


FIG. 4. Distribution of EMRIs in the plane $e\text{-}\log_{10}(p)$ for model M1. The color scale refers to the SNR of EMRIs computed through the simplified AK waveform. Note that louder EMRIs are generally characterized by smaller p_0 rather being very eccentric.

B. Background computation

Using the population described in the previous section, we generate the AAK waveform for each EMRI, inject them in the optimal LISA TDI variables, and analyze the data directly on those TDI data streams. For each model, we compute the initial background as the sum of all the signals for both channels A and E (in the frequency domain). We should remove the sources whose SNR exceeds a set threshold to obtain the final confusion noise. The detection threshold is adjusted for the different source types, taking into account search and parameter estimation studies. The Mock LISA Data Challenge results suggest that EMRIs with SNR as low as $\rho_0 = 20$ could be identified [21]. The overall SNR for a single source is computed from TDI data by taking the root of the sum of the square SNRs in each TDI channel, as defined in Eq. (1). However, in the presence of a GWB, the PSD in each channel becomes $S_n(f) = S_{\text{instr}}(f) + S_{\text{gwb}}(f)$, taking into account both the instrumental noise, $S_{\text{instr}}(f)$, and the confusion noise, $S_{\text{gwb}}(f)$. To solve this problem, we use the iterative foreground estimation (IFE) algorithm, developed by [28], which can be used for characterizing a background signal that originates from a combination of different types of sources. IFE is based on an iterative process to evaluate $S_n(f)$, whose main steps are the following:

- (1) After generating the total signal, $S_{\text{gwb}}(f)$, from the full sample of sources, the code computes the SNR of each source considering only the LISA instrumental noise. We call it SNR in isolation ρ_i^{iso} .
- (2) The code evaluates the new PSD, $S_{n,k}(f)$, as the sum of $S_{\text{instr}}(f)$ and the first estimation of $S_{\text{gwb}}(f)$. Then,

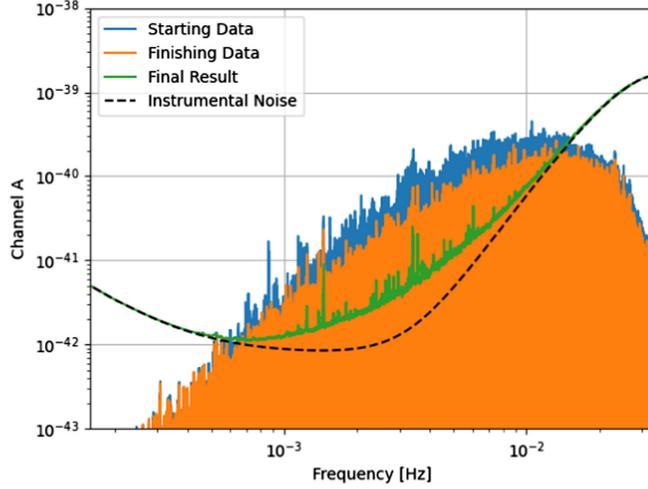


FIG. 5. Illustration of the procedure for estimating the unresolved GWB from EMRIs. This plot refers to M12, and it takes four iterations for the algorithm to converge. The starting data are represented in blue, while orange shows the result at the end of the procedure. The final combined instrument plus median confusion noise is represented in green.

the algorithm computes the new SNR ρ_i^{iso} , using $S_{n,k}(f)$ (k refers to the iteration number).

- (3) If $\rho_i^{\text{iso}} > \rho_0$, the code subtracts the source from the confusion noise $S_{\text{gwb}}(f)$. Instead, if $\rho_i^{\text{iso}} < \chi\rho_0$, the program skips the computation of the SNR for that source and adds it automatically to the final foreground. Here, χ is a prefixed factor, typically smaller than one.
- (4) After subtracting the bright sources in the previous step, the algorithm returns to step 2., estimating $S_{n,k+1}(f)$ and iteratively performing the same procedure.

The iteration proceeds until there are no more sources with $\rho_i^{\text{iso}} > \rho_0$ or alternatively when $S_{n,k+1}(f) \sim S_{n,k}(f)$, within a fractional tolerance limit. Typically, the algorithm requires one to five iterations to converge. An illustration of the procedure is represented in Fig. 5.

V. RESULTS

Combining the EMRI waveform modelization, the LISA TDI, and the GWB computation, we evaluated the EMRI background for the 12 catalogs as shown in Fig. 6. The GWB curves for all models lie between the optimistic and the pessimistic scenarios, which are models M12 and M11, respectively. Uncertainties in the EMRI GWB estimate span about 3 orders of magnitude, which is consistent with the uncertainty in the EMRI rates reported by [25]. One thing is noteworthy: part of the investigated models predicts a GWB comparable to or higher than the LISA instrumental noise, which therefore cannot be neglected when considering the detectability of other sources [23]. M12 and M7

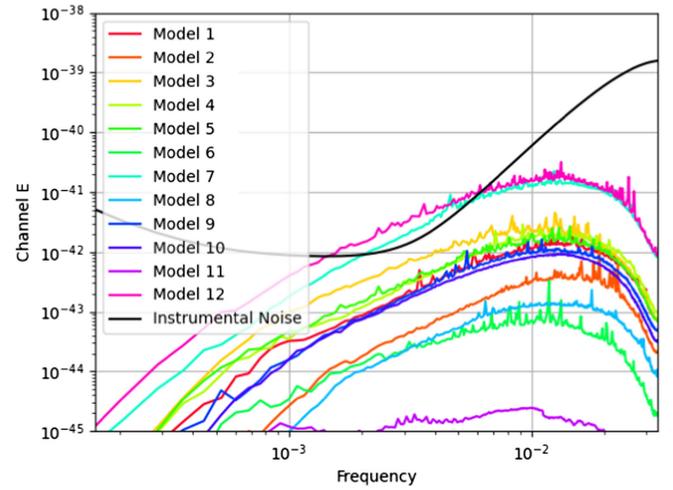
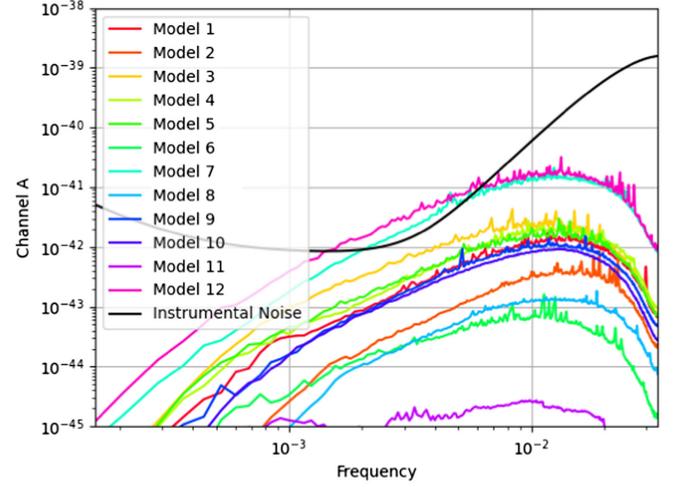


FIG. 6. The GWB generated in the 12 different EMRI formation scenarios represented in the TDI variables A (top panel) and E (bottom panel). The black line is the LISA noise in the correspondent channel.

overcome the LISA level, while M1, M3, M4, M5, M9, and M10 are less than an order of magnitude apart from the black line in the bucket, thus making a non-negligible contribution to the total PSD.

To characterize each single GWB, we compute its SNR. Since we are neglecting the T channel, the overall SNR is just the square root of the sum in quadrature of the A and E channels. The SNR of GWB in A and E TDI LISA channel has been computed through [43,44]

$$\text{SNR}_i = \sqrt{T_{\text{obs}} \int_0^\infty df \left(\frac{S_{\text{gwb},i}(f)}{S_{n,i}(f)} \right)^2} \quad (10)$$

where i can be A or E, and $S_{\text{GWB},i}$ is the PSD due to the GWB. Numbers are reported in Table V assuming 4 years of data collection. For most models, the SNR of the

TABLE V. Number of resolvable sources and residual GWB SNR for the investigated astrophysical models. The results have been obtained considering 4 years of observation and using the AAK waveform [26].

Model	Detections	SNR _{GWB}
M1	139	180
M2	42	40
M3	346	441
M4	516	235
M5	188	252
M6	13	21
M7	724	980
M8	19	22
M9	108	160
M10	97	136
M11	0	1.44
M12	891	1146

unresolved GWB is larger than 100, making the signal easily detectable. The GWB produced by models M2, M6, and M8 has a moderate SNR and might be hard to detect, whereas model M11 does not produce an appreciable signal.

A. Comparison with the literature

We can perform a comparison between our results and the findings of [23] by matching the SNRs shown in Table V with those in the second last column of Table II. The latter are computed following the same approach as [23], but applying the same sources selection described in Sec. IV A. We observe that our SNRs are lower with respect to those of [23] by a factor between 2 and 4, depending on the model. This difference may be due to some discrepancies between the two approaches used to estimate the EMRI GWB.

The first obvious reason can be ascribed to the choice of the waveform. Indeed, we perform our computation using the innovative AAK, which is more accurate (at the level of the numerical kludge) with respect to the AK waveform, used in [23]. As shown in Fig. 12 of [23], the AK model captures the salient features of the waveform, but is less precise in modeling the spectrum close to the final plunge, because it overestimates the frequency of the last stable orbit. This effect is enhanced when truncating the inspiraling at Kerr separatrix (green line), but it is also present in the Schwarzschild case. Thus, the excess power in the AK signal at higher frequencies can possibly induce a boost in the SNR. However, it is difficult to track this effect, since it can vary source by source given the complexity of the EMRIs signal.

A second cause contributing to the discrepancy can be related to the computation of the SNR itself. Indeed, in our study, the final SNRs are computed including in the noise spectral density both the instrumental and the EMRIs

background component, unlike Table II (and [23]) where the astrophysical noise has been neglected. Consequently, the SNR can be lower since the noise level is higher. Moreover, neglecting the EMRI GWB in the background computation can lead to an incorrect subtraction of the resolvable sources.

Finally, the GWB computation in [23] is based on the use of inclination-polarization averaged fluxes, while we automatically take into account the inclination of each individual system with respect to the observation frame injecting the signal in the TDI variables, which ultimately makes our estimation more realistic.

In Table V, we also report the number of resolvable sources which have been subtracted from the background using the iterative algorithm. We compare our results to the only other two estimates found in the literature, in [23] and [25]. In the last column of Table II, we report the EMRI detected in 4 years obtained with the same procedure as in [23], but adjusted with respect to our selection. Our detection rates are more than a factor of 3 smaller, except for M11, for which we confirm that no EMRI can be resolved. As mentioned, this discrepancy is twofold: (i) due to the AK waveform employed and on whether it is truncated at the Kerr or Schwarzschild last stable orbit, the truncation at the Kerr separatrix causes an excess power at high frequencies compared to the AAK waveform, inducing a boost in the SNR, and (ii) the instrumental noise level assumed in [25] is about 1.5 lower (above 2–3 mHz) compared to the currently adopted. In addition, some of the removed EMRIs, especially those with $p_0 < 10$, might be sufficiently loud to be resolved (without affecting the GWB). Still, when [23] add (in an approximate fashion) the corresponding EMRI GWB to the LISA sensitivity curve, we find rates of the same order (i.e., column labeled “AKSb” in Table II of [23]).

Compared to numbers obtained by [25] and reported in Table I, our detection rates are even lower. For example, considering the optimistic case M12, the EMRI detection rate reported in Table I is 4219 per year in the Kerr case and 2279 per year in the Schwarzschild case, while we obtain only 891 detections in 4 years, i.e., only ≈ 200 events per year. The difference can be ascribed to a number of causes: the selection of the sources in our catalogs, the choice of the waveform, and how the resolvable sources are subtracted, as explained in the previous comparison with [23]. A direct comparison with [25] is unfortunately not straightforward: [25] considers catalogs of EMRIs plunging over 10 years and computes SNR by considering only the last 2 years before the plunge (integrating backwards in time). They then divide the obtained numbers by 10 to get the detection rates per year. This procedure implies that the EMRI detection rate can somehow be independent of the mission duration; however, as shown in [23], this does not seem to be the case. The EMRI detection rate indeed depends

nonlinearly on the observation time, making any direct comparison flawed.

VI. CONCLUSIONS AND OUTLOOK

In this study, we take another stride towards a comprehensive assessment of the stochastic GWB originating from a population of EMRIs for LISA. Our approach entails incorporating the entire LISA response and utilizing a realistic gravitational waveform. Three main ingredients have been combined to achieve the goal: the generation of populations of EMRIs, the simulation of their GW signals, and their injection in LISA data stream.

We first built the EMRI population from the astrophysical models developed in [25], which include the most relevant ingredients driving EMRI formation in the standard relaxation-driven channel. Following the procedure devised in [23], we constructed from these models the corresponding EMRI catalogs. To speed up the foreground computation, we removed the sources contributing less to the final GWB by putting a minimum threshold in the single SNR pre-computed with a simplified, inclination-polarization averaged version of the AK waveform. Additionally, we removed sources with initial eccentricity exceeding 0.9 and semilatus smaller than 10 to avoid computational issues connected to the choice of the waveform.

We simulated the GW waveform of each EMRI assuming a 4 year mission and exploiting the innovative AAK waveform. We injected the GW signals in LISA using full TDI response. This technique combines the measurements along each LISA link in order to remove the laser frequency noise, which will be the dominant source of the noise. Finally, we estimated the final GWB using an iterative algorithm which removes the resolvable sources from the background. We found that all the astrophysical models, except for one (the pessimistic M11), result in a detectable EMRI GWB, with SNRs accumulated in 4 years of observation ranging from ≈ 20 up to > 1000 , cf. Table V. The background will act mainly in the range 10^{-3} – 10^{-2} Hz; in particular, GWBs from optimistic models (such as M7 and M12) are loud enough to dominate over the LISA instrumental noise in this frequency range. In the same timespan, LISA will allow the identification and characterization of several individual systems, up to ≈ 1000 in the more optimistic case (again, cf. Table V). We stress that since we excluded sources with initial semilatus rectum smaller than 10, which can contribute significantly to the signal, the aforementioned numbers are likely underestimated. Furthermore, given also the complexity of both the EMRI and the procedure computation of the GWB, it is difficult to quantify which ingredients contribute most to the discrepancy with previous results in the literature.

Our findings have several practical consequences for the forthcoming LISA mission. Specifically, EMRI GWB can significantly contribute to the LISA noise level around the sensitivity bucket, possibly jeopardizing the detectability of other interesting sources in the LISA window. In particular, the effect might be severe for two families of GWs sources that are of paramount importance for the mission, namely, the merger of massive black holes binaries seeds occurring at high redshift and the stellar origin black holes binaries. Because of their low mass and high redshift, seed binaries will populate the bottom of the LISA sensitivity bucket, overlapping with the EMRI GWB. SOBH binaries, instead, can be grouped into two categories: slow inspirals and multiband systems. The latter, which are perhaps the most interesting, are observed at $f > 0.01$ Hz and should be unaffected. Conversely, slow inspirals are mostly found at $f < 0.01$ Hz, where the GWB is more prominent and can greatly impact their SNR.

Finally, the EMRI GWB is not the only stochastic signal expected to fall into the LISA frequency band. Other sources of stochastic astrophysical foregrounds are anticipated to be GCBs and extragalactic neutron star and SOBH binaries. In our galaxy alone, it is estimated that there are roughly 250 million detached and 10 million interacting white dwarf binaries [45]. The sensitivity of LISA will allow the detection of thousands of these binaries as individual sources. The GWs from the other millions of binaries, especially below 1 mHz, will combine to form an unresolved background. Unlike EMRIs, these sources are anisotropically distributed in the sky and are concentrated in the disk of our Galaxy. As LISA orbits the sun, its orientation will change continuously, changing the antenna pattern of the detector and, consequently, the directions along which the detector will be most sensitive. The GW signals coming from the Galactic Plane, hosting the vast majority of WD binaries, will therefore fluctuate during a 1-year cycle, making the identification of the WD noise easier. Conversely, extragalactic SOBH and NS binaries will also be isotropically distributed, and their resulting GWB will overlap significantly with the one produced by EMRIs. The detection and characterization of these astrophysical backgrounds is going to be challenging, and further studies and dedicated analysis pipeline developments are needed in order to assess the potential of LISA to identify and separate them from each other.

ACKNOWLEDGMENTS

A. S. acknowledges financial support provided under the European Union’s H2020 ERC Consolidator Grant “Binary Massive Black Hole Astrophysics” (B Massive, Grant Agreement No. : 818691). S. B. acknowledges support from the French space agency CNES in the framework of LISA.

- [1] M. J. Rees, *Nature (London)* **333**, 523 (1988).
[2] J. G. Hills, *Nature (London)* **331**, 687 (1988).
[3] P. Amaro-Seoane, *Living Rev. Relativity* **21**, 4 (2018).
[4] A. King, *Mon. Not. R. Astron. Soc.* **515**, 4344 (2022).
[5] P. Amaro-Seoane, J. R. Gair, M. Freitag, M. C. Miller, I. Mandel, C. J. Cutler, and S. Babak, *Classical Quantum Gravity* **24**, R113 (2007).
[6] M. C. Miller, M. Freitag, D. P. Hamilton, and V. M. Lauburg, *Astrophys. J. Lett.* **631**, L117 (2005).
[7] R. Di Stefano, J. Greiner, S. Murray, and M. Garcia, *Astrophys. J. Lett.* **551**, L37 (2001).
[8] Y. Levin, *Mon. Not. R. Astron. Soc.* **374**, 515 (2007).
[9] P. Amaro-Seoane, H. Audley, S. Babak, J. Baker, E. Barausse *et al.*, [arXiv:1702.00786](https://arxiv.org/abs/1702.00786).
[10] M. Colpi, K. Holley-Bockelmann, T. Bogdanović, P. Natarajan, J. Bellovary, A. Sesana *et al.*, *Bull. Am. Astron. Soc.* **51**, 432 (2019).
[11] G. Nelemans, L. R. Yungelson, and S. F. Portegies Zwart, *Astron. Astrophys.* **375**, 890 (2001).
[12] A. Sesana, *Phys. Rev. Lett.* **116**, 231102 (2016).
[13] T. Hinderer and É. É. Flanagan, *Phys. Rev. D* **78**, 064028 (2008).
[14] P. C. Peters, *Phys. Rev.* **136**, B1224 (1964).
[15] F. D. Ryan, *Phys. Rev. D* **52**, 5707 (1995).
[16] L. Barack and C. Cutler, *Phys. Rev. D* **75**, 042003 (2007).
[17] J. R. Gair, M. Vallisneri, S. L. Larson, and J. G. Baker, *Living Rev. Relativity* **16**, 7 (2013).
[18] E. Barausse, L. Rezzolla, D. Petroff, and M. Ansorg, *Phys. Rev. D* **75**, 064026 (2007).
[19] E. Barausse, V. Cardoso, and P. Pani, in *J. Phys. Conf. Ser.* **610**, 012044 (2015).
[20] J. R. Gair, C. Tang, and M. Volonteri, *Phys. Rev. D* **81**, 104014 (2010).
[21] S. Babak *et al.*, *Classical Quantum Gravity* **27**, 084009 (2010).
[22] L. Barack and C. Cutler, *Phys. Rev. D* **70**, 122002 (2004).
[23] M. Bonetti and A. Sesana, *Phys. Rev. D* **102**, 103023 (2020).
[24] L. Barack and C. Cutler, *Phys. Rev. D* **69**, 082005 (2004).
[25] S. Babak, J. Gair, A. Sesana, E. Barausse, C. F. Sopuerta, C. P. L. Berry, E. Berti, P. Amaro-Seoane, A. Petiteau, and A. Klein, *Phys. Rev. D* **95**, 103012 (2017).
[26] M. L. Katz, A. J. K. Chua, L. Speri, N. Warburton, and S. A. Hughes, *Phys. Rev. D* **104**, 064047 (2021).
[27] S. Babak, H. Fang, J. R. Gair, K. Glampedakis, and S. A. Hughes, *Phys. Rev. D* **75**, 024005 (2007).
[28] N. Karnesis, S. Babak, M. Pironi, N. Cornish, and T. Littenberg, *Phys. Rev. D* **104**, 043019 (2021).
[29] LISA Science Study Team, LISA Science Requirements Document, Technical Report No. 1.0, ESA, 2018.
[30] M. Tinto and S. V. Dhurandhar, *Living Rev. Relativity* **24**, 4 (2005).
[31] N. J. Cornish and L. J. Rubbo, *Phys. Rev. D* **67**, 029905 (2003).
[32] M. Tinto and S. V. Dhurandhar, *Living Rev. Relativity* **17**, 6 (2014).
[33] M. Vallisneri, *Phys. Rev. D* **71**, 022001 (2005).
[34] M. Sasaki and H. Tagoshi, *Living Rev. Relativity* **6**, 6 (2003).
[35] S. A. Teukolsky, *Astrophys. J.* **185**, 635 (1973).
[36] L. Blanchet, *Living Rev. Relativity* **17**, 2 (2014).
[37] P. C. Peters and J. Mathews, *Phys. Rev.* **131**, 435 (1963).
[38] B. Carter, *Phys. Rev.* **174**, 1559 (1968).
[39] A. J. K. Chua, C. J. Moore, and J. R. Gair, *Phys. Rev. D* **96**, 044005 (2017).
[40] A. J. K. Chua and J. R. Gair, *Classical Quantum Gravity* **32**, 232002 (2015).
[41] R. Fujita and M. Shibata, *Phys. Rev. D* **102**, 064005 (2020).
[42] L. C. Stein and N. Warburton, *Phys. Rev. D* **101**, 064007 (2020).
[43] J. D. Romano and N. J. Cornish, *Living Rev. Relativity* **20**, 2 (2017).
[44] C. Caprini, D. G. Figueroa, R. Flauger, G. Nardini, M. Peloso, M. Pironi, A. Ricciardone, and G. Tasinato, *J. Cosmol. Astropart. Phys.* **11** (2019) 017.
[45] J. A. Edlund, M. Tinto, A. Królak, and G. Nelemans, *Phys. Rev. D* **71**, 122003 (2005).