

This is the accepted manuscript made available via CHORUS. The article has been published as:

Faster computation of adiabatic extreme mass-ratio inspirals using resonances

Rebecca Grossman, Janna Levin, and Gabe Perez-Giz

Phys. Rev. D **88**, 023002 — Published 8 July 2013

DOI: [10.1103/PhysRevD.88.023002](https://doi.org/10.1103/PhysRevD.88.023002)

Faster computation of adiabatic EMRIs using resonances

Rebecca Grossman,^{1,*} Janna Levin,^{2,3,†} and Gabe Perez-Giz^{1,4,‡}

¹*Physics Department, Columbia University, New York, NY 10027*

²*Department of Physics and Astronomy, Barnard College of Columbia University, 3009 Broadway, New York, NY 10027*

³*Institute for Strings, Cosmology and Astroparticle Physics, Columbia University, New York, NY 10027*

⁴*Center for Cosmology and Particle Physics, Department of Physics, New York University, New York, NY 10003*

Motivated by the prohibitive computational cost of producing adiabatic extreme mass ratio inspirals, we explain how a judicious use of resonant orbits can dramatically expedite both that calculation and the generation of snapshot gravitational waves from geodesic sources. In the course of our argument, we clarify the resolution of a lingering debate on the appropriate adiabatic averaging prescription in favor of torus averaging over time averaging.

PACS numbers: 97.60.Lf, 04.70.-s, 95.30.Sf, 95.75.Pq

I. INTRODUCTION

Stellar mass compact objects inspiraling into supermassive black holes (SMBHs) will be important astrophysical sources of gravitational waves (GWs) for future space-based detectors. Accurate GW templates for such extreme mass-ratio inspirals (EMRIs) require detailed knowledge of the motion of the source, so there has been a community effort to calculate EMRI trajectories. If we neglect the gravitational self-force of the small object, its orbit is a Kerr geodesic that, up to parameters specifying the initial position, is characterized by three constant orbital parameters: an energy E , an azimuthal angular momentum L_z , and the Carter constant Q . Determining the inspiral is tantamount to calculating how the self-force causes both the positional parameters and the orbital parameters to evolve in time.

Despite ongoing efforts, direct evaluation of the self-force in the Kerr case is still not possible. Accordingly, there have been parallel efforts to approximate its effects. The focus of this paper is the adiabatic approximation, which captures the slow secular evolution of E, L_z, Q by solving a system of ordinary differential equations (ODEs) of the form

$$\frac{dE}{dt} = \mathcal{F}_E(E, L_z, Q) \quad (1a)$$

$$\frac{dL_z}{dt} = \mathcal{F}_{L_z}(E, L_z, Q) \quad (1b)$$

$$\frac{dQ}{dt} = \mathcal{F}_Q(E, L_z, Q) \quad (1c)$$

For now, it suffices to know that the righthand sides (RHSs) of equations (1) are so costly to evaluate that these equations will have to be integrated using a numerical grid. More specifically, the EL_zQ velocity field will be pre-computed only on a dense mesh of points in




FIG. 1: Above is a heuristic depiction of two possible numerical grids that could be used to generate adiabatic approximations (dashed curves) to true inspirals (solid curves) in the orbital parameter space. The dots represent a set of non-resonant grid points and the plus signs a set of resonant grid points. The resulting adiabatic curves are the same in either case but significantly less costly to produce with the resonant grid. A true inspiral may evolve in a way that is not well-approximated adiabatically as it approaches a low-order resonance, as on the left. That divergence, if it occurs, happens regardless of whether the resonant point is used as part of the numerical grid.

EL_zQ -space. Real-time integration of (1) will then rely on derivative values interpolated off of that grid.

In this paper, we advocate building such grids using only points corresponding to resonant geodesics, for which the frequencies of the radial and polar motions are rationally related. As we will see, intermediate calculations that comprise the bulk of the computational expense can be recycled among several Fourier modes on resonant grid points but must be recomputed for every mode in the non-resonant case. We estimate that, compared to using non-resonant grid points, our prescription could reduce the computational cost of an EMRI grid by an order of magnitude or more. The resonant-grid prescription will also facilitate faster computation of GW

*Electronic address: becky@phys.columbia.edu

†Electronic address: janna@astro.columbia.edu

‡Electronic address: gabe@phys.columbia.edu

snapshots from geodesic sources.

We represent our proposal schematically in Figure 1. First, the RHSs of equations (1) are evaluated directly on a grid of either resonant points (dots) or non-resonant points (plus signs). At any other point, the RHS values can be interpolated from the values at the grid points. The adiabatic equations (1) are continuous and smooth, so regardless of which grid is used, integrating them produces the same adiabatic solutions (dashed curves). The only difference is that the resulting adiabatic curves cost significantly less to generate with the resonant grid.

Ref. [1] and more recently, Ref. [2] have noted that such adiabatic approximations may fail to capture important features of the true inspiral (solid curves) near low-order¹ resonances. Heuristically speaking, those authors argue that while an adiabatic solution may remain fairly faithful to an inspiral that steers clear of resonant points (upper right of the figure), those approximations may fare much worse for an inspiral that transits near a resonant point (middle left of the figure). To pre-empt possible confusion, we remark that there is no inconsistency between this observation and our proposal. The decision to include any particular EL_zQ point, resonant or not, in the numerical grid is unrelated to whether the resulting adiabatic curves will faithfully reflect EMRI motion near that same point. The ironic coincidence is that the points where the adiabatic approximation is most likely to fare poorly² are also the optimal grid points for generating adiabatic curves.

The rest of this paper is organized as follows. In Section II, we review some relevant features of resonant Kerr orbits in both physical space and phase space. In Section III, we summarize how one arrives at the adiabatic equations of motion and clarify why an averaging prescription required to derive those equations must be a torus average rather than a time average. While averaging procedures in the EMRI problem have been discussed in this and other contexts [1, 3–7], the rationale for using torus versus time averages in the adiabatic approximation has not previously been made explicit in the literature. Partly to help make the averaging argument and partly because we will focus on frequency-domain approaches to solving the adiabatic equations, Section III also provides some necessary mathematical background on Fourier analysis in both the non-resonant and resonant cases. With a clear view of the adiabatic program now in hand, Section IV presents the main result of the paper, namely a concrete prescription for computational savings that frequency-domain EMRI codes can leverage by using resonances. Finally, Section V speculates about how a more unorthodox use of resonances could

offer additional efficiencies provided it can be practically implemented.

II. RESONANT KERR ORBITS

The paramount role of resonant orbits was the central theme of an earlier series of papers [8–10]. (We use the terms “resonant”, “closed”, and “periodic” interchangeably.) A spectrum of closed orbits, which manifests as a spectrum of multi-leaf clovers, entirely structures black hole dynamics. Although completely closed orbits must return to their initial values³ of (r, θ, φ) simultaneously, only the r - θ periodicity detailed in Ref. [9] is relevant to the present work. The rational number associated with the r - θ frequencies determines the multi-leaf clover geometry. What’s more, that rational obediently stacks in energy monotonically: lower rationals correspond to lower energies than do higher rationals. Such orbits constitute a measure zero set but are nonetheless dense in the phase space. Every non-resonant orbit is arbitrarily close to some resonant orbit.

We first consider resonant orbits in *physical space* and then again in *phase space*. Since we will be concerned with functions that do not depend explicitly on either the azimuthal angle or on coordinate time, it will suffice for us to restrict attention to geodesic motion in two coordinates (r, θ) in physical space and to the projection of the motion into a 4D submanifold of the phase space spanned by (r, θ) and their conjugate momenta.

A. Resonant Orbits in Physical Space

The black hole is completely characterized by its mass M and spin a . The geodesic of the lighter companion is characterized by four dimensionless constants μ, E, L_z, Q . In Boyer-Lindquist coordinates and dimensionless units, which is equivalent to setting $M = \mu = 1$, the radial and polar Kerr equations of motion can be written as

$$\dot{r} = \pm \sqrt{R(r)} \quad (2)$$

$$\dot{\theta} = \pm \sqrt{\Theta(\theta)} \quad , \quad (3)$$

where

$$R(r) \equiv - (1 - E^2) r^4 + 2r^3 \quad (4)$$

$$- [a^2 (1 - E^2) + L_z^2] r^2 + 2(aE - L_z)^2 r - Q\Delta$$

$$\Theta(\theta) \equiv Q - \cos^2(\theta) \left\{ a^2 (1 - E^2) + \frac{L_z^2}{\sin^2 \theta} \right\} \quad (5)$$

and

$$\Delta \equiv r^2 - 2r + a^2 \quad . \quad (6)$$

¹ A resonance is low-order if the numerator and denominator of the rational frequency ratio are both small integers.

² To balance the argument, Ref. [1] also offers plausible reasons why the adiabatic approximation may still be valid near resonances.

³ φ need only return to its initial position mod 2π .

An overdot denotes differentiation with respect to Mino time [11], λ , related to proper time τ by

$$\frac{d\tau}{d\lambda} = \Sigma \equiv r^2 + a^2 \cos^2 \theta \quad . \quad (7)$$

The advantage of using Mino time is that the r and θ equations of motion decouple. To make connections with observations, we will often care about how certain quantities evolve with respect to coordinate time t . However, coordinate time turns out to be mathematically cumbersome, so throughout this paper, we perform all intermediate calculations related to such quantities by first changing variables to Mino time.

Solving equations (4) and (5) for the radial and polar turning points, we find that the radial coordinate varies between a periastron r_p and an apastron r_a and that the polar coordinate similarly varies between some minimum value θ_{min} and maximum value $\theta_{max} = \pi - \theta_{min}$. All turning points depend only on the constants E, L_z, Q . We introduce the simplifying notation

$$\vec{\mathcal{E}} \equiv (E, L_z, Q) \quad (8)$$

for those 3 orbital parameters and reserve the symbol \mathcal{E} to refer to any one of E, L_z, Q individually.

The radial and polar coordinates are each periodic with respective Mino periods

$$\Lambda_r = 2 \int_{r_p}^{r_a} \frac{dr}{\sqrt{R(r)}} \quad (9a)$$

$$\Lambda_\theta = 4 \int_{\theta_{min}}^{\pi/2} \frac{d\theta}{\sqrt{\Theta(\theta)}} \quad (9b)$$

and corresponding frequencies

$$\Omega_r = \frac{2\pi}{\Lambda_r} \quad (10a)$$

$$\Omega_\theta = \frac{2\pi}{\Lambda_\theta} \quad . \quad (10b)$$

The radial and polar velocities are also periodic with the same corresponding periods and frequencies. If the frequency ratio

$$1 + q_{r\theta} \equiv \frac{\Omega_\theta}{\Omega_r} \quad (11)$$

is a rational number, an r - θ projection of the resulting orbit closes after a finite time.

Note from equations (4), (5) and (9) that the frequencies and $q_{r\theta}$ depend only on the constants $\vec{\mathcal{E}}$. $q_{r\theta}$ is also a topological invariant and thus coordinate independent. A $q_{r\theta}$ for which the relatively prime numerator and denominator are both low-valued integers will be referred to as “low-order”. We arbitrarily call low-order resonant orbits those for which the numerator and denominator of $q_{r\theta}$ are each less than 10.

Projections of periodic orbits into the r - θ plane produce Lissajous figures. The top panel of Figure 2 shows




FIG. 2: Top: A low-order periodic orbit with $q_{r\theta} = \frac{1}{2}$, $a = 0.9$, $E = 0.954788$, $L_z = 2.65115$ and $Q = 0.944969$, projected on the r - θ plane (we plot $r \cos \theta$ to make the figure more viewable) with initial conditions $r_0 = r_a = 17.8148$ and $\theta_0 = \theta_{min} = 1.22079$. Bottom: A non-resonant orbit with $q_{r\theta} \approx \frac{125,857}{250,000}$, $a = 0.9$, $E = 0.956$, $L_z = 2.65115$ and $Q = 0.944969$, with initial conditions $r_0 = r_a = 18.4568$ and $\theta_0 = \theta_{min} = 1.22076$.

the Lissajous figure of a periodic orbit with a low-order $q_{r\theta}$, while the bottom panel shows the analogous projection of a neighboring orbit with an irrational $q_{r\theta}$.

The figures produced by projecting into the r - θ plane are less topologically insightful than the figures in an orbital plane, loosely defined in the Kerr system as the plane perpendicular to the orbital angular momentum [9]. In the orbital plane, the rational $q_{r\theta}$ has powerful topological information and can be interpreted as $q_{r\theta} = w + v/z$, where the integer w represents the number of nearly circular whirls near periastron, the integer z is the number of elliptical leaves in the multi-leaf clover pattern, and the integer v is the order in which the leaves are hit [8, 9]. To illustrate, the same two orbits of Figure 2 are plotted in the orbital plane in Figure 3. The orbit in the top panel of Figure 3 has $q_{r\theta} = 1/2$ and therefore corresponds to a 2-leaf clover, as is now evident. The bottom panel non-resonant orbit is close to a resonant orbit with $q_{r\theta} \approx \frac{125,857}{250,000}$ which would correspond to a 250,000-leaf clover that skips 125,857 leaves in the pattern each time. Notice that $\frac{125,857}{250,000} = \frac{1}{2} + \frac{857}{250,000}$ so that the orbit is really a tight precession of the 2-leaf clover through an angle of $\left(\frac{857}{250,000}\right) 2\pi \approx 0.02154$ radians per radial cycle.




FIG. 3: Top: The same orbits from Figure 2, projected into the orbital plane.

While they do fix the turning points, the constants $\vec{\mathcal{E}}$ do not uniquely fix the orbit [12]. An orbit that hits apastron at θ_{min} is not identical to the orbit that hits apastron at a different value of θ , as shown in Fig. 4. Since $q_{r\theta}$ depends only on constants, a $q_{r\theta} = 1/2$ orbit is always a 2-leaf clover in the orbital plane [9]. However, orbits with different r - θ initial conditions (r_0, θ_0) are rotated relative to each other in the orbital plane.

As Fig. 4 shows, the resulting orbits are genuinely distinct in 3D. Presumably, they could have distinct gravitational wave emissions. Interestingly, though, perihelion precession happens on a faster time scale than plane precession. It is therefore reasonable to suspect that all orbits with the same $q_{r\theta}$ generate similar waveforms and that the different plane precessions induce modest differences in the modulations of the amplitude [13]. We

remain agnostic on the relative importance of r - θ initial values on the waveform generated and instead focus on efficient calculation of adiabatic inspirals.

B. Resonant Orbits in Phase Space – Phase Space Tori

In the 4D space spanned by $(r, p_r, \theta, p_\theta)$, all orbits (resonant or non-resonant) lie on 2D tori that can be constructed as the Cartesian product of two closed curves. We obtain one of those closed curves if we project an orbit into the r - p_r plane. The area of the curve is the familiar action J_r used in action-angle coordinates. Analogously, the projection of the same orbit into the θ - p_θ plane yields another closed curve with area J_θ . We now consider that pair of curves as a locus of points on a 2D surface with the topology of the 2-torus $\mathbb{S}^1 \times \mathbb{S}^1 \equiv \mathbb{T}^2$. Every set of orbital parameters $\vec{\mathcal{E}}$ defines one such torus that we denote $\mathbb{T}_{\vec{\mathcal{E}}}^2$.

The use of Mino time as an evolution parameter furnishes one (but certainly not the only) coordinate system for $\mathbb{T}_{\vec{\mathcal{E}}}^2$, according to the following construction. As already mentioned, the motions $r(\lambda)$ and $\theta(\lambda)$ are each individually periodic in Mino time, with periods Λ_r and Λ_θ (and frequencies Ω_r and Ω_θ), respectively. Scaling the evolution parameter λ on each of the r - p_r and θ - p_θ curves by Ω_r and Ω_θ , respectively, leads to a natural definition of angle variables $\chi_r \equiv \Omega_r \lambda$ and $\chi_\theta \equiv \Omega_\theta \lambda$. We choose a specific trajectory $(r(\lambda), p_r(\lambda), \theta(\lambda), p_\theta(\lambda))$ in order to assign χ_r and χ_θ values, respectively, along the r - p_r and θ - p_θ curves, but the trajectory is only a device that we discard once the torus coordinate system is in place. The points at 0 and 2π in each of χ_r, χ_θ are identified, so the torus can be represented as a 2π -by- 2π square with opposite sides identified as in Fig. 5. We will make a simplifying choice that (r_a, θ_{min}) corresponds to the origin⁴ of the torus. Then, a reflection in the line $\chi_r = \pi$ corresponds to keeping r fixed and reversing the sign of p_r , and analogously for reflections in $\chi_\theta = \pi$. Note that each quadrant of the toroidal square therefore contains the same (r, θ) pairs but with all possible sign combinations for the momenta $(++, +-, -+, --)$.

Note that each χ_r corresponds to an ordered pair⁵

⁴ Many references, including [1, 3, 11, 14], instead tacitly choose the point $(r = r_p, p_r = 0, \theta = \theta_{min}, p_\theta = 0)$ as the origin of the torus coordinates. We say “tacitly” because they refer to the individual orbit with those initial conditions as a fiducial geodesic to use in their analyses. Another interpretation of that choice is that they are working not with one geodesic but with one *torus* and that they have instead chosen a fiducial origin for a χ_r - χ_θ coordinate system on that torus.

⁵ Some references describe the mapping of functions of the form $F(r, \theta)$ to corresponding functions $F(\chi_r, \chi_\theta)$. In fact, no function that enters an adiabatic EMRI calculation depends on r and θ alone. The notation $F(r, \theta)$ in those references is used because, once restricted to a torus, the value of each coordinate determines




FIG. 4: The above figures are all $q_{r\theta} = \frac{1}{2}$ orbits with $a = 0.9$ and orbital parameters $L_z = 2.65115$, $Q = 0.944969$ and $E = 0.954788$. The three figures in each row have the same initial coordinates. The column on the left shows an $r\text{-}\cos\theta$ projection of the orbit, the middle column is a projection in the orbital plane, and the right column is the 3D orbit. All three rows have $r_0 = r_a = 17.8148$. The first row has $\theta_0 = \theta_{min} = 1.22079$ and is the same orbit shown in the top panel of Figs. 2 and 3, the middle row has $\theta_0 = 1.39579$, and the bottom row has $\theta_0 = \frac{\pi}{2}$.

(r, p_r) and each χ_θ corresponds to an ordered pair (θ, p_θ) . We discuss alternative coordinate systems for \mathbb{T}_ε^2 in Appendix A and elsewhere in this article but will use only the (χ_r, χ_θ) coordinates for calculations.

On the compact (χ_r, χ_θ) square defined above, geodesic trajectories are lines of slope $\Omega_\theta/\Omega_r = 1 + q_{r\theta}$.

With respect to Mino time, those orbits are given parametrically by

$$\chi_r(\lambda) = \Omega_r \lambda + \chi_{r_0} \quad (12a)$$

$$\chi_\theta(\lambda) = \Omega_\theta \lambda + \chi_{\theta_0} \quad (12b)$$

Two different initial positions $\vec{\chi}_0$ and $\vec{\chi}'_0$ produce distinct orbits unless there exist real numbers x and y that simultaneously satisfy

$$\frac{y - \chi_{\theta_0}}{x - \chi_{r_0}} = \frac{\Omega_\theta}{\Omega_r}, \quad (13)$$

its conjugate momentum up to a sign. Still, the values of those signs affect the value of the function. We believe a notation such as $F(r, p_r, \theta, p_\theta)$ for these pre-torus phase space functions is more appropriate.




FIG. 5: The above picture shows a resonant torus mapped to a square with the path of two resonant orbits traced out. The solid line shows the path of a resonant orbit with $\chi_{r_0} = \chi_{\theta_0} = 0$ and the orbit traced out by the dotted line has $\chi_{r_0} = 0$ and $\chi_{\theta_0} = 0.7894$. The resonant torus and both resonant orbits have $\frac{\Omega_\theta}{\Omega_r} = 1 + q_{r\theta} = \frac{p}{z} = \frac{3}{2}$.

and

$$\begin{aligned} \chi'_{r_0} &= x \bmod 2\pi \\ \chi'_{\theta_0} &= y \bmod 2\pi \end{aligned} \quad (14)$$

If these conditions are met, then the two different initial positions produce time-translated versions of the same orbit.

When Ω_θ/Ω_r is irrational, we will call both the torus and any orbits on that torus **non-resonant**. Orbits on non-resonant tori never close and instead sample the entire torus ergodically: an orbit starting from any initial condition will pass arbitrarily close to every point in the torus after some finite (but possibly very long) time. Therefore, non-resonant orbits with different $\vec{\chi}_0 \equiv (\chi_{r_0}, \chi_{\theta_0})$ are arbitrarily close to time translations of every other non-resonant orbit with the same $\vec{\chi}_0$. We will alternately refer to such orbits as aperiodic or biperiodic.

When the frequency ratio Ω_θ/Ω_r is a rational number $\frac{p}{z}$, we will call both the underlying torus and orbits on that torus **resonant**. Orbits that live on resonant tori inherit the rational frequency ratio and thus always trace out closed curves. Since no single resonant orbit ergodically fills the torus, even after infinite time, two resonant orbits with the same $\vec{\mathcal{E}}$ but different $(\chi_{r_0}, \chi_{\theta_0})$ are not necessarily time translations of each other. *The set of all resonant orbits with the same $\vec{\mathcal{E}}$ does fill the entire*

torus. Because they return to their initial conditions after a finite time, we will alternately refer to these orbits as periodic or singly periodic.

Figure 5 shows two resonant orbits on the resonant torus defined by $E = 0.954788$, $L_z = 2.65115$, $Q = 0.944969$. Each is thus a $q_{r\theta} = \frac{1}{2}$ orbit, or 2-leaf clover in the orbital plane. These are the same two orbits illustrated in physical space in the top two rows of Figure 4. The two orbits are distinguished by their initial position $\vec{\chi}_0$ on the torus. The solid line orbit, which starts at $\chi_{r_0} = 0$, $\chi_{\theta_0} = 0$, corresponds to the physical space orbit with initial conditions $r_0 = r_a = 17.81477$ and $\theta_0 = \theta_{\min} = 1.220793$. The dotted line orbit with initial conditions $\chi_{r_0} = 0$ and $\chi_{\theta_0} = 0.7854$ has physical space initial conditions of $r_0 = r_a = 17.81477$ and $\theta_0 = 1.39579$. Notice that any two adjacent line segments belonging to a single orbit are separated in χ_r by $\frac{2\pi}{p}$ and in χ_θ by $\frac{2\pi}{z}$ but are not traced out sequentially for general $\frac{p}{z}$.

In the same way that the rational numbers have zero measure on the line, the set of resonant tori has zero measure in the 4D phase space. To date, most of the literature on the adiabatic EMRI problem has ignored resonant geodesics precisely for this reason. Nevertheless, as we will see, the judicious exploitation of this measure zero set leads to significant computational efficiencies in adiabatic EMRI calculations.

III. AVERAGING IN THE ADIABATIC APPROXIMATION

Given the background on geodesic dynamics, we now turn to the adiabatic approximation of EMRIs, an approximation that has seen substantial debate in the literature. As we elucidate below, most of that debate conflates the question of what kind of averaging procedure to use in the equations of motion (1) with other related but logically independent questions about the adiabatic approximation. In this section, we clarify why phase space averaging (as opposed to time averaging) is the correct averaging procedure, a conclusion first reached in Ref. [2]. We also establish the results we will need in Section IV to exploit resonant orbits for computational savings.

A. The adiabatic equations of motion

Let \vec{X} denote the Boyer-Lindquist coordinates of the inspiraling object along with its canonical radial and polar momenta. In the absence of radiation reaction, the equations of motion are

$$\frac{d\vec{X}}{dt} = \vec{G}(\vec{X}, \vec{\mathcal{E}}) \quad (15a)$$

$$\frac{d\vec{\mathcal{E}}}{dt} = 0 \quad , \quad (15b)$$

where the RHSs \vec{G} of the positional equations are some form of the equations for geodesic motion, e.g. Hamilton's equations for free-particle Kerr motion. Radiation reaction adds to the RHSs new functions

$$\frac{d\vec{X}}{dt} = \vec{G}(\vec{X}, \vec{\mathcal{E}}) + \vec{F}(\vec{X}, \vec{\mathcal{E}}) \quad (16a)$$

$$\frac{d\vec{\mathcal{E}}}{dt} = 0 + \vec{f}(\vec{X}, \vec{\mathcal{E}}) \quad (16b)$$

that are determined by the full gravitational self-force on the particle. Those unknown functions can be expanded in a perturbation series in powers of a natural small parameter: the system's mass ratio $\varepsilon \equiv \mu/M \ll 1$. Furthermore, at each order in ε , the functions above decompose into a sum of dissipative and conservative pieces:

$$\vec{F} = \varepsilon \left[\vec{F}_{\text{diss}}^{(1)} + \vec{F}_{\text{cons}}^{(1)} \right] + \varepsilon^2 \left[\vec{F}_{\text{diss}}^{(2)} + \vec{F}_{\text{cons}}^{(2)} \right] + \mathcal{O}(\varepsilon^3) \quad (17a)$$

$$\vec{f} = \varepsilon \left[\vec{f}_{\text{diss}}^{(1)} + \vec{f}_{\text{cons}}^{(1)} \right] + \varepsilon^2 \left[\vec{f}_{\text{diss}}^{(2)} + \vec{f}_{\text{cons}}^{(2)} \right] + \mathcal{O}(\varepsilon^3) \quad (17b)$$

See [7] and references therein for a fuller account.

We expect a natural separation of timescales in this system. The “fast” positional variables \vec{X} will change substantially on a short timescale equal to an orbital period $T_{\text{orb}} \sim M$, while the “slow” orbital parameters $\vec{\mathcal{E}}$ only change substantially on the much longer timescale $T_{\text{rad}} \sim M/\varepsilon$. Due to the coupling of the \vec{X} and $\vec{\mathcal{E}}$ equations, both $\vec{X}(t)$ and $\vec{\mathcal{E}}(t)$ should exhibit oscillations around a secularly trending central value, but the oscillations in $\vec{\mathcal{E}}$ should be $\mathcal{O}(\varepsilon)$. In such a system, a first-order averaging procedure seeks an approximate and hopefully more tractable set of equations for the slow variables from which the dependence on the fast variables (and thus the source of the small oscillations) is removed [15–17].

Averaging must therefore decouple the $\vec{\mathcal{E}}$ equations from the \vec{X} equations in (16) in order to isolate the secular trend in the former.⁶ One can even adopt the point of view that the desideratum of a preliminary averaging procedure is to decouple the equations for the slow and the fast variables from each other as much as possible. We represent those averaged, decoupled equations for the orbital parameters (equivalent to equation (1)) as

$$\begin{aligned} \frac{d\vec{\mathcal{E}}_{\text{secular}}}{dt} &= \left\langle \frac{d\vec{\mathcal{E}}}{dt} \right\rangle \\ &= \left\langle \vec{f}(\vec{X}, \vec{\mathcal{E}}) \right\rangle \\ &= \vec{\mathcal{F}}(\vec{\mathcal{E}}) \quad . \end{aligned} \quad (18)$$

⁶ Note that the converse is not possible, since the fast variables are coupled to the slow ones at zeroth order, where the slow variables appear as constant parameters.

Throughout this paper, we will represent averages of all sorts with angle brackets $\langle \rangle$ and use subscripts on the brackets to denote the type of average implied. Note that in (18), we have used $\langle \dots \rangle$ to denote an average without yet specifying which variables that average is to be taken over.

B. Flux balance and its relationship to averaging

Although we now have the general form of the adiabatic equations, we cannot write them explicitly because we still do not know how to evaluate the self-force. Mino showed [1, 11] that, under the assumption of non-resonance, the infinite time-averaged values⁷ of the functions \mathcal{F}_E and \mathcal{F}_{L_z} equal the sum of the infinite time-averaged fluxes of the corresponding orbital parameters at radial infinity and the central black hole horizon in GWs emitted by the system⁸. While there is no conserved Q -current to associate with a GW Q -flux, Mino likewise showed that there are analogous infinite time-averaged quantities at infinity and the horizon that sum to the infinite time-averaged value of \mathcal{F}_Q . Though not strictly physically accurate, we will henceforth refer to those quantities as fluxes of Q for ease. Subsequent work [3, 18, 19] has led to explicit formulae for these Q -fluxes.

Fortunately, we do know how to calculate the aforementioned time-averaged fluxes at infinity and the horizon via the computationally mature Teukolsky formalism. Various Teukolsky-based (TB) codes can compute the required fluxes from equatorial orbits [20], spherical orbits⁹ (constant r) [22], and now generic orbits of arbitrary inclination and eccentricity [14, 19, 23].

These developments have led to the following two-stage implementation of the adiabatic approximation. The first stage, usually called the radiative approximation¹⁰, keeps only the lowest-order contributions from the dissipative self-force (since the conservative contributions will average to zero). The second stage, called the flux-balance method, uses the time-averaged *nonlocal* fluxes (com-

⁷ The quantities of physical interest are averages over coordinate time t .

⁸ Mino's proof suggests that this equivalence is only true for non-resonant geodesics and possibly a small subset of resonant geodesics.

⁹ Many authors refer to orbits of constant radial coordinate r as “circular” even when they are nonequatorial. We prefer the term “spherical” for such orbits (as used in [21]) and reserve the term circular for constant r equatorial orbits.

¹⁰ There is some dispute about whether the neglected secular effects of the averaged conservative piece of the self-force manifest themselves at the same order in the small expansion parameter μ/M as the dissipative pieces [4–7]. That dispute does not concern us here. Whatever its limitations, the radiative approximation, is here to stay for at least the foreseeable future, if for no other reason than that it is the only relativistically correct approximation to the inspiral motion accessible to numerical calculation in the status quo.

putable) as proxies for the averaged *local* contributions of the dissipative self-force (not currently computable). The RHSs of equation (18) end up with nonlocal fluxes inside the brackets and an interpretation of those brackets as infinite time-averages.

There is a problem, however, with this prescription, which intertwines two logically distinct facets of the adiabatic approximation to EMRIs:

1. Is it mathematically appropriate to interpret the angle brackets in equation (18) as a time-average, or is some other sort of average required?
2. Given the answer to 1, can we evaluate the RHS of (18), either directly or by finding a numerically equivalent proxy?

After all, the fact that we *can* compute a time-averaged proxy does not imply that we *should* be time averaging in the first place.

The form of equation (18) suggests two ways to average the RHS in order to remove the dependence on the positional variables: for fixed $\vec{\mathcal{E}}$, we can either phase space average over the torus, or we can evaluate the RHS along a specific orbit on that torus and then average over time. In Section III E, we offer a definitive argument in favor of torus averaging instead of time averaging.

To arrive at that conclusion, we must first distinguish between torus functions and time functions. Torus functions assign a value to every point on a phase space torus, while time functions assign a value to points along an individual orbit that are labeled by the value of an evolution parameter (i.e. a time variable). Our conclusions about adiabatic averaging will be based on differences in how Fourier analysis is done on these two domains — a 2-dimensional compact position space for the torus-functions and a 1-dimensional noncompact time axis for the time-functions. Moreover, numerically accurate flux calculations require frequency-domain TB codes that separately compute fluxes from individual Fourier modes, and the aforementioned different domains also impact the details of the modewise flux calculation.

Before delving into those details, we must mention an important point. Average values, whether in the torus or time sense, are coordinate dependent, and in certain applications it matters which coordinates the average is taken over. The angle brackets in equation (18), for instance, will turn out to denote a torus average not over $\vec{\chi}$ but over a different set of torus coordinates $\vec{\gamma} \equiv (\gamma_r, \gamma_\theta)$ described in Appendix A. However, torus averages with respect to $\vec{\chi}$ are much easier to compute than those over $\vec{\gamma}$, in much the same way as Mino time averages are easier to compute than are coordinate time averages. Luckily, for every torus function $U(\vec{\gamma})$ and every time function $u(t)$, we can always construct different functions $V(\vec{\chi})$ and $v(\lambda)$ such that

$$\langle U \rangle_{\vec{\gamma}} = \langle V \rangle_{\vec{\chi}} \quad (19)$$

$$\langle u \rangle_t = \langle v \rangle_\lambda \quad (20)$$

The relationship between U and V (or between u and v) is highlighted in Appendix A 1. Sections III C and III D present the necessary Fourier analysis details.

We will always avail ourselves of the simplification in equation (19). Accordingly, throughout the rest of the paper, we focus exclusively on torus averages over $\vec{\chi}$ and time averages over Mino time λ with the understanding that they may merely be computation-friendly proxies for averages of different but related functions over different torus or time coordinates.

C. Torus averaging and Fourier analysis of torus-functions

We will call a **torus function** $f(\vec{\chi}; \vec{\mathcal{E}})$ any rule that assigns a complex number to every point on a phase space torus. Note that $\vec{\mathcal{E}}$ specifies both the torus function and the phase-space torus that serves as its domain. Usually, we will be discussing properties of torus functions evaluated at some definite value of $\vec{\mathcal{E}}$. We thus omit the explicit dependence on the orbital parameters $\vec{\mathcal{E}}$ for brevity, except where it might lead to confusion.

We assume that every such torus-function is continuous and differentiable in all its arguments (including $\vec{\mathcal{E}}$). We also require it to be single-valued on the torus, which implies 2π periodicity in each of the angle variables:

$$f(\chi_r, \chi_\theta) = f(\chi_r + 2\pi, \chi_\theta) = f(\chi_r, \chi_\theta + 2\pi) \quad (21)$$

Like any function that is independently periodic in two independent variables, a torus-function has a double Fourier series representation

$$f(\vec{\chi}) = \sum_{k,n} A_{kn} e^{-in\chi_r} e^{-ik\chi_\theta} \quad (22)$$

with the A_{kn} 's given as usual by

$$A_{kn} = \frac{1}{(2\pi)^2} \int_0^{2\pi} d\chi_r \int_0^{2\pi} d\chi_\theta f(\chi_r, \chi_\theta) e^{+in\chi_r} e^{+ik\chi_\theta} \quad (23)$$

In order to distinguish them from another set of double-indexed quantities we introduce later, we will refer to the A_{kn} 's as **spatial Fourier coefficients** or **torus Fourier coefficients**.

We now define, in the usual way, the following useful quantities. The **torus average** of f is

$$\begin{aligned} \langle f(\vec{\chi}) \rangle_{\vec{\chi}} &\equiv \frac{1}{(2\pi)^2} \int_0^{2\pi} d\chi_r \int_0^{2\pi} d\chi_\theta f(\chi_r, \chi_\theta) \\ &= A_{00} \end{aligned} \quad (24)$$

The **torus averaged Fourier power** of f is

$$\mathcal{P}_{\vec{\chi}} \equiv \frac{1}{(2\pi)^2} \int_0^{2\pi} d\chi_r \int_0^{2\pi} d\chi_\theta |f(\chi_r, \chi_\theta)|^2 \quad (25)$$

By Parseval's theorem, the torus-averaged power must also equal

$$\mathcal{P}_{\vec{\chi}} = \sum_{k,n} |A_{kn}|^2 \quad . \quad (26)$$

The **2D power spectrum of f** is the contribution to the torus-averaged Fourier power from each pair of spatial frequencies or wavenumbers $(\kappa_r, \kappa_\theta)$. Note that because the period in each of the χ_r and χ_θ directions is 2π , the corresponding fundamental spatial frequencies are $\kappa_r = \kappa_\theta = 1$, so we see power only at integer lattice points (k, n) in the 2D wavenumber space.

All statements above are standard results from the Fourier analysis of functions on a compact 2D spatial domain. They apply equally well on resonant and non-resonant tori.

D. Time averaging and Fourier analysis of time-functions

We can evaluate any torus function along a curve (12) on its associated torus that corresponds to an orbit. Since each orbit is specified by its initial position $\vec{\chi}_0$ on the torus, each torus function naturally induces a 2-parameter family of time functions, one for each $(\chi_{r0}, \chi_{\theta0})$ pair. Time functions, then, are grouped into 5-parameter families – 3 parameters $\vec{\mathcal{E}}$ to specify a torus, and 2 parameters $\vec{\chi}_0$ to specify an orbit on that torus.

We denote one member of such a family as $f(\vec{\chi}(\lambda); \vec{\chi}_0; \vec{\mathcal{E}})$. We will sometimes omit the explicit $\vec{\chi}_0$ and $\vec{\mathcal{E}}$ -dependence of a time function and simply write $f(\lambda)$, again except where clarity would suffer. Throughout this paper, we adopt the notational convention that a time function and the torus function from which it is derived are denoted by the same symbol (f , in the examples so far).

For time functions, non-resonant and resonant tori must be treated separately.

1. Non-resonant tori

When Ω_θ/Ω_r is irrational, every k, n pair leads to a distinct frequency

$$\Omega_{kn} \equiv n\Omega_r + k\Omega_\theta, \quad k, n \in \mathbb{Z} \quad . \quad (27)$$

Such a biperiodic time-function is not periodic: it is bounded on $(-\infty, \infty)$, but there is no finite time interval over which the function exactly repeats itself. Still, every biperiodic function has a unique discrete Fourier representation [24]

$$f(\lambda; \vec{\chi}_0) = \sum_{k,n} A_{kn;\lambda} e^{-i(n\Omega_r + k\Omega_\theta)\lambda} \quad . \quad (28)$$

Note that the harmonics are *not* equally spaced in frequency. The temporal Fourier coefficients $A_{nk;\lambda}$ (which we have named suggestively) are given by the limit

$$A_{kn;\lambda} = \lim_{\Lambda \rightarrow \infty} \frac{1}{\Lambda} \int_{b-\Lambda/2}^{b+\Lambda/2} d\lambda f(\lambda; \vec{\chi}_0) e^{i(n\Omega_r + k\Omega_\theta)\lambda} \quad , \quad (29)$$

which exists and is independent of b [24] (henceforth, we set $b = 0$ for convenience). Equivalently, we could say that the Fourier transform of $f(\lambda; \vec{\chi}_0)$ consists of a series

$$\tilde{f}(\Omega) = \sum_{k,n} A_{kn;\lambda} \delta(\Omega - (n\Omega_r + k\Omega_\theta)) \quad (30)$$

of delta-function impulses unequally spaced in frequency.

Paralleling the Fourier discussion of torus functions, we now define the time-averaged value, time-averaged Fourier power, and the 1D power spectrum of a time function associated with a non-resonant orbit. Biperiodic functions offer no single period over which to time-average in a natural way. Given the existence¹¹ of expressions like (29), averaging over all time seems like a sensible choice. The theory of almost-periodic functions states that such an infinite time-average indeed exists [25], so we define

$$\langle f(\lambda; \vec{\chi}_0) \rangle_\lambda = \lim_{\Lambda \rightarrow \infty} \frac{1}{\Lambda} \int_{-\infty}^{\infty} d\lambda f(\lambda; \vec{\chi}_0) \quad . \quad (31)$$

We will refer to $\langle f \rangle_\lambda$ simply as *the* time-average of f rather than as the infinite time-average value, as it is sometimes called. Comparing equations (29) and (31), the time-average equals

$$\langle f(\lambda; \vec{\chi}_0) \rangle_\lambda \equiv A_{00;\lambda} \quad . \quad (32)$$

We define the time-averaged Fourier power of $f(\lambda)$ as a special case of (31) by

$$\mathcal{P}_\lambda \equiv \lim_{\Lambda \rightarrow \infty} \frac{1}{\Lambda} \int_{-\frac{\Lambda}{2}}^{\frac{\Lambda}{2}} d\lambda |f(\lambda; \vec{\chi}_0)|^2 \quad . \quad (33)$$

Parseval's theorem also applies to biperiodic time-functions [25], so

$$\mathcal{P}_\lambda = \sum_{k,n} |A_{kn;\lambda}|^2 \quad . \quad (34)$$

¹¹ Technically, the logical presentation of Fourier coefficients and time-averages for biperiodic (or more general multiperiodic) functions goes in the reverse order. First, the existence of the infinite time-average in (29) is established for a biperiodic function $f(\lambda)$. The existence of the Fourier coefficients in (29) then follows from the existence of the average value and the fact that $f(\lambda)e^{i(n\Omega_r + k\Omega_\theta)\lambda}$ is also biperiodic. We have chosen this order to parallel the presentations in Sections III C and III D 2.

That fact allows us to define a 1D power spectrum for $f(\lambda)$ as the contribution to the time-averaged power from each temporal frequency Ω . The graph of $|A_{kn;\lambda}|^2$ over the 1D Ω -space would show power only at the discrete and unequally spaced set of frequencies (27).

The question now is how to evaluate these time averages in practice. Though equation (29) defines the $A_{kn;\lambda}$'s, such integrals over infinite intervals divided by infinite quantities do not lend themselves to simple evaluation, either analytically or numerically¹².

To compute the temporal Fourier coefficients of $f(\lambda; \vec{\chi}_0)$, we must instead proceed circuitously. Consider the spatial Fourier representation (22) of the torus-function $f(\vec{\chi})$ evaluated along the orbit (12), which yields

$$f(\lambda; \chi_{r0}, \chi_{\theta0}) = \sum_{k,n} A_{kn} e^{-i(n\chi_{r0} + k\chi_{\theta0})} e^{-i(n\Omega_r + k\Omega_\theta)\lambda} \quad (35)$$

By uniqueness¹³ of the Fourier representation of $f(\lambda; \vec{\chi}_0)$, and comparing equations (28) and (35), we conclude that the temporal Fourier coefficients of $f(\lambda; \vec{\chi}_0)$ and the spatial Fourier coefficients of $f(\vec{\chi})$ are related¹⁴ by

$$A_{kn;\lambda} \equiv A_{kn} e^{-i(n\chi_{r0} + k\chi_{\theta0})} \quad (36)$$

We note that each temporal coefficient differs from the corresponding spatial coefficient in (23) only by a complex phase determined by the initial conditions $\vec{\chi}_0$ of the orbit. Consequently, their magnitudes are identical, regardless of the initial position of the orbit:

$$|A_{kn;\lambda}| = |A_{kn}|, \quad \forall \vec{\chi}_0 \in \mathbb{T}_{EL_zQ}^2 \quad (37)$$

This fact is consistent with the ergodic property of these orbits. Every orbit eventually comes arbitrarily close to every point on the torus, so shifting initial conditions leads to a new orbit that is arbitrarily close to a time

translation of the original orbit. And, of course, time translations only affect the complex phase of temporal Fourier coefficients.

If we know the torus function $f(\vec{\chi})$, its spatial Fourier coefficients A_{kn} can be computed by any number of efficient numerical routines, without any of the difficulties that beset computation of the $A_{kn;\lambda}$'s via direct evaluation of the definition (28). This fact, combined with equation (36), leads to the only practical recipe for computing the $A_{kn;\lambda}$'s of the orbit with initial position $\vec{\chi}_0$, namely to compute instead the A_{kn} 's and then use equation (36). Ref. [26] introduced just such a technique in the specific context of functions of Kerr geodesics.

All the other quantities mentioned in this section are likewise determined from their torus function counterparts. From equations (32) and (36), $A_{00;\lambda} = A_{00}$. We thus conclude that on a non-resonant torus, the time average of $f(\lambda)$ equals the torus average of its associated torus-function $f(\vec{\chi})$. Moreover, since this is true for every time function on that torus, the time average of such a function is independent of the initial condition $\vec{\chi}_0$:

$$\langle f(\lambda; \vec{\chi}_0) \rangle_\lambda = \langle f(\vec{\chi}) \rangle_{\vec{\chi}}, \quad \forall \vec{\chi}_0 \in \mathbb{T}_{EL_zQ}^2 \quad (38)$$

Likewise, equations (37) and (26) imply that, on a non-resonant torus, the time-averaged Fourier power of $f(\lambda; \vec{\chi}_0)$ equals the torus-averaged Fourier power of $f(\vec{\chi})$ for every $\vec{\chi}_0$:

$$\mathcal{P}_\lambda \equiv \mathcal{P}_{\vec{\chi}} \quad \forall \vec{\chi}_0 \in \mathbb{T}_{EL_zQ}^2 \quad (39)$$

Equations (26), (37) and (39) together imply equation (34)¹⁵. By extension, the 1D power spectrum of $f(\lambda)$ can be derived from the 2D power spectrum of $f(\vec{\chi})$ by mapping wavenumber pairs to frequencies using eq. (27).

All of the above relationships between torus-function quantities and those of any biperiodic time function induced via (35) are well-established and well-known in the literature on almost-periodic functions [24, 25] and on integrable Hamiltonian systems [15]. Many of these facts, however, are used but not so clearly delineated in this way in the literature relating to EMRI calculations. We have gone through the trouble of including them here not only for completeness and clarity but also to emphasize that we can only execute the above recipes if we know the corresponding torus function $f(\vec{\chi})$.

This leads us to a crucial observation. If all we know is $f(\lambda)$, either as some closed-form expression in terms of λ or as a numerical time-series, there is no practical scheme for computing its temporal Fourier coefficients $A_{kn;\lambda}$, even though those coefficients are perfectly well-defined. In addition to the initial conditions $\vec{\chi}_0$ associated with $f(\lambda)$, we *must* also know the torus-function $f(\vec{\chi})$ (or at least its spatial Fourier coefficients A_{kn}) in

¹² To evaluate equation (29) numerically, larger and larger values of Λ would be required before converging to some accuracy. This is computationally impractical because such a process will in general converge extremely slowly. Thus, as the size of the integration interval grows, so will the required number of evaluations of the integrand, a particularly problematic development if the integrand is expensive to calculate. Moreover, the prefactor of $1/\Lambda$ can eventually become so small that there is loss of significance in the final answer, thus compromising accuracy.

¹³ The set of complex exponential functions $e^{-i\Omega\lambda}$ for all Ω are a basis for absolutely integrable functions on the space $\lambda \in (-\infty, \infty)$. $f(\lambda; \vec{\chi}_0)$ inherits absolute integrability from the associated torus function $f(\vec{\chi})$, which has a spatial double Fourier series representation and thus is absolutely integrable by assumption. Equation (28) is therefore a projection onto the complex exponential basis, and projections onto basis sets are unique.

¹⁴ We denoted the kn th temporal Fourier coefficient by $A_{kn;\lambda}$ in anticipation of its close relationship to the kn th spatial Fourier coefficient A_{kn} of the associated torus function and added the λ subscript to remind us of when we are dealing with spatial vs. temporal Fourier coefficients.

¹⁵ This proves Parseval's theorem for biperiodic functions.

order to compute the $A_{kn;\lambda}$'s. We summarize the implications of this fact for flux balancing in section III E.

All of the equivalences noted between torus-function quantities and their time-function counterparts followed from the assumption of non-resonance. On resonant tori, all of these equivalences break down, as we now show.

2. Resonant tori

Unlike time functions evaluated along non-resonant orbits, time functions on resonant orbits are singly periodic, with (possibly very long) period Λ_P and corresponding fundamental frequency $\Omega_P = 2\pi/\Lambda_P$. The single periodicity of time functions of resonant orbits means that all frequency-domain quantities have straightforward and familiar definitions.

Any time function evaluated on a resonant orbit has a Fourier series representation

$$f(\lambda; \vec{\chi}_0) = \sum_j C_{j;\lambda} e^{-ij\Omega_P \lambda} \quad (40)$$

whose coefficients¹⁶ are single-index objects

$$C_{j;\lambda} = \frac{1}{\Lambda_P} \int_{-\Lambda_P/2}^{\Lambda_P/2} d\lambda f(\lambda; \vec{\chi}_0) e^{+ij\Omega_P \lambda} \quad (41)$$

Like the $A_{kn;\lambda}$'s, each $C_{j;\lambda}$ varies with the initial condition $\vec{\chi}_0$. Unlike the $A_{kn;\lambda}$'s, the Fourier transform of $f(\lambda; \vec{\chi}_0)$ is a sequence

$$\tilde{f}(\Omega) = \sum_j C_{j;\lambda} \delta(\Omega - j\Omega_P) \quad (42)$$

of equally spaced delta-function impulses in frequency space.

In the resonant case, we define the time average of $f(\lambda)$ straightforwardly as

$$\langle f(\lambda; \vec{\chi}_0) \rangle_\lambda \equiv \frac{1}{\Lambda_P} \int_{-\Lambda_P/2}^{\Lambda_P/2} d\lambda f(\lambda; \vec{\chi}_0) = C_{0;\lambda} \quad (43)$$

Likewise, we define the time-averaged power as

$$\mathcal{P}_\lambda(\vec{\chi}_0) \equiv \frac{1}{\Lambda_P} \int_{-\Lambda_P/2}^{\Lambda_P/2} d\lambda |f(\lambda; \vec{\chi}_0)|^2 \quad (44)$$

By Parseval's theorem, the time-averaged power is also given by

$$\mathcal{P}_\lambda(\vec{\chi}_0) = \sum_j |C_{j;\lambda}|^2 \quad (45)$$

To flush out how the $C_{j;\lambda}$'s relate to the spatial A_{kn} 's and to the initial condition $\vec{\chi}_0$, we begin, as in the non-resonant case, by inducing a time function (35) from a torus function. In the resonant case, the frequency ratio Ω_θ/Ω_r is a rational number p/z , where p and z are relatively prime and $p > z$. (In terms of integers in the definition $q_{r\theta} = w + \frac{v}{z}$, $p = (w+1)z + v$.) The individual r and θ frequencies and periods are therefore related to the fundamental frequency and total period of the periodic orbit by

$$\Omega_r = z\Omega_P \quad (46a)$$

$$\Omega_\theta = p\Omega_P \quad (46b)$$

and

$$\Lambda_r = \frac{\Lambda_P}{z} \quad (47a)$$

$$\Lambda_\theta = \frac{\Lambda_P}{p} \quad (47b)$$

As a result, all kn combinations for which

$$nz + kp = j \quad (48)$$

lead to identical frequencies

$$n\Omega_r + k\Omega_\theta = nz\Omega_P + kp\Omega_P = j\Omega_P \quad (49)$$

in the arguments of the exponential functions on the RHS of equation (35).

The selection rule (48) maps every kn pair to some j . By the uniqueness of Fourier representations, we conclude from equations (35) and (40) that

$$C_{j;\lambda}(\vec{\chi}_0) = \sum_{\substack{k,n: \\ nz+kp=j}} A_{kn} e^{-i(n\chi_{r0} + k\chi_{\theta0})} \quad (50)$$

Note that equation (50) is really only a summation over a single index since the value of k in any term is determined by the value of n and the (fixed) value of j .

It is tempting to rewrite each term on the RHS of equation (50) as $A_{kn;\lambda}$, mimicking the notation for the non-resonant temporal Fourier coefficients. We refrain from doing so because we seek a clear distinction between spatial and temporal Fourier coefficients, and temporal double-index coefficients are not defined in the resonant case [25, 27]. Fourier representations are unique, so the familiar single-index representation (40) is the only such projection of $f(\lambda)$ onto a set of mutually orthogonal basis functions. If we were to write an expression like (28) on a resonant orbit, the different harmonics on the RHS would not all be orthogonal, and we would not have a bona fide Fourier expansion in hand until we collapsed all terms corresponding to the same frequency into a single term.

Equation (36) implied that, on non-resonant orbits, several quantities one can compute for a time-function $f(\lambda; \vec{\chi}_0)$ turn out to be independent of $\vec{\chi}_0$: the magnitudes of its Fourier coefficients, its time-averaged value,

¹⁶ Periodicity of $f(\lambda)$ implies that the integral in equation (41) has the same value taken over any interval of length Λ_P . We choose the symmetric interval $[-\Lambda_P/2, \Lambda_P/2]$ solely for aesthetic reasons.

its time-averaged Fourier power, and its power spectrum. In contrast, equation (50) implies that, on resonant orbits, each of those quantities *does* depend on the initial condition $\vec{\chi}_0$. Each $C_{j;\lambda}$ is a sum of spatial A_{kn} 's with $\vec{\chi}_0$ -dependent phases rather than just one such term (cf. equation (36)), so both the magnitudes

$$|C_{j;\lambda}|(\vec{\chi}_0) = \left| \sum_{\substack{k,n: \\ nz+kp=j}} A_{kn} e^{-i(n\chi_{r_0}+k\chi_{\theta_0})} \right| \quad (51)$$

and time-averaged value

$$\begin{aligned} \langle f(\lambda; \vec{\chi}_0) \rangle_\lambda &= C_{0;\lambda}(\vec{\chi}_0) \\ &= \sum_{\substack{k,n: \\ nz+kp=0}} A_{kn} e^{-i(n\chi_{r_0}+k\chi_{\theta_0})} \\ &= A_{00} + \sum_{\substack{n \neq 0, k \neq 0: \\ nz+kp=0}} A_{kn} e^{-i(n\chi_{r_0}+k\chi_{\theta_0})} \end{aligned} \quad (52)$$

$$\begin{aligned} |C_j(\chi_{r_0}, \chi_{\theta_0})|^2 &= \sum_{\substack{kn: \\ j=nz+kp}} A_{kn} e^{-i(n\chi_{r_0}+k\chi_{\theta_0})} \sum_{\substack{k'n': \\ j=n'z+k'p}} A_{k'n'}^* e^{i(n'\chi_{r_0}+k'\chi_{\theta_0})} \\ &= \sum_{\substack{k=k', n=n': \\ j=nz+kp}} |A_{kn}|^2 + \sum_{\substack{k \neq k', n \neq n': \\ j=nz+kp, \\ j=n'z+k'p}} A_{kn} A_{k'n'}^* e^{-i(n-n')\chi_{r_0}} e^{-i(k-k')\chi_{\theta_0}}. \end{aligned} \quad (53)$$

also depend on $\vec{\chi}_0$ through cross terms, and the time-averaged power and power spectra inherit this dependence via (45). Figure 6 illustrates this point for the test function $r \cos \theta$.

Via its $\vec{\chi}_0$ -dependence, equation (52) defines a torus function in the variables $\chi_{r_0}, \chi_{\theta_0}$. Complex exponentials have a zero average value, so averaging that torus function over all $\chi_{r_0}, \chi_{\theta_0}$ kills every term in the summation on the RHS of (52), leaving only A_{00} . But A_{00} is the torus averaged value of the associated torus function $f(\vec{\chi})$. We conclude that the torus-average over all initial conditions of the time average of a time function equals the torus-average of the underlying torus-function. An identical argument applies if we torus average the squared-magnitudes (53) of the coefficients over all $\vec{\chi}_0$ and, by extension, if we likewise torus-average the time-averaged power (44).

The upshot is that the parallels between torus functions and time functions obtained in the non-resonant case break down in the resonant case:

$$\langle f(\lambda; \vec{\chi}_0) \rangle_\lambda \neq \langle f(\chi_r, \chi_\theta) \rangle_{\vec{\chi}} \quad (54)$$

$$\mathcal{P}_\lambda(\vec{\chi}_0) \neq \mathcal{P}_{\vec{\chi}}. \quad (55)$$

retain $\vec{\chi}_0$ -dependence. The squared magnitudes

However, torus averages over initial conditions and torus averages of time averaged *are* equal for both a function f and its Fourier power:

$$\langle \langle f(\lambda; \vec{\chi}_0) \rangle_\lambda \rangle_{\vec{\chi}} = \langle f(\chi_r, \chi_\theta) \rangle_{\vec{\chi}} \quad (56)$$

$$\langle \mathcal{P}_\lambda(\vec{\chi}_0) \rangle_{\vec{\chi}} = \mathcal{P}_{\vec{\chi}}. \quad (57)$$

To clarify, equations (54) and (55) state that the time average of a time function and the torus average of its associated torus-function are not *identically equal* as they are in the non-resonant case. That does not, of course, preclude the possibility that the two could be *circumstantially equal* for some particular choice of initial condition $\vec{\chi}_0$. In fact, for real-valued functions f , the mean-value theorem guarantees that $\langle f(\lambda; \vec{\chi}_0) \rangle_\lambda = \langle f(\vec{\chi}) \rangle_{\vec{\chi}}$ for at least one $\vec{\chi}_0^{\text{mvt}} \in \mathbb{T}_{\vec{\chi}}^2$. In general, f will be complex-valued, and we have no such guarantee. The time-averaged Fourier power, however, is strictly real, so there is at least one $\vec{\chi}_0^{\text{mvt}}$ such that $\mathcal{P}_\lambda(\vec{\chi}_0^{\text{mvt}}) = \mathcal{P}_{\vec{\chi}}$. We explore some implications of this fact for adiabatic EMRI calculations in Section V.

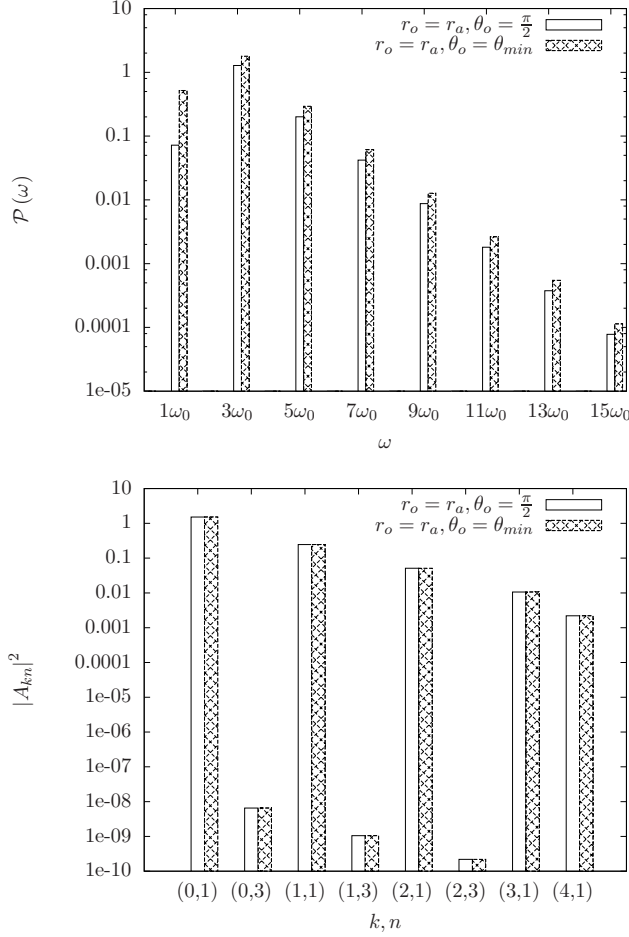


FIG. 6: Top: The Fourier power spectrum of the function $r \cos \theta$ for a $q_{r,\theta} = \frac{1}{2}$ periodic orbit for two different sets of initial coordinates, $(r_0 = r_a, \theta_0 = \frac{\pi}{2})$ and $(r_0 = r_a, \theta_0 = \theta_{min})$, but the same sets of orbital parameters, $a = 0.9$, $E = 0.954788$, $L_z = 2.65115$ and $Q = 0.944969$. Bottom: The magnitudes of some spatial Fourier coefficients for the same orbits in the top panel.

E. Why torus averaging is correct

To summarize, time averaging is equivalent to torus averaging for non-resonant orbits. Furthermore, torus averaging is the only practical recipe for computing Fourier coefficients and so torus averaging is the explicit computation instituted in practice.

However, time averaging is *inequivalent* to torus averaging for resonant orbits. Thus torus averaging wins out for two reasons. First, torus averaging along a resonant torus crucially washes away any $\vec{\chi}_0$ positional dependence, while time averaging does not. The $\vec{\chi}_0$ -dependence violates the spirit of averaging, namely to remove all dependence on the fast variables. Second, and more seriously, though the time-averaged equations are continuous in $\vec{\chi}_0$ for fixed $\vec{\mathcal{E}}$, they are in general discon-

tinuous in $\vec{\mathcal{E}}$ for fixed $\vec{\chi}_0$. The situation will resemble that of Thomae's modified Dirichlet function

$$D_M(x) = \begin{cases} 0 & \text{if } x \text{ is irrational} \\ \frac{1}{z} & \text{if } x = p/z, \text{ with } p \text{ and } z \text{ coprime} \\ 1 & \text{if } x = 0 \end{cases}, \quad (58)$$

which is continuous at the irrationals, discontinuous on the rationals, and nowhere differentiable¹⁷ [29]. Pragmatically speaking, even if a set of ODEs with such pathologically discontinuous and non-differentiable equations had a solution, it is unclear how one would numerically integrate them. Furthermore, the continuity furnished by torus-averaged fluxes is absolutely essential for the proper construction of a grid through which adiabatic trajectories are to be interpolated, as discussed in the introduction.

In short, torus-averaged equations are well-behaved, while time-averaged equations lose the continuity and differentiability that guarantee the resulting equations are well-posed and have unique solutions, the very basis of every standard numerical integration scheme.

The arguments made in favor of torus averaging apply to the radiative approximation, based on an average of the dissipative piece of the *local* self-force on the inspiraling particle. But when it comes to flux-balance as a specific implementation of the radiative approximation, this now leaves a logical gap. As acknowledged in [1, 11], the flux-balance arguments that allow the nonlocal fluxes of conserved quantities to be used as proxies for the local dissipative self-force have been derived on a time-averaged basis and under the assumption of non-resonance. Since time and torus averages agree for non-resonant orbits/tori, the time-averaged nonlocal fluxes are still good proxies for the torus-averaged local dissipative self-force in the non-resonant case.

What has never been made explicit is whether flux-balance is also valid in the resonant case, on either a time averaged or torus-averaged basis. We resolve this issue now: time averaged flux-balancing may not be true on resonant orbits in general, but that will be irrelevant since it will be true on a torus-averaged basis. Mino showed that, under the assumption of non-resonance, the time-averaged fluxes of $\vec{\mathcal{E}}$ at infinity and the horizon furnish proxies for the time averaged RHSs of equations (18). But then, by the arguments we have heavily exploited, the corresponding torus-averaged versions must

¹⁷ Such a function certainly seems unphysical. It violates the hypotheses of continuity and differentiability in all arguments required by the theorems bounding the error in a solution to time-averaged equations with almost-periodic dependence on time [17]. More pathologically, it violates the hypotheses for the well-posedness of an initial value problem and for the existence of solutions to systems of ODEs [28]. Such a function would not be Riemann integrable, and it is unclear whether a system of such functions would even be Lebesgue integrable.

also be equal. And since torus averages are insensitive to the resonance or non-resonance of the underlying torus *and are continuous in $\vec{\mathcal{E}}$* , the flux-balance prescription is valid in a torus-averaged sense for all orbits. If the torus-averaged fluxes were not good proxies for the torus-averaged local equations only at resonances, a discrete set of measure zero, then they could not be continuous in $\vec{\mathcal{E}}$. But we have shown torus averages are everywhere continuous. *De facto*, then, Mino's argument establishes the validity of torus-averaged flux-balancing generally.

It would thus seem that both flux-balancing as a general procedure and its specific implementation in a frequency-domain application of the Teukolsky formalism treat non-resonant and resonant tori equally, as stated in [7]. Flux-balancing is, in fact, thusly impartial, but interestingly, the Teukolsky formalism is not. As we will show below, a TB torus-averaged flux calculation can achieve computational savings of an order of magnitude or more on low-order resonant tori that are simply not available on non-resonant tori. These efficiencies follow from a simple observation about the Fourier integrals a TB code must evaluate and are independent of the specific implementation in code of the Teukolsky formalism. Thus, rather than disfavoring resonances, as has been commonly assumed, the Teukolsky formalism actually shows favoritism for resonances, and properly leveraged, that favoritism can substantially accelerate adiabatic inspiral calculations.

IV. COMPUTATIONAL SAVINGS ALONG RESONANCES

We now explain computational efficiencies that exploit resonant tori. Although specific to the Teukolsky formalism, the computational expedience can be understood without all details of that formalism. We simply assert some features and formulae from a TB flux calculation that we require to make our argument. For reference, Appendix B gives a somewhat more detailed overview of the Teukolsky formalism and offers at least skeletal derivations of the formulae listed below. For a fuller treatment of the Teukolsky formalism, which is beyond the scope of this work, we direct the reader to the references listed in Appendix B.

As a reminder, our argument is specific to *frequency-domain* Teukolsky calculations and corresponding codes. In the context of the EMRI problem, such codes compute a combined multipole and Fourier decomposition of the metric perturbations at infinity and the black hole horizon due to a geodesic source.

A. The fluxes of E, L_z, Q

The fluxes of the conserved quantities $\vec{\mathcal{E}}$ are usually reported as quantities averaged over *coordinate* time t

on non-resonant orbits¹⁸. By the arguments of Section IIID 1 and Appendix A, this is equivalent to the torus-averages of the fluxes on all tori over the torus coordinates $\vec{\gamma}$. We therefore report those same expressions here as torus-averaged fluxes. For E and L_z , those expressions are [3, 14]

$$\left\langle \frac{dE}{dt} \right\rangle_{\vec{\gamma}}^{H/\infty} = \sum_{lmkn} \frac{\alpha_{lmkn}^{\infty/H}}{4\pi\omega_{mkn}^2} \left| Z_{lmkn}^{\infty/H} \right|^2 \quad (59a)$$

$$\left\langle \frac{dL_z}{dt} \right\rangle_{\vec{\gamma}}^{H/\infty} = \sum_{lmkn} \frac{\alpha_{lmkn}^{\infty/H} m}{4\pi\omega_{mkn}^3} \left| Z_{lmkn}^{\infty/H} \right|^2 \quad (59b)$$

Based on Mino's argument in [11], Refs. [18, 30] worked out the corresponding expression for the time-averaged Q flux for non-resonant orbits, which we also report as the torus-averaged flux

$$\begin{aligned} \left\langle \frac{dQ}{dt} \right\rangle_{\vec{\gamma}}^{H/\infty} &= -2 \langle a^2 E \cos^2 \theta \rangle_{\vec{\gamma}} \left\langle \frac{dE}{dt} \right\rangle_{\vec{\gamma}}^{H/\infty} \\ &+ 2 \langle \cot^2 \theta \rangle_{\vec{\gamma}} \left\langle \frac{dL_z}{dt} \right\rangle_{\vec{\gamma}}^{H/\infty} \\ &- \sum_{lmkn} \frac{k\omega_{\theta}}{2\pi\omega_{mkn}^3} \alpha_{lmkn}^{\infty/H} \left| Z_{lmkn}^{\infty/H} \right|^2 \quad (59c) \end{aligned}$$

The prefactors in the first two lines of (59c) are computed only once for the entire torus. Thus, substituting equations (59a) and (59b) into the RHS of (59c) and combining like terms with those in the summation of the last line, the flux for Q has the same general form as the fluxes of E and L_z have. Our savings arguments will be based on that form, so although we will speak about E and L_z for concreteness, those arguments will apply to Q as well. Appendices B and C summarize the derivations of these expressions.

Before proceeding with those arguments, we clarify the notation in equations (59). First, the apparent discrepancy between the ordering of the H/∞ superscripts on the left- and righthand sides of the equations is not a typographical error. On the LHS, the superscript denotes fluxes at the black hole horizon and radial infinity, respectively. The somewhat backward notational choice to have the fluxes at infinity depend on a quantity labeled Z_{lmkn}^H and the horizon fluxes on Z_{lmkn}^{∞} is, at this point, ingrained in the literature. To maintain a modicum of notational uniformity, we have labeled the weighting factors $\alpha^{H/\infty}$ with the same backward superscript convention. The exact form of those weighting factors will not concern us. What matters for our purposes is that every factor α_{lmkn}^H for the fluxes at infinity is equal to 1 and that every factor α_{lmkn}^{∞} for the fluxes at the horizon is

¹⁸ See Appendix C for the time-averaged fluxes from resonant orbits.

real and depends on k, n only through ω_{kn} . All the arguments to follow apply equally to fluxes at infinity and at the horizon. We borrow the notation \star from Ref. [14] to denote either of H/∞ .

Continuing, the indices l, m are standard multipole indices¹⁹, with $l \geq 2, -l \leq m \leq l$. Our argument will focus on the Fourier analysis of each l, m term individually, so that, unless explicitly stated otherwise, l, m are taken to be fixed everywhere in this section, while k, n each run from $-\infty$ to ∞ . The frequencies

$$\omega_{mkn} \equiv m\omega_\varphi + \omega_{kn} = m\omega_\varphi + n\omega_r + k\omega_\theta \quad (60)$$

are the combined harmonics ω_{kn} of the r and θ fundamental frequencies (the coordinate time version of equation (27)) and the fundamental azimuthal frequency ω_φ . Note that the integer m is both a multipole index and the relative contribution of ω_φ to each frequency ω_{mkn} . Other than attaching itself as a label to frequencies in this way, however, m will not appear as a Fourier index in any sense below.

Finally, Appendix C explains why we have written the fluxes as average values over the $\tilde{\gamma}$ torus coordinates mentioned in Section III B and in Appendix A. We note here simply that if we seek adiabatic solutions in the form $\tilde{\mathcal{E}}(t)$ (as opposed to $\tilde{\mathcal{E}}(\lambda)$), then the angle brackets in (18) should also be averages over $\tilde{\gamma}$, so that (59) have the correct form to be proxies for $\tilde{\mathcal{F}}(\tilde{\mathcal{E}})$. The representations of the LHSs of the flux equations as averages over $\tilde{\gamma}$ is otherwise irrelevant, since in light of (19), we will always seek equivalent and easier to compute $\tilde{\chi}$ -averaged quantities.

B. Z_{lmkn}^\star as Fourier coefficients of a torus function

With these preliminaries out of the way, we are ready to list the features of the RHSs of (59) that we will need for our savings arguments both in this section and in Section V. For our principal argument, what matters is that for fixed l, m values, each Z_{lmkn}^\star takes the form of a Fourier coefficient of some torus function,

$$Z_{lmkn}^\star = \int_0^{2\pi} d\chi_r \int_0^{2\pi} d\chi_\theta e^{i(n\chi_r + k\chi_\theta)} f_{lm;\omega=\omega_{mkn}}^\star(\chi_r, \chi_\theta) \quad (61)$$

This form of the Z_{lmkn}^\star 's associated with a geodesic source of arbitrary eccentricity and inclination is detailed in several references (see, for instance, [3, 14, 18, 19, 26, 30]) and summarized in Appendix B.

Equation (61) parallels the form of equation (23) from Section III C, but there is one critical difference. For fixed l, m , the function $f_{lm;\omega}^\star(\chi_r, \chi_\theta)$ further depends on a continuous parameter ω that must be set to ω_{mkn} when evaluating Z_{lmkn}^\star for a given multipole mode. Postponing for the moment any details of the function $f_{lm;\omega}^\star(\chi_r, \chi_\theta)$ or its derivation, we remark that this dependence on the coordinate time harmonic frequencies of the source as an external parameter persists despite the fact that equation (61) is a spatial Fourier integral.

Thus, for fixed l, m , the Z_{lmkn}^\star are not the Fourier coefficients of a single function but rather isolated Fourier coefficients of several different functions²⁰. On a non-resonant torus, every k, n pair leads to a different value of ω_{mkn} , and every coefficient computed has a distinct function $f_{lm;\omega}^\star(\chi_r, \chi_\theta)$ in the integrand. On a resonant torus with associated frequency ratio $\omega_\theta/\omega_r = \Omega_\theta/\Omega_r = p/z$, all k, n pairs that satisfy the selection rule (48) for the same j lead to identical values of Ω_{mkn} , and some coefficients with different values of k, n will share the same integrand function $f_{lm;\omega}^\star(\chi_r, \chi_\theta)$. The practical implications of this asymmetry for a TB flux calculation constitute the basis of our savings argument.

In anticipation of later arguments, we also note that for each fixed value of the pair l, m , the resulting doubly infinite sum over k, n in (59) has the appearance of a torus-averaged Fourier power in the sense of Section III C with the identification

$$\begin{aligned} |A_{kn}|^2 &= \text{prefactor} \times |Z_{lmkn}^\star|^2 \\ A_{kn} &= \sqrt{\text{prefactor}} \times Z_{lmkn}^\star \end{aligned} \quad (62)$$

The prefactors in front of $|Z_{lmkn}^\star|^2$ turn out to be real-valued and non-negative for all values of the indices, so it is valid to subsume them into some new coefficients A_{kn} .

C. Recycling computations between Fourier modes

The complex-valued quantities Z_{lmkn}^\star are the backbone of a frequency-domain Teukolsky calculation, and a code that implements such a calculation spends by far the lion's share of its CPU budget on computing them. To explain how resonances can be leveraged to optimize that budget, we must look a bit more closely at the integrand functions $f_{lm;\omega}^\star(\chi_r, \chi_\theta)$.

The main ingredients in $f_{lm;\omega}^\star(\chi_r, \chi_\theta)$ are two separate functions that have the same sort of ω dependence described above: a radial Teukolsky function $R_{lm;\omega}^\star(r)$ and a spin-weighted spheroidal harmonic $_{-2}S_{lm}^{aw}(\theta)$. We imagine re-expressing the former as a torus function of χ_r alone and the latter as a torus function of χ_θ alone

¹⁹ The values $l = 0, 1$ are not relevant in GW calculations, for which the lowest non-vanishing moment is the $l = 2$ quadrupole.

²⁰ This is part of the reason why we cannot compute all the Z_{lmkn}^\star coefficients for a given l, m at once with, for instance, a 2-dimensional Fast Fourier Transform (FFT).

but will continue to write them as functions of r and θ , as they are in the rest of the literature. $f_{lm;\omega}^*(\chi_r, \chi_\theta)$ consists of a somewhat messy assortment of terms and factors involving these two functions, several of their derivatives, the coordinates and velocities of the particle (both of these are absorbed into the torus coordinates χ_r, χ_θ), and other elementary functions.

Each of $R_{lm;\omega}^*$ and ${}_{-2}S_{lm}^{a\omega}(\theta)$ satisfies an ODE that depends on l, m and ω in a nontrivial and partly implicit way (see Appendix B). No simple closed-form solutions to these equations exist that make the functional dependence of the solutions on those parameters explicit. As a result, for every distinct set of values (l, m, ω) , those ODEs must be solved from scratch to obtain the numerical representations of $R_{lm;\omega}^*$ and ${}_{-2}S_{lm}^{a\omega}$ needed to evaluate the integrand. Particularly in the case of $R_{lm;\omega}^*$, this operation is computationally costly.

Schematically, then, one calculates each Z_{lmkn}^* for fixed l, m via the following steps:

1. Determine the frequency $\omega = \omega_{mkn}$
2. Obtain a representation of ${}_{-2}S_{lm}^{a\omega}$
3. Obtain a representation of $R_{lm;\omega}^*$ for $\star = H, \infty$ (this step requires first determining an eigenvalue of the ${}_{-2}S_{lm}^{a\omega}$ ODE)
4. Evaluate $f_{lm;\omega}^*(\vec{\chi}_i)$ at whatever abscissae $\vec{\chi}_i$ are required by the specific numerical integration algorithm chosen
5. Compute whatever weights w_i the integration algorithm may require for those function values and tabulate the integral (61) as $\sum_i w_i f_{lm;\omega}^*(\vec{\chi}_i)$.

On a non-resonant torus, each k, n pair produces a different answer to step 1 and requires the execution from scratch of all the remaining steps as well. On a resonant torus, in contrast, the k, n pairs can be grouped by a common value of j in the selection rule (48). Steps 1–3 need only be performed once for an entire j -group. Depending on the integration algorithm selected, steps 4 and 5 may also only need to be performed once or a small number of times per j -group, with a total number of reusable function evaluations set by the Z_{lmkn}^* in the group requiring the greatest number of sample points to attain some target accuracy. We will make the reasonable assumption that steps 1 and 5, even if done several times per j -group, are a small fraction of the total cost of evaluating all the coefficients in that group, and we will take the cost of steps 2–4 as an estimate of the total cost of computing any single coefficient.

Consider now evaluating all the Z_{lmkn}^* on a low-order resonant torus with $\omega_\theta/\omega_r = \Omega_\theta/\Omega_r = p/z = 1 + q_{r\theta}$ and on a neighboring non-resonant torus with nearly identical orbital parameters. By the continuity of the Z_{lmkn}^* with respect to $\vec{\mathcal{E}}$, coefficient values will be nearly identical on those two tori. The integer values of n_{\max}, k_{\max} determined should also be identical or nearly identical on the

two tori (we assume for simplicity that they are identical). Let \mathcal{N}_2 and \mathcal{N}_1 denote, respectively, the number of separate times steps 2–4 above must be executed on the non-resonant torus and resonant torus. To make a more apples to apples comparison, one can instead let \mathcal{N}_2 represent the total number of distinct executions of steps 2–4 on the resonant torus if the resonance of that torus is not acknowledged from the outset. Roughly speaking, generating all the Z_{lmkn}^* with $|n| \leq n_{\max}, |k| \leq k_{\max}$ on the non-resonant torus will require $\mathcal{N}_2/\mathcal{N}_1$ times more computation than it will on the neighboring resonant torus. Symmetries in the underlying equations imply that the value of $Z_{l(-m)(-k)(-n)}^*$ is uniquely determined by the value of Z_{lmkn}^* . Thus, in practice, one of the indices n and k can be restricted to run over only nonnegative values, and the value of $\mathcal{N}_2/\mathcal{N}_1$ should take that fact into account.

Figure 7 estimates the savings factor $\mathcal{N}_2/\mathcal{N}_1$ for resonant tori with various values of $q_{r\theta}$ and for several representative hypothetical values of n_{\max} and k_{\max} consistent with the reported performance of the TB code for arbitrary eccentricities and inclinations described in Ref. [14]²¹. For simplicity, we have taken $n_{\max} = k_{\max}$.

We can see the following trends in the histograms. First, for a given k_{\max}, n_{\max} , if we fix the value of p and increase z or vice versa, the savings factor drops. Thus, the savings factor is largest when both p and z are as low as possible. The greatest savings (over an order of magnitude) accrue when $z = 1$. Second, the larger the values of k_{\max}, n_{\max} , i.e. the more slowly converging the expressions for the fluxes, the greater the savings factor for a given p/z . Generally speaking, the most slowly converging fluxes are for orbits with moderate to high eccentricities [14, 19], which typically have higher associated values of $q_{r\theta}$ since they are closer to the separatrix between plunging and non-plunging motion [9]. Thus, for instance, a rough approximation of the true savings factor in the top two panels of Figure 7 would be given by a roughly horizontal or slightly downward sloping line connecting the histogram bar with the lowest $q_{r\theta}$ at the lowest n_{\max} with the highest $q_{r\theta}$ at the highest n_{\max} . A good rough predictor for the expected savings would thus be the z value of a torus, yielding a savings factor of ~ 30 for $z = 1$ and ~ 7 for $z = 6$. The lowest savings factor on that graph of ~ 3 (corresponding to the not-so-low-order resonance with $p/z = 25/6$) is nothing to sneeze at, and more typically the savings factor from acknowledging resonance would appear to be around an order of magnitude on average.

While a detailed audit of comparative cost would have

²¹ We estimate n_{\max} and k_{\max} based on the code in [14] rather than the similar code in [23] only because the truncation rules used in [14] are more amenable to direct cost comparison with our proposal. Both codes seem to need to compute a total number of modes of similar order of magnitude to achieve high flux accuracy, and both are apt to profit from our proposal.

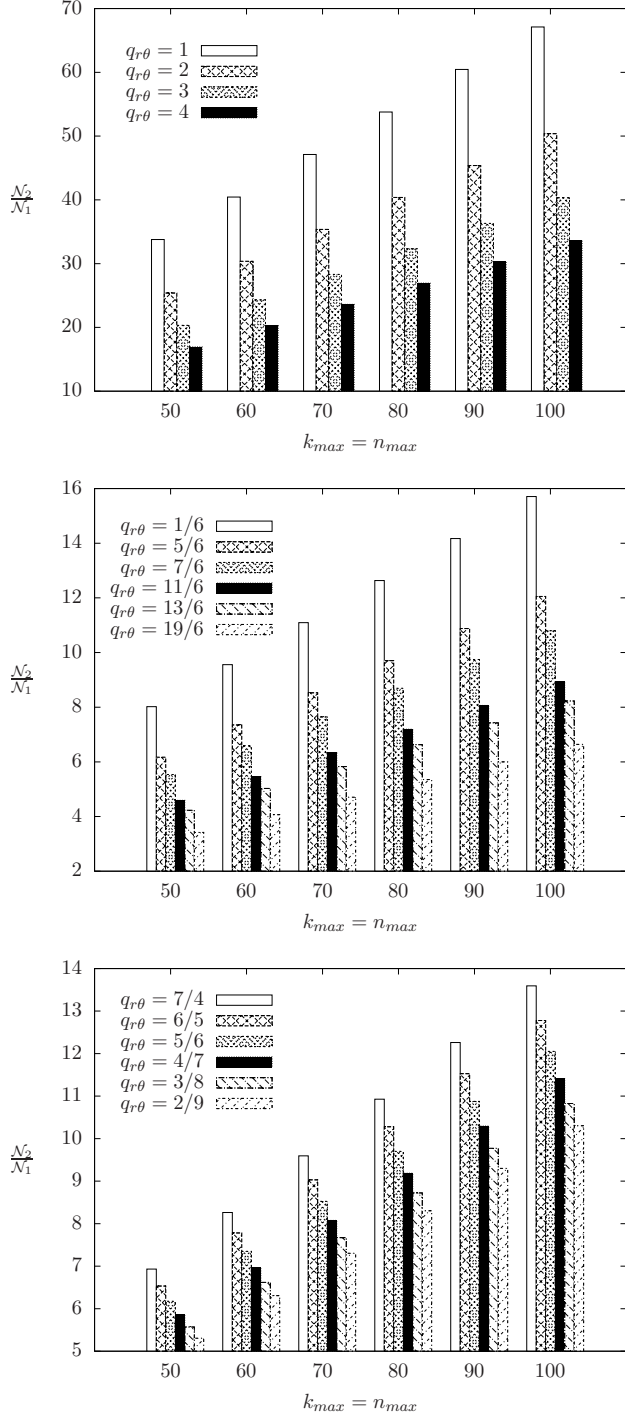


FIG. 7: The three histograms show the average number of Z_{lmnk}^* coefficients that pertain to a single frequency on a resonant torus, a number that corresponds to the savings factor $\mathcal{N}_2/\mathcal{N}_1$. We show the savings factor for a variety of $q_{r\theta}$ geodesics and a variety of n_{\max} and k_{\max} . Top: $q_{r\theta} = \text{Integers}$. Middle: $q_{r\theta}$ is a variety of values all with the same denominator, $z = 6$. Bottom: $q_{r\theta}$ is a variety of non-integer values all with $p = 11$ but different z .

to be done on a code-specific basis, the potential payoff of these observations makes a case for testing our proposal in existing codes.

D. Numerical EMRI grids

Even if we stipulate that fluxes can be computed more efficiently on low-order resonant tori than on non-resonant tori, is this fact necessarily useful? After all, points in $\vec{\mathcal{E}}$ -space corresponding to resonant tori, let alone low-order ones, are already a measure zero set, so to construct the inspiral curve $\vec{\mathcal{E}}(t)$, wouldn't the RHSs of the adiabatic ODEs have to be evaluated in general (and, formally, infinitely more often) on non-resonant tori than on resonant ones? Interestingly, while the answer to that question is “yes”, the actual calculation of TB fluxes itself need only ever be done on resonant tori, and at least predominantly (and possibly exclusively) on low-order ones, at least for the foreseeable future.

The reason has to do with the *absolute* computational cost of those fluxes, even on resonant tori. Simply inserting a TB frequency-domain flux routine into the RHS of, say, a standard Runge-Kutta ODE solver to generate inspiral curves in real-time is untenable, even with a large number of processors at one's disposal to parallelize the TB calculation. Instead, solution of the adiabatic ODEs will proceed as follows. For each value of the black hole parameters, one would build a numerical grid of flux values on some dense mesh of points in $\vec{\mathcal{E}}$ -space and then interpolate off of that grid to obtain the fluxes for arbitrary values of $\vec{\mathcal{E}}$. Once handed such a grid, those interpolated flux values would go into a standard ODE solver which could presumably generate inspiral curves very efficiently. The main expense to consider, then, is the construction of the grid.

Our proposal is that such a grid should be built using exclusively resonant grid points. More specifically, we propose a hierarchical population of such a grid, beginning with the low-order resonant points and then increasing the order of the resonance (or just increasing z , if our loose conjecture about horizontal lines in the top panels of Figure 7 proves to be correct) until some requisite grid density is obtained to minimize interpolation error. Those grid density requirements may force the evaluation of fluxes on some higher-order resonances, but no resonant grid point (whether low- or high-order) will ever be more expensive to populate with Fourier flux data than a nearby non-resonant grid point will be. The worst-case scenario near certain locations in the space would be to break even by using a resonant versus a non-resonant grid point. Our hierarchical approach would seem to at least lower if not minimize the total computational cost of such a grid.

We remark on two features of our proposal. First, the savings factor discussed depends only on the decision to use resonant tori for TB calculations and no other implementation-specific features of that calcula-

tion. Thus, those savings will multiply any additional savings that may stem from other algorithmic improvements in any such implementation or from the availability of more or faster processors to perform the TB calculations.

Second, its efficacy has nothing to do with interesting physical effects that may occur in the neighborhood of resonant orbital parameter values during a real inspiral [2]. The flux-balance method, and in fact the adiabatic approximation in general, may fail to capture these effects. Any such failure is immaterial to our argument, which rests not on *physical* properties of resonant tori but rather *mathematical* ones they have in specific relation to frequency-domain TB calculations. In other words, despite the fact that the adiabatic approximation might be least faithful to reality in and around resonances, leveraging resonances is nonetheless the most efficient means of attaining an adiabatic approximation for those regimes where it is likely to be faithful.

E. Gravitational waveform snapshots

As already argued, we are free to interpret the coefficients Z_{lmkn}^* either as spatial Fourier coefficients of a torus function of $\vec{\chi}$ or of a different torus function of $\vec{\gamma}$. As shown in Appendix B, the t -function versions of the Z_{lmkn}^* coefficients are used to build the Weyl scalar ψ_4 at radial infinity, from which the two polarizations of the waveform h are constructed. These waveform “snapshots” from geodesic sources [14] are useful for exploring how a known orbital motion impacts GW signals, and though they will quickly go out of phase with a true inspiral signal, they are still likely to play a pivotal role in hierarchical searches for GWs from EMRIs.

More specifically, with the $\vec{\gamma}$ -coefficients Z_{lmkn}^* in hand, h can be reconstructed from the associated t -coefficients. By analogy to equation (36), we get

$$Z_{lmkn;t}^* \equiv Z_{lmkn}^* e^{-i(n\gamma_{r_0} + k\gamma_{\theta_0})} \quad , \quad (63)$$

in the non-resonant case, and by analogy to equation (50) get

$$Z_{lmj;t}^*(\vec{\gamma}_0) = \sum_{\substack{k,n: \\ nz+kp=j}} Z_{lmkn}^* e^{-i(n\gamma_{r_0} + k\gamma_{\theta_0})} \quad (64)$$

in the resonant case if we know the initial conditions. Since the waveforms (or, rather, their Fourier representations) depend on the Z_{lmkn}^* coefficients, then like the fluxes, they will also probably need to be interpolated from a grid that stores the Z_{lmkn}^* values themselves instead of or in addition to the fluxes. The same arguments made above for the fluxes thus cross-apply to waveform snapshots.

V. SPECULATIONS ON FURTHER SAVINGS

In this section, we sketch a speculative but tantalizing possibility for further efficiencies in adiabatic EMRI grid construction beyond those discussed in Section IV. The idea centers around calculating time-averaged rather than torus-averaged fluxes on resonant tori. At first glance, that suggestion seems to fly in the face of earlier arguments that the RHSs of the adiabatic equations should be torus-averaged fluxes and that torus averages and time averages are not identical on resonant tori. The apparent incongruity disappears, however, in light of two facts:

1. On any resonant torus, the mean value theorem guarantees that torus-averaged fluxes equal time-averaged fluxes on certain special orbits.
2. For low-order resonances, those time-averaged fluxes are more accurate and cheaper to compute.

The additional savings are beyond the cost benefit of incorporating the proposal of Section IV.

We substantiate these claims below in turn. We caution the reader that, in contrast to the savings of Section IV, those discussed in this section may prove more elusive in practice because determining the special orbits mentioned in step 1 above could prove so difficult as not to be net-beneficial. We discuss such limitations and suggest fruitful avenues of numerical investigation to help further reduce the cost of generating adiabatic inspirals.

A. Using time-averages to compute torus-averages

The time-averaged fluxes from a single resonant orbit do not appear elsewhere in the literature. As we explain in Appendix C, the arguments of Section IIID 2 imply that those fluxes are (note, these are single-index objects in j)

$$\left\langle \frac{dE}{dt} \right\rangle_t^* = \sum_{lmj} \frac{\alpha_{lmj}^*}{4\pi\omega_{mj}^2} |Z_{lmj;\lambda}^*|^2 \quad (65)$$

$$\left\langle \frac{dL_z}{dt} \right\rangle_t^* = \sum_{lmj} \frac{\alpha_{lmj}^* m}{4\pi\omega_{mj}^3} |Z_{lmj;\lambda}^*|^2 \quad . \quad (66)$$

It remains to be shown whether the following would translate to $\left\langle \frac{dQ}{dt} \right\rangle_t$. We restrict attention in this section to E and L_z fluxes for the sake of exposition.

As before, we assume fixed l, m in everything below. In the fluxes, the frequencies

$$\omega_{mj} \equiv m\omega_\varphi + j\omega_P \quad , \quad (67)$$

the real-valued weight factors α_{lmj}^* , and the temporal Fourier coefficients

$$Z_{lmj;\lambda}^* = \frac{1}{\Lambda_P} \int_0^{\Lambda_P} d\lambda e^{ij\Omega_P \lambda} f_{lmj;\lambda}^*(\vec{\chi}(\lambda; \vec{\chi}_0)) \quad (68)$$

all become single-index quantities by the arguments of Section IIID 2. We recall from equation (62) that the torus-averaged fluxes have the form of a torus-averaged power of some unspecified torus-function. Likewise, for fixed values of l, m , the time-averaged fluxes (65) and (66) have the appearance of a time-averaged Fourier power in the sense of Section IIID 2 with the identification

$$\begin{aligned} |C_{j;\lambda}|^2 &= \text{prefactor} \times |Z_{lmj;\lambda}^*|^2 \\ C_{j;\lambda} &= \sqrt{\text{prefactor}} \times Z_{lmj;\lambda}^* \end{aligned} \quad (69)$$

Each time-averaged flux, like any time-averaged Fourier power, is real-valued. Therefore, by equations (69) and (57) and the mean-value argument made at the end of Section IIID 2, there exist initial positions $\bar{\chi}_{0;E}^{\text{mvt},\star}, \bar{\chi}_{0;L_z}^{\text{mvt},\star}$ on the torus such that

$$\left\langle \frac{dE}{dt} \right\rangle_t^* (\bar{\chi}_{0;E}^{\text{mvt},\star}) = \left\langle \frac{dE}{dt} \right\rangle_{\bar{\gamma}}^* \quad (70)$$

$$\left\langle \frac{dL_z}{dt} \right\rangle_t^* (\bar{\chi}_{0;L_z}^{\text{mvt},\star}) = \left\langle \frac{dL_z}{dt} \right\rangle_{\bar{\gamma}}^* \quad (71)$$

Actually, there must be at least two continuous 1-parameter families of special initial values $\bar{\chi}_{0;E}^{\text{mvt},\star}$ (one for each of $\star = H, \infty$) and two such families for $\bar{\chi}_{0;L_z}^{\text{mvt},\star}$: any two initial conditions that lie on the same orbit simply time-translate that orbit, and time-translation does not change time-averaged function values or time-averaged powers.

None of the values $\bar{\chi}_{0;E}^{\text{mvt},\star}$ and $\bar{\chi}_{0;L_z}^{\text{mvt},\star}$ need agree. Thus, if we sought to determine the torus-averaged fluxes indirectly by instead evaluating time-averaged fluxes, we might need to evaluate each coefficient $Z_{lmj;\lambda}^*$ as many as four²² times, once for each of the initial conditions $\bar{\chi}_{0;E}^{\text{mvt},\star}$ and $\bar{\chi}_{0;L_z}^{\text{mvt},\star}$.

We can, however, also apply the mean-value argument individually to each real-valued $|Z_{lmj;\lambda}^*|^2$. In this case, we would obtain a sequence of special initial conditions $\bar{\chi}_{0;j}^{\text{mvt},\star}$ that cause each $|Z_{lmj;\lambda}^*|^2$ to attain its torus-averaged value over all possible initial conditions. The different $\bar{\chi}_{0;j}^{\text{mvt},\star}$ would not necessarily agree for different values of j . Since the prefactors in (69) are independent of initial position, each $\bar{\chi}_{0;j}^{\text{mvt},\star}$ would simultaneously set the j th term in the power spectrum of every flux to its torus-averaged value. Evaluating the time-averaged fluxes for any of the individual initial conditions $\bar{\chi}_{0;j}^{\text{mvt},\star}$ would not produce a torus-averaged flux. However, since the average of a sum of terms must equal the sum of the individual averages of those terms, the sum of the resulting $|C_{j;\lambda}|^2$ (each evaluated at a possibly different

$\bar{\chi}_{0;j}^{\text{mvt},\star}$) would yield the torus-averaged value of all fluxes simultaneously. Recalling that each integrand in (68) is different anyway, there is no further waste in evaluating each one using a different initial condition $\bar{\chi}_{0;j}^{\text{mvt},\star}$.

It is important to note that we have simply made an existence argument for $\bar{\chi}_{0;E}^{\text{mvt},\star}, \bar{\chi}_{0;L_z}^{\text{mvt},\star}$ and every $\bar{\chi}_{0;j}^{\text{mvt},\star}$. What those values actually are would vary from problem to problem, and finding them for the Teukolsky problem may not be practical. The integrands in (68) are not especially analytically transparent, so it may be that they can only be determined by evaluating those integrands for several initial conditions $\bar{\chi}_0$, which would defeat the purpose of invoking the mean-value theorem in the first place. Still, we believe the potential added savings from knowing the $\bar{\chi}_{0;j}^{\text{mvt},\star}$ merits exploring whether the Teukolsky calculation harbors some structure or symmetries that would allow those initial conditions to be determined with little or no added expense. We turn to those additional potential savings now.

B. Relative cost of time-averaged vs. torus-averaged functions on low-order resonant tori

Assume that we have in hand the $\bar{\chi}_{0;j}^{\text{mvt},\star}$ for each j and agree to evaluate the coefficients $Z_{lmj;\lambda}^*$ using those special initial conditions. The added efficiency is twofold: each $C_{j;\lambda}$ should potentially be less expensive to compute than any given A_{kn} (by reducing a double integral to a single integral), and fewer such $C_{j;\lambda}$'s than A_{kn} 's will have to be computed in order to achieve a given target accuracy in the torus-averaged fluxes (by reducing a double sum to a single sum). In fact, the more efficient calculation might even increase the resulting flux accuracy. We justify those claims in turn below.

1. Cost of a coefficient

For ease of illustration, we will estimate the relative computational costs of a single $C_{j;\lambda}$ and of any single A_{kn} for which k, n satisfy the selection rule (48). To make the comparison more stark, we remap the integral (68) to the interval $[0, 2\pi]$ via a linear change of variable

$$\chi_P \equiv \Omega_P \lambda \quad (72)$$

to obtain

$$Z_{lmj;\lambda}^* = \frac{1}{2\pi} \int_0^{2\pi} d\chi_P e^{ij\chi_P} f_{lmj;\lambda}^*(\bar{\chi}(\chi_P; \bar{\chi}_0)) \quad (73)$$

The relative cost of the single integral (73) and its double-index counterpart (61) will depend on the specific numerical integration algorithms used to evaluate them and are difficult to estimate. However, we can sketch a crude argument that the single integral should be more cost efficient by considering the Fast Fourier transform (FFT) as the algorithm.

²² Or, rather, six times, since each $\bar{\chi}_{0;Q}^{\text{mvt},\star}$ would likely also be different from $\bar{\chi}_{0;E}^{\text{mvt},\star}$ and $\bar{\chi}_{0;L_z}^{\text{mvt},\star}$.

Consider first the 1D integral (73), which we can interpret as the j th Fourier coefficient of a periodic function on $[0, 2\pi]$. For a periodic function, an FFT will return all the Fourier coefficients from C_{-N_1} through C_{N_1} by sampling the integrand at $2N_1 + 1$ equally spaced abscissae²³. So to capture C_j , we would need $2|j| + 1$ evaluations of the integrand. However, the highest index coefficients computed via an FFT are heavily afflicted by aliasing error, while the lowest index coefficients computed are relatively free of such error. To minimize aliasing effects, we imagine increasing the number of sample points (and thus of coefficients computed) by some integer safety factor²⁴ \mathcal{S} so that C_j will be one of lowest index coefficients returned by the FFT and thus fairly free of aliasing error. The total number of integrand evaluations under this scheme for computing C_j would thus be $\mathcal{S}(2|j| + 1)$.

Now imagine evaluating the double integral (61) using a 2D FFT, which we (even more crudely) envision simply as nested 1D FFTs. Assuming the same safety factor \mathcal{S} throughout, we would need $\mathcal{S}(2|n| + 1)\mathcal{S}(2|k| + 1) = \mathcal{S}^2(2|n| + 1)(2|k| + 1)$ function evaluations. Re-expressing j in terms of n and k via the selection rule and using the number of integrand evaluations as a metric of numerical expense, the ratio of the cost of A_{kn} to the cost of C_j would be

$$\frac{\text{cost of } A_{kn}}{\text{cost of } C_j} = \mathcal{S} \frac{(2|n| + 1)(2|k| + 1)}{2|nz + kp| + 1} \quad (74)$$

Generally speaking, for small values of both p and z , the denominator in the cost ratio is smaller than the numerator since n and k more often than not have opposite signs for a given j . It is conceivable that a single A_{kn} could turn out less costly to evaluate than C_j , but the likelihood of that would become higher as both p and z became large, for which case a resonant torus would be barely distinguishable from a non-resonant torus in terms of all the aspects discussed in this paper.

The argument above artificially increases the true cost of evaluating both integrals and is not intended even to be fully convincing, let alone a proof. Rather, it is a heuristic illustration of a rule of thumb in numerical integration that, with similarly behaved integrands, 1D integrals are less costly to compute than 2D integrals.

2. Number of coefficients

In contrast to the relative cost of computing a coefficient, we can say more definitively that the total number of single-index coefficients needed to achieve some specified accuracy in the torus-averaged fluxes will be less

than the number of double-index coefficients needed to obtain the same accuracy.

Suppose achieving a certain flux accuracy for a given l, m pair requires computing all A_{kn} with indices up to n_{\max} and k_{\max} . Denote the total number of torus coefficients computed by $\mathcal{N}_{\bar{\chi}}$. The Z^* coefficients satisfy $|Z_{lm\omega}^*|^2 = |Z_{l(-m)(-\omega)}^*|^2$ [32], so one of n and k need only run over non-negative values to obtain all the coefficients with $|n| \leq n_{\max}, |k| \leq k_{\max}$. The total number of coefficients actually computed is therefore (having k run only non-negative)

$$\mathcal{N}_{\bar{\chi}} = (2n_{\max} + 1)(k_{\max} + 1) - n_{\max} \quad (75)$$

For comparison, we determine the number \mathcal{N}_{λ} of C_j coefficients (evaluated at $\bar{\chi}_{0,j}^{\text{mvt},*}$) that we would have to calculate so that, in light of the arguments of subsection V A, every $|A_{kn}|^2$ above would automatically be included in the sum of all $|C_j|^2$. As we showed in subsection V A, the maximum j index that needs to be included in the single-index series that will thusly catch every k, n pair is

$$j_{\max} = zn_{\max} + pk_{\max} \quad (76)$$

The symmetry of Z^* implies that j need not run both positive and negative, and the number of C_j 's we would need to calculate to ensure at least the same level of flux convergence as that attained with the A_{kn} coefficients is

$$\mathcal{N}_{\lambda} = zn_{\max} + pk_{\max} + 1 \quad (77)$$

Comparing equations (75) and (77), we see that we need a factor of

$$\begin{aligned} \mathcal{N}_{\text{savings}} &= \frac{\mathcal{N}_{\bar{\chi}}}{\mathcal{N}_{\lambda}} \\ &= \frac{(2k_{\max} + 1)(n_{\max} + 1) - k_{\max}}{zn_{\max} + pk_{\max} + 1} \end{aligned} \quad (78)$$

fewer coefficients. The reduction in the number of coefficients therefore depends on the order of the periodic orbit as well as on n_{\max} and k_{\max} . The lower the values of p and z , the greater the reduction factor.

Figure 7 showed the average number of kn modes on a resonant torus per distinct frequency. It also gives a general sense of how $\mathcal{N}_{\text{savings}}$ varies with k_{\max} and n_{\max} . The agreement between the two is not exact because, when computing all j coefficients up to the maximum j_{\max} , some additional frequencies will be included that do not correspond to any of the included kn frequencies with $|k| \leq k_{\max}, |n| \leq n_{\max}$. Therefore, Figure 7 overestimates $\mathcal{N}_{\text{savings}}$ but only slightly and gives a better estimate for the larger values of k_{\max}, n_{\max} . For example, for $q_{r\theta} = \frac{1}{6}$, equation (78) gives $\mathcal{N}_{\text{savings}} \approx 7.84$ for $n_{\max} = k_{\max} = 50$ and $\mathcal{N}_{\text{savings}} \approx 12.45$ for $n_{\max} = k_{\max} = 80$, both of which agree with the values in the histogram of Figure 7 within a few percent. On the basis of Figure 7, we can therefore conclude that focusing on temporal rather than

²³ To make the formulae that follow more intelligible, we are separately counting the value at 2π , even though it is the same as the value at 0

²⁴ Ref. [31] recommends a factor of at least 4 for most applications.

spatial Fourier coefficients and invoking the above mean-value arguments could reduce by a factor of an order of magnitude or so the total number of Z^* coefficients required to obtain accurate torus-averaged fluxes.

VI. CONCLUSION

Computation of adiabatic inspirals with a grid of resonant orbits could be an order of magnitude more efficient than the same computation with a non-resonant grid. If our speculations are verified and double sums can be collapsed to single sums (and double integrals to single integrals), there may be substantial additional savings since fewer and simpler coefficients will be required. To date, no accurate adiabatic EMRIs have been computed. Such a dramatic boost in speed would bring EMRIs more within computational reach.

Acknowledgments

We gratefully acknowledge valuable discussions with Scott Hughes. This work was supported by an NSF grant AST-0908365. JL gratefully acknowledges support of a KITP Scholarship, under Grant no. NSF PHY05-51164.

Appendix A: Torus coordinates and time coordinates

1. Mino time vs. coordinate time Fourier coefficients

Suppose $A_{kn;t}$ is a Fourier coefficient of some biperiodic coordinate time function $f(t)$,

$$A_{kn;t} = \lim_{T \rightarrow \infty} \frac{1}{T} \int_{-\frac{T}{2}}^{\frac{T}{2}} dt e^{i(k\omega_\theta + n\omega_r)t} f(t) \quad . \quad (\text{A1})$$

$$\begin{aligned} \lim_{T \rightarrow \infty} \frac{1}{T} \int_{-\frac{T}{2}}^{\frac{T}{2}} dt e^{i(k\omega_\theta + n\omega_r)t} f(t) &= \lim_{\Gamma\Lambda \rightarrow \infty} \frac{1}{\Gamma\Lambda} \int_{-\frac{\Gamma\Lambda}{2}}^{\frac{\Gamma\Lambda}{2}} d\lambda e^{i(k\omega_\theta + n\omega_r)(\Gamma\lambda + \Delta t(\lambda))} \frac{dt}{d\lambda}(\lambda) f(t(\lambda)) \\ &= \lim_{\Lambda \rightarrow \infty} \frac{1}{\Lambda} \int_{-\frac{\Lambda}{2}}^{\frac{\Lambda}{2}} d\lambda e^{i(k\Omega_\theta + n\Omega_r)\lambda} e^{i(k\omega_\theta + n\omega_r)\Delta t(\lambda)} \frac{dt}{d\lambda}(\lambda) f(t(\lambda)) \quad . \end{aligned} \quad (\text{A5})$$

In the second line above, we have absorbed²⁵ the Γ into the Λ and used the fact that the coordinate-time and

Then $A_{kn;t}$ is also the kn th Fourier coefficient of a different biperiodic Mino time function $g(\lambda)$,

$$\begin{aligned} A_{kn;t} &= \lim_{\Lambda \rightarrow \infty} \frac{1}{\Lambda} \int_{-\frac{\Lambda}{2}}^{\frac{\Lambda}{2}} d\lambda e^{i(k\Omega_\theta + n\Omega_r)\lambda} g(\lambda) \\ &\equiv A_{kn;\lambda} \quad . \end{aligned} \quad (\text{A2})$$

We prove equation (A2) by constructing the function g from f explicitly. We will need the fact (see Ref. [26] for details) that $dt/d\lambda$ depends on r and θ and is thus biperiodic when evaluated on a trajectory $r(\lambda), \theta(\lambda)$. $dt/d\lambda$ also has a nonzero average value Γ defined by

$$\Gamma \equiv \lim_{\Lambda \rightarrow \infty} \frac{1}{\Lambda} \int_{-\Lambda/2}^{\Lambda/2} d\lambda \frac{dt}{d\lambda} \quad . \quad (\text{A3})$$

Consequently, the function $t(\lambda)$ takes the form

$$t(\lambda) = \Gamma\lambda + \Delta t(\lambda) \quad (\text{A4})$$

where $\Delta t(\lambda)$ is biperiodic in λ and has zero average value.

We can now construct $g(\lambda)$. Since $\Delta t(\lambda)$ is biperiodic, it is also bounded. Thus, if we define $T \equiv t(\Lambda)$, then in the limit $T \rightarrow \infty$, we get $T \rightarrow \Gamma\Lambda$. We can now change variables in the integral in (A1):

Mino-time frequencies will be related [26] by

$$\omega_r = \frac{1}{\Gamma} \Omega_r \quad (\text{A6a})$$

$$\omega_\theta = \frac{1}{\Gamma} \Omega_\theta \quad , \quad (\text{A6b})$$

to convert to Mino frequencies in the argument of the exponential.

Comparing (A2) to (A5), we see that due to the dependence on k and n in the argument of the exponential

²⁵ All we need is for the denominator of the prefactor and the size of the integration interval to agree. Since there is no preferred size for that interval (the function f is not periodic) and we are taking the infinite limit, we are free to call the size of that interval $\Gamma\Lambda$ or Λ .

(the coordinate time frequencies ω_θ and ω_r appear here simply as parameters), there is actually a different function

$$g_{kn}(\lambda) = e^{i(k\omega_\theta + n\omega_r)\Delta t(\lambda)} \frac{dt}{d\lambda}(\lambda) f(t(\lambda)) \quad (\text{A7})$$

for each $A_{kn;\lambda}$. In other words, the $A_{kn;t}$ are Fourier coefficients of a single function f while each $A_{kn;\lambda}$ is the kn th Fourier coefficient of a different function $g_{kn}(\lambda)$. But that poses no problem — we only sought to show that every t -Fourier coefficient is also the λ -Fourier coefficient of *some* function of λ . The pragmatic importance of this fact has to do with the evaluation of coefficients of torus functions, which, though stated in slightly different language, is the crux of the original argument in Ref. [26] and which we discuss in Section A 2.

Analogous reasoning to the above in the resonant case with

$$\frac{\omega_\theta}{\omega_r} = \frac{\Omega_\theta}{\Omega_r} = \frac{p}{z} \quad (\text{A8})$$

and p, z relatively prime shows that each coefficient $C_{j;t}$ of $f(t)$ can likewise be considered a Mino time coefficient $C_{j;\lambda}$ of some different function

$$g_j(\lambda) = e^{ij\omega_P \Delta t(\lambda)} \frac{dt}{d\lambda}(\lambda) f(t(\lambda)) \quad (\text{A9})$$

The difference is that now each of $\Delta t(\lambda)$, $dt/d\lambda(\lambda)$ and $f(t(\lambda))$ is a singly periodic function of λ with period

$$\Lambda_P = z\Lambda_r = p\Lambda_\theta = \frac{2\pi}{\omega_P} \quad (\text{A10})$$

Temporal Fourier coefficients can be calculated in the resonant case without first having to convert to Mino time. We will nonetheless express and evaluate coefficients $C_{j;t}$ as coefficients $C_{j;\lambda}$ in order to parallel the non-resonant case.

2. λ -based vs. t -based torus coordinates

Figure 5 represents $\mathbb{T}_\mathcal{E}^2$ as a compact 2π -by- 2π square in the χ_r, χ_θ angle coordinates defined in Section II B. Kerr geodesics trace out lines on this torus-square at constant velocity. With respect to any other time parameter, geodesic curves continue to be lines on the torus-square, but their parametric representations are not in general linear in time, nor are their velocities on the torus constant.

For any choice of time parameter, however, there is always some set of coordinates on the torus such that geodesic motion on that torus is linear in that time parameter and has constant velocity²⁶. For instance, with

respect to coordinate time t , there will be coordinates $\vec{\gamma} \equiv (\gamma_r, \gamma_\theta)$ such that

$$\gamma_r(t) = \omega_r t + \gamma_{r0} \quad (\text{A11a})$$

$$\gamma_\theta(t) = \omega_\theta t + \gamma_{\theta0} \quad (\text{A11b})$$

Note that in any set of angle coordinates in which the trajectory velocities are constant with respect to some time parameter, orbit trajectories will all be lines with the same slope $1 + q_{r\theta}$. Such coordinate systems are nevertheless distinct: identical ordered pairs in two such coordinate systems will not, in general, correspond to the same point on the torus.

Though not unique, the χ_r - χ_θ coordinate system on $\mathbb{T}_\mathcal{E}^2$ is nevertheless uniquely *useful*. Since each of γ_r, γ_θ would be a combination of χ_r and χ_θ , each point on the projected r - p_r curve of an orbit would be labeled by a pair of values $(\gamma_r, \gamma_\theta)$ rather than by a single value χ_r , and likewise for the projected θ - p_θ curve. This mixing of radial and polar motions in each torus coordinate makes most calculations harder than they need to be, and the impetus behind χ_r - χ_θ coordinates is precisely the convenience that flows from torus coordinates that separately shadow radial and polar motion.

Still, we sometimes are interested in values of quantities averaged over the $\vec{\gamma}$ coordinates. Luckily, by the correspondence between temporal Fourier coefficients of biperiodic functions and spatial Fourier coefficients of torus functions, equations (A2) and (A7) further establish that if $A_{kn;\vec{\gamma}}$ is the kn th coefficient of the torus function $f(\vec{\gamma})$ associated with $f(t)$, then it is also the kn th coefficient $A_{kn;\vec{\chi}}$ of the torus function $g_{kn}(\vec{\chi})$ corresponding to $g_{kn}(\lambda)$. This is important for evaluating torus coefficients in practice: it is usually very difficult to go from a function $f(r(t), p_r(t), \theta(t), p_\theta(t))$ to a form $f(\vec{\gamma})$ explicitly while it is straightforward to go from $g(r(\lambda), p_r(\lambda), \theta(\lambda), p_\theta(\lambda))$ to $g(\vec{\chi})$.

Appendix B: A synopsis of the Teukolsky formalism

Here we summarize some relevant aspects of the Teukolsky formalism as applied to the EMRI problem. More details of this application can be found in numerous references, including [14, 18–20, 22, 30, 33]. Our goal in this appendix is to justify the expressions for the coefficients $Z_{lmkn}^{H/\infty}$ and $Z_{lmj;\lambda}^{H/\infty}$ in equations (61) and (68), respectively.

ton's equations with respect to any evolution parameter. If the system is integrable, there will then exist a transformation to angle variables on the torus that increase linearly with respect to that evolution parameter [15].

²⁶ Darboux's theorem guarantees that there is a way to write Hamil-

1. The Weyl scalar, ψ_4

In 1972 Teukolsky derived the master equation [34, 35], a separable partial differential equation (PDE) whose solution describes the propagation in the Kerr spacetime of small perturbations to fields of different spin-weights s : scalar, electromagnetic and gravitational. Each solution to the master equation is a separable function which can be written as a multipole expansion. There are two computational approaches to solving the master equation for each multipole mode: the time-domain approach, which solves the resulting PDE directly, and the frequency-domain approach, which further Fourier expands the solutions. For the purposes of extracting flux information from gravitational perturbations, frequency-domain codes are the accuracy standard and the ones to which our savings proposal applies. We thus restrict our attention to the frequency-domain approaches to solving the master equation.

The combined multipole-Fourier expanded perturbations take the form

$${}_s\psi(t, r, \theta, \varphi) = \sum_{lm} \int_{-\infty}^{\infty} d\omega R_{lm\omega}(r) {}_sS_{lm}^{a\omega}(\theta) e^{-i\omega t + im\varphi} \quad , \quad (\text{B1})$$

where ω denotes the coordinate-time frequency of the perturbations at the field point due to the source. Each ${}_s\psi$ is a function of the field point (t, r, θ, φ) at which we wish to evaluate the perturbation. The s marker in equation (B1) is a “spin-weight parameter” [34] which denotes the perturbation type. For gravitational radiation, $s = -2$, and ${}_{-2}\psi = \psi_4 \rho^{-4}$ where

$$\rho = -(r - ia \cos \theta)^{-1} \quad . \quad (\text{B2})$$

The functions $R_{lm\omega}(r)$ and ${}_{-2}S_{lm}^{a\omega}(\theta)$ (described in the next subsections) each depend on the parameter ω as a consequence of the separation of variables procedure. When the source is a geodesic, ω turns out to be a discrete variable composed of harmonics of the radial, polar and azimuthal frequencies of that geodesic. That discrete dependence can be expressed differently for non-resonant orbits,

$$\omega = \omega_{mkn} = m\omega_\varphi + k\omega_r + n\omega_\theta \quad (\text{B3a})$$

and for resonant orbits,

$$\omega = \omega_{mj} = m\omega_\varphi + j\omega_P \quad . \quad (\text{B3b})$$

Once ω becomes a discrete variable, we can replace the integral over all possible ω in equation (B1) with a sum over either m, k, n or m, j for non-resonant and resonant sources, respectively.

Because everything in this paper deals with gravitational spin-weighting, we henceforth omit all the -2 subscripts. The net result is that equation (B1) becomes

$$\psi_4(t, r, \theta, \varphi) = \rho^4 \sum_{lmnk} R_{lm\omega_{mkn}}(r) S_{lm}^{a\omega_{mkn}}(\theta) e^{-i\omega_{mkn}t + im\varphi} \quad , \quad (\text{B4a})$$

for a non-resonant source and

$$\psi_4(t, r, \theta, \varphi) = \rho^4 \sum_{lmj} R_{lm\omega_{mj}}(r) S_{lm}^{a\omega_{mj}}(\theta) e^{-i\omega_{mj}t + im\varphi} \quad (\text{B4b})$$

for a resonant source.

2. The Spheroidal Harmonics

The functions $S_{lm}^{a\omega}(\theta)$ with a spin-weight of $s = -2$ are the gravitational (tensor) spheroidal harmonics, a generalization of the likewise spin-weighted spherical harmonics. These functions satisfy [35]

$$\left[(a\omega)^2 \cos^2 \theta + 4a\omega \cos \theta - \left(\frac{m^2 - 4m \cos \theta + 4}{\sin^2 \theta} \right) + \mathcal{C}_{lm} \right] S_{lm}^{a\omega}(\theta) + \frac{1}{\sin \theta} \frac{d}{d\theta} \left(\sin \theta \frac{dS_{lm}^{a\omega}(\theta)}{d\theta} \right) = 0 \quad . \quad (\text{B5})$$

\mathcal{C}_{lm} are the eigenvalues for which equation (B5) has solutions. Solving for $S_{lm}^{a\omega}(\theta)$ for given l, m, ω requires simultaneously determining an eigenvalue \mathcal{C}_{lm} and the associated spheroidal harmonic. These eigenvalue-eigenfunction pairs can be computed in several different ways (see Refs. [19, 22, 23]).

The spheroidal harmonics satisfy several orthogonality relations. The one we will need in this paper is that, for fixed m and ω ,

$$\int_0^\pi S_{lm}^{a\omega}(\theta) \bar{S}_{l'm'}^{a\omega}(\theta) \sin \theta d\theta = \frac{1}{2\pi} \delta_{ll'} \quad , \quad (\text{B6})$$

where the overbar denotes complex conjugation. We have chosen a normalization of $\frac{1}{2\pi}$, as in Ref. [14].

3. The radial Teukolsky functions

Solving for the radial functions $R_{lm\omega}(r)$ is more difficult. $R_{lm\omega}(r)$ satisfy the inhomogeneous radial Teukolsky equation [35]

$$\mathcal{T}_{lm\omega}(r) = \Delta^2 \frac{d}{dr} \left(\frac{1}{\Delta} \frac{dR_{lm\omega}(r)}{dr} \right) - V_{lm\omega}(r) R_{lm\omega}(r) \quad . \quad (\text{B7})$$

The potential $V_{lm\omega}(r)$ depends in part on the eigenvalue \mathcal{C}_{lm} of $S_{lm}^{a\omega}(\theta)$, so equation (B5) must be solved before the homogeneous or inhomogeneous version of equation (B7) can be.

The source term $\mathcal{T}_{lm\omega}$ is built by, among other things, evaluating $S_{lm}^{a\omega}(\theta)$ and two homogeneous solutions²⁷ $R_{lm\omega}^{\text{in/up}}(r)$ to (B7) along the geodesic source. Two general methods are described for constructing $R_{lm\omega}^{\text{in/up}}(r)$. One approach integrates the homogeneous Teukolsky equation (or, equivalently, the better numerically behaved Sasaki-Nakamura equation [36]) outward from the horizon. The other expands $R_{lm\omega}^{\text{in/up}}(r)$ in terms of hypergeometric functions (see [23] and references therein) and evaluates $R_{lm\omega}^{\text{in/up}}(r)$ directly at certain points, possibly extrapolating its values to nearby points with series expansions. Both approaches are fairly computationally costly. For detailed explanations on these different approaches and how various numerical problems are circumvented, see [23, 36–40]. We elaborate a bit more on the structure of the source term below.

4. The quantities $Z_{lm\omega}^{H/\infty}$

With the homogeneous radial solutions in hand, the inhomogeneous Teukolsky equations can be solved using the method of variation of parameters²⁸ [41]. The radial functions $R_{lm\omega}(r)$ can be written as [14, 26, 34, 35, 42, 43]

$$R_{lm\omega}(r) = Z_{lm\omega}^H(r)R_{lm\omega}^{\text{up}}(r) + Z_{lm\omega}^\infty(r)R_{lm\omega}^{\text{in}}(r) \quad , \quad (\text{B8})$$

where $Z_{lm\omega}^H(r)$ and $Z_{lm\omega}^\infty(r)$ are defined as

$$\begin{aligned} Z_{lm\omega}^H(r) &= \int_{r_+}^r dr' \frac{R_{lm\omega}^{\text{in}}(r')\mathcal{T}_{lm\omega}(r')}{c} \\ Z_{lm\omega}^\infty(r) &= \int_r^\infty dr' \frac{R_{lm\omega}^{\text{up}}(r')\mathcal{T}_{lm\omega}(r')}{c} \quad . \end{aligned} \quad (\text{B9})$$

The constant c is related to the Wronskian of $R_{lm\omega}^{\text{in}}(r)$ and $R_{lm\omega}^{\text{up}}(r)$. r_+ is the larger root of Δ and is the radial coordinate of the black hole horizon (the smaller root is denoted r_-).

Because we are only interested in the radiation going into the black hole and being carried away to infinity, we are only concerned about the asymptotic behavior of $R_{lm\omega}(r)$ as $r \rightarrow r_+$ and $r \rightarrow \infty$. In fact, the homogeneous basis solutions have been chosen to have the simplifying feature that

$$\begin{aligned} Z_{lm\omega}^\infty(r \rightarrow r_+) &= Z_{lm\omega}^\infty \\ Z_{lm\omega}^H(r \rightarrow r_+) &= 0 \\ Z_{lm\omega}^\infty(r \rightarrow \infty) &= 0 \\ Z_{lm\omega}^H(r \rightarrow \infty) &= Z_{lm\omega}^H \quad . \end{aligned} \quad (\text{B10})$$

Note that we have used the same notation for the *functions* $Z_{lm\omega}^{H/\infty}(r)$ and for the *constants* $Z_{lm\omega}^{H/\infty}$ representing their asymptotic values at ∞ and r_+ , respectively. As mentioned in Section IV A, the literature seems stuck with the rather backward notational convention that $Z_{lm\omega}^H$ is nonvanishing at ∞ while $Z_{lm\omega}^\infty$ is nonvanishing at r_+ .

The radial functions as $r \rightarrow \infty$ and $r \rightarrow r_+$ thus become

$$\begin{aligned} R_{lm\omega}^\infty &= R_{lm\omega}(r \rightarrow \infty) \\ &= Z_{lm\omega}^H R_{lm\omega}^{\text{up}}(r \rightarrow \infty) \\ &= Z_{lm\omega}^H r^3 e^{i\omega r^*} \end{aligned} \quad (\text{B11a})$$

$$\begin{aligned} R_{lm\omega}^H &= R_{lm\omega}(r \rightarrow r_+) \\ &= Z_{lm\omega}^\infty R_{lm\omega}^{\text{in}}(r \rightarrow r_+) \\ &= Z_{lm\omega}^\infty \Delta^2 e^{-ikr^*} \quad , \end{aligned} \quad (\text{B11b})$$

where and $k \equiv \omega - ma/(2r_+)$ and r^* is the Kerr tortoise coordinate defined by

$$\begin{aligned} r^*(r) &= r + \frac{2r_+}{r_+ - r_-} \ln \frac{r - r_+}{2} - \frac{2r_-}{r_+ - r_-} \ln \frac{r - r_-}{2} \\ \frac{dr^*}{dr} &= \frac{r^2 + a^2}{\Delta} \quad . \end{aligned} \quad (\text{B12})$$

The coefficients $Z_{lm\omega}^{H/\infty}$ are found from

$$Z_{lm\omega}^{H/\infty} = \int_{r_+}^\infty dr' \frac{R_{lm\omega}^{\text{in/up}}(r')\mathcal{T}_{lm\omega}(r')}{c} \quad . \quad (\text{B13})$$

The source function $\mathcal{T}_{lm\omega}(r)$ in (B13) is an integral of the form

$$\begin{aligned} \mathcal{T}_{lm\omega}(r) &= \\ &= \int_{-\infty}^\infty dt \int d\Omega B_{m\omega}(t, r, \theta, \varphi) S_{lm}^{a\omega} e^{i\omega t} e^{-im\varphi} \quad . \end{aligned} \quad (\text{B14})$$

The source term is derived in [44], but we have written it borrowing the notation $B_{m\omega}$ from [14]. All that matters for our purposes is that $B_{m\omega}$ is built out of a series of operations on the null tetrad components of the energy-momentum tensor of the orbiting particle and thus contains delta functions and derivatives of delta functions centered on the source geodesic. Thus, the $d\Omega$ integral can be evaluated, resulting in every θ and φ in (B14) being replaced with the source trajectories $\theta_s(t), \varphi_s(t)$ (the subscript s here denotes “source” as opposed to a spin-weight as earlier in this Appendix).

Delta functions $\delta(r - r_s(t))$ and derivatives thereof still remain in (B14), along with the integration over t . When we plug (B14) into equation (B13), we can switch the order of integration for r' and t and use those remaining delta functions in r (we rename r' to r now, for simplicity) to replace every r with $r_s(t)$. The net result

²⁷ Other basis solutions to the homogeneous equation exist, e.g. the out/down basis $R_{lm\omega}^{\text{out/down}}(r)$. For a summary, see [3] and references therein.

²⁸ Most references use the method of Green functions, but variation of parameters works as well.

is that $Z_{lm\omega}^{H/\infty}$ takes the form

$$Z_{lm\omega}^{H/\infty} = \int_{-\infty}^{\infty} dt e^{i\omega t} e^{-im\varphi_s(t)} \mathcal{I}_{lm\omega}^{H/\infty}(r_s(t), \theta_s(t)) \quad . \quad (\text{B15})$$

The functions $\mathcal{I}_{lm\omega}^{H/\infty}$ depend on r and θ both directly and via a combination of elementary functions, the spheroidal harmonics, the homogeneous radial Teukolsky functions, and various derivatives thereof. Explicit expressions can be found in several sources (see, for instance, [14, 23]).

We will now use this form for $Z_{lm\omega}^{H/\infty}$ to define the coefficients $Z_{lmkn}^{H/\infty}$ and $Z_{lmj;\lambda}^{H/\infty}$ to which the efficiency arguments of Sections IV and V apply, respectively.

5. The quantities $Z_{lmkn}^{H/\infty}$

We now show that, when the source is a non-resonant orbit, the ω -dependence of $Z_{lm\omega}^{H/\infty}$ takes the form

$$Z_{lm\omega}^{H/\infty} = \sum_{kn} Z_{lmkn;\lambda}^{H/\infty} \delta(\omega - \omega_{mkn}) \quad , \quad (\text{B16})$$

where ω_{mkn} are the coordinate-time harmonic frequencies defined in equation (B3a). Note that even though equation (B16) has the form of the Fourier transform of an almost-periodic function with respect to coordinate time t , we are free to interpret the coefficients of the delta functions either as Fourier coefficients of a coordinate time function or as Fourier coefficients of a (different) Mino time function.

For the reasons stated in Appendix A, we opt for the latter and begin by rewriting (B15) as an integral over Mino time. Treating all source coordinates as functions of λ and using equation (A4), we get

$$Z_{lm\omega}^{H/\infty} = \int_{-\infty}^{\infty} d\lambda e^{i\omega(\Gamma\lambda + \Delta t(\lambda))} e^{-im\varphi_s(\lambda)} \mathcal{I}_{lm\omega}^{H/\infty}(r_s(\lambda), \theta_s(\lambda)) \quad . \quad (\text{B17})$$

We now use the fact (see Ref. [26]) that, like $dt/d\lambda$, $d\varphi/d\lambda$ depends on r and θ and is thus biperiodic when evaluated on a trajectory $r(\lambda), \theta(\lambda)$. $d\varphi/d\lambda$ has a nonzero average value Ω_φ defined by

$$\Omega_\varphi \equiv \lim_{\Lambda \rightarrow \infty} \frac{1}{\Lambda} \int_{-\Lambda/2}^{\Lambda/2} d\lambda \frac{d\varphi}{d\lambda} \quad . \quad (\text{B18})$$

Consequently, the function $\varphi_s(\lambda)$ takes the form

$$\varphi_s(\lambda) = \Omega_\varphi \lambda + \Delta\varphi_s(\lambda) \quad (\text{B19})$$

where $\Delta\varphi_s(\lambda)$ is biperiodic in λ and has zero average value. Like its radial and polar counterparts, Ω_φ is related to the coordinate-time frequency ω_φ via

$$\omega_\varphi = \frac{1}{\Gamma} \Omega_\varphi \quad . \quad (\text{B20})$$

More generally, coordinate-time and Mino-time frequencies are related by

$$\omega = \frac{1}{\Gamma} \Omega \quad . \quad (\text{B21})$$

In light of equations (B19)–(B21), equation (B17) becomes

$$Z_{lm\omega}^{H/\infty} = \int_{-\infty}^{\infty} d\lambda e^{i\Omega\lambda} e^{-im\Omega_\varphi\lambda} e^{i\omega\Delta t(\lambda)} e^{-im\Delta\varphi_s(\lambda)} \mathcal{I}_{lm\omega}^{H/\infty}(r_s(\lambda), \theta_s(\lambda)) \quad . \quad (\text{B22})$$

Note that the coordinate-time frequency ω still appears as a parameter in both $\mathcal{I}_{lm\omega}^{H/\infty}$ and in the argument of $e^{i\omega\Delta t(\lambda)}$.

Like $\Delta t(\lambda)$ and $\Delta\varphi(\lambda)$, the functions $\mathcal{I}_{lm\omega}^{H/\infty}$ are biperiodic in λ . We define the biperiodic function

$$f_{lm\omega}^{H/\infty}(\lambda) \equiv e^{i\omega\Delta t(\lambda)} e^{-im\Delta\varphi_s(\lambda)} \mathcal{I}_{lm\omega}^{H/\infty}(r_s(\lambda), \theta_s(\lambda)) \quad (\text{B23})$$

and Fourier expand it as

$$f_{lm\omega}^{H/\infty}(\lambda) \equiv \sum_{k,n} Z_{lm\omega;k n;\lambda}^{H/\infty} e^{-i(k\Omega_\theta + n\Omega_r)\lambda} \quad . \quad (\text{B24})$$

Inserting (B24) into equation (B22) yields

$$\begin{aligned} Z_{lm\omega}^{H/\infty} &= \sum_{k,n} Z_{lm\omega;k n;\lambda}^{H/\infty} \int_{-\infty}^{\infty} d\lambda e^{i\Omega\lambda} e^{-i(m\Omega_\varphi + k\Omega_\theta + n\Omega_r)\lambda} \\ &= \sum_{k,n} Z_{lm\omega;k n;\lambda}^{H/\infty} 2\pi \delta(\omega - \omega_{mkn}) \\ &= \sum_{k,n} Z_{lm\omega=\omega_{mkn};k n;\lambda}^{H/\infty} 2\pi \delta(\omega - \omega_{mkn}) \quad . \end{aligned} \quad (\text{B25})$$

The multipole index m and the Fourier indices k, n all do double duty by helping to specify the value of the parameter ω in $Z_{lm\omega=\omega_{mkn};k n;\lambda}^{H/\infty}$. Consequently, each $Z_{lm\omega=\omega_{mkn};k n;\lambda}^{H/\infty}$ is fully specified by the four integers l, m, k, n and we can define

$$Z_{lmkn;\lambda}^{H/\infty} \equiv 2\pi Z_{lm\omega=\omega_{mkn};k n;\lambda}^{H/\infty} \quad . \quad (\text{B26})$$

We will absorb the factor of 2π into the function $f_{lm\omega}^{H/\infty}$ in the integral defining $Z_{lmkn;\lambda}^{H/\infty}$.

Since each $Z_{lmkn;\lambda}^{H/\infty}$ is a temporal Fourier coefficient, then by the arguments of Section III D 1, we determine it by instead computing the corresponding spatial Fourier coefficient

$$\begin{aligned} Z_{lmkn}^{H/\infty} &\equiv \frac{1}{(2\pi)^2} \int_0^{2\pi} d\chi_r \int_0^{2\pi} d\chi_\theta e^{ik\chi_\theta} e^{in\chi_r} f_{lm\omega=\omega_{mkn}}^{H/\infty}(\chi_r, \chi_\theta) \end{aligned} \quad (\text{B27})$$

of the associated torus function in the $\vec{\chi}$ torus coordinates. Despite our earlier notation, the function $\mathcal{I}_{lm\omega}^{H/\infty}$ actually depends not only on r and θ but also on their conjugate momenta (since it depends on the r and θ velocities via the energy-momentum tensor of the particle). $f_{lm\omega}^{H/\infty}$ inherits this dependence, and it is thus appropriate to write it as a function of the torus coordinates $\vec{\chi}$ that need not have any special symmetries on the torus.

As we show in Appendix C, the averaged fluxes required to evaluate the RHS of the adiabatic equations for $\vec{\mathcal{E}}$ depend on the $Z_{lmkn}^{H/\infty}$ defined in (B27). Thus, equation (B27) verifies equation (61), on which the savings arguments of Section IV B are based.

6. The quantities $Z_{lmj;\lambda}^{H/\infty}$

We now show that when the source is a resonant orbit, the ω -dependence of $Z_{lm\omega}^{H/\infty}$ instead takes the form

$$Z_{lm\omega}^{H/\infty} = \sum_j Z_{lmj;\lambda}^{H/\infty} \delta(\omega - \omega_{mj}) \quad , \quad (\text{B28})$$

where ω_{mj} are now the coordinate-time harmonic frequencies defined in equation (B3b). As before, we are free to interpret the coefficients of the delta functions either as Fourier coefficients of a coordinate time function or as Fourier coefficients of a (different) Mino time function and take the computationally tractable latter option.

Equations (B17)–(B23) carry over to the resonant case with the difference that each of $\Delta t(\lambda)$, $\Delta\varphi(\lambda)$, $r(\lambda)$, $\theta(\lambda)$ and thus $\mathcal{I}_{lm\omega}^{H/\infty}$ is now singly periodic with period Λ_P . Thus $f_{lm\omega}^{H/\infty}$ can be Fourier expanded in harmonics of a single fundamental frequency Ω_P ,

$$f_{lm\omega}^{H/\infty}(\lambda; \vec{\chi}_0) \equiv \sum_j Z_{lm\omega;j;\lambda}^{H/\infty}(\vec{\chi}_0) e^{-ij\Omega_P\lambda} \quad (\text{B29})$$

with

$$Z_{lm\omega;j;\lambda}^{H/\infty}(\vec{\chi}_0) = \frac{1}{\Lambda_P} \int_0^{\Lambda_P} d\lambda e^{ij\Omega_P\lambda} f_{lm\omega}^{H/\infty}(\lambda; \vec{\chi}_0) \quad . \quad (\text{B30})$$

The functions $f_{lm\omega}^{H/\infty}(\lambda; \vec{\chi}_0)$ are induced from some torus function, so by the arguments of Section III D 2, both they and the coefficients $Z_{lm\omega;j;\lambda}^{H/\infty}(\vec{\chi}_0)$ depend on initial positions, which we can represent compactly as a dependence on initial position $\vec{\chi}_0$ on the phase space torus.

Inserting (B29) into the resonant version of equation (B22) yields

$$\begin{aligned} Z_{lm\omega}^{H/\infty} &= \sum_j Z_{lm\omega;j;\lambda}^{H/\infty} \int_{-\infty}^{\infty} d\lambda e^{i\Omega\lambda} e^{-i(m\Omega_\varphi + j\Omega_P)\lambda} \\ &= \sum_j Z_{lm\omega;j;\lambda}^{H/\infty} 2\pi\delta(\omega - \omega_{mj}) \\ &= \sum_j Z_{lm\omega=\omega_{mj};j;\lambda}^{H/\infty} 2\pi\delta(\omega - \omega_{mj}) \quad . \end{aligned} \quad (\text{B31})$$

The relevant quantities, then, are those in equation (B30) with the parameter ω in $f_{lm\omega}^{H/\infty}$ set to ω_{mj} .

Paralleling the non-resonant case, each $Z_{lm\omega=\omega_{mj};j;\lambda}^{H/\infty}$ is fully specified by the three integers l, m, j . By absorbing the factor of 2π into the functions $f_{lm\omega}^{H/\infty}$ in the integrand of (B30), we can define the notationally more compact coefficients

$$Z_{lmj;\lambda}^{H/\infty} \equiv 2\pi Z_{lm\omega=\omega_{mj};j;\lambda}^{H/\infty} \quad . \quad (\text{B32})$$

By the construction above, each such $Z_{lmj;\lambda}^{H/\infty}$ is given by

$$Z_{lmj;\lambda}^{H/\infty} = \frac{1}{\Lambda_P} \int_0^{\Lambda_P} d\lambda e^{ij\Omega_P\lambda} f_{lm\omega=\omega_{mj}}^{H/\infty}(\lambda; \vec{\chi}_0) \quad . \quad (\text{B33})$$

As we show in Appendix C, the time-averaged fluxes for $\vec{\mathcal{E}}$ from resonant orbits depend on the $Z_{lmj;\lambda}^{H/\infty}$ defined in (B33). Thus, equation (B33) verifies equation (68), on which the more speculative savings arguments of Section V are based.

Appendix C: Fluxes from the Teukolsky formalism

In this appendix, we review how the apparatus of Appendix B yields fluxes of conserved quantities. Several authors [3, 14, 19, 20, 22, 23] implement this Fourier-domain formalism in TB codes to calculate the radiative $\vec{\mathcal{E}}$ fluxes at radial infinity and the horizon to determine how the inspiral evolves. We show how expressions for time-averaged (as opposed to torus-averaged) fluxes differ between non-resonant and resonant orbits.

1. Overview of flux calculation

We will not refer to the Q flux, but restrict our discussion to E and L_z . To determine the evolution of an inspiral in orbital parameter space, we use the E and L_z fluxes at radial infinity and the horizon as proxies for the local self-force. This subsection gives an overview of the calculation for finding these fluxes.

From the Weyl scalar ψ_4 , the gravitational waveform and the E , L_z and Q radiation fluxes can be calculated. Specifically, we can calculate the polarizations h_+ and h_\times of the metric perturbations at infinity (i.e. the GWs) via

$$\psi_4 = \frac{1}{2} \frac{\partial^2}{\partial t^2} (h_+ + ih_\times) \quad . \quad (\text{C1})$$

After integrating equation (C1) twice to get $h_+ + ih_\times$ (the integration constants are set to zero), we can calculate an effective GW stress-energy tensor at infinity [45] as

$$T_{\alpha\beta}^\infty = \frac{1}{16\pi} \left\langle \frac{\partial h_+}{\partial x^\alpha} \frac{\partial h_+}{\partial x^\beta} + \frac{\partial h_\times}{\partial x^\alpha} \frac{\partial h_\times}{\partial x^\beta} \right\rangle_{\text{several wavelengths}} \quad . \quad (\text{C2})$$

The average over several wavelengths signifies the following [46]. The stress-energy tensor of GWs contributes to curvature in the analogous way that the stress-energy tensor of matter does. However, since GWs are fluctuations in the metric itself, we cannot define their stress-energy tensor at a field point because, with only information at one point, we cannot distinguish between the curvature of the background spacetime and the contributions to that curvature from fluctuations on the spacetime. To distinguish between the background and the effect of the fluctuations, there needs to be either a length or frequency scale separation between the two. In our case, as $r \rightarrow \infty$, the background curvature scale is much greater than the wavelength of the fluctuations. Therefore, we average over several gravitational wavelengths in order to smooth out the fluctuations and determine only the secular contribution of such GWs to the curvature.

While the background curvature scale is much greater than the GW wavelength, GW detectors look not for spatial fluctuations to the metric but rather temporal ones. Therefore, in the spirit of applicability to actual experiments, it is more useful to distinguish between a background frequency (i.e. a reciprocal of a curvature scale in the timelike direction) and the much larger frequency of the fluctuations. Therefore, rather than average over several wavelengths, we average over several periods to isolate the net effect of the GWs [46].

The energy and angular momentum fluxes carried to radial infinity by GWs are related to the components of $T_{\alpha\beta}^\infty$ by,

$$\begin{aligned} T_{tt}^\infty &= \frac{dE}{dt dA} \\ T_{t\varphi}^\infty &= \frac{dL_z}{dt dA} \end{aligned} \quad (C3)$$

The E and L_z fluxes are calculated by integrating equation (C3) over a 2-sphere of radius r on a constant t spacelike hypersurface

$$\begin{aligned} \left(\frac{dE}{dt} \right)^\infty &= \int T_{tt}^\infty r^2 d\Omega \\ \left(\frac{dL_z}{dt} \right)^\infty &= \int T_{t\varphi}^\infty r^2 d\Omega \end{aligned} \quad (C4)$$

The time-averaged fluxes from a given geodesic can then

be calculated by

$$\begin{aligned} \left\langle \frac{dE}{dt} \right\rangle_t^\infty &= \lim_{T \rightarrow \infty} \frac{1}{T} \int_{-\frac{T}{2}}^{\frac{T}{2}} \left(\frac{dE}{dt} \right)^\infty dt \\ \left\langle \frac{dL_z}{dt} \right\rangle_t^\infty &= \lim_{T \rightarrow \infty} \frac{1}{T} \int_{-\frac{T}{2}}^{\frac{T}{2}} \left(\frac{dL_z}{dt} \right)^\infty dt \end{aligned} \quad (C5)$$

Alternatively, we can calculate the torus-averaged fluxes by taking the average of the time-averages for all geodesics with a given set of orbital parameters over all possible initial conditions,

$$\begin{aligned} \left\langle \frac{dE}{dt} \right\rangle_{\vec{\gamma}}^\infty &= \frac{1}{(2\pi)^2} \int_0^{2\pi} d\gamma_{r_o} \int_0^{2\pi} d\gamma_{\theta_o} \left\langle \frac{dE}{dt} \right\rangle_t^\infty \\ \left\langle \frac{L_z}{dt} \right\rangle_{\vec{\gamma}}^\infty &= \frac{1}{(2\pi)^2} \int_0^{2\pi} d\gamma_{r_o} \int_0^{2\pi} d\gamma_{\theta_o} \left\langle \frac{dL_z}{dt} \right\rangle_t^\infty \end{aligned} \quad (C6)$$

An analogous procedure can be performed to calculate the fluxes at the horizon [38, 47].

2. Fluxes from ψ_4

Combining equations (B4) with (B11) we find that

$$\psi_4^\infty = \rho^{-4} \sum_{lmkn} Z_{lmkn;\lambda}^H r^3 e^{i\omega_{mkn} r^*} S_{lm}^{a\omega_{mkn}}(\theta) e^{-i\omega_{mkn} t + im\varphi} \quad (C7a)$$

for non-resonant orbits, and

$$\psi_4^\infty = \rho^{-4} \sum_{lmj} Z_{lmj;\lambda}^H r^3 e^{i\omega_{mj} r^*} S_{lm}^{a\omega_{mj}}(\theta) e^{-i\omega_{mj} t + im\varphi} \quad (C7b)$$

Similar expressions can be found for ψ_4^H , but for brevity, we will only proceed with the detailed computations for $r \rightarrow \infty$. Also, we only work out the details for the E flux, but the L_z flux follows exactly the same prescription.

Following the prescription laid out in section C1 we find that

$$\left(\frac{dE}{dt} \right)^\infty = \frac{1}{4\pi} \left\langle \sum_{lm\omega} \sum_{l'm'\omega'} \frac{1}{\omega\omega'} Z_{lm\omega;\lambda}^H \bar{Z}_{l'm'\omega';\lambda}^H e^{i(\omega-\omega')r^*} e^{-i(\omega-\omega')t} \int_0^\pi \sin\theta S_{lm}^{a\omega}(\theta) \bar{S}_{l'm'}^{a\omega'}(\theta) d\theta \int_0^{2\pi} d\varphi e^{i(m-m')\varphi} \right\rangle_{\text{several periods}} \quad (C8)$$

where ω, ω' denote the discrete variables that are two-

indexed for resonant orbits and three-indexed for non-

resonant orbits. Performing the φ integration in equation (C8) yields $2\pi\delta_{mm'}$. We thus set $m = m'$ everywhere.

For non-resonant orbits, equation (C8) becomes

$$\left(\frac{dE}{dt}\right)^\infty = \frac{1}{2} \left\langle \sum_{lmkn} \sum_{l'k'n'} \frac{1}{\omega_{mkn}\omega_{mk'n'}} Z_{lmkn;\lambda}^H \bar{Z}_{l'mk'n';\lambda}^H e^{-i\{(n-n')\omega_r + (k-k')\omega_\theta\}(t-r^*)} \int_0^\pi d\theta \sin\theta S_{lm}^{a\omega_{mkn}} \bar{S}_{l'm}^{a\omega_{mk'n'}} \right\rangle_{\text{several periods}} \quad (\text{C9})$$

We are interested in an infinite time-average of equation (C9). Therefore, we can drop the average over several periods because the time-averaging process will smooth out the fluxes so that the period averaging will have no further effect once we have time-averaged. The time-average of a function corresponds to the constant term in a Fourier expansion, so the argument of the exponential in t will need to be zero. Therefore, performing the infinite time-average yields the added conditions $k = k'$ and $n = n'$. This equates the frequencies everywhere, including in the spheroidal harmonics. We can therefore now perform the θ integration using (B6) and get $l = l'$. The result is that the infinite time average in the non-resonant case is

$$\left\langle \frac{dE}{dt} \right\rangle_t^\infty = \sum_{lmkn} \frac{1}{4\pi\omega_{mkn}^2} |Z_{lmkn;\lambda}^H|^2 \quad (\text{C10})$$

In Section IIID 1, we saw that infinite time averages over non-resonant orbits are the same as torus averages over non-resonant tori with respect to the corresponding torus coordinates. The time average over t on the LHS of (C10) thus corresponds to a torus average over the $\tilde{\gamma}$ torus coordinates. We also saw that $|Z_{lmkn;\lambda}^H|^2 = |Z_{lmkn}^H|^2$, where the Z_{lmkn}^H are spatial Fourier coefficients

(for fixed l, m) with respect to the $\tilde{\chi}$ torus coordinates. Therefore, the torus averaged energy flux is

$$\left\langle \frac{dE}{dt} \right\rangle_{\tilde{\gamma}}^\infty = \sum_{lmkn} \frac{1}{4\pi\omega_{mkn}^2} |Z_{lmkn}^H|^2 \quad (\text{C11})$$

This expression is true for all tori, as torus averages are insensitive to whether the orbits on that torus are resonant or non-resonant.

Analogous arguments lead to the angular momentum flux at infinity. Similar arguments to those above then lead to the corresponding fluxes at the horizon. The upshot is that all the torus-averaged fluxes are given by

$$\begin{aligned} \left\langle \frac{dE}{dt} \right\rangle_{\tilde{\gamma}}^{H/\infty} &= \sum_{lmkn} \frac{\alpha_{lmkn}^{H/\infty}}{4\pi\omega_{mkn}^2} |Z_{lmkn}^{H/\infty}|^2 \\ \left\langle \frac{dL_z}{dt} \right\rangle_{\tilde{\gamma}}^{H/\infty} &= \sum_{lmkn} \frac{\alpha_{lmkn}^{H/\infty} m}{4\pi\omega_{mkn}^3} |Z_{lmkn}^{H/\infty}|^2 \end{aligned} \quad (\text{C12})$$

where $\alpha_{lmkn}^\infty \equiv 1$ and the details of α_{lmkn}^H can be found in reference [38]. We note that there is no residual dependence on the initial conditions $\tilde{\chi}_0$.

We return to equation (C9) and evaluate it for a resonant orbit,

$$\left(\frac{dE}{dt}\right)^\infty = \frac{1}{2} \left\langle \sum_{lmj} \sum_{l'j'} \frac{1}{\omega_{mj}\omega_{mj'}} Z_{lmj;\lambda}^H Z_{l'mj';\lambda}^{*H} e^{-i(j-j')\omega_P(t-r^*)} \int_0^\pi d\theta \sin\theta S_{lm}^{a\omega_{mj}} \bar{S}_{l'm}^{a\omega_{mj'}} \right\rangle_{\text{several periods}} \quad (\text{C13})$$

As was the case with the non-resonant infinite time average, we can drop the averaging over several periods. Additionally, the resonant time-average picks out the constant term in the Fourier expansion, which results when $j = j'$. Therefore,

$$\left\langle \frac{dE}{dt} \right\rangle_t^\infty = \sum_{lmj} \frac{1}{4\pi\omega_{mj}^2} |Z_{lmj;\lambda}^H|^2 \quad (\text{C14})$$

The rest of the E and L_z time-averaged fluxes can be

found similarly and are,

$$\begin{aligned} \left\langle \frac{dE}{dt} \right\rangle_t^{H/\infty} &= \sum_{lmj} \frac{\alpha_{lmj}^{H/\infty}}{4\pi\omega_{mj}^2} |Z_{lmj;\lambda}^{H/\infty}|^2 \\ \left\langle \frac{dL_z}{dt} \right\rangle_t^{H/\infty} &= \sum_{lmj} \frac{\alpha_{lmj}^{H/\infty} m}{4\pi\omega_{mj}^3} |Z_{lmj;\lambda}^{H/\infty}|^2 \end{aligned} \quad (\text{C15})$$

We remark that unlike the torus-averaged fluxes, the

time-averaged fluxes of resonant orbits clearly depend on the initial conditions of the orbit, since as we saw in Section B 6, $|Z_{lmj;\lambda}^{H/\infty}|^2$ is not the same for all initial conditions.

Alternatively, we can write the time-averaged fluxes of equation (C15) explicitly in terms of the torus coefficients

$Z_{lmkn}^{H/\infty}$. From section III D 2 we know that

$$Z_{lmj;\lambda}^{H/\infty} = \sum_{\substack{kn: \\ pn+zk=j}} Z_{lmkn}^{H/\infty} e^{-in\chi_{r_0}} e^{-ik\chi_{\theta_0}} \quad . \quad (C16)$$

Therefore, we can rewrite equation (C15) as

$$\begin{aligned} \left\langle \frac{dE}{dt} \right\rangle_t^{H/\infty} &= \sum_{lmkn} \sum_{\substack{k'n': \\ zn+pk= \\ zn'+pk'}} \frac{\alpha_{lmkn}^{H/\infty}}{4\pi\omega_{mkn}^2} Z_{lmkn}^{H/\infty} Z_{lmk'n'}^{*H/\infty} e^{-i\{(n-n')\chi_{r_0} + (k-k')\chi_{\theta_0}\}} \\ \left\langle \frac{dL_z}{dt} \right\rangle_t^{H/\infty} &= \sum_{lmkn} \sum_{\substack{k'n': \\ zn+pk= \\ zn'+pk'}} \frac{\alpha_{lmkn}^{H/\infty} m}{4\pi\omega_{mkn}^3} Z_{lmkn}^{H/\infty} Z_{lmk'n'}^{*H/\infty} e^{-i\{(n-n')\chi_{r_0} + (k-k')\chi_{\theta_0}\}} \quad . \end{aligned} \quad (C17)$$

The explicit initial condition dependence is now evident. Notice that if we average the time-averaged flux expressions of equation (C17) over all possible initial conditions, we reproduce the torus-averaged fluxes of equation (C11)

$$\begin{aligned} \left\langle \frac{dE}{dt} \right\rangle_\gamma^{H/\infty} &= \frac{1}{(2\pi)^2} \int_0^{2\pi} d\chi_{r_0} \int_0^{2\pi} d\chi_{\theta_0} \left\langle \frac{dE}{dt} \right\rangle_t^{H/\infty} \\ &= \sum_{lmkn} \frac{\alpha_{lmkn}^{H/\infty}}{4\pi\omega_{mkn}^2} |Z_{lmkn}^{H/\infty}|^2 \quad . \end{aligned} \quad (C18)$$

-
- [1] Y. Mino, Prog. Theor. Phys. **113**, 733 (2005), gr-qc/0506003.
 - [2] E. E. Flanagan and T. Hinderer (2010), 1009.4923.
 - [3] E. F. S Drasco and S. A. Hughes, Class. Quant. Grav. **22**, 801 (2005), arXiv:gr-qc/0505075.
 - [4] A. Pound, E. Poisson, and B. G. Nickel, Phys. Rev. **D72**, 124001 (2005), gr-qc/0509122.
 - [5] A. Pound and E. Poisson, Phys. Rev. **D77**, 044012 (2008), 0708.3037.
 - [6] A. Pound and E. Poisson, Phys. Rev. **D77**, 044013 (2008), 0708.3033.
 - [7] T. Hinderer and E. E. Flanagan, Phys. Rev. **D78**, 064028 (2008), 0805.3337.
 - [8] J. Levin and G. Perez-Giz, Phys. Rev. D **77**, 103005 (2008), 0802.0459.
 - [9] R. Grossman, J. Levin, and G. Perez-Giz (2011), 1105.5811.
 - [10] J. Levin and R. Grossman, gr-qc/08093838 (2008).
 - [11] Y. Mino, Phys. Rev. D **67**, 084027 (2003), arXiv:gr-qc/0302075.
 - [12] J. Levin, Class. Quant. Grav. **26**, 235010 (2009), 0907.5195.
 - [13] J. Levin, S. T. McWilliams, and H. Contreras, ArXiv e-prints (2010), 1009.2533.
 - [14] S. Drasco and S. Hughes, Phys. Rev. D **73**, 024027 (2006), gr-qc/0509101.
 - [15] V. I. Arnold, V. V. Kozlov, and A. I. Neishtadt, *Mathematical Aspects of Classical and Celestial Mechanics* (Springer, 2002), 3rd ed.
 - [16] J. Kevorkian and J. D. Cole, *Multiple Scale and Singular Perturbation Methods* (Springer, 1996).
 - [17] F. Verhulst, *Methods and applications of singular perturbations*, vol. 50 of *Texts in Applied Mathematics* (Springer, New York, 2005), ISBN 978-0387-22966-9; 0-387-22966-3, boundary layers and multiple timescale dynamics, URL <http://dx.doi.org/10.1007/0-387-28313-7>.
 - [18] N. Sago, T. Tanaka, W. Hikida, K. Ganz, and H. Nakano, Prog. Theor. Phys. **115**, 873 (2006), gr-qc/0511151.
 - [19] K. Ganz, W. Hikida, H. Nakano, N. Sago, and T. Tanaka, Prog. Theor. Phys. **117**, 1041 (2007), gr-qc/0702054.
 - [20] K. Glampedakis and D. Kennefick, Phys. Rev. D **66**,

- 044002 (2002), gr-qc/0203086.
- [21] D. C. Wilkins, Phys. Rev. D **5**, 814 (1972).
 - [22] S. A. Hughes, erratum-ibid.d **63**, 049902 (2001), URL <http://www.citebase.org/abstract?id=oai:arXiv.org:gr-qc/9910091>.
 - [23] R. Fujita, W. Hikida, and H. Tagoshi, Prog. Theor. Phys. **121**, 843 (2009), 0904.3810.
 - [24] A. M. Samoilenko, *Elements of the Mathematical Theory of Multi-Frequency Oscillations* (Springer-Verlag, 1991).
 - [25] C. Corduneanu, *Almost Periodic Functions* (Interscience Publishers, 1968).
 - [26] S. Drasco and S. A. Hughes, Phys. Rev. D **69**, 044015 (2004), arXiv:astro-ph/0308479.
 - [27] S. S. F. Schilder, W. Vogt and H. M. Osinga, International Journal for Numerical Methods in Engineering **67**, 629 (2006).
 - [28] Boyce and DiPrima, *Elementary Differential Equations and Boundary Value Problems* (Wiley, 2005).
 - [29] J. W. R. Kevin Beanland and C. Stevenson, The American Mathematical Monthly **116**, 531 (2009), URL <http://www.jstor.org/stable/40391145>.
 - [30] N. Sago, T. Tanaka, W. Hikida, and H. Nakano, Prog. Theor. Phys. **114**, 509 (2005), gr-qc/0506092.
 - [31] W. T. V. William H. Press, Saul A. Teukolsky and B. P. Flannery, *Numerical Recipes* (Cambridge University Press, 2007), 3rd ed.
 - [32] P. L. Chrzanowski, Phys. Rev. **D13**, 806 (1976).
 - [33] S. A. Hughes, S. Drasco, E. E. Flanagan, and J. Franklin, Phys. Rev. Lett. **94**, 221101 (2005), gr-qc/0504015.
 - [34] S. A. Teukolsky, Phys. Rev. Lett. **29**, 1114 (1972).
 - [35] S. A. Teukolsky, Ap. J. **185**, 635 (1973).
 - [36] S. A. Hughes, Phys. Rev. D **62**, 044029 (2000), arXiv:gr-qc/0002043.
 - [37] W. H. Press and S. A. Teukolsky, Astrophys. J. **185**, 649 (1973).
 - [38] S. A. Teukolsky and W. H. Press, Astrophys. J. **193**, 443 (1974).
 - [39] S. Chandrasekhar, Royal Society of London Proceedings Series A **343**, 289 (1975).
 - [40] E. W. Leaver, Royal Society of London Proceedings Series A **402**, 285 (1985).
 - [41] R. Haberman, *Elementary Applied Partial Differential Equations with Fourier Series and Boundary Value Problems* (Prentice Hall, 1998), 3rd ed.
 - [42] K. Glampedakis, Class. Quant. Grav. **22**, S605 (2005), gr-qc/0509024.
 - [43] M. Sasaki and T. Nakamura, Phys. Lett. A **89A**, 68 (1982).
 - [44] R. A. Breuer, *Gravitational perturbation theory and synchrotron radiation*, vol. 44 of *Lecture notes in physics* (Springer-Verlag, the University of California, 1975), ISBN 0387075305, 9780387075303.
 - [45] R. A. Isaacson, Phys. Rev. **166**, 1272 (1968).
 - [46] M. Maggiore, *Gravitational Waves, Volume 1: Theory and Experiments* (Oxford University Press, 2008).
 - [47] S. W. Hawking and J. B. Hartle, Commun. Math. Phys. **27**, 283 (1972).