

Dynamics of spinning test particles in Kerr spacetime

Michael D. Hartl*

Department of Physics, California Institute of Technology, Pasadena, California 91125

(Received 14 October 2002; published 14 January 2003)

We investigate the dynamics of relativistic spinning test particles in the spacetime of a rotating black hole using the Papapetrou equations. We use the method of Lyapunov exponents to determine whether the orbits exhibit sensitive dependence on initial conditions, a signature of chaos. In the case of maximally spinning equal-mass binaries (a limiting case that violates the test-particle approximation) we find unambiguous positive Lyapunov exponents that come in pairs $\pm\lambda$, a characteristic of Hamiltonian dynamical systems. We find no evidence for nonvanishing Lyapunov exponents for physically realistic spin parameters, which suggests that chaos may not manifest itself in the gravitational radiation of extreme mass-ratio binary black-hole inspirals (as detectable, for example, by LISA, the Laser Interferometer Space Antenna).

DOI: 10.1103/PhysRevD.67.024005

PACS number(s): 04.70.Bw, 04.80.Nn, 95.10.Fh

I. INTRODUCTION

The presence of chaos (or lack thereof) in relativistic binary inspiral systems has received intense attention recently due to the implications for gravitational-wave detection [1–8], especially regarding the generation of theoretical templates for use in matched filters. There is concern that the sensitive dependence to initial conditions that characterizes chaos may make the calculation of such templates difficult or impossible [8]. In particular, in the presence of chaos the number of templates would increase exponentially with the number of wave cycles to be fitted. In addition to this important concern, the problem of chaos in general relativity has inherent interest, as the dynamical behavior of general relativistic systems is poorly understood.

Several authors have reported the presence of chaos for systems of two point masses in which one or both particles are spinning [3,1,6]. Our work follows up on [3], which studies the dynamics of a spinning test particle orbiting a nonrotating (Schwarzschild) black hole using the Papapetrou equations [Eqs. (2.7)]. We extend this work to a rotating (Kerr) black hole, motivated by the expectation that many astrophysically relevant black holes have nonzero angular momentum. Furthermore, the potential for chaos may be greater in Kerr spacetime since the Kerr metric has less symmetry and hence fewer integrals of the motion than the Schwarzschild metric. In addition, the decision to focus on test particles is motivated partially by the Laser Interferometer Space Antenna (LISA) gravitational wave detector [9], which will be sensitive to radiation from spinning compact objects orbiting supermassive black holes in galactic nuclei. Using the Kerr metric is appropriate since such supermassive black holes will in general have nonzero spin.

There are many techniques for investigating chaos in dynamical systems, but for the case at hand we favor the use of Lyapunov exponents to quantify chaos. Informally, if ϵ_0 is the phase-space distance between two nearby initial conditions in phase space, then for chaotic systems the separation grows exponentially (*sensitive dependence on initial condi-*

tions): $\epsilon(\tau) = \epsilon_0 e^{\lambda\tau}$, where λ is the Lyapunov exponent. (See Sec. III A for a discussion of issues related to the choice of metric used to determine the distance in phase space.) The value of Lyapunov exponents lies not only in establishing chaos, but also in providing a characteristic time scale $\tau_\lambda = 1/\lambda$ for the exponential separation.

By definition, chaotic orbits are bounded phase space flows with at least one nonzero Lyapunov exponent. There are additional technical requirements for chaos that rule out periodic or quasiperiodic orbits, equilibria, and other types of patterned behavior [10]. For example, unstable circular orbits in Schwarzschild spacetime can have positive Lyapunov exponents [5], but such orbits are completely integrable (see Sec. VI) and hence not chaotic. In practice, we restrict ourselves to generic orbits, avoiding the specialized initial conditions that lead to positive Lyapunov exponents in the absence of chaos.

The use of Lyapunov exponents is potentially dangerous in general relativity because of the freedom to redefine the time coordinate. Chaos can seemingly be removed by a coordinate transformation: simply let $\tau' = \log \tau$ and the chaos disappears. Fortunately, in our case there is a fixed background spacetime with a time coordinate that is not dynamical but rather is simply a reparametrization of the proper time. As a result, we will not encounter this time coordinate redefinition ambiguity (which plagued, for example, attempts to establish chaos in mixmaster cosmological models, until coordinate-invariant methods were developed [11]). Furthermore, we can compare times in different coordinate systems using ratios: if t_p is the period of a periodic orbit in some coordinate system with time coordinate t , and τ_p is the period in proper time, then their ratio provides a conversion factor between times in different coordinate systems [5]:

$$\frac{t}{\tau} = \frac{t_p}{\tau_p}. \quad (1.1)$$

For chaotic orbits, which are not periodic, we use the average value of $dt/d\tau$ over the orbit, so that

$$\frac{t_\lambda}{\tau_\lambda} = \left\langle \frac{dt}{d\tau} \right\rangle \quad (1.2)$$

*Electronic address: mhartl@tapir.caltech.edu

as discussed in Sec. VII D. [This more general formula reduces to Eq. (1.1) in the case of periodic orbits.] Since we want to measure the local divergence of trajectories, the natural definition is to use the divergence in local Lorentz frames, which suggests that we use the proper time τ as our time parameter. The Lyapunov time scale in any coordinate system can then be obtained using Eq. (1.2).

Lyapunov exponents provide a quantitative definition of chaos, but there are several common qualitative methods as well, none of which we use in the present case, for reasons explained below. Perhaps the most common qualitative tool in the analysis of dynamical systems is the use of Poincaré surfaces of section. Poincaré sections reduce the phase space by one dimension by considering the intersection of the phase space trajectory with some fixed surface, typically taken to be a plane. Plotting momentum vs position for intersections of the trajectory with this surface then gives a qualitative view of the dynamics. As noted in [4], such sections are most useful when the number of degrees of freedom minus the number of constraints (including integrals of the motion) is not greater than two, since in this case the resulting points fall on a one-dimensional curve for nonchaotic orbits, but are “dusty” for chaotic orbits (and in the case of dissipative dynamical systems lie on fractal attractors). Unfortunately, the system we consider has too many degrees of freedom for Poincaré sections to be useful. It is possible to plot momentum vs position when the trajectory intersects a section that is a plane in physical space (say $x=0$) [3], but this is not in general a true Poincaré section.¹

Other qualitative methods include power spectra and chaotic attractors. The power spectra for regular orbits have a finite number of discrete frequencies, whereas their chaotic counterparts are continuous. Unfortunately, it is difficult to differentiate between complicated regular orbits, quasiperiodic orbits, and chaotic orbits, so we have avoided their use. Chaotic attractors, which typically involve orbits asymptotically attracted to a fractal structure, are powerful tools for exploring chaos, but their use is limited to dissipative systems [10]. Nondissipative systems, including test particles in general relativity, do not possess attractors [12].

Following Suzuki and Maeda [3], we use the Papapetrou equations to model the dynamics of a spinning test particle in the absence of gravitational radiation. We extend their work in a Schwarzschild background by considering orbits in Kerr spacetime, and we also improve on their methods for calculating Lyapunov exponents. The most significant improvement is the use of a rigorous method for determining Lyapunov exponents using the linearized equations of motion for each trajectory in phase space (Sec. III A), which requires knowledge of the Jacobian matrix for the Papape-

trou system (Sec. V B). We augment this method with an implementation of an informal deviation vector approach, which tracks the size of an initial deviation of size ϵ_0 and uses the relation $\epsilon(\tau) = \epsilon_0 e^{\lambda\tau}$ discussed above. We are careful in all cases to incorporate the constrained nature of the Papapetrou equations (Sec. II A) in the calculation of Lyapunov exponents (Sec. IV B).

We use units where $G=c=1$ and sign conventions as in Misner, Thorne, and Wheeler (MTW) [13]. We use vector arrows for 4-vectors (e.g., \vec{p} for the 4-momentum) and bold-face for Euclidean vectors (e.g., $\boldsymbol{\xi}$ for a Euclidean tangent vector). The symbol \log refers in all cases to the natural logarithm \log_e .

II. SPINNING TEST PARTICLES

A. Papapetrou equations

The Papapetrou equations [14] describe the motion of a spinning test particle. Although Papapetrou first derived the equations of motion for such a particle, the formulation by Dixon [15] is the starting point for most investigations because of its conceptual clarity. Dixon writes the equations of motion in terms of the 4-momentum p^α and spin tensor $S^{\alpha\beta}$, which are defined by integrals of the particle’s stress-energy tensor $T^{\alpha\beta}$ over an arbitrary spacelike hypersurface Σ :

$$p^\alpha(\Sigma) = \int_\Sigma T^{\alpha\beta} d\Sigma_\beta \quad (2.1)$$

$$S^{\alpha\beta}(\vec{z}, \Sigma) = 2 \int_\Sigma (x^{[\alpha} - z^{[\alpha} T^{\beta]\gamma} d\Sigma_\gamma, \quad (2.2)$$

where \vec{z} is the coordinate of the center of mass. The equations of motion for a spinning test particle are then

$$\frac{dx^\mu}{d\tau} = v^\mu$$

$$\nabla_{\vec{v}} p^\mu = -\frac{1}{2} R^\mu{}_{\nu\alpha\beta} v^\nu S^{\alpha\beta}$$

$$\nabla_{\vec{v}} S^{\mu\nu} = 2p^{[\mu} v^{\nu]}, \quad (2.3)$$

where v^μ is the 4-velocity, i.e., the tangent to the particle’s worldline. It is apparent that the 4-momentum deviates from geodesic motion due to a coupling of the spin to the Riemann curvature.

1. Spin supplementary conditions

As written, the Papapetrou equations (2.3) are underdetermined, and require a *spin supplementary condition* to determine the rest frame of the particle’s center of mass. Following Dixon, we choose

$$p_\mu S^{\mu\nu} = 0, \quad (2.4)$$

¹In [3], they are aided by the symmetry of Schwarzschild spacetime, which guarantees that one component of the spin tensor (Sec. II A below) is zero in the equatorial plane. As a result, it turns out that all but two of their variables are determined on the surface, and thus their sections are valid. Unfortunately, the reduced symmetry of the Kerr metric makes this method unsuitable for the system we consider in this paper.

which picks out a unique worldline that we identify as the center of mass. In particular, in the zero 3-momentum frame defined by $p^i=0$, applying Eq. (2.4) to Eq. (2.2) yields

$$z^i = \frac{\int_{t=\text{const}} x^i T^{00} d^3x}{\int_{t=\text{const}} T^{00} d^3x}, \quad (2.5)$$

which is the proper relativistic generalization of the Newtonian center of mass. The frame defined by $p^i=0$ is thus the rest frame of the center of mass, and in this frame Eq. (2.4) implies that $S^{0j}=0$, i.e., the spin is purely spatial in the rest frame.

A second possibility for the supplementary condition is

$$v_\mu S^{\mu\nu} = 0. \quad (2.6)$$

This condition has the disadvantage that it is satisfied by a family of helical worldlines filling a cylinder with frame-dependent radius [15,16], centered on the worldline picked out by condition (2.4). As a result, we adopt $p_\mu S^{\mu\nu}=0$ as the supplementary condition.

We note that the difference between the conditions (2.4) and (2.6) is third order in the spin [which follows from Eq. (2.13) below], which means that it is negligible for physically realistic spins (Sec. II B). In particular, the two conditions are equivalent for post-Newtonian expansions [17], where condition (2.6) is typically employed [18].

2. A reformulation of the equations

For numerical reasons, we use a form of the equations different from Eqs. (2.3). (We discuss this and other numerical considerations in Sec. V A.) Following the appendix in [3], we write the equations in terms of the momentum 1-form p_μ and the spin 1-form S_μ .² The system under consideration is a spinning particle of rest mass μ orbiting a central body of mass M ; in what follows, we measure all times and lengths in terms of M , and we measure the momentum of the particle in terms of μ , so that $p_\nu p^\nu = -1$. In these normalized units, the equations of motion are

$$\begin{aligned} \frac{dx^\mu}{d\tau} &= v^\mu \\ \nabla_{\dot{v}} p_\mu &= -R_{\mu\nu}^*{}^{\alpha\beta} v^\nu p_\alpha S_\beta \\ \nabla_{\dot{v}} S_\mu &= -p_\mu (R^*{}^{\alpha\beta\gamma\delta} S_\alpha v^\beta p_\gamma S_\delta) \end{aligned} \quad (2.7)$$

where

$$R^*{}^{\alpha\beta\mu\nu} = \frac{1}{2} R^\alpha{}_{\beta\rho\sigma} \epsilon^{\rho\sigma\mu\nu}. \quad (2.8)$$

²The lowered indices are motivated by the Hamiltonian formulation for a nonspinning test particle, where it is the one-form p_μ that is canonically conjugate to x^μ [13].

The tensor and vector formulations of the spin are related by

$$S_\mu = \frac{1}{2} \epsilon_{\mu\nu\alpha\beta} u^\nu S^{\alpha\beta} \quad (2.9)$$

and

$$S^{\mu\nu} = -\epsilon^{\mu\nu\alpha\beta} S_\alpha u_\beta, \quad (2.10)$$

where $u_\nu = p_\nu/\mu$ ($=p_\nu$ in normalized units). In addition, the spin satisfies the condition

$$S_\mu S^\mu = \frac{1}{2} S_{\mu\nu} S^{\mu\nu} = S^2, \quad (2.11)$$

where S is the spin of the particle measured in units of μM (see Sec. II B).

Because of the coupling of the spin to the Riemann curvature, the 4-momentum p^μ [Eq. (2.1)] is not parallel to the tangent v^μ . The supplementary condition (2.4) allows for an explicit solution for the difference between them (see [19] for a derivation):

$$v^\mu = N(p^\mu + w^\mu), \quad (2.12)$$

where

$$w^\mu = -{}^*R^{\mu\alpha\beta\gamma} S_\alpha p_\beta S_\gamma \quad (2.13)$$

and

$${}^*R^{\alpha\beta\mu\nu} = \frac{1}{2} R^{\alpha\beta\rho\sigma} \epsilon_{\rho\sigma}{}^{\mu\nu}. \quad (2.14)$$

The normalization constant N is fixed by the constraint $v_\mu v^\mu = -1$. We see from Eq. (2.13) that the difference between p^μ and v^μ is $\mathcal{O}(S^2)$, so that the difference between Eqs. (2.4) and (2.6) is $\mathcal{O}(S^3)$.

The spin 1-form satisfies two orthogonality constraints:

$$p^\mu S_\mu = 0 \quad (2.15)$$

and

$$v^\mu S_\mu = 0. \quad (2.16)$$

These two constraints are equivalent as long as v^μ is given by Eq. (2.12), since $w^\mu S_\mu \propto {}^*R^{\mu\alpha\beta\gamma} S_\mu S_\alpha \equiv 0$. When parameterizing the initial conditions, we enforce Eq. (2.15); since we use Eq. (2.12) in the equations of motion, Eq. (2.16) is then automatically satisfied.

3. Range of validity

We note that the Papapetrou equations include effects due only to the mass monopole and spin dipole (the pole-dipole approximation). In particular, the tidal coupling, which is a mass quadrupole effect, is neglected. It is also important to note that the Papapetrou equations are conservative and hence ignore the effects of gravitational radiation. For a thorough and accessible general discussion of the Papapetrou

equations and related matters, including a comprehensive literature review, see Semerák [19].

B. Comments on the spin parameter

It is crucial to note that, in our normalized units, the spin parameter S is measured in terms of μM , not μ^2 . The system we consider in this paper is a compact spinning body of mass μ orbiting a large body of mass M , which we take to be a supermassive Kerr black hole satisfying $M \approx 10^5 - 10^6 M_\odot$. We will show that physically realistic values of the spin must satisfy $S \ll 1$ for the compact objects (black holes, neutron stars, and white dwarfs) most relevant for the test particles described by the Papapetrou equations.³ The case of a black hole is simplest: a maximally spinning black hole of mass μ has spin angular momentum $s = \mu^2$, so a small black hole μ orbiting a large black hole of mass $M \gg \mu$ has a small spin parameter S :

$$S = \frac{s}{\mu M} \leq \frac{\mu^2}{\mu M} = \frac{\mu}{M} \ll 1.$$

The limit is similar for neutron stars: most models of neutron stars have a maximum spin of $s_{\max} \approx 0.6 \mu^2$ [20], which gives $S \leq 0.6 \mu/M$.

1. Bounds on S for stellar objects

The bound on S is relatively simple for black holes and neutron stars, but the situation is more complicated for compact stellar objects such as white dwarfs. The maximum spin of a stellar object is typically determined by the mass-shedding limit, i.e., the maximum spin before the star begins to break up. The spin in the case of the break-up limit is the moment of inertia times the maximum (break-up) angular velocity: $s_{\max} = I \Omega_{\max}$. If we write $I = \alpha \mu R^2$ and $\Omega_{\max} = \beta \sqrt{G \mu / R^3}$ for some constants $\alpha, \beta \leq 1$, then we have

$$s_{\max} = \alpha \beta (G \mu^3 R)^{1/2}. \quad (2.17)$$

The values of α and β depend on the stellar model; if we use the values for an $n=1.5$ polytrope, we get $\alpha = 0.2044$ and $\beta = 0.5366$ [21], so that $s_{\max} = 0.110 (G \mu^3 R)^{1/2}$.

The limit in Eq. (2.17) depends on the mass-radius relation for the object in question. Since most neutron stars have masses and radii in a narrow range, the estimate of $s_{\max} \approx 0.6 \mu^2$ discussed above is sufficient, but for white dwarfs the value of s_{\max} can depend strongly on the mass. An analytical approximation for the mass-radius relation for nonrotating white dwarfs is [22]⁴

$$\frac{R}{R_\odot} = 0.01125 \left(\frac{\mu}{\mu_{\max}} \right)^{-1/3} f(\mu)^{1/2} \quad (2.18)$$

³Recall that the Papapetrou equations ignore tidal coupling, so they are inappropriate for modeling more extended objects.

⁴The mean molecular weight $\bar{\mu}$ is set equal to 2, corresponding to helium and heavier elements, which is appropriate for most astrophysical white dwarfs.

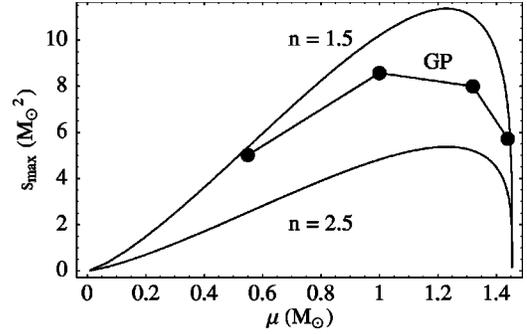


FIG. 1. The maximum spin angular momentum s_{\max} vs mass μ for a rigidly rotating white dwarf. We plot curves for $n=1.5$ and $n=2.5$ polytropic approximations using Eq. (2.22), together with four points derived using a more realistic numerical white dwarf model (Geroyannis and Pappasotiriou [23]).

where

$$f(\mu) = 1 - \left(\frac{\mu}{\mu_{\max}} \right)^{4/3} \quad (2.19)$$

and

$$\mu_{\max} = 1.454 M_\odot. \quad (2.20)$$

We could plug Eq. (2.18) into Eq. (2.17) to obtain an order-of-magnitude estimate, but [21] tabulates a constant \bar{J} equivalent to the product $\alpha\beta$ (which increases as the angular velocity of the star increases). They write $J = \bar{J} (GM^3 R_0)^{1/2}$ for a rotating white dwarf, where \bar{J} depends on the polytropic index n of a *nonspinning* white dwarf of the same mass, and R_0 is the nonspinning radius. In our notation, this reads

$$s_{\max} = \bar{J} (G \mu^3 R)^{1/2}. \quad (2.21)$$

White dwarfs with $\mu > 0.6 M_\odot$ are not well approximated by polytropes (the effective polytropic index varies from near 3 in the core to near 1.5 in the outer parts), but useful bounds can be obtained by substituting R from Eq. (2.18), which is more accurate for white dwarfs than a pure polytrope model. Plugging Eq. (2.18) into Eq. (2.21) and converting to geometric units gives

$$s_{\max} = 77.68 \bar{J} \mu^{4/3} M_\odot^{2/3} f(\mu)^{1/4}. \quad (2.22)$$

From Table 3 in [21], we have $\bar{J} = 0.1660$ for a maximally rotating $n=1.5$ polytrope (vs $\alpha\beta = 0.110$ for a slowly rotating one) and $\bar{J} = 0.0785$ for $n=2.5$. As illustrated in Fig. 1, the values for a more realistic numerical model [23] lie between these curves, as expected.

Note from Eq. (2.22) that $s_{\max}/\mu^2 \propto \mu^{-2/3}$ for $\mu \ll \mu_{\max}$, so that the spin per unit mass squared is unbounded as $\mu \rightarrow 0$.⁵ Nevertheless, the spin parameter S_{\max} is bounded, since

⁵Equation (2.22) is valid only for $\mu \geq 0.01 M_\odot$, but s_{\max}/μ^2 continues to increase with decreasing μ for equations of state appropriate for brown dwarfs and planets.

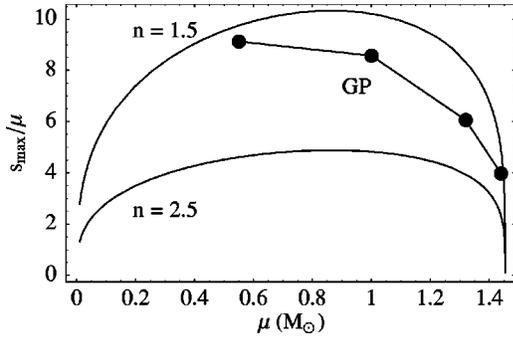


FIG. 2. s_{\max}/μ vs mass for a white dwarf. As in Fig. 1, we plot curves for $n=1.5$ and $n=2.5$ polytropes and the numerical model from [23]. The corresponding spin parameter S_{\max} is obtained simply by dividing s_{\max}/μ by the mass M of the central black hole.

$S_{\max} \propto s_{\max}/\mu \propto \mu^{1/3}$ in the low mass limit. We plot s_{\max}/μ vs μ in Fig. 2, which shows that the maximum value of s_{\max}/μ is approximately $9M_{\odot}$ (corresponding to a $\mu=0.5M_{\odot}$ white dwarf). For a central black hole of mass $M=10^6M_{\odot}$, we then have

$$S \leq S_{\max} = \frac{s_{\max}}{\mu M} = 9 \times 10^{-6}, \quad (2.23)$$

which is small compared to unity.

2. Tidal disruption

We can obtain a higher value of S if the central black hole mass is smaller, but it is important to bear in mind that such lower-mass black holes may tidally disrupt the white dwarf companion, thereby violating a necessary condition for the validity of the Papapetrou equations. In order of magnitude, a white dwarf orbiting at radius r will be disrupted when the tidal acceleration due to the central body overcomes its self-gravity, i.e.,

$$\frac{GM}{r^3} R \geq \frac{G\mu}{R^2}. \quad (2.24)$$

For the white dwarf to be undisrupted down to the horizon at $r=M$, we must have $M \leq R^{3/2} \mu^{-1/2}$, so that [using Eq. (2.18)] the minimum mass not to disrupt is $M_{\min} \propto \mu^{-1}$. We could evaluate the proportionality constant using Eq. (2.18), but we can obtain a more accurate result by adopting a constant based on a more realistic tidal disruption model. Tables I and II of [24] give the value of the variable $\hat{r} \equiv (r/R) \times (\mu/M)^{1/3}$, which is approximately 2.0 for the white dwarfs of interest here. This gives

$$M_{\min} = 2.0^{-3/2} R^{3/2} \mu^{-1/2}, \quad (2.25)$$

as illustrated in Fig. 3. For a $1.0M_{\odot}$ white dwarf, which (based on [23]) has $s_{\max}=8.57M_{\odot}^2$, the central black hole must satisfy $M_{\min}=8.2 \times 10^4 M_{\odot}$, so that the spin parameter S can be no bigger than $S_{\max}=s_{\max}/(\mu M_{\min})=1.0 \times 10^{-4}$ in order to avoid tidal disruption.

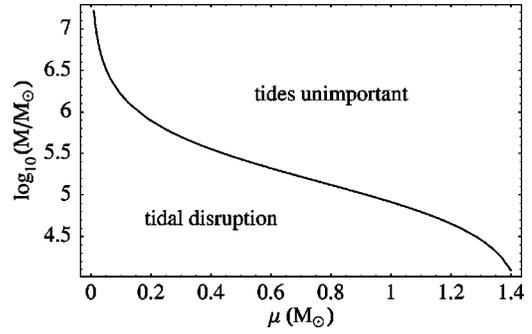


FIG. 3. The minimum black hole mass M required not to disrupt an inspiraling corotating white dwarf before the last stable (prograde) circular orbit around a maximally rotating Kerr black hole, as a function of white dwarf mass μ .

3. The $S=1$ limit

We have shown that all physically realistic cases satisfy $S \ll 1$, but we nevertheless consider the limit of $S=1$ (corresponding to $\mu=M$) in order to investigate more thoroughly the dynamics of the Papapetrou equations, and to compare our results with [3], which investigates the $S=1$ limit in detail. The $S=1$ limit introduces no singularities into the equations of motion, and the resulting orbits are valid solutions of the equations. On the other hand, in this limit the Papapetrou equations are not physically realistic, since they are derived in the limit of spinning test particles, which must satisfy $\mu \ll M$. We thus cannot draw reliable results about the behavior of astrophysical systems from the $S=1$ limit.

C. Symmetries and the parametrization of initial conditions

In the approximation represented by the Papapetrou equations there is still a constant of the motion associated with each Killing vector $\vec{\xi}$ of the spacetime [15]:

$$C_{\xi} = \xi^{\mu} p_{\mu} - \frac{1}{2} \xi_{\mu;\nu} S^{\mu\nu}. \quad (2.26)$$

[For brevity, we write the constant in terms of the spin tensor $S^{\mu\nu}$ [Eq. (2.10)].] Since Kerr spacetime is stationary and axially symmetric, it has the Killing vectors $\vec{\xi}^t = \partial/\partial t$ and $\vec{\xi}^{\phi} = \partial/\partial \phi$, so the energy E and z angular momentum J_z are conserved:

$$E = -p_t + \frac{1}{2} g_{t\mu,\nu} S^{\mu\nu} \quad (2.27)$$

and

$$J_z = p_{\phi} - \frac{1}{2} g_{\phi\mu,\nu} S^{\mu\nu}. \quad (2.28)$$

(We write J_z in place of the orbital angular momentum L_z since the spin also contributes to the angular momentum of the system.) In contrast to the energy and momentum integrals, the Carter constant Q is no longer present when the test particle has nonvanishing spin [25].

In our problem there are twelve variables, four each for position, momentum, and spin. For the purposes of finding orbits by numerical integration, we may parametrize the initial conditions by providing $\tau=0$, r , θ , $\phi=0$, p_r , E , J_z , S , and any two of the spin components. The normalization conditions $p_\mu p^\mu = -1$ and $S_\mu S^\mu = S^2$ allow us to eliminate one component each of momentum and spin. The constraint $p^\mu S_\mu = 0$ and the integrals of the motion then give three equations in three unknowns:

$$0 = p_\mu S_\nu g^{\mu\nu} \quad (2.29)$$

$$E = -p_t + \frac{1}{2} g_{t\mu,\nu} S^{\mu\nu} \quad (2.30)$$

$$J_z = p_\phi - \frac{1}{2} g_{\phi\mu,\nu} S^{\mu\nu}. \quad (2.31)$$

We must solve these equations for the two remaining components of p_μ and one remaining component of S_μ . In Schwarzschild spacetime these can be solved explicitly due to the greater symmetry [3], but in the Kerr case of interest here the problem requires numerical root finding.

We also use a related parametrization method starting with the Kerr geodesic orbital parameters: eccentricity e , inclination angle ι , and pericenter r_p . We derive the corresponding energy, angular momentum, and relevant momenta, and then proceed as above. This method is discussed further in Sec. VII A 3.

III. LYAPUNOV EXPONENTS

A. General discussion of Lyapunov exponents

Our method for calculating Lyapunov exponents is well-established in the literature of nonlinear dynamical systems [10,12], but accessible treatments are hard to find in the physics literature, so we summarize the method here. Our discussion is informal and oriented toward practical calculation, based on Ref. [10]; for a more formal, rigorous presentation see Eckmann and Ruelle [26].

First we give an overview of the methods for calculating Lyapunov exponents most commonly used in physics. Given an initial condition, a set of differential equations determines a solution (the *flow*), which is a curve in the phase space. The *Lyapunov exponents* of the flow measure the rate at which nearby trajectories separate. As discussed in the Introduction, an orbit is chaotic if a nearby phase-space trajectory separated by an initial distance ϵ_0 separates exponentially: $\epsilon(\tau) = \epsilon_0 e^{\lambda\tau}$, where λ is the Lyapunov exponent.

Implicit in the definition of chaos above is a notion of a distance function on the phase space [or, more properly, the tangent space to the phase space, as in Eq. (3.3) below]. It is conventional to use a Euclidean metric to define such lengths [10,12], but any positive-definite nondegenerate metric will do [26]. While the magnitude of the resulting exponent obviously depends on the particular metric used, the signs of the Lyapunov exponents are a property of the dynamical system and do not rely on any underlying metric structure. We discuss these issues further in Sec. IV A and Sec. VII D.

This informal definition of Lyapunov exponents leads to a practical method for calculating λ : given an initial condition, consider a nearby initial condition a distance ϵ_0 away, where ϵ_0 is “small,” typically 10^{-5} – 10^{-7} of the relevant physical scales. (Values of ϵ_0 much smaller than this can result in a loss of numerical precision.) Keeping track of the deviation vector between the two points yields a numerical approximation of λ . (It is important to rescale the deviation vector if it grows too large, since for any bounded phase space flow even a tiny deviation can grow to at most the size of the bounded region.) We call this approach the *deviation vector method*.

There are two primary limitations to the approach outlined above. First, the method yields only the largest Lyapunov exponent, which is sufficient to establish the presence of chaos but paints a limited picture of the dynamics. Second, the deviation vector approach is most appropriate when an analytical expression for the Jacobian matrix is unknown; by choosing ϵ_0 small enough [and by keeping $\epsilon(\tau)$ small by rescaling if necessary], the method essentially takes a numerical derivative. Among other complications, the value of the exponent depends both on the maximum allowable size ϵ_{\max} (the size at which the deviation is rescaled) and the initial value ϵ_0 (the size of the deviation after each rescaling).

The principal virtue of the deviation vector approach compared to the more complicated Jacobian method (discussed below) is speed, since it requires solving only the equations of motion. (As we discuss in Sec. III B 1, the Jacobian method involves the time-consuming evolution of the Jacobian matrix in parallel with the equations of motion.) It also provides a valuable way to verify the validity of the Jacobian method.

The *Jacobian method* is a more thorough and rigorous approach to the calculation of Lyapunov exponents, which makes precise the notion of “infinitesimally” separated vectors. The general method proceeds as follows: consider a phase space with variables $\mathbf{y} = \{y_i\}$ and an autonomous set of differential equations

$$\frac{d\mathbf{y}}{d\tau} = \mathbf{f}(\mathbf{y}). \quad (3.1)$$

(Here we use τ instead of t in anticipation of the application of these results to general relativity, where we will be using proper time as our time parameter.) If $\delta\mathbf{y}$ represents a small deviation vector, then the distance between the two trajectories is

$$\frac{d(\delta\mathbf{y})}{d\tau} = \mathbf{f}(\mathbf{y} + \delta\mathbf{y}) - \mathbf{f}(\mathbf{y}) = \mathbf{Df} \cdot \delta\mathbf{y} + \mathcal{O}(\|\delta\mathbf{y}\|^2), \quad (3.2)$$

where \mathbf{Df} is the Jacobian matrix $[(\mathbf{Df})_{ij} = \partial f_i / \partial x^j]$.

We can clarify the notation and make the system easier to visualize if we introduce ξ as an element of the tangent space at \mathbf{y} , so that

$$\frac{d\xi}{d\tau} = \mathbf{Df} \cdot \xi, \quad (3.3)$$

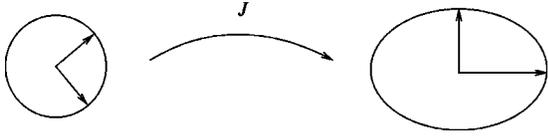


FIG. 4. The Jacobian matrix maps a tangent space ball to an ellipsoid.

which is equivalent to taking the limit $\|\delta \mathbf{y}\| \rightarrow 0$. We visualize ξ as a perfectly finite vector (as opposed to an “infinitesimal”). Since it lives in the tangent space, not the physical phase space, ξ can grow arbitrarily large with time. This means that instead of the frequent rescaling required in the deviation vector approach, ξ must be rescaled only when it grows so large that it approaches the floating point limit of the computer. This is a rare occurrence, and in practice the tangent vector almost never needs rescaling.

Although following the evolution of an arbitrary initial tangent vector ξ yields the largest Lyapunov exponent, we can do even better by following the evolution of a family of n tangent vectors, which allows us to determine all n Lyapunov exponents. The essence of the method is as follows: for a system of differential equations with n variables, we consider a set of n vectors that lie on a ball in the tangent space. We represent this ball using a matrix whose columns are n normalized, linearly independent tangent vectors, conventionally taken to be orthogonal. This set of orthonormal vectors then spans a unit ball in the tangent space. The action of the Jacobian matrix, which is a linear operator on the tangent space, is to map the ball to an ellipsoid under the time-evolution of the flow, as shown in Fig. 4.

For a dynamical system with n degrees of freedom, there are n Lyapunov numbers that measure the average growth of the n principal axes $\{r_i(\tau)\}_{i=1}^n$ of the ellipsoid. More formally, the Lyapunov numbers L_i are given by

$$L_i = \lim_{\tau \rightarrow \infty} [r_i(\tau)]^{1/\tau}, \quad (3.4)$$

where $r_i(\tau)$ is the length of the i th principal axis of the ellipsoid. The corresponding Lyapunov exponents are the natural logarithms of the Lyapunov numbers, so that

$$\lambda_i = \lim_{\tau \rightarrow \infty} \frac{\log[r_i(\tau)]}{\tau}. \quad (3.5)$$

These limits exist for a broad class of dynamical systems [26].

The principal axes of the tangent space ellipsoid indicate the directions along which nearby initial conditions separate or converge, which we may call the *Lyapunov directions*. In particular, consider a principal axis that is stretched under the time evolution. Such a vector has one component for each dimension (position or momentum) in the phase space; a nonzero component in any direction indicates an exponential divergence in the corresponding coordinate. For example, if a system has two spatial coordinates (r, ϕ) and corresponding momenta (p_r, p_ϕ) , then a typical tangent vector will have components $\xi = (\xi_r, \xi_\phi, \xi_{p_r}, \xi_{p_\phi})$. If the only tangent

vector with nonzero Lyapunov exponent is, for example, $\xi = (1/\sqrt{3})(1, 1, 1, 0)$, then nearby initial conditions separate exponentially in r , ϕ , and p_r , but nearby values of p_ϕ do not separate exponentially. This is potentially relevant to the present study since, in the limit of a point test particle, the gravitational radiation depends on the spatial variables but not the spin. If the principal axes along expanding directions have nonzero components only in the spin directions, the system could be formally chaotic without affecting the gravitational waves.

In summary, the method for visualizing the Lyapunov exponents of a dynamical system is to picture a ball of initial conditions—an infinitesimal ball if visualized in the phase space, or a unit ball if visualized in the tangent space—and watch it evolve into an ellipsoid under the action of the Jacobian matrix. After a sufficiently long time, the ellipsoid will be greatly deformed, stretched out along the expanding directions and compressed along the contracting directions. The directions of the principal axes are the Lyapunov directions, and their lengths give the Lyapunov numbers through the relation $L_i \approx [r_i(\tau)]^{1/\tau}$.

B. Numerical calculation of Lyapunov exponents

In order to implement a numerical algorithm based on the considerations above, we must bear two things in mind. First, since the vectors spanning the initial unit ball are arbitrary, they will all be stretched in the direction of the largest exponent: in general every initial vector has some nonzero component along the direction of greatest stretching, which dominates as $\tau \rightarrow \infty$. In order to find the other principal axes, we must periodically produce a new orthogonal basis. We will show that the Gram-Schmidt procedure is appropriate. Second, the lengths of the vectors could potentially overflow or underflow the machine precision, so we should periodically normalize the ellipsoid axes.

1. The algorithm in detail

To simplify the notation, we denote the (time-dependent) Jacobian matrix \mathbf{Df} by J_τ and the ellipsoid (whose columns are the tangent vectors) by U . The algorithm then proceeds as follows.

(i) Construct a set of n orthonormal vectors (which span an n -dimensional ball in the tangent space of the flow). Represent this ball by a matrix U whose columns are the tangent vectors ξ_i .

(ii) Equation (3.3), applied to each tangent vector, implies that U satisfies the matrix equation

$$\frac{dU}{d\tau} = J_\tau U, \quad (3.6)$$

which constitutes a set of *linear* differential equations for the tangent vectors. Since J_τ depends on the values of \mathbf{y} , these equations are coupled to our system of nonlinear differential equations $\dot{\mathbf{y}} = \mathbf{f}(\mathbf{y})$, so they must be solved in parallel with Eq. (3.1).

(iii) Choose some time T big enough to allow the expanding directions to grow but small enough so that they are not

too big. Numerically integrate Eqs. (3.1) and (3.6), and every time T apply the Gram-Schmidt orthogonalization procedure. The vectors resulting from the Gram-Schmidt procedure approximate the semiaxes of the evolving ellipsoid. Record the log of the length $\log[r_i(\tau_n)]$ of each vector after each time T , where $\tau_n = nT$. Finally, normalize the ellipsoid back to a unit ball.

(iv) At each time τ , the sum

$$\lambda_i \approx \frac{1}{\tau} \sum_{n=1}^N \log[r_i(\tau_n)] \equiv \frac{\log[r_i(\tau)]}{\tau} \quad (3.7)$$

is a numerical estimate for the i th Lyapunov exponent.

2. Gram-Schmidt and Lyapunov exponents

The use of the Gram-Schmidt procedure is crucial to extracting all n Lyapunov exponents. Let us briefly review this important construction. Given n linearly independent vectors $\{\mathbf{u}_i\}$, the Gram-Schmidt procedure constructs n orthogonal vectors $\{\mathbf{v}_i\}$ that span the same space, given by

$$\mathbf{v}_i = \mathbf{u}_i - \sum_{j=1}^{i-1} \frac{\mathbf{u}_i \cdot \mathbf{v}_j}{\|\mathbf{v}_j\|^2} \mathbf{v}_j. \quad (3.8)$$

To construct the i th orthogonal vector, we take the i th vector from the original set and subtract off its projections onto the previous $i-1$ vectors produced by the procedure.

The use of the Gram-Schmidt procedure in dynamics comes from observing that the resulting vectors approximate the semiaxes of the tangent space ellipsoid. After the first time T , all of the vectors point mostly along the principal expanding direction. We may therefore pick any one as the first vector in the Gram-Schmidt algorithm, so choose $\xi_1 \equiv \mathbf{u}_1$ without loss of generality. If we let \mathbf{e}_i denote unit vectors along the principal axes and let r_i be the lengths of those axes, the dynamics of the system guarantees that the first vector \mathbf{u}_1 satisfies

$$\mathbf{u}_1 = r_1 \mathbf{e}_1 + r_2 \mathbf{e}_2 + \dots \approx r_1 \mathbf{e}_1 \equiv \mathbf{v}_1$$

since \mathbf{e}_1 is the direction of fastest stretching. The second vector \mathbf{v}_2 given by Gram-Schmidt is then

$$\mathbf{v}_2 = \mathbf{u}_1 - \frac{\mathbf{u}_1 \cdot \mathbf{v}_1}{\|\mathbf{v}_1\|^2} \mathbf{v}_1 \approx \mathbf{u}_1 - r_1 \mathbf{e}_1 = r_2 \mathbf{e}_2,$$

with an error of order r_2/r_1 . The procedure proceeds iteratively, with each successive Gram-Schmidt step (approximately) subtracting off the contribution due to the previous semiaxis direction.

It is important to choose time T long enough to keep errors of the form r_2/r_1 small but short enough to prevent numerical under- or overflow. In practice, the method is quite robust, and it is easy to find valid choices for the time T , as discussed in Sec. VII.

IV. RELATIVITY AND PAPAPETROU SUBTLITIES

The algorithm described above is of a general nature, designed with a generic dynamical system in mind. The Papapetrou equations and the framework of general relativity present additional complications. Here we discuss some refinements to the algorithm necessary for the present case.

A. Phase space norm

In the context of general relativistic dynamical systems, the meaning of trajectory separation in phase space is somewhat obscured by the time variable. We can skirt the issue of trajectories “diverging in time” by using a 3+1 splitting of spacetime, and consider trajectory separation in a spacelike hypersurface [27]. This prescription reduces properly to the traditional method for classical dynamical systems in the nonrelativistic limit.

In Kerr spacetime, we use the zero angular-momentum observers (ZAMOs), and project 4-dimensional quantities into the ZAMO hypersurface using the projection tensor $P^\mu_\nu = \delta^\mu_\nu + U^\mu U_\nu$, where U^μ is the ZAMO 4-velocity. In this formulation, spatial variables obey $x^\mu \rightarrow \tilde{x}^i = P^i_\mu x^\mu$ and momenta obey $p_\mu \rightarrow \tilde{p}_i = P^\mu_i p_\mu$ (and similarly for S_μ) [27]. The relevant norm is then a Euclidean distance in the 3-dimensional hypersurface.

We should note that we use the projected norm for conceptual clarity, not necessity. The naive use of a Euclidean norm using unprojected components yields the same sign for the exponents, as noted in Sec. III A. The magnitudes of the resulting exponents are also similar (Sec. VII D).

B. Constraint complications

Although the Lyapunov algorithm is fairly straightforward to implement for a general dynamical system, the constrained nature of the Papapetrou equations adds a considerable amount of complexity. The fundamental problem is that the tangent vector ξ cannot have arbitrary initial components for the Papapetrou system, as it can for an unconstrained dynamical system. Each ξ must correspond to some deviation $\delta \mathbf{y}$ which is not arbitrary: the deviated point $\mathbf{y} + \delta \mathbf{y}$ must satisfy the constraints.

1. Constraint-satisfying deviations

Recall that the dynamical variables in the Papapetrou equations must satisfy normalization and orthogonality constraints (Sec. II A): $p^\nu p_\nu = -1$ (normalized units), $S^\nu S_\nu = S^2$, and $p^\nu S_\nu = 0$. To make the constraint condition on $\delta \mathbf{y}$ clearer, let $\mathbf{C}(\mathbf{y}) = 0$ represent the constraints rearranged so that the right hand side is zero. For example, with $\mathbf{y} = (t, r, \mu, \phi, p_t, p_r, p_\mu, p_\phi, S_t, S_r, S_\mu, S_\phi)$,⁶ we can write

$$C_1(\mathbf{y}) = p^\nu p_\nu + 1, \quad (4.1)$$

⁶Recall that we write the equations of motion in terms of $\mu = \cos \theta$.

so that $C_1(\mathbf{y})=0$ for a constraint-satisfying \mathbf{y} . The other constraints are then

$$C_2(\mathbf{y})=S^v S_v - S^2 \quad (4.2)$$

and

$$C_3(\mathbf{y})=p^v S_v. \quad (4.3)$$

A deviation $\delta\mathbf{y}$ is *constraint-satisfying* if $\mathbf{C}(\mathbf{y} + \delta\mathbf{y}) = \mathbf{0}$ when $\mathbf{C}(\mathbf{y}) = \mathbf{0}$.

We may construct a constraint-satisfying deviation $\delta\mathbf{y}$ as follows. Begin with a 12-dimensional vector \mathbf{y} that satisfies the constraints. Add a random small deviation to eight of its components to form a new vector \mathbf{y}' . (We need not add a deviation to t ; see Sec. IV B 2 below.) Determine the remaining three components of \mathbf{y}' using the constraints, using the same technique used to set the initial conditions. Finally, set $\delta\mathbf{y} \equiv \mathbf{y}' - \mathbf{y}$. The corresponding ξ is then simply $\delta\mathbf{y}/\|\delta\mathbf{y}\|$.

The prescription above glosses over an important detail: the inference of tangent vector components from the constraints is not unique. Solving the constraint equations involves taking square roots in several places, so there are a number of sign ambiguities representing different solution branches. The implementation of the component-inference algorithm must compare each component of \mathbf{y} with the corresponding component of \mathbf{y}' to ensure that they represent solutions from the same branches. Enforcing the constraints in this manner, and thereby inferring the full tangent vector ξ , is especially important for the algorithm described in the next section.

2. A modified Gram-Schmidt algorithm

A spinning test particle has an apparent twelve degrees of freedom—four each for position, momentum, and spin—so *a priori* there is the potential for twelve nonzero exponents. Since the Papapetrou equations have no explicit time-dependence, we can eliminate the time degree of freedom. The three constraints (momentum and spin normalization, and momentum-spin orthogonality) further reduce the number of degrees of freedom by three. We are left finally with eight degrees of freedom.

Eliminating the four spurious degrees of freedom from the tangent vectors presents a formidable obstacle to the implementation of the phase space ellipsoid method described in Sec. III B 1. The crux of the dilemma is that the axes of the ellipsoid must be orthogonal, but must also correspond to constraint-satisfying deviation vectors—mutually exclusive conditions. Solving this problem requires a modification of the Gram-Schmidt algorithm.

(i) Instead of a 12×12 ball (i.e., $n=12$ in the original algorithm), consider an 8×8 ball by choosing to eliminate the t , p_t , p_ϕ , and S_t components. The time component ξ_t of each tangent vector is irrelevant since nothing in the problem is explicitly time dependent; the first column of the Jacobian matrix is zero, so ξ_t is not necessary to determine the

time-evolution.⁷ The other three components are determined by the constraints as described above.

(ii) Given eight initial random tangent vectors, apply the Gram-Schmidt process to form an 8×8 ball. For each vector, determine the three missing components using the constraints, and then evolve the system using

$$\frac{dU}{d\tau} = J_\tau U$$

as before. (Now U represents a 12×8 matrix instead of a 12×12 ball.)

(iii) At each time T , extract the relevant eight components from each vector to form a new 8×8 ellipsoid, apply the Gram-Schmidt process, and then fill in the missing components using the constraints, yielding again a 12×8 matrix. The projected norms of the eight tangent vectors contribute to the running sums for the Lyapunov exponents as in the original algorithm.

The algorithm above yields eight Lyapunov exponents for the Papapetrou system of equations.

In order to implement this algorithm, we must have a method for constructing a full tangent vector ξ from an eight-component vector $\tilde{\xi}$. The method is as follows.

(i) Let $\tilde{\mathbf{y}}' = \mathbf{y} + \epsilon \tilde{\xi}$ for a suitable choice of ϵ .

(ii) Fill in the missing components of $\tilde{\mathbf{y}}'$ using the constraints to form \mathbf{y}' , taking care that \mathbf{y} and \mathbf{y}' have the same constraint branches.

(iii) Infer the full tangent vector using $\xi = (\mathbf{y}' - \mathbf{y})/\epsilon$.

This technique depends on the choice of ϵ , and fails when ϵ is too small or too large. Using the techniques discussed in the next section to calibrate the system, we find that $\epsilon \approx 10^{-5} - 10^{-6}$ works well in practice.

3. Two rigorous techniques

It should be clear from the discussion above that extracting all eight Lyapunov exponents is difficult, and in practice the techniques are finicky, depending (among other things) on the choice of ϵ as described in Sec. IV B 2 above. How, then, can we be confident that the results make sense? Fortunately, there are two techniques that give rigorous Lyapunov exponents by managing to sidestep the constraint complexities entirely.

First, it is always possible to calculate the single largest exponent using the Jacobian method without considering the constraint subtleties. The complexity of the main Jacobian approach involves the competing requirements of Gram-Schmidt orthogonality and constraint satisfaction, but in the case of only one vector these difficulties vanish. Since the equations of motion preserve the constraints, an initial constraint-satisfying tangent vector retains this property throughout the integration. Thus, we begin with a vector constructed as in Sec. IV B 1 and evolve it (without rescaling) along with the equations of motion. Other than the require-

⁷Also, the time piece is discarded in the projected norm formalism in any case (Sec. IV A).

ment of constraint satisfaction, its initial components are arbitrary, so it evolves in the direction of largest stretching and eventually points in the largest Lyapunov direction. The logarithm of its projected norm then contributes to the sum for the largest Lyapunov exponent.

Second, we can implement a deviation vector approach as described in Sec. III A. Given an initial condition \mathbf{y}_0 , we construct a nearby initial condition \mathbf{y}'_0 as in Sec. IV B 1 and then evolve them both forward. In principle, an approximation for the largest Lyapunov exponent is then

$$\frac{1}{\tau} \log \left(\frac{\|\mathbf{y}' - \mathbf{y}\|}{\|\mathbf{y}'_0 - \mathbf{y}_0\|} \right) \equiv \frac{1}{\tau} \log \left(\frac{\|\delta \mathbf{y}\|}{\|\delta \mathbf{y}_0\|} \right).$$

In practice (for chaotic systems) the method saturates: for a given initial deviation, say $\|\delta \mathbf{y}_0\| \sim 10^{-6}$, once the initial conditions have diverged by a factor of $\sim 10^6$ the method breaks down.⁸ (The traditional solution to the saturation problem is to rescale the deviation before it saturates, but such a rescaling in this case violates the constraints.) Despite its limitations, this unrescaled deviation vector technique is valuable, since it tracks the correct solution until the saturation limit is reached, and avoids the subtleties associated with the constraints.

With these two techniques in hand, we have a powerful method for verifying that the largest Lyapunov exponent produced by the Gram-Schmidt method is correct. This, in turn, gives us confidence that the other Lyapunov exponents produced by the main algorithm are meaningful as well.

V. IMPLEMENTATION DETAILS

A. Some numerical comments

Finally, we discuss some specialized issues related to integrating the Papapetrou equations on a computer. The primary subjects are the formulation of the equations, optimization techniques, and error checking.

Our choice to write the Papapetrou equations using the spin vector is motivated partially by numerical considerations. The spin vector approach has nice properties compared to the tensor approach as $S \rightarrow 0$. Comparing their covariant derivatives is instructive:

$$\nabla_{\nu} S_{\mu} = -p_{\mu} (R^{*\alpha}_{\beta}{}^{\gamma\delta} S_{\alpha} v^{\beta} p_{\gamma} S_{\delta})$$

$$\nabla_{\nu} S^{\mu\nu} = p^{\mu} v^{\nu} - p^{\nu} v^{\mu} = 2p^{[\mu} v^{\nu]}.$$

Though simpler in form, the derivative of $S^{\mu\nu}$ has unfortunate numerical properties for small S , since in the limit $S \rightarrow 0$ we have $p^{\mu} \rightarrow v^{\mu}$: the difference $p^{\mu} v^{\nu} - p^{\nu} v^{\mu}$ goes to zero in principle but in practice is plagued by numerical

roundoff errors. Since $S \ll 1$ is the most physically interesting limit, the vector approach is more convenient for our purposes.

Calculating the many tensors and derivatives which go into the Papapetrou equations and the corresponding Jacobian matrix is a considerable task. As a first step, we use GRTENSOR for MAPLE to calculate all relevant quantities, and we use MAPLE'S optimized C output to create C code automatically. Due to the symmetries of the Riemann tensor and the metric, many terms are identically zero, which significantly reduces the number of required operations. For example, in order to calculate $R^{*\alpha}_{\beta}{}^{\gamma\delta} S_{\alpha} v^{\beta} p_{\gamma} S_{\delta}$ we need four loops, which constitutes $4^4 = 256$ evaluations, but in fact $R^{*\alpha}_{\beta}{}^{\gamma\delta}$ has only 80 nonzero components. Performing loop unrolling by writing these terms to an optimized derivatives file consisting of explicit sums speeds up calculation by an order of magnitude compared to nested for loops.

Another optimization involves the choice of coordinates used in the metric, which has significant consequences for the size of the tensor files and the number of floating point operations required. Simply using $\mu = \cos \theta$ in the Kerr metric reduces the size of the Riemann derivatives by at least a factor of 2.⁹ Since these derivatives are the bottleneck in the calculation of the Jacobian matrix, we can get more than a 50% improvement in performance with even this simple variable transformation.

All integrations were performed using a Bulirsch-Stoer integrator adapted from *Numerical Recipes* [28]. Occasional checks with a fifth-order Runge-Kutta integrator were in agreement. We verified the Papapetrou integration by checking errors in the constraints and conserved quantities; for an orbit such as that shown in Fig. 6, all errors are at the 10^{-11} level after $\tau = 10^5 M$.

As should be clear from Sec. V B below, the Jacobian matrix of the Papapetrou equations has a large number of terms, and it is essential to verify its correctness by using a diagnostic that compares $\mathbf{Df} \cdot \delta \mathbf{y}$ with the difference $\mathbf{f}(\mathbf{y} + \delta \mathbf{y}) - \mathbf{f}(\mathbf{y})$ for a suitable constraint-satisfying $\delta \mathbf{y}$. It is not sufficient for the difference merely to be small: we must calculate the quantity $\mathbf{f}(\mathbf{y} + \delta \mathbf{y}) - \mathbf{f}(\mathbf{y}) - \mathbf{Df} \cdot \delta \mathbf{y}$ for several values of $\delta \mathbf{y}$ and verify that each component scales as $\|\delta \mathbf{y}\|^2$. An early implementation of the Jacobian matrix, which gave nearly identical results for $\mathbf{f}(\mathbf{y} + \delta \mathbf{y}) - \mathbf{f}(\mathbf{y})$ and $\mathbf{Df} \cdot \delta \mathbf{y}$, nevertheless had an undetected $\mathcal{O}(S^2)$ error. The unrescaled deviation vector approach showed a discrepancy with the Jacobian method,¹⁰ which showed spurious chaotic behavior. The $\|\delta \mathbf{y}\|^2$ scaling method described above eventually diagnosed the problem, which resulted from a missing term in $\partial \dot{S}_{\mu} / \partial S_{\nu}$ (Sec. V B).

⁹Warning: This variable substitution changes the handedness of the coordinate system, since the unit vector $\hat{\mu}$ points opposite to $\hat{\theta}$. This in turn introduces an extra minus sign in the Levi-Civita tensor $\epsilon^{\alpha\beta\gamma\delta}$, which appears many times in the Papapetrou equations and the corresponding conserved quantities. The author discovered this subtlety the hard way.

¹⁰This illustrates the value of calculating the Lyapunov exponents using two different methods.

⁸This underscores the point that chaos is essentially a *local* phenomenon. Any unrescaled deviation vector approach must saturate, since no bounded system can have trajectories that diverge for arbitrarily long times.

B. The Jacobian matrix

For reference, we write out explicit equations for part of the Jacobian matrix of the Papapetrou equations.

The Jacobian matrix of a system of differential equations, specialized to the case at hand, is as follows:

$$\begin{pmatrix} \frac{\partial \dot{x}^\mu}{\partial x^\nu} & \frac{\partial \dot{x}^\mu}{\partial p_\nu} & \frac{\partial \dot{x}^\mu}{\partial S_\nu} \\ \frac{\partial \dot{p}_\mu}{\partial x^\nu} & \frac{\partial \dot{p}_\mu}{\partial p_\nu} & \frac{\partial \dot{p}_\mu}{\partial S_\nu} \\ \frac{\partial \dot{S}_\mu}{\partial x^\nu} & \frac{\partial \dot{S}_\mu}{\partial p_\nu} & \frac{\partial \dot{S}_\mu}{\partial S_\nu} \end{pmatrix}. \quad (5.1)$$

Once we calculate $\partial \dot{x}^\mu / \partial x^\nu = v^\mu$, all the other derivatives can be expressed in terms of the derivatives of v^μ , the tensors and connection coefficients, and Kronecker δ 's.

Written out in full, the Papapetrou equations are as follows:

$$\dot{x}^\mu = v^\mu \quad (5.2)$$

$$\dot{p}_\mu = -R_{\mu\nu}^* \alpha\beta v^\nu p_\alpha S_\beta + \Gamma_{\beta\mu}^\alpha p_\alpha v^\beta \quad (5.3)$$

$$\dot{S}_\mu = -p_\mu (R^* \alpha\beta\gamma\delta S_\alpha v^\beta p_\gamma S_\delta) + \Gamma_{\beta\mu}^\alpha S_\alpha v^\beta. \quad (5.4)$$

We measure τ and r in units of M (the mass of the central body), p_μ in units of the particle rest mass μ , and S_μ in terms of the product μM . The overdot is an ordinary derivative with respect to proper time: $\dot{x} \equiv dx/d\tau$.

The unusual placement of indices on R^* is motivated by the form of the Jacobian matrix. The index placement shown above brings the equations into a form where the indices on p_μ and S_μ are always lowered, which simplifies the Jacobian matrix since (for example) $\partial p_\mu / \partial x^\mu = 0$. Otherwise the Jacobian matrix is unnecessarily complicated; for example, if p^μ appeared anywhere on the right hand side then we would have $\partial p^\mu / \partial x^\nu \neq 0$, which would contribute to J_τ .

As discussed in Sec. II A, the supplementary condition $p_\mu S^{\mu\nu} = 0$ [Eq. (2.4)] leads to the equation for v^μ in terms of p^μ :

$$v^\mu = N(p^\mu + w^\mu) = N\tilde{v}^\mu, \quad (5.5)$$

where

$$\tilde{v}^\mu = p^\mu + w^\mu \quad (5.6)$$

and

$$w^\mu = - *R^* \mu\alpha\beta\gamma S_\alpha p_\beta S_\gamma. \quad (5.7)$$

N is a normalization constant fixed by $v_\mu v^\mu = -1$.

The calculation of the partial derivatives \dot{x}^μ in Eq. (5.1) proceeds as follows. From the relation for $v^\mu = N\tilde{v}^\mu$, we have

$$\frac{\partial \dot{x}^\mu}{\partial x^\nu} = v^\mu{}_{,\nu} = N\tilde{v}^\mu{}_{,\nu} + N_{,\nu}\tilde{v}^\mu.$$

Now, $\tilde{v}^\mu{}_{,\nu} = p^\mu{}_{,\nu} + w^\mu{}_{,\nu} = p_\alpha g^{\alpha\mu}{}_{,\nu} - *R^* \mu\alpha\beta\gamma{}_{,\nu} S_\alpha p_\beta S_\gamma$, so the first term is easy. The second term is trickier: from the expression for v^μ , we have that $-1 = v^\mu v_\mu = N^2(p^\mu p_\mu + 2w^\mu p_\mu + w^\mu w_\mu) = N^2(-1 + 2w^\mu p_\mu + w^\mu w_\mu)$, so we have

$$N = (1 - 2w^\mu p_\mu - w^\mu w_\mu)^{-1/2}.$$

Differentiating gives

$$\begin{aligned} N_{,\nu} &= N^3(p_\alpha w^\alpha{}_{,\nu} + w^\alpha{}_{,\nu} w_\alpha + \frac{1}{2} w^\alpha w^\beta g_{\alpha\beta,\nu}) \\ &= N^3(\tilde{v}_\alpha w^\alpha{}_{,\nu} + \frac{1}{2} w^\alpha w^\beta g_{\alpha\beta,\nu}) \end{aligned}$$

where we have relabeled the dummy index ($\mu \rightarrow \alpha$). Summing the various terms, we have

$$v^\mu{}_{,\nu} = N[p_\alpha g^{\alpha\mu}{}_{,\nu} + w^\mu{}_{,\nu} + v^\mu(v_\alpha w^\alpha{}_{,\nu} + \frac{1}{2} N w^\alpha w^\beta g_{\alpha\beta,\nu})]. \quad (5.8)$$

The expression for $\partial \dot{x}^\mu / \partial p_\nu$ is similar to $v^\mu{}_{,\nu}$, but it is simpler because the derivative of the metric with respect to the momentum is zero. As before, we use the product rule:

$$\frac{\partial v^\mu}{\partial p_\nu} = N \frac{\partial \tilde{v}^\mu}{\partial p_\nu} + \frac{\partial N}{\partial p_\nu} \tilde{v}^\mu.$$

The first term requires

$$\begin{aligned} \frac{\partial \tilde{v}^\mu}{\partial p_\nu} &= \frac{\partial p^\mu}{\partial p_\nu} + \frac{\partial w^\mu}{\partial p_\nu} = g^{\mu\nu} - *R^* \mu\alpha\nu\beta S_\alpha S_\beta \\ &\equiv g^{\mu\nu} + W^{\mu\nu}. \end{aligned}$$

Note that $W^{\mu\nu}$ is symmetric. The second term requires

$$\begin{aligned} \frac{\partial N}{\partial p_\nu} &= N^3(W^{\alpha\nu} p_\alpha + w^\alpha \delta_\alpha{}^\nu + W^{\alpha\nu} w_\alpha) \\ &= N^3(w^\nu + \tilde{v}^\alpha W^{\alpha\nu}). \end{aligned}$$

Summing the terms gives

$$\frac{\partial \dot{x}^\mu}{\partial p_\nu} = N(g^{\mu\nu} + W^{\mu\nu} + N v^\mu w^\nu) + N v^\mu v_\alpha W^{\alpha\nu} \quad (5.9)$$

with

$$\frac{\partial w^\mu}{\partial p_\nu} \equiv W^{\mu\nu} = - *R^* \mu\alpha\nu\beta S_\alpha S_\beta. \quad (5.10)$$

Finally, we calculate $\partial \dot{x}^\mu / \partial S_\nu$. With

$$\frac{\partial \tilde{v}^\mu}{\partial S_\nu} = -S_\alpha p_\beta (*R^* \mu\alpha\beta\nu - *R^* \mu\nu\alpha\beta) \equiv V^{\mu\nu},$$

and

$$\frac{\partial N}{\partial S_\nu} \tilde{v}^\mu = N v^\mu v_\alpha V^{\alpha\nu},$$

we have

$$\frac{\partial \dot{x}^\mu}{\partial S_\nu} = N V^{\mu\nu} + N v^\mu v_\alpha V^{\alpha\nu}. \quad (5.11)$$

We calculate the derivatives of \dot{p}_μ and \dot{S}_μ using $v^\mu_{,\nu}$, the product rule, and the derivatives of the various tensors in the problem. The full results appear in Appendix A.

VI. INTEGRABILITY AND CHAOS

A. Phase space and constants of the motion

Having laid the foundation for the numerical calculation of Lyapunov exponents, we now discuss some general aspects of dynamical systems relevant to our study. A dynamical system with n coordinates has a $2n$ dimensional phase space, typically consisting of generalized positions and their corresponding conjugate momenta. Motion in the phase space is arbitrary in general, but when there are integrals of the motion then the flow is confined to a surface on which the integral is constant. This can be seen most easily by transforming to angle-action coordinates, where the surface is an invariant (multidimensional) torus.

A system with n coordinates and n constants of the motion is *integrable* and cannot have chaos (though the motion can still be quasiperiodic or exhibit other complicated behavior). For example, we can consider geodesic orbits around a Kerr black hole to have eight degrees of freedom ($n=4$) and four constants of the motion—particle rest mass μ , energy E , axial or z angular momentum L_z , and Carter constant Q —which are enough to integrate the equations of motion explicitly. Alternatively, we may look at Kerr spacetime as having a 6-dimensional phase space by eliminating time (which is simply a reparametrization of the proper time) and using rest mass conservation to eliminate one momentum coordinate. Then the three integrals E , L_z , and Q are sufficient to integrate the motion. (In practice, we allow all four

momenta to evolve freely; the normalization is then a constraint which can be checked for consistency at the end of the integration.)

In the case of a spinning test particle, the extra spin degrees of freedom create the possibility for chaotic behavior. Moreover, since Q is not conserved in the case of nonzero spin, even without the extra spin degrees of freedom the potential for chaos would exist. Kerr spacetime has just enough constants to make the system integrable; losing Q reduces the number of analytic integrals below the critical level required to guarantee integrability.¹¹

B. Hamiltonian systems

1. Lyapunov exponents for Hamiltonian flows

The phase space flow of Hamiltonian systems is constrained by more than the integrals of the motion. In particular, the Lyapunov exponents of a Hamiltonian system come in pairs $\pm\lambda$; i.e., if λ is a Lyapunov exponent then so is $-\lambda$ [26]. Geometrically, this means that if one semimajor axis of the phase-space ellipsoid stretches an amount $e^{\lambda\tau} = L$, another axis must shrink by an amount $e^{-\lambda\tau} = 1/L$. One consequence of this property is that the product of the lengths of the axes is 1. Since the ellipsoid volume is proportional to this invariant product, Liouville's theorem on the conservation of phase space volume follows as a corollary.

The $\pm\lambda$ property of Hamiltonian flows results from the symplectic nature of the Jacobian matrix for Hamiltonian dynamical systems.¹² But a naive analysis of the Jacobian matrix of the Papapetrou equations shows that it is not symplectic in the canonical sense. Nevertheless, the Papapetrou equations can be derived from a Lagrangian [30], and can be cast in Hamiltonian form by use of a free Hamiltonian with added constraints (following the method of Dirac [31] as discussed in [32]). As a consequence, we could in principle find coordinates in which the Jacobian matrix is symplectic with respect to the canonical symplectic matrix. Fortunately, this is an unnecessary complication, since the underlying dynamics are independent of the coordinates.

2. Exponents for spinning test particles

As discussed in Sec. IV B 2, the lack of explicit time dependence independence and the three constraints reduce the

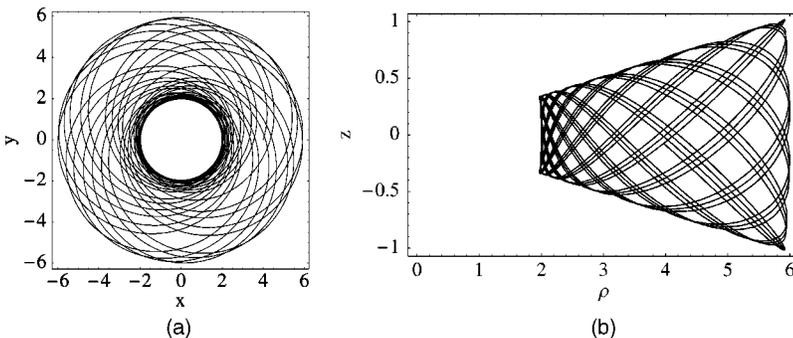


FIG. 5. The orbit of a nonspinning ($S=0$) test particle in maximal ($a=1$) Kerr spacetime, plotted in Boyer-Lindquist coordinates. (a) $y = r \sin \theta \sin \phi$ vs $x = r \sin \theta \cos \phi$; (b) z vs $\rho = \sqrt{x^2 + y^2}$. The orbital parameters are $E = 0.8837\mu$ and $J_z = 2.0667\mu M$, with pericenter $2.0M$ and apocenter $6.0M$.

¹¹It is possible that deformations of Kerr geometry that destroy Q nevertheless possess a numerical integral that preserves integrability, in analogy with some galactic potentials [29], but the loss of Q certainly ends the *guarantee* of integrability.

¹²A matrix S is symplectic with respect to the canonical symplectic matrix J if $S^T J S = J$, where $J = \begin{pmatrix} 0 & -I \\ I & 0 \end{pmatrix}$ and I is the $n \times n$ identity matrix.

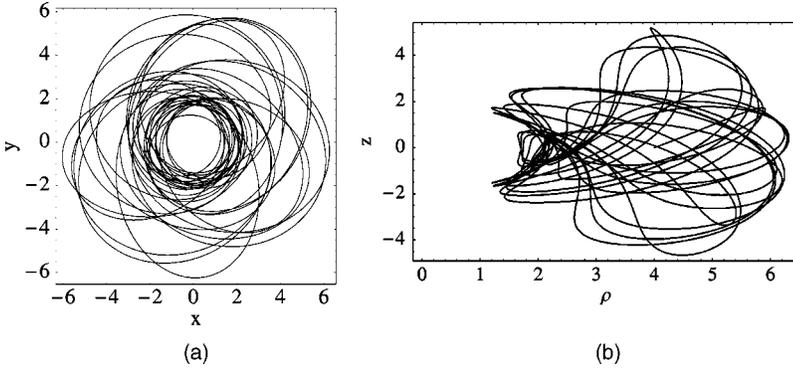


FIG. 6. The orbit of a maximally spinning ($S=1$) particle in maximal Kerr spacetime, for $E=0.8837\mu$ and $J_z=2.0667\mu M$ (the same values as in Fig. 5). The spin has initial values of $S^r=S^\mu=0.1$, corresponding to an initial angle of 54° with respect to the vertical in the particle's rest frame. As in Fig. 5, we plot y vs x in (a) and z vs ρ in (b). The spin causes significant deviations from geodesic orbits.

degrees of freedom from twelve to eight, which leaves the possibility of eight nonzero Lyapunov exponents. The phase space flow is further constrained by the constants of the motion, energy and z angular momentum; corresponding to each constant should be a zero Lyapunov exponent, since trajectories that start on an invariant torus must remain there. This leaves six exponents potentially nonzero. Since the exponents must come in pairs $\pm\lambda$, there should be at most three independent nonzero exponents.

VII. RESULTS

First we give results for the dynamics of the Papapetrou equations in the extreme (and unphysical) limit $S=1$, which represents a violation of the test-particle approximation but is still mathematically well-defined. We find the presence of chaotic orbits (in agreement with [3]). We next examine the effects of varying S , including the limit $S\ll 1$. Finally, we investigate more thoroughly the dynamics for physically realistic spins.

A. Chaos for $S=1$

1. Maximally spinning Kerr spacetime

In a background spacetime of a maximally spinning Kerr black hole ($a=1$) (see Fig. 5) there are unambiguous posi-

tive Lyapunov exponents for a range of physical parameters when $S=1$. We show a typical orbit that produces nonzero Lyapunov exponents in Fig. 6. The orbit has energy $E=0.8837\mu$, z angular momentum $J_z=2.0667\mu M$, and the radius ranges from pericenter $r_p=1.7M$ to apocenter $r_a=6.7M$. The Lyapunov integrations typically run for $10^4 M$, which corresponds approximately to 400ϕ -orbital periods.

We can illustrate the presence of a chaotic orbit by plotting the natural logarithm of the i th ellipsoid axis $\log[r_i(\tau)]$ vs τ [Eq. (3.7)], so that the slope is the Lyapunov exponent, as shown in Fig. 7.¹³ There appear to be two nonzero Lyapunov exponents; the third largest exponent is consistent with zero, as shown in Fig. 8. The reflection symmetry of the figure is a consequence of the exponent pairing: for each line with slope λ , there is a second line with slope $-\lambda$.

The main plot in Fig. 7(a) is generated by the modified Gram-Schmidt (GS) algorithm (Sec. IV B 2). Recall that this method depends on the value of ϵ used to infer the tangent vector; we find a valid ϵ by calibrating it using the rigorous Jacobian method, which must yield an exponent that matches the largest exponent from the modified Gram-Schmidt method. The plot in Fig. 7(a) represents the case $\epsilon=10^{-6}$; it is apparent that the two methods agree closely. The unres-

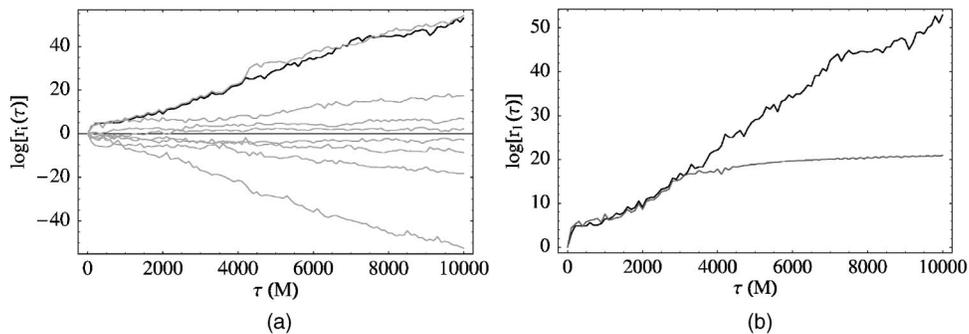


FIG. 7. Natural logarithms of the phase space ellipsoid axes vs proper time in Kerr spacetime with $S=1$. The slopes of the lines are the Lyapunov exponents; the largest exponent is approximately $\lambda_{\max}=5\times 10^{-3}M^{-1}$. The initial conditions are the same as in Fig. 6, and one point is recorded at each time $T=100M$ (Sec. III B 1). (a) Full Gram-Schmidt Jacobian method (light) with rigorous Jacobian method (dark). The full GS method is rescaled at each time T according to the algorithm in Sec. III B, while the rigorous Jacobian method is unrescaled. The two methods agree closely on the value of the largest Lyapunov exponent. (b) Rigorous Jacobian method compared to unrescaled deviation vector method. Note that the latter method, which started with a deviation of size 10^{-7} , saturates at ~ 16 . This corresponds to a growth of $e^{16}\approx 9\times 10^6$, which means that the separation has grown to a size of order unity.

¹³It is traditional to plot $\log[r_i(\tau)]/\tau$, which converges to the Lyapunov exponent as $\tau\rightarrow\infty$, but it is much easier to identify the linear growth of $\log[r_i(\tau)]$ than to identify the convergence of $\log[r_i(\tau)]/\tau$. The $\pm\lambda$ property is also clearer on such plots.

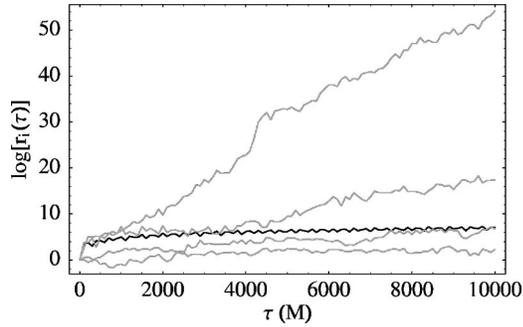


FIG. 8. Ellipsoid axis lengths from the upper half of Fig. 7(a) (light), compared to an integration with zero spin and hence zero Lyapunov exponent (dark). Only two of the four lines represent exponents distinguishable from zero.

caled deviation vector method provides an additional check on the validity of the largest exponent, as shown in Fig. 7(b). As expected, the unrescaled approach closely tracks the full Jacobian approach until it saturates.

The numerical values of the exponents are shown in Table I. The $\pm\lambda$ property is best satisfied by $\pm\lambda_{\max}$, the exponents with the largest absolute value. The exponents are least-squares fits to the data, with approximate standard errors of 1%. These errors are not particularly meaningful since the exponents themselves can vary by $\sim 10\%$ depending on the initial direction of the deviation vector. Moreover, even exponents that appear nonzero may be indistinguishable from zero in the sense of Fig. 8; for such exponents a “1%” error on the fit is meaningless.

For initial conditions considered in Fig. 6, and other orbits in the strongly relativistic region near the horizon, the typical largest Lyapunov exponents are on the order of a few $\times 10^{-3}/M$. For the particular case illustrated in Fig. 6, we have $\lambda_{\max} \approx 5 \times 10^{-3} M^{-1}$, which implies an e -folding time scale of $\tau_{\lambda} \equiv 1/\lambda \approx 2 \times 10^2 M$. This is strongly chaotic, with a significant divergence in approximately eight ϕ -orbital periods.

Based on integrations in the case of zero spin, which corresponds to no chaos (Lyapunov exponents all zero), we can determine how quickly the exponents approach zero

TABLE I. Lyapunov exponents in Kerr spacetime in units of $10^{-3} M^{-1}$, using a least squares fit. The exponents correspond to the semimajor axis evolution shown in Fig. 7(a). As is typical with the Gram-Schmidt Jacobian method, the exponents with the largest magnitudes are determined most accurately, and thus show the $\pm\lambda$ property most clearly. The standard errors on the fit are $\sim 1\%$ for each exponent, but these errors are dominated by two systematic errors: (i) the variation due to different choices of initial (random) tangent vectors; (ii) nonzero numerical values even for exponents that converge to zero eventually. In particular, the four smallest exponents (in absolute value) are indistinguishable from zero (see Fig. 8).

	5.5	1.5	0.56	0.25
$+\lambda$	5.5	1.5	0.56	0.25
$-\lambda$	5.3	1.6	0.76	0.072

numerically.¹⁴ Figure 8 compares the four apparently positive exponents with a known zero exponent. Only two of the four exponents are unambiguously distinguishable from zero, consistent with the argument in Sec. VI B that there should be at most three independent nonzero exponents.

Finally, we note that the components of the direction of largest stretching are all nonzero in general. The chaos is not confined to the spin variables alone, but rather mixes all directions. This indicates that chaos could in principle manifest itself in the gravitational waves from extreme mass-ratio binaries—but see Sec. VII C below.

2. Schwarzschild spacetime revisited

We now reconsider the case of a spin $S=1$ particle in Schwarzschild spacetime, as investigated in Ref. [3]. Figure 9 shows an orbit similar to a chaotic orbit considered there [Fig. 4(d) in [3]]. A plot of $\log[r_i(\tau)]$ vs τ (Fig. 10) shows behavior similar to that in Fig. 7. In particular, the $\pm\lambda$ symmetry is present, apparently with two positive exponents. (The other lines are indistinguishable from zero, again using $S=0$ orbits as a baseline.) The largest exponent of $1.5 \times 10^{-3} M^{-1}$ agrees closely with the value from Ref. [3], which reported an exponent of $\sim 2 \times 10^{-3} M^{-1}$ for a similar orbit. (This agreement is somewhat surprising, since [3] ap-

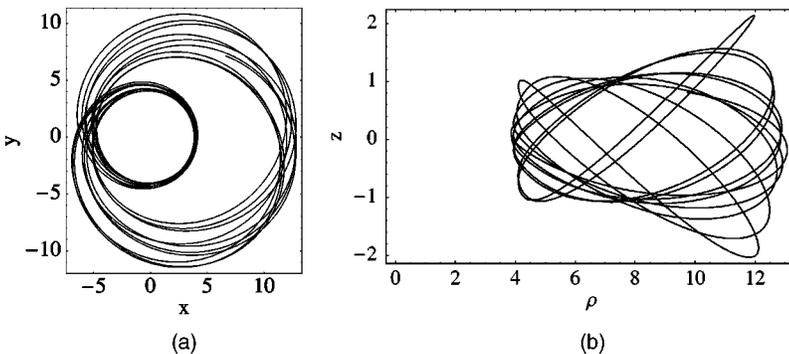


FIG. 9. The orbit of a maximally spinning ($S=1$) test particle in Schwarzschild spacetime for $E=0.94738162\mu$ and $J_z=4.0\mu M$. As before, we plot (a) y vs x and (b) z vs $\rho = \sqrt{x^2 + y^2}$.

¹⁴As noted in the Introduction, it is possible for integrable but unstable orbits to have positive Lyapunov exponents. We avoid this issue by choosing a baseline orbit that is not unstable.

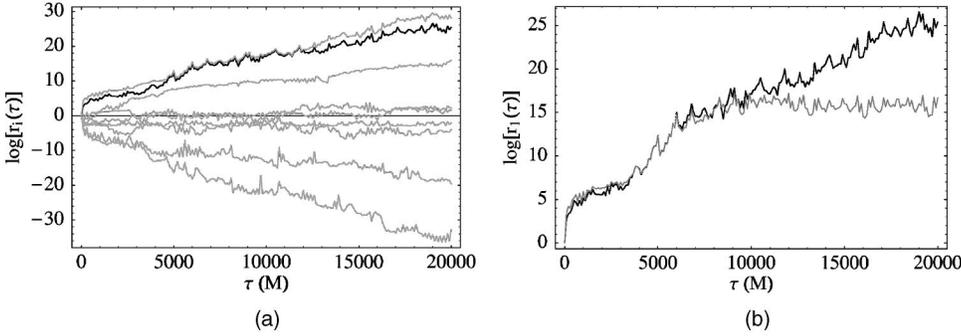


FIG. 10. Natural logarithms of the phase space ellipsoid axes vs proper time in Schwarzschild spacetime with $S=1$. The largest exponent is $\lambda_{\max} \approx 1.2 \times 10^{-3} M^{-1}$. The initial conditions are the same as in Fig. 9. (a) Full Gram-Schmidt Jacobian method (light) with rigorous Jacobian method (dark). (b) Rigorous Jacobian method compared to unrescaled deviation vector method. As in Fig. 7(b), the unrescaled method eventually saturates.

appears not to have taken the constrained nature of the deviation vectors into account. Luckily, the exponents are robust, and even unconstrained deviation vectors give nearly correct results.)

3. Kerr and Schwarzschild orbits compared

The Kerr and Schwarzschild Lyapunov exponents of the previous two sections are not all that different; both are 10^{-2} – $10^{-3} M^{-1}$ in order of magnitude (see Table II). Nevertheless, the two systems prove to be quite different: chaotic orbits are easy to find in Kerr spacetime for nearly any initial condition that explores the strongly relativistic region near the horizon, whereas nearly all analogous orbits in Schwarzschild spacetime are not chaotic.

Figure 11 compares Kerr and Schwarzschild orbits with the same inclination angle $\iota = 10^\circ$ and eccentricity $e = 0.5$ but varying pericenters r_p . (Details of this parametrization method, mentioned above in Sec. II C, appear in [33].) We insure that the systems are analogous by using orbits of $S = 1$ particles with the same values of r_p/r_{ms} , where r_{ms} is the radius of the marginally stable orbit in the corresponding $S=0$ (geodesic) case. We use a Kerr geodesic integrator developed by Hughes [34] to find r_{ms} , which is the smallest pericenter that still yields a stable orbit. For the values of ι and e considered, $r_{\text{ms}} = 1.0M$ for Kerr orbits and $r_{\text{ms}} = 4.67M$ for Schwarzschild orbits.

It is evident from Fig. 11 that the Kerr orbits are chaotic for a broad range of pericenters, with the maximum Lyapunov λ_{\max} generally decreasing as the pericenter increases. In contrast, the Schwarzschild orbits are not chaotic

TABLE II. Lyapunov exponents in Schwarzschild spacetime in units of $10^{-3} M^{-1}$, using a least squares fit. The exponents correspond to the semimajor axis evolution shown in Fig. 10(a), which is similar to the orbit in Fig. 4(d) of Ref. [3]. As with the Kerr case (Table I), the standard errors on the fit are $\sim 1\%$ for each exponent, and the same caveats apply. The four smallest exponents (in absolute value) are indistinguishable from zero in the sense of Fig. 8.

$+\lambda$	1.2	0.67	0.21	0.0063
$-\lambda$	1.5	0.57	0.10	0.00023

anywhere over the entire range of valid initial conditions. In fact, we are unable to find any chaotic orbits in Schwarzschild spacetime other than the types identified by Suzuki and Maeda [3], which were exceptional cases of orbits on the edge of a generalized effective potential. In Kerr, on the other hand, chaotic orbits appear to be the rule for pericenters near r_{ms} .

B. Dependence on S

Since chaos must disappear as $S \rightarrow 0$, we expect to see the largest Lyapunov exponent approach zero in this limit. This is indeed the case: in Fig. 12, which shows the variation of λ_{\max} with S for two different orbits, we see that the chaos unambiguously present when $S=1$ is not present for smaller values of S . In particular, the largest Lyapunov exponent is indistinguishable from zero over the entire range $10^{-6} \leq S \leq 10^{-1}$. (The far left of the plots have data points for each decade in this range.)

Although the strength of the chaos generally decreases with S , one remarkable feature of Fig. 12(a) is the return of

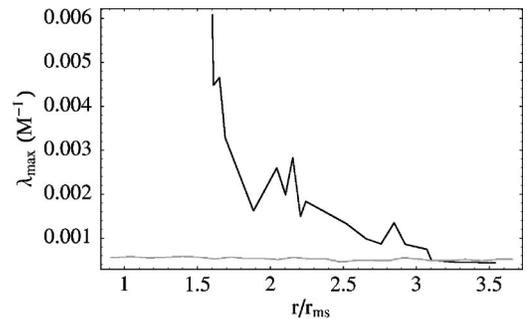


FIG. 11. Comparison of maximally spinning ($S=1$) Kerr particle orbits (dark) and Schwarzschild particle orbits (light). We plot the largest Lyapunov exponent versus pericenter (normalized by the marginally stable radius). The Kerr initial conditions for the innermost orbits are essentially as in Fig. 6. The Schwarzschild orbits are identical to their Kerr counterparts in inclination (10°) and eccentricity ($e=0.5$) but have the Kerr parameter a set to zero. The Schwarzschild orbits have exponents indistinguishable from zero over the entire range of parameters.

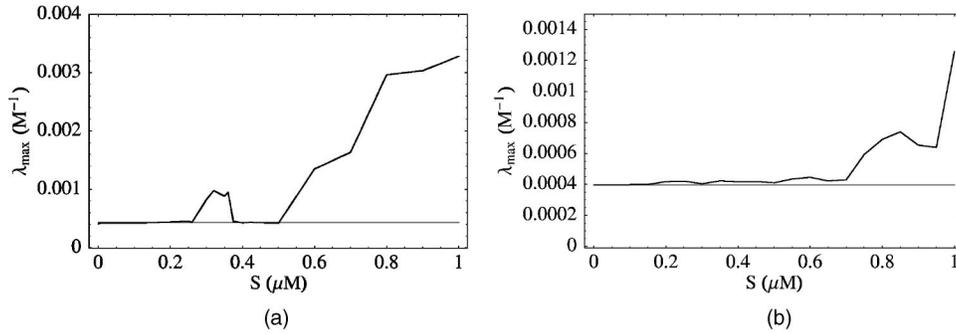


FIG. 12. Variation of the largest Lyapunov exponent vs S . (a) The spin $S=1$ initial conditions are the same as in Fig. 6. (b) Another $S=1$ case with a different inclination angle (20°) and pericenter ($2.5 M$). As the spin decreases, we hold fixed the Kerr orbital parameters: inclination angle, eccentricity, and pericenter. Note that in (a) the chaos disappears below $S \sim 0.5$, but returns in a region centered on $S \sim 0.3$. The horizontal line in both plots is the value of λ_{\max} calculated for the baseline $S=0$ orbit. In both (a) and (b) the Lyapunov exponent is indistinguishable from zero for physically realistic spins.

chaotic orbits between $S \sim 0.25$ and 0.4 after their disappearance at $S \sim 0.5$. The effect is qualitatively clear in Fig. 13. This chaotic “bump” in λ_{\max} vs S illustrates an important theme in nonlinear dynamical systems: the *only* way to determine whether an orbit is chaotic is to do the calculation. Though we certainly expect the strength of chaos to be smaller for $S \ll 1$ than for $S \approx 1$, it is impossible, in general, to determine *a priori* whether a particular set of parameters will lead to chaotic behavior.

C. Physically realistic spins

The Papapetrou equations are only realistic in the test-particle limit, so physically realistic spins must satisfy $S \ll 1$ (Sec. II B). This corresponds to likely sources of gravitational waves for LISA [35–37], e.g., maximally spinning $\mu = 10 M_\odot$ black holes spiraling into supermassive $M = 10^6 M_\odot$ Kerr black holes, which have spin parameters of $S = \mu/M = 10^{-5}$. Because of their likely importance as emitters of gravitational waves, it is essential to understand the dynamics of such systems.

1. Vanishing Lyapunov exponents

We would like to be able to make a definitive statement about the presence or absence of chaos for “small” spins, e.g., values of S in the range 10^{-2} – 10^{-6} . Unfortunately,

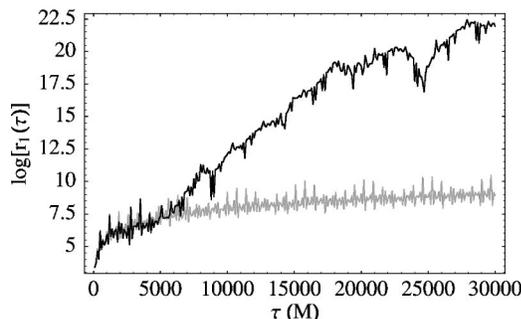


FIG. 13. Two orbits from the “bump” in Fig. 12(a). The $S=0.4$ orbit (light) is not chaotic, but the $S=0.3$ orbit (dark) is chaotic, despite having a smaller value of the spin.

when determining Lyapunov exponents numerically, it is impossible to conclude definitively that an orbit is or is not chaotic, since to do so would require an infinite-time integration. On the other hand, for suspected nonchaotic orbits, we can provide an approximate bound on the e -folding time scale.

The numerical values of exponents suspected to be zero depend strongly on the time of the integration. For example, for values of S in the range $10^{-2} \leq S \leq 10^{-6}$, the exponent in Fig. 12 appears to be $\lambda_{\max} \approx 5 \times 10^{-4} M^{-1}$, but this plot represents an integration time of only $10^4 M$. Longer integration times give correspondingly smaller estimates for the suspected zero exponents (Fig. 14). For the system shown in Fig. 12, an integration of $10^7 M$ yields an estimate of $\lambda_{\max} \approx 3.0 \times 10^{-7} M^{-1}$ for all spins in the range $10^{-2} \leq S \leq 10^{-7}$. In this case, the relevant Lyapunov timescales are at least $3 \times 10^6 M$, and are probably much longer; the size of the bound is limited only by our patience and computer budget. It seems highly likely that such orbits are not chaotic.

2. Spin-induced phase differences

Even if their Lyapunov exponents are zero, small spins affect the relative phase of the orbits, and since phase differ-

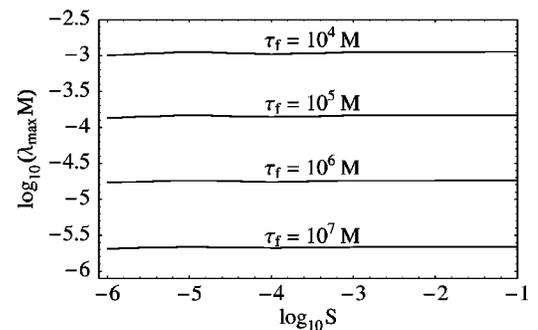


FIG. 14. The variation of the dimensionless quantity $\lambda_{\max} M$ with final integration time τ_f for spin parameter S in the range $10^{-2} \leq S \leq 10^{-6}$. From top to bottom, the total integration time varies from $10^4 M$ to $10^7 M$. It is likely that the true Lyapunov exponent is zero.

TABLE III. Phase shifts $\Delta\phi = \phi_{\text{geodesic}} - \phi_{\text{spin}}$ in radians as a function of orbital inclination angle ι and pericenter r_p for $a=0.5$ and $S = 10^{-5}$. Inclination angle $\iota=0^\circ$ is prograde equatorial and $\iota=180^\circ$ is retrograde equatorial. The geodesic orbits and their corresponding spin orbits start with the same initial 4-velocity v^μ , and the integrations are performed using Boyer-Lindquist coordinate time t , with $t_{\text{max}} \approx (2000 \text{ times the average radial orbital period})$. The pericenters are scaled by the marginally stable radius r_{ms} , and we start at $r_p/r_{\text{ms}} = 1.5$ to guarantee the existence of valid initial conditions for the nongeodesic orbit. The spin has fixed initial values of $S^{\hat{r}} = S^{\hat{\mu}} = 0.1S$ (with hats indicating an orthonormal basis), corresponding to initial angles of 9° to 30° with respect to the vertical in the particle's rest frame, increasing with decreasing pericenter.

ι	r_p/r_{ms}							
	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5
10°	1.50×10^2	5.69×10^3	4.32×10^3	2.13×10^3	2.02×10^3	1.27×10^3	1.14×10^3	7.77×10^4
45°	2.79×10^2	1.23×10^2	1.01×10^2	4.34×10^3	4.60×10^3	2.24×10^3	1.66×10^3	1.83×10^3
85°	4.36×10^2	2.92×10^3	1.48×10^3	8.24×10^4	1.00×10^3	2.20×10^3	1.86×10^3	1.26×10^3
135°	-9.02×10^3	-6.25×10^3	-2.34×10^3	-1.30×10^3	-1.73×10^3	-8.17×10^4	-6.76×10^4	-7.72×10^4
170°	8.40×10^4	2.85×10^4	1.84×10^4	7.31×10^5	1.25×10^4	1.12×10^4	3.35×10^5	3.07×10^5

ences accumulate secularly [38], the spin can still affect the gravitational wave signal. It is therefore useful to have a sense of the orders of magnitude of such spin-induced phase-shifts. Tables III and IV show typical values for the phase difference $\Delta\phi = \phi_{\text{geodesic}} - \phi_{\text{spin}}$ for $S = 10^{-5}$, where the geodesic and spin systems start with the same initial 4-velocity v^μ . The most useful quantity in practice is the phase shift as measured by observers at infinity, so we integrate in terms of the Boyer-Lindquist coordinate time t in place of τ . (This involves multiplying the differential equations by $d\tau/dt$ at each time step.) As is apparent from the tables, the phase shifts range broadly, from 10^{-1} to 10^{-5} radians after 2000 radial orbital periods, but tend to decrease in magnitude with increasing inclination angle or pericenter.

Reference [37] shows that the number of orbital periods in a full inspiral from $r \approx 4M$ to the final plunge is $N \sim M/\mu$, which is 10^5 for the systems in Tables III and IV. Since the table represents values of $\Delta\phi$ for 2000 times the average radial orbital period, this means that the total phase shift during the inspiral is $50 \Delta\phi_{\text{table}}$. For a 10° inclination angle the total phase shift is on the order of a tenth of a radian to a radian. Slightly more realistic values of the number of orbits can be obtained using Fig. 2 in [37], which gives $N \sim 2 \times 10^4$ orbital periods from $r = 4M$ to the plunge at $r \approx M$ for $a = 0.998$, $\iota = 10^\circ$, and $M/\mu = 10^5$. Since the orbit spends most of its time between $4 r_{\text{ms}}$ and $2 r_{\text{ms}}$, interpolating in Table IV gives $\Delta\phi_{\text{total}} \approx 10 \times \Delta\phi_{r=3.0} = 2$

$\times 10^{-2}$. This is only a rough estimate, since the orbits in [37] are circular, while the orbits we consider are eccentric.

D. Comments on time, rescaling, and norms

In this paper, we have elected to use τ as the time parameter, a rescaling time T of $100M$, and a projected norm (Sec. IV A). Here we discuss the effects of varying these choices.

First, we consider the effects of using coordinate time t in place of τ . In Fig. 15, we plot the natural logarithm of the largest ellipsoid axis $\log[r_1(\tau)]$ vs τ together with $\log[r_1(t)]$ vs t . (We use the unrescaled deviation vector approach for simplicity, since the Jacobian approach requires a new Jacobian matrix for each coordinate change.) The exponents are $\lambda_\tau = 5.05 \times 10^{-3} M^{-1}$ and $\lambda_t = 2.51 \times 10^{-3} M^{-1}$, implying Lyapunov time scales of $\tau_\lambda = 1.98 \times 10^2 M$ and $t_\lambda = 3.98 \times 10^2 M$. The average value of $dt/d\tau$ over the orbit is 2.06, whereas $t_\lambda/\tau_\lambda = 2.01$, so the relationship

$$\frac{t_\lambda}{\tau_\lambda} = \left\langle \frac{dt}{d\tau} \right\rangle \quad (7.1)$$

discussed in the Introduction is well satisfied.

Second, we discuss the effects of varying the rescaling time T . We find that choosing T to be a moderate fraction of the shortest Lyapunov time scale (corresponding to the largest Lyapunov exponent) works best, giving each axis enough

TABLE IV. Phase shifts $\Delta\phi = \phi_{\text{geodesic}} - \phi_{\text{spin}}$ in radians as a function of orbital inclination angle ι and pericenter r_p for $a=1$ and $S = 10^{-5}$. As in Table III, the pericenters are scaled by r_{ms} , and the spin has fixed initial values of $S^{\hat{r}} = S^{\hat{\mu}} = 0.1S$ (corresponding in this case to initial angles of 28° to 61° , again decreasing with increasing pericenter).

ι	r_p/r_{ms}							
	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5
10°	7.21×10^2	4.58×10^2	2.41×10^2	1.83×10^2	1.10×10^2	9.46×10^3	6.56×10^3	7.43×10^3
45°	2.37×10^1	5.56×10^2	2.59×10^2	1.83×10^2	1.73×10^2	1.52×10^2	1.08×10^2	7.83×10^3
85°	1.96×10^2	6.21×10^3	2.82×10^3	2.13×10^3	2.66×10^2	3.64×10^3	6.47×10^4	3.48×10^3
135°	-1.04×10^2	-3.17×10^3	-3.21×10^3	-1.41×10^3	-1.12×10^3	-8.46×10^4	-8.82×10^4	-5.59×10^4
170°	3.89×10^4	1.48×10^4	6.68×10^5	5.97×10^5	8.09×10^5	9.55×10^5	3.06×10^5	1.66×10^5

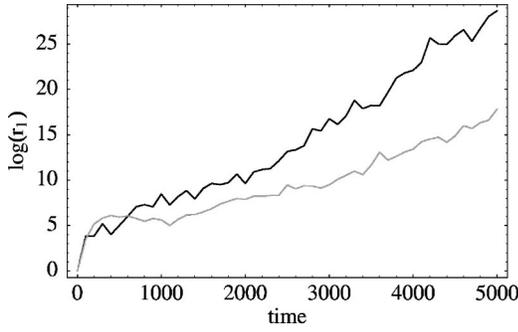


FIG. 15. The natural logarithm of the largest ellipsoid axis vs time for proper time τ (dark) and coordinate time t (light). The exponents are clearly different, but the Lyapunov time scales $\tau_\lambda = 1/\lambda_\tau$ and $t_\lambda = 1/\lambda_t$ are related by Eq. (7.1).

time to grow before rescaling while still keeping the negative exponents from underflowing and preventing the largest axis from dominating. Rescaling times between $50M$ and $100M$ work best for the systems we consider, which have Lyapunov time scales ranging from 10^2M to 10^3M . A comparison of results for $T=50M$ and $T=100M$ appears in Fig. 16.

Third, we compare the projected norm used here to a naive Euclidean norm for determining the length of the phase-space tangent vectors ξ_i . As shown in Fig. 17, even using a 12-dimensional Euclidean norm changes the resulting exponent very little (approximately 15% in this example). Given its conceptual advantages, we choose to use the projected norm with the confidence that the Lyapunov exponent order of magnitude is robust.

VIII. CONCLUSIONS

A spinning test particle, as described by the Papapetrou equations, appears to be chaotic in Kerr spacetime, with maximum e -folding time scales of a few $\times 10^2 M$. The applicability of this result is limited by three main factors: (i) chaos appears only for physically unrealistic values of the spin parameter; (ii) other effects, such as tidal coupling, may be important for some astrophysical systems, violating the pole-dipole approximation implicit in the Papapetrou equations; and (iii) we neglect gravitational radiation. The third limitation is not fatal, since the radiation time scales can be long enough that chaos, if present in the conservative limit,

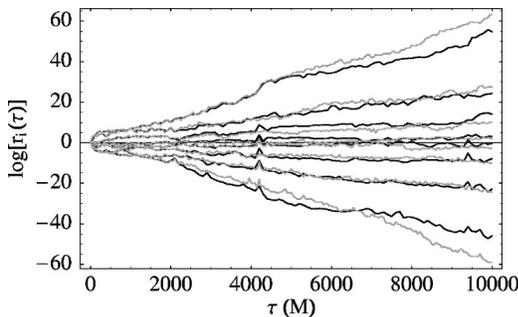


FIG. 16. Natural logarithms of the ellipsoid axes vs τ for rescaling time $T=100M$ (dark) and time $T=50M$ (light).

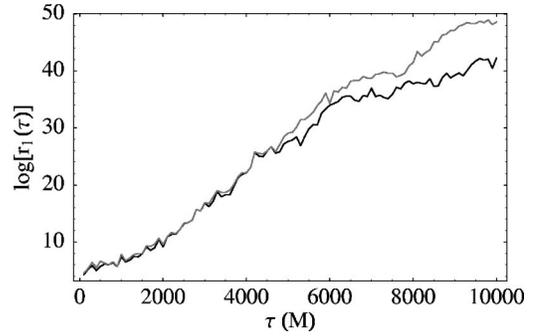


FIG. 17. The natural logarithm of the largest ellipsoid axis vs τ for the Euclidean norm (top) and the projected norm from Sec. IV A (bottom).

would have time to manifest itself in the gravitational radiation of extreme mass-ratio systems.

In the unphysical $S=1$ limit, the Lyapunov exponents exhibit characteristics expected of a Hamiltonian system, appearing in pairs $\pm\lambda$ (Sec. VI B). There are zero Lyapunov exponents which correspond to the constants of the motion, but the other exponents are in general nonzero. (For the Kerr orbits considered in this paper, we find that two of the three independent exponents are nonzero, as illustrated in Fig. 8.) Typical orders of magnitude for the largest Lyapunov exponents are a few $\times 10^{-3}M^{-1}$ for unphysical spins ($S=1$). For physically realistic spin parameters (Sec. VII C), we find that $\lambda_{\max} \leq \text{few} \times 10^{-7}M^{-1}$, corresponding to e -folding time scales of a few $\times 10^6M$. Even this bound appears to be limited only by the total integration time; in all physically realistic cases considered, λ_{\max} is indistinguishable from zero (using $S=0$ integrations as a baseline).

From the perspective of gravitational radiation detection, our most important conclusion is that chaos seems to disappear for physically realistic values of S , i.e., values of S for which the test-particle approximation and hence the Papapetrou equations are valid. We are unable to comment on the dynamics of comparable mass-ratio binaries, since such systems are not accurately modeled by the Papapetrou equations, but for *extreme* mass-ratio binaries it appears unlikely that chaos will present a problem for the calculation of theoretical templates for use in matched filters. A more thorough exploration of parameter space is needed to reach a firmer conclusion [33].

ACKNOWLEDGMENTS

I would like to thank Sterl Phinney for his support, encouragement, and excellent suggestions. Thanks also to Scott Hughes for contributing through his ideas and enthusiasm for this project. I would especially like to thank Janna Levin for her careful reading of the paper and insightful comments. Finally, I would like to thank Kip Thorne for teaching me general relativity and Jim Yorke for teaching me dynamical systems theory.

APPENDIX A: FULL JACOBIAN

For reference, we list the derivatives needed to calculate the full Jacobian matrix.

From Sec. V B, we have the following:

$$\frac{\partial \dot{x}^\mu}{\partial x^\nu} = N \left[p_\alpha g^{\alpha\mu}{}_{,\nu} + w^\mu{}_{,\nu} + v^\mu \left(v_\alpha w^\alpha{}_{,\nu} + \frac{1}{2} N w^\alpha w^\beta g_{\alpha\beta,\nu} \right) \right] \quad (\text{A1})$$

$$\frac{\partial \dot{x}^\mu}{\partial p_\nu} = N (g^{\mu\nu} + W^{\mu\nu} + N v^\mu w^\nu) + N v^\mu v_\alpha W^{\alpha\nu} \quad (\text{A2})$$

with

$$W^{\mu\nu} = - * R^{*\mu\alpha\nu\beta} S_\alpha S_\beta \quad (\text{A3})$$

$$\frac{\partial \dot{x}^\mu}{\partial S_\nu} = N V^{\mu\nu} + N v^\mu v_\alpha V^{\alpha\nu} \quad (\text{A4})$$

with

$$V^{\mu\nu} = - S_\alpha p_\beta (* R^{*\mu\alpha\beta\nu} - * R^{*\mu\nu\alpha\beta}). \quad (\text{A5})$$

Now we simply apply the product rule many times:

$$\begin{aligned} \frac{\partial \dot{p}_\mu}{\partial x^\nu} &= - p_\alpha S_\beta (R^{*\alpha\beta}{}_{\mu\gamma}{}_{,\nu} v^\gamma + R^{*\alpha\beta}{}_{\mu\gamma} v_{,\nu}^\gamma) + p_\alpha (\Gamma^\alpha_{\beta\mu,\nu} v^\beta \\ &\quad + \Gamma^\alpha_{\beta\mu} v^\beta{}_{,\nu}) \end{aligned} \quad (\text{A6})$$

$$\begin{aligned} \frac{\partial \dot{p}_\mu}{\partial p_\nu} &= - S_\beta \left(R^{*\nu\beta}{}_{\mu\gamma} v^\gamma + R^{*\alpha\beta}{}_{\mu\gamma} p_\alpha \frac{\partial v^\gamma}{\partial p_\nu} \right) + \Gamma^\nu{}_{\beta\mu} v^\beta \\ &\quad + \Gamma^\alpha{}_{\beta\mu} p_\alpha \frac{\partial v^\beta}{\partial p_\nu} \end{aligned} \quad (\text{A7})$$

$$\frac{\partial \dot{p}_\mu}{\partial S_\nu} = - R^{*\alpha\nu}{}_{\mu\gamma} v^\gamma p_\alpha - R^{*\alpha\beta}{}_{\mu\gamma} \frac{\partial v^\gamma}{\partial S_\nu} p_\alpha S_\beta + \Gamma^\alpha{}_{\beta\mu} p_\alpha \frac{\partial v^\beta}{\partial S_\nu} \quad (\text{A8})$$

$$\begin{aligned} \frac{\partial \dot{S}_\mu}{\partial x^\nu} &= - p_\mu S_\alpha p_\gamma S_\delta (R^{*\alpha\gamma\delta}{}_{\beta\nu} v^\beta + R^{*\alpha\gamma\delta}{}_{\beta\nu} v^\beta{}_{,\nu} \\ &\quad + S_\alpha (\Gamma^\alpha{}_{\beta\mu,\nu} v^\beta + \Gamma^\alpha{}_{\beta\mu} v^\beta{}_{,\nu})) \end{aligned} \quad (\text{A9})$$

$$\begin{aligned} \frac{\partial \dot{S}_\mu}{\partial p_\nu} &= - S_\alpha S_\delta v^\beta (\delta_\mu{}^\nu R^{*\alpha\gamma\delta}{}_{\beta\gamma} p_\gamma + p_\mu R^{*\alpha\gamma\delta}{}_{\beta\nu}) \\ &\quad - p_\mu R^{*\alpha\gamma\delta}{}_{\beta\nu} S_\alpha \frac{\partial v^\beta}{\partial p_\nu} p_\gamma S_\delta + \Gamma^\alpha{}_{\beta\mu} \frac{\partial v^\beta}{\partial p_\nu} S_\alpha \end{aligned} \quad (\text{A10})$$

$$\begin{aligned} \frac{\partial \dot{S}_\mu}{\partial S_\nu} &= - p_\mu p_\gamma v^\beta (R^{*\nu\gamma\delta}{}_{\beta\delta} S_\delta + R^{*\alpha\gamma\nu} S_\alpha) \\ &\quad - p_\mu R^{*\alpha\gamma\delta}{}_{\beta\nu} S_\alpha \frac{\partial v^\beta}{\partial S_\nu} p_\gamma S_\delta + \Gamma^\nu{}_{\beta\mu} v^\beta + \Gamma^\alpha{}_{\beta\mu} \frac{\partial v^\beta}{\partial S_\nu} S_\alpha. \end{aligned} \quad (\text{A11})$$

Accidentally leaving off the final term in $\partial \dot{S}_\mu / \partial S_\nu$ led to the robust but spurious chaotic behavior mentioned in Sec. V A.

-
- [1] J. Levin, Phys. Rev. Lett. **84**, 3515 (2000).
[2] J. Levin, gr-qc/0010100.
[3] S. Suzuki and K. Maeda, Phys. Rev. D **55**, 8 (1997).
[4] J.D. Schnittman and F.A. Rasio, Phys. Rev. Lett. **87**, 121101 (2001).
[5] N.J. Cornish, Phys. Rev. D **64**, 084011 (2001).
[6] N.J. Cornish and J. Levin, gr-qc/0207016.
[7] N.J. Cornish and J. Levin, Phys. Rev. Lett. **89**, 179001 (2002).
[8] S.A. Hughes, Class. Quantum Grav. **18**, 4067 (2001).
[9] See <http://lisa.jpl.nasa.gov>
[10] K.T. Alligood, T.D. Sauer, and J.A. Yorke, *Chaos: An Introduction to Dynamical Systems* (Springer, New York, 1997).
[11] N. Cornish and J. Levin, Phys. Rev. Lett. **78**, 998 (1997).
[12] E. Ott, *Chaos in Dynamical Systems* (Cambridge University Press, Cambridge, England, 1993).
[13] C.W. Misner, K.S. Thorne, and J.A. Wheeler, *Gravitation* (Freeman, San Francisco, 1973).
[14] A. Papapetrou, Proc. R. Soc. London **A209**, 248 (1951).
[15] W.G. Dixon, Proc. R. Soc. London **A314**, 499 (1970).
[16] C. Møller, *Theory of Relativity* (Oxford University Press, London, 1972).
[17] B.M. Barker and R.F. O'Connell, Gen. Relativ. Gravit. **5**, 539 (1974).
[18] L.E. Kidder, Phys. Rev. D **52**, 821 (1995).
[19] O. Semerák, Mon. Not. R. Astron. Soc. **308**, 863 (1999).
[20] G.B. Cook, S.L. Shapiro, and S.A. Teukolsky, Astrophys. J. **424**, 823 (1994).
[21] D. Lai, F.A. Rasio, and S.L. Shapiro, Astrophys. J., Suppl. Ser. **88**, 205 (1993).
[22] M. Nauenberg, Astrophys. J. **175**, 417 (1972).
[23] V.S. Geroyannis and P.J. Pappasotiriou, Astrophys. J. **534**, 359 (2000).
[24] P. Wiggins and D. Lai, Astrophys. J. **532**, 530 (2000).
[25] T. Tanaka, Y. Mino, M. Sasaki, and M. Shibata, Phys. Rev. D **54**, 3726 (1996).
[26] J.-P. Eckmann and D. Ruelle, Rev. Mod. Phys. **57**, 617 (1985).
[27] V. Karas and D. Vokrouhlický, Gen. Relativ. Gravit. **24**, 729 (1992).
[28] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery, *Numerical Recipes in C* (Cambridge University Press, Cambridge, England, 1992).
[29] J. Binney and S. Tremaine, *Galactic Dynamics* (Princeton University Press, Princeton, NJ, 1987).
[30] S. Hojman, Ph.D. thesis, Princeton University (1975).
[31] P.A.M. Dirac, *Lectures on Quantum Mechanics* (Belfer Gradu-

- ate School of Science, Yeshiva University, New York, 1964).
- [32] A.J. Hanson and T. Regge, *Ann. Phys. (N.Y.)* **87**, 498 (1974).
- [33] M. D. Hartl, in preparation.
- [34] S. Hughes, http://www.tapir.caltech.edu/~listwg1/EMRI/hughes_geodesic.html
- [35] L.S. Finn and K.S. Thorne, *Phys. Rev. D* **62**, 124021 (2000).
- [36] S.A. Hughes, *Phys. Rev. D* **61**, 084004 (2000).
- [37] S.A. Hughes, *Phys. Rev. D* **64**, 064004 (2001).
- [38] T.A. Apostolatos, C. Cutler, G.J. Sussman, and K.S. Thorne, *Phys. Rev. D* **49**, 6274 (1994).