Computer simulations, exact trajectories, and the gravitational *N*-body problem

Wayne Hayes
Institute for Physical Science and Technology, Building 224, University of Maryland, College Park, Maryland 20742

(Received 11 December 2003; accepted 30 April 2004)

Many physical systems of current interest are chaotic, which means that numerical errors in their simulation are exponentially magnified with the passage of time. This could mean that a numerical solution of a chaotic system is the result of nothing but magnified noise, which calls into question the value of such simulations. Although this fact has been well known for a long time, its impact on the validity of simulations is not well understood. The study of *shadowing* may provide an answer. A shadow is an *exact* trajectory of a chaotic map or ordinary differential equation that remains close to an approximate solution for a nontrivial duration of time. If it can be shown that a numerical solution has a shadow, then the validity of the solution is strong, in the sense that it can be viewed as an experimental observation of the shadow, which is an exact solution. We present a discussion of shadowing, including an algorithm to find shadows, using the gravitational *N*-body problem as an example.
[DOI: 10.1119/1.1764561]

I. INTRODUCTION AND MOTIVATION

> ''The art of prophecy is very difficult, especially with respect to the future,'' attributed to Mark Twain.[1]

Gravitational *N*-body systems are chaotic, which means that they depend sensitively on the initial condition, so that except for a few special cases, the phase-space distance between two solutions whose initial conditions differ by an arbitrarily small amount will increase exponentially with time. Because computers constantly make small errors in the computation of such solutions, it is guaranteed that a computed solution will diverge exponentially from the exact solution with the same initial conditions. Thus, it is possible that numerical solutions for chaotic systems are overwhelmed by the exponential magnification of small errors, which would imply that the computed solutions are worthless, even if quantities such as the total energy or momentum are conserved to arbitrary accuracy. That is, there are infinitely many solutions whose energy is exactly the same, but that have vastly different phase space trajectories.

A. Simulations

Numerical simulation is a standard tool in the study of complex systems. For example, the astronomical literature is brimming with the results of large gravitational *N*-body simulations. Examples include studies of the formation, evolution and structure of galaxies and clusters of galaxies,[2] and the cosmos at large.[3] Because such simulations have been used to invalidate theories,[3] establishing their validity is critical. Like many dynamical systems, however, a gravitational system displays chaotic behavior: two solutions whose initial conditions differ by an arbitrarily small amount generally diverge exponentially from each other.[4] Because numerical methods introduce errors, it is virtually guaranteed that a numerically computed solution diverges exponentially from the exact solution with the same initial conditions. The phenomenon has been described (see, for example, Ref. 4) as the ''exponential magnification of small errors,'' leaving open the possibility that trajectories of such simulations are the result of nothing but magnified noise.

Although much effort has been devoted to many aspects of the reliability of simulations, and although sensitive dependence on initial conditions is widely known to be one of these aspects, its impact on simulation reliability is not well understood. Can these simulation results really be trusted? What conditions must a simulation meet for its accuracy to be assured? Is there a limit on the duration of the time a system can be simulated accurately? What measures can be used to ascertain the accuracy of a simulation? More fundamentally, what do we mean by ''accuracy'' and ''error'' in these simulations, given that we know the numerical solution diverges exponentially from the exact solution?

B. History of exponential divergence in gravitational systems

Miller[5] was the first to show that small changes in the initial conditions of a gravitational system result in exponentially diverging solutions. Lecar[6] coordinated a study between many researchers, each of whom independently computed the solution to an *N*-body problem with identical initial conditions. They found that different algorithms and computers gave results in which some measures differed by as much as 100%. More recent work on the growth of errors includes Kandrup and Smith,[7] who showed that under a large range of parameters, the time scale over which small perturbations grow by a factor of e, called the e-folding or *Lyapunov* time, is comparable to the *crossing time*, the average time it takes a particle to cross the system once. Goodman, Heggie, and Hut[4] developed a detailed theory of the growth of small perturbations, and verified it with simulations to show that the exponential increase of small errors results mostly from close encounters, which occur infrequently. This result is interesting because it says that, even though the full phase-space solutions may experience exponential error growth, the growth is dominated by the particles that undergo stronger collisions, while the error growth of

less strongly interacting particles is slower. This result suggests that perhaps the simulations of collisionless systems can be trusted for longer than collisional ones, because close encounters have a much smaller effect in the former. Kandrup, Smith, and Willmes[8] showed that as softening is increased (that is, the system is made less collisional—see Sec. II), the Lyapunov exponent decreases, so errors are magnified more slowly. Reference 8 agrees with Ref. 4 that the error magnification is due more to the rare individual particles whose errors grow much more quickly than the average, although they claim the collective gravitational potential of all the particles, called the ''global'' potential, also plays a role.

Because the time scale for the growth of errors is so short (the errors can be magnified by a factor of ~ 10 each crossing time), the results of all N-body simulations may be suspect. If the relative error per crossing time for a simulation is 10^{-p}, then after about p crossing times, a particle's position will have an error comparable to the size of the system—in other words, all information will have been lost about the particle's position. A typical simulation has an error per crossing time with p between 4 and 8. Because such systems are generally simulated for tens, hundreds, or thousands of crossing times, we can be sure that the positions of individual particles bear no resemblance to the positions they would have in the exact system, starting with the same initial condition.

C. The kinds of errors made in N-body simulations

1. Input and output errors

I first distinguish between two general types of errors. *Input errors* can be controlled directly while devising and implementing models of N-body systems. These errors in N-body systems may be divided into modeling approximations and implementation approximations. Modeling approximations simplify the system being simulated, and include the following.

(i) Finite N sampling, also called *discreteness noise*, because the N used is generally several orders of magnitude less than the N of the real system being modeled. The consensus is that this error is the limiting source of error in current large N-body simulations.[9–15]

(ii) Force softening, that is, replacing r^2 by $(r^2+\epsilon_{\text{soft}}^2)$ in the denominator of Newton's gravitational force equation for some small constant ϵ_{soft}, usually chosen to approximate the average inter-particle separation. Softening is used because it allows a smaller N to approximate a larger N, and also to eliminate the singularity at $r=0$.[16]

Implementation approximations measure how well the implementation simulates the model, and includes machine roundoff error, numerical integration truncation error, and approximate force computation algorithms like the Barnes–Hut tree code or the fast multipole method.[17] Hernquist, Hut, and Makino[9] have tried to show that the effect of this latter error is negligible by showing that the energy of each individual particle is conserved to a high degree regardless of whether the Barnes–Hut or the direct $O(N^2)$ algorithm is used. However, they used the leapfrog integrator with a constant time step, which guarantees that the energy error for the entire system is bounded.[18,19]

The *output error* measures the difference between the output and a real system, and results from the cumulative effect of all the input errors. A simulation with small output errors can be said to have high accuracy. Given that N-body systems are chaotic, and that their simulation introduces the input errors we discussed, we must now ask precisely what we mean by the ''accuracy'' of a simulation. Amazing as it may seem, there is currently no clear definition of simulation accuracy.[20] Obviously, attempting to follow the individual paths of all N particles is impractical; Goodman, Heggie, and Hut[4] show that this would require $O(N)$ digits of precision. On the other hand, most astronomical publications quote energy conservation as their only measure of output error, even though there are infinitely many solutions with equal energy but vastly different phase-space trajectories. Some of these simulations even use an energy-conserving integrator such as the leapfrog algorithm, in which case quoting energy conservation is of dubious merit, because the integrator conserves energy no matter how big other errors become!

2. Macroscopic statistics versus microscopic details

In large N-body simulations, one is not usually concerned with the precise evolution of individual particles, but instead with the evolution of the distribution of particles.[21] Most practitioners know that the exponential magnification of errors means they cannot possibly trust the microscopic details, but they believe that the statistical results are independent of the microscopic errors, although little work has been done to test this belief.[4] Barnes and Hut[10] claim that gravitational N-body simulations require only ''modest'' accuracy levels, but also concede that quoting energy conservation is not enough, and that more stringent tests are needed.

An example of conservation of macroscopic properties is given by Kandrup and Smith.[7] They show that a histogram of the Lyapunov times of individual particles stays constant within statistical uncertainties, even though the phase-space distribution of those particles is vastly different for different initial conditions.

II. SHADOWING

> ''What shadows we are, and what shadows we pursue,'' Edmund Burke.[22]

We have noted that many dynamical systems display sensitive dependence on initial conditions. This sensitive dependence guarantees that a typical numerical solution diverges exponentially from the exact solution with the same initial conditions. Fortunately, most studies of dynamical systems do not aim to predict the precise evolution of a particular choice of initial conditions. Instead, the dynamics of the system is sampled in order to study its general behavior. In such cases, we typically choose initial conditions from a random distribution and would be happy if our numerical solution exhibited behavior typical of *any* valid choice of initial conditions from our distribution. In particular, we would be satisfied if our numerical solution closely follows some exact solution whose initial conditions are close to those that we chose. The study of *shadowing* provides just such a property: a *shadow* is an exact solution to a given set of equations that remains close to a numerically computed solution of the

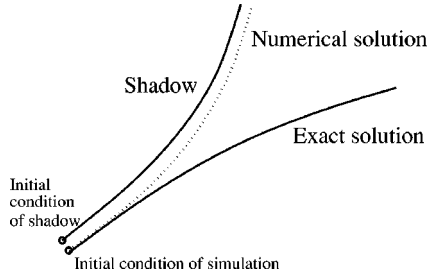

Fig. 1. A schematic diagram of shadowing. Because a chaotic system has a sensitive dependence on initial conditions, and because a numerical solution contains errors that perturb it from the exact solution, the numerical solution diverges exponentially from the exact solution with the same initial condition. A shadow, if one exists, has a nearby initial condition, and remains close to the numerical solution for a nontrivial duration of time.

same set of equations. A shadow, if it exists, has initial conditions close to those of the numerical solution, and also remains close to the numerical solution for a nontrivial duration of time. By ''nontrivial duration,'' we mean that the shadow remains close to the numerical solution significantly longer than the exact solution that starts at the same initial condition as the numerical solution, because *that* exact solution is known to diverge exponentially away from the numerical one (see Fig. 1).

A. Definitions

When referring to mathematical variables, boldface will refer to vectors, and *italic* will refer to scalars, matrices, and functions. Scalars are written in lower case and matrices in upper case. Some of the following definitions are taken, with minor modifications, from Grebogi, Hammel, Yorke, and Sauer (GHYS).[23] The terms *trajectory*, *orbit*, and *solution* are used interchangeably throughout this paper.

Although most physical systems we model evolve continuously in time, a numerical solution is a discrete sequence of points $\{\mathbf{y}_i\}_{i=a}^{b}$, representing snapshots of the system occurring at successive small time intervals. The relationship between the continuous system and the discrete sequence of points can be discussed by defining the *time-h solution operator* φ_h to be the function that takes a point $\mathbf{x}$ in phase space and gives us the *exact* solution passing through $\mathbf{x}$ at a time h later. For convenience, we will drop the subscript h. Then we say that φ generates an exact solution. A numerical solution, on the other hand, is a discrete sequence of points generated using an *approximation* to φ. We now define these ideas more formally.

Definition: An *exact trajectory* $\{\mathbf{x}_i\}_{i=a}^{b}$ of φ satisfies $\mathbf{x}_{i+1}=\varphi(\mathbf{x}_i)$ for $a\leqslant i<b$. We are interested in the case where a and b are finite integers. For a chaotic map, φ may be a simple equation, such as the logistic equation $\varphi(x)=1-2x^2$, which always maps the interval $[-1,1]$ onto itself. For a system described by ordinary differential equations like the N-body problem, $\varphi(\mathbf{x})$ represents the exact solution of integrating the phase-space coordinates $\mathbf{x}$ for one time step.

Definition: $\{\mathbf{y}_i\}_{i=a}^{b}$ is a *δ-pseudo-trajectory*, also called a *noisy orbit*, for φ if $\|\mathbf{y}_{i+1}-\varphi(\mathbf{y}_i)\|<\delta$ for $a\leqslant i<b$. We call δ the *noise amplitude*.

Definition: For $a\leqslant i<b$, the *1-step error* made between step i and step $i+1$ of the pseudo-trajectory $\{\mathbf{y}_i\}_{i=a}^{b}$ is the vector $\mathbf{e}_{i+1}=\mathbf{y}_{i+1}-\varphi(\mathbf{y}_i)$. Thus, an exact trajectory is one whose 1-step error is identically zero.

Definition of shadowing: The exact trajectory $\{\mathbf{x}_i\}_{i=a}^{b}$ *ϵ-shadows* the δ-pseudo-trajectory $\{\mathbf{y}_i\}_{i=a}^{b}$ on $a\leqslant i\leqslant b$ if $\|\mathbf{x}_i-\mathbf{y}_i\|<\epsilon$ for $a\leqslant i\leqslant b$.

Definition: The pseudo-trajectory $\{\mathbf{y}_i\}_{i=a}^{b}$ has a *glitch* at point $i=g_0<b$ if for some relevant ϵ, there exists an exact trajectory that ϵ-shadows $\{\mathbf{y}_i\}_{i=a}^{b}$ for $a\leqslant i\leqslant g_0$, but no exact trajectory that ϵ-shadows it for $a\leqslant i\leqslant g$, for $g>g_0$.

The first group of chaotic systems for which it was proven that shadow orbits exist are *hyperbolic* systems.[24,25] In a two-dimensional hyperbolic system, there are two special directions called the *unstable* (or *expanding*) and the *stable* (or *contracting*) directions, which are generally not orthogonal. Small perturbations along the stable direction decrease exponentially in time, while small perturbations in the unstable direction increase exponentially in time. The two directions reverse roles if the time is reversed. A trajectory for such a system can be imagined as a point moving in a plane.

For such a system it was shown that, if the angle between the stable and unstable directions is uniformly bounded away from 0, then a noisy trajectory can be shadowed for all time.[24,25] For nonhyperbolic systems, it appears that shadows may exist only for a finite time.[23] The most important question in this regard is, how long can a noisy orbit be shadowed? If the time is at least as long as most typical simulations of chaotic nonhyperbolic systems, then the simulations have great validity; if the shadowing time turns out to be too short, then we cannot use shadowing as a justification for the validity of the simulation.

Refinement is an iterative process that, on each iteration, simultaneously perturbs all points of a noisy orbit in an attempt to produce a nearby orbit with less noise. A refinement iteration is *successful* if the trajectory before the iteration has noise δ^0 and the trajectory after the iteration has noise δ^1, where $\delta^1<\mu\delta^0$ for some reasonable $\mu\in(0,1)$. Otherwise, the iteration is *unsuccessful*.

The refinement algorithm that interests us in this paper was first given in two dimensions by GHYS, and generalized to handle arbitrary Hamiltonian systems by Quinlan and Tremaine (QT).[20] Refinement is not the only way to look for shadows; see Refs. 26 or 27 for a review. QT make the distinction between *dynamical noise* and *observational noise*. Observational noise does not affect the future evolution of the system. Laboratory measurements of a macroscopic system are usually of this type; another example is computer output that prints fewer digits than are represented internally. In contrast, dynamical noise affects the future evolution of the system. The noise introduced by the numerical solution of a system of ordinary differential equations is dynamical.

B. The refinement procedure of GHYS

1. The algorithm

The refinement procedure of GHYS and QT can be likened to Newton's method for finding a zero of a function. The basic idea is as follows. Let $\mathbf{Y}=\{\mathbf{y}_i\}_{i=0}^{S}$ be a trajectory with S steps that has noise $\delta>\eta\epsilon_{\text{mach}}$, where ϵ_{mach} is the machine precision, and η is some constant significantly greater than 1 which allows room for improvement toward

the machine precision. Let $\mathbf{e}_{i+1}=\mathbf{y}_{i+1}-\varphi(\mathbf{y}_i)$ be the 1-step error at step $i+1$, where $\|\mathbf{e}_{i+1}\|<\delta$ for all i. The set of 1-step errors is represented by $\mathbf{E}=\{\mathbf{e}_i\}_{i=1}^{S}$, and is estimated by a numerical technique that has higher accuracy than used to compute $\mathbf{Y}$. (Note that the list of errors has one less element than the list of points in the noisy trajectory, because a point on the noisy trajectory only has an error defined if has a predessor point in the trajectory.) The process of computing the one-step errors across the entire noisy trajectory defines a function, call it g, that takes as input the entire orbit $\mathbf{Y}$ and produces an output equal to the set of 1-step errors $\mathbf{E}$, that is, $g(\mathbf{Y})=\mathbf{E}$. Because the 1-step errors are assumed to be small, $\|\mathbf{E}\|$ is small. That is, $\mathbf{Y}$ may be close to a zero of g, if one exists. A zero of g would represent an orbit with zero 1-step error, that is, an exact orbit. This situation is an ideal one in which to use Newton's method. If Newton's method converges, then an exact orbit has been found. Refinement is not exactly an example of Newton's method, but it is very similar.

Assume that we have a noisy n-dimensional orbit with S steps, $\mathbf{Y}=\{\mathbf{y}_i\}_{i=0}^{S}$, $\mathbf{y}_i\in\mathbf{R}^n$, and it has a shadow $\{\mathbf{x}_i\}_{i=0}^{S}$, $\mathbf{x}_i\in\mathbf{R}^n$. Then $\mathbf{x}_{i+1}=\varphi(\mathbf{x}_i)$ and $\mathbf{y}_{i+1}=\tilde{\varphi}(\mathbf{y}_i)=\varphi(\mathbf{y}_i)+\mathbf{e}_{i+1}$, where $\tilde{\varphi}$ is an approximation to φ with noise bounded by δ. Now suppose we approximate the one-step errors $\mathbf{e}_{i+1}=\mathbf{y}_{i+1}-\varphi(\mathbf{y}_i)$ using an integration method with noise significantly less than δ. Let $\hat{\mathbf{c}}_i\equiv\mathbf{x}_i-\mathbf{y}_i$ represent a correction term that perturbs $\mathbf{y}_i$ toward $\mathbf{x}_i$. Then

$$\hat{\mathbf{c}}_{i+1}=\mathbf{x}_{i+1}-\mathbf{y}_{i+1}=\varphi(\mathbf{x}_i)-\varphi(\mathbf{y}_i)-\mathbf{e}_{i+1}$$
$$=D\varphi(\mathbf{y}_i)\hat{\mathbf{c}}_i-\mathbf{e}_{i+1}+O(\|\hat{\mathbf{c}}_i\|^2). \tag{1}$$

In the spirit of Newton's method, we ignore the $O(\|\hat{\mathbf{c}}_i\|^2)$ term, and so one refinement iteration defines the corrections along the entire orbit:

$$\mathbf{c}_{i+1}:=D\varphi(\mathbf{y}_i)\mathbf{c}_i-\mathbf{e}_{i+1}. \tag{2}$$

For a discrete map, $D\varphi(\mathbf{y}_i)$ is just the Jacobian of the map at step i. For a system of ordinary differential equations, $D\varphi(\mathbf{y}_i)$ is the Jacobian of the solution of the ordinary differential equation from step i to step $i+1$.[28] In other words, let

$$\mathbf{y}'=\mathbf{f}(\mathbf{y}(t)) \tag{3}$$

be the first-order ordinary differential equation. Note that $\mathbf{y}_{i+1}=\varphi(\mathbf{y}_i)$ is the solution of Eq. (3) using $\mathbf{y}_i$ as the initial condition and integrating $\mathbf{f}$ to time t_{i+1}. The Jacobian $D\mathbf{f}(\mathbf{y}_i)$ measures how $\mathbf{y}'$ changes if $\mathbf{y}$ is changed by a small amount. The *resolvent* $R(t_{i+1},t_i)$ is the integral of $D\mathbf{f}(\mathbf{y})$ along the path $\mathbf{y}(t)$, and describes how a small perturbation $\delta\mathbf{y}$ of $\mathbf{y}_i$ at time t_i is mapped to a perturbation of $\mathbf{y}_{i+1}$ at time t_{i+1}. $R(t_{i+1},t_i)$ is the solution of the *variational equation*,

$$\frac{\partial R}{\partial t}=D\mathbf{f}(\mathbf{y}(t))R(t,t_i),\quad R(t_i,t_i)=I, \tag{4}$$

where I is the identity matrix. The reason the arguments to R seem to be reversed is for notational convenience: they satisfy the identity $R(t_2,t_0)=R(t_2,t_1)R(t_1,t_0)$, and so a perturbation $\delta\mathbf{y}$ at time t_0 gets mapped to a perturbation at time t_2 by the matrix–matrix and matrix–vector multiplication $R_2\delta\mathbf{y}=R_1R_0\delta\mathbf{y}$.[29] Finally, the linear map in the GHYS refinement procedure is $D\varphi(\mathbf{y}_i)=R(t_{i+1},t_i)$, where φ is the time-h solution operator for Eq. (3).

For brevity, we set $L_i\equiv D\varphi(\mathbf{y}_i)$, and for simplicity, we assume an $n=2$ dimensional problem for the remainder of this section.

If the system were not chaotic, the correction terms $\mathbf{c}_i$ could be computed directly from Eq. (2). But because L_i will amplify any errors in $\mathbf{c}_i$ that occur near the unstable direction, computing the $\mathbf{c}_i$'s by iterating Eq. (2) quickly produces nothing but noise; iterating backward suffers the same problem in the stable direction. Therefore, GHYS split the error and correction terms into components in the stable ($\mathbf{s}_i$) and unstable ($\mathbf{u}_i$) directions at each time step:

$$\mathbf{e}_i=e_{u_i}\mathbf{u}_i+e_{s_i}\mathbf{s}_i \tag{5}$$

$$\mathbf{c}_i=c_{u_i}\mathbf{u}_i+c_{s_i}\mathbf{s}_i. \tag{6}$$

Because it is not known *a priori* which direction is unstable at each time step, the unstable vector $\mathbf{u}_0$ at time t_0 is initialized to an arbitrary unit vector. The linearized map is then iterated forward with

$$\overline{\mathbf{u}}_{i+1}=L_i\mathbf{u}_i,\quad \mathbf{u}_{i+1}=\overline{\mathbf{u}}_{i+1}/|\overline{\mathbf{u}}_{i+1}|. \tag{7}$$

Because L_i magnifies any component that lies in the unstable direction, and assuming we are not so unlucky as to choose a $\mathbf{u}_0$ that lies precisely along the stable direction, then after a few Lyapunov times, $\mathbf{u}_i$ will point roughly in the actual unstable direction. Similarly, the stable unit direction vectors $\mathbf{s}_i$ are computed by initializing $\mathbf{s}_S$ to an arbitrary unit vector and iterating backward,

$$\overline{\mathbf{s}}_i=L_i^{-1}\mathbf{s}_{i+1},\quad \mathbf{s}_i=\overline{\mathbf{s}}_i/|\overline{\mathbf{s}}_i|, \tag{8}$$

where L_i^{-1} is most efficiently computed by inverting L_i. (We also could integrate the resolvent backward as described in Ref. 28, but this procedure is more expensive.)

If we substitute Eqs. (5) and (6) into Eq. (2), we obtain

$$c_{u_{i+1}}\mathbf{u}_{i+1}+c_{s_{i+1}}\mathbf{s}_{i+1}=L_i(c_{u_i}\mathbf{u}_i+c_{s_i}\mathbf{s}_i)$$
$$-(e_{u_{i+1}}\mathbf{u}_{i+1}+e_{s_{i+1}}\mathbf{s}_{i+1}). \tag{9}$$

For the same reason that L_i magnifies errors in the unstable direction, it damps errors in the stable direction. Likewise, L_i^{-1} damps errors in the unstable direction and magnifies errors in the stable direction. Thus the c_u terms should be computed in reverse order, and the c_s terms in forward order. If we take components of Eq. (9) in the unstable direction at step $i+1$ (recall that $L_i\mathbf{u}_i=\overline{\mathbf{u}}_{i+1}$ lies in the same direction as $\mathbf{u}_{i+1}$), we can iterate backward on

$$c_{u_i}=(c_{u_{i+1}}+e_{u_{i+1}})/|\overline{\mathbf{u}}_{i+1}|. \tag{10}$$

Similarly, if we take components in the stable direction [cf. Eq. (2)], we can iterate forward on

$$c_{s_{i+1}}=|L_i\mathbf{s}_i|c_{s_i}-e_{s_{i+1}}. \tag{11}$$

The initial choices for c_{s_0} and c_{u_S} are arbitrary as long as they are small—smaller than the desired maximum allowable shadowing distance because Eq. (11) damps initial conditions, and Eq. (10) damps final conditions. QT and GHYS choose them both as 0. This choice is probably as good as any, but it can be seen here that if one shadow exists, there are infinitely many of them. Another way of looking at these initial choices for c_{s_0} and c_{u_S} is that they "pinch" the growing components at the final end point and the backward-

growing components at the initial point, to be small so that $(\mathbf{y}_0-\mathbf{x}_0)\cdot\mathbf{s}_0=0$ and $(\mathbf{y}_S-\mathbf{x}_S)\cdot\mathbf{u}_S=0$. That is, the boundary conditions are being forced on the problem so that the exponential divergence is forcibly masked, if possible, making the solution of Eq. (2) numerically stable.

Note that these boundary conditions allow the initial conditions for the shadow and noisy orbits to differ along the unstable direction. In fact, this *must* be the case if the change in initial conditions is to have any effect. That is, when looking for a shadow, the perturbations in the stable direction would die out, leading the ''shadow'' to follow the exact orbit that passes through the initial conditions—the one that is already known to diverge exponentially from the computed (noisy) orbit.

2. Discussion of the GHYS algorithm

There is no guarantee that refinement converges toward an exact orbit; if there were, then all noisy orbits would be shadowable! In fact, even if some refinements are successful, numerical refinement alone does not prove rigorously that an *exact* shadow exists; it only proves the existence of a numerical shadow, that is, a trajectory that has less noise than the original. Furthermore, the 1-step error $\mathbf{e}_{i+1}$ computed by any numerical technique measures the difference between the noisy and more accurate solutions at time step $i+1$, where both start from the same position at time step i. This is different than computing the actual error per step. One must therefore have faith in the accuracy of the ''more accurate'' integrator to have faith in the numerical shadow.

Furthermore, a numerical integration routine that is requested to integrate to a tolerance close to the machine precision might not achieve it, because it might undetectably lose a few digits near the machine precision. Thus, even when a numerical shadow is found with 1-step errors claimed to be as small as 10^{-15} by the ''accurate'' integrator, the actual 1-step errors may be closer to 10^{-12}. GHYS provide a method called *containment* that can prove rigorously when an exact shadow exists, but containment (generalized for arbitrary dimension ordinary differential equations in Refs. 27 and 26) is far too expensive for the N-body problem. As a surrogate to containment, QT did experiments on simple chaotic maps with 100-digit accuracy (using the Maple symbolic manipulation package[30]) showing that if the GHYS refinement procedure refined the trajectory to 1-step errors of about 10^{-15}, then successful refinements could be continued down to 10^{-100}. It is probably reasonable to assume that refinement would continue to decrease the noise, converging in the limit to a noiseless (exact) trajectory. For these reasons, we are confident that convergence to a numerical shadow implies, with high probability, the existence of an exact shadow.

There also is no guarantee that, even if the refinement procedure converges, that it converges to a reasonable shadow of $\mathbf{Y}$; in principle, it could converge to an exact orbit that is far from $\mathbf{Y}$, in which case the term ''shadow'' would be inappropriate. However, in practice it appears that, if refinement fails, it always fails due to the 1-step errors becoming large. This error explosion occurs when the 1-step errors are so large that the linearized map becomes invalid for computing corrections. Because the method is global (that is, each correction depends on all the others), inaccuracies in the computation of the corrections can quickly amplify the noise rather than decreasing it. Thus, within 1 or 2 refinement iterations, the 1-step errors can grow by many orders of magnitude, resulting in failed refinements.

An important open question is whether shadows are typical of exact orbits chosen at random. Simple examples exist of shadows that are atypical,[31–33] although it seems unlikely that atypical shadows are common—lest the numerical solutions we compute would be commonly atypical as well. An extensive discussion of this issue can be found in Ref. 27.

III. SHADOWING THE GRAVITATIONAL *N*-BODY PROBLEM

> ''It is in vaine to goe about to make the shadowe straite, if the bodie whiche giveth the shadowe bee crooked,'' Stefano Guazzo.[34]

Although not all numerical simulations are likely to be shadowable, the existence of a shadow is a strong property: it asserts that a numerical solution can be viewed as an *experimental observation* of an exact solution. As such, within the ''observational'' error, the dynamics observed in a numerical solution that has a shadow represent the dynamics of an exact solution. We consider the existence of a shadow to be the ''gold standard'' of reliability for simulations of chaotic systems. For the systems we consider, a shadow lasting several tens of crossing times is sufficient. For example, because its formation about 10^{10} years ago, our Milky Way galaxy has rotated only about 40 times at the orbital radius of our Sun; a shadow of a numerical simulation of our galaxy lasting as long would be more than sufficient.

If we think of shadowing as a measure of the error of a simulation—with the relevant measures being how close the shadow is and how long it lasts—then we also must emphasize when shadowing is an appropriate measure of the error. The answer is that it only makes sense to use shadowing if the model being simulated accurately reflects the system being studied. Shadowing is thus best applied to systems in which the governing equations are well known, and virtually all error is introduced by the imprecise knowledge of initial conditions or by numerical error in the computation of the solution. It is less applicable to systems in which the model only approximate. For example, shadowing is an appropriate measure of error for the gravitational N-body problem, because the equations of motion are extremely well understood, and an exact solution of the model very closely approximates the behavior of a real system under the assumed conditions of the model. Conversely, shadowing is an inappropriate measure of error for a weather simulation, because the models are known to be only rough approximations and having a numerical solution that closely follows an exact solution of such an approximate model is of a dubious value.

To apply refinement to the N-body problem, the method needs to be generalized to work in more than two dimensions. The extensions were introduced by QT and are described in Refs. 35 and 27. QT undertook the first study of shadows of numerical simulations of the N-body problem, and found that a single particle moving in the potential of 100 fixed particles can be shadowed for a few tens of crossing times. Hayes[36] demonstrated that the distribution of shadowing lengths for such a system is fit well by an exponential curve, suggesting that glitches in one-particle trajectories are encountered as a Poisson process in an unsoftened system. This observation is intriguing because a Poisson pro-

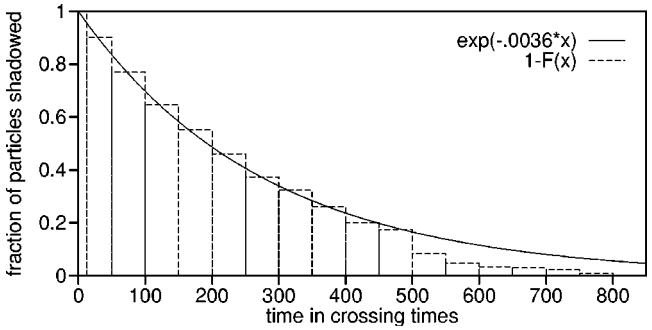

Fig. 2. The estimated fraction of particles that would be shadowed as a function of time in a system of 100 particles in which the gravitational force-softening parameter ϵ is 1/4 of the mean inter-particle spacing. Noisy trajectories were generated using a leapfrog integrator with a time step that allows particles to travel an average distance of $\epsilon/5$ per time step. Each particle has an expected shadow duration of 280 crossing times, equivalent to having a glitch probability of 0.36% per crossing time. The number of valid (nonglitched) particles decays exponentially, assuming each particle encounters glitches as a Poisson process independent of all the other particles.

cess is memoryless,[37] which means that the history of the orbit has no effect on the glitch probability. This result is consistent with the view of Ref. 23, that a glitch is a *sudden* occurrence independent of the history of the orbit, and not the result of a long-term buildup of error.

Shadowing large systems is extremely expensive. Straightforward integration of the variational equation for an N-body system requires $O(N^3)$ per time step. Because a typical galaxy or cosmological simulation today contains millions of particles, we would not be capable of shadowing the trajectories of such a simulation for a long time. However, shadowing can still provide significant insight into the reliability of such simulations for two reasons. First, in a large system, the motion of particles is governed far more by the global potential than by the positions of nearby particles,[16] and second, individual particles appear to encounter glitches independently of one another.[36,38] These observations suggest that if one particle encounters a glitch, then it will have a negligible effect on the motion of the others, at least at first, because one errant particle does not appreciably change the global gravitational potential created by the (millions of) other particles. In fact, the global potential will remain substantially valid as long as most particles are on valid trajectories. This interpretation implies that the validity of a simulation slowly degrades as the number of glitched particles increases, infecting the field and the motion of other particles. The question is how many particles remain on valid trajectories as the simulation time progresses?

If each particle encounters glitches independently as a Poisson process, then we would expect the number of nonglitched particles to decay exponentially with time. Figure 2 plots this fraction of nonglitched (that is, shadowed) particles as a function of time for a set of parameters described more fully in Ref. 38. As Fig. 2 shows, for these simulation parameters over 90% of particles still have valid trajectories after about 50 crossing times. Thus, a simulation of a galaxy for 50 crossing times (which is about the age of our Milky Way Galaxy) with these parameters would probably have more than 90% of its particles in the "correct" place, within "observational" (that is, shadow distance) error. This is a very encouraging sign that perhaps such galaxy simulations are reliable in the shadowing sense.

IV. EXERCISES

> "On voit courir après l'ombre Tant de fous qu'on n'en sait pas." ("One sees chasing after shadows more fools than one can count.")—La Fontaine.[39]

(1) Write a program to implement refinement in two dimensions. Your program should first create a noisy trajectory of a two-dimensional map, and then pass the entire trajectory to the refinement subroutine, which returns, if possible, a new trajectory with the noise as small as possible. You can find the longest shadow by trying to shadow longer and longer segments of the noisy trajectory. You will need to calculate (by hand) the Jacobian of the map to compute the variational equation. The following two-dimensional hyperbolic maps whose trajectories you can try to shadow are taken from Chapter 2 of Ref. 40. The second problem exercise is suitable as an assignment for lower year courses; the second is more appropriate for upper years; and the third is more suitable for graduate students.

(2) The Hénon map,
$$(x_{i+1},y_{i+1}) := (a - x_i^2 + by_i, x_i). \tag{12}$$
The classical parameter values used by Hénon are $a=1.28, b=0.3$; another interesting set is $a=1.4, b=-0.3$. A good initial condition is $(x_0,y_0)=(0,0)$.

(3) The map,
$$(r_{i+1},\theta_{i+1}) := (r_i^2, 2\theta), \tag{13}$$
is derived from the polar coordinate form of the complex map $f(z)=z^2$. It can be shown that there exist initial conditions that produce chaotic trajectories of this map that are bounded, have a positive Lyapunov exponent, and do not converge to the point (0,0).[40] Find such an initial condition (using either the paper and pencil method described in Ref. 40, or by experimenting with a large number of random initial conditions), and see how long you can shadow it.

(4) The time-2π map of the forced-damped pendulum,
$$(\theta_{i+1},\dot{\theta}_{i+1}) = F(\theta_i,\dot{\theta}_i), \tag{14}$$
is found by integrating
$$\ddot{\theta} = -c\dot{\theta} - \sin\theta + \rho\sin t, \tag{15}$$
from 0 to 2π using an integer number of time steps of size $2\pi/k$ for some k. One integration from 0 to 2π constitutes one iteration of the map; a new integration must be performed for each iteration of the map. This integration allows us to use the map shadowing method described in this paper to shadow trajectories of an ordinary differential equation. We see that shadowing an ordinary differential equation is much more expensive than shadowing a map, because a map such as the Hénon map does not require an ordinary differential equation integration at each step. Note that to compute the Jacobian of this map, you actually need to compute a resolvent of the ordinary differential equation for one time step as described in Ref. 28. Interesting parameter values include $c=0.05$, $\rho=2.5$) and $c=0.2$, $\rho=1.66$. Try various randomly generated initial conditions in $(\theta_0,\dot{\theta}_0) \in ([-\pi,\pi],[-2,4])$.

[1]This quotation is frequently attributed to Mark Twain, but in the Chronicle of Higher Education (4 September 1991, p. A8), the group at the University of California, Berkeley (⟨http://bancroft.berkeley.edu/MTP/⟩), which is editing all of Twain's papers, put this one into the list of quotations that Twain did not say.
[2]D. A. Clarke and M. J. West, in *The 12th Kingston Meeting: Computational Astrophysics* (ASP Conference Series, Astron. Soc. of the Pacific, 1997), Vol. 123.
[3]E. Bertschinger, "Simulations of structure formation in the universe," Annu. Rev. Astron. Astrophys. **36**, 599–654 (1998).
[4]J. Goodman, D. C. Heggie, and P. Hut, "On the exponential instability of N-body systems," Astrophys. J. **415**, 715–733 (1993).
[5]R. H. Miller, "Irreversibility in small stellar dynamical systems," Astrophys. J. **140**, 250–256 (1964).
[6]M. Lecar, "A comparison of eleven numerical integrations of the same gravitational 25-body problem," Bull. Astron. **3**, 91 (1968).
[7]H. E. Kandrup and H. Smith, "On the sensitivity of the N-body problem to small changes in initial conditions," Astrophys. J. **374**, 255–265 (1991).
[8]H. E. Kandrup, H. Smith, and D. Willmes, "On the sensitivity of the N-body problem to small changes in initial conditions. III," Astrophys. J. **399**, 627–633 (1992).
[9]L. Hernquist, P. Hut, and J. Makino, "Discreteness noise versus force errors in N-body simulations," Astrophys. J. Lett. **402**, L85–L88 (1993).
[10]J. E. Barnes and P. Hut, "Error analysis of a tree code," Astrophys. J., Suppl. Ser. **70**, 389–417 (1989).
[11]J. A. Sellwood, "The art of N-body building," Annu. Rev. Astron. Astrophys. **25**, 151-86 (1987).
[12]J. P. Singh, J. L. Hennessy, and A. Gupta, "Implications of hierarchical N-body methods for multiprocessor architecture," Technical Report CSL-TR-92-506, Computer Systems Lab, Stanford University, Stanford, CA 1992.
[13]D. Pfenniger and D. Friedli, "Computational issues connected with 3D N-body simulations," Astron. Astrophys. **270**, 561–572 (1993).
[14]J. G. Jernigan and D. H. Porter, "A tree code with logarithmic reduction of force terms, hierarchical regularization of all variables, and explicit accuracy controls," Astrophys. J., Suppl. Ser. **71**, 871–893 (1989).
[15]J. Barnes and P. Hut, "A hierarchical $O(N \log N)$ force-calculation algorithm," Nature (London) **324**, 446–449 (1986).
[16]J. Binney and S. Tremaine, *Galactic Dynamics* (Princeton University Press, Princeton, NJ, 1987).
[17]L. Greengard and V. Rokhlin, "A fast algorithm for particle simulation," J. Comput. Phys. **73**, 325–348 (1987).
[18]P. J. Channell and C. Scovel, "Symplectic integration of Hamiltonian Systems," Nonlinearity **3**, 231–259 (1990).
[19]J. M. Sanz-Serna, "*Symplectic Integrators for Hamiltonian Problems: An Overview*," Acta Numerica 1992, edited by A. Iserles (Cambridge University Press, Cambridge, 1992), pp. 243–286.
[20]G. D. Quinlan and S. Tremaine, "On the reliability of gravitational N-body integrations," Mon. Not. R. Astron. Soc. **259**, 505–518 (1992).
[21]This is in stark contrast to N-body simulations of our solar system, in which case we are interested in the precise evolution of the particular solar system we live in, not a hypothetical one very similar to it.
[22]Speech at Bristol on Declining the Poll Vol. ii, p. 420. (1780) from *Macmillan Book of Proverbs, Maxims, and Famous Phrases* (Macmillan, New York, 1965), p. 2080 (quotes about shadows).
[23]C. Grebogi, S. M. Hammel, J. A. Yorke, and T. Sauer, "Shadowing of physical trajectories in chaotic dynamics: Containment and refinement," Phys. Rev. Lett. **65**, 1527–1530 (1990).
[24]D. V. Anosov, "Geodesic flows and closed Riemannian manifolds with negative curvature," Proc. Steklov Inst. Math. **90**, 1–235 (1967).
[25]R. Bowen, "ω-limit sets for axiom A diffeomorphisms," J. Diff. Eqns. **18**, 333–339 (1975).
[26]W. B. Hayes and K. R. Jackson, "Rigorous shadowing of numerical solutions of ordinary differential equations by containment," SIAM (Soc. Ind. Appl. Math.) J. Numer. Anal. **41**, 1948–1973 (2003).
[27]W. B. Hayes, "Rigorous shadowing of numerical solutions of ordinary differential equations by containment," Ph.D. thesis, Department of Computer Science, University of Toronto, 2001. Available at ⟨http://www.cs.toronto.edu/NA/reports.html#hayes-01-phd⟩.
[28]In other words, let $\mathbf{y}' = \mathbf{f}(\mathbf{y}(t))$ be a first-order ordinary differential equation. Note that $\mathbf{y}_{i+1} = \varphi(\mathbf{y}_i)$ is the solution of Eq. (3) using $\mathbf{y}_i$ as the initial condition and integrating $\mathbf{f}$ to time t_{i+1}. The Jacobian $D\mathbf{f}(\mathbf{y}_i)$ measures how $\mathbf{y}'$ changes if $\mathbf{y}$ is changed by a small amount. The *resolvent* $R(t_{i+1}, t_i)$ is the integral of $D\mathbf{f}(\mathbf{y})$ along the path $\mathbf{y}(t)$, and describes how a small perturbation $\delta\mathbf{y}$ of $\mathbf{y}_i$ at time t_i is mapped to a perturbation of $\mathbf{y}_{i+1}$ at time t_{i+1}. $R(t_{i+1}, t_i)$ is the solution of the *variational equation*,

$$\frac{\partial R}{\partial t} = D\mathbf{f}(\mathbf{y}(t))R(t, t_i), \quad R(t_i, t_i) = I,$$

where I is the identity matrix. The reason the arguments to R seem to be reversed is for notational convenience: they satisfy the identity $R(t_2, t_0) = R(t_2, t_1)R(t_1, t_0)$, and so a perturbation $\delta\mathbf{y}$ at time t_0 gets mapped to a perturbation at time t_2 by the matrix–matrix and matrix–vector multiplication $R_2\delta\mathbf{y} = R_1 R_0 \delta\mathbf{y}$.[29] Finally, the linear map in the GHYS refinement procedure, if φ is the time-h solution operator for Eq. (3), is $D\varphi(\mathbf{y}_i) = R(t_{i+1}, t_i)$.
[29]E. Hairer, S. P. Nørsett, and Gerhard Wanner, *Solving Ordinary Differential Equations*, 2nd ed. (Springer-Verlag, Berlin, 1993), Vols. 1 and 2.
[30]B. W. Char, *First Leaves: A Tutorial Introduction to Maple V* (Springer-Verlag, Berlin, 1993).
[31]J. D. Farmer and J. J. Sidorowich, "Optimal shadowing and noise reduction," Physica D **47**, 373–392 (1991).
[32]S. T. Fryska and M. A. Zohdy, "Computer dynamics and the shadowing of chaotic orbits," Phys. Lett. A **166**, 340–346 (1992).
[33]R. M. Corless, "What good are numerical simulations of chaotic dynamical systems?," Comput. Math. Appl. **28**, 107–121 (1994).
[34]S. Guazzo, *Civile Conversation*, Book ii (1574), p. 231.
[35]W. Hayes, "Efficient shadowing of high dimensional chaotic systems with the large astrophysical n-body problem as an example," Master's thesis, Department of Computer Science, University of Toronto, January 1995.
[36]W. Hayes, "Shadowing high-dimensional Hamiltonian systems: The gravitational n-body problem," Phys. Rev. Lett. **90**, 054104-1–054104-4 (2003).
[37]J. A. Rice, *Mathematical Statistics and Data Analysis* (Wadsworth & Brooks/Cole, 1988).
[38]W. Hayes, "Shadowing-based reliability decay in softened n-body simulations," Astrophys. J. Lett. **587**, L59–L62 (2003).
[39]La Fontaine, *Fables*, Book vi, Fable 17, 1668.
[40]K. T. Alligood, T. D. Sauer, and J. A. Yorke, *Chaos: An Introduction to Dynamical Systems* (Springer-Verlag, Berlin, 1996).