APS/123-QED

Shadowing high-dimensional Hamiltonian systems: the gravitational N-body problem

Wayne B. Hayes*
Department of Computer Science, University of Toronto, Toronto, Ontario, M5S 3G4 Canada

(Dated: December 21, 2002)

A *shadow* is an exact solution to a chaotic system of equations that remains close to a numerically computed solution for a long time. Using a variable-order, variable-timestep integrator, we numerically compute solutions to a gravitational N-body problem in which many particles move and interact in a fixed potential. We then search for shadows of these solutions with the longest possible duration. We find that in "softened" potentials, shadow durations are sufficiently long for significant evolution to occur. However, in unsoftened potentials, shadow durations are typically very short.

PACS numbers: 02.60.-x, 02.70.-c, 05.45.-a, 05.45.Jn, 95.75.P, 05.40.-a, 98.10.+z

The astronomical literature is brimming with the results of gravitational N-body simulations. Examples include studies of the formation, evolution and structure of galaxies and clusters of galaxies [1], and the cosmos at large [2]. Considering that such simulations have been used to invalidate theories [2], establishing their trustworthiness is critical. Like many dynamical systems, however, a gravitational system displays *sensitive dependence on initial conditions* (SDIC): two solutions whose initial conditions differ by an arbitrarily small amount generally diverge exponentially away from each other [3]. Since numerical methods introduce errors, it is virtually guaranteed that a numerically computed solution diverges exponentially away from the exact solution with the same initial conditions. This remains true even if integrals of motion such as energy and momentum are conserved to arbitrary precision. The phenomenon has been described (*eg.,* [3]) 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 simulation reliability, and although SDIC is widely known to be one of these aspects, the impact on simulation reliability of SDIC is not well understood.

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 behaviour. In such cases, we typically choose initial conditions from a random distribution and would be happy if our numerical solution exhibited behaviour typical of *any* valid choice of initial conditions from our distribution. In particular, we may 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 same set of equations for a nontrivial duration of time, *i.e.,* significantly longer than would the exact solution starting at the same initial condition as the numerical solution. Although not all numerical simulations are likely to be shadowable [4–6], 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. There are only two remaining questions (both beyond the scope of this paper): (1) Whether the mathematical model being simulated accurately reflects the system being studied. This is certainly *not* the case for systems such as the weather (and thus shadowing is probably an inappropriate measure of error), but can be assumed to be the case for others, such as the unsoftened gravitational N-body problem. (2) Whether shadows are typical of exact solutions chosen at random. Simple examples exist of shadows that are atypical [7–9], although it seems unlikely that atypical shadows are common—lest the numerical solutions we compute would be commonly atypical as well [18].

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* (the time it takes a particle to cross the system from one side to another) is sufficient. For example, since 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.

The first proof of the existence of a shadow of a computer-generated numerical trajectory was provided by [10]. The first study of shadows of numerical simulations of the N-body problem was undertaken by [11], who found that a single particle moving in the potential of 100 fixed particles could be shadowed for a few tens of crossing times. However, [11] were unable to predict the behaviour of systems with more moving particles because shadowing is computationally very expensive, taking time $O(M^3)$ for M moving particles.

In this paper, we extend the results of [11] by increas-

*Electronic address: `wayne@cs.toronto.edu`

ing the number of moving particles, having greatly optimized their shadowing algorithm. We find the following. First, that if the gravitational potential is "softened" and the number of moving particles M is increased from 1 to 25, shadow durations tend to decrease with M slowly enough that shadows of large N-body simulations may exist for nontrivial durations of time. Second, we demonstrate that M *one*-moving-particle systems can be used to approximately predict the shadow duration of an M-moving-particle system. This is done by postulating that a *glitch* (the point beyond which no shadow can be found) in the large system occurs when *one* of its moving particles encounters a glitch local to its own trajectory, and then showing that the same process can be approximated by superimposing M one-moving-particle trajectories and taking the *minimum* shadow duration of those trajectories. The approximation is excellent for unsoftened systems, and reasonable for softened ones. This reduces the amount of computation required to study shadows of large systems from $O(M^3)$ to $O(M)$, which *greatly* facilitates the study of large systems.

The trajectories that we will attempt to shadow belong to a slightly simplified gravitational N-body problem in which there are N total particles, only M of which move. We do this because the shadowing algorithm we use takes time $O(M^3)$ and we want to simulate a large system with a complex potential while keeping the time to compute a shadow tractable. We will refer to such a system as an M-body system. Each moving particle interacts both with fixed particles and with other moving particles via Newton's gravitational force law, $F_{ij} = -(Gm_i m_j)/(r_{ij}^2 + \varepsilon^2)$, where F_{ij} is the force on particle i from particle j, m_i and m_j are their masses, r_{ij} is the distance between them, and ε is the *gravitational softening parameter* which, if non-zero, artificially smoothens the gravitational potential in order to approximately emulate a system with more particles than are actually present and to avoid the singularity at $r_{ij} = 0$. We use normalized units [12] in which each particle has mass $1/N$, the system has diameter of order unity, and the crossing time is of order unity. We use a variable-order, variable timestep integrator [13] for all integrations. We generate noisy trajectories with local errors of about 10^{-5} per crossing time. To find shadows, we use an algorithm described in [11], optimized to run between two and three orders of magnitude faster. Called *iterative refinement*, we use the same integrator as the noisy trajectory with tighter tolerance to estimate the full phase-space vector of local errors of the noisy trajectory, and then use a Newton-like correction to refine the trajectory until it has local errors as small as possible. For simple systems, the errors of the refined trajectory can be as small as the machine precision (10^{-16}), but the minimum local error achievable with refinement increases as the number of dimensions increases, due to numerical errors in computing the Newton corrections. A trajectory produced by refinement is called a *numerical shadow*. The existence of a numerical shadow is expected to indicate the existence of an exact shadow of comparable duration [11].

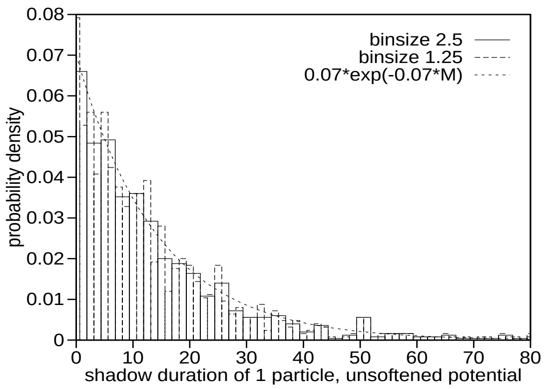

FIG. 1: Histogram of shadow durations of 1000 unsoftened M-body systems. Each system has 99 fixed particles and one moving particle (*i.e.,* $M = 1$). Noisy orbits have local error of about 10^{-5} per crossing time; numerical shadows were required to have a maximum local error no bigger than 10^{-14}. The horizontal axis is in crossing times; the vertical axis is the measured probability density. The distribution fits an exponential curve with a mean glitch rate of about 0.07 per crossing time, indicating that the moving particle encounters glitches as a Poisson process in an unsoftend system.

If a shadow is viewed as a measure of error of a numerical solution, then the relevant measures are the phase-space distance between corresponding points on the "noisy" and exact trajectories (smaller is better), and the duration over which they remain close together (longer is better). Generally, the smaller the local error in the trajectory, the closer and longer the shadow [10, 11]. In this paper, the noisy-shadow distance is typically less than 1% of the size of the system.

We present our results of the unsoftened case first, because it is instructive and generalizes more elegantly to $M > 1$ than the softened case. Fig. 1 introduces a histogram of shadow durations for 1000 systems each with $N = 100, M = 1, \varepsilon = 0$. The distribution is well-fit by an exponential curve, suggesting that glitches in one-particle trajectories are encountered as a Poisson process in an unsoftened system [14]. This is intriguing because a Poisson process is memoryless, which means that the history of the orbit has no effect on the glitch probability. This is consistent with the view of [10], 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.

Fig. 2 introduces how the average shadow duration scales as the number of moving particles is increased. We study two cases: (i) a $99 + M$ particle system in which the motion of all particles is simulated simultaneously as a single system but the moving particles do not interact, and (ii) a 100-particle system in which M particles move and interact. Surprisingly, the two curves are statistically indistinguishable, suggesting that although particles interact in motion, they do not interact to cause glitches at a rate any different than that of the system with non-interacting particles. A possible explanation is that if

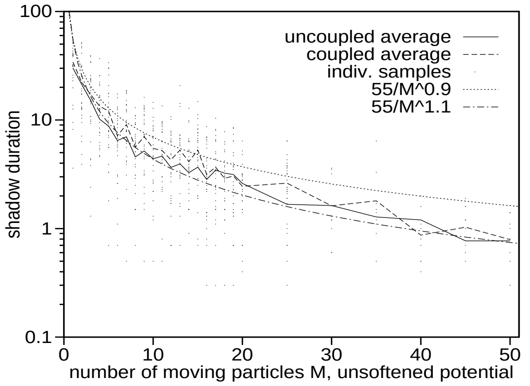

FIG. 2: Average shadow durations in crossing times of an unsoftened gravitational M-body system in which there are 100 particles, M of which move, as a function of M. The noisy orbits have local error of 10^{-5} per crossing time; each numerical shadow was required to have a maximum local error no bigger than 10^{-12}. The dots represent sample shadow durations, 30 samples each for $M = 1, 2, \ldots, 19, 20, 25$ and 10 samples each for $M = 30, 35, \ldots, 50$. The "coupled average" line joins their averages, while $55/M^{0.9}$ and $55/M^{1.1}$ are plotted for comparison. The "coupled average" line is statistically indistinguishable from the "uncoupled average" one, in which the gravitational interaction between moving particles is deleted. This suggests that even coupled particles encounter glitches independently of one another.

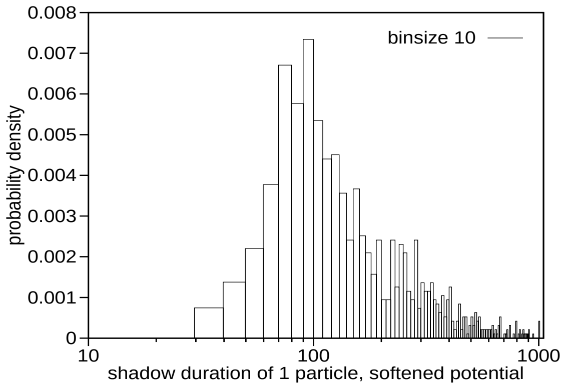

FIG. 3: Histogram of shadow durations of 1000 systems identical to those in Fig. 1 except with softening $\varepsilon = 0.1$. Note that, in contrast to Fig. 1, the x axis is logarithmic and extends to 1000 crossing times, and that the histogram height is zero near a shadow duration of zero, meaning that *no* particles undergo glitches until several tens of crossing times have occurred; in fact, of the 1000 systems sampled, the shortest shadow lasts 33 crossing times.

$1 < M \ll N$, then the coupling between the M moving particles is weak on average and we can still view the system as the superposition of M one-particle systems. Furthermore, we note that the cross-section for close approaches is not altered as M increases (which simply changes fixed particles into moving ones), so the argument still holds independent of whether $M \ll N$.

Fig. 3 introduces the distribution of shadow durations

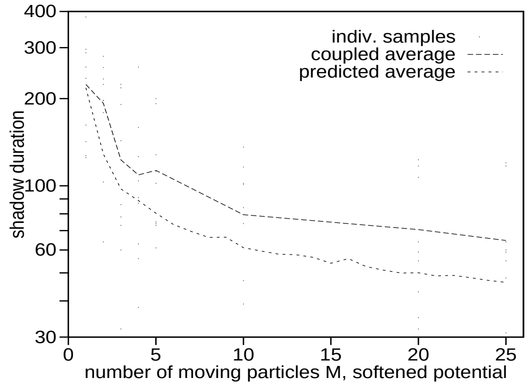

FIG. 4: Average shadow durations in crossing times of a system with softening parameter $\varepsilon = 0.1$ as a function of M. The points represent sample shadow durations, 9 samples each for M = 1–5, 10, 20, 25, while the "coupled average" line joins their averages. Following the assumption that particles encounter glitches independently of one another, the "predicted average" is artificially constructed for each $M = 1, \ldots, 25$ by superimposing M samples chosen at random from Fig. 3 and taking the minimum shadow duration of those samples. The resulting graph resembles the "coupled average" graph reasonably well, except for the surprising effect of *under*-estimating shadow duration of the real system by about $(20\pm10)\%$.

for 1000 $M = 1$ systems in which softening has been set to $\varepsilon = 0.1$, which is approximately half the average inter-particle separation. The differences from Fig. 1 are quite marked. First, the distribution is peaked near 100 crossing times and has a long tail going out to hundreds of crossing times. Second, even though the local errors of the noisy and shadow orbits are identical to those from Fig. 1, the average shadow length has increased by more than an order of magnitude from 14 to 218. Although this is roughly equivalent to the increase in the Lyapunov timescale [11], the distribution is far from exponential. In fact, the most striking difference from Fig. 1 is that the distribution has a vanishingly small density near zero shadow duration, in striking contrast to a Poisson process. In other words, *virtually no particles undergo glitches until several tens of crossing times have occurred.* If this remains true even for $M > 1$, then shadowing of softened gravitational systems would be feasible even for large M, because the trajectories of all particles in the simulation would remain valid for many crossing times. This question is addressed in Fig. 4, where we introduce the average shadow duration for softened systems as a function of M, along with the shadow duration predicted by superimposing M one-particle systems. We see that although the duration of shadows for coupled systems decreases as M increases, they decrease *much* more slowly than $1/M$, and appear to be leveling off at about 50 crossing times. This is consistent with superimposing one-particle trajectories all of which have shadows that last several tens of crossing times, although it is surprising that the shadow durations are slightly *longer* than that predicted by superimposing one-particle trajectories.

The difference between the shadow durations for softened *vs.* unsoftened systems undoubtedly is related to fluctuating Lyapunov exponents [4–6]. Although we measured no Lyapunov exponents, we measured a related quantity, namely the expansion and contraction factors across a timestep of the vectors that span the locally expanding and contracting spaces. In a uniformly hyperbolic system, these factors will always be greater and less than 1, respectively. An event of "non-hyperbolicity" can be observed by looking for areas along the trajectory where directions which were previously expanding over long periods instead start to contract, or vice versa. If we plot the expansion and contraction amounts along a trajectory for either a softened or unsoftened potential, we find that "non-hyperbolic" events correlate well with the occurrence of glitches. If we plot an $M = 1$ particle orbit in three dimensions, we also observe that these events loosely correlate with times when the particle's orbital geometry changes in an obvious way. We postulate that the locally expanding and contracting directions of a particle in the system are closely related to the geometry of the particle's orbit, so that changing the geometry of the orbit can cause these local vectors to become inconsistent as time progresses. In an unsoftened system, the geometry of a particle's orbit can be suddenly and violently changed by a close encounter. In softened systems, however, there is no precise, short-duration "event" which triggers non-hyperbolicity; instead, the geometry of the orbit of a particle changes slowly, so that many crossing times occur before a glitch is likely. This helps to explain Figs. 1 and 3, and is consistent with [11]'s finding that glitches tend to occur near close approaches in unsoftened systems, but less so in softened systems.

Now, consider an N-body simulation with one million particles, which is not uncommon today. Technically, the first particle to encounter a glitch in its own individual 6-dimensional phase-space trajectory causes a glitch in the full 6-million-dimensional phase space solution. However, the motion of particles in such a system is governed far more by the global potential than by the position of any one particle [15]. This raises the question of whether a glitch in the trajectory of just one particle out of millions is sufficient to invalidate a simulation. The answer is almost certainly "no". More likely, the validity of a simulation from the shadowing perspective probably degrades slowly, as the number of "locally-glitched" particles slowly increases, "infecting" the motion of the remaining particles [19]. Thus, we can hypothesize that reasonable statistical results may be acquired from long simulations of large softened systems as long as only a few particles have undergone glitches, and the statistics taken depend on large numbers of particles. Thus, for example, the global spatial distribution of matter in a simulated galaxy may be correct, but the number of escapers from a simulated globular cluster may be incorrect if the stars that escape happen also to be the stars that underwent glitches before escaping.

To conclude, we believe there is no feasible integration accuracy which will give long shadows in an unsoftened potential. This does not necessarily mean that such simulations are untrustworthy, only that shadowing may be too stringent a measure of error. In contrast, we believe that there exists a feasible integration accuracy for which softened systems are shadowable for many crossing times even for large N.

The author's software [16] is available on request.

Acknowledgments: We thank S. Tremaine, K. Jackson, and J. Yorke for discussions, and the University of Toronto's Computing Disciplines Facility (*CDF*) for cluster computer time. This work was supported in part by NSERC Canada, and ITRC Ontario.

[1] D. A. Clarke and M. J. West, eds., *The 12th "Kingston Meeting": Computational Astrophysics*, vol. 123 of *ASP Conference Series* (Astron. Soc. of the Pacific, 1997).
[2] E. Bertschinger, Annu. Rev. Astron. Astrophys. **35**, 599 (1998).
[3] J. Goodman, D. C. Heggie, and P. Hut, The Astrophysical Journal **415**, 715 (1993).
[4] S. Dawson, C. Grebogi, T. Sauer, and J. A. Yorke, Physical Review Letters **73**, 1927 (3 Oct 1994).
[5] T. Sauer, C. Grebogi, and J. A. Yorke, Physical Review Letters **79**, 59 (7 July 1997).
[6] T. D. Sauer, Physical Review E **65**, 036220 (2002).
[7] J. D. Farmer and J. J. Sidorowich, Physica D **47**, 373 (1991).
[8] S. T. Fryska and M. A. Zohdy, Physics Letters A **166**, 340 (1992).
[9] R. M. Corless, Computers Math. Applic. **28**, 107 (1994).
[10] C. Grebogi, S. M. Hammel, J. A. Yorke, and T. Sauer, Physical Review Letters **65**, 1527 (1990).
[11] G. D. Quinlan and S. Tremaine, Monthly Notices of the Royal Astronomical Society **259**, 505 (1992).
[12] D. C. Heggie and R. D. Mathieu, in *The Use of Supercomputers in Stellar Dynamics*, edited by P. Hut and S. L. W. McMillan (Springer-Verlag, 1986), pp. 233–235.
[13] A. C. Hindmarsh, ACM-SIGNUM Newsletter **15**, 10 (1980).
[14] J. A. Rice, *Mathematical Statistics and Data Analysis* (Wadsworth & Brooks/Cole Advanced Books & Software, 1988).
[15] J. Binney and S. Tremaine, *Galactic Dynamics*, Princeton Series in Astrophysics (Princeton University Press, 1987).
[16] W. Hayes, Master's thesis, Dept. of Computer Science, University of Toronto (1995).
[17] J. Wisdom and M. Holman, The Astronomical Journal **104**, 2022 (1992).
[18] This question is not unique to shadowing; *eg.,* is the exact solution to a nearby Hamiltonian problem produced by a symplectic integrator typical of nearby problems?
[19] This effect will be quantified in an upcoming paper.