THE ASTROPHYSICAL JOURNAL, 587:L59–L62, 2003 April 20


# SHADOWING-BASED RELIABILITY DECAY IN SOFTENED *n*-BODY SIMULATIONS

WAYNE B. HAYES
Department of Computer Science, University of Toronto, 10 King's College Road, Room 3302, Toronto, ON M5S 3G4, Canada; wayne@cs.toronto.edu
*Received 2002 November 7; accepted 2003 March 11; published 2003 March 19*

## ABSTRACT

A shadow of a numerical solution to a chaotic system is an *exact* solution to the equations of motion that remains close to the numerical solution for a long time. In a collisionless *n*-body system, we know that particle motion is governed by the global potential rather than by interparticle interactions. As a result, the trajectory of each individual particle in the system is independently shadowable. It is thus meaningful to measure the *number* of particles that have shadowable trajectories as a function of time. We find that the number of shadowable particles decays exponentially with time as $e^{-\mu t}$ and that for $\epsilon \in [\sim 0.2, 1]$ (in units of the local mean interparticle separation $\bar{n}$), there is an explicit relationship among the decay constant $\mu$, the time step $h$ of the leapfrog integrator, the softening $\epsilon$, and the number of particles $N$ in the simulation. Thus, given $N$ and $\epsilon$, it is possible to precompute the time step $h$ necessary to achieve a desired fraction of shadowable particles after a given length of simulation time. We demonstrate that a large fraction of particles remain shadowable over ~100 crossing times even if particles travel up to about $\frac{1}{3}$ of the softening length per time step. However, a sharp decrease in the number of shadowable particles occurs if the time step increases to allow particles to travel farther than $\frac{1}{3}$ the softening length in 1 time step or if the softening is decreased below $\sim 0.2\bar{n}$.

*Subject heading:* methods: *n*-body simulations

## 1. INTRODUCTION

Numerical simulations of the softened gravitational *n*-body problem are used to gain insight into the formation, evolution, and structure of gravitational systems ranging from galaxies and clusters of galaxies to the large-scale structure of the universe (Clarke & West 1997; Bertschinger 1998). Since such simulations have been used to invalidate theories (Bertschinger 1998), establishing their trustworthiness is critical. These simulations have several sources of error, including the use of several (sometimes many) orders of magnitude fewer particles than the system being modeled, or discreteness noise; the use of approximate force-computation methods (the latter two errors are compared in Hernquist, Hut, & Makino 1993); the use of a softened potential; the use of finite time step numerical integration to evolve the system of ordinary differential equations; and machine round-off error. These errors are aggravated by the fact that gravitational *n*-body systems are chaotic and display sensitive dependence on initial conditions: two solutions with nearby initial conditions diverge exponentially away from each other on about a crossing timescale (Goodman, Heggie, & Hut 1993), so that *any* error results in the numerical trajectory diverging exponentially away from the exact solution with the same initial conditions. The phenomenon has been described (e.g., Goodman et al. 1993) as the "exponential magnification of small errors," implying the possibility that trajectories of such simulations are the result of nothing but magnified noise.

Fortunately, the purpose of a softened *n*-body simulation is not to follow the evolution of a particular choice of initial conditions but instead to *sample* the evolution of large systems whose initial conditions are drawn from a random distribution. As such, we would likely be more than satisfied if our numerical solution closely followed the evolution of a nearby set of initial conditions. The study of shadowing provides just such a property: a shadow of a numerical, or noisy, solution is an exact solution whose initial conditions and subsequent evolution remain nearby, in phase space, to the numerical solution. Thus, a numerical solution that has a shadow is essentially an experimental observation of that shadow, which is an exact trajectory of the mathematical system being modeled. Although this observation does not alleviate errors introduced between the physical system and the mathematical model (such as discreteness noise and force softening), it *does* say that the numerical simulation is faithfully solving the mathematical model. Quinlan & Tremaine (1992) showed that a single particle moving among 99 fixed particles is shadowable for several tens of crossing times and that glitches (the point beyond which a shadow cannot be found) tend to occur near close encounters. Hayes (2003) demonstrated that if $M > 1$ particles move in a softened system with $100 - M$ fixed particles, then very few particles encounter glitches within the first few tens of crossing times. However, both of these studies used highly accurate integrators to generate the "noisy" trajectories. Although high accuracy is commonly used for simulations of unsoftened systems, softened simulations most often use the second-order symplectic and time-symmetric leapfrog integrator.

In this Letter, we use the leapfrog integrator to generate noisy trajectories of systems that have $M$ particles moving and interacting in a softened potential among a background of $N - M$ fixed particles (Quinlan & Tremaine 1992; Hayes 1995, 2003). We use normalized units (Heggie & Mathieu 1986) in which each particle has mass $1/N$ and the system diameter, crossing time, and average velocity all have order unity. We then lead the reader through the following observations. First, we observe that glitches in the trajectory of a single particle occur as a Poisson process (Fig. 1). Next, we demonstrate that as $M$ increases, shadow durations scale roughly as $1/M^{0.8}$. (The physical significance of 0.8 is unclear and may be dependent on other parameters such as $N$, the softening $\epsilon$, and the time step $h$.) More importantly, this scaling can be experimentally reproduced by superimposing trajectories of $M$ single moving particle systems and taking the *shortest* shadow of those $M$ systems. In other words, particles appear to encounter glitches independently of one another (Fig. 2). Now, the glitching and subsequent errant behavior of just one particle (the first to undergo a glitch) in a large simulation is unlikely to have a

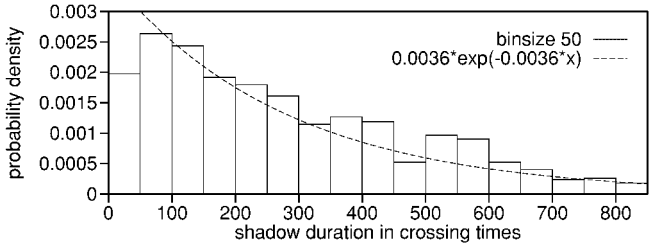


FIG. 1.—Histogram of shadow durations of 1000 systems in which one particle moves in the potential of 99 fixed particles with softening $\epsilon = 0.054$ or $\frac{1}{4}$ of the mean interparticle spacing. Noisy trajectories were generated using a leapfrog integrator with time step $h = 0.011 = \epsilon/5$. After an initial transient, the distribution fits an exponential curve with a mean shadow duration of 280 crossing times, indicating that the moving particle encounters glitches as a Poisson process with a glitch probability of 0.36% per crossing time.

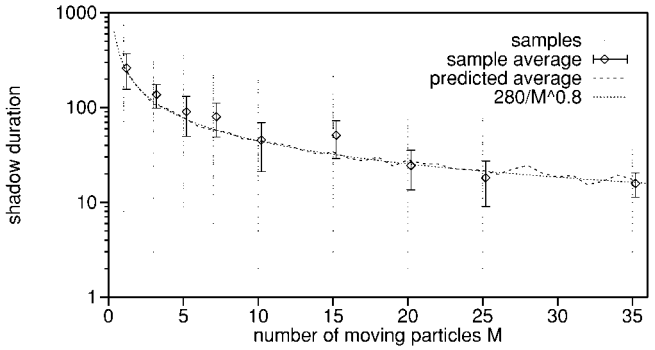


FIG. 2.—How the average shadow duration scales as the number of moving particles $M$ is increased. Each system is identical to that described in Fig. 1, except now $M$ takes on the values 1, 3, 5, 7, 10, 15, 20, 25, and 35. The dots represent sample shadow durations, 40 samples for each $M$, while the "sample average" is plotted with sample error bars of full width 1 $\sigma$. The "predicted average" is artificially constructed for each $M = 1, \ldots, 35$ by superimposing $M$ samples chosen at random from Fig. 1 and taking the minimum shadow duration of those samples. This demonstrates that the average shadow duration of an $M$ moving particle system is well predicted using an $M$ single moving particle system and suggests that particles encounter glitches independently of one another; $280/M^{0.8}$ is plotted for comparison, although the physical significance of 0.8 is unclear.

large effect on the reliability of that simulation; in fact, as long as only a small fraction of particles have glitched, then the reliability of the simulation probably remains high. Then, assuming that particles encounter glitches independently of one another, we can use the distribution of shadow durations of $M = 1$ systems to predict the fraction of shadowable particles as a function of time. We find that this fraction is a decaying exponential with some exponent $\mu$ (Figs. 3 and 4). Finally, we demonstrate an explicit relationship between $\mu$, $N$, $\epsilon$, and $h$ that holds as long as $\epsilon$ is in the range $\sim$[0.2, 1] times the mean interparticle separation and $h \lesssim (\epsilon/3)$ (Fig. 5). This means that given $N$, $\epsilon$, and the expected duration of the simulation, one can *precompute* the time step $h$ necessary to have a desired fraction of shadowable particles remain at the end of the simulation.

## 2. METHODS

Initial positions for all particles and velocities for moving particles were chosen uniformly at random from the unit cube. Although not corresponding to any real astrophysical system, we join the good company of Goodman et al. (1993) in believing that the detailed initial conditions do not affect the chaotic nature of the problem, which is what we focus on in this Letter. Furthermore, we repeated a small set of these experiments with several spherically symmetric distributions, with virtually no changes to our results.

Newtonian gravitational forces were summed directly, with softening implemented by adding $\epsilon^2$ to the denominator of Newton's equations. Shadow computations are described in detail elsewhere (Quinlan & Tremaine 1992; Hayes 1995). Essentially, a highly accurate integrator is used to directly measure the error made by leapfrog across each step, and then Newton's method is applied in an effort to force these errors to zero. The duration of the longest shadow is estimated by running Newton's method on longer and longer segments of the noisy orbit until it fails.

## 3. RESULTS

Figure 1 introduces a histogram of shadow durations for 1000 softened systems with $N = 100, M = 1$. After an initial transient (explained in the discussion of Fig. 3 in Hayes 2003), the distribution fits an exponential curve, suggesting that glitches occur as a Poisson process.

Figure 2 introduces how the average shadow duration scales as the number of moving particles is increased. For various values of $M$, we perform 40 experiments in which $M$ particles move and interact among $100 - M$ fixed particles and plot the mean and standard deviation of the shadow durations. For $M > 1$, the shadow formally ends when *any* particle encounters a glitch, and this is the duration plotted in Figure 2. We make the following observations: (1) a glitch in the local six-dimensional phase-space trajectory of any one particle will cause a glitch in the full $6M$-dimensional phase-space trajectory of the $M$ moving particle system; and (2) in a large collisionless system, the gravitational potential is governed more by the global potential than by interparticle interactions (Binney & Tremaine 1987), and so it is reasonable to expect that particles encounter glitches independently of each other. So, perhaps the mean shadow duration of an $M$ moving particle system can be predicted by the mean shadow duration of a system with $M$ completely uncoupled one moving particle systems (Hayes 2003). We test this hypothesis with the "predicted average" of Figure 2 by taking $M$ samples at random from Figure 1 and taking the *shortest* shadow duration. We see that the predicted curve is well within the error bars of the "real" $M$ moving particle system. Formally, this suggests that the average duration before the occurrence of the first glitch in the system is statistically independent of whether particles interact or not.

Once one particle in the system encounters a glitch, its trajectory after that point is incorrect, and it will presumably start to "infect" the motion of other particles. However, by observation (2) above, we can hope that in a large collisionless system, one errant particle will, for a time, have a negligible effect on the trajectories of other particles. In fact, we can guess that, for a time, any small fraction of errant particles will have little effect on the global behavior of system. The goal would then be to minimize, at reasonable cost, the number of particles that have glitched by the end of a given simulation.

We now take as a working assumption that particles encounter glitches independently of one another and that the first small fraction of particles that encounter glitches have a negligible effect on the others. That is, we reinterpret Figure 1 to represent the distribution of shadow durations for *all* the individual particles in a *single* many–moving particle system. Of course, the figure is likely to be valid only for a duration much

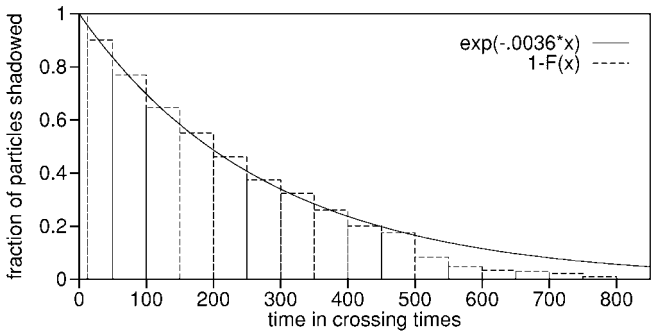


FIG. 3.—Estimated fraction of particles that would be shadowed as a function of time, in a system similar to that described in Fig. 1, except with *all* particles moving. This is computed simply by taking the cumulative distribution function $F(t)$ of Fig. 1, where $F(t)$ represents the fraction of particles that have undergone a glitch, and then plotting $1 - F(t)$.

shorter than 800 crossing times, as the earlier glitched particles "infect" the motion of the remainder, but let us assume that it is a reasonable approximation for some shorter period. This allows us, as a first approximation, to estimate the fraction of glitched particles in a real simulation at a given time by computing the fraction of one moving particle systems that have glitched by that time. Figure 3 plots the opposite—the fraction of nonglitched (i.e., shadowed) particles as a function of time—derived from Figure 1 by taking its cumulative distribution function $F(t)$ and "flipping" it to $1 - F(t)$. As expected from the Poisson process in Figure 1, the fraction of shadowable particles decays exponentially with a rate corresponding to a 0.36% glitch probability per particle per crossing time.

Figure 3 is interesting but of little use, because it does not tell us how the shape of this curve varies with the total number of particles $N$, the softening $\epsilon$, or the time step $h$. However, for reasons that we discuss later, we have found that if the time step $h$ is scaled as

$$h^2 \propto \epsilon^2 N^{1/3}, \tag{1}$$

then each of Figures 1, 2, and 3 are preserved if $\epsilon$ is not too small. That is, if $N$ and $\varepsilon$ are changed in a given simulation but the initial conditions are drawn from the same distribution, then using equation (1) to scale the time step will preserve the same degree of simulation reliability from the standpoint of shadowing. Intuitively, the scaling of $h \propto \epsilon$ is not surprising and is a commonly used time step criterion. The $N^{1/3}$ is more surprising, telling us that we can *increase* the time step as $N$ increases at a fixed softening; intuitively, this is because the gravitational potential becomes smoother with increasing $N$.

To demonstrate this scaling, we have performed many experiments with various values of $h$, $N$, and $\epsilon$. Figure 4 summarizes the results. Each line represents the fraction of shadowable particles as a function of time for some ($N$, $\epsilon$) pair. Each closely clustered set of four lines represents sets of runs with various ($N$, $\epsilon$) and the time step $h$ scaled using equation (1) to give the same "shadowing reliability." Finally, decreasing the time step (via decreasing $k$) increases reliability by decreasing the decay rate of the fraction of shadowed particles. We also performed similar experiments with all combinations of parameters $N = 10^2, 10^3, 10^4$ and $\epsilon = \{2, 1, \frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}\}N^{-1/3}$. The scaling with $N$ works well up to $N = 10^4$ and presumably beyond. All curves are very similar for $\epsilon \geq \frac{1}{4}N^{-1/3}$. However, as seen in the figure, shadows are significantly shorter for $\epsilon = \{\frac{1}{8}, \frac{1}{16}\}N^{-1/3}$, even if the smallest time step is used. This is consistent with the observations of Quinlan & Tremaine (1992) and may be related to unphysical results obtained with a too-small softening (Splinter et al. 1998).

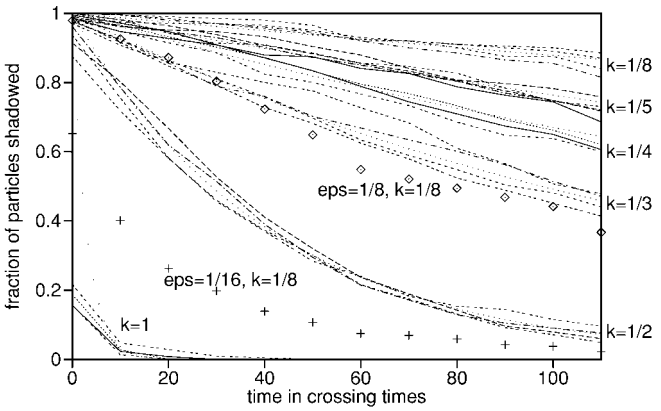


FIG. 4.—Curves suggesting that (1) eq. (1) preserves simulation reliability from a shadowing standpoint, (2) reliability increases as $h$ decreases, and (3) the scaling breaks down if $\epsilon$ is too small. Each line represents a curve similar to that in Fig. 3, derived from a 200 sample set of simulations similar to Fig. 1 but with different $N$ and $\epsilon$. There are six clusters of four lines each. The four lines in each cluster come from four sets of simulations using all four combinations of $N = 100,1000$ and $\epsilon = \{\frac{1}{2}, \frac{1}{4}\}N^{-1/3}$. Each system was integrated using leapfrog with time step $h$ scaled according to eq. (1). Each cluster represents a particular choice of constant $k$ in the scaling, namely, $h^2 = k^2\epsilon^2(N/100)^{1/3}$, for the displayed values of $k$. Decreasing $k$ increases accuracy, giving longer shadows. Embedded in each cluster is a dotted line representing $e^{-\mu_k t}$, discussed later. Finally, the diamond and plus sign curves have $\epsilon = \{\frac{1}{8}, \frac{1}{16}\}N^{-1/3}$, respectively, and suggest that shadows are much shorter if the softening is too small, even though they use the smallest time step factor $k = \frac{1}{8}$.

Finally, we come to the crux. Armed with the scaling equation (1) and the knowledge that the fraction of shadowable particles decays exponentially as $e^{-\mu_k t}$, we would like to find a relationship between the time step proportionality constant $k$ of Figure 4 and $\mu_k$. An eyeball fit of exponential curves to each of the clusters in Figure 4 gives values of the decay constant $\mu_k$ for each cluster. These values, along with the curve $\mu_k$ versus $k$, are plotted in Figure 5. As can be seen, there is evidence that for $k \lesssim \frac{1}{3}$, the curve settles to a power law of approximately

$$\mu_k = (0.047 \pm 0.015)k^{1.75 \pm 0.25}. \tag{2}$$

However, the shape of the curve is also consistent with a slowly decreasing slope as $k$ decreases.

## 4. DISCUSSION

The relations described above offer an a priori algorithm for choosing a time step for a softened $n$-body simulation, viz., given $N$, $\epsilon$, the expected simulation duration $T$ in crossing times, and a desired fraction $F$ of shadowable particles remaining at time $T$, solve for $\mu$ in $F = e^{-\mu T}$, and then solve for $k$ using equation (2). Of course, these relationships will need to be scaled to appropriate units for the simulation. The job should be easiest for a simulation of one galaxy; for simulations of clusters of galaxies or a cosmological simulation, we would scale to the smallest subsystems that we expect to accurately integrate. We are unsure of the effects of dynamically changing the softening based on the local mean particle density but suspect that some reasonable interpretation may be possible whereby a softening and time step (modulo the discussion of the next paragraph) are chosen based upon local mean particle density.

The fact that constant time step leapfrog is symplectic is

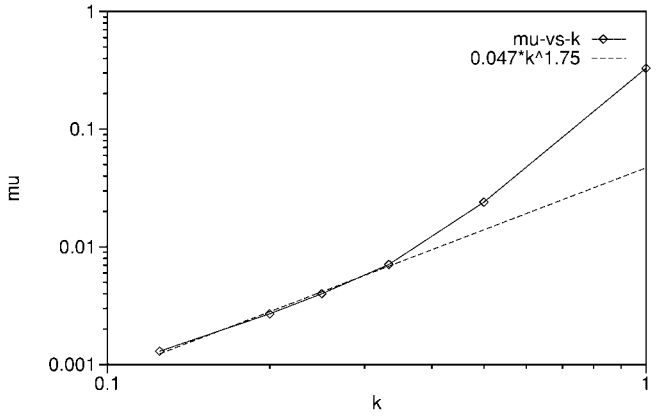


FIG. 5.—For $k = 1, \frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \frac{1}{5}$, and $\frac{1}{8}$ and $\epsilon \geq \frac{1}{4}N^{-1/3}$ of Fig. 4, the values of $\mu_k$ were fitted by eye and are, respectively, $\mu = 0.33, 0.024, 0.0071, 0.0040, 0.0027$, and 0.0013. These are plotted and, for $k \lesssim \frac{1}{3}$, fit the curve $\mu = (0.047 \pm 0.015)k^{1.75 \pm 0.25}$.

probably significant to these results. We experimented briefly with a dynamically changing time step but found that shadows were virtually destroyed if the time step changes "too often." However, we found that if the time step was decreased as a particle entered a high-density region but never increased the time step again, these results were preserved. This may be a reasonable choice for simulations of clusters if most particles that enter a high-density region remain there for the remainder of the simulation. Alternatively, perhaps a particle's time step could be reincreased only after the dynamic time step criterion says the particle's time step should be *significantly* increased, say, by an order of magnitude. This will ensure that the particle has left the high-density region far behind and preserve the internal reliability of high-density regions. Quinn et al. (1997) discuss time steps in great detail.

More detailed arguments deriving equation (1) show that the forward global error of a softened $n$-body simulation scales as $h^2\epsilon^{-2}N^{-1/3}$, where the forward global error is defined as the total error accumulated over a fixed duration of simulation time (i.e., if the time step is halved, then the number of time steps must be doubled to fix the duration). The $h^2$ scaling is due to leapfrog being a globally second-order integrator,[1] the $N^{-1/3}$ scaling has been seen before (Goodman et al. 1993), and the $\epsilon^{-2}$ is new and can be easily verified by plotting the forward global error as a function of $\epsilon$ for fixed $N$ and $h$, for the values of $\epsilon$ used in this Letter. Thus, equation (1) simply holds the forward global error constant; to our knowledge, this Letter is the first to demonstrate that holding the forward global error constant results in a constant distribution of shadow lengths, although this may be related to a shadowing concept known as brittleness (Dawson et al. 1994).

Since the scaling of $h$ with $N^{1/3}$ is so weak, and we usually increase $N$ in order to increase reliability, it certainly will not hurt to ignore the scaling with $N$ and simply scale $h \propto \epsilon$. The remaining questions are what value of $\epsilon$ to use and what fraction of $\varepsilon$ should a particle be allowed to travel in 1 time step. The first question is answered by the fact that shadows do not appear to exist for very long if $\epsilon \lesssim \frac{1}{8}N^{-1/3}$; concerning the second question, we believe that Figures 4 and 5 suggest that the particle should be allowed to travel *at most* $\frac{1}{3}\epsilon$ in 1 time step, and possibly much less, if a shadow duration of 100 crossing times encompassing most particles is desired.

We thank A. Melott and D. Weinberg for many helpful presentation suggestions.

[1] Note that leapfrog is $O(h^3)$ *per step*, but if the duration is fixed and the time step is halved, then twice as many steps are required, making the forward global error $O(h^2)$.

## REFERENCES

Bertschinger, E. 1998, ARA&A, 36, 599
Binney, J., & Tremaine, S. 1987, Galactic Dynamics (Princeton: Princeton Univ. Press)
Clarke, D. A., & West, M. J., eds. 1997, ASP Conf. Ser. 123, The 12th "Kingston Meeting": Computational Astrophysics (San Francisco: ASP)
Dawson, S., Grebogi, C., Sauer, T., & Yorke, J. A. 1994, Phys. Rev. Lett., 73, 1927
Goodman, J., Heggie, D. C., & Hut, P. 1993, ApJ, 415, 715
Hayes, W. 1995, Master's thesis, Univ. Toronto
Hayes, W. 2003, Phys. Rev. Lett., 90, 054104
Heggie, D. C., & Mathieu, R. D. 1986, in The Use of Supercomputers in Stellar Dynamics, ed. P. Hut & S. L. W. McMillan (Berlin: Springer), 33
Hernquist, L., Hut, P., & Makino, J. 1993, ApJ, 402, L85
Quinlan, G. D., & Tremaine, S. 1992, MNRAS, 259, 505
Quinn, T., Katz, N., Stadel, J., & Lake, G. 1997, ApJ, submitted (astro-ph/9710043)
Splinter, R. J., Melott, A. L., Shandarin, S. F., & Suto, Y. 1998, ApJ, 497, 38