

COMPUTATIONAL PHYSICS

The Computational Physics Section publishes articles that help students and their instructors learn about the physics and the computational tools used in contemporary research. Most articles will be solicited, but interested authors should email a proposal to the editors of the Section, Jan Tobochnik (jant@kzoo.edu) or Harvey Gould (hgould@clarku.edu). Summarize the physics and the algorithm you wish to include in your submission and how the material would be accessible to advanced undergraduates or beginning graduate students.

Adaptive Runge–Kutta integration for stiff systems: Comparing Nosé and Nosé–Hoover dynamics for the harmonic oscillator

William Graham Hoover^{a)}

Ruby Valley Research Institute, Highway Contract 60, Box 601, Ruby Valley, Nevada 89833

Julien Clinton Sprott

Department of Physics, University of Wisconsin, Madison, Wisconsin 53706

Carol Griswold Hoover

Ruby Valley Research Institute, Highway Contract 60, Box 601, Ruby Valley, Nevada 89833

(Received 26 February 2016; accepted 13 July 2016)

We describe the application of adaptive (variable time step) integrators to stiff differential equations encountered in many applications. Linear harmonic oscillators subject to nonlinear thermal constraints can exhibit either stiff or smooth dynamics. Two closely related examples, Nosé's dynamics and Nosé–Hoover dynamics, are both based on Hamiltonian mechanics and generate microstates consistent with Gibbs' canonical ensemble. Nosé's dynamics is stiff and can present severe numerical difficulties. Nosé–Hoover dynamics, although it follows exactly the same trajectory, is smooth and relatively trouble-free. We emphasize the power of adaptive integrators to resolve stiff problems such as the Nosé dynamics for the harmonic oscillator. The solutions also illustrate the power of computer graphics to enrich numerical solutions. © 2016 American Association of Physics Teachers. [<http://dx.doi.org/10.1119/1.4959795>]

I. INTRODUCTION

Simulations of classical many-particle systems are frequently done by solving Newton's equations of motion numerically using leapfrog or Runge–Kutta algorithms.¹ This method is termed molecular dynamics.^{2,3} These simulations are necessarily done in the microcanonical ensemble for which the number of particles, volume, and energy are each held fixed. However, most experiments are done at constant temperature. Hence, it is desirable to modify Newton's equations so that the temperature rather than the energy is fixed.

In 1984 Nosé modified constant energy molecular dynamics so that the simulations were done at a fixed temperature.^{4–9} His goal was to simulate the canonical ensemble directly from the dynamics. Nosé based his work on Hamiltonian mechanics. Because the resulting equations of motion are typically stiff and difficult to solve, Nosé introduced the idea of time-scaling to cope with these difficulties.^{4,5} Hoover showed that an improved set of equations of motion could be based on Liouville's theorem without the need for time-scaling or a Hamiltonian basis.^{2,3,7} Dettmann later provided a Hamiltonian basis that linked both equations of motion.^{8,9}

Both the original Nosé dynamics and the improved Nosé–Hoover dynamics share a common flaw; although their solutions are consistent with the canonical distribution

and maintain it, they are not necessarily ergodic. Their phase-space flows cover only a fraction of the available states, even for the harmonic oscillator.^{6–13} It was not until 2015 that an ergodic analog of Nosé dynamics was developed,¹² thus achieving the original goal of Nosé. The key to this success was including a cubic dependence on momentum in formulating thermostating forces.

We will discuss applications of the Nosé and Nosé–Hoover algorithms to the harmonic oscillator. Both algorithms yield exactly the same trajectories in phase space but at different rates and are related to each other by time scaling. Only Nosé's original dynamics exhibits the stiffness that is the focus of our paper.

Nosé dynamics has both regular (stable to small perturbations) and chaotic solutions, coexisting in a constant-energy three-dimensional volume within a four-dimensional phase space. We will see that within this four-dimensional space there are infinitely many regular solutions (concentric tori and stable periodic orbits). In addition, there is a single chaotic sea that stretches to infinity and occupies about six percent of the stationary measure defined by Liouville's theorem. That stationary measure is a three-dimensional product of Gaussians in the coordinate, the momentum, and the “friction coefficient” ζ . Random sampling reveals that only six percent of this Gaussian product lies in the chaotic sea. The rest is occupied by regular toroidal solutions. Although the simpler regular stable tori and the periodic orbits they enclose are

easy to solve, the motion in Nosé's version of the chaotic sea is sufficiently stiff to require the special numerical methods that we will describe.

Typical fixed-time-step numerical integrators such as leap-frog and fourth-order Runge–Kutta^{1–3} are ineffective in Nosé's chaotic sea. We will discuss a family of useful alternatives, simple variable-time step adaptive algorithms,¹ and explore Nosé's dynamics with their aid. Adaptive techniques vary the time step to compensate for time-dependent changes in the stiffness of the underlying ordinary differential equations.

Our goal is to introduce readers to applications of Nosé and Nosé–Hoover dynamics to the harmonic oscillator and to their solution using adaptive integration. We connect stiffness with time scaling, the adaptive timestep, and Lyapunov instability. While exploring these numerical features of the dynamics, we will find intricate topologies with their roots in simple quadratic differential equations.

The application of Nosé dynamics to the harmonic oscillator is an excellent testbed for numerical integrators, computer graphics, and numerical methods, and provides challenges in visualizing the fascinating topology of knots and island chains in readily accessible three- and four-dimensional phase space. The model's simple structure makes it an ideal introduction to dynamical systems.

In Sec. II, we discuss the statistical mechanics background that links Nosé and Nosé–Hoover dynamics. We introduce a family of flexible integrators capable of accurate solutions of stiff problems in Sec. III. In Sec. IV, we apply a common initial condition for our investigations of time-scaling and stiffness. In Sec. V, we summarize what we have learned and suggest areas for further study using our new tools.

II. NOSÉ AND NOSÉ–HOOVER DYNAMICS FOR THE HARMONIC OSCILLATOR

A. Nosé dynamics

In 1984 Shuichi Nosé proposed a temperature-dependent Hamiltonian \mathcal{H}_N with a corresponding dynamics that generates a temperature-dependent trajectory. The evolution is governed by the time-scaling control variable $0 < s < 1$ that controls the kinetic temperature T .^{4,5} This s -dependent dynamics generates microstates that are consistent with the canonical ensemble rather than the more usual constant energy microcanonical ensemble. In conventional mechanics, the total energy of the one-dimensional harmonic oscillator, $\mathcal{H} = (1/2)(q^2 + p^2)$, is constant, where q and p are the displacement and momentum, respectively. This dynamics generates oscillatory solutions with an amplitude proportional to the square root of the kinetic temperature p^2 . In one dimension we have

$$\langle q^2 \rangle = \langle p^2 \rangle = T = \langle \mathcal{H} \rangle, \quad (1)$$

where we have chosen units so that the mass $m = 1$ and the spring constant $\kappa = 1$. For simplicity, and without loss of generality in the isothermal equilibrium case, we also set Boltzmann's constant equal to unity, $k = 1$.

Nosé's time-scaled Hamiltonian, with s the time-scaling factor and ζ its conjugate momentum, for the isothermal one-dimensional harmonic oscillator is

$$2\mathcal{H}_N = q^2 + (p/s)^2 + \ln(s^2) + \zeta^2. \quad (2)$$

The Hamiltonian in Eq. (2) governs the evolution of four time-dependent variables (q, p, s, ζ) . (We will discuss in Sec. II C why a Hamiltonian such as the one in Eq. (2) is isothermal.) The usual oscillator coordinate and momentum variables are still q and p , but the usual link between velocity and momentum is changed. The relation $\dot{q} = p$ is replaced by $\dot{q} = (\partial\mathcal{H}_N/\partial p) = (p/s^2)$. Here and in what follows we choose the initial conditions so that the Hamiltonian vanishes, $\mathcal{H}_N = 0$, for reasons that we will explain in the following.

Because the canonical distribution includes microstates with all energies, $0 < \mathcal{H} < \infty$, Nosé included the time-scale factor s and its conjugate momentum ζ in the Hamiltonian, making it possible for the scaled momentum (p/s) to cover the infinite range required by the distribution function of the canonical ensemble

$$f(q, p) \equiv e^{-q^2/2} e^{-p^2/2} / 2\pi. \quad (3)$$

Hamilton's equations of motion for any (q, p) coordinate-momentum pair are $\dot{q} = +\partial\mathcal{H}/\partial p$ and $\dot{p} = -\partial\mathcal{H}/\partial q$. In addition to (q, p) , \mathcal{H}_N also has the time-scaling variable s and its conjugate momentum $p_s = \zeta$, making a second coordinate-momentum pair (s, ζ) . The equations of motion in the four-dimensional phase space (q, p, s, ζ) follow from \mathcal{H}_N

$$\dot{q} = p/s^2 \quad (\text{Nosé dynamics}), \quad (4a)$$

$$\dot{p} = -q, \quad (4b)$$

$$\dot{s} = \zeta, \quad (4c)$$

$$\dot{\zeta} = p^2/s^3 - 1/s. \quad (4d)$$

Because the value of the Hamiltonian is constant (the reader can check this by computing $\dot{\mathcal{H}}_N$), the motion takes place in the (unbounded) three-dimensional volume where $\mathcal{H}_N = 0$. The unbounded logarithmic potential makes it possible to reach arbitrarily large values of the coordinate and scaled momentum, as is necessary to cover the entire canonical distribution.

In Nosé's mechanics, the temperature is a measure of the kinetic energy, $T \propto (p/s)^2$, where p/s is the scaled momentum. Nosé's approach uses s to scale the momentum over a broad range. As a result, Nosé's equations of motion in Eq. (4) present numerical challenges. The scale factor s varies from its maximum of unity to less than 10^{-9} over a 10^9 time step adaptive simulation, where the mean time step is about 0.002 (see the shorter 10^6 time step oscillator histories shown in Fig. 1). Fortunately, there is a way to temper the singular behavior of $\dot{\zeta}$.⁹

B. Dettmann and Morriss' contribution

Dettmann and Morriss' contribution^{8,9} is an alternative Hamiltonian description of the same trajectories but with all the rates given in Eq. (4) multiplied by s . Their Hamiltonian, \mathcal{H}_D , and the new, generally slower and smoother rates it generates, are

$$\mathcal{H}_D = s\mathcal{H}_N = \frac{1}{2} \left[sq^2 + (p^2/s) + s\ln(s^2) + s\zeta^2 \right] = 0. \quad (5)$$




Fig. 1. The variation of the time-scaling factor s found from the numerical solution of Eq. (4) (left) and Eq. (6) (right). In principle, exactly the same values of (q, p, s, ζ) are generated by both sets of equations. Equations (4) and (6) were solved for 10^6 time steps using the adaptive fourth-order Runge–Kutta algorithm. The time step is doubled when the root-mean-square discrepancy between a single time step (Δt and two half steps $\Delta t/2$) is less than 10^{-12} . The time step is halved if the discrepancy is greater than 10^{-10} . The minimum value of s during the time shown is about $e^{-10} \simeq 0.00005$. All four rates in Eq. (6) are smaller than in Eq. (4) by a factor of $s(t)$. The initial conditions for Eqs. (4) and (6) are given in Eq. (14) and are chosen so that the Hamiltonian vanishes.

$$\dot{q} = p/s \quad (\text{Nose–Hoover1}), \quad (6a)$$

$$\dot{p} = -qs, \quad (6b)$$

$$\dot{s} = \zeta s, \quad (6c)$$

$$\dot{\zeta} = p^2/s^2 - 1. \quad (6d)$$

Multiplying \mathcal{H}_N (and thus the four rates) by s tames the singular behavior of Nosé mechanics when s is small and is equivalent to a close relative of Nosé–Hoover mechanics, which we refer to as Nosé–Hoover1 and which is identical to Dettmann and Morriss’ dynamics.

This improvement can be simplified further to obtain the usual Nosé–Hoover equations of motion by replacing the scaled momentum (p/s) by p

$$\dot{q} = p \quad (\text{Nose–Hoover2}), \quad (7a)$$

$$\dot{p} = -q - \zeta p. \quad (7b)$$

$$\dot{\zeta} = p^2 - 1. \quad (7c)$$

$$\dot{s} = s\zeta. \quad (7d)$$

A numerical advantage of Eq. (7) is the irrelevance of the scaling variable s , which means Eq. (7d) can be ignored. The evolution of $\{q, p, \zeta\}$ in a three-dimensional space extending to infinity can be determined without any consideration of s provided that \mathcal{H} vanishes.

The choice of the unusual condition $\mathcal{H}_D = 0$ is necessary for the derivation of the equations of motion in Eq. (6) directly from Hamilton’s equations. With this trick, Nosé’s multiplication of all the rates by s is not required. Time-scaling is absent in the Dettmann–Morriss approach. The vanishing of the Hamiltonian makes it possible to simplify the expression for $\dot{\zeta}$

$$\dot{\zeta} = -\frac{\partial \mathcal{H}_D}{\partial s} = -\frac{1}{2} \left[q^2 - (p/s)^2 + \ln(s^2) + \zeta^2 + 2 \right]. \quad (8)$$

If $\mathcal{H}_D = 0$, Eq. (8) reduces to the much simpler, but numerically equivalent Eq. (6d).

C. Equations of motion using Liouville’s theorem

We now discuss the physics behind Hoover’s much simpler three equation subset of Eq. (7).⁷ We begin with an

augmented set of equations of motion for a harmonic oscillator with a friction coefficient ζ that acts over a characteristic relaxation time τ

$$\dot{q} = p, \quad (9a)$$

$$\dot{p} = -q - \zeta p, \quad (9b)$$

$$\dot{\zeta} = [(p^2/T) - 1]/\tau^2. \quad (9c)$$

Here, ζ is the friction coefficient or control variable. If the kinetic temperature p^2 exceeds the target temperature T , the friction increases, slowing p . If p^2 is less than T , the friction is reduced and can become negative, accelerating the oscillator. Provided only that a stationary state results, the long-time-averaged value $\langle \dot{\zeta} \rangle$ is necessarily zero, so that the kinetic temperature p^2 eventually reaches its target

$$\langle \dot{\zeta} \rangle = \langle [(p^2/T) - 1]/\tau^2 \rangle = 0, \quad (10)$$

which implies that $\langle p^2 \rangle = T$. By design the control variable responds to $(p^2/T) - 1$ in a characteristic time τ .

A remarkable feature of the $(\dot{q}, \dot{p}, \dot{\zeta})$ equations of motion is that they leave Gibbs’ canonical distribution function (or probability density) unchanged. Suppose that

$$f(q, p, \zeta, T) = (\tau/T) (2\pi)^{-3/2} e^{-q^2/2T} e^{-p^2/2T} e^{-(\zeta\tau)^2/2}, \quad (11)$$

and consider the rate of change of the probability density in $\mathbf{r} = (q, p, \zeta)$ space as a result of the continuity equation for the flow $\mathbf{v} = (\dot{q}, \dot{p}, \dot{\zeta})$

$$\frac{\partial f}{\partial t} = -\nabla_{\mathbf{r}} \cdot (f\mathbf{v}) \quad (12)$$

$$\begin{aligned} &= -f \left(\frac{\partial \dot{q}}{\partial q} + \frac{\partial \dot{p}}{\partial p} + \frac{\partial \dot{\zeta}}{\partial \zeta} \right) - \dot{q} \frac{\partial f}{\partial q} - \dot{p} \frac{\partial f}{\partial p} - \dot{\zeta} \frac{\partial f}{\partial \zeta} \\ &= -f \left\{ 0 - \zeta + 0 - \frac{q}{T} p - \frac{p}{T} (-q - \zeta p) \right. \\ &\quad \left. - \zeta \tau^2 \left[\left(\frac{p^2}{T} - 1 \right) / \tau^2 \right] \right\} = 0. \end{aligned} \quad (13)$$

The vanishing rate of change throughout (q, p, ζ) space implies that the Gibbs’ distribution is left unchanged by the flow.

We modified Nosé's original work^{4,5} and replaced his factor of $(2/s)$ by $(1/s)$ and his p_s by ζ to match Eq. (6). The wide-ranging values of the time-scaling variable s (see Fig. 1) make the Nosé equations so stiff that an accurate, fixed-time-step fourth-order-Runge–Kutta solution of the equations for a typical duration of 10^6 would require 10^{13} time steps with $\Delta t = 10^{-7}$.

In brief, we have learned that by scaling the time and redefining the momentum, $(p/s) \rightarrow p$, we can obtain a more manageable set of equations consistent with the canonical distribution.⁶ The same equations of motion can be derived directly from the phase-space continuity equation if we insist that the friction coefficient ζ imposes the kinetic temperature T on the dynamics. The situation is ideal because we have several ways to check our work. Apart from the time, the stiff Nosé dynamics in Eq. (4) and the smooth Nosé–Hoover equations in Eqs. (6) and (7) have identical solutions. Because our goal is learning to solve the stiff set in Eq. (4), we turn next to developing suitable adaptive integrators.

III. STIFF OSCILLATOR SOLUTIONS USING ADAPTIVE RUNGE–KUTTA

The original work on the stiff oscillator problem^{6,7} was frustrated by huge and rapid variations of p and s . Stiff equations were unusual in the molecular dynamics simulations of the 1980s. Researchers with a background in simulation were not familiar with adaptive integrators. In contrast, researchers in control theory and heat transfer often used the Runge–Kutta–Fehlberg integrator¹⁴ to solve their stiff equations. The Fehlberg modification of the classic Runge–Kutta methods of the early 20th century compares the RK4 and RK5 results of a trajectory integration for a single time step. The two trajectories over one time step provide a criterion for increasing or decreasing Δt on the next time step. If the discrepancy is too large, Δt is reduced. If it is too small, Δt is increased. In this way the discrepancy between the two estimates can be restricted to a band. If the discrepancy is too large, an alternative to proceeding with the better RK5 estimate is to repeat the current step with smaller values of Δt until the discrepancy falls within the acceptable limits. Suitable limits¹⁵ for double precision and quadruple precision are 10^{-12} to 10^{-10} and 10^{-24} to 10^{-20} .

Rather than comparing the RK4 and RK5 trajectories, we can just as well compare an iteration with a single time step to two successive iterations with half the time step. We have adopted this choice here. We compare two RK4 integrations [Δt versus $(\Delta t/2) + (\Delta t/2)$] of the vector (q, p, ζ) or $\mathbf{x} = (q, p, s, \zeta)$ for a single time step from t to $t + \Delta t$, with the result of two half-time steps of the initially identical vector. The comparison of the single-time step \mathbf{x}_1 solution to the double-time step \mathbf{x}_2 solution provides the criterion for increasing or decreasing the time step. In either case, the less-accurate vector \mathbf{x}_1 is set equal to the more accurate vector \mathbf{x}_2 before the next time step is undertaken. The error is the square root of the quantity $(\mathbf{x}_1 - \mathbf{x}_2)^2$. If the error is greater than 10^{-10} , $\Delta t \rightarrow \Delta t/2$; if the error is less than 10^{-12} , then $\Delta t \rightarrow 2\Delta t$.

As an example, we choose for our standard initial condition

$$(q, p, s, \zeta) = (2.4, 0, e^{-2.88}, 0), \quad (14)$$

because this choice corresponds to $\mathcal{H}_D = 0$ for which the Nosé and Nosé–Hoover equations of motion provide identical trajectories (apart from numerical errors):

$$2\mathcal{H}_D = q^2 + (p/s)^2 + \ln(s^2) + \zeta^2 = 2.4^2 + 0 - 5.76 + 0 = 0.$$

Figure 2 shows a short history of the root-mean-square discrepancy between the single and double-step versions of Eq. (4). The first step is taken with $\Delta t = 0.001$ with the value of the discrepancy checked at each full time step and Δt changed accordingly. Readers are encouraged to reproduce Fig. 2 and to consider modifications of the algorithm.

There is a downside to existing adaptive integrators. Because the past is different than the future, adaptive integrators are not reversible, which means that the accuracy of such algorithms cannot be checked directly. We can easily check a conserved quantity, such as the total energy, equally well in both time directions, but trajectory accuracy requires a more elaborate investigation.¹⁶ One criterion is the crossing of the $p = 0$ section. We should obtain the same number of crossings forward and backward with a good adaptive integrator.

Because atomistic mechanics problems are typically time-reversible, as are all the problems we will discuss, we have developed a time-reversible adaptive integrator by averaging the forward and backward errors at each time step to define an error independent of the direction of time, and then using this error to decide on the magnitude of the local time step. Although this averaging doubles the computational time, it is a reasonable price to pay for a robust integrator useful for problems with time-reversible dynamics. Because our approach is only approximate, developing a time-reversible adaptive integrator remains a worthy research goal.

Both the Nosé and Nosé–Hoover equations of motion have stable and unstable chaotic solutions. In the latter case, the effects of small changes in the initial conditions increase exponentially in time. Choosing an initial condition (see Fig. 3) that yields a chaotic trajectory for Nosé dynamics and the initial time step $\Delta t = 0.001$ generates an adaptive trajectory for 10^9 time steps without any numerical difficulties. Figure 3




Fig. 2. Example of the variation of the root-mean-square integration error within the band 10^{-12} to 10^{-10} for $0.001 < \text{time} < 10$. The standard initial condition is given in Eq. (14). Although both double-precision and quadruple-precision results are shown here, so as to confirm the accuracy of the double-precision work, the difference is barely visible on the scale of this plot.




Fig. 3. The variation of $\log_2(\Delta t)$ with time using double-precision adaptive integration with the root-mean-square error confined to lie between 10^{-12} and 10^{-10} for Nosé (bottom) and Nosé–Hoover (top) dynamics. Crossings of the $p=0$ plane are indicated by filled circles. The Nosé data cover about 88.7×10^3 variable time steps, and the Nosé–Hoover data correspond to the range from 80.2×10^3 to 84.4×10^3 time steps. In both cases the initial condition is given in Eq. (14).

shows the variation of the adaptive time step for a chaotic trajectory. As a check of such trajectories other integrators can be used and compared. Although chaos prevents us from ever finding a nearly exact solution to an initial-value dynamics problem, it does not prevent us from finding a reasonable solution.

IV. CHARACTERIZATION OF STIFFNESS AND CHAOS

We now turn to the numerical characterization of stiffness and chaos. These results are new and are indicative of many novel and promising research directions. To begin we quantify the Lyapunov instability of a chaotic trajectory by measuring the rate at which two neighboring trajectories diverge. If the distance between the neighboring trajectories is denoted by δ , we can characterize its single-time step tendency toward divergence by the local (time-dependent) value of the largest (indicated by the subscript 1) Lyapunov

exponent, $\lambda_1(t) \equiv (\dot{\delta}/\delta)$. The procedure for measuring λ_1 is described in Sec. IV B. Figure 4 shows that the dependence of the time-averaged value $\lambda_1 \equiv \langle \lambda_1(t) \rangle$ on the temperature disappears for long times. The fluctuations in $\lambda_1(t)$ have a strong dependence on temperature. Let us explain the details of such simulations.

A. Nosé dynamics

Nosé’s dynamics for the harmonic oscillator has two kinds of solutions, smooth and stable periodic or toroidal solutions, and chaotic solutions in which nearby trajectories separate from one another exponentially fast with $\delta \simeq e^{\lambda t}$. Figure 5 shows a cross section of the chaotic sea. All of the points in the chaotic sea are accessible by a single unstable trajectory. All other points lie on stable tori or periodic orbits. The chaotic points are plotted whenever a trajectory passes through the plane $p=0$ or equivalently $(p/s)=0$; they therefore reflect the flux through the section, which is the product of the probability density and the speed normal to the plane, $|\dot{p}|$

$$|q|f(q, 0, s, \zeta) \propto |q|f(q, 0, \zeta) \propto |q|e^{-q^2/2}e^{-\zeta^2/2}. \quad (15)$$

The stationary distribution satisfying $(\partial f/\partial t) = 0$ is Gaussian in all state variables (q, p, ζ) , as is shown in Eq. (13). The cross section in Fig. 5 and the flux through it are exactly the same in Nosé and Nosé–Hoover dynamics because the trajectory is the same, with the same velocity at the $p=0$ plane, independent of s ; that is, $\dot{p}_N = \dot{p}_{NH} = -q$. Solving either set of equations, Nosé or Nosé–Hoover, generates exactly the same cross section. Although their trajectory speeds are generally different, they agree exactly in the $p=0$ cross section.

Let us delve into the details of the chaotic sea from the perspective of Nosé dynamics, using the initial condition in Eq. (14), which implies that $\mathcal{H}_N = (1/2)[q^2 + (p/s)^2 + \ln(s^2) + \zeta^2] = 0$. One way (there is no consensus) to quantify the oscillator’s stiffness is to record the range of the time-scaling factor s that is responsible for the stiffness. Another way is to record the range over which Δt must be varied in order to solve the equations with a given integrator. We have used the classic RK4 integrator, increasing or decreasing Δt as needed for accuracy. The variables $(s$ and $|p|)$ each depend linearly on the variable time step Δt as shown in the log-log plot of Fig. 6. The stiffness gives a




Fig. 4. The time variation of $\langle \lambda_1 \rangle$ and $\langle \lambda_1^2 \rangle$ for Nosé dynamics using $\delta = 10^{-6}$ and double-precision integration. The root-mean-square error is confined to the band 10^{-12} to 10^{-10} . The time-averaged Lyapunov exponent is temperature-independent, but the fluctuations of the exponent vary with temperature. The curves correspond to equally spaced temperatures from $T=1$ to 10, where the highest temperature is the shortest curve. Because the equations are stiffer at higher temperatures, the length of the curves shortens as the temperature is increased. The time reached is a uniformly decreasing function of the temperature due to the stiffer nature of the higher temperature simulations.




Fig. 5. The $(q, p, s, \zeta) = (q, 0, s, \zeta)$ cross section for Nosé dynamics in the constant energy three-dimensional chaotic sea. The energy is zero. A point in this zero momentum cross section is determined whenever the product of two successive momentum values is negative. Stable families of tori cut by this $p=0$ section occupy the infinitely many holes in the distribution. See Ref. 6 for many examples. The largest time-averaged Lyapunov exponent in the sea is $\lambda_1 = 0.046$, and thus the chaos is weak. The same cross section is found for the isothermal Nosé–Hoover equations, for which $\lambda_1 = 0.0145$. Despite the similar Lyapunov exponents and identical cross sections, there are huge fluctuations in the rate of exploration of phase-space for Nosé dynamics, which requires an adaptive integrator to reproduce the cross section shown here. The abscissa reflects the scaling of the coordinate q and momentum p [or (p/s) for Nosé–Hoover dynamics] with the square root of the temperature T . This million-point cross section applies to both the Nosé and the Nosé–Hoover trajectories as they share a common chaotic sea when $\mathcal{H} = 0$.

rough correspondence between $\ln(\Delta t)$ and $\ln(s)$. Because the scaled momentum (p/s) is proportional to \sqrt{T} , and the temperature T is unity in the figure, $|p|$ is roughly proportional to s and to Δt . To solve this same problem for Nosé dynamics with a fixed time step would require a tiny time step, $\Delta t \approx 10^{-7}$.

In contrast, the smooth equations of motion in Eq. (7) can be solved with the classic RK4 integrator for 2×10^8 double-precision time steps with $\Delta t = 0.01$, reaching a duration of 2×10^7 without any difficulty. The three-dimensional initial condition, $(q, p, \zeta) = (2.4, 0, 0)$, corresponds to the four-dimensional initial condition used in Fig. 2 and given in Eq. (14). The Nosé–Hoover solution generates more than 5×10^6




Fig. 6. The dependence of the speed $|p|$ (points) and scale factor s (horizontal lines) on the variable time step Δt for Nosé dynamics. Both variations are close to linear. The adaptive fourth-order Runge–Kutta integrator maintains a single-step root-mean-square error between 10^{-12} and 10^{-10} . The plot includes 10^6 double-precision time steps in the chaotic sea with $\mathcal{H} = 0$.

penetrations of the $p=0$ cross-sectional plane. With this fixed time step the Nosé oscillator’s progress is roughly one hundred thousand times slower, while generating exactly the same $(q, p/s, \zeta)$ states.

B. The local Lyapunov exponent $\lambda(t)$ as a stiffness criterion

An alternative measure of stiffness can be based on the local (instantaneous) Lyapunov exponent. The local exponent describes the rate $\lambda_1(t)$ at which two nearby trajectories tend to separate. To measure this rate, we choose a satellite trajectory x_s constrained to be a distance $\delta = 10^{-6}$ from the reference trajectory x_r . The distance is rescaled after each time step Δt by multiplying the separation by a factor g that is close to unity

$$x_s = x_r + g(x_s - x_r), \quad (16)$$

where

$$g = \frac{\delta}{\sqrt{(x_s - x_r)^2}} \simeq e^{-\lambda_1 \Delta t}. \quad (17)$$

The largest local Lyapunov exponent is thus

$$\lambda_1(t) = -\ln(g)/\Delta t.$$

Figure 7 shows two sections of a typical time history of λ from an adaptive fourth-order Runge–Kutta solution of Nosé dynamics in Eq. (4). The central Nosé peak near $t = 5.0862$ was resolved by an adaptive integrator that chose a time step of 2×10^{-9} at this peak. The corresponding time for Nosé–Hoover dynamics in Eq. (7) is $t = 16.42$ and corresponds to a broad minimum in s . The amplitudes and the required number of time steps differ by about six orders of magnitude for the two equivalent representations of the oscillator trajectory at constant temperature.

Nosé avoided the stiffness of his equations of motion by arbitrarily multiplying the right-hand sides of each of his four equations of motion by s . This trick doesn’t change the four-dimensional trajectory, but the rate at which the trajectory is followed is changed by the factor s . Nosé termed this change as “scaling the time.”^{4,5} We now can solve his equations directly with adaptive integrators, and have attained a good picture of the Lyapunov instability and stiffness of the original Nosé equations. Figure 8 shows the structure of a typical Nosé–Hoover peak for this same oscillator problem.




Fig. 7. The largest local Lyapunov exponent of the Nosé oscillator in the chaotic sea as a function of time with $\delta = 10^{-6}$ and the quadruple-precision error confined to be between 10^{-28} and 10^{-24} . (left) Data with time steps from 1.63 to 3.07×10^6 . (right) Data between 2.01 and 2.16×10^6 time steps. The standard chaotic initial condition in Eq. (14) was used.




Fig. 8. The variation of the local Lyapunov exponent with $\delta = 10^{-6}$ for Nosé–Hoover dynamics in the chaotic sea with a closeup showing the resolution of a typical large peak. The maximum time shown, 440, corresponds to a total of 286,575 time steps with Δt confined to the interval 10^{-16} to 10^{-14} using double-precision arithmetic.

V. SUMMARY AND OUTLOOK

A. Adaptive integrators and the use of mappings

Present day molecular-dynamics simulations of systems such as liquid water at constant temperature involve a host of practical computational issues.¹⁷ One that we have not discussed is the need for integration techniques that deal with discontinuous terms. The discontinuous jumps of a control variable provide an example. Special precautions need to be taken when the dynamics is singular as in hard-sphere or square-well dynamics, and the smooth separation rate $\lambda_1(t)$ is replaced by a singular map¹⁸ as collisions occur. Although this approach has been commonly used in molecular dynamics simulations since the 1980s, the mapping technique relating post-collision momenta to their pre-collision predecessors is unfamiliar to most workers in dynamical systems simulations. In hard-sphere molecular dynamics, going back to Alder and Wainwright’s pioneering work,¹⁹ with thousands to millions of degrees of freedom, it is usual to integrate up to the time of the next collision, change the momenta of the colliding particles at the collision (this change is the “mapping” referred to previously), and then continue until the next collision.

B. Characterizing chaos with Lyapunov instability

If chaos is present, which is typical of interesting problems, Lyapunov instability is most easily quantified by following two neighboring trajectories. The relative motion of the two trajectories can be constrained by using a Lagrange multiplier.²⁰ Alternatively, the distance between the two trajectories can be rescaled. Either way the original separation length is recovered at the end of each time step. As we discussed, the logarithm of the scale factor required for the rescaling is simply related to the local largest Lyapunov exponent. These two approaches provide the largest Lyapunov exponent $\lambda_1 = \langle \lambda_1(t) \rangle$, the mean separation rate over a long simulation. The Lagrange multiplier approach can also be applied to a linearized version of the equations at the expense of additional algebra. All three of these techniques are useful tools that can benefit from adaptive integration.

Joseph Ford²¹ emphasized that numerical methods are unable to follow a chaotic trajectory accurately for very long. The only convincing test of accuracy is the reproducibility of the trajectory itself. The simple reversal of a trajectory or conservation of energy are not reliable criteria for accuracy.¹⁶ Numerical chaotic trajectories can provide weak averages,

which are accurate despite the lack of global accuracy of the computed trajectory. This good fortune is likewise typical of event-driven dynamical systems. No numerical methods are capable of precise solutions of such problems. Inaccurate though they may be, adaptive integrators are a useful tool for producing reasonable chaotic trajectories.

C. Unique features of Nosé dynamics

Because Eqs. (4), (6), and (7) share common trajectories the elapsed times of their trajectories are related to their Lyapunov exponents in an intriguing way, unique in our experience with dynamical systems. For $T = 1$ the exponents and the elapsed times for the three models are related by $\lambda_N = 3.26\lambda_{NH}$ and $t_{NH} = 3.26t_N$. These ratios are long-time limits for which the exponents and the speeds have converged to their limiting values. The Nosé trajectory is the fastest of the three. The differences are the result of the slower rates (by a common factor $s < 1$) in the two versions of Nosé–Hoover dynamics. The travel times for a long trajectory in (q, p, s, ζ) space are necessarily related by the same ratio, with $t_{NH} = 3.26t_N$. The long time averaged growth rate of the slower Nosé–Hoover trajectories leads to a smaller common Lyapunov exponent $\lambda_{NH1} = \lambda_{NH2} = \lambda_N/3.26$. This relation between the time-averaged exponents does not hold for the local exponents $\lambda_{NH}(t)$ and $\lambda_N(t)$.

Why not? It is evident that both dynamics share the same reference trajectory, so it might appear that the neighboring satellite trajectory determining $\lambda(t)$ would also be the same for both dynamics. But because the local exponent responds to the time rate of change of the scale factor $s(t)$, there is no simple relation linking the local exponents. Nevertheless, the relation between the rates, $(d/dt)_{NH} = s(d/dt)_N$, holds for any (q, p, s, ζ) trajectory segment with a vanishing Hamiltonian. The relation linking all the rates leads directly to the useful relations

$$t_{NH} = \langle 1/s \rangle_N t_N, \quad (18)$$

$$t_N = \langle s \rangle_{NH} t_{NH}. \quad (19)$$

As an example (see Fig. 9), consider the simplest stable periodic orbit, with initial values $(q, p, s, \zeta) = (1.2145, 0.0, e^{-q^2/2}, 0)$, traced out in times of $t_{NH} = 5.5781$ and $t_N = 2.1655$. The ratios of the two times are related to the mean values of s and $1/s$ averaged over the orbit




Fig. 9. (Left) The coordinate q is plotted as a function of s , which is always positive, and ζ . (Right) A plot of p as a function of q . The inner ellipse corresponds to Eq. (4); the outer ellipse represents Eq. (7). The two Nosé–Hoover curves have a period of 5.578, and the Nosé period is 2.1655, useful test cases for a reader’s program. The extreme values of the coordinate q are ± 1.2145 .

$$t_{\text{NH}}/t_{\text{N}} = 5.5781/2.1655 = 2.5759 = \langle 1/s \rangle_{\text{N}}, \quad (20)$$

$$t_{\text{N}}/t_{\text{NH}} = 2.1655/5.5781 = 0.3882 = \langle s \rangle_{\text{NH}}. \quad (21)$$

An interesting feature of the oscillator problems is that the largest Lyapunov exponent λ_1 is independent of temperature [see Fig. 4(a)]. Figure 4(b) for $\langle \lambda_1^2(t) \rangle$ does show a temperature dependence. The fluctuations of the exponents, which are two orders of magnitude larger than $\langle \lambda_1(t) \rangle^2$, and therefore hardly distinguishable from $\langle \lambda_1^2(t) \rangle$, increase with T for $T > 2$.

We can prove the temperature independence shown in Fig. 4. If we consider the temperature-dependent Nosé–Hoover dynamics in Eq. (7) and introduce the new variables

$$Q \equiv (q/\sqrt{T}), \quad P \equiv (p/\sqrt{T}), \quad Z = \zeta, \quad (22)$$

the temperature-independent result shows that the largest Lyapunov exponent (but not its fluctuations) is temperature independent because a linear change of scale leaves the time-averaged logarithmic growth rate unchanged. That is, $\dot{q} = p$ implies that $\dot{Q} = P$, $\dot{p} = -q - \zeta p$ implies that $\dot{P} = -Q - ZP$, and $\dot{\zeta} = (p^2/T) - 1$ implies that $\dot{Z} = P^2 - 1$.

In Sec. II B, we saw that the temperature-dependent Nosé–Hoover equations provide a stationary solution from Liouville’s theorem. That is, $f \propto e^{-q^2/2T} e^{-p^2/2T} e^{-\zeta^2 \tau^2/2}$

implies that Eq. (9) holds. Just as in the simpler case with T and τ equal to unity, we can find a corresponding \mathcal{H}_D with any temperature T as well as an arbitrary relaxation time for the thermostating

$$2\mathcal{H}_D = sq^2 + p^2/s + T \ln(s^2) + s\zeta^2/T\tau^2 \equiv 0. \quad (23)$$

We remarked on the disparity of the Lyapunov exponents for Nosé and Nosé–Hoover dynamics. The largest local Lyapunov exponent is $\lambda_1 = 0.046 \pm 0.001$ for Nosé and $\lambda_1 = 0.0145 \pm 0.0001$ for the two Nosé–Hoover dynamics. Figure 7 illustrates, on two very different scales, a typical excursion of the local Lyapunov exponent for Nosé dynamics. On the scale of Fig. 7(a) the variation looks singular, but it is actually smooth when adaptive integration is used to control the scale of the smallest time step, which is 10^{-9} . Figure 8 shows similar data for the Nosé–Hoover version of the same problem. In that case no special precautions need be taken. A fixed time step of 0.001 is perfectly adequate for accurate estimates of the local Lyapunov exponent.

D. Challenges and ideas for future work

There is a significant variation in the topology of the chaotic and regular solutions for our oscillator models if the




Fig. 10. (Color online) (a) The variation of $\lambda_1(t)$, (b) the adaptive integrator time step Δt , and (c) the time scale factor s . These are three possible criteria for the stiffness of the Nosé oscillator as described by Eq. (4) and shown in the intersection of the chaotic sea with the $p=0$ plane. In general, the larger Lyapunov exponents are found near the outer boundary of the plots with the smaller Lyapunov exponents closer to the origin. (In the online version, red indicates regions of maximum stiffness, blue indicates regions of least stiffness, and green indicates regions of intermediate stiffness.)

thermostating is speeded up by introducing a quick response of ζ in a short time τ

$$\dot{\zeta} = [p^2 - 1] \rightarrow \dot{\zeta} = [p^2 - 1]/\tau^2. \quad (24)$$

In particular, fast thermostating, which speeds up the convergence of the temperature to its target value, generates infinitely many intricately knotted trajectories, of which the trefoil knot (an overhand knot in a closed loop) is the simplest. Many such knots are described and illustrated in Refs. 22 and 23. These simple oscillator models, with quadratic ordinary differential equations (and their nonequilibrium fractal relatives, generated by the motion equations when temperature is a function of the coordinates), can trace out interlinked rings in phase space! An example is described in Ref. 24. These models' chaotic trajectories, with their knots and interlinked rings could easily fill an entertaining and profusely illustrated book on the subject. We encourage readers to explore these topics. In this spirit, we include in Fig. 10 three Nosé $p=0$ sections illustrating the dependence of $\lambda_1(t)$, Δt , and s on (q, ζ) . See either the online version or the arXiv versions for color.

ACKNOWLEDGMENTS

The authors thank Piotr Pieranski for introducing us to his artistic and computational work on knots,²⁵ Xiaosong Yang for extending his work on Nosé–Hoover knots to include the case $\tau = 1$ treated in the present work, and Julyan Cartwright for pointing out his related work on stiffness and Lyapunov exponents.²⁶

^{a)}Electronic mail: hooverwilliam@yahoo.com

¹J. C. Sprott, *Chaos and Time-Series Analysis* (Oxford U.P., Oxford, 2003).

²Wm. G. Hoover, *Molecular Dynamics*, Lecture Notes in Physics (Springer-Verlag, New York, 1986).

³Wm. G. Hoover, *Computational Statistical Mechanics* (Elsevier, New York, 1991).

⁴S. Nosé, "A unified formulation of the constant temperature molecular dynamics methods," *J. Chem. Phys.* **81**, 511–519 (1984).

⁵S. Nosé, "A molecular dynamics method for simulations in the canonical ensemble," *Mol. Phys.* **52**, 255–268 (1984).

⁶H. A. Posch, W. G. Hoover, and F. J. Vesely, "Canonical dynamics of the Nosé oscillator: Stability, order, and chaos," *Phys. Rev. A* **33**, 4253–4265 (1986).

⁷W. G. Hoover, "Canonical dynamics: Equilibrium phase-space distributions," *Phys. Rev. A* **31**, 1695–1697 (1995).

⁸Wm. G. Hoover, "Mécanique de Nonéquilibre à la Californienne," *Physica A* **240**, 1–11 (1997).

⁹C. P. Dettmann and G. P. Morriss, "Hamiltonian reformulation and pairing of Lyapunov exponents for Nosé–Hoover dynamics," *Phys. Rev. E* **55**, 3693–3696 (1997).

¹⁰W. G. Hoover and B. L. Holian, "Kinetic moments method for the canonical ensemble distribution," *Phys. Lett. A* **211**, 253–257 (1996).

¹¹D. Kusnezov, A. Bulgac, and W. Bauer, "Canonical ensembles from chaos," *Ann. Phys. (NY)* **204**, 155–185 (1990).

¹²W. G. Hoover, J. C. Sprott, and P. K. Patra, "Ergodic time-reversible chaos for Gibbs' canonical oscillator," *Phys. Lett. A* **379**, 2935–2940 (2015).

¹³P. K. Patra and B. Bhattacharya, "A deterministic thermostat for controlling temperature using all degrees of freedom," *J. Chem. Phys.* **140**, 064106 (2014).

¹⁴See, for example, <en.wikipedia.org/wiki/Runge-Kutta-Fehlberg_method> for a discussion of the Runge–Kutta–Fehlberg method.

¹⁵The Fortran compiler from the GNU Project is routinely used in numerical work with double precision arithmetic (about 16 decimal digits). For stiff problems quadruple precision (with about 34 decimal digits) is useful. To invoke the latter (which is about fifty times slower), the program code.f producing the executable code xcode should be compiled as gfortran -O -o xcode -freal-8-real-16 code.f rather than gfortran -O -o xcode code.f.

¹⁶Wm. G. Hoover and C. G. Hoover, "Comparison of very smooth cell-model trajectories using five symplectic and two Runge–Kutta integrators," *Comput. Methods Sci. Technol.* **12**, 109–116 (2015).

¹⁷R. L. Davidchack, "Discretization errors in molecular dynamics simulations with deterministic and stochastic thermostats," e-print arXiv:1412.7067.

¹⁸Ch. Dellago and H. A. Posch, "Lyapunov exponents of systems with elastic hard collisions," *Phys. Rev. E* **52**, 2401–2406 (1995).

¹⁹B. J. Alder and T. E. Wainwright, "Molecular motions," *Sci. Am.* **201**, 113–126 (1959).

²⁰W. G. Hoover and H. A. Posch, "Direct measurement of equilibrium and nonequilibrium Lyapunov spectra," *Phys. Lett. A* **123**, 227–230 (1987).

²¹T. Uzer, B. Chirikov, F. Vivaldi, and G. Casati, "Joseph Ford," *Phys. Today* **48**(10), 88 (1995).

²²L. Wang and X.-S. Yang, "The invariant tori of knot type and the interlinked invariant tori in the Nosé–Hoover system," *Eur. Phys. J. B* **88**, 78–82 (2015).

²³L. Wang and X.-S. Yang, "A vast amount of invariant tori in the Nosé–Hoover oscillator," *Chaos* **25**, 123110 (2015).

²⁴P. K. Patra, W. G. Hoover, C. G. Hoover, and J. C. Sprott, "The equivalence of dissipation from Gibbs' entropy production with phase-volume loss in ergodic heat-conducting oscillators," *Int. J. Bifurcation Chaos* **26**, 1650089-1-11 (2016).

²⁵Piotr Pieranski's personal website is located at <etacar.put.poznan.pl/piotr.pieranski/Personal.html>.

²⁶J. H. E. Cartwright, "Nonlinear stiffness, Lyapunov exponents, and attractor dimension," *Phys. Lett. A* **264**, 298–302 (1999).