PHYSICAL REVIEW E VOLUME 52, NUMBER 3 SEPTEMBER 1995

Thermostatted molecular dynamics: How to avoid the Toda demon hidden in Nosé-Hoover dynamics

Brad Lee Holian and Arthur F. Voter
Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545

Ramon Ravelo
Department of Physics, University of Texas, El Paso, Texas 79968

(Received 1 May 1995)

The Nosé-Hoover thermostat, which is often used in the hope of modifying molecular dynamics trajectories in order to achieve canonical-ensemble averages, has hidden in it a Toda "demon," which can give rise to unwanted, noncanonical undulations in the instantaneous kinetic temperature. We show how these long-lived oscillations arise from insufficient coupling of the thermostat to the atoms, and give straightforward, practical procedures for avoiding this weak-coupling pathology in isothermal molecular dynamics simulations.

PACS number(s): 05.40.+j, 02.70.Ns, 05.70.Ln, 05.20.Gg

I. INTRODUCTION

Since the pioneering work of Nosé [1], which was made more practical and accessible for general molecular dynamics (atomistic) simulations by Hoover [2], a large number of constant-temperature studies have been performed [3,4]. The Nosé-Hoover (NH) thermostatting method is motivated by imagining that the system is put into contact with a thermal reservoir, which can contribute to, or take away from, the thermal motion of the sample being studied in a homogeneous way, rather than at a physical boundary (discontinuity). At the heart of the thermostatting techniques using feedback, such as Nosé-Hoover, are two principles: (1) the temperature in a classical system is simply related to the kinetic energy of the atoms; (2) the motion of each atom is altered by an extra term in the acceleration equation that is linear in its velocity, and which bears a strong resemblance to simple viscous damping.

In simple viscous damping, every particle is slowed down by a force opposite to its velocity, with a proportionality constant $\gamma > 0$, so that the temperature drops exponentially with time to zero. In order to deal with nonzero temperatures via viscous damping, Berendsen and co-workers proposed direct feedback [5], where the damping coefficient is no longer strictly a constant, but is instead, a heat-flow *variable*, $\gamma(1-T_0/T)$, that can be either positive or negative, depending on the instantaneous kinetic temperature T of the sample being thermostatted (T_0 is the temperature of the thermal bath in which the sample is homogeneously immersed). While Berendsen's method reduces to simple viscous damping when $T_0 \equiv 0$, it is ill-posed when (for $T_0 > 0$) the initial temperature is chosen to be $T(0) \ll T_0$.

Hoover (and independently, Evans) proposed another approach [6]—differential feedback—where the heat-flow variable satisfies an even more complex constraint, namely, that of holding the kinetic energy of the sample at a constant value. (In the work of Hoover and Evans, the term "Gauss thermostat" is employed for this isokinetic dynamics in honor of Gauss' principle of least constraint.) In order to achieve constant kinetic energy, the heat-flow variable is found to be proportional to the time derivative of the potential energy divided by the temperature T_0. For this reason, the Gauss thermostat is ill-posed if the thermostat temperature $T_0=0$; moreover, if for $T_0>0$, the initial temperature is chosen to be $T(0)\equiv 0$, the Gauss thermostat will forever spin its wheels, to no avail—the temperature will remain at zero. The Gauss thermostat is equivalent, in the limit of small time steps, to scaling particle velocities at each time step—for several years prior to the Gauss thermostat, this was referred to as "ad hoc velocity rescaling."

The Nosé-Hoover thermostat goes one step further in complexity by introducing a new dynamical variable with its own equation of motion—hence, integral feedback. The purpose of this extra dynamical variable is to allow the instantaneous kinetic temperature T to fluctuate, while assuring that its long-time average is equal to T_0. In many ways, the NH thermostat is attractive because it is easy to implement, having only one coupling parameter (the rate of thermostatting), and because it shares with the Gauss thermostat the property of time reversibilty [7], as contrasted with the irreversible approaches of Berendsen and standard viscous damping. The NH thermostat suffers from the same zero-temperature dysfunctionalities as the Gauss thermostat—you cannot achieve zero temperature, and you create something from nothing (i.e., start from zero temperature).

For all three of these examples of deterministic feedback, nonequilibrium steady states are guaranteed to achieve a steady-state temperature equal to T_0, although the Berendsen and Gauss thermostats both require an additional term in the heat-flow variable proportional to the

1063-651X/95/52(3)/2338(10)/$06.00

external driving force. Empirical studies show that transport coefficients, obtained from nonequilibrium steady states using any of these thermostats, are essentially indistinguishable [3].

In this paper, we show how one extra equation of motion (for the heat-flow degree of freedom) in the Nosé-Hoover thermostat can lead to surprising—even disturbing—results if care is not taken in its implementation. That is, the NH thermostat can easily cause the temperature of a many-body classical system to exhibit persistent noncanonical oscillations, as though it were driven by a "demon"—a Toda oscillator [8]. This pathology can be avoided by coupling the thermostat sufficiently strongly to the particles. We motivate this strong coupling by considering the fundamental nature of the thermostatting process. Here, we will be interested primarily in equilibrium states—mean values, fluctuations about the mean, and time correlation functions, or their power spectra.

II. BACKGROUND

We begin by introducing the instantaneous kinetic temperature of a system, whose long-time average—either under standard molecular dynamics (microcanonical ensemble) or under Nosé-Hoover dynamics (canonical ensemble)—gives the thermodynamic temperature; in the thermostatted case, that temperature will be guaranteed to be T_0, while under standard Newtonian mechanics, the temperature will, in general, be different. For N particles in a box of volume V in d spatial dimensions, the total mass M, total linear momentum $M\mathbf{u}$, and the instantaneous temperature T in a molecular dynamics computer simulation can be written as

$$\begin{aligned} M &= \sum_{i=1}^{N} m_i \ , \\ M\mathbf{u} &= \sum_{i=1}^{N} m_i \mathbf{u}_i \ , \\ dNkT &= \sum_{i=1}^{N} m_i |\mathbf{u}_i|^2 \ . \end{aligned} \tag{1}$$

(The atomic mass for particle i is m_i and its velocity is $\mathbf{u}_i$.)

Henceforth in this paper, we will use the term "temperature" and "instantaneous kinetic temperature" interchangeably, even though we wish to emphasize that "thermodynamic temperature" and "long-time average of the instantaneous kinetic temperature at equilibrium" are the only strictly correct equivalences. In nonequilibrium situations, the instantaneous kinetic temperature is the *only* meaningful definition [9]. (Notice that we have not bothered to subtract the center-of-mass velocity $\mathbf{u}$ from the particle velocities in the expression for the temperature. Most often, $\mathbf{u}$ is set to zero, and occasionally one will therefore see the factor of N in the temperature expression replaced by N-1.)

The Nosé-Hoover equations of motion for the N particles and the dimensionless heat-flow variable ζ are given by

$$\begin{aligned} m_i \dot{\mathbf{u}}_i &= \mathbf{F}_i - \nu\zeta m_i \mathbf{u}_i \\ \dot{\zeta} &= \nu \left[\frac{T}{T_0} - 1 \right] , \end{aligned} \tag{2}$$

where ν is the rate of coupling of particles to the NH heat bath, whose temperature is T_0. If ν is set to zero, ζ becomes irrelevant and the particles obey Newton's equation of motion. Common sense suggests that the temperature responds on a time scale of $1/\nu$, so that choosing too small a value of ν would result in a long time for equilibration, while too large a value could perturb the "natural" dynamics of the system [1]. (In the Appendix, we present the Størmer finite-difference approximation—time step δ—for these equations of motion.)

The theoretical groundwork for the Nosé-Hoover thermostat was most clearly stated by Hoover [2,7]: the dynamics should be such that any initial condition leads to a sampling of the canonical phase space distribution function $\rho_0(\mathbf{x})$ after a computationally "reasonable" length of time (here, $\mathbf{x}$ represents the coordinates $\mathbf{q}$ and momenta $\mathbf{p}$ of the N particles, as well as the heat-flow variable ζ):

$$\begin{aligned} \rho_0(\mathbf{x}) &= \frac{1}{Q_0} \exp[-\beta E(\mathbf{x})] \ , \quad \mathbf{x} = (\mathbf{q},\mathbf{p},\zeta) \ , \\ Q_0(N,V,T_0) &= \int d\mathbf{x}\, e^{-\beta E} \ , \quad \beta = 1/kT_0 \ , \\ E(\mathbf{x}) &= K(\mathbf{p}) + \Phi(\mathbf{q}) + \frac{d}{2} NkT_0 \zeta^2 \ . \end{aligned} \tag{3}$$

The canonical partition function Q_0 is the integral over all phase space $\mathbf{x}$ of the Boltzmann factor of the internal energy E, which includes a contribution from the thermostat heat-flow variable (ζ is postulated to be Gaussian distributed with an average contribution of $kT_0/2$, compared to order N particle contributions from kinetic K and potential Φ energies).

The flow of an ensemble of trajectories of independent (noninteracting) many-body systems, begun from different initial conditions, can be represented by a general, time-dependent phase-space distribution function $\rho(\mathbf{x},t)$, which obeys the Liouville (continuity) equation [7]:

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial \mathbf{x}}(\rho \mathbf{u}) = 0 \ , \quad \mathbf{u} = \dot{\mathbf{x}}(\mathbf{x},t) \ , \tag{4}$$

where now $\mathbf{u}$ represents the equations of motion—the "fluid" velocity of the ensemble at each point $\mathbf{x}$ in phase space. Under equilibrium Nosé-Hoover dynamics [Eq. (2)], the canonical distribution $\rho_0(\mathbf{x})$ is the stationary solution of the Liouville equation. This means that if the system is sufficiently mixing, each and every trajectory, begun from any initial condition whatsoever (except for a set of measure zero), visits every small box of volume $\Delta\mathbf{x}$, located at point $\mathbf{x}$ in phase space, with probability $\rho_0(\mathbf{x})\Delta\mathbf{x}$. Consequently, a long-time average over a single trajectory will be equal to a canonical ensemble average. The Nosé-Hoover equations of motion are a necessary, but not sufficient, condition for the dynamics to exhibit canonical behavior at equilibrium. The sufficient condition is that the equations of motion be Lyapunov unstable (that is, excluding integrable systems), so that the motion is chaotic and therefore, ergodic or "mixing." (To be precise, more complex equations of motion—variations upon Nosé-Hoover—have been proposed [10]; while they, too, satisfy this necessary condition, the original NH equations are the simplest, one-parameter set that fill the bill.)

The equations of motion contain the coupling parameter ν, which, if set to zero, gives Newtonian equations of motion appropriate to the microcanonical ensemble. That is, if we set up a canonical distribution of initial conditions for the trajectories, and then propagate them with standard molecular dynamics, the above argument applies: the canonical distribution for the whole ensemble will be preserved. But each trajectory will *not* sample—independently—the canonical distribution; only the entire ensemble will do that. A Nosé-Hoover trajectory *will* sample the canonical distribution— *provided* that the interparticle forces are nonlinear and the motion is sufficiently chaotic.

The Nosé-Hoover equations of motion are deceptively simple—as we have said, the simplest, one-parameter way to achieve canonical dynamics—and from the equation of motion of the heat-flow variable, it is clear that the long-time average of T is *guaranteed* to be the prescribed value T_0 (even under nonequilibrium steady driving by an external force). If ζ is bounded (a reasonable assumption at equilibrium, where we expect it to oscillate about a value of zero—no heat pumped into or out of the system—or at the nonequilibrium steady state, where it fluctuates about some nonzero value), then the long-time average of its equation of motion is zero, and as a consequence, the average temperature is that of the thermostat:

$$\bar{\zeta} = \lim_{t\to\infty} \frac{1}{t}\int_0^t ds\, \dot{\zeta}(s) = \lim_{t\to\infty} \frac{\zeta(t)-\zeta(0)}{t} = 0 = \nu\left[\frac{\bar{T}}{T_0}-1\right]$$
$$\Longrightarrow \bar{T} = T_0 \; . \qquad (5)$$

However, we actually have to do thermostatting simulations to see whether or not the temporal behavior of the temperature is reasonable—the Liouville formalism only says that Nosé-Hoover dynamics makes it *possible* to achieve canonical-ensemble behavior along a trajectory.

At equilibrium, the temperature in a finite system fluctuates about the mean value: $\delta T = T - \langle T\rangle$, where $\langle T\rangle = T_0$ in the NH thermostatted system. If the NH system exhibits canonical fluctuations (within a computationally "reasonable" time), as it should if it is sufficiently mixing, then the variance of the temperature, as determined from the canonical distribution function in Eq. (3), is

$$\langle(\delta T)^2\rangle_0 = \frac{2}{d}\frac{\langle T\rangle^2}{N} \; , \qquad (6)$$

where the subscript "0" indicates the canonical-ensemble value. In contrast, a microcanonical [standard molecular dynamics (MD)] simulation is characterized by a much smaller variance [11]:

$$\langle(\delta T)^2\rangle_{\rm MD} = \frac{2}{d}\frac{\langle T\rangle^2}{N}\frac{C_v^{\varphi}}{C_v} = \langle(\delta T)^2\rangle_0 \frac{C_v^{\varphi}}{C_v} \; , \qquad (7)$$

where the subscript "MD" refers to the microcanoical-ensemble value (remember, in standard microcanonical MD simulations, only by iteration can one achieve $\langle T\rangle = T_0$). The total constant-volume heat capacity is C_v—the kinetic part is $C_v^{\kappa} = (d/2)Nk$; the potential part is C_v^{φ}. For a harmonic system, the ratio C_v^{φ}/C_v is one-half; for a dense fluid, it is somewhat smaller. Note that while the variance of temperature in both microcanonical and canonical ensembles decreases like $1/N$, the ratio is always C_v^{φ}/C_v.

The higher moments of the temperature distribution are the skewness

$$\frac{\langle(\delta T)^3\rangle_0}{\langle(\delta T)^2\rangle_0^{3/2}} = \left[\frac{8}{dN}\right]^{1/2} , \qquad (8)$$

and the kurtosis

$$\frac{\langle(\delta T)^4\rangle_0}{\langle(\delta T)^2\rangle_0^{2}} = 3\left[1+\frac{4}{dN}\right] , \qquad (9)$$

where again, the subscript "0" indicates the canonical result. These higher moments are subject to larger statistical uncertainties than the variance. In practice, only the variance and skewness are really useful for determining whether a system is canonical, and therefore where the NH thermostat has achieved its purpose of producing a mixing system with a computationally reasonable length of time.

III. THE TODA DEMON IN THE NOSÉ-HOOVER THERMOSTAT

We chose to study, as an example of an inherently mixing system, the two-dimensional (2D) Lennard-Jones (LJ) fluid of 90 particles at the normal solid density and at a temperature of about 2.5 times the melting temperature. (The LJ unit of distance is the minimum of the potential r_0; the unit of energy is the well depth of the potential ε; the unit of mass is the atomic mass m; from these, the unit of time t_0 is defined by the relation $\varepsilon = mr_0^2/t_0^2$. In these simulations, we used a time step of $\delta = 0.01t_0$, which was found to give essentially the same results as $\delta = 0.005$; see the Appendix for a discussion of the finite-difference approximations to the equations of motion and the integration errors.) The rate of coupling of the thermostat to the particles was chosen to be $\nu = 1/t_0$.

In Fig. 1, we show the time history of the temperature and the heat-flow variable under Nosé-Hoover thermostatting in the 2D LJ fluid. The fluctuation in the temperature in this simulation is enormous, overshooting the mean value by a factor of 2, and undershooting by half. There is, moreover, a profoundly persistent period of oscillation of about 2π. The initial temperature was set equal to the thermostat temperature $T(0) = T_0 = 1 (= \varepsilon/k)$, and the NH heat-flow variable was set to zero (particles were placed on triangular lattice sites and given random velocities appropriate to the Maxwell-Boltzmann distribution [12]). All these initialization

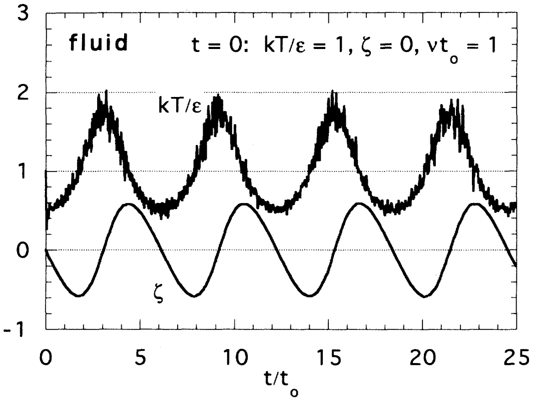

FIG. 1. The temperature T and heat-flow variable ζ for the 2D, periodic, 90-atom LJ fluid at normal density ($\rho=\rho_0$), thermostatted by Nosé-Hoover dynamics at a temperature $T_0=1$ (about 2.5 times melting), with a coupling rate of particles to the thermal bath of $\nu=1$ (LJ units). The initial atomic positions were triangular lattice sites; velocities were chosen from a Maxwell-Boltzmann distribution at temperature $T(0)=T_0$; the heat-flow variable was set at $\zeta(0)=0$.

steps would seem to be perfectly reasonable and certainly within the bounds of common practice for such a MD simulation. However, the observed temperature variance is 17 times the canonical expectation. This is clearly an unwelcome result for a supposedly "thermostatted" system.

Changing the initial temperature from $T(0)=T_0$ to $T(0)=2T_0$ might seem to offer an improvement, based on the rough equipartition of energy expected for a dense fluid or solid. The result, shown in Fig. 2, is clearly preferable to that of Fig. 1. Now, the fluctuations are considerably reduced, both for the temperature and the heat-flow variable; however, the persistent oscillatory period of about 2π (frequency $\omega\approx 1$) is still apparent. Unfortunately, upon closer examination, we find that the temperature variance is *too* much reduced: 0.52 times canonical. Thus, the system is not canonical at all—it is as though the thermostat were not even turned on. This is just as terrible—even though not quite so obvious—as the temperature fluctuations in the first initial condition $[T(0)=T_0]$.

Preparing the sample in a way designed to make the system forget entirely its peculiar initial conditions—that is, by first thermostatting using Berendsen's method ($\gamma=2/t_0$) for a time of $5t_0$, followed by standard microcanonical MD for another $5t_0$, and then applying Nosé-Hoover—also does not solve the problem, as shown in Fig. 3(a) for early times, and in Fig. 3(b) for later times, where the oscillations have finally built up again to the same apparent value as in Fig. 2. Indeed, the variance accumulated over the time span of $t=10$ to $25t_0$ in Fig. 3(a) is 0.41 times canonical, while over the span of $t=135$ to $160t_0$, the variance is 0.53 times the canonical value. The microcanonical variance is 0.46 times the canonical one, while under Berendsen thermostatting, the variance is even lower, 0.41 (the Berendsen temperature depends nonlinearly on the coupling rate: at $\gamma=10$, the value is reduced further to 0.36).

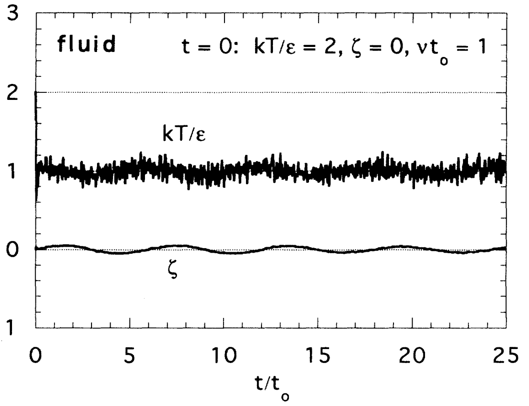

FIG. 2. Temperature and Nosé-Hoover heat-flow variable for the 2D LJ fluid (nominally at the same thermodynamic state as in Fig. 1); initial conditions were $T(0)=2T_0$ and $\zeta(0)=0$.

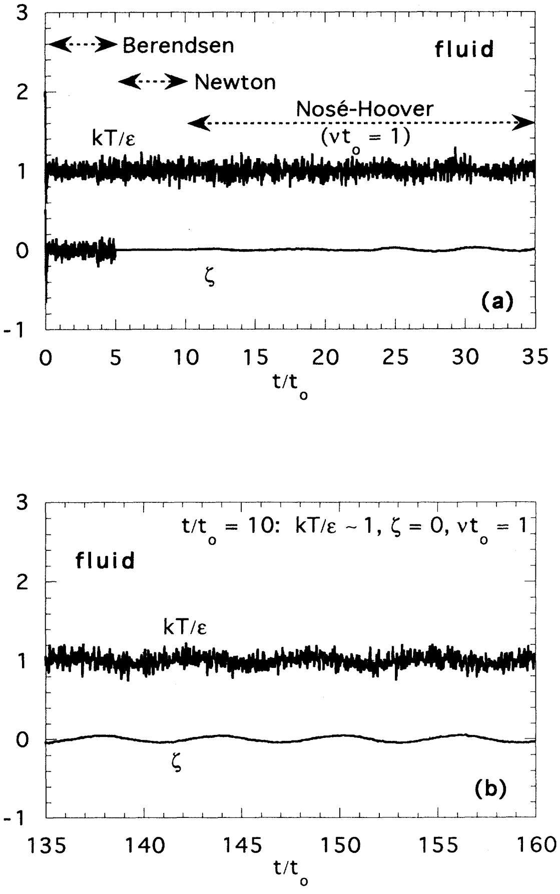

FIG. 3. Careful preparation of 2D LJ fluid by equilibration using Berendsen thermostat, followed by Newtonian mechanics, before thermostatting by Nosé-Hoover (same thermodynamic state as in Figs. 1 and 2): (a) early times; (b) late times (compare with Fig. 2).

In order to determine what becomes of these initial conditions and answer the question—does the "thermostatted" system eventually exhibit canonical behavior over some extremely long-time scale?—we ran these simulations for as much as $10^4 t_0$(10^6 time steps). In the case of Fig. 1, where early fluctuations are too large, the envelope of the temperature slowly damps down to the expected canonical fluctuations after a time of about $2000 t_0$ ($\sim$5000 vibrational periods). Similarly, for Figs. 2 and 3, where early fluctuations are too small, the temperature fluctuations build up to the canonical value on about the same long-time scale. (Apparently, there is no recurrence of the early noncanonical fluctuations, i.e., a superperiodic behavior, such as that seen in the famous Fermi-Pasta-Ulam problem [13], at least on these very long-time scales.) For smaller temperatures, such as $T_0 = 0.1$ (about $\frac{1}{4}$ the melting temperature of the 2D LJ system), we find the time constant for this relaxation of fluctuations toward canonical to be longer by a factor of more than 2. From the point of view of practical molecular dynamics simulations, we would say that this behavior of the thermostat violates the "Principal of Exhaustion of the Observer": it is unreasonable to wait more than a few vibrational periods to attain canonical behavior.

How are we to explain these profoundly disturbing results, namely, the persistent low-frequency oscillation of $\omega \approx 1$ and the failure to attain the canonical temperature variance, by nearly a factor of 2, regardless of initial conditions? Part of the answer lies in a perturbation analysis of the equations of motion, provided that we coarse-grain time, so as to eliminate the high-frequency ($\omega \sim \omega_0$) oscillations, leaving exposed the low-frequency undulation in the temperature that so concerns us. The rate of change in the internal energy of the particles, $E_0 = K + \Phi$, can be related to the heat flow into the system from the thermal bath:

$$\begin{aligned} E_0 &= \sum_{i=1}^{N} \tfrac{1}{2} m_i |\mathbf{u}_i|^2 + \Phi(\{\mathbf{q}\}) , \\ \dot{E}_0 &= \sum_{i=1}^{N} m_i \mathbf{u}_i \cdot \dot{\mathbf{u}}_i + \sum_{i=1}^{N} \frac{\partial \Phi}{\partial \mathbf{q}_i} \cdot \mathbf{u}_i \\ &= \sum_{i=1}^{N} \mathbf{u}_i \cdot (\mathbf{F}_i - \nu \zeta m_i \mathbf{u}_i) - \sum_{i=1}^{N} \mathbf{F}_i \cdot \mathbf{u}_i \\ &= -\nu \zeta \sum_{i=1}^{N} m_i |\mathbf{u}_i|^2 \\ &= -dNk \nu \zeta T . \end{aligned} \tag{10}$$

Now, suppose that we average over the high-frequency oscillations in the system; then, the internal energy can be decomposed into the zero-temperature potential energy, Φ_0, and the thermal part, given by the heat capacity times the temperature:

$$\begin{aligned} \bar{E}_0 &= \Phi_0 + C_v \bar{T} , \\ \dot{\bar{E}}_0 &= C_v \dot{\bar{T}} , \end{aligned} \tag{11}$$

where the bar over the symbols indicates this coarse-grained average over roughly one vibrational period.

The coarse-grained average of Eq. (10) involves the average of ζT, which itself can be obtained in the same manner as the average temperature [recall Eq. (5)]:

$$\begin{aligned} \overline{\frac{d}{dt} \tfrac{1}{2} \zeta^2} = \overline{\zeta \dot{\zeta}} = \nu \left[\frac{\overline{\zeta T}}{T_0} - \bar{\zeta} \right] = 0 \\ \Longrightarrow \overline{\zeta T} = \bar{\zeta} T_0 = \bar{\zeta}\, \bar{T} . \end{aligned} \tag{12}$$

That is, the average of the product is the product of the averages. Of course, this result is strictly true only in the long-time limit, while we wish to use it for the intermediate time scale between the fast atomic vibrations and the slow oscillations due to the thermostat. Therefore, we use the latter form of the equality in Eq. (12); that is, the coarse-grained average temperature is not strictly equal to T_0, but rather executes smoothed oscillations about T_0. Then, we can combine Eqs. (10) and (11), dropping the bar symbol:

$$C_v \dot{T} = -dNk \nu \zeta T \Longrightarrow -\frac{\dot{T}}{T} = \frac{dNk}{C_v} \nu \zeta . \tag{13}$$

Having coarse-grained the variables ζ and T over the high-frequency vibrations of the system, we can introduce a collective variable for the system, which we will call the coordinate of a pseudoparticle, along with its velocity and acceleration:

$$\begin{aligned} r &= -\ln \frac{T}{T_0} \Longrightarrow \frac{T}{T_0} = e^{-r} , \\ \dot{r} &= -\frac{\dot{T}}{T} = \frac{dNk}{C_v} \nu \zeta , \\ \ddot{r} &= \frac{dNk}{C_v} \nu \dot{\zeta} = \frac{dNk}{C_v} \nu^2 \left[\frac{T}{T_0} - 1 \right] = \frac{dNk}{C_v} \nu^2 (e^{-r} - 1) \\ &= -\omega^2 r + \cdots , \end{aligned} \tag{14}$$

where the fundamental harmonic frequency of oscillation of this pseudoparticle is given by

$$\omega = \left[\frac{dNk}{C_v} \right]^{1/2} \nu . \tag{15}$$

For most condensed-phase systems (including anharmonic solids and dense fluids), $C_v \approx dNk$, so that $\omega \approx \nu$.

If time is now scaled by this new frequency, i.e., $s = \omega t$, then the equation of motion of this pseudoparticle is such that it moves in a potential well—the Toda potential:

$$\ddot{r}(s) = e^{-r} - 1 = -\varphi'(r) , \quad \varphi(r) = e^{-r} + r - 1 . \tag{16}$$

This describes the motion of the famous Toda oscillator, which in the nonlinear regime is the temporal analog to the soliton in the Toda chain [8]. As the pseudoparticle goes up the repulsive wall ($r < 0$), r exhibits a sharp peak as a function of time—the temperature of the many-body system is greater than T_0; in the attractive tail ($r > 0$), r goes through a broad valley—the temperature of the many-body system is less than T_0 (see Fig. 1, for example).

Now, it should be clear why Fig. 1 exhibits undulations in the temperature and the heat-flow variable: a

Toda demon (pseudoparticle) has been created in the Nosé-Hoover thermostatted system, and it causes the undulations at frequency $\omega \approx \nu = 1$ to persist for a very long time. Clearly, such behavior is unacceptable, with early thermal fluctuations that are far larger than would be expected in an isothermal system. From simulations showing this pathology, we would expect abnormal predictions for thermally activated processes, such as diffusion. For more gentle initial conditions, as in Figs. 2 and 3, the early thermal fluctuations are hardly bigger than standard MD, that is, a factor of 2 too small. In either case, the thermostat has failed miserably, since the relaxation towards canonical behavior takes an enormous length of time. What can be done to correct this bizarre, noncanonical behavior?

IV. OPTIMIZING THE THERMOSTAT COUPLING

The one free parameter (other than the desired temperature T_0) in the Nosé-Hoover thermostat is the coupling parameter ν, which, as we saw in the preceding section, can induce an unwanted undulation in the instantaneous kinetic temperature of the system at a frequency $\omega \sim \nu$. So as not to greatly perturb the "natural" (microcanonical) time dependence, a seemingly reasonable (and not at all uncommon) choice for ν is a value much smaller than the characteristic frequency ω_0, which can be computed from the second derivative of the pair potential φ'' at the equilibrium zero-temperature, zero-pressure spacing r_0 (the sound speed is $c_0 = r_0\omega_0$). For the LJ system,

$$r_0^2\varphi''(r_0) = 72\varepsilon = m\omega_0^2 r_0^2 (= mc_0^2)$$
$$\Longrightarrow \omega_0 t_0 = \sqrt{72} = 8.49 \ . \qquad (17)$$

From this, one might have imagined that $\nu = 1$ would have been a reasonable choice for the thermostat coupling in the LJ dense fluid. However, our results show that such a choice fails to provide sufficient coupling. What, then, is a reasonable choice that will make the system a mixing one under NH dynamics, on a reasonable time scale?

This question forces us to think more deeply about the very nature of the process of thermostatting a finite system, whereby we wish to mimic a *canonical sample* of N atoms in volume V drawn at random from an infinite system at temperature T_0. In standard molecular dynamics simulations of dense fluids and solids, surface effects, such as density fluctuations near hard walls, are reduced by imposing periodic boundary conditions. That is, we surround the MD sample of N atoms in volume V by a universe of copies of itself. No free surfaces or walls enclose the sample, but the replicated dynamics of neighboring regions imposes restrictions on the power spectrum, or distribution of frequencies (collisional for fluids, vibrational for solids). For example, the maximum wavelength of disturbances allowed in the standard MD system is governed by the sidelength of the periodic box, $L = V^{1/d}$. In contrast, rather than exhibiting discrete peaks broadened by anharmonicity, the frequency spectrum of a canonical sample is continuous, like that of the infinite system. (We have already mentioned other distinctions between a finite-size MD system and a canonical sample: a MD simulation is isoenergetic, while the canonical sample experiences fluctuations in energy about the mean value; the temperature fluctuations in the MD system are roughly half that of the canonical sample.)

Ideally, a thermostatted system, like the hypothetical canonical sample, should experience the full spectrum of frequencies of an infinite sample, in contrast to the restricted set available to a periodic MD system. In order to visualize this important feature of thermostatting, imagine that we try to mimic an infinite, one-dimensional, harmonic chain of atoms by a small sample—let us take, as an extreme example, a single thermostatted atom. What should the thermostat coupling be in order to achieve something like a canonical sample? We will consider first of all the one-dimensional harmonic oscillator, thermostatted by Nosé-Hoover [14]. Then we will relate this to a thermostatted Einstein cell model of a solid, and compare the resulting temperature power spectrum with that of a large chain.

In Fig. 4, we show the power spectrum of the velocity for the unthermostatted ($\nu = 0$) and thermostatted harmonic oscillator ($\nu = 0.05$ and 1). The theoretical spec-

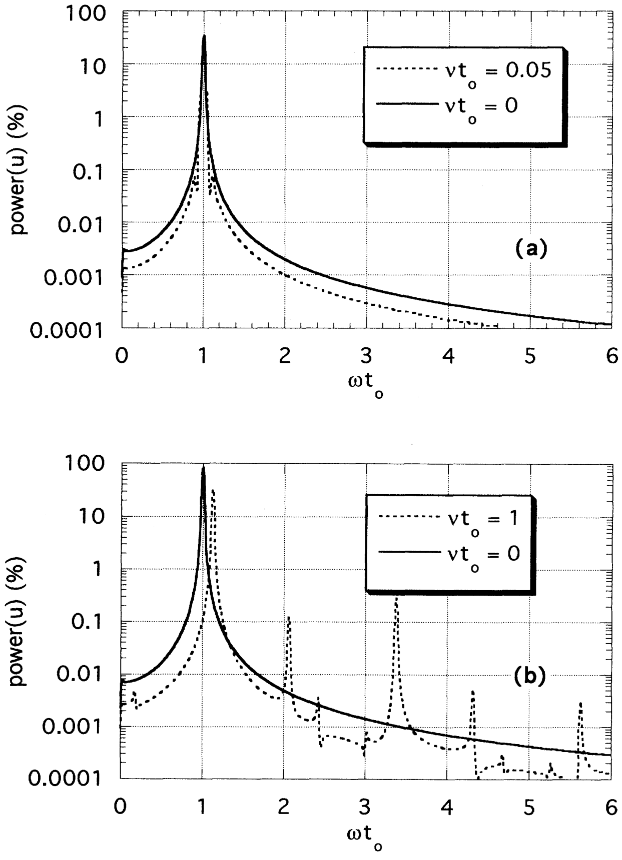

FIG. 4. Power spectrum of velocity in the single-particle 1D harmonic oscillator. Solid curves, no thermostatting; dotted curves, NH thermostatting: (a) weakly coupled ($\nu = 0.05$); (b) strongly coupled ($\nu = 1$). The fundamental frequency $\omega_0 = 1$.

trum for $\nu=0$ is a δ function located at $\omega=\omega_0=1$, which is broadened by numerical errors associated with the finite length of time of the trajectory. In Fig. 4(a) we show the weak-coupling case, where the main peak is only slightly perturbed by the thermostat, producing splittings at $\omega_0 \pm \nu$, $\omega_0 \pm 2\nu$, and so forth. In contrast, Fig. 4(b) shows that setting the thermostat coupling to $\nu=\omega_0=1$ strongly perturb the spectrum by shifting the main peak to slightly higher frequency, as well as producing higher harmonics.

The Einstein cell model, where one atom moves in the potential field of all other atoms frozen at their lattice sites, is characterized by the root-mean-square frequency of the full quasiharmonic crystal, $\omega_E = \langle \omega^2 \rangle^{1/2} \approx (d+1)^{1/2}\omega_0$, for nearest-neighbor interactions in d dimensions. Thus, in 1D, the Einstein frequency is $2^{1/2}\omega_0$. For the cell model, the temperature power spectra—both unthermostatted ($\nu=0$) and thermostatted at the Einstein frequency ($\nu=\omega_E=2^{1/2}$)—are shown together in Fig. 5(a). For the many-body systems we have considered in this paper, we have chosen to study the temperature time series, rather than the more computationally intensive velocity autocorrelation function, from which one could extract the full frequency distribution; the Einstein model example shows that, for our purposes, the simpler temperature-series case exhibits the principal features of interest.

The temperature power spectrum of our hypothetical canonical sample, the thermostatted cell model, mimics that of the full harmonic chain much more closely than the unthermostatted cell model, as shown in Fig. 5(b). Not too surprisingly, the main peak of the thermostatted cell model is almost identical to that of the full chain, and the lower-frequency secondary peaks are similar. The highest frequency in the 1D chain occurs for the binary mode at twice the fundamental frequency (that is, $\omega_{\max}=2\omega_0=2$), so that for the *temperature* power spectrum, the maximum frequency is 4. On the other hand, in the thermostatted Einstein model, even higher frequencies are excited, although (other than the first harmonic of the main peak) they are of relatively minor importance and tend to follow the high-frequency envelope of the chain (power is represented logarithmically).

While the 1D harmonic chain is far from an interesting thermodynamic system, its temperature power spectrum and that of a one-particle approximation—the thermostatted Einsten cell model—provide an important clue for the resolution of the question: What is the "right"

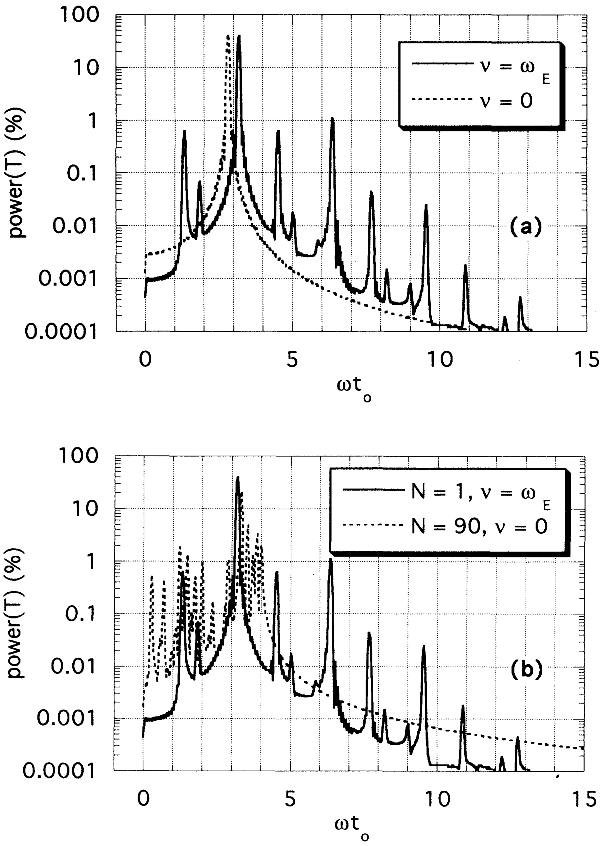

FIG. 5. Power spectrum of temperature in the 1D harmonic Einstein cell model. Solid curves, NH thermostatting, strong coupling ($\nu=\omega_E$=Einsten frequency); dotted curves, no thermostatting ($\nu=0$): (a) cell model ($N=1$); (b) comparison of cell model with periodic harmonic chain ($N=90$).

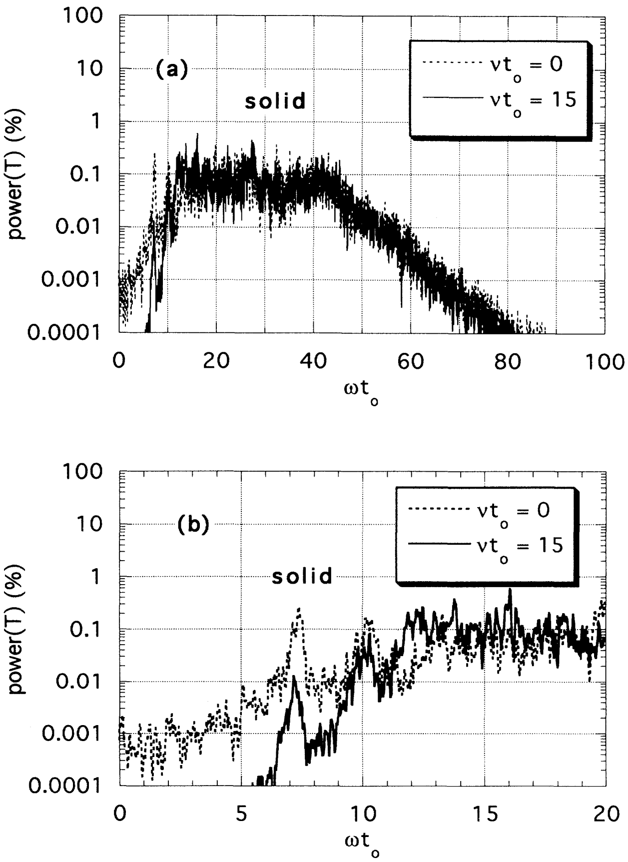

FIG. 6. Power spectrum of temperature in the 2D LJ solid at normal density. Solid curves, NH thermostatting at $T_0=0.1$ (about 0.25 times melting) and $\nu=15\approx\omega_E$; dotted curves, no thermostatting: (a) full frequency range; (b) low frequencies.

amount of coupling of the thermostat to the particles, such that canonical-ensemble fluctuations in temperature are obtained? The answer would appear to be to couple strongly by setting $\nu=\omega_E$.

In order to avoid the pathologies of either 1D or purely harmonic systems, we return now to the 2D LJ system at about 0.25 and 2.5 times the melting temperature: the temperature power spectra are shown in Figs. 6 and 7, respectively. We have set the coupling rate of the thermostat to the system of 90 atoms equal to the Einstein frequency, which, for the 2D LJ solid at normal density (ρ_0) is given approximately by $3^{1/2}\omega_0=14.7$. The temperature power spectra of both thermostatted and unthermostatted systems are remarkably similar for both solid and fluid, with principal differences in the lower-frequency domain, where the thermostatted systems have less power than the standard MD counterparts.

In Fig. 8, we show the distribution of temperature for the LJ fluid at about 2.5 times the melting temperature for microcanonical ($\nu=0$) and Nosé-Hoover thermostatting, both weakly coupled ($\nu=1$) and strongly coupled ($\nu=15$). It is very clear that the temperature distribution in the weakly coupled thermostatted case is hardly any different from that of standard MD, while the strongly coupled thermostat gives a significantly different distribution. In Table I, we show that, in fact, strongly coupled thermostatted systems are canonical, while weak coupling fails to produce canonical temperature fluctuations, at least on a practical computational time scale. (Statistics for temperature fluctuations were typically gathered over a time of about $100t_0$, or ~ 250 vibrational periods.)

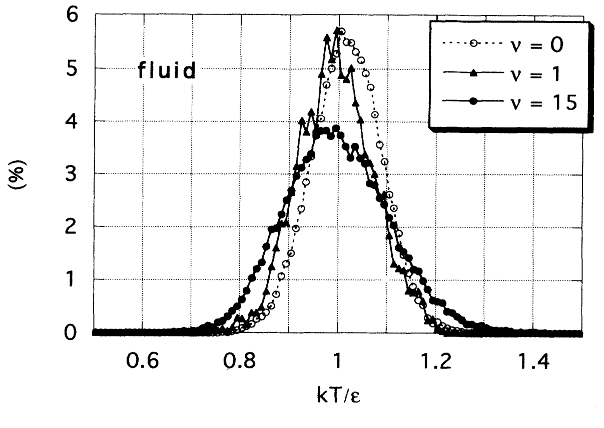

FIG. 8. Temperature distribution for the 2D LJ fluid ($T_0=1$) for thermostatting rates $\nu=0$, 1, and 15.

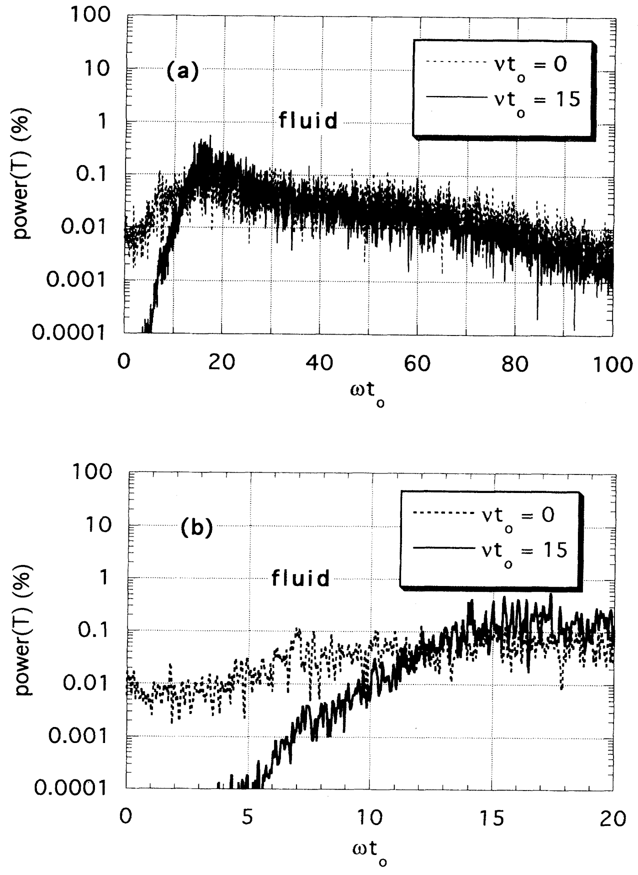

FIG. 7. Power spectrum of temperature in the 2D LJ fluid at normal density. Solid curves, NH thermostatting at $T_0=1$ (about 2.5 times melting) and $\nu=15\approx\omega_E$; dotted curves, no thermostatting: (a) full frequency range; (b) low frequencies.

An estimate for the minimum frequency at which coupling could be effective is that of the longest wavelength disturbance allowed by the periodic boundary conditions we have imposed in both thermostatted and unthermostatted 90-atom LJ systems: $\nu=\omega_{\min}=2\pi c_0/L\approx 6$ (for comparison, the equivalent system size in 3D is $90^{3/2}=854$ particles). As we see in Table I, this rate is only marginally capable of producing a canonical sample. Hence, one should use strong coupling, like $\nu=\omega_E$, to ensure that the Nosé-Hoover thermostat produces—within a couple of vibrational periods—a mixing system with canonical fluctuations. This advice applies even for the

TABLE I. Temperature variance and skewness for the 2D LJ system (90 atoms at normal density). T_0 is thermostat temperature, ν is coupling rate of thermostat to particles. (From comparison of values in independent trajectory segments, the reported variance results have an error of about 3%, while skewness errors are about 10%.) The Einstein frequency is 14.7.

T_0	ν	Variance[a]	Skewness[a]
1.0	0	0.46	−0.1
	1	0.53	0.2
	6	0.81	0.5
	10	1.07	0.6
	15	1.00	1.0
0.1	0	0.50	0.1
	1	0.54	0.2
	10	1.05	1.1
	15	1.03	1.2
0.01	1	0.72	−0.1
	10	1.31	0.9

[a] Observation divided by canonical expectation [see Eqs. (6) and (8)].

thermostatting of low-temperature (nearly harmonic) solids, though the inherent limitation to $T_0 > 0$ is unavoidable for the NH thermostat.

To see how far this strong coupling can be pushed at low temperatures, we also studied the 2D LJ solid ($N=90$) at a temperature $T_0=0.01$ ($\frac{1}{40}$ of the melting temperature), with $\nu=10$. There, we found the temperature variance to be 1.31 times the canonical value (compared to 1.05 at $T_0=0.1$). For classical harmonic systems at very low temperatures, there is a tendency for the trajectories to avoid the region of the so-called "Hoover hole" where all the momenta are equal to zero [10]. Hence, as temperature drops somewhere between 0.1 and 0.01, fluctuations begin to be enhanced, compared to the canonical value, under NH thermostatting.

So far, we have only considered the low-frequency case ($\nu \ll \omega_E$) of inefficient coupling of thermostat to particles. Interestingly enough, the case of overdriving by the thermostat ($\nu \gg \omega_E$) can also be analyzed by the method outlined in Sec. III: again, a Toda demon emerges. In this other extreme limit of poor coupling, the interatomic forces become unimportant, so that the system behaves like a thermostatted ideal gas. Using the ideal-gas heat capacity in Eq. (15), we predict that the temperature oscillations should be driven at a frequency of $\omega=2^{1/2}\nu$. For the 2D LJ fluid at $T_0=1$ and with $\nu=100$, we find that the temperature exhibits a regular oscillatory period of 0.043, which compares favorably with the predicted value of 0.044. Once again, the early temperature fluctuations are not canonical. This overdriving limit is only of academic, rather than practical interest, since the thermostat dominates the dynamics and forces us to reduce the time step.

A final comment is in order for molecular systems: if a molecule has internal vibrational frequencies that are quite high, a global heat-flow variable acting on the atomic velocities may well not equilibrate the bond-vibrational, torsional (bond-bending, etc.), rotational, and molecular center-of-mass vibrational components. Each of these should then be thermostatted by separate heat-flow variables, so as to achieve better coupling to the high-frequency internal modes [15].

V. CONCLUSIONS

The Nosé-Hoover thermostat, because it represents integral feedback, has inherent in it the possibility of unwanted oscillations in the instantaneous kinetic temperature. We have shown that the characteristic period of these undulations and their amplitude are the predictable consequence of a Toda demon hidden in the Nosé-Hoover dynamics, at coupling that is either too low or too high. Their presence in a thermostatted MD simulation, along with temperature fluctuations differing from the canonical expectation, are symptomatic of poor coupling of the thermostat to the system. Such undesirable effects can be eliminated by choosing a coupling rate ν that is equal to the characteristic vibrational (or, in dense fluids, collisional) frequency of the N-body system of interest. For dense fluids and solids, even down to rather low temperatures, the resulting equilibrium averages will then be those of a canonical-ensemble system at the desired thermostat temperature T_0, showing canonical fluctuations about the mean.

ACKNOWLEDGMENTS

We are grateful to Bill Hoover and Harold Posch for stimulating discussions on this subject and for helpful comments on the manuscript. Work done at Los Alamos is supported by the U.S. Department of Energy.

APPENDIX: STØRMER INTEGRATION

The Nosé-Hoover equations of motion can be integrated by Størmer finite differences, where coordinates and the heat-flow variable are evaluated at integral values of the time step δ (as are forces, which depend on coordinates) while velocities are staggered by $\delta/2$ (as is the temperature):

$$x(t)=x(t-\delta)+u(t-\tfrac{1}{2}\delta)\delta+O(\delta^3)\ ,$$

$$\zeta(t)=\zeta(t-\delta)+\nu\left[\frac{T(t-\frac{1}{2}\delta)}{T_0}-1\right]\delta+O(\delta^3)\ , \qquad \text{(A1)}$$

$$u(t+\tfrac{1}{2}\delta)=\frac{u(t-\frac{1}{2}\delta)[1-\frac{1}{2}\nu\zeta(t)\delta]+\dfrac{F(t)}{m}\delta}{1+\frac{1}{2}\nu\zeta(t)\delta}+O(\delta^3)\ .$$

(For simplicity, we have dropped vector notation and particle indices.) These central-difference equations [Eq. (A1)] are perfectly time reversible, as are the Nosé-Hoover equations of motion themselves [Eq. (2)]. By this we mean that reversing time, velocities, and the heat-flow variable results in the same equations of motion; thus, a movie of the motion when run backwards through a projector satisfies the same mathematical description, even though it may violate our sensibilities of macroscopic irreversibility (Second Law of Thermodynamics) [7].

In order to achieve a stable and sufficiently accurate representation of a realistic trajectory, the time step δ, under most circumstances (densities near normal and temperatures up to about twice the melting temperature), should be chosen to be $\frac{1}{30}$ to $\frac{1}{60}$ of the Einstein vibrational period, which is a straightforward calculation, given the bulk modulus and its pressure dependence. For example, at normal density in d dimensions, the Einstein frequency ω_E is approximately the square root of $d+1$ times the fundamental frequency ω_0 at the minimum of the pair potential. The individual local errors in the Størmer solutions shown above for coordinates, velocities, and the heat-flow variable, are formally of order δ^3, but when combined to produce a time series of coordinates (velocities being secondary in that they can be differenced in a variety of ways), they give a local error of δ^4 for the second-order ordinary differential equations of motion. Integrating (i.e., finite differencing) the coordinate for a 1D harmonic oscillator (fundamental frequency ω_0), for example, over a finite time interval t for times before complete dephasing, $\omega_0 t \ll 24/(\omega_0\delta)^2$, gives a global error of the order $\omega_0^3\delta^2 t/24$. The Størmer method is usually referred to as "second-order accurate." Should higher accuracy be desired for some reason—and therefore a smaller time step needed than that suggested above—higher-order methods (such as fourth-order accurate Runge-Kutta) could be used, but they become computationally faster only for time steps of order $\frac{1}{120}$ of the Einstein period or smaller, and also require more computer memory than Størmer. In the work presented here, we have used the Størmer method and checked the results by halving the time step.

---

[1] S. Nosé, J. Chem. Phys. **81**, 511 (1984).

[2] W. G. Hoover, Phys. Rev. A **31**, 1695 (1985).

[3] D. J. Evans and B. L. Holian, J. Chem. Phys. **83**, 4069 (1985). These authors failed to include the external force dependence in the heat-flow variable for the Berendsen thermostat under Couette shear flow; consequently, the temperature achieved at the steady state was not quite equal to the one desired. Nevertheless, they were able to show that the Berendsen thermostat, like Gauss and Nosé-Hoover, produces the same shear viscosity for the same strain rate, at the observed temperature and density.

[4] See references in (a) *Molecular Dynamics Simulations of Statistical Mechanical Systems,* edited by G. Ciccotti and W. G. Hoover (North-Holland, Amsterdam, 1986); (b) *Microscopic Simulations of Complex Hydrodynamic Phenomena,* edited by M. Mareschal and B. L. Holian (Plenum, New York, 1992).

[5] H. J. C. Berendsen, J. P. M. Postma, W. F. van Gunsteren, A. DiNola, and J. R. Haak, J. Chem. Phys. **81**, 3684 (1984).

[6] W. G. Hoover, A. J. C. Ladd, and B. Moran, Phys. Rev. Lett. **48**, 1818 (1982); D. J. Evans, J. Chem. Phys. **78**, 3297 (1983).

[7] B. L. Holian, W. G. Hoover, and H. A. Posch, Phys. Rev. Lett. **59**, 10 (1987); B. L. Holian, G. Ciccotti, W. G. Hoover, B. Moran, and H. A. Posch, Phys. Rev. A **39**, 5414 (1989); B. L. Holian, H. A. Posch, and W. G. Hoover, *ibid.* **42**, 3196 (1990). These three papers constitute a rather complete account of the fundamentals of modern nonequilibrium statistical mechanics, with particular emphasis on issues important to nonequilibrium molecular dynamics computer simulations, including the macroscopic irreversibility of the Second Law of Thermodynamics arising from time-reversible microscopic equations of motion and the consequences for nonlinear response theory.

[8] M. Toda, J. Phys. Soc. J **22**, 431 (1967).

[9] W. G. Hoover, B. L. Holian, and H. A. Posch, Phys. Rev. E **48**, 3196 (1993).

[10] G. J. Martyna, M. L. Klein, and M. Tuckerman, J. Chem. Phys. **97**, 2635 (1992). The authors (who invented the term "Hoover hole") proposed an extension to Nosé-Hoover dynamics, namely, a hierarchy of thermostats; unfortunately, this additional complexity means that the nonequilibrium temperature can no longer be guaranteed to be the desired value.

[11] J. L. Lebowitz, J. K. Percus, and L. Verlet, Phys. Rev. **153**, 250 (1967).

[12] See, for example, B. L. Holian, O. E. Percus, T. T. Warnock, and P. A. Whitlock, Phys. Rev. E **50**, 1607 (1994).

[13] E. Fermi, J. G. Pasta, and S. M. Ulam, Los Alamos Scientific laboratory Report No. LA-1940, 1955, reprinted in *Collected Works of Enrico Fermi, Vol. II,* edited by E. Amaldi, H. L. Anderson, E. Persico, F. Rasetti, C. S. Smith, A. Wattenberg, and E. Segré (Chicago University Press, Chicago, 1965), p. 978.

[14] H. A. Posch, W. G. Hoover, and F. J. Vesely, Phys. Rev. A **33**, 4253 (1986); B. L. Holian and W. G. Hoover, *ibid.* **34**, 4229 (1986).

[15] See B. L. Holian, in *Molecular Dynamics Simulations of Statistical Mechanical Systems* [Ref. [4(a)], p. 241.