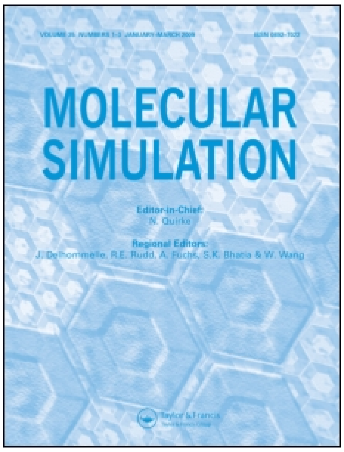

Molecular Simulation
Publication details, including instructions for authors and subscription information:
http://www.informaworld.com/smpp/title~content=t713644482

An Iterative Variable-timestep Algorithm for Molecular Dynamics Simulations

Steven J. Stuart[a]; Jacob M. Hicks[a]; Michael T. Mury[a]
[a] Department of Chemistry, Clemson University, Clemson, SC, USA

To cite this Article Stuart, Steven J. , Hicks, Jacob M. and Mury, Michael T.(2003) 'An Iterative Variable-timestep Algorithm for Molecular Dynamics Simulations', Molecular Simulation, 29: 3, 177 — 186
To link to this Article: DOI: 10.1080/0892702031000089650
URL: http://dx.doi.org/10.1080/0892702031000089650

Molecular Simulation, 2003 Vol. 29 (3), pp. 177–186

An Iterative Variable-timestep Algorithm for Molecular Dynamics Simulations

STEVEN J. STUART*, JACOB M. HICKS and MICHAEL T. MURY

Department of Chemistry, Clemson University, Clemson, SC 29634, USA

(Received July 2002; In final form September 2002)

A method for performing variable-timestep molecular dynamics integration is described, in which an iterative algorithm is used to select the largest timestep consistent with the desired simulation accuracy. Accuracy in this context is defined in terms of energy conservation, rather than trajectory correctness. Specifically, a timestep-independent measure of the rate of "diffusion" of the total energy is used. This variable timestep approach is compared to fixed-timestep integration for three different hydrocarbon systems (polyethylene, liquid benzene and ethylene), which are modeled with a reactive bond-order potential. These systems represent both equilibrium and highly non-equilibrium systems at temperatures ranging from 298 to 2500 K. The variable-timestep method is found to be approximately twice as computationally efficient as fixed-timestep integration for the non-equilibrium sputtering of polyethylene, and the two methods were competitive for the equilibrium systems. The algorithm requires the specification of two parameters controlling the rates of timestep growth and decay, but it is found that one set of values is appropriate for all three systems studied, and there is reason to believe that the parameters are transferable to other systems. The algorithm was developed specifically for simulations involving disparate timescales, such as are encountered with the reactive bond-order model used here, but it should also prove beneficial for a wide range of molecular dynamics applications.

Keywords: Variable-timestep algorithm; Molecular dynamic simulations; Diffusion; Reactive bond-order model

INTRODUCTION

Classical molecular dynamics (MD) simulations involve the approximation of continuous dynamics by a discrete-time trajectory, using finite timesteps. This discretization of the continuous trajectory is one of the largest sources of numerical error in an MD simulation, and the chaotic nature of most systems results in rapid growth of these integration or trajectory errors with time. Because MD is often used as a sampling technique, however, these integration errors only become a problem when they result in sampling from an ensemble other than the one intended. Consequently, failure to conserve energy is widely used as a measure of the accuracy of a simulation in the microcanonical ensemble, as it indicates deviations from the constant-energy ensemble that would be sampled under rigorously correct dynamics. Even for simulations in the canonical or other ensembles, the energy conservation of the underlying integration algorithm is usually of interest. These ensembles are frequently generated either by using an extended Lagrangian or Hamiltonian that can still generate conservative dynamics [1–4], or by applying stochastic perturbations (of velocity, volume, etc.) to an otherwise microcanonical simulation [1,5]. The correct ensemble may not be sampled in these simulations if the underlying integration method does not generate very nearly microcanonical dynamics. Energy conservation is also important for purely practical reasons, as poor energy conservation can lead to numerically induced heating, temperature drift, and even instability.

Because the energy error scales with timestep for any numerical integrator, the simulation timestep is typically determined by what is considered to be an acceptably small integration error for the physical system being studied. Consequently, the acceptable timestep (and thus the maximum practical length of a simulation) is determined largely by the fastest

---

*Corresponding author. E-mail: ss@clemson.edu

ISSN 0892-7022 print/ISSN 1029-0435 online © 2003 Taylor & Francis Ltd
DOI: 10.1080/0892702031000089650

dynamics in the system. This can be a serious problem for highly energetic or non-equilibrium systems, in which the fast dynamics occur only for a small fraction of the simulation, but the computational penalty of using a small timestep is incurred over a much longer time. Examples include simulations of sputtering or energy deposition, in which an initially highly local concentration of energy is eventually thermalized across the entire system; as well as simulations of reactive dynamics, in which much of the simulation is spent integrating between infrequent reaction events that are accompanied by large and quickly fluctuating forces. Even for equilibrium systems, the energy and trajectory errors can be dominated by fairly rare fluctuations into steeply repulsive regions of configuration space, so that the simulation timestep is still determined by the dynamics in a relatively small portion of the simulation.

Variable-timestep integrators can be used to circumvent this problem, by varying the timestep to maintain a roughly constant level of error. In these approaches, the timestep is changed adaptively as the simulation proceeds, and can be viewed as a function of phase space [6]. The timestep is sometimes chosen predictively, based on either heuristics [7–12] or an intimate knowledge of the characteristics of the integrator and the potential [13–15]. Alternatively, the timestep can be selected via an iterative procedure, in which a timestep is rejected if it results in errors that are larger than some predefined tolerance [12,13]. This latter approach is used quite commonly in numerical analysis to integrate ordinary differential equations, but is used much less often in molecular dynamics simulations of complex physical systems. This is due, in part, to the common maxim that, because the force evaluation is by far the most expensive part of the calculation, it does not pay to perform multiple force evaluations per integration step. This rule fails, of course, if the additional work allows the use of sufficiently large timesteps to compensate for the extra force evaluation(s). It is generally true that an extra force evaluation per step does not provide sufficient improvement in accuracy to allow for a doubling of the timestep, thus reinforcing this rule of thumb. The purpose of this paper, however, is to demonstrate the utility of an extremely simplistic iterative variable-timestep integrator, in which an extra force evaluation is needed only in a small fraction of the timesteps. We demonstrate that this algorithm is considerably more efficient than constant-timestep integration for non-equilibrium systems, where a variable-timestep integrator is most appropriate. Somewhat surprisingly, the efficiency is also competitive with fixed-timestep integration for equilibrium systems, at least for the reactive, condensed-phase hydrocarbon systems studied here.

In the following section, the integrator is described, along with the means of parameterizing and characterizing it. The integrator is then compared to the standard fixed-timestep velocity Verlet integrator for three different hydrocarbon systems, ranging from equilibrium to highly non-equilibrium. Finally, we present some discussion and conclusions regarding the effectiveness of the integrator and its applicability to other systems and other potentials.

METHOD

Molecular dynamics simulations make use of finite-difference integration algorithms. Most commonly used (explicit) integrators can be derived as discretizations of the classical time propagator, $U(t)$,

$$U(t) = e^{-iLt} = \exp\left[\left(v\frac{\partial}{\partial x} + \frac{F}{m}\frac{\partial}{\partial v}\right)t\right] \tag{1}$$

which is the operator that displaces state points $\boldsymbol{\Gamma}(t)$ forward in time,

$$\boldsymbol{\Gamma}(t_0 + t) = U(t)\boldsymbol{\Gamma}(t_0) \tag{2}$$

Here, L is the Liouvillian; x, v and F are the Cartesian position, velocity and forces; and we assume a one-dimensional system for simplicity of notation. The propagation over time t can be factored without error into N individual timesteps $\Delta t \equiv t/N$, so that

$$U(t) = [U(\Delta t)]^N. \tag{3}$$

The approximation in a numerical integrator arises only when the single-step propagator $U(\Delta t)$ is factored in a way that ignores the coupling of position and velocity in Eq. (1). For example, using a Trotter factorization of Eq. (1),

$$\begin{aligned} G(\Delta t) &= \exp\left(\frac{\Delta t}{2}\frac{F}{m}\frac{\partial}{\partial v}\right)\exp\left(v\frac{\partial}{\partial x}\right)\exp\left(\frac{\Delta t}{2}\frac{F}{m}\frac{\partial}{\partial v}\right) \\ &\approx U(\Delta t), \end{aligned} \tag{4}$$

generates [16] the well-known velocity Verlet integrator [17], equivalently represented by the algorithm

$$v\left(t_0 + \frac{\Delta t}{2}\right) = v(t_0) + \frac{F(t_0)}{m}\frac{\Delta t}{2} \tag{5a}$$

$$x(t_0 + \Delta t) = x(t_0) + v\left(t_0 + \frac{\Delta t}{2}\right)\Delta t \tag{5b}$$

$$v(t_0 + \Delta t) = v\left(t_0 + \frac{\Delta t}{2}\right) + \frac{F(t_0 + \Delta t)}{m}\frac{\Delta t}{2}. \tag{5c}$$

The decoupling of position and velocity occurs by assuming that velocity updates occur under constant force, and position updates occur under constant velocity, an assumption which only becomes valid in the limit of small Δt. A variety of integrators can be generated by different decompositions of the time propagator [16].

All such integrators can be extended quite easily to include variable timesteps. The time propagator is still written as a product of single-step propagators, but the restriction that the timesteps be equal is relaxed,

$$U(t) = \prod_{i=1}^{N} U(\Delta t_i), \tag{6}$$

$$t = \sum_{i=1}^{N} \Delta t_i \tag{7}$$

The way that these variable-timestep integrators are implemented in practice is to update the positions and velocities as usual for one timestep using the velocity Verlet or other integration algorithm. Then, between timesteps, the value of Δt_i can be changed to a new value before beginning the next timestep. This is straightforward for any integrators except those (such as the leapfrog algorithm) that can not be represented by discrete timesteps as in Eq. (6). The only practical problem, of course, is in specifying the timestep Δt_i at each step.

The goal in a simulation with a variable-timestep integrator is to choose the timestep at each step such that the integration error has a predictable size, or is at least bounded by some limit. At a finite timestep size, the actual trajectory or energy errors incurred during a particular step depend on the current state, as well as higher-order derivatives of the potential energy surface; that is, the errors are functions of phase space. The problem of predicting the relationship between timestep size and energy error is a very well-studied problem in numerical analysis, and has no solution for the general case [13]. Consequently, predictive variable-timestep methods are used most frequently with simple, well-understood potentials, such as the harmonic oscillator or few-body gravitational systems [14,18].

In molecular dynamics simulations with more complex potentials, *ad hoc* rules for timestep selection may be used instead. Frequently, magnitudes of the force, $\mathbf{F}$, and its time derivative, $\dot{\mathbf{F}}$, are assumed to correlate well with numerical integration errors, and these are used to predict the simulation timestep [9,11]. Although analytical expressions for $\dot{\mathbf{F}}$ (or, equivalently, the third time derivative of the position) are typically not practical for any but the simplest potentials [19,20] finite-difference approximations can be used to obtain $\dot{\mathbf{F}}$. In a typical method, the timestep can be chosen such that $\|\dot{\mathbf{F}}\|\,\Delta t$ is approximately constant. In our experience, however, it is extremely difficult to find correlations between $\dot{\mathbf{F}}$ and energy errors that are strong enough to allow efficient prediction of timestep size in a variable-timestep integrator. This is true for a variety of measures of $\|\dot{\mathbf{F}}\|$, including the Euclidean norm (the square root of the sum of the squares) across all particles, or only for the particle with the maximum value of $\dot{\mathbf{F}}$. So although accurate prediction of the energy errors is, in principle, possible using the local features of the potential energy surface, this prediction is impractical for all but the simplest systems.

Rather than attempting to predict the energy errors for dynamics on complex potentials, we will instead find it useful to consider the energy errors to be the result of a random process. When a simulation is run with a fixed timestep size, a series of single-step energy errors results. These errors are not stochastic, of course, and are completely determined by the potential, the initial conditions, and the integration algorithm. But for simulations much longer than the autocorrelation time of this series, the energy conservation characteristics can be determined (in a statistical sense) by knowing only the probability distribution of energy errors, rather than the full time series. It is also important to recognize that the energy conservation behavior of a simulation is dominated by the tails of this distribution; much of the energy error in a fixed-timestep simulation is often due to large errors in a rather small fraction of the steps.

Thus, rather than attempting to accurately predict the error at each point in the trajectory, we attempt to alter the overall distribution in such a way that the tails can be diminished. For the algorithm described here, this is done in a rather simple-minded manner. Beginning with an initial timestep (chosen arbitrarily), the dynamics are integrated as usual for one step. (The particular integrator used—Verlet, velocity Verlet, etc.—does not matter here). If the energy of the resulting state differs from that of the initial state by more than some tolerance ϵ to be defined below, then the system is restored to the initial state, the trial timestep is multiplied by some factor β (where $0 < \beta < 1$), and continued attempts are made at integration until the step can be taken with a tolerable energy error. Once a timestep is found that does provide sufficient energy conservation, the integration step is accepted and the timestep is multiplied by a factor α (where $\alpha > 1$). This is done to prevent the timesteps from decreasing monotonically.

More formally, we define the discretized time in the simulation as

$$t_n = t_0 + \sum_{i=1}^{n} \Delta t_i, \tag{8}$$

where t_n is the time at the end of the nth timestep. To determine the value of Δt_{j+1}, we make use of an energy error function $\Delta H_j(\Delta t)$,

$$\Delta H_j(\Delta t) = H(t_j + \Delta t) - H(t_j), \tag{9}$$

which specifies the error in energy conservation made by integrating with a timestep Δt after step j, with $H(t) = H(\boldsymbol{\Gamma}(t))$ specifying the value of the system Hamiltonian at time t. Then we choose

$$\Delta t_j = \alpha\beta^k \Delta t_{j-1}, \tag{10}$$

where k is the smallest non-negative integer satisfying the constraint that

$$\Delta H_{j-1}(\alpha\beta^k \Delta t_{j-1}) < \epsilon, \tag{11}$$

The two constants α and β are required to fully specify the algorithm. The most efficient values of α and β will be system-dependent, but it is expected that β should typically be substantially less than one, so that the timestep can quickly "ratchet down" when encountering a portion of phase space that is especially prone to numerical errors. And α will typically be only slightly larger than one, so that the timestep does not grow so quickly that it requires frequent downward corrections in non-problematic areas of phase space.

Because this algorithm involves multiple force evaluations on those steps that require re-integration, it would appear to be less efficient than standard algorithms. However, the expectation is that (for well-chosen α and β), the overhead associated with discarding steps will be more than compensated by the decrease in computation time that results from using large timesteps when possible, and small timesteps only when necessary. In fact, we note that the number of repeated iterations can be predicted in advance, for a sufficiently long simulation. In order for the average timestep to remain approximately constant over a long simulation, the rates of timestep growth, due to α, and decay, due to β, must cancel one another. Thus, for a simulation of N steps with a fraction f of the steps requiring re-integration, we have

$$\frac{\Delta t_N}{\Delta t_0} = \alpha^{(1-f)N}\beta^{fN}, \tag{12}$$

so that

$$\alpha^{1-f}\beta^f = \left(\frac{\Delta t_N}{\Delta t_0}\right)^{1/N} \approx 1. \tag{13}$$

Specifying α and β thus determines this fraction,

$$f = \frac{\ln\alpha}{\ln\alpha - \ln\beta}. \tag{14}$$

It is interesting to note that this equation must hold (for large N), regardless of the type of system, energy conservation requirements, or individual character of the dynamics.

The largest benefit from this approach is expected for non-equilibrium systems, or systems undergoing infrequent, energetic processes. In those cases, the timestep will become small as needed, and will remain small only until the system exits the region of phase space with large and quickly fluctuating forces. For equilibrium systems, the repeated steps may make this method less efficient than fixed-timestep dynamics at an appropriately chosen Δt. Even so, this method may have advantages if the distribution of energy errors is sufficiently broad.

To specify the energy tolerance that is considered sufficient at any particular timestep, we must define more precisely what we mean by energy conservation. In principle, dynamics in the microcanonical ensemble should preserve the total energy exactly, since the Hamiltonian is a constant of the motion. In practice, the energy errors ΔE between successive steps will not be identically zero, but are considered here to behave as a random variable with a particular distribution, $P(\Delta E)$.

We assume that the energy conservation properties of a particular simulation are described fully by this distribution, and thus by the set of moments of the distribution. The first two moments demonstrate the most significant types of errors in a simulation. A trajectory that exhibits an energy drift over long times will have a nonzero first moment, $\langle\Delta E\rangle$. Trajectories with a zero first moment but nonzero second central moment $\sigma^2_{\Delta E} = \langle(\Delta E - \langle\Delta E\rangle)^2\rangle$ exhibit fluctuations in the energy, but with no long-term drift. Which is more important depends to some extent on the application. Because we are interested in long-time microcanonical integration, drift is a larger problem than fluctuations. Luckily, drift can be reduced to negligible levels by using a symplectic integrator. For example, symplectic integrators such as velocity Verlet [17], position Verlet [16], or higher-order examples [21,22], tend to drift much less than non-symplectic, predictor-corrector-style integrators [13], whereas the latter may exhibit smaller fluctuations. Once the drift is eliminated, the energy fluctuations represent the next largest contribution to the long-time energy errors. Consequently, we will focus on the root mean square (rms) energy errors, $\sigma_{\Delta E}$, as the single measure of energy conservation, and will assume the use of a symplectic integrator to minimize drift.

It is not sufficient to simply set a fixed value of ϵ for the energy tolerance in the iterative variable-timestep integration algorithm (Eq. (11)); ϵ should be timestep-dependent, to ensure that decreasing the timestep results in better energy conservation per unit time, rather than per single timestep. In deriving an appropriate, timestep-dependent value for ϵ, we begin by assuming that, for a particular timestep size

Δt and (symplectic) integrator, $P(\Delta E)$ is Gaussian with zero mean and variance $\sigma^2_{\Delta E}(\Delta t)$. This assumption is at least approximately true for equilibrium systems with reasonable timesteps. If the successive energy errors are independent of one another, the distribution of cumulative energy errors after N such timesteps will also be Gaussian, with zero mean and variance $N\sigma^2_{\Delta E}$. (This also holds if the energy errors are correlated in time, as long as the cumulative energy errors are measured over times that are long compared to the correlation time, although N in this case should be replaced by $N\tau/\Delta t$, where τ is the correlation time). So if we desire that the distribution of energy errors after a total simulation time of t have a standard deviation of $\sigma_{\Delta E}(t)$, we can require that the rms deviation of each individual step satisfy

$$\frac{t}{\Delta t}\sigma^2_{\Delta E}(\Delta t) \le \sigma^2_{\Delta E}(t). \qquad (15)$$

Equivalently,

$$\frac{\sigma^2_{\Delta E}(\Delta t)}{\Delta t} \le \frac{\sigma^2_{\Delta E}(t)}{t} \equiv D_E, \qquad (16)$$

where we have defined $\sigma^2_{\Delta E}(t)/t$ as our scaled, timestep-independent measure of energy conservation, and specified D_E as a target value. We have assumed that the energy is undergoing a Gaussian random walk, so D_E is effectively a diffusion constant of the energy for its Brownian motion in the one-dimensional energy space. In practice, trajectory errors scale as Δt^p for an integrator of order p. The resulting energy errors depend on the potential used, but also typically scale as Δt^p or higher. Thus, since all integrators in common use have $p \ge 2$ we see (reassuringly) from Eq. (16) that the scaled energy errors do decrease with decreasing timestep.

In the iterative variable-timestep scheme described here, a target level of energy "diffusion" is chosen, and individual timesteps are rejected when they exceed the resulting scaled energy error in Eq. (16). The energy tolerance ϵ used in Eq. (11) thus depends on both the size of the timestep and the targeted value of the energy diffusion constant, D_E,

$$\epsilon = \sqrt{D_E \Delta t} \qquad (17)$$

For example, in the simulations to be described below, a reasonable level of energy conservation is to require that the expected rms deviation in the total energy be within 1 eV (i.e. 10^{-3} eV/atom for a 1000-atom system, or a temperature change of $\lesssim 5$ K) after a 100-ps simulation. This implies a targeted energy diffusion constant of $D_E = 0.01\,\mathrm{eV}^2/\mathrm{ps}$, or expected rms deviations of 0.003 eV in the energy for a 1-fs timestep. In the iterative variable-timestep method described above, an absolute limit is imposed on every single-step energy error, $|\Delta E| \le \sqrt{D_E \Delta t}$, rather than limiting the rms deviation (which would require more than one step to evaluate). This absolute limit is more straightforward to implement, and is also more stringent, resulting in an actual rms deviations smaller than $D_E/3$.

RESULTS

The new iterative variable timestep selection algorithm has been tested, and compared to fixed-timestep integration, for three different systems. All of the systems considered are hydrocarbon systems, and are modeled with the AIREBO (adaptive intermolecular reactive empirical bond-order) potential [23]. This potential belongs to the family of empirical bond-order models originated by Tersoff [24,25], and also represented by the popular Brenner potential for hydrocarbons [26,27]. All of these potentials allow reasonable treatment of covalent bonding reactions for various systems, in the context of a purely classical potential. The AIREBO potential differs from the Brenner potential in that it also allows for treatment of both the torsional and van der Waals interactions needed for accurate simulations of condensed-phase hydrocarbons. Full details on the AIREBO potential are available elsewhere [23]. Although all of the systems examined here are hydrocarbons, and all are treated with the same model, they were chosen to test the efficiency of the iterative timestep selection procedure under as wide a range of simulation conditions as possible. The integrator used in both the fixed- and variable-timestep simulations was the symplectic velocity Verlet algorithm [17].

The first system considered is a simulation of energetic particle bombardment, in which a 100-eV Ar atom strikes a 4961-atom surface of amorphous polyethylene, initially at a density of 0.854 g/cm^3 and a temperature of 298 K. The polymer is treated with periodic boundary conditions in the two directions parallel to the surface. The surface area is 1265 Å^2, and the slab has a depth of 35.6 Å. After the incident Ar atom strikes the surface, its initial kinetic energy of 100 eV is dissipated into thermal energy via a series of energetic collisions and sputtering events that persist for ~1 ps. The variable-timestep algorithm is expected to be considerably more efficient than fixed-timestep algorithms for this system, due to the nonequilibrium energy distribution in the system; small timesteps are needed during the initial collisions, increasing gradually to more standard timesteps for the thermalized system.

The second system consists of 128 molecules of ethylene at a liquid-like density of 1.40 g/cm^3 and an elevated temperature of 2500 K. This system is in a metastable equilibrium at this temperature and at sub-nanosecond timescales, punctuated by relatively infrequent collisions that result in strongly fluctuating covalent bonding forces. (On nanosecond timescales, these collisions, some of which result in reactions, will lead to polymerization [28]). This system was chosen as an example of a quasi-equilibrium system in which there is a need for a rather broad distribution of timesteps, due to the expected wide distribution of energy errors.

The final system considered comprises 128 molecules of liquid benzene at $0.873\,\mathrm{g/cm^3}$ and 298 K. This is an equilibrium simulation under mild conditions with no reason to expect a particularly heterogeneous distribution of timesteps or energy errors. This system was chosen to present the most challenging test for the variable-timestep algorithm, and it was not expected that variable timesteps would be substantially more efficient than standard fixed-timestep integrators for this system.

As mentioned above, the timestep in a variable-timestep integration method can be viewed as a function of phase space [6]. Every different state point $\mathbf{\Gamma}$ will result in a different choice of timestep, $\Delta t(\mathbf{\Gamma})$. In the current algorithm, for example, the timestep is chosen as (approximately) the largest value that will result in a scaled energy error less than some predefined value of D_E. The distribution of timesteps, $P(\Delta t)$, required to achieve this level of energy conservation in a given simulation thus gives some indication of the need for a variable-timestep integrator; if the distribution is narrowly peaked, with small tails, then a single, fixed timestep can be chosen with only moderate inefficiency. If the distribution is broad, and has a significant tail at low timestep values, a fixed-timestep method will be quite inefficient compared to an optimal variable-timestep method.

This distribution $P(\Delta t)$ can be obtained directly from a simulation using the variable-timestep integration method described above. Because of the iterative nature of the timestep selection, however, the timestep used will always underestimate the maximum permissible value. This underestimation can be minimized by using the algorithm with a very large α and a value for β that is very close to one. In this way, the chosen timestep can be guaranteed not to differ from the optimal value by more than a factor of $1-\beta$, and the distribution $P(\Delta t)$ can be obtained to arbitrary accuracy.

Figure 1 shows the distribution $P(\Delta t)$ that is obtained for each of the three test systems using an energy tolerance of $D_E = 0.01\,\mathrm{eV^2/ps}$. The values of α and β used were $2\,\mathrm{fs}/\Delta t$ and 0.99, respectively, so that the timestep was always within 1% of the maximum possible value. (Note that this results in an extremely inefficient algorithm, with many timesteps requiring several hundred force evaluations. This is done only to characterize the method; the parameters α and β will be optimized for efficiency below).

The distribution of timesteps at constant energy conservation in Fig. 1 indicates that the sputtered polyethylene system (top panel) requires a number of small timesteps of 0.2 fs and below to integrate the energetic collisions, but that more practical timesteps near 1 fs suffice for much of the simulation. All but one of the timesteps below 0.1 fs occurred within 1 ps of impact, during the initial collision cascade. The 2500 K ethylene system (middle panel) also displays a bimodal distribution, with timesteps of near 0.1 fs required for the more energetic collisions, and step sizes of ~0.8 fs for the majority of the simulation time (but a minority of the integration steps). The bimodal nature of this curve is most likely related to the nature of the AIREBO potential; the most energetic collisions involve a damping of the Lennard–Jones interaction at short distances, which results in quickly fluctuating forces, and thus necessitates a very small timestep [23]. This feature of the potential, rather than a specific

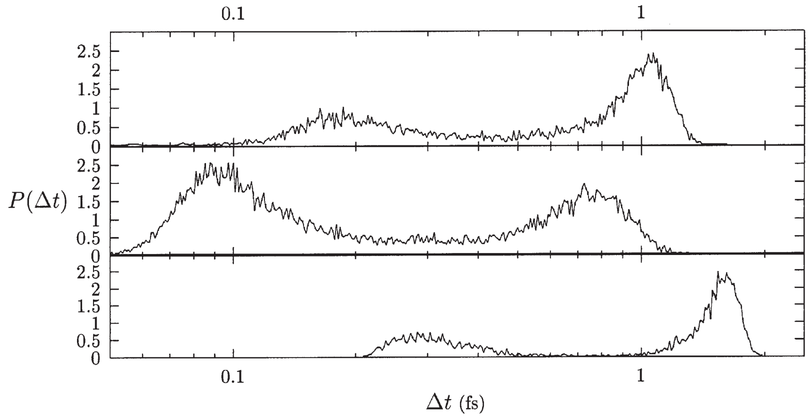

FIGURE 1 Distribution of timesteps required to maintain energy conservation at a level of $D_E = 0.01\,\mathrm{eV^2/ps}$ for three different systems. The upper panel is for the simulation of 100 eV particle bombardment of polyethylene, the middle panel is for the 2500 K ethylene system, and the lower panel is for room-temperature liquid benzene. All probability distributions are normalized to have an area of one.

characteristic of the dynamics, is most likely responsible for the bimodal nature of all three curves. Note that the small-Δt peak is largest for the hot ethylene system, where there are a substantial number of potentially reactive collisions throughout the simulation. The cooler benzene system (bottom panel) can use much larger timesteps than the other systems, as expected, with $\Delta t \approx 1.6 fs$ sufficient for most of the dynamics.

Common practice with the AIREBO potential is to use a fixed timestep of 0.25 or 0.5 fs at ambient temperatures to obtain reasonable energy conservation—quite a bit smaller than the ~1 fs that is standard for potentials without bond flexibility or reactivity. The usual explanation is that this is necessary to account for the high-frequency vibrations of the C–H bonds, but Fig. 1 clearly indicates that these small timesteps are necessary only for a rather small fraction of the trajectory. This suggests that perhaps variable timestep integrators can result in increased efficiency even for equilibrium simulations. Although the bimodal character of these curves is most likely specific to AIREBO and related potentials, the large-Δt peaks appear broad enough to indicate that some benefit could be obtained from a variable timestep for non-reactive potentials as well.

The breadth of the distributions in Fig. 1 indicates that variable-timestep integration methods may potentially be more efficient than fixed-timestep methods, but only if the timestep can be predicted accurately with few force evaluations. The iterative timestep selection algorithm employed here will be computationally efficient only if each timestep is sufficiently predictive of the timestep in the succeeding step—i.e. if the timesteps are well correlated in time. This is not guaranteed, but would seem to be a reasonable assumption: a collision or other event requiring a smaller than average timestep should typically last for several of the timesteps required to integrate it accurately. Measurement of the timestep autocorrelation time is complicated by the fact that the timesteps are not uniform. But informal investigation indicates that the correlation time is in fact quite short; only about 50% of the small timesteps (below 0.4 fs) in the sputtered polyethylene system are followed directly by another small timestep.

To evaluate the efficiency of the algorithm, we tested a range of α and β values for each of the three systems studied, with the goal of minimizing the average number of force evaluations needed per unit of simulation time. The optimum values differed somewhat for each system, with $(\alpha, \beta) = (1.1, 0.5)$ for the sputtering system, (1.1, 0.55) for the hot ethylene system, and (1.05, 0.55) for the room-temperature benzene system. Fortunately, however, consensus values of $\alpha = 1.1$ and $\beta = 0.5$ proved to be near-optimal for all systems (decreasing the efficiency by only 8% for the benzene system and 3% for the hot ethylene system), and these values were used in all of the simulations described below. We note that this pair of α and β values results in 12.1% of the iterations being rejected and repeated for each system, in agreement with Eq. (14). It is especially encouraging that the dependence of the optimal α and β on the simulated system is so slight, and that similar values can be used effectively for such a wide range of systems and thermodynamic state points. We thus expect that these values would prove reasonable for other carbon and hydrocarbon systems modeled with the AIREBO potential, although the optimal α and β values may well prove to be different for other (non-bond-order) classes of potentials.

A quantitative comparison of the computational efficiency of the fixed- and variable-timestep methods is displayed in Fig. 2. This plot demonstrates that the variable-timestep method is considerably more efficient than the fixed-timestep algorithm for the non-equilibrium sputtering simulation, and of comparable efficiency for the two equilibrium simulations. At a fixed level of accuracy, the sputtered system (left panel, Fig. 2) can be integrated with 50–55% fewer force evaluations with the variable-timestep integrator as are required using fixed-timestep integration. Alternatively, at a fixed computational cost, the variable-timestep integrator generates energy "diffusion" constants that are only 4–6% of those obtained with fixed timesteps. For the high-temperature simulation of reactive ethylene (middle panel, Fig. 2), the efficiency is nearly identical for both integrators. The average timestep is larger by 12% for the variable timestep runs, but this is almost exactly compensated by the need for additional force evaluations on 12% of the steps (cf. Eq. (14)). For the room-temperature simulation of liquid benzene (right panel, Fig. 2), the variable-timestep simulation is less efficient by approximately 5–10% at a fixed level of energy conservation, or delivers energy diffusion values approximately 10–60% larger at a fixed computational cost. We note that the 5–10% decrease in efficiency could (coincidentally) be counteracted almost exactly by using the truly optimal α and β values for this system. These simulations were run with the suggested consensus values, however, to ensure a fair comparison.

We conclude that the variable-timestep integration should provide similar improvements for other non-equilibrium systems, and comparable, or only marginally decreased efficiencies for homogenous, equilibrium systems. These conclusions are most certain for other systems treated with AIREBO or related potentials, of course, because all of the

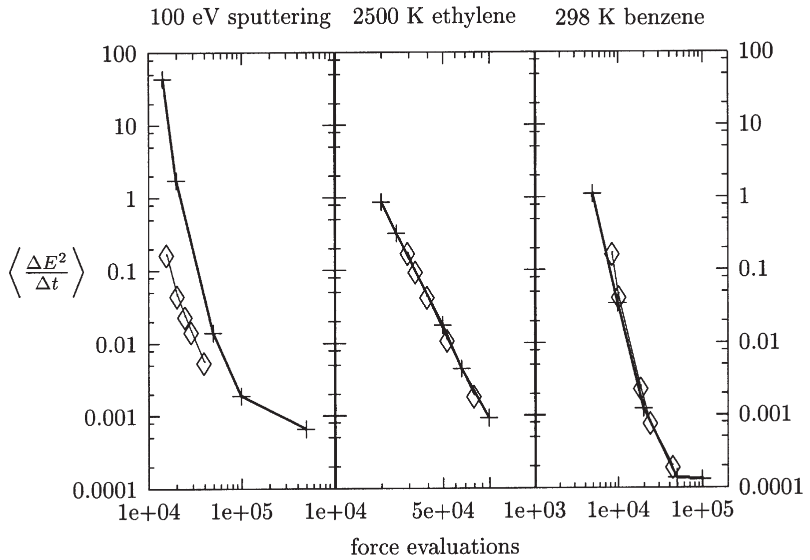

FIGURE 2 Computational efficiency of the variable-timestep algorithm (◇, thin line) and fixed-timestep algorithms (+, heavy line) for three different physical systems. The left panel is for the 100-eV bombardment of polyethylene, the middle panel is for the 2500 K ethylene system, and the right panel is for room-temperature benzene. Energy conservation, measured as the observed mean scaled energy error, $\langle \Delta E^2/\Delta t \rangle$, is plotted against computational cost, measured as the number of force evaluations needed per 5 ps of dynamics. Efficient simulations are thus represented by curves towards the left (reduced cost) and towards the bottom (increased accuracy).

systems considered here were modeled with this potential. It is unknown to what degree the current results depend on the bimodal character of the distributions in Fig. 1 (which are hypothesized to be a consequence of the adaptive treatment of the LJ repulsive terms), as opposed to the breadth of the individual peaks in the distributions, or the autocorrelations of the energy errors. Investigations with a wider variety of potentials would be interesting, but have not been pursued here.

DISCUSSION

An iterative method for choosing the timestep size in molecular dynamics simulations has been presented, in which integration steps are rejected and repeated with smaller timesteps if the resulting error $|\Delta E|$ in the total energy exceeds a scaled energy error tolerance, $(D_E \Delta t)^{1/2}$. This approach has been shown to be more efficient than fixed-timestep integration for hydrocarbons modeled with the AIREBO potential under non-equilibrium conditions. Somewhat surprisingly, the algorithm also demonstrates efficiency comparable to that of fixed-timestep integration in equilibrium simulations of condensed phase hydrocarbons. Another convenient result is that the single set of timestep scaling parameters $\alpha = 1.1$ and $\beta = 0.5$ can be used advantageously across a wide range of conditions. Consequently, this algorithm can be applied confidently to new hydrocarbon systems, without any need for extensive parameterization or testing.

Variable-timestep integration algorithms are much more commonly used in numerical analysis than in computational chemistry or physics[13]. Even in that field, however, variable-timestep integration has fallen out of favor to some degree recently, due to the discovery that fixed-timestep symplectic integrators can be much more efficient than either fixed- or variable-step nonsymplectic integrators [29,30], and that variable timesteps damage the performance of a symplectic integrator [29–32]. In this study, however, variable-timestep integration outperforms the fixed-timestep, symplectic velocity Verlet algorithm, despite the fact that the variable-timestep approach has destroyed the symplectic property of the integrator. The resolution of this apparent contradiction lies in the choice of metric for the accuracy of an integration method. Numerical analysts are primarily concerned with the total, integrated trajectory error. This is quite appropriate for few-body systems with simple potentials; celestial mechanicians are quite justifiably concerned with the total error in the position of the interacting bodies at the end of a long simulation of the solar system. Condensed-phase chemistry and physics applications, on the other hand, model many more particles, often use potentials that are substantially more complex, and are typically quite chaotic. While accuracy in the short-time dynamics may be important, the purpose of long-time molecular dynamics simulation is often simply a means of sampling phase space or trajectory space. In these circumstances, the total trajectory error is almost irrelevant, as the chaotic dynamics prevent numerical integration to within a specified tolerance in any case. But the stability of the trajectory and validity of the sampled ensemble are crucial. Thus it is not inconsistent that a non-symplectic, variable-timestep integrator is found to be more efficient and accurate in the current study of large-scale systems (using rms energy errors as the metric), but less so for simulations using gravitational potentials [29,30] or harmonic oscillators [32] (using total trajectory error as the metric). Another significant difference between the current method and variable-timestep integration methods commonly used in numerical analysis [33–35] is that this method only requires multiple force evaluations on a small fraction of the integration steps (those that do not satisfy the error criterion with $k = 0$ in Eq. (11)). For typical ODE applications, the system can not be assumed to be Hamiltonian, so multiple force evaluations are required at every step in order to obtain error estimates. The energy conservation guaranteed by strictly Hamiltonian, microcanonical dynamics provides a much simpler method to obtain energy errors in molecular dynamics simulations.

The algorithm presented here requires the specification of an energy "diffusion constant", D_E, in units of energy^2/time. This is a rather unorthodox means for describing energy conservation, but is in some ways quite natural; it is perhaps more instructive and appropriate to plan numerical simulations in terms of the acceptable level of energy error after a certain simulation time, rather than by specifying a fixed timestep (often obtained only after trial-and-error experience with a particular combination of potential and system). One consequence of this, however, is that it becomes more difficult to predict the amount of computer time required for a long simulation when the average timestep is not known in advance (without the same trial-and-error experience).

The method described here is meant to be used to generate a simulation with a specifiable amount of energy error, and the means for doing so were derived under the assumption of zero-mean, normally distributed single-step energy errors. These simplifying assumptions were necessary in order to obtain a method that could be used to make decisions on timestep size based only on the error from a single step. Ideally, timesteps could be rejected based on more complex criteria, over times much longer than Δt. This is impractical, of course, due to the overhead associated with calculating and rejecting long portions of the trajectory. In order for the single-step energy error to have predictive value at longer times, the noise need not be strictly Gaussian, but the assumption of a zero-mean distribution is crucial. If there is any net drift in the energy, the long-time energy errors will increase linearly with time, rather than as $\sqrt{t}$, and will exceed the specified tolerance. Luckily, this is not a problem in practice for continuous, conservative potentials integrated with symplectic integrators; we have observed no such difficulties in any of the systems we have examined.

As with any variable-timestep algorithm, the integration timestep depends on the system's location in phase space [6]. Consequently, the resulting integrator is not reversible. In addition, the iterative nature of this integrator gives it some history dependence; the choice of timestep depends not only on the system state, but also the previous timestep value, via Eq. (10). (This is essentially a course-resolution optimization, where the history dependence appears in the optimization grid). The timestep thus can not be viewed as a purely local function of phase space. In a practical sense, we not that this affects the reproducibility of restarted trajectories, and for this purpose the current value of the timestep should be viewed as part of the state of the computational system.

It is possible that the performance of the current method could be improved by imposing an approximate time symmetry on the variable timesteps [36,37]; such extensions have not been considered here.

Because this integration algorithm is designed to deliver a specified level of energy conservation, it is most useful for simulations in the microcanonical ensemble. Implementations for other ensembles are possible, however. For example, when using a stochastic thermostat to integrate in the canonical ensemble, one could be careful to integrate the thermostat force separately from the internal forces, and evaluate ΔE only for the portion of the dynamics that occurs before the perturbation due to the thermostat. Likewise, this approach can be used with other external forces or perturbations, as long as care is taken to integrate the perturbation separately. Typically, however, energy conservation is much less important in thermostatted simulations, and we expect that the primary use for this algorithm will be in microcanonical simulations, particularly those of long duration.

Acknowledgements

Financial support for this work by the Department of Defense, the Petroleum Research Fund, and the South Carolina Commission on Higher Education is gratefully acknowledged. The authors also thank Brad Dickson and Barbara Garrison for helpful discussions.

References

[1] Andersen, H.C. (1980) "Molecular dynamics simulations at constant pressure and/or temperature", *J. Chem. Phys.* **72**, 2384–2393.

[2] Hoover, W.G. (1985) "Canonical dynamics: equilibrium phase-space distributions", *Phys. Rev. A* **31**, 1695–1697.

[3] Martyna, G.J., Klein, M.L. and Tuckerman, M. (1992) "Nosé-Hoover chains: the canonical ensemble via continuous dynamics", *J. Chem. Phys.* **97**, 2635–2643.

[4] Tuckerman, M.E. and Martyna, G.J. (2000) "Understanding modern molecular dynamics: techniques and applications", *J. Phys. Chem. B* **104**, 159–178.

[5] Adelman, S.A. and Doll, J.D. (1976) "Generalized Langevin equation approach for atom/solid-surface scattering: general formulation for classical scattering off harmonic solids", *J. Chem. Phys.* **64**, 2375–2388.

[6] Stoffer, D.M. and Nipp, K. (1991) "Invariant curves for variable step size integrators", *BIT* **31**, 169–180.

[7] Heyes, D.M. (1982), *Inf. Q. Comput. Simul. Condens. Phases* **5**, 26.

[8] Whittle, M. (1984) "A variable timestep algorithm", *Inf. Q. Comput. Simul. Condens. Phases* **14**, 47–49.

[9] Gibson, K.D. and Scheraga, H. (1990) "Variable step molecular dynamics: an exploratory technique for peptides with fixed geometry", *J. Comput. Chem.* **11**, 468–486.

[10] Beardmore, K.M. and Grønbech-Jensen, N. (1999) "An efficient molecular dynamics scheme for predicting dopant implant profiles in semiconductors", *Nucl. Instr. Meth. Phys. Res. B* **153**, 391–397.

[11] Garrison, B.J. (2001) "Molecular dynamics simulations, the theoretical partner to static SIMS experiments", In: Vickerman, J.C. and Briggs, D., eds, *ToF-SIMS: Surface Analysis by Mass Spectrometry* (Surface Spectra and IM Publications, Manchester, UK), pp 223–257.

[12] Henkelman, G. and Jónsson, H. (2001) "Long time scale kinetic Monte Carlo simulations without lattice approximation and predefined event table", *J. Chem. Phys.* **115**, 9657–9666.

[13] Gear, C.W. (1971) *Numerical Initial Value Problems in Ordinary Differential Equations* (Prentice-Hall, Englewood Cliffs, NJ).

[14] Barth, E., Leimkuhler, B. and Reich, S. (1999) "A time-reversible variable-stepsize integrator for constrained dynamics", *SIAM J. Sci. Comput.* **21**, 1027–1044.

[15] Holder, T., Leimkuhler, B. and Reich, S. (2001) "Explicit variable step-size and time-reversible integration", *Appl. Numer. Math.* **39**, 367–377.

[16] Tuckerman, M., Berne, B.J. and Martyna, G.J. (1992) "Reversible multiple time scale molecular dynamics", *J. Chem. Phys.* **97**, 1990–2001.

[17] Swope, W.C., Andersen, H.C., Berens, P.H. and Wilson, K.R. (1982) "A computer simulation method for the calculation of equilibrium constants for the formation of physical clusters of molecules: application to small water clusters", *J. Chem. Phys.* **76**, 637–649.

[18] Huang, W. and Leimkuhler, B. (1997) "The adaptive Verlet method", *SIAM J. Sci. Comput.* **18**, 239–256.

[19] Toxvaerd, S. (1982) "A new algorithm for molecular dynamics calculations", *J. Comput. Phys.* **47**, 444–451.

[20] Heyes, D.M. and Singer, K. (1982) "A very accurate molecular dynamics algorithm", *Inf. Q. Comput. Simul. Condens. Phases* **6**, 11–23.

[21] Hoover, W.G., Kum, O. and Owens, N.E. (1995) "Accurate symplectic integrators via random sampling", *J. Chem. Phys.* **103**, 1530–1532.

[22] Gray, S.K., Noid, D.W. and Sumpter, B.G. (1994) "Symplectic integrators for large scale molecular dynamics simulations: a comparison of several explicit methods", *J. Chem. Phys.* **101**, 4062–4072.

[23] Stuart, S.J., Tutein, A.B. and Harrison, J.A. (2000) "A reactive potential for hydrocarbons with intermolecular interactions", *J. Chem. Phys.* **112**, 6472–6486.

[24] Tersoff, J. (1986) "New empirical model for the structural properties of silicon", *Phys. Rev. Lett.* **56**, 632–635.

[25] Tersoff, J. (1989) "Modeling solid-state chemistry: interatomic potentials for multicomponent systems", *Phys. Rev. B* **39**, 5566–5568.

[26] Brenner, D.W. (1990) "Empirical potential for hydrocarbons for use in simulating the chemical vapor deposition of diamond films", *Phys. Rev. B* **42**, 9458–9471.

[27] Brenner, D.W., Shenderova, O.A., Harrison, J.A., Stuart, S.J., Ni, B. and Sinnott, S. (2002) "A second-generation reactive empirical bond order (REBO) potential energy expression for hydrocarbons", *J. Phys.: Condens. Matter* **14**, 783–802.

[28] Stuart, S.J., Dickson, B.M., Sumpter, B.J. and Noid, D.W. (2001) "Computational study of polymerization in carbon nanotubes", *Mat. Res. Soc. Symp. Proc.* **651**, T1.8.1–T1.8.6.

[29] Gladman, B., Duncan, M. and Candy, J. (1991) "Symplectic integrators for long-term integrations in celestial mechanics", *Celest. Mech. Dyn. Astron.* **52**, 221–240.

[30] Calvo, M.P. and Sanz-Serna, J.M. (1993) "The development of variable-step symplectic integrators, with application to the two-body problem", *SIAM J. Sci. Comput.* **14**, 936–952.

[31] Skeel, R.D. and Gear, C.W. (1992) "Does variable step size ruin a symplectic integrator?", *Physica D* **60**, 311–313.

[32] Skeel, R.D. (1993) "Variable step size destabilizes the Störmer/leapfrog/Verlet method", *BIT* **33**, 172–175.

[33] Ceschino, F. and Kuntzmann, J. (1966) *Numerical Solutions of IVP* (Prentice Hall, Englewood Cliffs, NJ).

[34] Dormand, J.R. and Prince, P.J. (1980) "A family of embedded Runge–Kutta formulae", *J. Comp. Appl. Math.* **6**, 19–26.

[35] Hairer, E., Norsett, S.P. and Wanner, G. (1987) *Solving Ordinary Differential Equations I* (Springer, Berlin).

[36] Hut, P., Makino, J. and McMillan, S. (1995) "Building a better leapfrog", *Astrophys. J.* **443**, L93–L96.

[37] Stoffer, D. (1995) "Variable steps for reversible integration methods", *Computing* **55**, 1–22.