MOLECULAR PHYSICS, 1996, VOL. 89, NO. 5, 1265–1276

Generalized valence bond molecular dynamics at constant temperature

By DOUGLAS A. GIBSON and EMILY A. CARTER
Department of Chemistry and Biochemistry, University of California, Los Angeles, 405 Hilgard Avenue, Los Angeles, CA, 90095-1569, USA

(*Received 7 December 1995; accepted 28 May 1996*)

The Nosé–Hoover chain thermostat of Martyna, Klein and Tuckerman has been adapted to work with our generalized valence bond *ab initio* molecular dynamics (AIMD) method, allowing us to perform AIMD in the canonical ensemble (constant number of particles, temperature, and volume). The canonical ensemble is ideal for simulated annealing, which can be used to find global potential energy minima. The efficacy of constant-temperature-based AIMD simulated annealing is demonstrated for a Si_6 cluster and is compared with that of a non-constant-temperature-based method. It is found that in comparison with the velocity-scaling approach to annealing, constant-temperature-based annealing gives better control over the annealing process and, at least in the case under study, results in attainment of the preferred geometry of the system more efficiently.

1. Introduction

In a standard molecular dynamics (MD) simulation, trajectories in phase space are generated by numerically integrating Newton's or Hamilton's classical equations of motion. The forces on the nuclei in the system are obtained by differentiating a potential energy function which should be an adequate approximation to the true potential energy of the system. For many systems of interest, however, not even the functional form of the true potential energy is known; in such cases a somewhat different approach is required. One tool for obtaining the derivative of the potential energy of such a system at an arbitrary geometry is an *ab initio* electronic structure calculation. We refer to an MD simulation which takes its forces from such a calculation as an *ab initio* molecular dynamics (AIMD) simulation. The main drawback of this method is that it is computationally very expensive. Nevertheless, since AIMD requires only that the forces on the nuclei be computable at any given point and not that the entire potential energy surface be known, it may be the only feasible way of performing molecular dynamics with reasonably accurate forces for some systems.

A standard MD simulation, as described above, treats its system in the microcanonical (NVE) ensemble. It is often desirable to work in other ensembles, particularly the canonical (NVT) ensemble. One way to achieve this goal is to add an additional dynamic degree of freedom which corresponds to a heat bath. This type of extended system method was formulated originally by Nosé [1, 2]. In that form, the bath degree of freedom corresponds to a time scaling factor between a set of real variables and a set of virtual variables. It has been proved that the physical portion of the extended system reflects the canonical distribution so long as the system is ergodic. Nosé's equations of motion have been transformed by Hoover into a form in which the propagation of the system can be done entirely in terms of real variables without any

0026–8976/96 $12·00

need for a set of virtual variables; the resulting simplified set of equations of motion is known as the Nosé–Hoover thermostat [3].

However, Nosé–Hoover dynamics are not ergodic (and thus do not generate a canonical distribution) for small systems, which are precisely those which we are able to simulate currently with AIMD methods. The Nosé–Hoover equations of motion have been modified in a straightforward manner by Martyna, Klein and Tuckerman to address this very problem [4]. Their approach is to attach a 'chain' of bath degrees of freedom to the system instead of a single bath 'particle'; in this chain, one bath particle is coupled directly to the physical system and every other bath particle is connected only to the degrees of freedom immediately before and after it in the chain. They find that even for very small systems, such as a single harmonic oscillator, which are handled very poorly by a single Nosé–Hoover thermostat, the canonical distribution is reached with the addition of a small number of additional bath particles. Essentially, these extra bath particles act to thermostat the original thermostat in a hierarchical fashion.

We have adapted the Nosé–Hoover chain thermostat to work with our generalized valence bond (GVB) [5] AIMD method. The implementation of GVB-AIMD has been described in previous work [6, 7]. Our current implementation of GVB-AIMD uses exact forces at each time step (so-called Born–Oppenheimer dynamics) and no fictitious wavefunction dynamics [7]. For those interested in implementation of a Car–Parrinello (CP) type GVB-MD, note that references [6] and [7] omitted an orthogonality constraint term in the nuclear equations of motion ([6], equation (2) and [7], equation (1)) that should be present. We are currently revisiting the efficacy of the CP-GVB-MD method when this term is accounted for [8]; thus far the conclusions of [7] remain unchanged. We will show in the next section why it is desirable to perform AIMD at constant temperature in the context of simulated annealing.

2. Methods

Simulated annealing is one method for finding the global energy minimum of a system. It is most useful for systems for which chemical intuition regarding the interatomic interactions is lacking, such as metal and semiconductor clusters. The basic idea of simulated annealing is to allow the system to propagate for some time at a relatively high temperature and then slowly cool the system until it ceases to change significantly. If the initial temperature is sufficiently large (enough to overcome any relevant barriers on the potential energy surface) and the cooling of the system is sufficiently slow, this approach will reach the global energy minimum [9]. In practice, a number of trajectories should be run to ensure that the actual cooling schedule is sufficiently slow.

To implement such a scheme in a molecular dynamics simulation, we must first decide exactly what it is we mean by cooling the system. Remember that in the simplest form of MD, temperature is not an explicit parameter and thus cannot be controlled directly. Instead, we must settle for reducing the kinetic energy periodically by multiplying the nuclear velocities by a scaling factor. This approach does work, but it has certain disadvantages. The kinetic energy can vary widely over the course of a relatively small amount of simulation time, making the exact amount of energy lost in any given cooling step essentially random, since it is proportional to the instantaneous kinetic energy. Furthermore, as the nuclei move to parts of the potential energy surface which are lower in energy, they gain a corresponding amount of kinetic energy. The

kinetic energy gained in this manner may be large compared with that possessed by the system before going down such a slope, resulting in a 'runaway' system which has a temperature much higher than is desired. This runaway behaviour is actually quite common at the beginning of a trajectory if a random starting geometry is chosen because, in the region of interest (bound geometries), potential energy surfaces generally tend to be steep over a larger area than they are flat. Both of these features of the velocity scaling approach to simulated annealing make controlling the annealing process difficult.

A temperature-based approach to simulated annealing, however, does not suffer from these defects. If we can treat temperature explicitly, with the Nosé–Hoover chain thermostat, we can control the temperature directly instead of indirectly as was the case with velocity scaling. This direct control of the system's temperature permits us to perform uniform cooling steps; instead of getting an unpredictable change in the kinetic energy, we simply change the temperature which appears in the equations of motion to a new value which has been changed in exactly the desired manner. The Nosé–Hoover chain thermostat also prevents rapid accumulation of kinetic energy by forcing the system to remain relatively near the temperature which has been set. The Nosé–Hoover chain thermostat has been used also by Tobias and co-workers to equilibrate a protein, for which it has been found that NVT based dynamics are preferable [10]. We thus expect that annealing with the Nosé–Hoover chain thermostat should have much better performance than can be obtained with an annealing scheme based on scaling velocities. This constant temperature annealing method may be thought of as corresponding physically to a seeded gas expansion through a nozzle. In this paradigm, the overall temperature reduction over time represents the expansion process, while the instantaneous effect of the bath is analogous to collisions with the inert gas atoms in which the system of interest is seeded.

The Nosé–Hoover thermostat adds a term to one of the equations of motion of the physical system; instead of

$$\dot{p}_i = -\frac{\partial V(\mathbf{q})}{\partial q_i} \tag{1}$$

we have

$$\dot{p}_i = -\frac{\partial V(\mathbf{q})}{\partial q_i} - p_i \frac{p_\eta}{Q_1}, \tag{2}$$

where $\{q_i\}$ are the positions of the real particles, $\{p_i\}$ are the momenta of the real particles, V is the potential energy of the real system, p_η is the momentum of the bath particle, and Q_1 is the mass of the bath particle. This second term in equation (2) couples the physical system to the bath. The bath particle itself is subject to an equation of motion which is simply a function of the kinetic energy of the system and the temperature:

$$\dot{p}_\eta = \sum_{i=1}^{N} \frac{p_i^2}{m_i} - NkT, \tag{3}$$

where N is the number of degrees of freedom (normally three times the number of particles) in the real system, k is Boltzmann's constant, and T is the temperature of the system.

The extension of the Nosé–Hoover thermostat to the Nosé–Hoover chain thermostat is straightforward; it involves simply adding a term to the first bath particle

which contains the momentum of the second and so on, resulting in the following equation for the first bath particle:

$$\dot{p}_{\eta_1} = \left[\sum_{i=1}^{N} \frac{p_i^2}{m_i} - NkT\right] - p_{\eta_1}\frac{p_{\eta_2}}{Q_2}. \tag{4}$$

The velocity Verlet integrator [11] has been used to integrate the equations of motion for the extended system. It has the advantage that the velocities are explicit and do not have to be derived numerically as needed. This advantage leads to a drawback in the Nosé–Hoover chain implementation, however, because the equation of motion for the momentum contains that same momentum, and because of this we must iterate to a consistent solution at every time step. A simple iterative scheme, as hinted at by Martyna, Klein and Tuckerman [4], was tried initially and found to work for simple systems but diverged for the complicated potential energy surface of a semiconductor cluster. It was therefore necessary to implement a different iterative scheme. We have rewritten the equations of motions in the form *expression* = 0; for example, equation (2) becomes

$$\dot{p}_i + \frac{\partial V(\mathbf{q})}{\partial q_i} + p_i\frac{p_{\eta}}{Q_1} = 0. \tag{5}$$

We then minimized the sum of the squares of the various expressions thus derived, since this sum is always non-negative and goes to zero when the equations of motion are satisfied exactly. An adaptation of the direct inversion in the iterative subspace (DIIS) [12] method was used to accomplish this minimization in an efficient and robust manner. We test for convergence by taking the sum of the squares of the differences in velocities from one iteration to the next and comparing this sum to a specified convergence criterion.

The Nosé–Hoover and Nosé–Hoover chain thermostats have been applied to AIMD simulations before, but in a different context. In AIMD using forces derived from density functional theory, several groups interested primarily in controlling the electronic adiabaticity problem have applied such thermostats to the electronic and ionic degrees of freedom, but not for studying annealing [13–17]. In this work, we will be concerned entirely with the application of the Nosé–Hoover chain thermostat to the nuclear degrees of freedom for the purpose of annealing.

3. Results

We have performed a number of test simulations on a Si_6 cluster, varying the time step, Nosé–Hoover particle mass, and the convergence criterion for the velocities in the iterative scheme used to obtain them. The time step was varied from 10 atomic time units (0•24 fs) to 50 atu (1•2 fs). The Nosé–Hoover particle masses were chosen such that the first particle, which interacts directly with the physical system, has a larger mass than the others, which do not; those masses ranged from 10 to 500 and 2 to 100, respectively, all in atomic units. The exact choice of fictitious mass is not critical; it has been demonstrated that variations of up to 4 orders of magnitude do not change the ability to generate canonical distributions [4, 10]. We have chosen our masses smaller than prescribed in [4] (10 600 au and 590 au, respectively) for our starting temperature, since we wish to demonstrate performance in the context of an annealing simulation where the temperature at the end will be much lower, and the mass is prescribed to be directly proportional to temperature. It may be preferable in the future, however, to

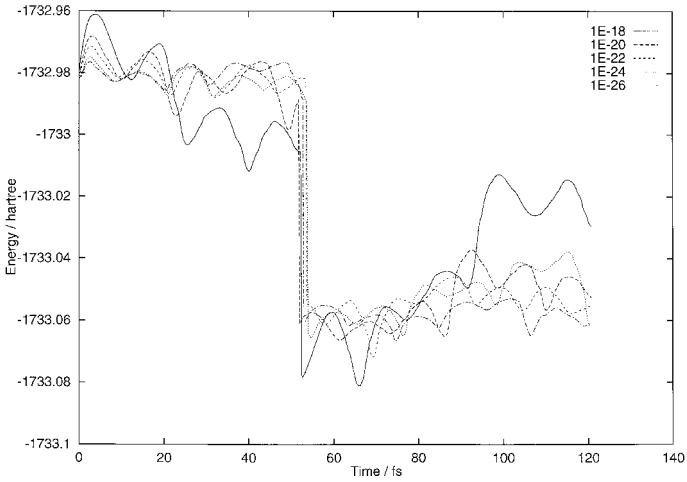

Figure 1. The conserved quantity (similar to the total energy in a constant energy MD simulation, with additional terms) in these Si_6 simulations as a function of time for a range of Nosé–Hoover propagation velocity convergence criterion values (listed in the key). Four bath particles are used with masses (in atomic units) of 500 for the first bath particle and 100 for each of the other three, and the time step is 10 atu. The discontinuity at approximately 50 fs is due to a switch between electronic states.

use larger masses for the high temperature portion of an annealing run and smaller masses for the low temperature portion, since the larger masses will allow the use of a longer time step. In all cases, four Nosé–Hoover particles were used. The velocity convergence criterion was varied from 10^{-16} to 10^{-26}. In all cases except for that shown in figure 3 the temperature was set initially at 1000 K. The electronic wavefunction used was a generalized valence bond with perfect pairing (GVB-PP) [5] wavefunction with 12 pairs of electrons in 24 orbitals. The electronic wavefunction was converged at every time step to ensure adherence to the Born–Oppenheimer potential energy surface. We also used a multiple time step approach [18, 19] to cut down on the number of computations of the forces on the nuclei, since the natural time scale of the nuclear motion is longer than that of the bath particles due to the mass disparity.

Conservation of energy is a common measure of the accuracy of an MD simulation. In Nosé–Hoover chain constant temperature MD, energy is not conserved, yet there exists an analogous quantity which is conserved:

$$H' = H + \sum_{i=1}^{M} \frac{p_{\eta_i}^2}{2Q_i} + NkT\eta_1 + \sum_{i=2}^{M} kT\eta_i, \tag{6}$$

where H is the Hamiltonian for the real system alone, H' is the Hamiltonian for the extended system, M is the number of bath particles, and all other symbols are as before. This conserved quantity is plotted versus time in figure 1 for a series of trajectories which differ only in the level of convergence of the velocities, which ranges from 10^{-18} to 10^{-26}. In this series, the time step is set at 10 atu and the Nosé–Hoover particle masses are 500 for the first bath particle and 100 for each of the others. Results for a convergence criterion of 10^{-16} are not shown; it is sufficient to note that they are of significantly lower quality than those obtained with a convergence criterion of 10^{-18}. The large drop in the middle of each curve is due to the electronic wavefunction converging to a new state and has nothing to do with the propagation itself. From these curves, it appears that there is only marginal improvement when the convergence

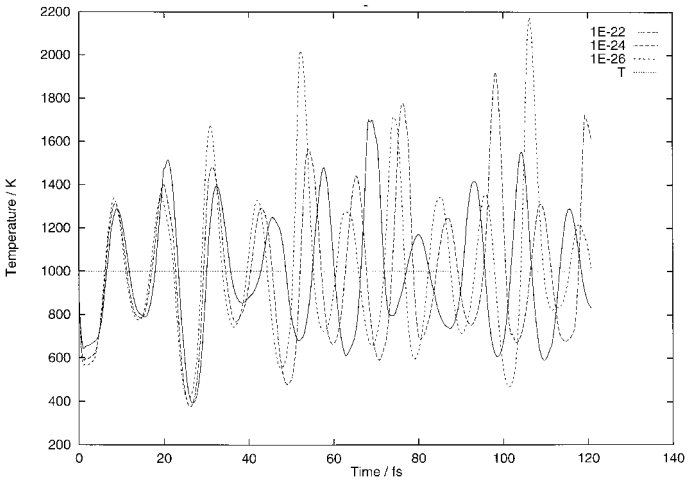

Figure 2. Instantaneous temperature (kinetic energy in units of degrees Kelvin) for the three trajectories with the smallest convergence criteria from figure 1. The line labelled 'T' is the T which appears in the equations of motion.

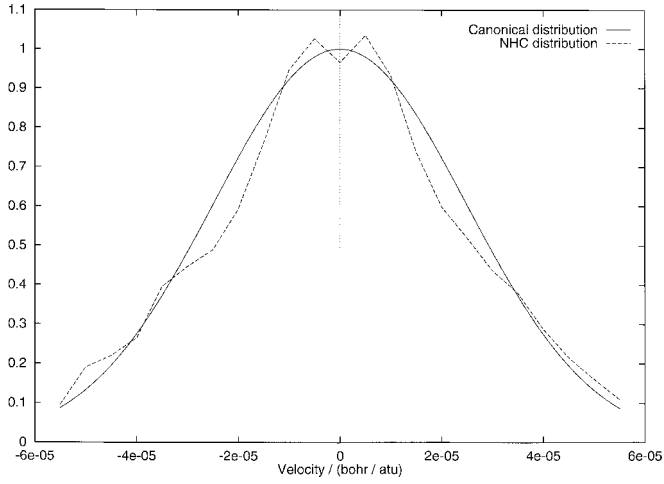

Figure 3. Velocity distribution for a 10 K trajectory with bath particle masses of 500 and 100. The solid line is the exact canonical distribution, and the dashed line is the velocity distribution observed in the simulation.

criterion is set smaller than about 10^{-22}. In figure 2, we show the instantaneous 'temperature' (kinetic energy) of the physical system as time passes for three smallest convergence criteria tested. We see that in all three cases the distribution is qualitatively correct. We note that smaller convergence criteria have a slightly shorter (and presumably more accurate) period of oscillation and have fewer points where the curve peaks very sharply, leading us to believe that it is in fact worth the slightly increased cost to converge to a tolerance of 10^{-26}.

In order to check further on the quality of the NHC thermostat, we calculated a velocity distribution for all the atoms in our Si_6 cluster. It is impractical to obtain a velocity distribution with good statistics at 1000 K; the rapid rearrangement which occurs at this temperature requires the sampling of a very large amount of phase space,

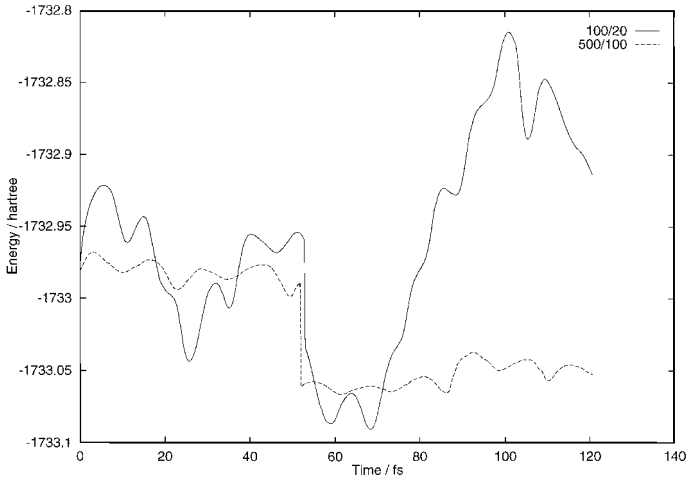

Figure 4. Conserved quantity for two Si_6 simulations with a time step of 10 atu and Nosé–Hoover convergence criterion of 10^{-20}; one uses masses of 500 and 100 for the bath particles while the other uses masses of 100 and 20. Conservation is significantly better for the trajectory with the heavier bath particles. As before, the discontinuity is due to a sudden change in the electronic state.

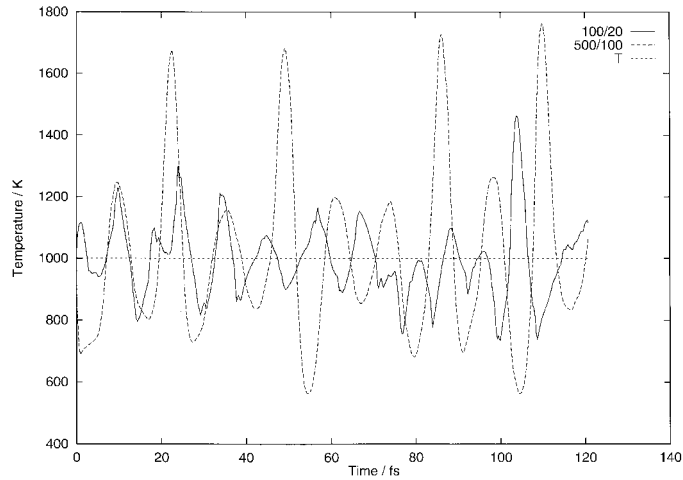

Figure 5. Instantaneous temperature along the trajectories from figure 4. The relative jaggedness of the curve for the trajectory with lighter bath particles probably indicates that the integration is not very good.

which means that the simulation time must be very long. We have instead obtained a velocity distribution at 10 K, which is shown in figure 3. Again, bath particle masses of 500 au for the first bath particle and 100 au for each of the others was used. We find that the velocity distribution is quite close to the canonical distribution.

In figure 4, we have plotted the conserved quantity versus time for two simulations using different Nosé–Hoover particle masses; one uses 100 for the first and 20 for the others, while the other uses 500 for the first and 100 for the others. In both cases, the time step is 10 atu and the velocity convergence criterion is 10^{-20}. We see that the trajectory with larger masses deviates much less than does the trajectory with smaller masses. Figure 5 shows the nuclear kinetic energy for those same trajectories. The

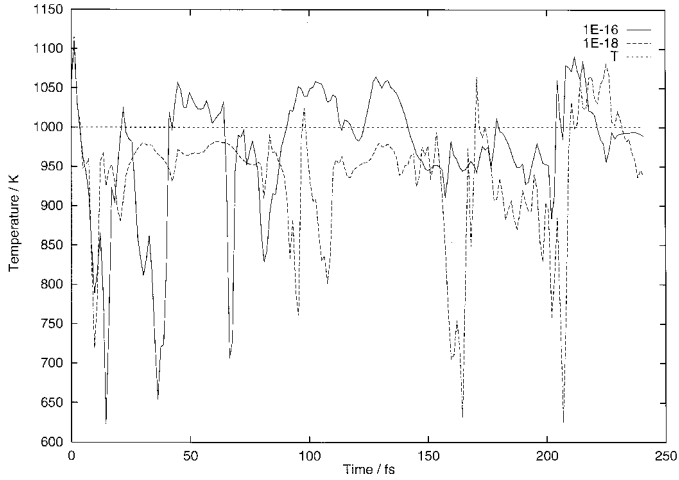

Figure 6. Instantaneous temperature for two trajectories with a time step of 50 atu, and bath particle masses of 100 for the first and 20 for each of the other three; the integration convergence criterion is given in the key. Note the excessive jaggedness of these curves, which shows that the numerical integration is quite poor, as well as the poor distribution of energies (which should be centred around the dotted line labelled 'T' at 1000 K).

amplitude of the energy fluctuations is smaller for the smaller Nosé–Hoover particle masses, but the relative jaggedness of the kinetic energy curve in that simulation leads us to believe that the integration is starting to break down noticeably for this combination of time step and masses. It is possible to use the smaller masses with a shorter time step, but this would increase the cost of the simulation significantly. It is also possible to use larger masses, and choosing the masses as prescribed in [4] for this system at 1000 K (as above, 10 600 au and 590 au) results in an essentially identical kinetic energy curve to that observed with masses of 500 au and 100 au, respectively.

We have also performed a set of simulations with different time steps. In general, the time step of 10 atu was the only one which performed well at all. As an example, see figure 6, where a time step of 50 atu with Nosé–Hoover masses of 100 and 20 results in an energy distribution which is not only very jagged (suggesting poor integration) but is not well centred around the target temperature of 1000 K.

After coming up with a set of parameters which gave reasonably good results, we performed a number of simulations to compare the velocity scaling approach to simulated annealing to a temperature-based approach. Initial conditions were chosen randomly, with only the constraint that the starting total energy must be below that of six isolated Si atoms. In order to simulate the system for any appreciable length of time, it was necessary to use a multiple time step approach to do the integration. We used a short time step of 10 atu for the Nosé–Hoover particles and a long time step of 100 atu for the nuclei. The Nosé–Hoover force computation is performed every short time step, proceeding to the halfway point within the long time step before computing the physical forces on the nuclei, after which the Nosé–Hoover forces are again computed every short time step until the end of the long time step has been reached; this approach is based on the reversible RESPA (r-RESPA) method of Tuckerman, Berne and Martyna [17], and is similar to our r-RESPA implementation for constant energy AIMD [18]. Nosé–Hoover particle masses of 500 for the first particle and 100

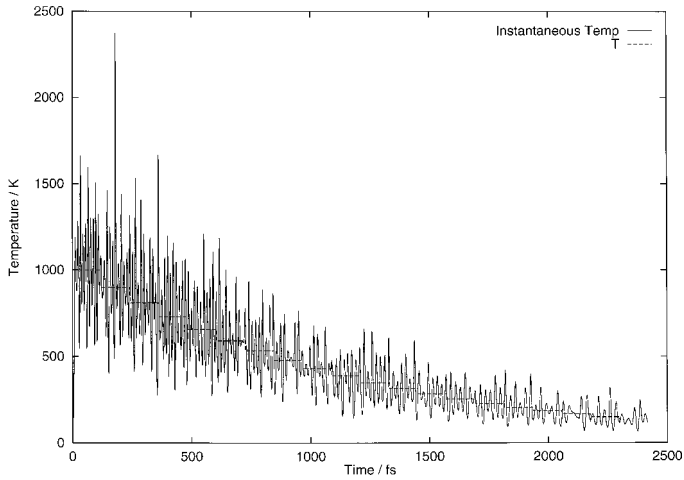

Figure 7. Instantaneous 'temperature' of an Si_6 constant-temperature-based annealing trajectory. The dotted line is the target temperature which appears in the equations of motion.

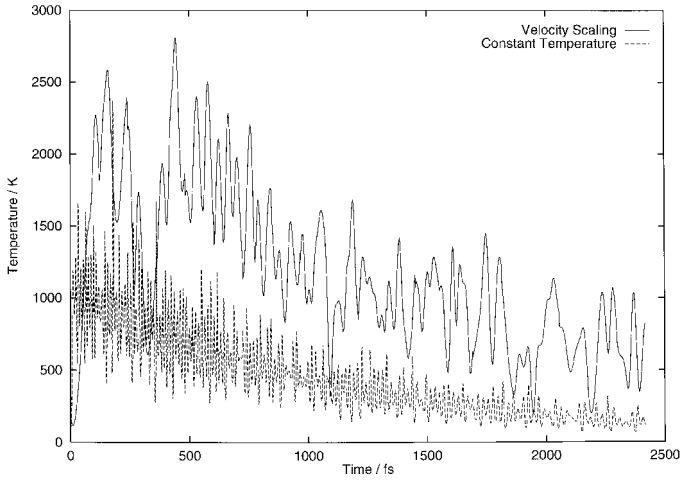

Figure 8. Instantaneous 'temperatures' of two Si_6 annealing trajectories, one with temperature-based annealing and one with velocity-based annealing.

for the other three were used. Every 120 fs, the system was cooled by multiplying the velocities by $0{\cdot}9^{1/2}$ in the velocity scaling simulations or by multiplying the temperature by 0·9 in the temperature-based simulations; these cooling steps are as equivalent as can be performed for the two types of simulation. The amount of CPU time required to perform this multiple time step simulation is 20% greater than if a single time step (for both the bath and real particles) of 100 atu is used.

Figure 7 shows the kinetic energy (instantaneous temperature) versus time for one Nosé–Hoover chain simulation. We see that even as the temperature is changed the physical system always observes a kinetic energy distribution which is centred around the temperature, and the width of the distribution is approximately proportional to the temperature. In figure 8 we compare this energy distribution to that of a velocity

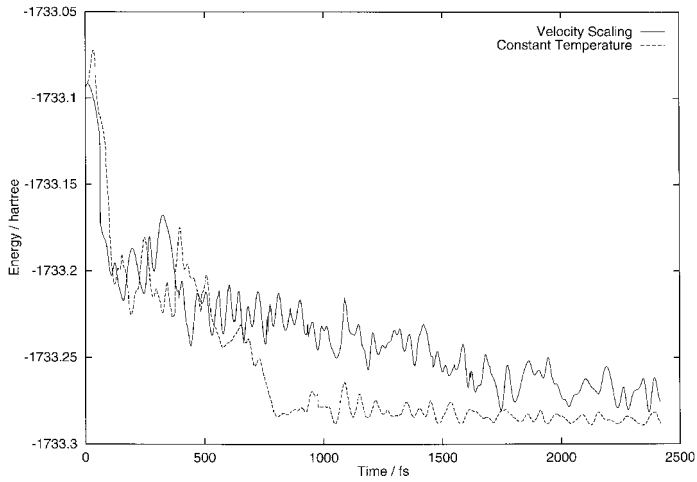

Figure 9. Potential energy of the trajectories from figure 8. Discontinuities observed are due to sudden changes in the electronic wavefunction.

scaling simulation starting from the same initial conditions. Note that in the velocity scaling trajectory a large amount of kinetic energy is accumulated early in the simulation. This energy buildup will often occur when starting from a randomly selected initial geometry. The accumulated kinetic energy takes a long time to be bled off in cooling steps, such that even after a picosecond has passed it has more energy than we desired it to have at the beginning. Note also the large variations in kinetic energy on relatively short time scales: these variations make the exact amount of kinetic energy removed at any given cooling step essentially random. Both of these factors make velocity scaling simulated annealing difficult to control, while neither of these effects is observed in the Nosé–Hoover chain simulation.

Figure 9 plots the potential energy of these two simulations against time. We see that the energy buildup in the velocity scaling simulation is not the result of a larger drop in potential energy than occurred in the Nosé–Hoover chain simulation, but rather is an undesirable effect which the Nosé–Hoover chain prevents. For the first 700 fs or so the potential energies of the two simulations are comparable. After that point, however, the Nosé–Hoover chain simulation drops to a lower part of the potential energy surface. The velocity scaling simulation will eventually reach the same region of the potential energy surface, but a couple of picoseconds of simulation time and numerous CPU hours later; inadequate control over the annealing process can be costly.

In figure 10 we show a comparison of the conserved quantity of this annealing simulation (which uses time steps of 10 and 100 atu for the bath and nuclei, respectively) with another which uses a single time step of 25 atu for both the nuclei and the bath particles. The periodic discontinuities observed in the conserved quantity are due to cooling, not inaccuracy in the dynamics. Though the conservation of this quantity is not perfect in the multiple time step simulation, it is much better than is observed for the single time step. Furthermore, the simulation which uses a single time step of 25 atu requires over three times as much CPU time to integrate for the same amount of simulation time than is necessary with two time steps of 10 atu and 100 atu.

Raghavachari [20] observed an edge-capped trigonal bipyramid to be the lowest

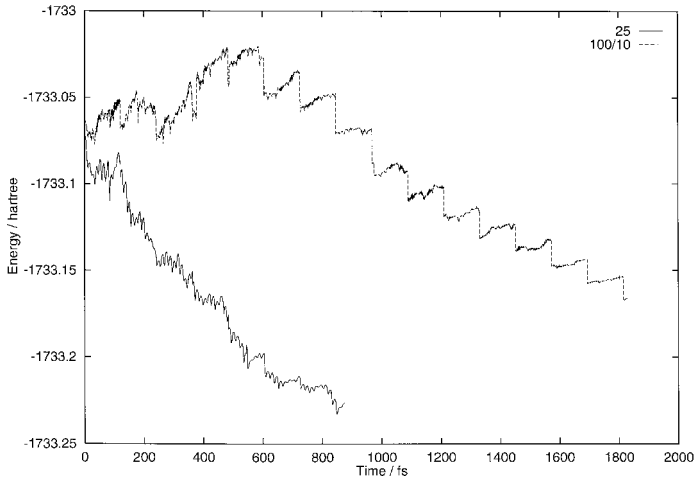

Figure 10. Conserved energy quantity for two annealing trajectories, one with a single time step of 25 atu and one with a time step of 10 atu for the bath particle and 100 atu for the nuclei. Discontinuous drops in energy here are due to periodic temperature reductions.

energy structure at the Hartree–Fock level of theory. Fournier, Sinnott, and DePristo find that the edge-capped trigonal bipyramid and the face-capped trigonal bipyramid are nearly degenerate in a density functional calculation, with the face-capped structure being slightly lower in energy [21]. The structure which we have obtained has approximately C_2 symmetry; it is thought of most easily as a distorted octahedron where the equatorial atoms are buckled both in and out of the equatorial plane. This structure and others will be discussed in more detail in a forthcoming publication; the purpose here is simply to point out that AIMD annealing has led to the prediction of a novel structure for Si_6.

4. Conclusion

We have adapted the Nosé–Hoover chain thermostat to work with our wave-function-based AIMD methods, including a quick and robust iterative scheme for the velocities. Though the small masses (relative to the rest of the system) of the bath particles require a shorter time step to be used in the integration of the equations of motion, the use of a multiple time step approach alleviates this problem dramatically; we find that only $\sim 20\%$ more CPU time is required for a simulation with time steps of 10 atu for the bath and 100 atu for the nuclei than is needed for the same simulation with a single time step of 100 atu, yet the quality of the results is improved greatly for the small added expense.

We also compare the results obtained from a temperature-based annealing simulation with those obtained from a velocity scaling annealing simulation. We find that while the velocity scaling simulation spends a large amount of simulation time at higher energies than are desired, the temperature-based simulation cools in a steady, controlled fashion. We not only have better control over the annealing process when the temperature can be controlled explicitly, but we find also that, at least in the particular simulation set we have performed, the system also reaches its preferred geometry faster. We attribute this to be a result of better control over the temperature of the system, since the velocity scaling simulation spends much more time at higher temperatures and, as a result, will escape from the region of the preferred geometry

even if it finds itself in that region early in the simulation. We therefore conclude that for optimum performance in annealing, a temperature-based annealing approach should be preferred over a velocity scaling approach.

This work was supported by the Office of Naval Research. E.A.C. also acknowledges support from the Camille and Henry Dreyfus Foundation and the Alfred P. Sloan Foundation through their Teacher-Scholar and Research Fellow Awards. D.A.G. thanks the National Science Foundation for a predoctoral fellowship.

References

[1] Nosé, S., 1984, *Molec. Phys.*, **52**, 255.
[2] Nosé, S., 1984, *J. chem. Phys.*, **81**, 511.
[3] Hoover, W. G., 1985, *Phys. Rev.* A, **31**, 1695.
[4] Martyna, G. J., Klein, M. L., and Tuckerman, M., 1992, *J. chem. Phys.*, **97**, 2635.
[5] Bobrowicz, F. W., and Goddard, W. A. III, 1977, *Methods of Electronic Structure Theory* (New York: Plenum), p. 79.
[6] Hartke, B., and Carter, E. A., 1992, *J. chem. Phys.*, **97**, 6569.
[7] Gibson, D. A., Ionova, I. V., and Carter, E. A., 1995, *Chem. Phys. Lett.*, **240**, 261.
[8] Gibson, D. A., and Carter, E. A., unpublished.
[9] Kirkpatrick, S., Gelatt, C. D. Jr., and Vecchi, M. P., 1983, *Science*, **220**, 671.
[10] Tobias, D. J., Martyna, G. J., and Klein, M. L., 1993, *J. phys. Chem.*, **97**, 12959.
[11] Swope, W. C., Andersen, H. C., Berens, P. H., and Wilson, K. R., 1982, *J. chem. Phys.*, **76**, 637.
[12] Pulay, P., 1980, *Chem. Phys. Lett.*, **73**, 393.
[13] Fois, E. S., Penman, J. I., and Madden, P. A., 1991, AIP Conference Proceedings. *Advances in Biomolecular Simulation*, Vol. 239, p. 33.
[14] Blöchl, P. E., and Parrinello, M., 1992, *Phys. Rev.* B, **45**, 9413.
[15] Fois, E. S., Penman, J. I., and Madden, P. A., 1993, *J. chem. Phys.*, **98**, 6361.
[16] Tuckerman, M., and Parrinello, M., 1994, *J. chem. Phys.*, **101**, 1301.
[17] Martyna, G. J., Tuckerman, M. E., Tobias, D. J., and Klein, M. L., 1996, *Molec. Phys.*, **87**, 1117.
[18] Tuckerman, M. E., Berne, B. J., and Martyna, G. J., 1992, *J. chem. Phys.*, **97**, 1990.
[19] Gibson, D. A., and Carter, E. A., 1993, *J. phys. Chem.*, **97**, 13429.
[20] Raghavachari, K., 1986, *J. chem. Phys.*, **84**, 5672.
[21] Fournier, R., Sinnott, S. B., and DePristo, A. E., 1992, *J. chem. Phys.*, **97**, 4149.