Simple bias potential for boosting molecular dynamics with the hyperdynamics scheme

M. M. Steiner and P.-A. Genilloud
Institute for Theoretical Physics, ETH-Honggerberg–HPZ, Zurich 8093, Switzerland

J. W. Wilkins
Department of Physics, Ohio State University, Columbus, Ohio 43210-1106

(Received 3 December 1997)

Two modifications of Voter's hyperdynamics scheme offer significant speedup of molecular dynamics simulations. (1) A simple construction of the bias potential—a few tens of lines of code—is validated for three systems. (2) A local construction of the bias potential permits the use of intuition to further improve the statistical error. These results suggest widespread applicability and the possibility of overcoming the unfavorable exponential scaling of molecular dynamics simulations as the temperature is lowered.
[S0163-1829(98)03917-4]

The simplicity of molecular dynamics (MD) simulations has led to their wide usage in chemistry and physics; one ''simply'' integrates the equations of motion governing the system forwards, measuring observables at suitable time intervals. A central assumption is that the path followed by the system is in some way ''typical,'' and that averages over this path, or many paths, are truly representative of the system. While high-frequency, short-ranged vibrations can typically be well described by MD, one is often interested in observing the system through a sequence of infrequent transitions, such as hopping from one potential energy minimum to another or surface diffusion of an adatom. Such simulations often require long run times, and due to the exponential scaling with inverse temperature rapidly become prohibitively expensive.

Voter recently proposed the hyperdynamics scheme[1,2] to speed up those MD simulations whose goal is to study the long-time behavior of systems. The scheme modifies the potential energy surface specifically (1) to reduce the CPU time spent evolving the system when it is trapped in high-frequency oscillations in potential minima and (2) to obtain an undistorted long-time description of the system's evolution. The cost is that detailed vibrational information is lost.

To understand qualitatively how a modified potential surface can lead to a boost in the computational speed, consider two potential surfaces in multidimensional configuration space $\mathbf{R}$: $V(\mathbf{R}) \leqslant V^*(\mathbf{R})$. The average time $\tau(\mathbf{R})$ spent at $\mathbf{R}$ is proportional to the probability of being at $\mathbf{R}$, which in turn is proportional to the Boltzmann factor. Hence $\tau(\mathbf{R})$ and $\tau^*(\mathbf{R})$ are proportional to $e^{-\beta V(\mathbf{R})}$ and $e^{-\beta V^*(\mathbf{R})}$, respectively. Defining the *bias* potential as $\Delta V_b = V^* - V$, we have

$$\tau(\mathbf{R}) = \tau^*(\mathbf{R}) e^{\beta \Delta V_b(\mathbf{R})}. \tag{1}$$

In an MD simulation on V^* a time interval Δt_{MD} spent in the neighborhood of $\mathbf{R}$ corresponds to a larger time interval $\Delta t_{\mathrm{MD}} \exp[\beta \Delta V_b(\mathbf{R})]$ for an MD simulation on V.

This qualitative argument has a more quantitative form leading to a *boost* factor of the computational time over conventional MD simulations:

$$\text{boost factor} = \left\langle e^{\beta \Delta V_b(\mathbf{R})} \right\rangle_{V^*}, \tag{2}$$

where the expectation value is taken over the classical canonical ensemble for the V^* potential. Within the framework of importance sampling a formal argument[1] can determine the escape rate $k_{A\rightarrow}$ from a well-defined region of configuration space A. In particular the escape rate for the biased potential surface V^* is related to the desired escape rate for V by

$$k_{A\rightarrow} = \left\langle |v_\perp| \, \delta_A(\mathbf{R}) \right\rangle_V = \frac{\left\langle |v_\perp| \, \delta_A(\mathbf{R}) e^{\beta \Delta V_b(\mathbf{R})} \right\rangle_{V^*}}{\left\langle e^{\beta \Delta V_b(\mathbf{R})} \right\rangle_{V^*}}, \tag{3}$$

where $v_\perp$ is the velocity perpendicular to the surface of A, which surface is defined by $\delta_A(\mathbf{R})$. We explicitly indicate on which potential surface the MD averaging is performed.

Central to substantially decreasing the computational time is the bias potential $\Delta V_b(\mathbf{R})$, which is added onto V to decrease the number of time steps spent in oscillatory motion in potential minima. In Refs. 1 and 2 the bias potential, which fills in the minima, was required to go to zero smoothly near a saddle point so that the evolution of the system through a saddle point region was unaffected.[4] Our approach abandons this reassuring feature to speed up the construction of $\Delta V_b(\mathbf{R})$.

For our simpler construction, in any region where $V(\mathbf{R})$ falls below a carefully chosen boost energy E_B, $V(\mathbf{R})$ is replaced by the constant E_B. In practice, the transition between $V(\mathbf{R})$ and E_B depends on $\chi \equiv \chi(\mathbf{R}) = V(\mathbf{R}) - E_B$ so as, on the one hand, to produce a smooth crossover and, on the other hand, to restrict the range of the smoothing region. Specifically

$$\Delta V_b = f(\chi) + E_B - V = \chi g(\chi) + h(\chi) + E_B - V, \tag{4}$$

where $g(x) = 1/(1+e^{-\alpha x})$ and $h(x) = e^{-\gamma^2 \alpha^2 x^2}/\alpha$. The choice of $\gamma = 1/3$ and α between 10 and 50 achieves the joint goals of a restricted, but smooth, crossover. We stress that this bias potential is not zero on the surface of A and that the ''numerator correction'' (NC) in Eq. (3) may need to be computed for accurate results.[3]

The forces due to V^* are directly related by the chain rule to those for V, and the boost factor [see Eq. (2)] is easily computed—in fact the total modification to a regular MD code is a few tens of lines of code. Further, in the ''flat'' region forces need not be computed, only the potential energy; typically we find that for up to half of the time steps the system is in the ''flat'' region—the force on V^* is less than 0.001 that on V. Additionally, physical intuition can be used to concentrate on a particular subcomponent of the potential energy. At the risk of introducing an unwanted bias, this improves the statistics and results. Finally, by lowering E_B we can systematically approach the traditional MD result.

Results. Our examples are closely related to the simulations in Ref. 1, the main difference being that we use unbiased MD as the reference point, rather than transition-state theory. The second-order Langevin-Verlet algorithm of Allen and Tildesley[5] is used. Our choices of Δt_{MD} and α are conservative, and the samples are initially well thermalized.

The first example is a simple two-dimensional potential energy surface which is periodic in the x direction:

$$V(x,y)=\cos(x)\left(1+\frac{4}{\pi}y\right)+y^2. \tag{5}$$

With energies at the minimum and saddle point being -1.41 and 0.59, $\Delta t_{\mathrm{MD}}=0.1$, and the Langevin damping being 0.4, we determine the hopping rate in the x direction (k_{hop}).

For a wide temperature range, Fig. 1(a) compares the hopping rates from MD simulations with hyper-MD results for different E_B's. The agreement is well within the statistical noise. The NC, which arises because of the smoothening procedure, is neglected—it is largest for $E_B=0.2$ and $T=0.3$, where it is 1.07; otherwise, it is less than 1.02.

Figure 1(b) shows the corresponding boost factor in Eq. (2) for the hyper-MD simulations. At lower temperature one obtains very large boosts, but even with low, conservative E_B's one obtains one to three orders of magnitude.

In a related test, we added a $0.1\cos(x/2)$ term to $V(x,y)$ and measured the left- and right-hopping rates ($k_{L/R}$)—it is important that relative bias (k_L/k_R) is not introduced. The difference in the left- and right-barrier heights is 0.2, and so assuming that the entropic prefactors cancel, we expect from the Arrhenius equation $k_L/k_R=\exp(-0.2/T)$. Figure 2 shows how well the temperature dependence of $k_L/k_R\exp(0.2/T)$ equals unity: within 3–10 % statistical noise—the slow (left) escape rate dominates the noise. Generally we can neglect the NC, but for the very aggressively biased $E_B=0.2$, including the NC yields results in statistical agreement with the expected result.

A final case is the diffusion of a Lennard-Jones (LJ) atom along a (100) terrace—we use the same configuration as the last problem in Ref. 1. We took the strength and the width parameters in the LJ potential both to be unity, and the lattice constant was chosen to be 1.10 to account for thermal expansion. The Langevin damping was taken to be 0.4, and the time step was usually 0.02, but doubled for the low-temperature reference MD simulations. The minimum potential energy of the system was $V_{\mathrm{min}}=-15.30$ and the lowest saddle point was at -14.93.

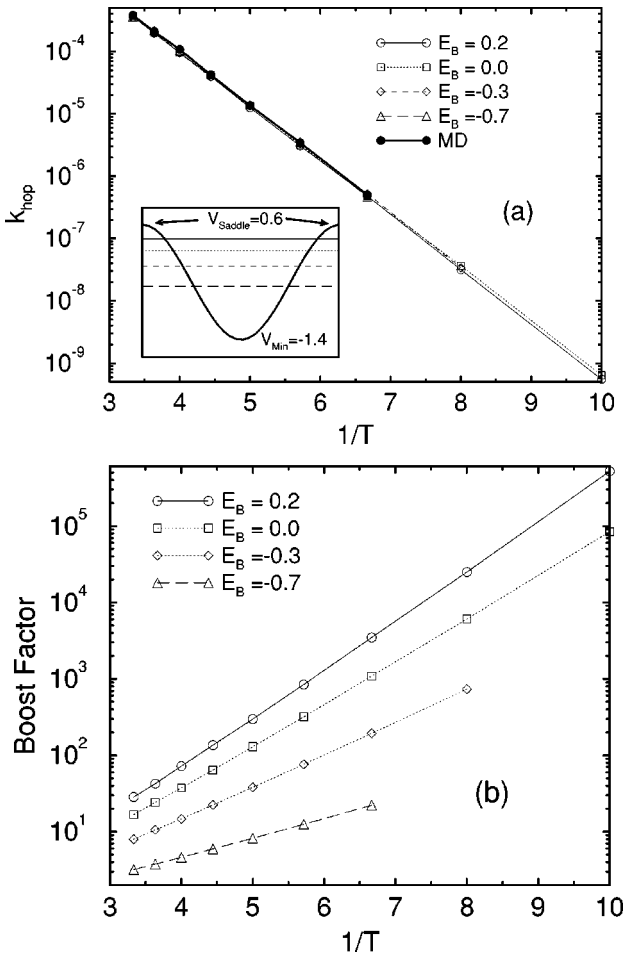

FIG. 1. Comparison of the temperature-dependent hopping rate (k_{hop}) of MD simulations with hyper-MD results for different boost energies E_B's. The inset schematically represents both the simple two-dimensional potential (3) from the minimum to the saddles and the four values of E_B. (a) shows good agreement between the rates for all values of E_B; we have neglected the NC—see text for details. (b) shows the boost factor [see Eq. (2)] which at low temperatures can become very large, corresponding to an enormous reduction in CPU time. For each point 1024 samples were taken, producing statistical errors less than the size of the symbols.

We used two different constructions of the bias potential: (i) *global*, which refers to the construction described in Eq. (4), and (ii) *local*, which concentrates on the atom which is hopping—the event we are interested in. We define V_{local} as the potential energy of the hopping atom, and then to determine the bias potential we replace V by V_{local} everywhere in Eq. (4). The determination of the forces is slightly trickier, as only the forces associated with the hopping atom are scaled.

Figure 3 compares the temperature dependence of the hopping rate of MD simulations with the results from hyper-MD calculations using both the global and local constructions. To obtain a substantial boost in the global approach one requires E_B^{glob} be above the potential energy of the lowest saddle point. The estimate of E_B^{glob} is as follows: Most of the potential energy is associated with displacements in the $N_{\mathrm{DF}}=3\times9-1=26$ degrees of freedom that are perpendicular to the direction of interest. For the temperature range in Fig. 3 we find that $N_{\mathrm{DF}}T/2$ lies between 0.3 and 0.6. Adding 0.6 to V_{min} gives -14.5. We present results for

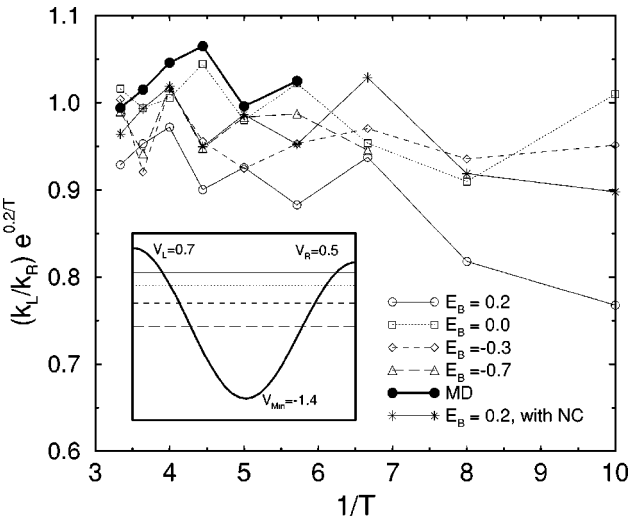

FIG. 2. Comparison of the temperature dependence of the ratio of the left- and right-hopping rates (k_L/k_R) of MD simulations with hyper-MD results for different E_B's. The potential in Eq. (5) has an additional $0.1\cos(x/2)$ term. The inset schematically represents both the potential from the minimum to the saddles and the four values of E_B. We multiply k_L/k_R by the inverse of the expected result to emphasize the errors. We neglect the NC, but for the case where it is largest: $E_B=0.2$ for which we show the effect of including it (see line with stars). The 4096 samples lead to a statistical error at the 3% (10%) level at high (low) temperature.

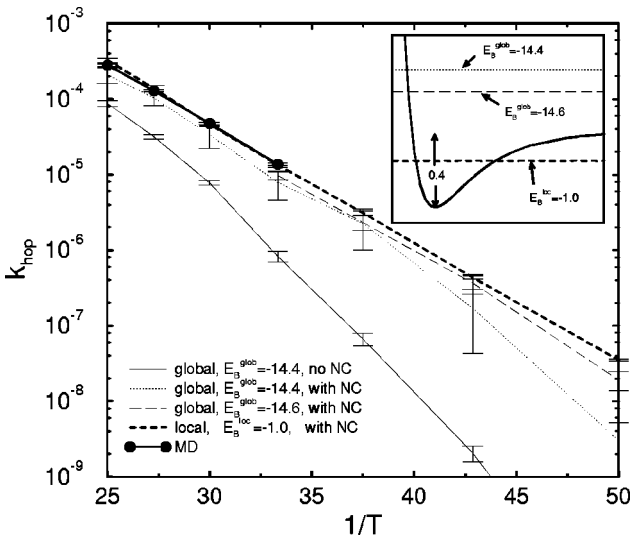

FIG. 3. Effect of global and local constructions of the bias potential on hyper-MD results for the hopping rate (k_{hop}). The inset schematically represents the global or local potential, and shows the relative positions of the global or local E_B's. The unprejudiced global approach exhibits large statistical and systematic errors, which can be reduced at the cost of more computer time (i.e., by reducing E_B^{glob}). (For $E_B^{\mathrm{glob}}=-14.6$ we show only the low temperature results for clarity.) The local approach shows much smaller statistical and systematic error, but runs the risk of introducing unforeseen biases. We used 1000 samples except for the $T=0.03$ reference MD simulations where 500 samples were used.

$E_B^{\mathrm{glob}}=-14.4$ and -14.6. For $E_B^{\mathrm{glob}}=-14.4$ the effect of the NC is substantial, indicating that on the surface defining a hop there can be a significant boost factor—in contrast to Voter's construction of the bias potential. To determine the boost factor the ratio of the ''bare'' boost factor to the NC [see Eq. (3)] must be taken. For $E_B^{\mathrm{glob}}=-14.4$, at the lowest temperature the average effective boost is $1.4\times10^7/470=3.6\times10^4$, while for $E_B^{\mathrm{glob}}=-14.6$ we find a lower boost $10^5/7.3=1.4\times10^4$. The results for the lower boost energy $E_B^{\mathrm{glob}}=-14.6$ yield $\ln k$ more nearly linear in $1/T$, and with respect to the MD results the systematic error is half that of the $E_B^{\mathrm{glob}}=-14.4$ results.

The local construction of the bias potential is a simple attempt to build some physics into the bias. Of course the risk is that other mechanisms will not be boosted and therefore suppressed—for example, atom exchange with the substrate will be suppressed by this construction. As with importance sampling there is significant variance reduction—the statistical errors are much smaller than with the global construction. This is related to the fact that the NC is much smaller (between 1.2 and 1.3) and much less noisy, which we expect to be the case as $E_B^{\mathrm{loc}}=-1.00$ is *below* the ''local'' potential energy of the saddle point (-0.89). The local results are also more linear and in good agreement with the reference MD calculations. At high temperature the effective boost factor was 23, while at low temperature it was 1.6×10^4.

Runs on double- and quadruple-length terraces further demonstrate the viability of the local approach; the high boost factors and low noise levels seen in the short terrace persist. On the other hand, when using the global construction, to maintain the noise level at that of the $E_B^{\mathrm{glob}}=-14.6$ level, $E_B^{\mathrm{glob}}-V_{\mathrm{min}}$ had to be lowered, reducing the boost factor by a factor of 3.

Conclusions. The hyper-MD scheme proposed by Voter is a powerful approach for reducing the CPU requirements of MD simulations of rare events with a concomitant loss of short time vibrational information. Our simple construction of the bias potential [see Eq. (4)] regularly produces boosts of three to five orders of magnitude, without a significant loss of accuracy. For simple systems with only a few degrees of freedom the very simple *global* construction is found to work well. For larger systems the slightly more complex *local* construction of the bias potential builds on physical intuition to reduce the statistical error with respect to the simpler *global* construction—however, the risk of suppressing unforeseen mechanisms must be recognized. Additionally, reducing the bias potential systematically improves the results.

It is possible to generalize the local scheme to systems (a) where the interesting events involve the motion of more than one atom or (b) the local potential is not readily constructed. For example, in a tight-binding-scheme MD approach, where a local potential energy is not defined, one can use any reasonable approximation to it—importance sampling will take care of the rest. Finally, Fig. 1(b) suggests that raising the bias potential or lowering the temperature can totally overcome the exponential slowdown of traditional MD simulations.[6]

We thank Art Voter and Shiwei Zhang for useful discussions and a critical reading of the manuscript. M.M.S. and P.A.G. were supported by the Swiss National Foundation. M.M.S. was on leave of absence from the Ohio State University, where the latter part of this work was performed and where he was supported by NSF Grant No. DMR-9520319. J.W.W. was supported by U.S.DOE Basic Energy Sciences, Division of Materials Sciences.

[1] A. F. Voter, J. Chem. Phys. **106**, 4665 (1997).
[2] A. F. Voter, Phys. Rev. Lett. **78**, 3908 (1997).
[3] A partial understanding of the NC is obtained by considering $V^* = V + \text{const}$. In this case the NC exactly cancels the boost factor in Eq. (3), leading to identical dynamics on V^* and V.
[4] For Voter's construction of the bias potential, the exponential in the numerator in Eq. (3) is unity.
[5] M. P. Allen and D. Tildesley, *Computer Simulations of Liquids* (Oxford, New York, 1987), p. 263.
[6] W. Wang and J. W. Wilkins (private communication).