

Temperature-accelerated dynamics for simulation of infrequent events

Mads R. Sørensen and Arthur F. Voter^{a)}

Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545

(Received 17 February 2000; accepted 18 February 2000)

We present a method for accelerating dynamic simulations of activated processes in solids. By raising the temperature, but allowing only those events that should occur at the original temperature, the time scale of a simulation is extended by orders of magnitude compared to ordinary molecular dynamics, while preserving the correct dynamics at the original temperature. The main assumption behind the method is harmonic transition state theory. Importantly, the method does not require any prior knowledge about the transition mechanisms. As an example, the method is applied to a study of surface diffusion, where concerted processes play a key role. In the example, times of hours are achieved at a temperature of 150 K. © 2000 American Institute of Physics.

[S0021-9606(00)70618-0]

I. INTRODUCTION

A long-standing obstacle to the understanding of condensed-phase systems is that many important processes occur on a time scale that is not easily accessible with conventional simulation methods. The standard method for dynamic simulations of materials at the atomistic level, molecular dynamics (MD), is generally limited to nanoseconds because of the small time step required for the numerical integration of the equations of motion. However, relevant processes that are *activated*, i.e., *infrequent events*, often take place on a time scale of microseconds or longer. Examples include the evolution of the surface morphology during crystal or film growth, the diffusion of point defects in solids, and the migration of grain boundaries during plastic strain.

For some problems, such as crystal growth, the kinetic Monte Carlo method¹ has been widely used to simulate very long time scales.² However, kinetic Monte Carlo is limited by the requirement that a complete catalog of all relevant processes and their rate constants must be specified. To simplify the task, it is often assumed that the atoms occupy perfect lattice positions and that the kinetics can be well described by a small number of elementary processes. The key problem with this approach is that *in general not all relevant transition mechanisms are known*. The mechanisms may be complicated processes where several atoms move in a concerted manner, and such processes are difficult to find by intuition. Moreover, in many cases, the assumption of lattice positions is not appropriate, such as on a reconstructed surface, at a grain boundary, or in an inhomogeneous strain field. The construction of a complete catalog can then become a nearly impossible task. Thus, there is a strong motivation for developing new methods that do not rely on any prior assumptions about the relevant mechanisms.

In this paper we present a novel method, temperature-accelerated dynamics (TAD), that can extend the accessible time scale by orders of magnitude for dynamic simulation of infrequent events in solid-state systems.³ For an infrequent-

event system, the dynamic evolution of the system on a long time scale can be described as a sequence of *transitions* between distinct *states*. After each transition, the system equilibrates in a state before making the next transition to another state. It is the purpose of TAD to advance the system from state to state at an accelerated pace in a way that preserves the transition rates without imposing any restrictions on the transition mechanisms that can occur.

The TAD method is similar in spirit to the previously proposed hyperdynamics method⁴ in that it applies to infrequent event dynamics and is based on transition state theory (TST).⁵⁻⁷ The challenging task in hyperdynamics is the construction of a bias potential that modifies the potential energy surface without affecting the transition state regions. For systems with many degrees of freedom, it is nontrivial to construct a successful bias potential and the bias potential may require some tuning for a specific system. TAD makes the additional assumption of harmonic TST (HTST) and is thus more approximate than hyperdynamics, but it should be easier to implement in many cases. Moreover, TAD offers a natural way to reduce the so-called *low barrier problem*—the problem of the system remaining trapped in a set of states connected by low barriers. Like hyperdynamics, TAD may be combined with the parallel replica method⁸ to achieve an event higher acceleration on parallel computers.

In the following, the method is derived, its implementation is discussed, and its performance is illustrated with examples. The method's accuracy is investigated and its advantages and limitations are discussed.

II. METHOD

We are interested in systems for which the potential energy surface consists of basins separated by barriers that are high compared to the thermal energy, $k_B T$. The system spends a long time vibrating in a state before escaping to another state by an infrequent thermal fluctuation. The rates for transitions between states depend sensitively on the temperature; when the temperature is raised, all processes happen more frequently. However, merely raising the tempera-

^{a)}Electronic mail: afv@lanl.gov

ture would be an inappropriate way of accelerating the dynamics, since an increase in temperature changes not only the rate constants but also the ratios of the rate constants. As the temperature is raised, the ratio of a high-barrier rate constant to a low-barrier rate constant increases. Thus, the kinetics may be very different at different temperatures. The essence of the TAD method is to raise the temperature and correct for the temperature-induced bias by filtering out some of the transitions and allowing only those transitions that should occur at the original temperature.

A. Derivation

The method applies to first-order kinetics, as expected for an infrequent-event system.⁸ The waiting time, t , i.e., the time spend in a state before a particular transition occurs, is then exponentially distributed with a probability distribution

$$f_i(t)dt = k_i e^{-k_i t} dt, \quad (1)$$

where k_i is the rate constant for the process i .

We assume that TST applies to the system under study. In TST, one defines a dividing surface in configuration space separating the reactant state and product states, and it is assumed that every time a trajectory crosses the dividing surface, a truly reactive transition has occurred, i.e., successive crossings are uncorrelated. We further assume that the rate constant, k_i , at a temperature, T , is given by an Arrhenius expression

$$k_i = \nu_i e^{-\beta E_{a,i}}, \quad (2)$$

resulting from the harmonic approximation to TST.⁹ Here, $E_{a,i}$ is the energy barrier (the difference in energy between the saddle point and the minimum), ν_i is a temperature-independent pre-exponential factor (prefactor), and β equals $1/(k_B T)$, where k_B is the Boltzmann constant.

We introduce a special kind of MD simulation, called *basin-constrained MD*, where the trajectory is confined to a particular potential energy basin. When the system tries to escape through a dividing surface to another state, it is reflected back into the original state. (As explained below, this can be implemented without knowing the position of the dividing surface in advance. The only requirement is that transitions can be detected.) During a basin-constrained MD simulation, we record the waiting time for each transition pathway, i.e., the time of the first attempted escape along the pathway.

From a basin-constrained MD simulation, a sequence of waiting times $\{t_i\}$ is obtained. As an example, assume that a simulation for a case with three different escape pathways gave the sequence $t_3 < t_1 < t_2$, i.e., the first attempted escape was along pathway 3, then at a later time an escape was attempted along pathway 1, and finally an escape along pathway 2 was attempted. [Obviously, had the simulation been an ordinary, unconstrained MD simulation, the system would have moved to another state by the first attempted pathway (pathway 3 in the example).] According to our assumptions, each of the waiting times $\{t_i\}$ can be thought of as a statistical sample from an exponential distribution [Eq. (1)]. This implies that if the transition pathways and rate constants had been known in advance, we could have generated waiting

times simply by drawing random numbers from these exponential distributions. These considerations lead to the following interpretation, suitable for infrequent events: *An ordinary MD simulation is equivalent to drawing a random waiting time from the statistical distribution for each transition pathway and then discarding all times except for the shortest time.* In the following it is shown that even when the pathways and rate constants are *a priori* unknown, one can still use the idea of generating random waiting times for various pathways and accepting the transition with the shortest waiting time.

The idea is to use an MD simulation at a higher temperature to find transition pathways and simultaneously generate correctly distributed waiting times at the original temperature. To see how this works, consider a process, i , at two different temperatures, T_{low} and T_{high} ($T_{\text{low}} \leq T_{\text{high}}$), with corresponding rate constants, $k_{i,\text{low}}$ and $k_{i,\text{high}}$. Since the waiting times, $t_{i,\text{low}}$ and $t_{i,\text{high}}$, at the two temperatures are exponentially distributed, it follows that the products $k_{i,\text{low}} t_{i,\text{low}}$ and $k_{i,\text{high}} t_{i,\text{high}}$ have identical distributions. (This can be seen from the fact that if t has the probability distribution $f(t)dt = k \exp(-kt)dt$, then the product $x = kt$ is distributed as $f(x)dx = \exp(-x)dx$, for any value of k .) In other words, the exponentially distributed stochastic variables $t_{i,\text{low}}$ and $t_{i,\text{high}}$ are related as $k_{i,\text{low}} t_{i,\text{low}} = k_{i,\text{high}} t_{i,\text{high}}$, which can be rewritten using Eq. (2) as

$$t_{i,\text{low}} = t_{i,\text{high}} e^{E_{a,i}(\beta_{\text{low}} - \beta_{\text{high}})}. \quad (3)$$

Equation (3) implies that if the energy barrier, $E_{a,i}$, can be determined, one can simulate the system at a high temperature, T_{high} , where transitions can be directly observed with MD, and then *extrapolate* the observed waiting time to obtain a waiting time at the original temperature, T_{low} . Given the above assumptions, the distribution from which these waiting times are generated is identical to the distribution of waiting times in a direct simulation at T_{low} .

Now we have the ingredients for the TAD method, which works as follows: The system starts out in some state. A basin-constrained MD simulation is carried out at a high temperature and whenever a process is attempted for the first time, the energy barrier is determined and the high-temperature waiting time is extrapolated to the low temperature, which is the temperature of interest. After some time, the high-temperature MD simulation is terminated, the transition with the shortest extrapolated waiting time is accepted, and the system is moved to the state that this transition lead to. (A criterion for stopping the high-temperature MD simulation and accepting a transition will be derived below.) The whole procedure is repeated in the new state, and in this way the system is moved from state to state.

The method is illustrated pictorially in Fig. 1. The figure is similar to an Arrhenius plot; however, the ordinate is the inverse of an actual escape time observed in a simulation, and not the inverse of an average escape time (i.e., a rate constant). There is a vertical time line for each of the temperatures T_{high} and T_{low} , where the points move downwards as time progresses. The extrapolation from T_{high} to T_{low} corresponds to a straight line with slope $-E_{a,i}$. It is important to realize that the extrapolation lines may cross, so that the




FIG. 1. Pictorial representation of the TAD method. Time moves down the vertical time lines shown. Each attempted transition at the high temperature is extrapolated to the low temperature along a line with a slope given by the negative of the activation energy. The stop line (dashed) connects the points $(0, \ln[\nu_{\min}/\ln(1/\delta)])$ and $[\beta_{\text{low}}, \ln(1/t_{\text{low,short}})]$, where $t_{\text{low,short}}$ is the current, shortest low-temperature escape time. The MD simulation can be stopped when the high-temperature time reaches the intersection with the stop line. The transition corresponding to $t_{\text{low,short}}$ is accepted and the clock is advanced by $t_{\text{low,short}}$.

order of the escape times differs at the two temperatures. It is therefore essential to continue the MD simulation at T_{high} until one can be confident that the shortest waiting time at T_{low} has been found, as we discuss next.

We assume there is a lower bound, ν_{\min} , on the prefactors in the system. This is a good approximation for many problems of practical interest, e.g., surface diffusion prefactors tend to be around 10^{12} – 10^{13} s $^{-1}$. With this assumption, we can in a quantifiable way choose the MD simulation time, $t_{\text{high,stop}}$, so that there is only a low probability that if the MD simulation had been continued further, a new transition could occur with an extrapolated escape time shorter than $t_{\text{low,short}}$. The argument is the following: It follows from Fig. 1 that if a transition that occurs later than $t_{\text{high,stop}}$ is to have an extrapolated escape time shorter than $t_{\text{low,short}}$, the transition must have an energy barrier lower than $E_x = \ln(t_{\text{low,short}}/t_{\text{high,stop}})/(\beta_{\text{low}} - \beta_{\text{high}})$. Transitions with energy barriers higher than E_x extrapolate to times longer than $t_{\text{low,short}}$ and are not a concern—we need only consider energy barriers lower than E_x . Also, we are only concerned with prefactors higher than ν_{\min} , since we have assumed that processes with lower prefactors do not exist. We now require that a transition with energy barrier E_x and prefactor ν_{\min} has a low probability, δ , of occurring for the first time at a time later than $t_{\text{high,stop}}$. (If the prefactor is higher than ν_{\min} or the energy barrier is lower than E_x , the transition is even more unlikely to occur for the first time after $t_{\text{high,stop}}$.) By integrating Eq. (1) and using Eq. (2) and the expression for E_x , it can be shown that the required MD simulation time, $t_{\text{high,stop}}$, can be written as

$$t_{\text{high,stop}} = \frac{\ln(1/\delta)}{\nu_{\min}} \left(\frac{\nu_{\min} t_{\text{low,short}}}{\ln(1/\delta)} \right)^{\beta_{\text{high}}/\beta_{\text{low}}} \quad (4)$$

If, for instance, $\delta=0.001$, $t_{\text{high,stop}}$ indicates the simulation time required to be 99.9% confident that continuing the simulation would not give any new transition with an extrapolated escape time shorter than $t_{\text{low,short}}$. As seen from Eq. (4), ν_{\min} and δ could be replaced by a single parameter,

but we find it more transparent to keep both. Figure 1 illustrates how $t_{\text{high,stop}}$ can be found by a simple graphical construction.

The effect of stopping the MD simulation after a relatively short time is that high-barrier transitions tend to be accepted too frequently compared to low-barrier transitions. The parameters ν_{\min} and δ are chosen to ensure that an incorrect transition is accepted only rarely. For example, if the value ν_{\min} is chosen too high, it increases the uncertainty that the correct transition is accepted. The extreme case is when ν_{\min} is so high that the first attempted high-temperature transition is always accepted. An excessively low value of ν_{\min} will give a very small error, but there is a price to be paid in terms of efficiency, since more computational work is required for each accepted transition.

We stress that the method does not require knowledge of the prefactor for any process; the only assumption regarding the prefactor is that there is a lower bound, ν_{\min} . This is advantageous for larger systems since an explicit calculation of the prefactor within HTST⁹ would require computational work that scales as N^3 , where N is the number of moving atoms.

To summarize the approximations, we assume that the system obeys HTST, exhibits first-order kinetics, and that a minimum value, ν_{\min} , for the prefactors exists. The user specifies ν_{\min} and a parameter, δ , the level of uncertainty that each accepted transition is correct. In comparison with TAD, the hyperdynamics method⁴ assumes only TST, but requires that a valid bias potential be constructed.

B. Implementation

As discussed, the behavior of the system at a temperature, T_{low} , is studied by carrying out an MD simulation at a higher temperature, T_{high} . During the MD simulation, we check at regular intervals whether the system has made a transition from the current state, A , to another state. For the systems we have studied, this can be accomplished by following steepest descent or conjugate gradient steps towards the nearest local potential energy minimum (see Fig. 2). If the minimization leads to the original minimum (so that every atom is within a short distance from its position at the minimum), the system is considered to be in the same state. If, on the other hand, an atom is significantly displaced from its position at the minimum in A when the minimization is well converged, it is concluded that a transition to another state has occurred. It is conceivable that for other systems there may be more efficient ways to detect transitions. In the applications presented below, we checked for transitions every 2 ps and if a transition had occurred in a 2 ps interval, the time of the escape was determined to within 0.1 ps by minimizing intermediate MD configurations that were previously stored. When a transition has occurred, the saddle point for the transition is found in a manner discussed below, the energy barrier, $E_{a,i}$, is calculated, and the escape time, $t_{i,\text{high}}$, is extrapolated to $t_{i,\text{low}}$ using Eq. (3). The MD simulation is now resumed *in state A* by continuing the trajectory from a point just before the crossing of the dividing surface, but with the velocities reversed for all atoms (Fig. 2). The MD simulation employs a Langevin thermostat,¹⁰ which provides




FIG. 2. Illustration of the implementation of the TAD method. A two-dimensional potential energy surface is indicated by the thin contour lines. The thin dashed line labeled DS marks the dividing surface separating a state A from a neighboring state B. The thick line shows an MD trajectory starting in state A. The crossing of the dividing surface that takes place between the two trajectory points X_1 and X_2 is detected using local energy minimization. A steepest descent minimization initiated at X_1 leads towards the potential energy minimum X_A , whereas a minimization from X_2 goes to X_B . To search for the saddle point, a discretized path is constructed as a chain of configurations connecting the points X_A , X_1 , X_2 , and X_B (shown as the black dots connected by a thin line). The path is then optimized to a minimum energy path using the nudged elastic band method. The original MD trajectory continued into state B via point X_2 (dashed line, T_{MD}). To impose the basin constraint, the MD simulation is continued from X_1 , but with all velocities reversed (T_{BC}). The trajectory does not retrace its previous path since a Langevin thermostat injects noise into the trajectory.

canonical sampling at the desired temperature and injects noise in the trajectory. This noise ensures that the time-reversed trajectory does not simply retrace its path backwards. The MD simulation is stopped when the MD time has reached the time $t_{\text{high,stop}}$, which is determined by the current shortest extrapolated escape time, $t_{\text{low,short}}$ [Eq. (4)]. Then, the TAD clock at T_{low} is advanced by $t_{\text{low,short}}$, and the system is moved to the state corresponding to the accepted transition.

A key requirement of the method is that the saddle point that corresponds to a given crossing of a dividing surface in an MD simulation must be found. We have experimented with different techniques and found that the following procedure works well for the cases we have studied. A discretized transition path is constructed as a chain of approximately equally spaced configurations that connects the potential energy minima for the initial and final states via two configurations from the MD trajectory on each side of the dividing surface (Fig. 2). Linear interpolation is used to insert configurations in the chain. This transition path is then optimized to a minimum energy path using the nudged elastic band method.¹¹ The saddle point is the highest energy point on the minimum energy path.

III. SIMPLE EXAMPLES

A. A single process

To demonstrate the acceleration that can be obtained with the TAD method, we have applied the method to the diffusion of an adatom on a Ag(100) surface, as modeled by an embedded atom method (EAM) potential.¹² To keep the example simple, the tendency for the adatom to migrate via

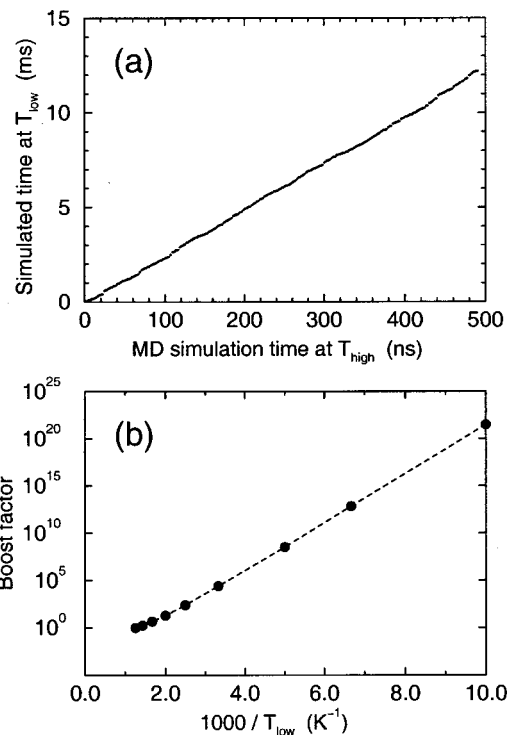


FIG. 3. Results from TAD simulations of an adatom hopping on a Ag(100) surface. Simulation parameters: $T_{\text{high}}=800$ K, $\nu_{\text{min}}=2 \times 10^{12} \text{ s}^{-1}$, and $\delta=0.001$. (a) Simulated time at $T_{\text{low}}=300$ K as a function of the MD simulation time at T_{high} . A total of 398 adatom hops was recorded during a time of 12.2 ms at T_{low} . (b) Boost factor as a function of the reciprocal of T_{low} .

an exchange mechanism was suppressed by allowing only the adatom and the top two surface layers to move (51 moving atoms), while two deeper layers were held fixed. The energy barrier for adatom hop on the surface was 0.524 eV and the prefactor within HTST was found to be $4.8 \times 10^{12} \text{ s}^{-1}$ using a Vineyard analysis.⁹ For these test simulations, we set $\delta=0.001$ and used our knowledge of the prefactor to choose $\nu_{\text{min}}=2 \times 10^{12} \text{ s}^{-1}$. We used an elevated temperature of $T_{\text{high}}=800$ K for the MD simulation, which was carried out using a Langevin algorithm¹⁰ with a friction coefficient of $1.0 \times 10^{12} \text{ s}^{-1}$ and a time step of $5.0 \times 10^{-15} \text{ s}$.

In Fig. 3(a) is shown the time at $T_{\text{low}}=300$ K as a function of the time spent in the high-temperature MD simulation. The slope of the line in Fig. 3(a) indicates a *raw boost factor* of around 2.5×10^4 . Including a computational overhead of around 27% due to detection of transitions and searches for saddle point, the *computational boost factor* is 2.0×10^4 (defined as the simulated time at T_{low} divided by the time that would have been simulated by direct MD with the same amount of computational work). The boost increases dramatically when T_{low} is reduced, as shown in Fig. 3(b); the boost climbs exponentially as a function of $1/T_{\text{low}}$.

B. Two competing processes

When there are several processes with various energy barriers, the kinetics may be very different at different temperatures. We have studied a simple example with two competing processes to check that the TAD method gives the correct statistics of the low-temperature transitions, i.e., that

the various transitions are accepted in the right proportions. To construct a simple example, we strained the system from the previous example in the surface plane by -1.5% in one close-packed direction and $+1.5\%$ in the other, creating two inequivalent hop directions, x and y . The energy barriers for hop in the x and y directions were 0.498 and 0.548 eV, respectively, and the prefactors within HTST were $4.8 \times 10^{12} \text{ s}^{-1}$ and $4.9 \times 10^{12} \text{ s}^{-1}$. We studied adatom diffusion at $T_{\text{low}} = 300 \text{ K}$ using the same values for T_{high} , δ , and ν_{min} as in the previous example.

In the TAD simulation, a time of $1.78 \times 10^{-2} \text{ s}$ was accumulated at the low temperature. During this time, 921 x -hops and 128 y -hops were accepted. This gives low-temperature rate constants of $2.58 \times 10^4 \text{ s}^{-1}$ and $3.60 \times 10^3 \text{ s}^{-1}$, taking into account two equivalent hop directions (forward and backward) for each kind of hop. For comparison, the high-temperature rate constants determined from counting attempted transitions in the basin-constrained MD simulation are $4.35 \times 10^9 \text{ s}^{-1}$ and $2.03 \times 10^9 \text{ s}^{-1}$. These high-temperature rate constants should extrapolate to the low temperature as $k_{i,\text{low}} = k_{i,\text{high}} \exp(-E_a(\beta_{\text{low}} - \beta_{\text{high}}))$. The rate constants obtained by this extrapolation are $2.56 \times 10^4 \text{ s}^{-1}$ and $3.56 \times 10^3 \text{ s}^{-1}$. The fact that these extrapolated rate constants agree within error bars with the rate constants of the accepted transitions confirms that transitions are accepted correctly at the low temperature. The simulation also demonstrates that the kinetic behavior obtained with the TAD method at the low temperature is significantly different from the results that would have been obtained if we had simply carried out ordinary MD at the high temperature: The ratio of the two hopping rate constants, k_x/k_y , is 7.2 at T_{low} , as opposed to 2.1 at T_{high} .

C. Anharmonic effects

The accuracy of TAD depends on the accuracy of HTST at the high temperature used in the MD simulation. To test the accuracy for realistic applications, we have compared rate constants determined from MD simulations at temperatures of 500–1000 K to rate constants predicted from HTST for various atomic mechanisms in metals.

As the first example, we investigate adatom hopping and exchange processes on the Ag(100) surface. In this case, the atoms in the top three substrate layers were allowed to move, to reduce the energy barrier for the exchange mechanism, as compared to the previous examples. The energy barriers for hop and exchange were 0.542 and 0.555 eV and the prefactors within HTST were $4.9 \times 10^{12} \text{ s}^{-1}$ and $4.5 \times 10^{12} \text{ s}^{-1}$, respectively.

The results are shown in Fig. 4. As the temperature is reduced, the MD rate constant approaches the HTST prediction. (It was not feasible to calculate the MD rate constant at temperatures below 500 K because of the long time between transitions.) At higher temperatures, the discrepancy between the MD and TST rate constant increases, as expected for anharmonic effects. The error in the harmonic approximation is larger for the exchange than for the hop mechanism. At 1000 K, the HTST rate constant is accurate within 60%, but for the exchange, the HTST rate constant is off by a factor of 4.

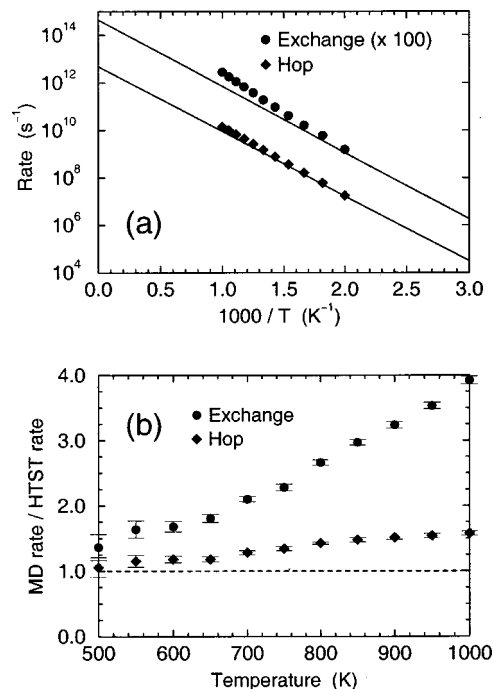


FIG. 4. Anharmonic effects for adatom hop and exchange mechanisms on a Ag(100) surface. (a) Arrhenius plot of the rate constant calculated by direct MD simulation (symbols) and the rate constant predicted from HTST (lines). For clarity the rate constant for the exchange mechanism is multiplied by a factor of 100. (b) Ratio of the MD rate constant to the HTST rate constant as a function of temperature.

We have carried out similar tests for other diffusion processes at surfaces and in the bulk of metals using EAM potentials. Examples include surface vacancy hop on Ag(100) and Ni(100), glide of a hexagonal 7-atom cluster on Ag(111) and Cu(111), and bulk vacancy hop in Ag and Cu. For these examples we find that the MD rate constants at suitable temperatures (700–1000 K) are within a factor of 2 of the HTST rate constants. For previous investigations of the effect of anharmonicity on bulk vacancy hopping rates, see Ref. 13.

We think that this level of accuracy is sufficient for the method to be useful, since at 300 K, a factor of 4 in the rate constant only corresponds to a change in the energy barrier of around 35 meV, so for many practical purposes, errors due to anharmonic effects will be smaller than errors due to the approximate interatomic potential.

We note that if T_{high} is not too much higher than T_{low} , the anharmonicity error is less severe, since anharmonic effects are expected to be similar at the two temperatures. In other words, the rate constant observed in this kind of a TAD simulation could be more accurate than the HTST prediction. Thus, by raising the temperature moderately TAD may offer an appreciable boost with almost no sacrifice of accuracy.

Thermal expansion is a special kind of anharmonic effect that can give rise to non-Arrhenius behavior, especially for systems with periodic boundary conditions or fixed boundaries. It is not clear if there exists a generally remedy that will work well in all cases. A tractable improvement in some cases may be to allow the system (i.e., the simulation supercell) to expand during high-temperature MD and contract during energy minimization and saddle-point searches.

Along similar lines, we note that the transition rate constants will be strongly non-Arrhenius if the system undergoes a second-order phase transition as the temperature is raised. In terms of the rate constants, this can be thought of as an anharmonic effect that makes the method less accurate.

IV. THE LOW-BARRIER PROBLEM

We have extended the method to address a situation that often arises for more complicated and realistic systems. The problem shows up when the system reaches a set of states that are connected by pathways with energy barriers that are low compared to energy barriers for pathways leading out of this subset of states. At low temperatures, in particular, the system will make very many transitions among these states before moving on to other states. Since every transition requires some amount of MD work, these repeated low-barrier transitions can dramatically slow the overall progress of the simulation.

A. Extension of method

In the extended method, we keep a cumulative history of every attempted transition out of every state and keep track of the total time spent in each state. When a particular pathway has been attempted a sufficient number of times (N_{attempts}), the high-temperature rate constant for this process can be estimated fairly accurately as the number of occurrences of the process divided by the total time spent in this state. Within the TAD approximations, the low-temperature rate constant can then be obtained by extrapolating the high-temperature rate constant as $k_{i,\text{low}} = k_{i,\text{high}} \exp(-E_a(\beta_{\text{low}} - \beta_{\text{high}}))$. Hereafter, the process is assigned to a special *synthetic* class. There is no need for using MD to generate waiting times for the synthetic processes, since random waiting times for these processes can be simply drawn from exponential distributions [Eq. (1)] using the estimated rate constants. We can use this to further accelerate the state-to-state evolution. First, during a high-temperature, basin-constrained MD simulation, attempted escapes along pathways in the synthetic class are excluded from the time-extrapolation scheme, so that a valid, low-temperature waiting time for a slower, nonsynthetic transition can be found. Second, the time gap until the nonsynthetic transition is to occur can then be filled up by drawing random waiting times for the synthetic transitions from the proper exponential distributions. We now have a list of all the escapes that shall occur from this state for some time into the future. When the state is revisited, the next transition in the list can be chosen without the need for any MD simulation. Only when the list has been used up and the nonsynthetic transition has been accepted is a new MD simulation required.

B. Simple example

We now demonstrate the gain in computational boost that can be obtained using the procedure for reducing the low-barrier problem. For this purpose we have simulated the diffusive motion of a three-atom island (trimer) on a Ag(100) surface. A system was created by adding two extra atoms to the system studied in Sec. III A. The lowest-energy state is a

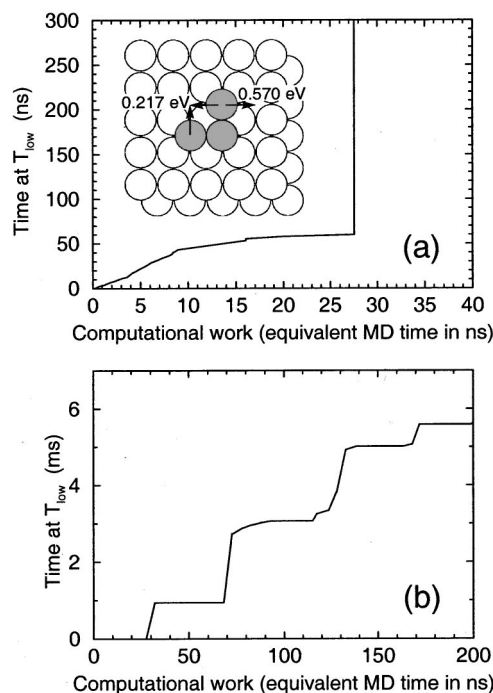


FIG. 5. Results from a TAD simulation of the diffusion of a trimer on a Ag(100) surface at $T_{\text{low}}=300$ K. Simulation parameters: $\nu_{\text{min}}=2 \times 10^{12} \text{ s}^{-1}$, $\delta=0.01$, and $N_{\text{attempts}}=10$. In (a) and (b) the simulated time at T_{low} versus the computational work, indicated by the equivalent MD time, is shown on two different scales. Insert in (a) shows the lowest energy trimer configuration and some atomic processes.

triangular trimer on the Ag(100) surface (see insert in Fig. 5). There is a low-energy barrier of 0.217 eV for each of the two weakly bound trimer atoms to hop to an equivalent site. By repeated occurrences of these hops, the island can switch between four equivalent configurations that essentially differ by a 90° rotation. Processes that allow the trimer to diffuse on the surface have significantly higher energy barriers, the lowest being 0.570 eV. A TAD simulation was carried out at $T_{\text{low}}=300$ K with $\nu_{\text{min}}=2 \times 10^{12} \text{ s}^{-1}$, $\delta=0.01$, and $N_{\text{attempts}}=10$. The temperature T_{high} was adjusted automatically for each state to achieve a value on the order of 6 for the ratio of the lowest energy barrier to $k_B T_{\text{high}}$ (excluding synthetic processes). This was accomplished in an *ad hoc* manner by initially choosing a high temperature (1100 K in this case). If an energy barrier, $E_{a,i}$ was found to be lower than $4k_B T_{\text{high}}$, the temperature was reduced to $E_{a,i}/(6k_B)$. (Of course, T_{high} was not allowed to be lower than T_{low} .) On the other hand, if the high-temperature basin-constrained MD simulation continued for too long without acceptance of a transition, the temperature was raised by 50%.

In Fig. 5, the time at T_{low} is shown as a function of the time that would have been accumulated in an ordinary MD simulation with the amount of computational work used in the TAD simulation (including detection of transitions and saddle-point searches). The time evolved in a highly nonuniform manner. Initially, the low-barrier processes had not yet been identified and the simulation was in the standard mode. The computational boost factor was low (around 5) because the low-barrier transitions would occur so frequently at T_{low} that there was little to be gained by raising T_{high} . Later, the




FIG. 6. Selected snapshots illustrating the structure of 1/4 monolayer of Cu on Cu(111) at different times during a TAD simulation at 150 K. The frame indicates the size of the simulation cell. Between pictures 5 and 6, two dimer shears have taken place (shown by arrows). Between pictures 6 and 7, the upper island has glided twice.

low-barrier processes were put in the special synthetic mode and after some additional MD simulation time, with T_{high} increased to 1100 K, a waiting time for a high-barrier transition could be found for each of the four equivalent trimer configurations. After this point, the time leaped forward with very little computational work since roughly a million synthetic transitions could be accepted without the need for any MD simulation (it only required generating random numbers and updating a few variables). In Fig. 5 this stage appears as a nearly vertical line segment indicating a rapid increasing in time. Eventually, one of the high-barrier transitions was accepted, and the trimer briefly changed shape and moved to a new site on the surface. The first move was completed after 0.94 ms and by then the computational boost, averaged over the simulation, had increased by four orders of magnitude. Hereafter, the above procedure repeated itself.

V. APPLICATION: FILM RIPENING ON CU(111)

As a more complex demonstration, we have applied the TAD method to simulate the evolution of 1/4 monolayer of Cu atoms on a Cu(111) surface at 150 K. Initially, 14 adatoms were deposited randomly on a Cu(111) slab of size $7 \times 4\sqrt{3}$ with the three top layers free to move (182 moving atoms) while two deeper layers were held fixed. The interatomic interactions were modeled by an EAM potential.⁸ The following TAD parameters were used: $\nu_{\text{min}} = 5 \times 10^{11} \text{ s}^{-1}$, $\delta = 0.01$, and $N_{\text{attempts}} = 10$. The temperature T_{high} was adjusted automatically as in the example in Sec. IV B. The evolution of the 1/4 monolayer in the simulation is shown in Fig. 6.

At 150 K the motion of single adatoms and dimers on the Cu(111) surface is fast enough to be simulated with or-

inary MD. After 38 ps, the atoms have formed two islands on the surface, containing 6 and 8 atoms. At this point, nothing further happens on a ns time scale, and the TAD method takes over. The 6-atom island diffuses on a μs time scale by a glide mechanism,¹⁴ i.e., a process where all atoms in the island move simultaneously, and at $t = 0.67 \text{ ms}$ it reaches the 8-atom island, where it picks off an atom, leaving a hexagonal 7-atom island. The other 7-atom island transforms into a hexagonal shape via two dimer shears.¹⁵ The hexagonal islands glide on a ms time scale, and at $t = 0.12 \text{ s}$ they coalesce to form one island which is stable on a time scale of seconds. On a longer time scale, corner atoms break loose and diffuse along the edge. By this edge running mechanism,¹ the island reaches a compact island shape at $t = 17 \text{ min}$, which remains stable for hours.

In this simulation there is a progression towards increasing energy barriers as the system evolves towards more stable clusters, giving an increasing boost. The increasing energy barriers allow T_{high} to be raised from an initial value of 150 K to a maximum of 1000 K. When the TAD simulation has reached 17 min, the computational boost, averaged over the simulation, is 2.4×10^9 ; an ordinary MD simulation with the same computational work would have reached only $0.41 \mu\text{s}$. In the TAD simulation a total of 2793 transitions occurred and 48 different states were visited; the vast majority of the transitions occurred in the special synthetic mode. As indicated above, the boost in TAD depends strongly on temperature. In the specific case, a simulation at 300 K, starting from the same configuration as above, showed the formation of a compact 14-atom cluster after only $50 \mu\text{s}$, giving an average computational boost of 260.

VI. CONCLUSION

We have presented a method for dynamic simulations of activated processes that can extend the simulation time by orders of magnitude compared to ordinary MD. There are several appealing features of the method: (1) It is relatively simple to implement, (2) it uses first derivatives only, and (3) the detailed information about available pathways and energy barriers, obtained automatically during the simulation, is extremely useful.

The shortcomings of the method appear for systems where the assumptions underlying the method do not strictly apply. For instance, transitions that do not obey TST, such as a double jump of an adatom, composed of two single jumps in rapid succession, are beyond the scope of the method. Also, the method will not be accurate for systems with strong anharmonic effects. There may also be problems for systems where the saddle point corresponding to a reactive trajectory cannot be uniquely determined.

We also point out that the idea of a basin-constrained MD simulation combined with saddle searches provides an automated procedure for finding all possible transition pathways with reasonably low energy barriers. By assuming a lower bound on the prefactors, one can estimate the confidence that one has found any process that has an energy barrier lower than some value.

In conclusion, we think that the examples presented in this paper show that the method looks promising for solid-

state systems where transitions can be easily detected and anharmonic effects play a minor role.

ACKNOWLEDGMENTS

We thank J. C. Hamilton and S.-C. Ying for useful comments and A. Redondo, T. C. Germann, R. LeSar, and T. Rasmussen for critical reading of various versions of the manuscript. This work was supported by the United States Department of Energy, Office of Basic Energy Sciences, under DOE Contract No. W-7405-ENG-36.

¹A. F. Voter, Phys. Rev. B **34**, 6819 (1986).

²For an example of an application, see J. Jacobsen, K. W. Jacobsen, P. Stoltze, and J. K. Nørskov, Phys. Rev. Lett. **74**, 2295 (1995).

³A preliminary account of the TAD method has appeared in A. F. Voter and M. R. Sørensen, Mater. Res. Soc. Symp. Proc. **538**, 427 (1999).

⁴A. F. Voter, J. Chem. Phys. **106**, 4665 (1997); Phys. Rev. Lett. **78**, 3908 (1997).

⁵C. H. Bennett, in *Algorithms for Chemical Computations*, edited by R. E. Christofferson (American Chemical Society, Washington, DC, 1977).

⁶D. Chandler, J. Chem. Phys. **68**, 2959 (1978).

⁷J. B. Anderson, Adv. Chem. Phys. **91**, 381 (1995).

⁸A. F. Voter, Phys. Rev. B **57**, R13985 (1998).

⁹G. H. Vineyard, J. Phys. Chem. Solids **3**, 121 (1957).

¹⁰M. P. Allen and D. J. Tildesley, *Computer Simulation of Liquids* (Clarendon, Oxford, 1987).

¹¹H. Jónsson, G. Mills, and K. W. Jacobsen, in *Classical and Quantum Dynamics in Condensed Phase Simulations*, edited by B. J. Berne, G. Ciccotti, and D. F. Coker (World Scientific, Singapore, 1998).

¹²A. F. Voter, in *Modeling of Optical Thin Films*, edited by M. R. Jacobson, Proc. SPIE **821**, 214 (1987).

¹³G. DeLorenzi and G. Jacucci, Phys. Rev. B **33**, 1993 (1986).

¹⁴C. L. Liu and J. B. Adams, Surf. Sci. **268**, 73 (1992).

¹⁵Z. P. Shi, Z. Zhang, A. K. Swan, and J. F. Wendelken, Phys. Rev. Lett. **76**, 4927 (1996).