Long time scale kinetic Monte Carlo simulations without lattice approximation and predefined event table

Graeme Henkelman
Department of Chemistry, University of Washington, Seattle, Washington 98195-1700

Hannes Jónsson
Department of Chemistry, University of Washington, Seattle, Washington 98195-1700 and Faculty of Science, VR-II, University of Iceland, 107 Reykjavík, Iceland

We present a method for carrying out long time scale dynamics simulations within the harmonic transition state theory approximation. For each state of the system, characterized by a local minimum on the potential energy surface, multiple searches for saddle points are carried out using random initial directions. The dimer method is used for the saddle point searches and the rate for each transition mechanism is estimated using harmonic transition state theory. Transitions are selected and the clock advanced according to the kinetic Monte Carlo algorithm. Unlike traditional applications of kinetic Monte Carlo, the atoms are not assumed to sit on lattice sites and a list of all possible transitions need not be specified beforehand. Rather, the relevant transitions are found on the fly during the simulation. A multiple time scale simulation of Al(100) crystal growth is presented where the deposition event, occurring on the time scale of picoseconds, is simulated by ordinary classical dynamics, but the time interval in between deposition events, on the order of milliseconds, is simulated by the long time scale algorithm. The Al(100) surface is found to grow remarkably smooth, even at 30 K because of concerted displacements of multiple atoms with significantly lower activation energy than adatom diffusion on the flat terrace. © *2001 American Institute of Physics.* [DOI: 10.1063/1.1415500]

I. INTRODUCTION

A common problem in theoretical chemistry, condensed matter physics and materials science is the calculation of the time evolution of an atomic scale system where, for example, chemical reactions and/or diffusion occur. In either case, the configuration of atoms is changed in some way as time evolves. The interaction between the atoms can be obtained from an (approximate) solution of the Schrödinger equation describing the electrons, or from a potential energy function determined in some empirical way. Most often, it is sufficient to treat the motion of the atoms using classical mechanics. Quantum mechanical effects in the motion of atoms are important only in exceptional cases. Since the classical equations of motion can easily be solved numerically the simulation of dynamical evolution is, therefore, in principle quite simple. The problem is that the transitions of interest are typically many orders of magnitude slower than vibrations of the atoms, so a direct simulation of the classical dynamics ends up being of little use. This ''rare event'' problem is best illustrated by an example. A typical, low activation energy for a chemical reaction or diffusion event is 0.5 eV. Such an event can occur thousands of times per second at room temperature and would typically be important in the time evolution of the system. But, the atoms vibrate on the order of 10^{10} times before a fluctuation of thermal energy occurs in the right degree of freedom for a transition to take place. A direct classical dynamics simulation which necessarily has to faithfully track all this vibrational motion would take thousands of years of computer calculations on the fastest present day computer before a single transition can be expected to occur! It is clear that meaningful studies of chemical reactions and/or diffusion cannot be carried out by simply simulating the classical dynamics of the atoms. It is essential to simulate the system on a much longer time scale. This time scale problem is one of the important challenges in computational research on atomic scale systems.

Fortunately, the separation of time scales between vibrations and transitions makes it possible to use a statistical approach to estimate transition rates. If a bottleneck region through which the system must pass in order to make the transition can be identified, the so-called transition state, then transition state theory (TST) (Refs. 1–7) can be used to calculate the average amount of time the system will spend in a given state. Short time dynamical trajectories can subsequently be used both to correct for the approximations made in TST and to identify the possible final states of transitions. In general, the free energy of the transition state needs to be evaluated to estimate the rate. But, for solid systems, where the atoms are relatively tightly held in place by their neighbors, it is often possible to assume that the transition state can be characterized by a few saddle points on the potential energy rim surrounding the initial state basin and that the partition function of the system near each saddle point and near the energy minimum can be approximated by a product of harmonic partition functions. In this case, TST simplifies to the harmonic transition state theory (hTST) and the rate of escape, k, through each of the saddle point regions can be

related to properties of the initial state energy minimum and the saddle point[8,9] as

$$k^{\mathrm{hTST}} = \frac{\Pi_i^{3N} \nu_i^{\mathrm{init}}}{\Pi_i^{3N-1} \nu_i^{\ddagger}} e^{-(E^{\ddagger} - E^{\mathrm{init}})/k_B T}. \qquad (1)$$

Here, $E^{\ddagger}$ is the energy of the saddle point, E^{init} is the energy of the local minimum corresponding to the initial state, and the ν_i are the corresponding normal mode frequencies. The symbol $\ddagger$ refers to the saddle point. All the quantities can be evaluated directly from the potential energy surface without dynamical calculations. Entropic and thermal effects are included through the harmonic partition functions.

With the use of TST or hTST, a long time scale simulation consists of identifying states of the system and finding the mechanism and rate of transitions from a current state to a new state. The key thing is to find the relevant mechanism and not just assume a mechanism. Often, preconceived notions of the transition mechanism have turned out to be incorrect. One example is the mechanism of adatom diffusion on Al(100) where a two atom concerted displacement process turned out to have a lower barrier than the simple hop mechanism.[10] When hTST is used, the most challenging part of the calculation is the search for the low lying saddle points without knowledge of the possible final states. A typical simulation system includes a hundred atoms or more, which means that saddle points need to be found in a space with at least several hundred degrees of freedom. The large number of degrees of freedom makes this a challenging problem. We will now review briefly various approaches that have been taken to address this issue.

A. Bias potentials

One approach is to decrease the probability that the system is found in an initial state by adding a repulsive potential energy, a bias potential, to the actual interaction potential in such a way as to increase the probability of finding the system at a transition state.[11] This approach is more powerful if the bias potential can be constructed to increase the potential energy of minima without altering the potential of unknown transition states so that the dynamics of the system can be simulated on the modified potential.[12–17] The most accurate and efficient formulation of such an approach is the hyperdynamics method developed by Voter.[13,14] It is important to make sure the bias potential vanishes at the transition state. Within the hTST approximation, this means that the bias potential needs to vanish at the potential energy rim near the relevant saddle points. Voter's bias potential is designed to accomplish this even if it exceeds the saddle point energy within the initial state energy basin. It is, in fact, very important to allow the bias potential to greatly exceed the saddle point energy within the basin in order to get any appreciable acceleration of transitions in a system with many degrees of freedom. We will illustrate this with an example.

One might expect that just a slight filling of the potential well, reducing the energy difference between the bottom of the well and the saddle point, could lead to significant acceleration. This is true when the system has only very few degrees of freedom. But, each degree of freedom brings in $\frac{1}{2} kT$ of kinetic energy and a system with enough degrees of free-
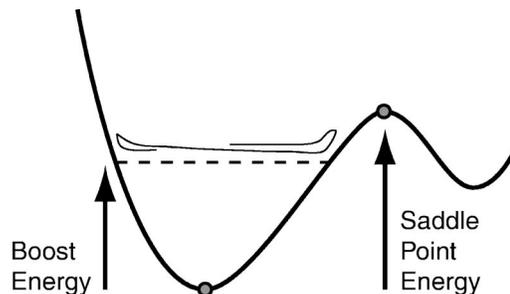

FIG. 1. A simple bias potential has been constructed to accelerate the dynamics of a system. The true potential is replaced by a constant equal to the "boost energy" whenever the true potential drops below the boost energy. The boost energy has to be set less than the saddle point energy so as not to alter the transition state. A stylized trajectory is shown for illustration. With the bias potential, the potential energy of the system never drops below the boost energy.

dom almost always has more kinetic energy than the saddle point energy. It is still a rare event to find enough energy focused in the right degree of freedom to bring the system through the saddle point region. The effect of dimensionality is best illustrated by applying a bias potential that leads to a flat-bottom potential. This method has, in fact, been proposed as a way of accelerating dynamics simulations.[15,16] When the true potential energy of the system is less than a certain value, the boost energy, the system evolves on a constant potential energy surface equal to the boost energy, but when the true potential energy of the system is higher than the boost energy, the system evolves on the unbiased potential surface. This type of bias potential is illustrated in Fig. 1. Results of calculations using this approach on a two-dimensional potential surface are shown in Fig. 2. The functional form and parameters of the potential function are given in the Appendix. In these simulations, the time step was adjusted when the system went from the biased to the unbiased regions, and vice versa, to ensure that energy was conserved, despite the discontinuity in the force. The escape rate for the system was calculated by counting the number of times the system left the basin shown in the inset in a fixed amount of simulation time. The figure shows how the escape rate varies with increasing boost energy. As the boost energy is raised up to the saddle point energy (gray line), the calculated escape rate remains constant and agrees well with the exact rate, k^{exact}, calculated by Voter.[13] The factor by which the dynamics are accelerated as compared with direct classical dynamics, the so-called boost factor,[13] is 1400 when the boost energy equals the saddle point energy. Remarkably, the calculated escape rate is still valid for this aggressive bias potential. As soon as the boost energy is increased over the saddle point, however, incorrect results are obtained for the calculated escape rate.

While a boost factor of 1400 is quite respectable, the problem arises when the number of degrees of freedom increases. The probability of finding the potential energy of the system below a saddle point energy becomes vanishingly small. To demonstrate this, a potential energy surface for a system with ten degrees of freedom was constructed by adding up five of the two-dimensional potential surfaces described above (see the Appendix). The potential energy and

J. Chem. Phys., Vol. 115, No. 21, 1 December 2001

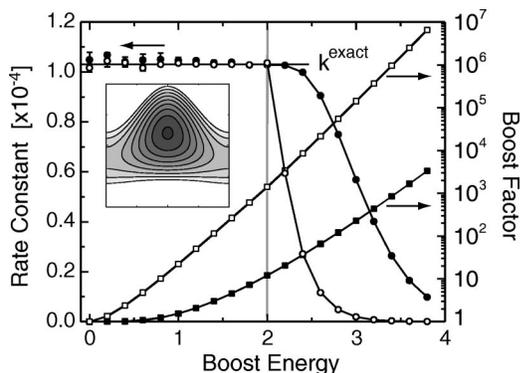Kinetic Monte Carlo simulations   9659

FIG. 2. The escape rate from a two-dimensional potential well (shown in the inset) and a ten-dimensional well (see Appendix) calculated with accelerated dynamics using the bias potential illustrated in Fig. 1. The acceleration obtained, i.e., effective time in the evolution of the system divided by the time of the simulated, boosted trajectory is referred to as the boost factor. The escape rate for $k_B T = 0.2$ is shown with open circles for the two-dimensional system and filled circles for the ten-dimensional system. The calculated rate is the same as the exact value, k^{exact}, for boost energy up to the saddle point energy of 2.0 (gray vertical line). When the boost energy is increased above the saddle point energy, the calculated escape rate drops rapidly below the exact value. The boost factor is also shown with open squares for two-dimensional system and filled squares for ten-dimensional system. With the maximum allowable boost energy, equal to the saddle point energy, the bias potential is accelerating the dynamics by a factor of 1400 in the two-dimensional system, but only by a factor of 12 in the ten-dimensional system. This illustrates the poor scaling of the boost factor with the number of degrees of freedom when this kind of bias potential is used. In the higher-dimensional system, the total potential energy is larger than the saddle point energy most of the time so the system spends little time in the region of constant potential, resulting in small boost.

the kinetic energy of this ten-dimensional system is on average five times higher than that of the two-dimensional system, but the activation energy, i.e., the saddle point energy minus the minimum energy, is the same. The calculated escape rate as a function of the boost energy for this larger system is also shown in Fig. 2 along with the boost factor. The higher-dimensional system spends less time in the region where the potential energy equals the bias and so the boost factor is smaller than for the two-dimensional case. The maximum boost factor is only 12, and this is still a very small system for practical applications. For a 20-dimensional system, the maximum boost factor is less than 2. This illustrates that a flat bias potential can only be effective in trivially small systems and that an effective bias potential must greatly exceed the saddle point energy in the interior of the potential energy basin but then drop to zero at the potential energy rim.

Voter's bias potentials are efficient in that the boost factor scales well with system size. However, the evaluation of his bias potentials is complicated and complex potential surfaces with small ripples reduce their effectiveness. Possibly, simpler bias potentials can be constructed that still scale well with dimensionality—that remains to be seen. Recently, a bias potential based on energy per atom has been proposed.[17] Here a strong bias is applied except in regions where the energy of any one of the atoms exceeds a predefined value. Although the energy per atom is not a well defined quantity for complicated potential functions (for example, interaction between water molecules that include induction) and *ab ini-*
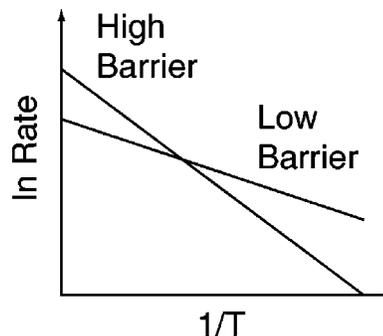

FIG. 3. The temperature dependence of the rate of two different transitions from the same initial state is shown. The transition with the lower energy barrier corresponds to the line with smaller slope. A typical situation is shown in which the higher barrier process has a higher prefactor (intercept with vertical axis). At low temperature, on the right side of the graph, the rate of the low barrier process will be higher than that of the high barrier process. At high temperature, the rate of the higher energy, higher entropy process becomes larger. This illustrates why a system cannot simply be heated to use a short classical dynamics simulation to find out which events would occur at low temperature on a longer time scale.

tio calculations, it can be defined for some simple potential functions. This method, however, seems to be applicable only to certain types of processes and simple systems. For example, in the Al/Al(100) system described in Sec. III, the lowest energy process for adatom diffusion is via a concerted displacement process involving two atoms. In this process, the energy of the adatom does not increase from the minimum to the saddle point, but rather drops as the adatom moves into the surface. Surface and bulk atoms which are more highly coordinated than adatoms and have lower energy should not be subject to the same boost energy as the adatom. A different bias potential would, therefore, need to be applied to atoms in different environment and during a concerted displacement process the bias would need to be changed continuously as the environment of the atoms changes. This makes the simulation of all but the simplest systems impractical with this method. The point is that constructing bias potentials of general applicability is nontrivial.

B. Heating

Another, perhaps simpler approach for identifying transitions is to increase the rate of rare events by simply heating the system. If the atoms have more energy, they will more likely undergo transitions. One should, however, not expect the favored transition mechanism to be the same at the higher temperature. This situation is illustrated in Fig. 3. The temperature dependence of the rate of two possible processes the system can undergo from a given initial state is shown. One of the processes has a low activation energy and small prefactor while the other has a high activation energy and large prefactor. At low temperature, the low barrier process will have a higher rate and dominate the dynamics. At high temperature, entropy becomes more important and the process with higher prefactor dominates even though the energy barrier is higher. An even more serious problem arises when the thermodynamic state of the system changes as it is heated up high enough to make transitions observable in a short, classical dynamics simulation. An example of this is diffu-

sion of admolecules on an ice surface. If an ice slab is heated up to a temperature at which the diffusion events occur frequently enough, the surface melts and the diffusion mechanism becomes quite different from the low temperature, long time scale diffusion mechanism.

High temperature dynamics can, however, in favorable cases be used to search for the relevant mechanism if many searches are carried out. Repeated simulations from the same initial state can be used to identify several possible transition mechanisms. Within the hTST the task is then to find the activation energy and prefactor for each mechanism and predict which one(s) would be relevant at the lower temperature of interest. Transitions can be detected from the high temperature dynamics by periodic minimization to the nearest local minimum on the energy surface and the nudged elastic band method[18–20] can then be used to find the minimum energy path connecting the given initial state with the final state found from the high temperature dynamics. The maximum energy on the minimum energy path gives the activation energy. This approach was used by Sørensen et al. to study the mechanism of contact formation as a metal tip is brought up to a metal surface on a laboratory time scale at room temperature.[21] While earlier classical dynamics simulations had seen a sudden jump-to-contact when the tip-surface distance became very small, the method described above could be used to show that a sequence of thermally activated processes would occur at a larger tip-surface distance on the time scale of the experiments.

By assuming the system is harmonic also at the high temperature and given a minimum value for the prefactors, Sørensen and Voter formulated a concise prescription for the length of time for which the high temperature dynamics need to be run in order to safely advance the system to a new state.[22] This eliminates the need to find all possible transition mechanisms. In their method, the escape time calculated from hTST at the low temperature is recorded for each transition observed in the high temperature dynamics. After the simulation has been run long enough at the high temperature, the system is advanced by the first process that would have been observed at the low temperature. They refer to the method as temperature accelerated dynamics (TAD). An application of this method to metal crystal growth has been presented.[23]

In this work we present a method that is quite different from the ones described above. Classical dynamics are not used in any form but instead the system is pushed up the potential energy surface using the so-called dimer method[24] to find saddle points. The rate of transitions through the vicinity of each saddle point is then estimated within hTST and a kinetic Monte Carlo method used to simulate the evolution of the system over long time scales.[25] This method is easy to implement and, compared to existing methods, may require less computational time for small systems. We expect it to be applicable to dynamics in ice, for example, a system where a high temperature search for transition mechanisms is likely to become problematic. While the use of hTST means that it can only be applied to solids, we expect it to be applicable to glasses and other amorphous solids.

II. THE LONG TIME DYNAMICS METHOD

Simulations of systems over long time scale are carried out by combining kinetic Monte Carlo (KMC) and saddle point searches. KMC relies on knowing the rate and mechanism of all the relevant transitions from a given initial state. Within hTST this corresponds to finding all the low energy saddle points on the rim of the potential energy basin corresponding to the initial state. The dimer method is used to search for these saddle points.

A. The dimer method

The dimer method has been described in detail in a previous publication.[24] Only a brief overview will be given here. The dimer is made up of two images (replicas) of the system. These images are separated in space by a finite distance displacement along a vector $\hat{N}$. For an empirical potential this can be small. Here we have used 0.005 Å.

There are two parts to each dimer move. The first part is dimer rotation. The lowest energy orientation of the dimer is along the lowest curvature mode. If the dimer is free to rotate, the forces acting on the two images will pull the dimer to the lowest curvature mode. This is done by defining a rotational force which is the difference in the force on the two images. Minimizing the energy of the dimer with respect to this rotational force aligns the dimer with the lowest curvature mode (this feature was used by Voter in his construction of bias potentials in hyperdynamics[14]). A modified Newton's method can be used to make this rotation efficient.[24] An important aspect of the dimer method is that it only requires the first derivative of the energy, not the second derivatives.

The second part of the algorithm is translation of the dimer. A first order saddle point on a potential surface is at a maximum along the lowest curvature direction and a minimum in all other directions. In order to converge to a saddle point, the dimer is moved up the potential along the lowest curvature mode, and down the potential in all other directions. This is done by defining an effective force on the dimer, in which the true force due to the potential acting at the center of the dimer has the component along the dimer inverted. Minimizing with respect to this effective force moves the dimer to a saddle point. With empirical potentials, minimization using the conjugate gradient method works well.

It is not necessary to fully converge the dimer orientation at each translational step. We have found it most efficient to pick a certain tolerance for the rotational force. This sets how precisely the dimer is oriented along the lowest curvature mode before it is moved. The dimer is rotated at least once and possibly a few times until the rotational force is less than the specified tolerance. In these calculations the maximum rotational force was set at 5 meV/Å.

For the Al adatom on an Al(100) surface, the dimer requires on average 400 force evaluations to converge on a saddle point.[24] A very important feature of the dimer method is the slow increase in the number of required force evaluations as more degrees of freedom are added to the system. This is to be contrasted with mode following approaches,[26–32] where the matrix of second derivatives is

constructed and diagonalized at each step in the search for a saddle point. Other methods for finding saddle points using only first derivatives of the energy have been proposed,[33,34] but it is unclear how their efficiency compares with the dimer method.

In the long time scale calculations presented here, several dimer searches were launched for each state visited. Dimer searches were started from configurations close to the potential minimum. The easiest way to choose a starting position is to make a random displacement away from the minimum. If all the atoms in the system are included in the random displacement, the dimer search can be biased towards finding higher energy processes involving many atoms. We have found that it is better to displace only atoms within a local region. For each dimer search in the Al(100) ripening calculations (described in Sec. IV), an undercoordinated surface atom was chosen at random to be the center of the initial, local displacement. The displacements had a Gaussian distribution with a mean of 0.2 Å in each degree of freedom. The region consisted of the central atom and its first and second neighbors. We have found that a continuous distribution in the magnitude of the displacements is better than a fixed magnitude because it increases the variety of saddle points found. In principle, a scheme like this can eventually find all saddle points around a minimum simply because the starting points of the dimer search can be at any point in space. In practice, there is no guarantee that a complete list of saddle points can be found in a reasonable amount of time. Our experience from studies of adatom diffusion processes[24] and island rearrangement processes[35] is that a dimer search started in a random direction most likely finds one of the low energy saddle points. It is important to keep in mind, though, that any calculation based on a finite number of dimer searches is not guaranteed to find the most relevant transition mechanism.

B. Kinetic Monte Carlo

Kinetic Monte Carlo[36–40] is a powerful method that can be used to extend the time scale of simulations far beyond the vibrational time scale. If a list of possible transitions for a given initial state is available, a random number can be used to choose one of the processes and evolve the system to a new state. The probability of choosing a certain transition is proportional to its rate, r_i. On average, the amount of time that would have elapsed in order for this process to occur is

$$\Delta t = \frac{1}{\Sigma r_i}, \qquad (2)$$

which is independent of the chosen transition. It may also be important to include the appropriate distribution of escape times. For random uncorrelated processes, this is a Poisson distribution. If μ is a random number from 0 to 1, the elapsed time for a particular transition is given by

$$\Delta t = \frac{-\ln \mu}{\Sigma r_i}. \qquad (3)$$

The system is then advanced to the final state of the chosen transition and the process is repeated.
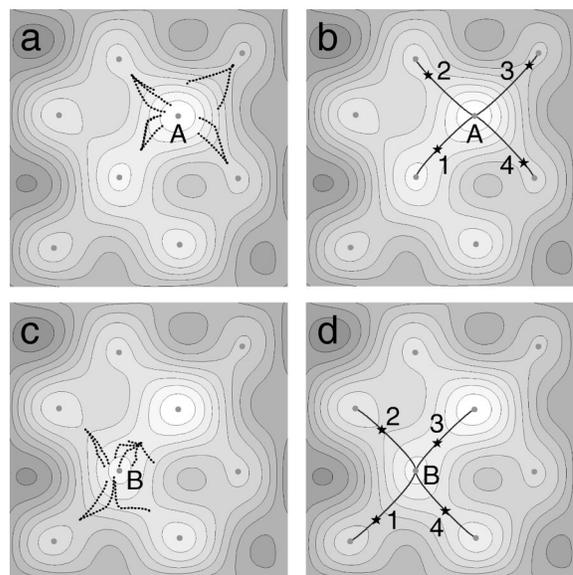

FIG. 4. Application of the long time scale simulation method to a model two-dimensional potential surface. The system is initially in state A. (a) Ten dimer searches are started from random positions around the minimum. They converge on four distinct saddle points (two of the searches practically overlap). (b) The system is then made to slide down the minimum energy path (gray lines) on either side of the saddle points which are indicated with ∗. Here, all four saddle points have a minimum energy path starting at the initial state minimum A, but this does not have to be the case. The rate of each process is then calculated using harmonic transition state theory. A process is chosen at random using the kinetic Monte Carlo algorithm. In this case, process 1 gets chosen. The system is moved to the final state of this process, to minimum B. (c) Dimer searches are run from the new minimum, again four distinct saddle points are found. (d) Minimum energy paths are traced out, and the process repeated.

In a traditional kinetic Monte Carlo simulation, all transitions that can ever occur in the system, along with their rates, must be known before the simulation starts. Ideally, the rates are estimated from some description of the atomic interactions,[41] such as an empirical interaction potential or *ab initio* calculations, but the problem is to know in advance the mechanism of the relevant transitions for each possible configuration of the atoms. The requirement of knowing and tabulating the relevant transitions ahead of time limits the method to simple systems. Systems which can undergo complicated transitions involving several atoms, such as the aluminum system described in Sec. III, or where atoms do not sit at lattice sites are extremely difficult to model with traditional KMC. The approach presented here relaxes these limitations.

C. Combined dimer and kinetic Monte Carlo

If the dimer method is used to find possible transitions, there is little limitation on the complexity in terms of the number of atoms or the spatial extent of the transition. Also, the energy barriers are found at each state the system evolves to and do not need to be known before the calculation is started. Furthermore, the atoms do not need to be mapped onto a lattice and it is not necessary to anticipate all possible states of the system.

The method is illustrated for a two-dimensional model potential in Fig. 4. The system is started at a potential mini-

mum, A. When a new state is visited, a swarm of dimer searches is sent out from the vicinity of the potential energy minimum. For the calculations described in Sec. IV, either 25 or 50 dimer searches were used. In this example, ten random displacements from the position of the minimum were chosen as starting points of dimer searches. Figure 4(a) shows the path of the ten dimer searches. In this calculation, four distinct saddle points ($\star$) were found. The system is then quenched on either side of each saddle point in order to verify that it lies on a minimum energy path (shown in gray) from the given initial state minimum. All the saddle points found in this case did connect back to the initial minimum. If not, the saddle point is discarded from the list of possible transitions. In the same way as described in Sec. II B, a transition is chosen from the list, the system is advanced to the final state of that transition, and the time interval associated with the transition is added to the accumulated time. In this example, transition 1, which corresponds to the lowest barrier was chosen, and the system was advanced to state B. From the new minimum the process is repeated. New dimer searches are sent out [Fig. 4(c)], the saddle points verified [Fig. 4(d)] and then one is chosen for the next transition.

III. THE Al(100) SURFACE

We have chosen dynamics of Al adatoms on an Al(100) surface as a test problem for the long time scale simulation method. There are two reasons for that. First, an accurate embedded atom potential of the Voter and Chen form exists for aluminum.[42] Second, Al(100) is a rich system because there are many different transitions with a rather low energy barrier even for just a single adatom on the Al(100) surface. We have previously studied this system extensively with the dimer method.[24] The four lowest energy processes found are shown in Fig. 5. A particularly interesting aspect of the Al(100) system is that a concerted displacement process has a lower energy barrier than the direct hop process. This was shown by Feibelman with density functional theory calculations.[10] In our simulations, the system consists of a six layer slab with 50 atoms per layer. The bottom two layers are held frozen and the top surface is left open to vacuum. The dimer calculations revealed 60 different transitions for adatom diffusion in 1000 dimer searches.[24] On average the low energy transitions, the concerted displacement involving two and four atoms, and the hop, were found three quarters of the time. One quarter of the time a wide collection of higher energy processes were found. This is a very important result because it indicates that a dimer search started in a random direction has high probability of finding a saddle point for one of the low barrier transitions.

In the KMC scheme, it is assumed that all relevant transitions have been found. Transitions that have a high activation energy and/or low prefactor are unlikely to occur and are typically not important. Given that the two-atom concerted displacement process is found a quarter of the time, there is an 80% certainty of finding all the four transitions that are equivalent by symmetry in 50 dimer searches. The fact that all of the equivalent processes may not be found is not so serious since only one gets chosen at random in the KMC simulation. The important thing is that a representative
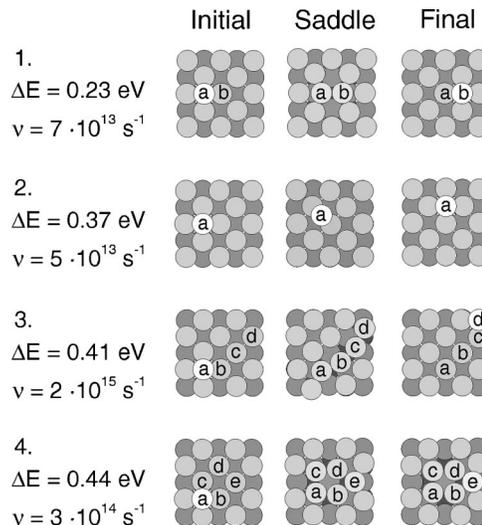

FIG. 5. The four lowest energy transitions found with the dimer method for the diffusion of an Al adatom on Al(100). For each transition, the initial state, the saddle point configuration, and the final state are shown. Atoms are shaded by depth and the atoms that move the most in each transition are labeled. The energy of each transition is given in eV. The lowest energy transition is the two-atom concerted displacement (1). The hop (2) has similar activation energy as a four-atom (3) concerted displacement process and a complex process (4) involving local hex reconstruction. Note the large prefactor for the multiatom processes. Because Al adatoms can so easily displace atoms in the surface, a growing Al(100) surface can undergo a great variety of transition which would be hard to find by guesswork.

sample of the transitions is found over the relevant range of activation energy and that no category of transitions is excluded. The error mainly shows up in the time scale of the simulation which is dominated by the low barrier transitions. If only half the low barrier transitions are found, the simulation clock will be running two times too fast. Since the challenge of reaching long time scale is a matter of spanning several orders of magnitude, a factor of 2 is typically not a serious issue. With a modest number of dimer searches, the method can give a good qualitative idea of how the system will behave at low temperature over long time scale. For an accurate simulation, it is clear that many dimer searches need to be carried out. Fortunately, the different searches can easily be carried out in parallel on loosely connected cluster of computers.

IV. APPLICATION TO ISLAND RIPENING

The results of a kinetic Monte Carlo simulation coupled with dimer searches is shown in Fig. 6. Initially 20 atoms, a coverage of 0.4, were randomly deposited on the Al(100) surface consisting of 50 atoms using classical dynamics. Then the system was quenched to the nearest local energy minimum. This configuration is shown in the first panel ($n=1$). After that, the time evolution of the system at 300 K was simulated. On average, 17 distinct processes were found when 25 dimer searches were carried out. When 50 dimer searches were used, this number rose to 23, indicating that some, but not too many, transitions were missed with 25 searches. During the first ten transitions, the adatoms diffused via the concerted displacement or hop processes to form clusters. After 344 transitions, a large island has

J. Chem. Phys., Vol. 115, No. 21, 1 December 2001
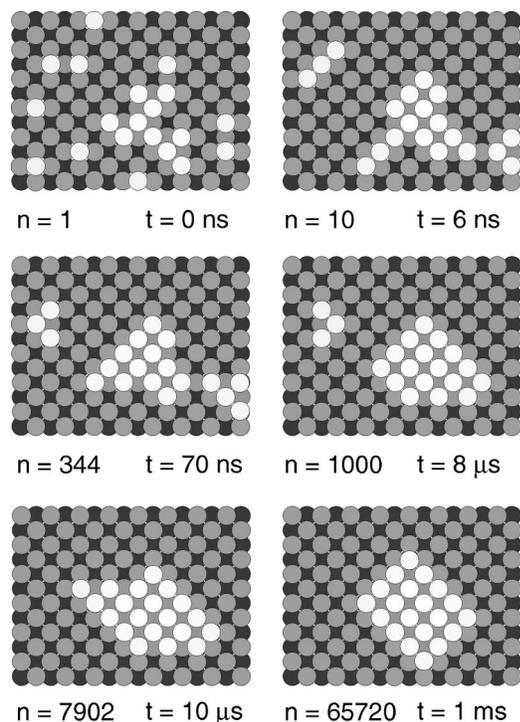
Kinetic Monte Carlo simulations   9663

FIG. 6. Snapshots from a simulation of ripening of Al adatoms on an Al(100) surface. Initially 20 atoms were deposited at random on the surface, a coverage of 0.4. After 6 ns (10 transitions) all the adatoms have merged to form clusters. After 70 ns (344 transitions) a large, compact island has formed, but there are still two outlying islands. The trimer in the upper left has four possible rearrangements with a low activation barrier. Many of the 344 transitions correspond to these rearrangements but old configurations are stored during the simulation so the repeat processes do not require new dimer searches. After 8 μs (1000 transitions) the large island has taken a more compact shape and merged with one of the smaller islands. Finally after 10 μs (7902 transitions) a single large island has formed, and at 1 ms (65 720 transitions) the island has rotated and taken its most compact shape.
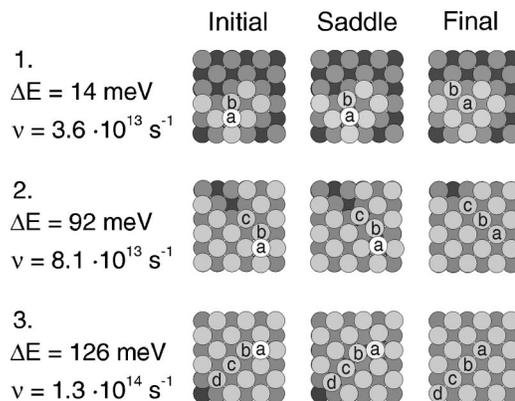

FIG. 7. Three processes observed during a simulation of Al atom deposition on Al(100) at 100 K. The deposition rate was one monolayer per millisecond. Ten monolayers were grown, each consisting of 32 atoms. The processes involve descent of an adatom into a lower layer. Most often an adatom descends by a two-atom concerted displacement process as shown in (1). But, processes involving concerted displacement of three and four atoms where the adatom starts two or three sites away from the island edge were also often observed in the simulations. While the activation energy of these long range processes is higher, it is still small compared with the 0.23 eV activation energy for adatom diffusion on the flat (100) terrace, and the prefactor is larger. The remarkable ease by which an adatom can descend from atop an island on Al(100) leads to layer-by-layer growth in this small system even at temperature as low as 30 K.

formed, but two smaller islands also exist. After 1000 transitions, one of the small clusters has merged with the large island which has taken a more compact shape. Finally after 7902 transitions, a single compact island is formed, and at 1 ms, (65 720 transitions) the island reaches its lowest energy shape. The mechanism for atoms crossing a corner or filling in a kink site invariably involves a two-atom displacement process with a substrate atom.

Of the 100 000 transitions in this simulation, only 300 brought the system to a new state. Many processes involved rearrangements of the three atom cluster or an adatom moving along the edge of an island. After each transition, the new state is compared to a table of all the old configurations. If the new state has been seen before, no dimer searches are performed. Rather, a new transition is chosen from the old list of processes. In this way, low barriers which are seen frequently do not contribute significantly to computational time. The total computer time for the simulation was approximately one week on a PC.

V. DEPOSITION AND SURFACE GROWTH

Multiple time scale simulations of Al(100) crystal growth were carried out by combining classical dynamics of deposition events with the long time scale simulation of the time intervals between deposition events. Interesting processes can take place during the deposition event,[43,44] but the excess energy released as the incoming atom binds to the surface dissipates quickly, in 1 or 2 ps.[43,44] This period of time can easily be simulated by direct classical dynamics. After the "hot-spot" has cooled down and the system has thermalized, transition state theory can be applied to the activated diffusion processes. The long time scale calculation was done in the same way as described in Sec. IV, but at each step a deposition process was added to the table of possible events. A smaller system was used than in the ripening simulations, the cell consisted of 32 atoms per layer. Atoms were deposited at a rate of one monolayer per millisecond (32×10^3 s^{-1}). A deposition process was simulated by placing an aluminum atom at a random position 10 Å above the surface and giving it an initial velocity towards the surface characteristic of the simulation temperature. Classical dynamics were run until the component of the velocity of the deposited atom perpendicular to the surface had changed sign twice. At this point the system was considered to be equilibrated and the KMC simulation took over. In this way, the short deposition time scale was simulated accurately as well as the longer time scale of thermally activated diffusion processes.

Growth at 100 K was first simulated. A total of 9.5 layers were deposited in 9.8 ms. This took 500 events in the KMC simulation of which 302 were deposition events. An important question in crystal growth is how smoothly the surface grows, i.e., how well a layer gets completed before the next layers starts forming. A key issue is how an adatom that lands on top of an island manages to descend down to the more stable edge site. Typically, an adatom descends by a two-atom concerted displacement process, as is shown in Fig. 7, but the simulation also revealed interesting three- and
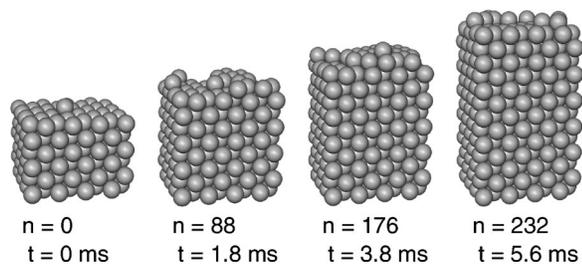
FIG. 8. Seven monolayers of Al were grown at 30 K on the Al(100) substrate. The deposition rate was one monolayer per millisecond. Even at this low temperature, the layers grown are free of defects and the surface is smooth.

FIG. 10. The average barrier height of events observed during the Al(100) growth simulations drops with temperature, as expected. In each simulation the deposition rate is one monolayer per millisecond. As the temperature is lowered, it gets harder and harder to surmount high barriers within the time interval between deposition events. At 100 K there is on average one surface diffusion process for every deposition event. At 30 K this ratio drops to one in eight. Even so, the Al(100) surface still grows smoothly at 30 K because as multiatom cascade events such as those shown in Fig. 9. Note that all the processes that occur in these growth simulations have significantly lower activation energy than the adatom diffusion on a flat terrace (indicated by the dashed vertical line at 0.23 eV).

four-atom descent processes where the adatom started out two or three sites away from the edge. The transitions involving three and four atoms have somewhat larger activation barrier, 92 meV and 126 meV, but were frequently observed in the simulation. The ease of these types of descent processes in the Al(100) system make it easy for aluminum to grow layer-by-layer, especially for this small 32 atom per layer cell.

The growth simulation was then repeated for a lower temperature, 30 K. A total of 7.3 layers were grown in 5.6 ms. The KMC simulation involved 264 events of which 232 were deposition events. Snapshots of the system during the growth simulation are shown in Fig. 8. The range of possible surface processes is now much more limited. The long range descent events observed at 100 K were not observed at 30 K. Since individual adatoms were less likely to descend, rougher surface morphology could form, but, after structures with several undercoordinated atoms had built up, some even not sitting at well-defined lattice sites, concerted processes involving several atoms would typically occur and smoothen out the surface. Two of these transitions are shown in Fig. 9.
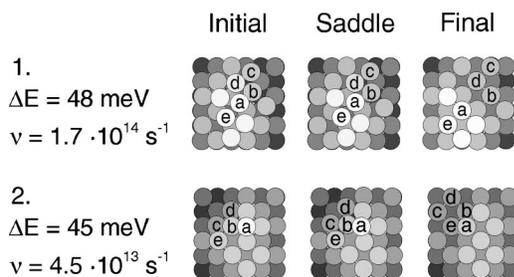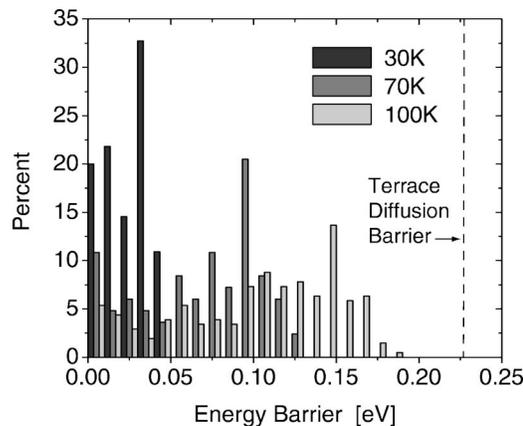

FIG. 9. Two processes observed during the 30 K deposition simulation of Al atoms on Al(100) shown in Fig. 8. At this low temperature only very low energy processes can occur. The surface tends to grow rough temporarily and regions with islands stacked on top of islands form. The initial state of the first process resulted from classical dynamics simulations of the deposition of atom (a) on top of a four atom island consisting of atoms (b), (c), (d), and (e). The five atoms are not sitting at lattice sites. In a concerted, activated process all five atoms move to lattice sites, two in the top layer [(a) and (e)] and three in the lower layer [(b), (c), and (d)]. In the second process, five atoms slide in a concerted way down the side of an island stack. The adatom (a) in the third layer moves down to the second layer, while two atoms in the second layer, (e) and (b), move to the first layer and two atoms in the first layer get pushed into different sites in the first layer. The barrier for these two multiatom, sliding processes is remarkably low, less than 0.05 eV, and the prefactors are large. The processes are active at 30 K even though adatom diffusion on the flat (100) terrace, which has a barrier of 0.23 eV, is not active.

In the initial state of the first process, an atom has been deposited on an island consisting of four atoms resulting in a structure where five atoms are not sitting on lattice sites. Then, a transition occurs which moves all five atoms simultaneously into lattice sites, two in the top layer and the other three in a lower layer. In the other process shown in Fig. 9, a deposited atom lands on top of an island, starting a third incomplete layer. In the subsequent transition, five atoms slide down in a concerted way. The third layer atom ends up in the second layer, the two second layer atoms end up in the first layer, and two first layer atoms get pushed out into different sites. The energy barrier of these two five-atom displacement processes is only 48 meV and 45 meV. Evidently, when the surface becomes rough, low barrier cascading events where many atoms move into more stable sites become possible. Even at this low temperature, the surface morphology remains smooth, although the basic adatom diffusion process on the flat Al(100) terrace, which has an activation energy of 0.2 eV, is not active.

A simulation was also carried out at an intermediate temperature, 70 K. As the temperature was lowered from 100 K to 30 K, the average barrier height of observed transitions also dropped, as expected. This is shown in Fig. 10. At 100 K the system was able to cross barriers of 0.18 eV, but at 30 K the highest barrier that was overcome was less than 50 meV.

Animations of the growth and ripening simulations can be viewed on the web at http://eon.chem.washington.edu.

VI. EFFICIENCY

It is important to know how a simulation method scales with system size. To determine this, we can consider how the number of force evaluations changes when the system is doubled in size. Force evaluations are considered rather than total computational time, because the evaluation of the force
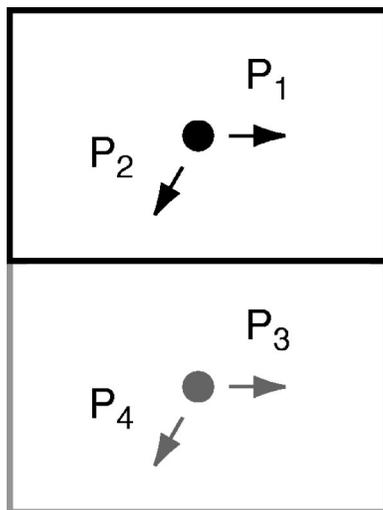
J. Chem. Phys., Vol. 115, No. 21, 1 December 2001

Kinetic Monte Carlo simulations   9665

FIG. 11. In a system of twice the size, there are on average twice as many possible transitions. If P_1 and P_2 are transitions in the original system, there are two more transitions, P_3 and P_4, possible in the larger system. In the simplest implementation of the simulation method, this means that the number of dimer searches increases linearly with system size. But, if the transitions are local, only a limited region of the system is effected by each transition and only a part of the system will require new dimer searches. Also, saddle points for transitions close to the transition region may only be slightly effected and previously found saddle point configurations can be reconverged with just a few dimer iterations. The number of dimer searches needed per transition does not, therefore, increase with system size. However, the time increment for each transition is only half as large in the larger system.

can scale differently depending upon the complexity of the interaction potential. In favorable situations, evaluation of the force scales linearly with system size.

Figure 11 shows an illustration of a system which has been doubled in size by joining two identical replicas of the system. In the larger system, there are twice as many possible transitions. In the crudest implementation of the simulation method, the number of dimer searches would then increase in proportion to the system size. If, however, the transitions are local, i.e., they only involve a small subset of the atoms, then large portions of the system will be unaffected by any one transition. Transitions that take place in these unaffected regions will not require repeated saddle point searches. Rather, only the transitions which are affected by the last transition need to be updated and new transitions need to be added only for the region where the last transition took place. With the dimer method, it is also relatively straightforward to reconverge a set of known saddle points which have changed only slightly because of a nearby transition. Therefore, if the transitions are local, the number of new dimer searches that need to be carried out after a transition has occurred will not change as a function of the system size.

Another aspect to the scaling is the time increment at each transition. When the system size is doubled, there are on average twice as many transitions possible. The time increment for each transition will then be half as long, as can be seen from Eq. (3). If the simulation is carried out in such a way as to cover a fixed time interval, then twice as many transitions need to be simulated in a system that is double in

size. The number of force evaluations required to simulate a fixed length of time, therefore, scales linearly with system size. These scaling arguments only apply to systems that are large enough that all relevant processes have been included.

This method for carrying out long time scale dynamics simulations is efficient enough for it to be implemented with first principles calculations of atomic interactions such as density functional theory. We have implemented the method in the VASP code.[45–48] More information about the implementation can be found on the web site http://ikazki01.chem.washington.edu/vasp/. We are currently applying the technique to the formation and break-up of boron clusters in silicon[49] and the diffusion of self-trapped excitons in silica glass.[50]

ACKNOWLEDGMENTS

APPENDIX: THE MODEL POTENTIALS

The bias potential algorithm described in the Introduction was first of all tested on the two-dimensional potential described by Voter,[13]

$$V(x,y) = \cos(2\pi x)(1+4y) + \tfrac{1}{2}(2\pi y)^2 + V_0. \quad (A1)$$

This potential is periodic in x. A contour plot is shown in the inset in Fig. 2. We chose to set $V_0 = -1.203$, so that the minimum potential is zero and occurs at the point $(x,y) = (k+0.5, 0.1013)$, where k is any integer. The saddle point has a potential $V_{SP} = 2$ and is located at $(k, -0.1013)$.

The ten-dimensional potential was constructed by adding up five of the two-dimensional potentials given by Eq. (A1),

$$\widetilde{V}(x,y) = \sum_{i=1}^{5} V(x_{2i-1}, x_{2i}). \quad (A2)$$

The results of the accelerated dynamics calculations where a repulsive bias potential is added to these potential functions so as to create a flat-bottom potential are shown in Fig. 2.

[1] H. Eyring, J. Chem. Phys. **3**, 107 (1935).
[2] E. Wigner, Trans. Faraday Soc. **34**, 29 (1938).
[3] J. C. Keck, Adv. Chem. Phys. **13**, 85 (1967).
[4] P. Pechukas, in *Dynamics of Molecular Collisions*, edited by W. Miller (Plenum, New York, 1976), Part B.
[5] D. Chandler, J. Chem. Phys. **68**, 2959 (1978).
[6] A. F. Voter and D. Doll, J. Chem. Phys. **80**, 5832 (1984).
[7] A. F. Voter and D. Doll, J. Chem. Phys. **82**, 80 (1985).
[8] C. Wert and C. Zener, Phys. Rev. **76**, 1169 (1949).
[9] G. H. Vineyard, J. Phys. Chem. Solids **3**, 145 (1957).
[10] P. J. Feibelman, Phys. Rev. Lett. **65**, 729 (1990).
[11] E. K. Grimmelman, J. C. Tully, and E. Helfand, J. Chem. Phys. **74**, 5300 (1981).
[12] H. Grübmuller, Phys. Rev. E **52**, 2893 (1995).
[13] A. F. Voter, J. Chem. Phys. **106**, 4665 (1997).
[14] A. F. Voter, Phys. Rev. Lett. **78**, 3908 (1997).
[15] M. M. Steiner, P. A. Genilloud, and J. W. Wilkins, Phys. Rev. B **57**, 10236 (1998).
[16] S. Pal and K. A. Fichthorn, Chem. Eng. J. **74**, 77 (1999).
[17] J. Wang, S. Pal, and K. A. Fichthorn, Phys. Rev. B **63**, 085403 (2001).
[18] H. 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).

[19] G. Henkelman and H. Jónsson, J. Chem. Phys. **113**, 9978 (2000).

[20] G. Henkelman, B. P. Uberuaga, and H. Jónsson, J. Chem. Phys. **113**, 9901 (2000).

[21] M. R. Sørensen, K. W. Jacobsen, and H. Jónsson, Phys. Rev. Lett. **77**, 5067 (1996).

[22] M. R. Sørensen and A. F. Voter, J. Chem. Phys. **112**, 9599 (2000).

[23] F. Montalenti, M. R. Sørensen, and A. F. Voter, Phys. Rev. Lett. **87**, 126101 (2001).

[24] G. Henkelman and H. Jónsson, J. Chem. Phys. **111**, 7010 (1999).

[25] G. Henkelman and H. Jónsson, Mater. Res. Soc. Symp. Proc. **677**, AA8.1 (2001).

[26] G. M. Crippen and H. A. Scheraga, Arch. Biochem. Biophys. **144**, 462 (1971).

[27] R. L. Hilderbrandt, Comput. Chem. (Oxford) **1**, 179 (1977).

[28] C. J. Cerjan and W. H. Miller, J. Chem. Phys. **75**, 2800 (1981).

[29] J. Simons, P. Jørgensen, H. Taylor, and J. Ozment, J. Phys. Chem. **87**, 2745 (1983).

[30] J. Baker, J. Comput. Chem. **9**, 465 (1988).

[31] D. J. Wales, J. Chem. Phys. **91**, 7002 (1989).

[32] J. Nichols, H. Taylor, P. Schmidt, and J. Simons, J. Chem. Phys. **92**, 340 (1990).

[33] R. Malek and N. Mousseau, Phys. Rev. E **62**, 7723 (2000).

[34] L. J. Munro and D. J. Wales, Phys. Rev. B **59**, 3969 (1999).

[35] G. Henkelman, G. Jóhannesson, and H. Jónsson, in *Progress on Theoretical Chemistry and Physics*, edited by S. Schwartz (Kluwer Academic, New York, 2000).

[36] A. B. Bortz, M. H. Kalos, and J. L. Lebowitz, J. Comput. Phys. **17**, 10 (1975).

[37] D. T. Gillespie, J. Comput. Phys. **22**, 403 (1976).

[38] D. T. Gillespie, J. Phys. **81**, 2340 (1977).

[39] D. T. Gillespie, J. Comput. Phys. **28**, 395 (1978).

[40] G. H. Gilmer, Science **208**, 335 (1980).

[41] A. F. Voter, Phys. Rev. B **34**, 6819 (1986).

[42] A. F. Voter and S. P. Chen, Mater. Res. Soc. Symp. Proc. **82**, 2384 (1987).

[43] M. Villarba and H. Jónsson, Phys. Rev. B **49**, 2208 (1994).

[44] H. Jónsson, Annu. Rev. Phys. Chem. **51**, 623 (2000).

[45] G. Kresse and J. Hafner, Phys. Rev. B **47**, 558 (1993).

[46] G. Kresse and J. Hafner, Phys. Rev. B **49**, 14251 (1994).

[47] G. Kresse and J. Furthmüller, Comput. Mater. Sci. **6**, 16 (1996).

[48] G. Kresse and J. Furthmüller, Phys. Rev. B **54**, 11169 (1996).

[49] G. Henkelman, B. P. Uberuaga, and H. Jónsson (unpublished).

[50] K. Tsemekhman, G. Henkelman, and H. Jónsson (unpublished).