Accelerated molecular dynamics simulation of thin-film growth with the bond-boost method

2009 J. Phys.: Condens. Matter 21 084212

(http://iopscience.iop.org/0953-8984/21/8/084212)

IOP Publishing Journal of Physics: Condensed Matter

J. Phys.: Condens. Matter **21** (2009) 084212 (9pp) doi:10.1088/0953-8984/21/8/084212

Accelerated molecular dynamics simulation of thin-film growth with the bond-boost method

Kristen A Fichthorn[1,2], Radu A Miron[2,3], Yushan Wang[1] and Yogesh Tiwary[1]

[1] Department of Chemical Engineering, The Pennsylvania State University, University Park, PA 16802, USA

[2] Department of Physics, The Pennsylvania State University, University Park, PA 16802, USA

E-mail: fichthorn@psu.edu

Received 8 July 2008, in final form 19 August 2008
Published 30 January 2009
Online at stacks.iop.org/JPhysCM/21/084212

Abstract

We review the bond-boost method for accelerated molecular dynamics (MD) simulation and we demonstrate its application to kinetic phenomena relevant to thin-film growth. To illustrate various aspects of the method, three case studies are presented. We first illustrate aspects of the bond-boost method in studies of the diffusion of Cu atoms on Cu(001). In these studies, Cu interactions are described using a semi-empirical embedded-atom method potential. We recently extended the bond-boost method to perform accelerated *ab initio* MD (AIMD) simulations and we present results from preliminary studies in which we applied the bond-boost method in AIMD to uncover diffusion mechanisms of Al adatoms on Al(110). Finally, a problem inherent to many rare-event simulation methods is the 'small-barrier problem', in which the system resides in a group of states connected by small energy barriers and separated from the rest of phase space by large barriers. We developed the state-bridging bond-boost method to address this problem and we discuss its application for studying the diffusion of Co clusters on Cu(001). We discuss the outlook for future applications of the bond-boost method in materials simulation.

(Some figures in this article are in colour only in the electronic version)

1. Introduction

A significant challenge in materials simulation is to conduct long-time and large length simulations of structural evolution, while accurately retaining atomic detail. Molecular dynamics (MD) simulations can provide accurate detail at the atomic scale. However, MD is not practical for simulating times or lengths much beyond the nanoscale. In many materials, dynamical evolution occurs through a series of 'rare events', in which the system spends a long time (compared to the vibrational timescale) in one potential-energy minimum before escaping and moving on to another. Transition-state theory (TST) [1] can predict the ensemble-averaged escape time for a system to progress from one minimum to another. With this information, dynamical evolution can be simulated as a series of long-time jumps based on the TST jump times between minima. This is the aim of kinetic Monte Carlo (KMC) simulations [2–6]. In principle, if a KMC simulation can incorporate all minima of a system and the TST rates of all possible long-time jumps between the minima, then this technique can reach macroscopic times and (possibly, with recent innovations in parallel KMC [7, 8]) lengths, while retaining the accuracy of MD. In practice, however, a significant limitation of KMC lies in determining all the important rate processes that govern dynamical evolution.

Traditionally, TST search algorithms have been employed to quantify the rates and mechanisms of rare events and a number of efficient algorithms have been proposed [9–14]. However, these algorithms often require knowledge of the final state, which is usually the unknown quantity that one would

[3] Present address: Software Design Ahnert GmbH, Arkonastrasse 45-49, 13189 Berlin, Germany.

like to find. Even if the final state is known, most TST search algorithms require an initial guess for the mechanism. This can limit convergence to pathways near the initial guess, which may not reflect the most important mechanism. TST search algorithms have been developed based only on knowledge of the initial state [15]; however, these do not inherently detect rate processes with the correct canonical frequency and many searches, along with auxiliary calculations, may be necessary to obtain a correct description of the dynamics. If it were possible to conduct a very long MD simulation, then the kinetic processes responsible for dynamical evolution would arise naturally, with the correct canonical frequency based on the phase space dictated by the underlying potential surface. The philosophy of accelerated MD methods is that it is possible to achieve such simulations by using various means to accelerate the MD trajectory for the escape of a system from a given potential minimum.

To date, there have been several different types of proposed accelerated MD methods [16–20]. In this paper, we focus on the bond-boost method [21, 22], which is a variant of Voter's hyperdynamics scheme [16, 17]. We review the bond-boost method and demonstrate its application to kinetic phenomena relevant to thin-film growth in three case studies designed to illustrate various aspects of the method. We discuss the outlook for future applications of the bond-boost method in materials simulation.

2. The bond-boost method

We consider an N-particle system in the canonical ensemble evolving on a potential-energy surface composed of several local minima. The TST rate $k^{\mathrm{TST}}_{i\to j}$ of a particular transition from state i to state j is given by the ensemble-averaged flux through the dividing hypersurface that separates i and j:

$$k^{\mathrm{TST}}_{i\to j} = \frac{1}{2}\left(\frac{2k_{\mathrm{B}}T}{\pi m}\right)^{1/2} \frac{\int_i \delta^{\dagger}_{ij}\Theta_i \exp(-\beta V\{\mathbf{x}\})\,\mathrm{d}\mathbf{x}}{\int_i \Theta_i \exp(-\beta V\{\mathbf{x}\})\,\mathrm{d}\mathbf{x}}. \quad (1)$$

Here $\beta = 1/k_{\mathrm{B}}T$, V is the potential energy, $\Theta_i = 1$ if the system is in state i and zero otherwise, and $\delta^{\dagger}_{ij}$ is the delta function defining the location of the dividing hypersurface. We note that MD simulations yield these rates naturally, by sampling the canonical distribution via integration of Newton's equations of motion—although the timescales associated with rare events often greatly exceed those accessible with MD.

Further considering the TST rate expression, we note that the domain of the integral in the denominator of equation (1) is over the entire phase space of state i, while the numerator is non-zero only at the dividing surface between i and j. The *relative* frequencies for the system to evolve from state i to neighboring states $j, l, \ldots$ are governed by the relative rates $k_{i\to j}/k_{i\to l}$, which depend only on the numerator in equation (1). On the other hand, the denominator in equation (1) is needed to obtain the physical time for escaping from minimum i. The philosophy behind hyperdynamics is that, by manipulating the denominator in equation (1), i.e., the partition functions of the local minima, without affecting the transition states, we can achieve the same dynamical evolution as on the original potential surface and we can accelerate the simulation timescale. The challenge with hyperdynamics approaches is to devise an effective means of altering the partition functions of potential minima without influencing transition states.

In most accelerated MD schemes based on hyperdynamics, acceleration is achieved by adding a boost potential $\Delta V\{\mathbf{x}\}$ to the original potential $V\{\mathbf{x}\}$ [16, 17]. This boost potential must satisfy the condition that $\Delta V^{\dagger} = 0$ at any dividing hypersurface (†) between local minima. MD simulations are then run on a new, 'boosted' potential surface $V^*\{\mathbf{x}\} = V\{\mathbf{x}\} + \Delta V\{\mathbf{x}\}$. On the boosted potential surface V^*, equation (1) becomes

$$k^{\mathrm{TST}}_{i\to j} = \frac{1}{2}\left(\frac{2k_{\mathrm{B}}T}{\pi m}\right)^{1/2} \frac{\int_i \delta^{\dagger}_{ij}\Theta_i \exp(-\beta V\{\mathbf{x}\})\,\mathrm{d}\mathbf{x}}{\int_i \Theta_i \exp(-\beta V^*\{\mathbf{x}\})\,\mathrm{d}\mathbf{x}}. \quad (2)$$

If the new denominator is smaller, which can be ensured by having $\Delta V > 0$, the new rates $k^*_{i\to j}$ are higher and the system escapes more quickly out of the local minima. The connection to the physical system is established by observing that

$$k^{\mathrm{TST}}_{i\to j} = k^*_{i\to j}\frac{1}{\langle \exp(\beta\Delta V)\rangle_*}, \quad (3)$$

where $\langle\,\rangle_*$ represents a canonical average on the boosted potential surface. The average escape time $\tau^{\mathrm{TST}}_{i\to j}$ from the local minimum via process $i \to j$ is the inverse of the corresponding rate, so

$$\tau^{\mathrm{TST}}_{i\to j} = \tau^*_{i\to j}\langle \exp(\beta\Delta V)\rangle_*. \quad (4)$$

In translation, the simulation time on the boosted potential surface V^* is not equal to the physical time. The latter becomes a statistical average and can potentially become orders of magnitude higher than the simulation time.

A variety of different schemes have been proposed for devising a boost potential ΔV such that $\Delta V > 0$ at a potential minimum and $\Delta V^{\dagger} = 0$ at the dividing hypersurfaces (†) [16, 17, 21–27]. Our group has developed the bond-boost method [21, 22], in which an empirical form for the boost potential is derived from the bond-breaking picture of atomic cohesion in the solid state. Without assuming specific processes, we tag all atoms that are able to participate in configurational changes. In the case of a surface, this may include all atoms within the first one or two layers. The boost potential $\Delta V\{\mathbf{x}\}$ is a function of all nearest-neighbor bond lengths associated with the tagged atoms and is given by

$$\Delta V\{\mathbf{x}\} \equiv \Delta V(r\{\mathbf{x}\}) = A\{r_1 \cdots r_{N_{\mathrm{b}}}\}\sum_{i=1}^{N_{\mathrm{b}}} \delta V(r_i). \quad (5)$$

Here, r_i are the lengths of the tagged bonds, N_{b} is the number of tagged bonds, $A\{r_1 \cdots r_{N_{\mathrm{b}}}\}$ is an envelope function, and $\delta V(r_i)$ is a boost potential applied to each bond i. To enforce the condition that $\Delta V^{\dagger} = 0$ at the dividing hypersurfaces (†) we specify that $\Delta V \to 0$ when one of the tagged bonds 'breaks'. If we monitor the relative stretch/compression of bond i, given by

$$|\epsilon_i| = \left|\frac{r_i - r_i^{\mathrm{eq}}}{r_i^{\mathrm{eq}}}\right| \quad , \quad (6)$$

where r_i is the instantaneous length of bond i and $r_i^{\rm eq}$ is the equilibrium value, we say that a bond has broken when ϵ_i of any of the bonds exceeds an empirical threshold value of q. We choose q to be smaller than the maximum fractional bond change at the transition states of elementary processes. The value of q for a given application can be determined from a few trial calculations and a conservative estimate of q should be sufficient to enforce $\Delta V^\dagger = 0$ for all processes of interest. In the discussion below, we focus on applications for which a single value of the threshold is sufficient. However, it is possible to define a boost based on multiple values of q (e.g., for different bond types) [28]. Empirical functional forms are assigned to $A\{r_1 \cdots r_{N_b}\}$ and $\delta V(r_i)$ in equation (5). These are continuous functions, whose form can be adapted to the problem at hand—specific choices will be presented in the sections below. The bond-boost potential $\delta V(r_i)$ is maximum when a bond is at its equilibrium value and goes to zero when the fractional change in bond length reaches the threshold q:

$$\delta V(r_i) \equiv \delta V(\epsilon_i) = \begin{cases} \max & \text{if } \epsilon_i = 0 \\ 0 & \text{if } |\epsilon_i| > q, \end{cases} \tag{7}$$

with ϵ_i defined by equation (6). The bond-boost potential effectively lowers the binding energy of each bond that is close to its equilibrium value.

The envelope function $A\{r_1 \cdots r_{N_b}\}$ depends only on the maximally stretched or compressed bond, and it is a continuous function fulfilling the boundary conditions:

$$A\{r_1 \cdots r_{N_b}\} \equiv A(\epsilon^{\max}) = \begin{cases} 1 & \text{if } \epsilon^{\max} = 0 \\ 0 & \text{if } \epsilon^{\max} > q, \end{cases} \tag{8}$$

with $\epsilon^{\max} = \mathrm{MAX}\{|\epsilon_i|\}$. The envelope function couples all the tagged bonds, so that when the system approaches a transition state, all individual boost potentials go to zero. Since the envelope function depends only on the maximally stretched bond, it also channels the boost into the maximally stretched bond that is about to break.

Equation (5) provides the boost potential when the system is in a known local minimum. When the system leaves the local minimum, one of the bonds exceeds the threshold q and $\Delta V \to 0$. A new equilibrium configuration needs to be found to define the new bond-stretch parameters $r_i^{\rm eq}$. There are efficient ways to find energy minima, for example we employ conjugate-gradient minimization [29]. In our implementation, we define sequential time windows and check if the system is boosted during those times. If no boost occurs, a transition to a new state is assumed and the equilibrium configuration is reset by finding the new minimum.

In summary, the bond-boost algorithm is as follows. The bonds belonging to atoms that may participate in transitions are tagged. We find the equilibrium configuration for the local minimum and store the equilibrium values of the tagged bond lengths. Standard MD is done on the boosted potential surface defined by $V^* = V + \Delta V$, with ΔV defined by equation (5). When the system progresses to a new minimum, the boost potential goes to zero and remains there. We detect that and after some predefined time period, we reset the reference equilibrium configuration accordingly. The physical time increment δt corresponding to each MD time step δt^* (where δt^* is in the fs range) is calculated according to equation (4):

$$\delta t = \delta t^* e^{\beta \Delta V}. \tag{9}$$

The usefulness of the bond-boost method resides in its simplicity and general applicability. At each MD step, the only additional computation involves the forces associated with ΔV, since bond lengths are typically calculated in standard MD algorithms. Thus, a boosted MD step requires only a fraction more than the computation needed for a regular MD step. By construction, the computational cost of the method scales as the number of atoms included in the boost, and therefore the relative computational overhead should remain negligible for any system size. Additional work is necessary to find the local energy minima. However, since the transition rates are slow with respect to the timescale of one MD step, many thousands of MD steps will correspond to one minimization. Using conjugate-gradient algorithms, this minimization should generally require only a few tens of force calculations, which is a small relative overhead. In addition, the method can be applied to systems modeled with any type of force field, from pair potentials to first principles, since the form of the boost potential can be tailored to the problem at hand.

3. Applications

We have applied the bond-boost method to a variety of problems involving thin-film growth [21, 22, 30]. Below, we demonstrate its application to kinetic phenomena relevant to thin-film growth in three case studies designed to illustrate various aspects of the method. We first illustrate aspects of the bond-boost method in studies of the diffusion of Cu atoms on Cu(001) [21]. In these studies, Cu interactions are described by a semi-empirical embedded-atom method (EAM) potential [31], which provides good computational efficiency and reasonable accuracy for metals with filled d shells. For many applications it is desirable to achieve a more accurate and detailed description of chemical bonding than that provided by (semi-) empirical force fields. To address this concern, we recently extended the bond-boost method to perform accelerated *ab initio* MD (AIMD) simulations [32]. We discuss results from preliminary studies in which we applied the bond-boost method to uncover diffusion mechanisms of Al adatoms on Al(110) using accelerated AIMD simulations. Finally, a problem inherent in many rare-event simulation methods (including KMC) is the 'small-barrier problem', in which the system resides in a group of states connected by small energy barriers and separated from the rest of phase space by large barriers. We developed the state-bridging bond-boost method to address this problem and we discuss its application to study the diffusion of Co clusters on Cu(001) [22, 30].

3.1. Diffusion of Cu atoms on Cu(001): demonstration of the bond-boost method

We can illustrate various basic aspects of the bond-boost method in simulations of Cu atom, dimer, and vacancy

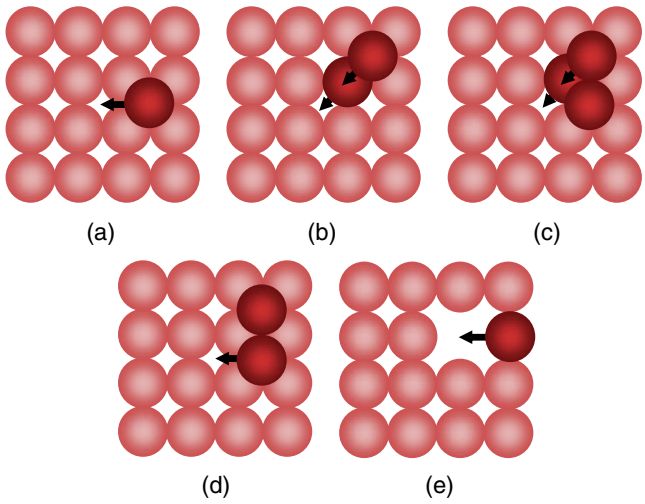

Figure 1. Mechanisms by which Cu atoms, dimers, and a vacancy can diffuse on Cu(001): (a) adatom hop; (b) adatom exchange; (c) dimer exchange; (d) dimer hop; (e) vacancy hop. Atoms, dimers, and first-layer surface atoms directly involved in the transition are dark, while the rest of the first-layer surface atoms are light.

diffusion on the Cu(100) surface. The mechanisms for these processes are depicted in figure 1. The interatomic interactions in these studies are modeled with the EAM potential of Foiles *et al* [31]. To model the Cu(001) surface, we use a six-layer slab with 50 atoms in each layer. The two bottom layers are immobile, while the third layer acts as a heat bath. Constant temperature MD is performed using the Nosé–Hoover algorithm [33] for atoms in the third layer. Atoms in the top three layers follow Newton's equations of motion for the microcanonical ensemble. All atoms in the first layer, as well as the adatom(s), are tagged for boosting. The nearest-neighbor distance in the bulk is 2.556 Å, and we boost all bonds that belong to the tagged atoms and are shorter than a cutoff of 3.0 Å in the local equilibrium configuration. This means that about 300 bonds are boosted simultaneously (cf, equation (5)).

For the bond-boost potential (equation (5)), we employ empirical functions for $\delta V(r_i)$ and $A\{r_1 \cdots r_{N_b}\}$ with the form

$$\delta V(r_i) = \delta V(\epsilon_i) = \frac{\Delta V^{\max}}{N_b}\left[1 - \left(\frac{\epsilon_i}{q}\right)^2\right], \qquad (10)$$

and

$$A(\epsilon^{\max}) = \left[1 - \left(\frac{\epsilon^{\max}}{q}\right)^2\right] \frac{1 - (\epsilon_{\max}/q)^2}{1 - P_1^2\,(\epsilon_{\max}/q)^2}, \qquad (11)$$

for $\epsilon \in [-q, q]$ and $\epsilon^{\max} \in [0, q]$. As stipulated by equation (7), $\delta V(\epsilon_i)$ is maximum when a bond is at its equilibrium value and goes to zero when the fractional change in bond length reaches the threshold q. To determine specific values for q, we used static calculations to find the maximal relative bond stretch at transition states, $q^\dagger$, which ranges from 0.35 for adatom hopping to 0.45 for exchange and vacancy hopping. We require $q < q^\dagger$, so we set $q = 0.3$. We note that δV also depends on a maximum boost parameter $\Delta V^{\max}$. This parameter sets the magnitude of the total boost potential ΔV and it needs to be on the scale of natural energy barriers in the system (≈0.5 eV in this case). We tested values from 0–0.5 eV for $\Delta V^{\max}$. The envelope function $A(\epsilon^{\max}) = A\{r_1 \cdots r_{N_b}\}$ depends only on the maximally stretched or compressed bond and has the properties indicated in equation (8). The form of equation (11) ensures that, as the system approaches a transition state, all individual boost potentials go to zero together with their first derivatives. This is necessary to maintain a continuous derivative of the potential-energy surface. The parameter $P_1 \lesssim 1$ in equation (11) controls the curvature near $\epsilon^{\max} = q$, and we choose values in the range 0.9–0.98. A large value means the force will change rapidly close to the boundary, while a small value corresponds to a 'softer' boundary.

Table 1. Prefactors (Γ_0) and activation energies (E_A) for elementary diffusion processes of Cu/Cu(100). The (*) values are from the MD study of Boisvert and Lewis [34].

Process	Γ_0 (THz)	Γ_0^* (THz)	E_A (eV)	E_A^* (eV)
Adatom hop	$40(\times e^{\pm 0.5})$	$20(\times e^{\pm 0.2})$	0.52 ± 0.03	0.49 ± 0.01
Adatom exchange	$270(\times e^{\pm 0.6})$	$437(\times e^{\pm 0.7})$	0.73 ± 0.04	0.70 ± 0.04
Vacancy hop	$54(\times e^{\pm 0.5})$	$27(\times e^{\pm 0.7})$	0.44 ± 0.03	0.47 ± 0.05
Dimer hop	$30(\times e^{\pm 0.7})$	$13(\times e^{\pm 0.5})$	0.47 ± 0.03	0.48 ± 0.03
Dimer exchange	$190(\times e^{\pm 0.8})$	$320(\times e^{\pm 0.8})$	0.71 ± 0.06	0.73 ± 0.05

Using accelerated MD, we calculated rates for all the processes shown in figure 1 for temperatures ranging between 230 and 600 K [21]. For each temperature, we performed several simulations of 30 ns (MD time) to obtain the rates of these various processes. The prefactors and barriers obtained from Arrhenius plots of the various rates as a function of temperature are shown in table 1. The values for the barriers agree well with static, TST estimates obtained using the step-and-slide method [13]. Table 1 also contains energy barriers and prefactors from the MD study of Boisvert and Lewis [34], which both fall within the uncertainties of our results. Using the same EAM potential as us, Boisvert and Lewis probed the temperature range between 650 and 900 K with regular MD simulations—below this range, the various rate processes are too slow to occur to any appreciable extent over the timescale of a regular MD simulation. Because of the acceleration possible with the bond-boost method, we were able to probe lower temperatures between 230 and 600 K.

To quantify the speed-up possible with the bond-boost method, we plot the boost as a function of temperature in figure 2. The boost is defined as the ratio of the physical time to the simulation time. According to equation (4), it is given as the average of the boost potential over the boosted potential surface, i.e.,

$$\text{Boost} = \langle e^{\beta \Delta V} \rangle_*. \qquad (12)$$

The boosts in figure 2 are obtained as averages over all runs for all the different types of rate processes shown in figure 1, with $\Delta V^{\max} = 0.4$ eV. In this study the computational overhead was less than 10% and therefore the boosts shown in figure 2 were not adjusted for it. The boost factor shows Arrhenius behavior, increasing exponentially with inverse temperature, with values of up to 10^6 at the low temperatures.

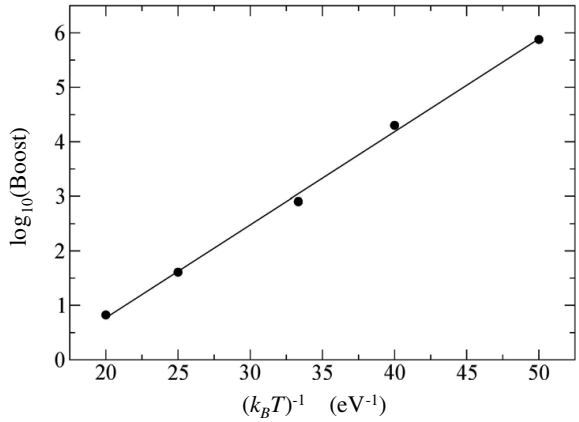

Figure 2. The average boost, given by equation (12), as a function of temperature for the various diffusion processes shown in figure 1.

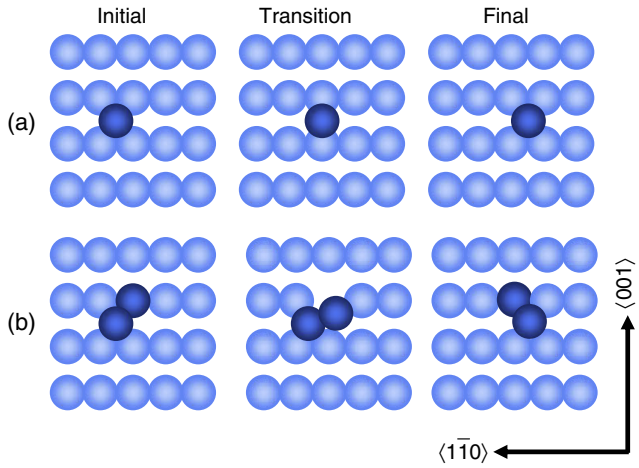

Figure 3. Mechanisms by which Al can diffuse in the in-channel direction on Al(110), indicating the initial, transition, and final states for: (a) adatom hop and (b) adatom exchange. Atoms and first-layer surface atoms directly involved in the transition are dark, while the rest of the first-layer surface atoms are light.

It is evident from equation (12) that the boost declines with increasing temperature, since β is inversely proportional to temperature. Additionally, the average displacement of the atoms from equilibrium grows with increasing temperature and sampling of high-boost regions of the potential surface is reduced. On the other hand, at very low temperatures, transition times may increase faster than the boost and no transitions are observed over the simulation timescale. For example, although very large boosts are possible below 230 K in our simulations, we could not observe any transitions over the simulation timescale at temperatures below about 230 K. The method thus achieves peak efficiency in the mid-range between the high-temperature domain of standard MD and very low temperatures.

3.2. Diffusion of Al on Al(110): ab initio accelerated MD with the bond-boost method

Achieving accurate multi-scale models of thin-film epitaxy based on first principles is an important objective. As discussed above, a popular method for simulating thin-film growth is to employ KMC simulations based on mechanisms and rates derived from TST searches on semi-empirical or first-principles-derived potential surfaces. However, static TST search algorithms are not guaranteed to converge to the most salient kinetic mechanisms. In this example, we highlight the utility of accelerated *ab initio* MD (AIMD) simulations in uncovering the mechanisms underlying the diffusion of Al on Al(110), which mediates the formation of nanohuts during Al(110) homoepitaxy [35–37]. We use the Vienna *ab initio* simulation package (VASP) [38–40], which is based on density-functional theory (DFT), to characterize the mechanisms and energy barriers by which Al diffusion may occur. We perform two sets of calculations: TST searches using the climbing-image nudged elastic band (CI-NEB) method [12] to elucidate minimum-energy pathways and, hence, diffusion mechanisms and energy barriers; and accelerated MD simulations with the bond-boost method [21]. The CI-NEB method is implemented in the VASP package and we modified the VASP source code to execute accelerated AIMD simulations based on the bond-boost method. We show how the combination of these two methods is beneficial in understanding mechanisms of Al adatom diffusion.

We consider the diffusion of a single Al adatom on Al(110). In the static, CI-NEB calculations, we constructed periodic supercells with a slab thickness of 10 atomic layers and 5 × 4 surface atoms in each layer. We use ultrasoft Vanderbilt pseudo-potentials [41], as supplied by Kresse and Hafner [42], the GGA-PW91 [43], Fermi–Dirac smearing [44] with a width of 0.2 eV, and an energy cutoff of 129.2 eV. To sample the Brillouin zone, we use a 5 × 4 × 1 Monkhorst–Pack [45] **k**-point mesh. We tested the convergence of the slab thickness, **k**-point mesh, etc in a previous study [46].

To probe diffusion mechanisms, we performed two sets of calculations. Initially, we performed CI-NEB calculations using three images in addition to the initial and final states. The positions of the images are optimized to yield a minimum-energy pathway in which (by virtue of the CI-NEB method) the image with the highest energy resides at the transition state. To confirm the transition states of the exchange mechanisms, we obtained minimum-energy pathways utilizing our original initial states and the transition states from the first calculation as the final state. Because of the symmetry of the diffusion mechanisms, it is sufficient to probe only half of the pathway. These second calculations confirmed that no new minima or transition states arise in the original pathways. In both sets of calculations, we relax adatom positions (and the top five surface layers) until the force on each unconstrained atom is smaller than 0.01 eV Å^{-1}.

One possible mechanism by which an Al adatom can diffuse on Al(110) involves a hop between neighboring binding sites along the [1$\bar{1}$0], or 'in-channel' direction, as shown in figure 3(a). Using the CI-NEB method, we find the transition state indicated in figure 3(a) and an energy barrier of 0.47 eV for this mechanism. A more favorable mechanism for in-channel diffusion is when the adatom changes places with a surface atom in the exchange mechanism, partially shown

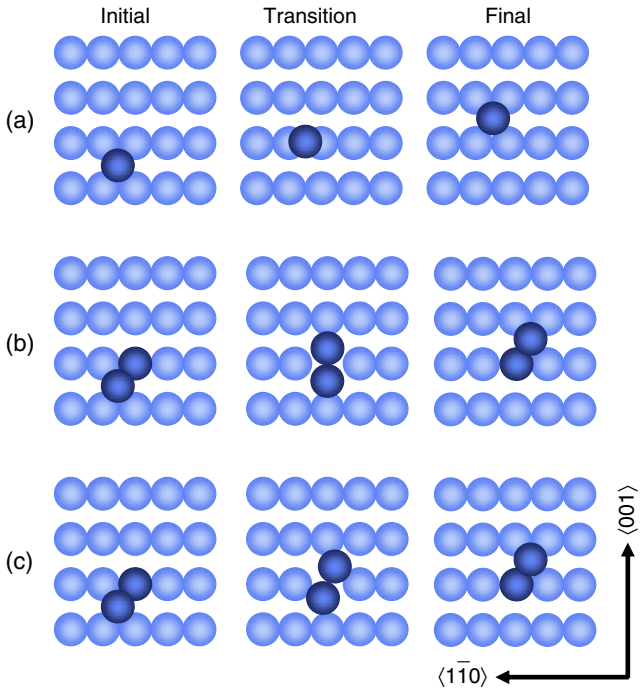

Figure 4. Mechanisms by which Al can diffuse in the cross-channel direction on Al(110) indicating the initial, transition, and final states for: (a) adatom hop; (b) adatom exchange with a 'linear' intermediate state and; (c) adatom exchange with the 'diagonal' transition state. Atoms and first-layer surface atoms directly involved in the transition are dark, while the rest of the first-layer surface atoms are light.

in figure 3(b). This mechanism involves two steps: a move over a first transition state (shown in the figure) to a shallow, intermediate state (not shown), over a second transition state (a mirror image of that shown in the figure) to the final state. We can calculate an effective barrier for the net move from the initial state through the intermediate state to the final state shown in figure 3(b) [32]. In the temperature range between 200 and 500 K, this barrier ranges between 0.39 and 0.40 eV. The atom may also move in the [100], or 'cross-channel' direction. The energy barrier for the cross-channel hop shown in figure 4(a) is 0.71 eV, making this move unfavorable compared to the in-channel moves. However, if we consider exchange diffusion via the mechanism shown in figure 4(b), we find that with an energy barrier of 0.38 eV, cross-channel exchange is on par with in-channel exchange.

In the accelerated AIMD simulations, we represent the surface as a five-layer slab with 16 (4×4) atoms in each layer. Atoms in the top two layers are allowed to move, while the bottom three layers are fixed to the Al bulk lattice positions. We use ultrasoft Vanderbilt pseudo-potentials [41], as supplied by Kresse and Hafner [42], the GGA-PW91 [43], Methfessel–Paxton smearing with a width of 0.2 eV, and an energy cutoff of 129.2 eV. A $2 \times 2 \times 1$ Monkhorst–Pack [45] **k**-point mesh is employed to sample the Brillouin zone. To ensure that our simulations are run within the canonical ensemble, we implemented and utilized the Andersen thermostat [47] within the VASP code. This thermostat robustly maintains the desired temperature in our simulations and we also verified that we could achieve the Maxwell–Boltzmann distribution of velocities with this thermostat.

To enhance the computational efficiency of the AIMD simulations, we use fewer **k** points and a thinner slab than we do in the static TST calculations. This enhanced efficiency comes at the expense of accuracy. Since we are using the AIMD simulations to uncover mechanistic trends without acquiring specific numbers, a certain loss of accuracy is acceptable. Care must be taken, however, to ensure that the accuracy of the AIMD simulations is sufficient to yield the same qualitative results as the static TST calculations. To this end, we used the CI-NEB method to calculate the energy barriers for all of the diffusion mechanisms using the slabs that we used for the AIMD simulations. Although key energy barriers can differ by as much as 70 meV between a five-layer slab with 4 **k** points and a ten-layer slab with 20 **k** points, we nevertheless find that both slabs exhibit the same trends regarding the relative values of the diffusion barriers.

For the accelerated AIMD simulations, our strategy is different than that in our studies of the diffusion of Cu on Cu(001), discussed above in section 3.1. In the studies of Cu/Cu(001), our use of a semi-empirical EAM potential enabled long MD simulations. Thus, the quantities shown in table 1 and figure 2 for Cu/Cu(001) represent averages over several hundred to several thousand hops [21]. AIMD simulations have considerably more computational overhead than simulations based on semi-empirical potentials and it is difficult to compile sufficient statistics to obtain meaningful estimates of average hop rates. Thus, here we exploit the fact that we can accurately sample the canonical ensemble and, even with limited statistics, we can find the most likely rate process(es) that occur at a fixed temperature. We compare these mechanisms to those we find with (static) CI-NEB calculations to ensure we have found the most important kinetic processes.

We consider the diffusion of an Al adatom on Al(110) at temperatures ranging from 300 to 650 K. To perform accelerated AIMD, we use a boost potential of the form given by equations (5), (10), and (11) and we apply a boost to the adatom, as well as to all the atoms in the top two surface layers. We tested several different values of $\Delta V^{\max}$, ranging from 0.1 to 0.4 eV and we employ a threshold of $q = 0.6$. We begin a simulation by running a short, 1 ps trajectory to equilibrate the adatom and surface atoms to the desired temperature. Subsequently the boost is 'turned on' and we monitor the value of $\epsilon^{\max}$ as a function of time. As this value approaches and exceeds the threshold q, the boost goes to zero and we run the unboosted AIMD simulation for an additional 1–2 ps to allow the system to equilibrate in its new minimum. By following the trajectories we can ascertain the mechanisms by which the Al adatoms diffuse.

Figure 4(c) shows the only mechanistic sequence observed in these simulations—cross-channel exchange. Close comparison of this mechanism to the one initially uncovered in static TST calculations (cf, figure 4(b)) reveals that the transition state is different than that for the mechanism in figure 4(b)—the accelerated AIMD simulations imply a 'diagonal' transition state formed by the adatom and the

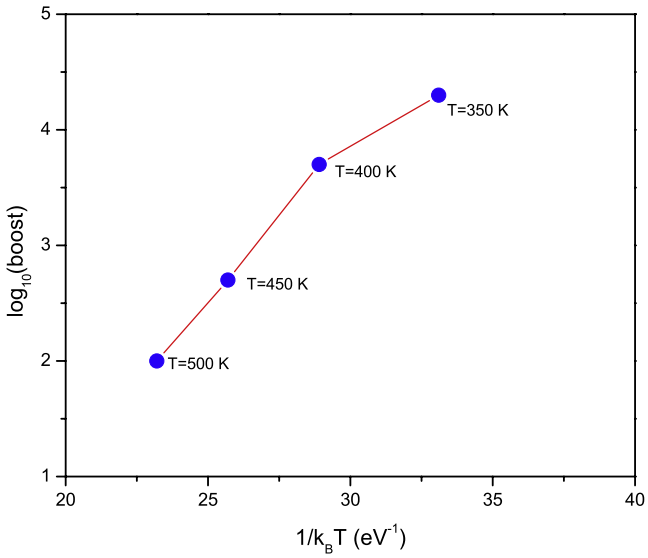

Figure 5. Boost achieved in *ab initio* accelerated MD simulations as a function of temperature for $\Delta V^{\max} = 0.25$ eV. At a fixed temperature, the boost is obtained from one run only.

exchanged surface atom. Static estimates of the minimum-energy path with the CI-NEB method will converge close to the initial guess for the minimum-energy pathway. By performing an additional CI-NEB search using an initial guess based on the accelerated AIMD trajectories, we confirm that a second transition state for cross-channel exchange occurs—that shown in figure 4(c). The energy barrier associated with this mechanism (0.33 eV) is 50 meV lower than the cross-channel exchange mechanism shown in figure 4(b) and hence this mechanism is more likely.

It is of interest to know the boost achievable in these accelerated AIMD simulations. Ideally, the boost is obtained as an ensemble average (using equation (12)) over many different runs. However, the high computational overhead associated with these simulations makes it difficult to obtain ensemble averages. Instead, we obtain the boost from a single trajectory run at each of the temperatures probed. The initial and the final (unboosted) equilibration periods are excluded in calculating the boost. The results from simulations with $\Delta V^{\max} = 0.25$ eV are shown in figure 5, where we see an increase in the boost with decreasing temperature, similar to that shown in figure 2 for Cu/Cu(001). The AIMD simulation time (excluding equilibration periods) associated with these different trajectories ranges from 5.8 ps at 350 K to 0.71 ps at 500 K. Thus, we were able to attain physical times ranging from 71 ps to 73 ns in these simulations. The boost associated with Al/Al(110) diffusion is smaller that that for Cu/Cu(001) because of the smaller energy barriers for diffusion on Al(110) (cf, equation (12))—for the 5-layer slab used in the accelerated AIMD simulations, the static barrier is only 0.26 eV.

3.3. Co/Cu(001) heteroepitaxy: the state-bridging bond-boost method

The heteroepitaxial growth of Co on Cu substrates has been studied extensively during the last decade. Systems consisting of thin Co/Cu layers exhibit interesting properties stemming from reduced dimensionality, strained atomic structure, and magnetic interlayer coupling. Electronic properties of heteroepitaxial thin-film structures are very sensitive to their atomic-scale morphology, hence the importance of accurate characterization of the growth morphology. In a series of studies [22, 30], we probed kinetic mechanisms associated with the growth of Co on Cu(001) using both accelerated MD and TST search algorithms. To describe Co and Cu interactions in these studies, we employ an empirical tight-binding second-moment approximation (TBSMA) potential that was originally developed by Levanov and colleagues [48] and modified by us [30] to provide better agreement with the results of first-principles calculations by Pentcheva and Scheffler [49]. To model the Cu(001) substrate, we employ a simulation cell with a slab consisting of five Cu(001) layers with 1296 atoms/layer. We apply periodic boundary conditions in the directions parallel to the surface. In our accelerated MD simulations, the bottom two layers are immobile and fixed to ideal bulk positions. The middle layer is connected to a Nosé–Hoover thermostat [33] for controlling the temperature, while the top two layers follow Newton's equations of motion for the microcanonical ensemble.

In the initial stages of low-temperature, thin-film growth in this system, Co adatoms aggregate to form small islands on top of the Cu(001) surface. These islands can diffuse and rearrange through atoms hopping along their edges, as well as through shearing mechanisms. Figure 6 shows typical low-temperature kinetic mechanisms involved in the diffusion of Co adatoms and small islands on Cu(001) and table 2 indicates the energy barriers for these processes that we obtained in TST searches based on the step and slide method [13], as well as with the accelerated MD simulations discussed below. This system is a prime example of the 'small-barrier problem' because diffusion along island edges is much faster than isolated adatom hopping. With a static barrier of 0.30 eV, edge diffusion (figure 6(c)) is about 10^6 times faster than adatom hopping (figure 6(a), $\Delta E^{\text{static}} = 0.63$ eV) at room temperature and trimer rotation (figure 6(d), $\Delta E^{\text{static}} = 0.10$ eV) is about 10^8 times faster than adatom hopping. Fast trimer rotation leads to a pool of 24 different states corresponding to different trimer configurations connected by small barriers. Without special measures, even accelerated MD simulations are limited by the short timescale of repetitive edge hops and trimer rotations. Here we review the state-bridging bond-boost method [22] to address the small-barrier problem in this system.

Figure 7 illustrates the essence of the small-barrier problem and the general concepts associated with the State-Bridging bond-boost method. In the bond-boost method, the time boost is controlled by the magnitude of the boost potential $\Delta V^{\max}$ (cf, equation (10)), which can, in principle, be tuned. However, there is an upper limit on the achievable boost, as a strong boost that exceeds the transition-state barrier does not preserve the correct dynamics of the low-barrier processes. As illustrated in figure 7(a), the shallow states A, B, C, and D become local peaks in this case and the minima are not effectively sampled, as the system is pushed toward the

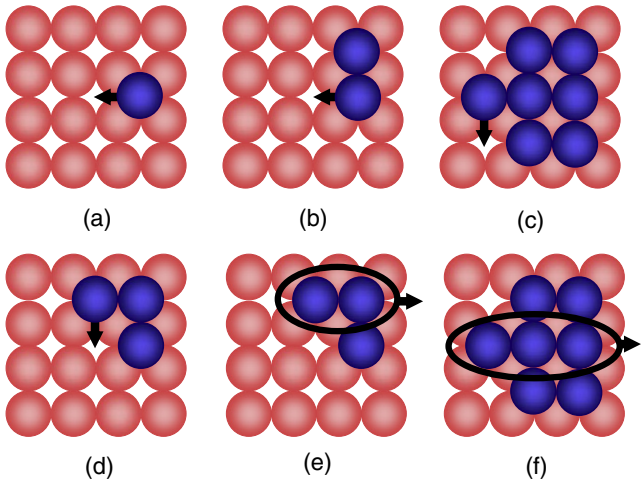

Figure 6. Low-temperature mechanisms by which Co adatoms and small islands diffuse on Cu(001): (a) adatom hop; (b) dimer hop; (c) adatom hop at a Co step edge; (d) trimer rotation; (e) trimer hop via a concerted, two-atom shearing mechanism; (f) heptamer hop via a concerted, three-atom shearing mechanism. Co atoms are shown in blue (dark) and Cu surface atoms are shown in red (light).

Table 2. Energy barriers ΔE obtained from accelerated MD simulations of the diffusion processes of Co/Cu(001) shown in figure 6. Static values are obtained with the step-and-slide method [13].

Process	ΔE^{static} (eV)	ΔE^{MD} (eV)
Adatom hop	0.63	0.63
Dimer hop	0.62	0.63
Adatom edge hop	0.30	—
Trimer rotation	0.10	—
Trimer shear	0.64	0.65
Heptamer shear	0.56	0.57

transition states. With a small boost, appropriate for shallow minima (cf, figure 7(b)), the system will rapidly and repeatedly cycle between A, B, C, and D and escape to state E only over a much longer timescale. The bulk of the simulation time is spent on these repeated transitions and evolution of the system is limited. Our proposed solution to this problem is to combine the large boost shown in figure 7(a) with 'bridge potentials' ΔV^{bridge}, which span the transition states between states A–D, as shown in figure 7(c). In doing this, we consolidate the shallow states A–D into a single, coarse state. This procedure rests on the observation that equilibrium between the shallow states is reached long before any slow event D $\rightarrow$ E occurs. Since, on the timescale of the slow escape, the 'fast' dynamics becomes irrelevant, we drop the requirement that $\Delta V(\mathbf{x}) = 0$ (cf, section 2) at the 'fast' transition states. Thus, for the escape rate $k_{\text{D}\rightarrow\text{E}}$ the entire set A $\cup$ B $\cup$ C $\cup$ D acts as the 'initial state'. Below, we outline the essential elements of the state-bridging bond-boost method—details can be found elsewhere [22].

The first issue in implementing the scheme shown in figure 7(c) is definition and detection of fast, recurrent processes as those having barriers below a threshold value ΔE^{th}. When an event M $\rightarrow$ N occurs, we determine if the

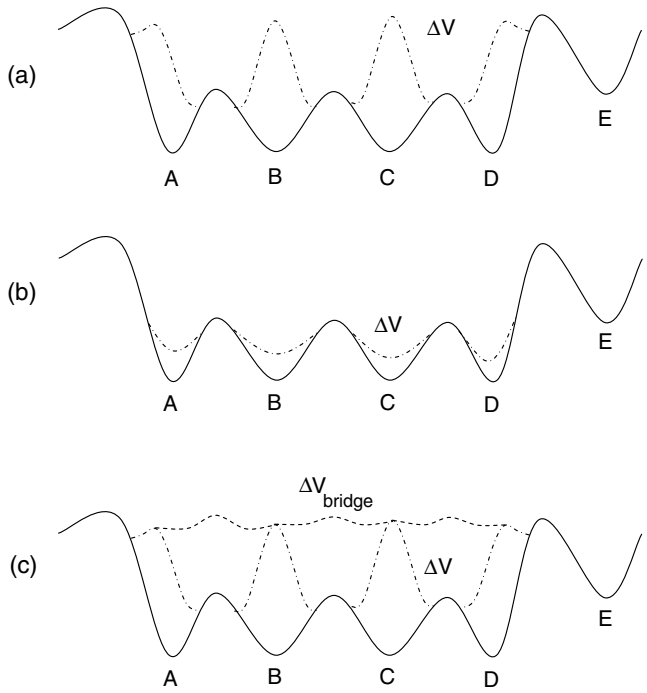

Figure 7. Illustration of the 'small-barrier' problem: (a) too large of a boost (large ΔV^{max}) leads to the creation of artificial minima at transition states; (b) an appropriate (small) boost will lead to rapid cycling of the system between shallow minima, with low acceleration; (c) combination of a large boost with bridging potentials to consolidate the shallow states into one, large state.

barrier $\Delta E^{\dagger}_{\text{M}\rightarrow\text{N}}$ is less than ΔE^{th}. This is done with little overhead using the step-and-slide method [13]. If a process is tagged as 'fast', the initial and final states M and N are stored for pattern matching, i.e. when the system revisits state M it should 'recognize' its transition to N as a fast process and activate the appropriate bridge potential. The second element of the State-Bridging method is to construct a bridge potential between M and N (see figure 7(c)) The complete boost $\Delta V(\mathbf{x})$ is obtained by taking the envelope of all bond and bridge terms which are active at the particular instantaneous configuration.

The algorithm proceeds as follows: after each event, we adopt a low boost, i.e. $\Delta V^{\text{max}} \lessapprox \Delta E^{\text{th}}$, that preserves the correct fast dynamics. The low boost allows for accurate dynamics of the fast events, albeit with low acceleration. As each new shallow state is encountered, it is stored and a bridge potential is constructed between the new state and its preceding state. In this way, we construct a chain of connected states. Eventually, we sample all of the shallow states. If no new state is encountered during a predefined waiting time t^{wait}, the boost strength is increased to the desired high value, the bridges are activated, and the simulation switches to the timescale of the slow events. We choose t^{wait} larger than the average waiting time for a process having ΔE^{th}. In each new state, the simulation code performs local pattern matching against stored states to find the applicable bridge potential terms. With efficient state recognition algorithms, the computational overhead is generally less than 10% of the normal simulation time for our test cases, and should remain low for a finite number of fast processes.

Using the State-Bridging bond-boost method, we studied the diffusion of the small Co/Cu(001) islands that are shown in figure 6. Energy barriers from static calculations and bridged, accelerated MD simulations are shown in table 2 and exhibit excellent agreement. Hopping along island edges is much faster than events that lead to center-of-mass motion of small islands, which are mainly collective shearing mechanisms. Dimers, trimers and heptamers have a high mobility, comparable to that of the isolated adatom. The trimer hops via a concerted shearing mechanism of two atoms, while the heptamer hops via concerted shearing of three atoms in its middle row. A less favorable mechanism for heptamer hopping involves an edge atom climbing on top of the cluster and descending again. By overcoming the small-barrier limitation, we achieve boosts ranging from 10^4 at 450 K to 10^8 at 250 K and slow island diffusion is correctly captured.

4. Conclusions and future outlook

Thus, we reviewed the bond-boost method and its application to thin-film epitaxy. There are a number of future prospects for this method in materials, catalysis, and even possibly in biology. Accelerated MD methods can be useful for finding salient kinetic processes and even for simulating certain experiments [30]. Although the efficiency of accelerated MD will improve as computing capabilities advance from the terascale to the petascale, perhaps an ideal solution in terms of efficiency is to combine these methods with KMC simulations. This marriage will improve our capabilities for multi-scale materials simulation.

Acknowledgments

This work has been supported by the National Science Foundation grants Nos. ECC 0085604 and DMR 0514336.

References

[1] Hänggi P, Talkner P and Borkovec M 1990 *Rev. Mod. Phys.* **62** 251
[2] Fichthorn K A and Weinberg W H 1991 *J. Chem. Phys.* **95** 1090
[3] Voter A F 1986 *Phys. Rev.* B **34** 6819
[4] Henkelman G and Jónsson H 2001 *J. Chem. Phys.* **115** 9657
[5] Trushin O, Karim A, Kara A and Rahman T S 2005 *Phys. Rev.* B **72** 9
[6] Chatterjee A and Vlachos D G 2007 *J. Comput.-Aided Mater. Des.* **14** 253
[7] Shim Y and Amar J G 2005 *Phys. Rev.* B **71** 115436
[8] Merrick M L and Fichthorn K A 2007 *Phys. Rev.* E **75** 011606
[9] Ulitsky A and Elber R 1990 *J. Chem. Phys.* **92** 1510
[10] Ionova I V and Carter E A 1993 *J. Chem. Phys.* **98** 6377
[11] Henkelman G and Jónsson H 1999 *J. Chem. Phys.* **111** 7010
[12] Henkelman G, Uberuaga B P and Jónsson H 2000 *J. Chem. Phys.* **113** 9901
[13] Miron R A and Fichthorn K A 2001 *J. Chem. Phys.* **115** 8742
[14] W E, Ren W and Vanden-Eijnden E 2002 *Phys. Rev.* B **66** 052301
[15] Olsen R A, Kroes G J, Henkelman G, Arnaldsson A and Jónsson H 2004 *J. Chem. Phys.* **121** 9776
[16] Voter A F 1997 *J. Chem. Phys.* **106** 4665
[17] Voter A F 1997 *Phys. Rev. Lett.* **78** 3908
[18] Voter A F 1998 *Phys. Rev.* B **57** R13985
[19] Sörensen M R and Voter A F 2000 *J. Chem. Phys.* **112** 9599
[20] Voter A F, Montalenti F and Germann T C 2002 *Ann. Rev. Mater. Res.* **32** 321
[21] Miron R A and Fichthorn K A 2003 *J. Chem. Phys.* **119** 6210
[22] Miron R A and Fichthorn K A 2004 *Phys. Rev. Lett.* **93** 138201
[23] Steiner M M, Genilloud P-A and Wilkins J W 1998 *Phys. Rev.* B **57** 10236
[24] Pal S and Fichthorn K A 1999 *Chem. Eng. J.* **74** 77
[25] Wang J-C, Pal S and Fichthorn K A 2001 *Phys. Rev.* B **63** 85403
[26] Chen L Y and Horing N J M 2007 *J. Chem. Phys.* **126** 224103
[27] Zhou X, Jiang Y, Kremer K, Ziock H and Rasmussen S 2006 *Phys. Rev.* E **74** 035701
[28] Mignogna M, Hammerschmidt T, Kratzer P, Scheffler M and Fichthorn K A 2008 in preparation
[29] Press W H, Teukolsky S A, Vetterling W T and Flannery B P 1992 *Numerical Recipes in C: the Art of Scientific Computing* 2nd edn (Cambridge: Cambridge University Press)
[30] Miron R A and Fichthorn K A 2005 *Phys. Rev.* B **72** 035415
[31] Foiles S M, Baskes M I and Daw M S 1986 *Phys. Rev.* B **33** 7983
[32] Wang Y, Tiwary Y and Fichthorn K A 2008 in preparation
[33] Hoover W G 1985 *Phys. Rev.* A **31** 1695
[34] Boisvert G and Lewis L J 1997 *Phys. Rev.* B **56** 7643
[35] de Mongeot F B, Zhu W, Molle A, Buzio R, Boragno C, Valbusa U, Wang E G and Zhang Z Y 2003 *Phys. Rev. Lett.* **91** 016102
[36] Zhu W, Buatier de Mongeot F, Valbusa U, Wang E G and Zhang Z Y 2004 *Phys. Rev. Lett.* **92** 106102
[37] Fichthorn K A and Scheffler M 2004 *Nature* **429** 617
[38] Kresse G and Hafner J 1993 *Phys. Rev.* B **47** 558
[39] Kresse G and Furthmüller J 1996 *Comput. Mater. Sci.* **6** 15
[40] Kresse G and Furthmüller J 1996 *Phys. Rev.* B **54** 11169
[41] Vanderbilt D 1990 *Phys. Rev.* B **41** 7892
[42] Kresse G and Hafner J 1994 *J. Phys.: Condens. Matter* **6** 8245
[43] Perdew J P and Wang Y 1992 *Phys. Rev.* B **45** 13244
[44] Mermin N D 1965 *Phys. Rev.* **137** A1441
[45] Monkhorst H J and Pack J D 1976 *Phys. Rev.* B **13** 5188
[46] Tiwary Y and Fichthorn K A 2007 *Phys. Rev.* B **75** 235451
[47] Andersen H C 1980 *J. Chem. Phys.* **72** 2384
[48] Levanov N A, Stepanyuk V S, Hergert W, Bazhanov D I, Dederichs P H, Katsnelson A and Massobrio C 2000 *Phys. Rev.* B **61** 2230
[49] Pentcheva R 2000 *Theory of Adsorption, Diffusion and Growth of Cobalt on Cu(001)* (Berlin: Freie Universität)