Localized saddle-point search and application to temperature-accelerated dynamics

Yunsic Shim, Nathan B. Callahan, and Jacques G. Amar

Citation: J. Chem. Phys. **138**, 094101 (2013); doi: 10.1063/1.4793218
View online: http://dx.doi.org/10.1063/1.4793218
View Table of Contents: http://jcp.aip.org/resource/1/JCPSA6/v138/i9
Published by the AIP Publishing LLC.

Localized saddle-point search and application to temperature-accelerated dynamics

Yunsic Shim,[1] Nathan B. Callahan,[1,2] and Jacques G. Amar[1]
[1]*Department of Physics and Astronomy, University of Toledo, Toledo, Ohio 43606, USA*
[2]*Department of Physics, Indiana University, Bloomington, Indiana 47405, USA*

(Received 28 November 2012; accepted 7 February 2013; published online 1 March 2013)

We present a method for speeding up temperature-accelerated dynamics (TAD) simulations by carrying out a localized saddle-point (LSAD) search. In this method, instead of using the entire system to determine the energy barriers of activated processes, the calculation is localized by only including a small chunk of atoms around the atoms directly involved in the transition. Using this method, we have obtained N-independent scaling for the computational cost of the saddle-point search as a function of system size N. The error arising from localization is analyzed using a variety of model systems, including a variety of activated processes on Ag(100) and Cu(100) surfaces, as well as multiatom moves in Cu radiation damage and metal heteroepitaxial growth. Our results show significantly improved performance of TAD with the LSAD method, for the case of Ag/Ag(100) annealing and Cu/Cu(100) growth, while maintaining a negligibly small error in energy barriers. © *2013 American Institute of Physics*. [http://dx.doi.org/10.1063/1.4793218]

I. INTRODUCTION

Recently there has been a great deal of interest in the development of methods for extending the time-scales of non-equilibrium simulations, including accelerated dynamics,[1–13] adaptive kinetic Monte Carlo,[14,15] self-learning kinetic Monte Carlo,[16–21] and Metropolis Monte Carlo.[22] Other methods based on metadynamics[23–26] have also been proposed to speed up slow dynamics, typically for systems trapped in deep energy minima. In addition, a considerable amount of effort has been devoted to the development of efficient computational methods for calculating the minimum energy path and corresponding energy barriers[27–32] in condensed matter systems.

One particular accelerated dynamics method is temperature-accelerated dynamics[8,9] (TAD) in which a high-temperature basin-constrained molecular dynamics (MD) simulation is used to speed-up the search for activated events, in order to determine the first transition which will occur at the desired low temperature $T_{\rm low}$. As part of this method, an efficient saddle-point search algorithm, such as the climbing image nudged elastic band (cNEB) method,[28] is used to determine the minimum energy paths and corresponding energy barriers for all high-temperature transitions from a given energy basin which occur before a well-defined high-temperature "stop-time" is reached. Accordingly, the efficiency with which activation barriers can be determined plays an important role in determining the speed of TAD simulations.

Here, we present a method by which a localized saddle-point (LSAD) search may be used to reduce the cost of calculating energy-barriers in TAD simulations, thus improving the scaling behavior of TAD as a function of system size without significantly reducing the accuracy of the resulting calculation. Our method is based on the assumption that only atoms which are sufficiently close to those involved in the transition have a significant effect on the potential energy surface along the minimum energy path. As a result, by carefully choosing the range of atoms to be included in the saddle-point calculation, the error resulting from localization can be minimized while the computational cost of saddle-point calculations is significantly reduced. Since in many TAD simulations, saddle-point calculations constitute a significant portion of the computational time, the development of such a method may be considered to be the first step in the development of a computationally efficient form of TAD with improved scaling as a function of system size.

To demonstrate the improved performance of saddle-point calculations with our LSAD method, we have first applied it to single- and multiatom moves found in Cu/Cu(100) and Cu bulk simulations. In particular, for a Cu monomer hopping on a Cu(100) substrate, the LSAD method yields N-independent scaling, instead of $N^{1.8}$ based on the whole system. In addition, we have carried out TAD simulations of Ag/Ag(100) annealing and Cu/Cu(100) growth with and without the LSAD method. As expected, the performance of TAD with the LSAD method is significantly better than without LSAD, while the relative performance increases with system size.

The organization of this paper is as follows. In Sec. II, we describe the LSAD method in more detail and also present results for the dependence of the error in the activation barrier as a function of chunk size along with results for the dependence of the cutoff radius on the number of atoms involved in a transition for both homoepitaxial and heteroepitaxial systems. In Sec. III, we present additional results for the performance of the LSAD method in activation barrier calculations for a variety of multiatom moves on the Cu(100) surface and Cu bulk TAD simulations. We then discuss the performance of TAD simulations of Ag/Ag(100) annealing and Cu/Cu(100) growth carried out with our LSAD method.

Finally, Sec. IV contains a summary of our results and brief discussion.

II. LOCALIZED SADDLE-POINT SEARCH

The underlying assumption on which our LSAD method is based is that atoms sufficiently far away from those directly involved in an activated event will have a minimal effect on the minimum energy path and activation energy. As a result, the corresponding saddle-point calculations may be accelerated without significant loss of accuracy by only including a small group of atoms (consisting of all atoms which are sufficiently close to those directly involved in the transition) in our calculations. We now describe our implementation of the LSAD method in more detail.

In order to determine the atoms to be included in our saddle-point calculations, we first identify those atoms which "participate" in a transition (green atoms in Fig. 1) as corresponding to those atoms for whom the displacement (in the minimized final configuration) compared to the initial minimized configuration is larger than a critical distance d_c, where typically d_c is of the order of a few tenths of Angstroms (here we have used $d_c = 0.3$ Å). The additional moving atoms to be included in LSAD calculations (blue atoms in Fig. 1) then correspond to those whose minimized positions are within a distance r_m from either the initial or final positions of the atoms participating in the transition. Since in our calculations we have typically used embedded-atom-method (EAM) potentials[33] for which there is a well-defined cutoff distance r_{cut}, an additional group of fixed atoms, corresponding to all atoms which are within a distance r_{cut} from any of the moving atoms, is also included (yellow atoms in Fig. 1 with $w_{fixed} = r_{cut}$).

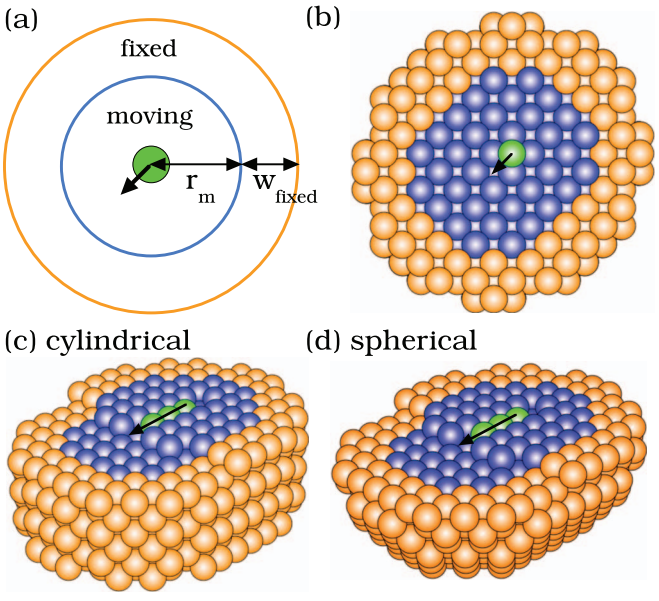

FIG. 1. (a) Schematic diagram of LSAD method describing moving and fixed regions around a monomer (green atom). (b) Top view of the chunk constructed by the LSAD method where blue and orange colors represent the moving and fixed regions, respectively. (c) Cylindrical and (d) spherical chunks for a three-atom slide-out move near a step.

It is natural to express the moving atom cutoff in the form,

$$r_{\mathrm{m},i} = d_m \times r_{\mathrm{cut},i}, \tag{1}$$

where d_m is typically slightly less than 2 and i represents the atom-type. In Eq. (1), one may use different values of $r_{\mathrm{cut,\,i}}$ for different types of atoms, but for simplicity if there is more than one type of atom present, we have used the maximum value of $r_{\mathrm{cut,\,i}}$. Since we are interested in surface problems such as thin-film growth as well as bulk problems, we have considered both a cylindrical cutoff (where the distance r_m is the distance to the cylinder axis which is perpendicular to the surface) as well as a spherical cutoff. We now consider the accuracy and performance of our LSAD calculations.

III. RESULTS

In order to test the accuracy and performance of our LSAD method, we have examined a variety of activated processes observed in homoepitaxial and heteroepitaxial simulations on fcc metal(100) surfaces and have also carried out TAD annealing and thin-film growth simulations with and without LSAD. In addition, we have carried out spatially parallel TAD (parTAD) simulations of radiation damage in Cu, using a three-dimensional extension of our parTAD method[34,35] based on our synchronous sublattice algorithm.[36] In all of our simulations, EAM potentials were used, including potentials developed by Voter-Chen (VC),[37] Mishin *et al.*[38] (Mishin EAM), Sheng *et al.*[39] (Sheng EAM), Adams *et al.*[40] (AFW EAM), and Mishin-Voter-Bonney (MVB) EAM[41] for Cu-Ni potential.

In our growth and annealing simulations, we used a substrate composed of five or six layers with the top two or three moving and bottom three fixed and lateral sizes in the range $5a$–$12a$ in units of the lattice constant a (where $a = 4.09$ Å, 3.615 Å, and 3.52 Å for Ag, Cu, and Ni, respectively), while periodic boundary conditions (PBCs) were applied parallel to the substrate. In contrast, PBCs were applied in all three directions in our three-dimensional parallel TAD simulations of Cu radiation damage.[34,36]

In our TAD simulations, the minimum prefactor was set to $\nu_{\min} = 5 \times 10^{11}$ s^{-1} and the uncertainty $\delta = 0.05$ was used. Our high-temperature MD simulations were carried out using a Langevin thermostat[42] with a friction coefficient of 10^{12} s^{-1} and a time step of 4 fs. Energy barriers for attempted events were calculated using the cNEB method with 11 images while Vineyard prefactors[43] were also measured for some representative cases. In our TAD simulations of Cu/Cu(100) growth at $T_{low} = 80$ K, atoms were deposited randomly with deposition rate $F = 1$ ML/s, while an initial kinetic energy of 0.2 eV was used for the depositing atoms. While most of our calculations were carried out using a 2 GHz Intel dual-core Mac mini, our annealing and growth simulations were carried out using the 2.6 GHz IBM 1350 Glenn cluster at the Ohio Supercomputing Center (OSC).

A. Dependence of activation barrier on r_m

Since our LSAD calculations use only a small portion of the whole system, it is important to determine to what

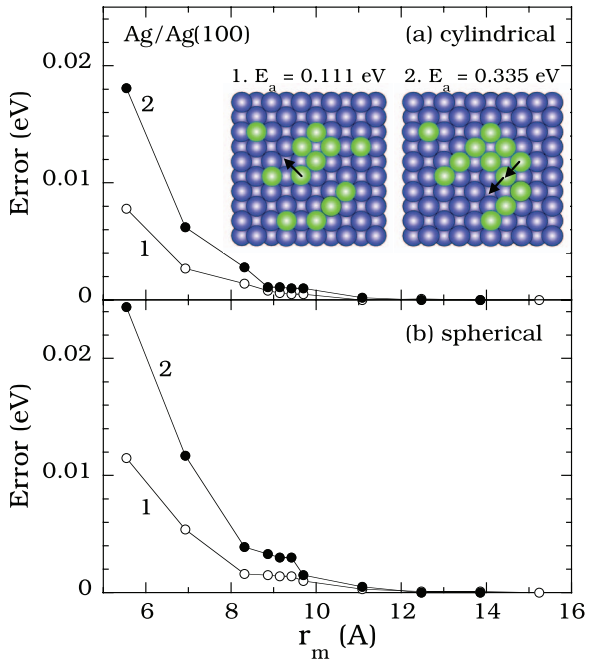

FIG. 2. Error in energy barrier obtained using VC EAM potential ($r_{\rm cut}$ = 5.542 Å) as a function of $r_{\rm m}$ for two Ag diffusion moves on Ag(100) shown for (a) cylindrical and (b) spherical cutoff. The barriers E_a = 0.111 eV and 0.335 eV in (a) are obtained from the whole system.

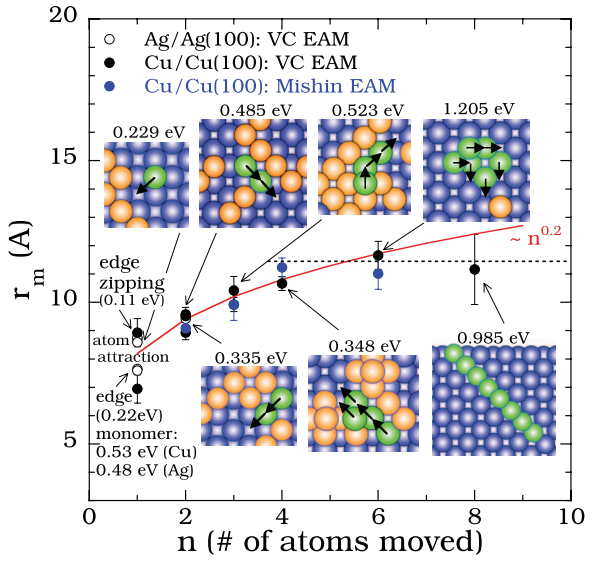

FIG. 3. Dependence of r_m (corresponding to an error in the activation barrier of 10^{-3} eV) on number of atoms n involved in transition for various Ag/Ag(100) and Cu/Cu(100) activated processes, assuming a cylindrical cutoff. Green atoms are those directly involved in transitions (arrows indicate diffusion directions), while neighboring atoms are orange. Dotted horizontal line corresponds to r_m = 11.4 Å.

extent this may lead to errors in the calculation of activation barriers. We first consider the two relatively low-barrier Ag/Ag(100) diffusion processes shown in Fig. 2(a)—one involving a single-atom "edge-zipping" move and the other involving a two-atom move—using the Voter-Chen EAM potential. As can be seen, restricting the number of moving atoms via a cutoff radius r_m typically leads to an activation barrier which is somewhat higher than that obtained using the entire system. This is not surprising since one expects "pinning" effects[44] to lead to an increase in the energy away from the local minimum. In addition, for a given value of $r_{\rm m}$, the cylindrical cutoff tends to yield a smaller error than the spherical cutoff since the cylindrical cutoff includes more atoms. However, in both cases the error in the activation barrier decays exponentially with increasing cutoff radius r_m. As a result, for both transitions (and for either a cylindrical or a spherical cutoff) we find that for r_m = 9.14 Å (corresponding to $d_{\rm m}$ = 1.65), the error in the activation barrier is smaller than 4×10^{-3} eV.

To determine how the cutoff r_m corresponding to a given error in the activation barrier depends on the number of atoms n involved in the transition, we have also examined a variety of additional activated processes observed in simulations of metal homoepitaxial and heteroepitaxial growth. In particular, Fig. 3 shows the value of the chunk radius $r_{\rm m}$ as a function of n (with n ranging from 1 to 8) needed to ensure an error of 10^{-3} eV or less in the activation barrier E_a obtained for a variety of Ag/Ag(100) and Cu/Cu(100) surface diffusion processes, including monomer diffusion, edge diffusion, two types of edge-zipping for Ag/Ag(100) (see Fig. 2(a).1 and Fig. 4.1 in Ref. 45), atom attraction (E_a = 0.229 eV for Ag/Ag(100)), and several multiatom moves such as the formation of a misfit dislocation on a Cu substrate by a Cu monomer. As can be seen in Fig. 3, the radius $r_{\rm m}$ appears to increase with n. In particular, for Cu/Cu(100), the dependence may be described by $r_m \sim n^{0.2}$ (i.e., $d_{\rm m} = c \times n^{0.2}$) for $n \leq 6$ with a constant $c \simeq 1.65$ for an error of 10^{-3} eV. For the Mishin *et al.* Cu EAM potential, which has a larger cutoff than the Voter-Chen potential, the radius r_m shows a similar increase with n but with slightly different energy barriers in each case. For Ag/Ag(100), the value of c is slightly smaller with $c \simeq 1.55$ due to a larger $r_{\rm cut}$.

The increase of r_m with n may be generally explained by an increase in the local distortions of atoms surrounding the transition with increasing n. However, in some cases the local distortion can be quite large even if a relatively small number of atoms (based on the initial and final configurations) are involved in the transition, and in these cases a somewhat larger chunk radius r_m may be needed even though n is small. One example is Ni monomer exchange with a Ni(100) substrate atom (corresponding to n = 2) for which the VC and Sheng Ni EAM potentials result in energy barriers of 1.17 eV and 0.85 eV with chunk sizes of r_m = 12.5 Å and 12.1 Å, respectively, for an error of 10^{-3} eV, while for the AFW EAM potential, r_m = 13.8 Å with an energy barrier of 1.37 eV. Similarly, for Cu/Cu(100), the same process yields r_m = 13.9 Å and a barrier of 0.80 eV with VC EAM potential. However, in all of these cases the distortion at the saddle-point is relatively large, and as a result, additional atoms (corresponding to those which have been displaced a distance larger than d_c at the saddle-point) also participate in the transition.

This suggests that in general it may be desirable to include the saddle-point configuration as well as the initial and final configurations in defining the number of atoms n involved in a transition. For example, an initial LSAD

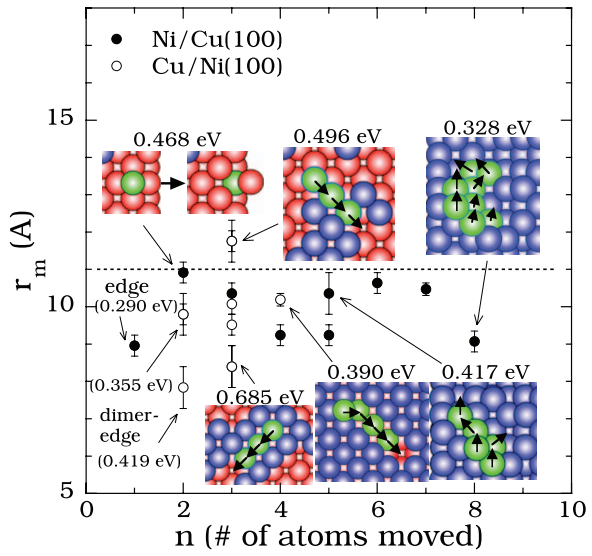

FIG. 4. Chunk radius r_m as function of number of atoms n involved in transitions corresponding to an error of 10^{-3} eV in the activation barrier for Ni/Cu(100) and Cu/Ni(100), where blue (red) atoms correspond to film (substrate) and green atoms represent atoms involved in the transition. Here we used the cylindrical chunk method along with MVB EAM. The horizontal dotted line represents a chunk radius of r_m, yielding an error smaller than 1.5×10^{-3} eV.

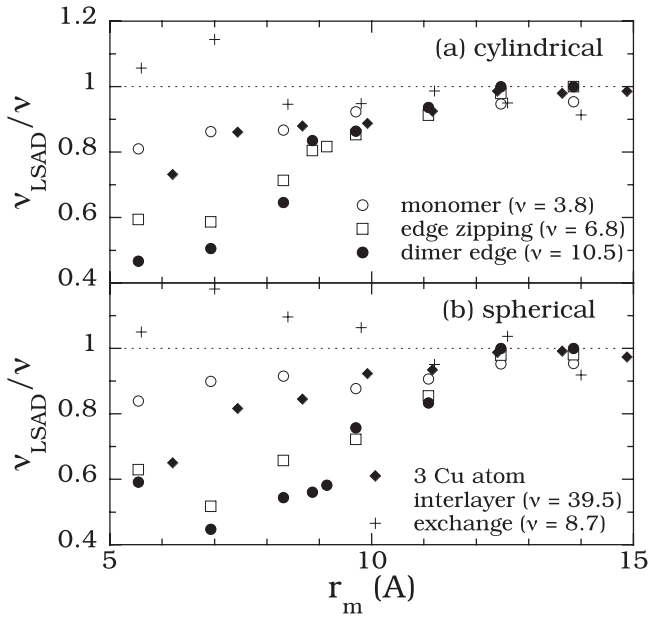

FIG. 5. Comparison of prefactors measured as function of chunk radius r_m with VC EAM, where ν (ν_{LSAD}) are prefactors in units of THz for whole (localized) system, respectively.

calculation may be carried out using a value of r_m based on the displacements of the initial and final configurations, followed by a more accurate LSAD calculation which takes into account any additional atoms which are displaced at the saddle-point. However, we note that in the applications of the LSAD method to low-temperature TAD simulations considered here, the dominant attempted events observed are single atom moves, while high-barrier processes such as the monomer exchange discussed above are unlikely to be accepted. Thus, we expect that at such low temperatures one may still use the power-law form for the chunk size to further improve the performance of TAD by carrying out LSAD calculations.

We now consider the dependence of r_m on n in heteroepitaxial systems. In particular, Fig. 4 shows the dependence of r_m corresponding to an error in the calculated activation barrier of 10^{-3} eV or less on the number of atoms n participating in the transition, for a variety of activated processes observed in Ni/Cu(100) and Cu/Ni(100) growth using the MVB EAM potential. Some of the representative diffusion moves are also shown along with their barriers. As can be seen, in this case the value of r_m needed to ensure an error of less than 10^{-3} eV in the calculated activation barrier shows no clear dependence on the number n of participating atoms. The lack of any trend with n in these heteroepitaxial systems may be due to the fact that the displacement near the saddle-point produces a relatively small distortion compared to the longer-range strain-field due to misfit strain.

B. Dependence of prefactor on r_m

While the prefactor for activated events is not typically calculated in TAD simulations (since the low-temperature time is determined by the high-temperature time combined with the activation energy), in some cases it may be of interest to efficiently calculate prefactors in the course of a TAD simulation. Accordingly, we have also examined the effects of the cutoff r_m on the prefactor for five different activated processes: (i) Ag/Ag(100) monomer, edge-zipping, and dimer edge diffusion where the latter two cases correspond to the Ag adatom moves shown in Fig. 2(a), (ii) a three Cu-atom interlayer diffusion move with activation energy 0.523 eV shown in Fig. 3, and (iii) a two-atom exchange move in Ni/Cu(100) with activation barrier 0.468 eV shown in Fig. 4. The prefactor ν was calculated using the Vineyard formula,[43] $\nu = \prod_{i=1}^{3N_p} \nu_i^i / \prod_{i=1}^{3N_p-1} \nu_i^s$, where N_p is the number of atoms used in the prefactor calculation, and ν^i and ν^s denote the vibration frequencies at the initial and transition states, respectively.

Figure 5 shows the ratio of ν_{LSAD} to ν as a function of the cutoff radius r_m, where ν and ν_{LSAD} are the prefactors measured using the whole system and "reduced" system for both a cylindrical and a spherical cutoff. We note that measuring a vibration frequency is a time consuming calculation whose computational cost increases as N_p^3. Thus, given the time-consuming nature of Vineyard calculations, the LSAD method should provide an efficient means for estimating the prefactor, especially for large systems. As can be seen in Fig. 5, the ratio ν_{LSAD}/ν converges somewhat more slowly with increasing r_m than the activation barrier while the rate of convergence depends on the activated process. In particular, for a cutoff $r_m = 11$ Å (corresponding to an error of 2×10^{-3} eV in the energy barrier), the corresponding prefactor ratio $\nu_{\mathrm{LSAD}}/\nu \simeq 0.9(0.8)$ with the cylindrical (spherical) chunk method. In addition, just as was found for the activation barrier, the minimum value of r_m needed so that the ratio is close to 1 also appears to increase with n.

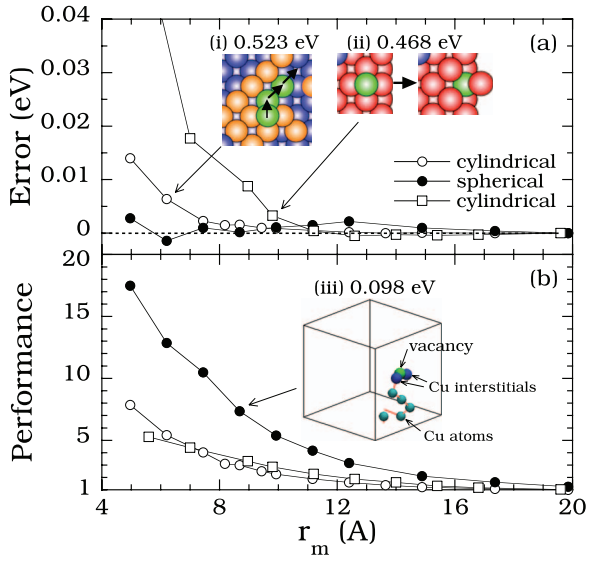

FIG. 6. Error in energy barrier and performance of LSAD method as function of r_m for (i) three Cu atom move on Cu(100) shown in Fig. 3, (ii) Ni adatom exchange move with a Cu substrate atom shown in Fig. 4, and (iii) seven Cu atom move found in TAD simulations of Cu radiation damage, where red lines represent displacements from their initial positions.

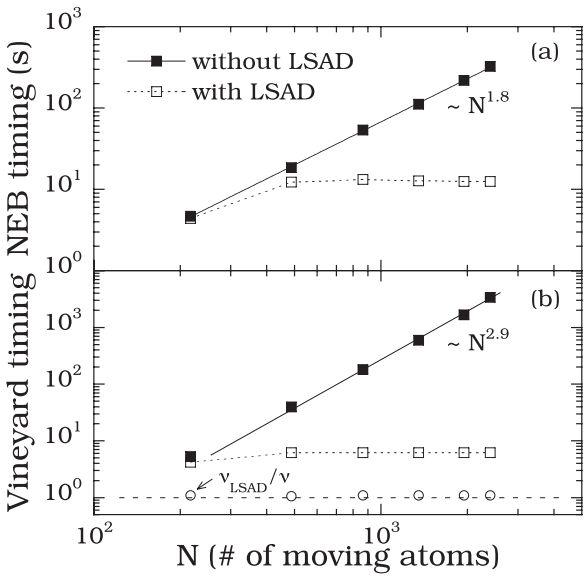

FIG. 7. Comparison of NEB and Vineyard timings measured in seconds and performance as a function of the number of moving atoms with and without the LSAD method, for the case of Cu monomer diffusion on a Cu(100) substrate with a cylindrical cutoff $r_m = 10.725$ Å.

C. Performance

We first consider the performance of the LSAD method in speeding up saddle-point calculations, where the performance can be defined as

$$\text{Performance} = \frac{t}{t_{\text{LSAD}}}, \tag{2}$$

and t and t_{LSAD} are the execution times obtained without and with the LSAD method, respectively. In particular, we have considered three representative cases, Cu surface diffusion on a $9a \times 9a$ substrate, a Ni adatom exchange move with a Cu substrate atom on a $12a \times 12a$ substrate, and a complex Cu atom move triggered by Cu interstitials found in our simulations of Cu radiation damage with a system size of $16a \times 16a \times 16a$. We note that for this bulk Cu system, the spherical cutoff is more appropriate.

Figure 6 shows the error in the energy barrier for these cases along with the corresponding performance as a function of the chunk radius r_m. As can be seen, while the performance decreases with increasing r_m, a performance factor ranging from 3 (for the first two cases) to 9 (for the case of Cu radiation damage) can be obtained with an error in the activation barrier of less than 3×10^{-3} eV. The superior performance for the case of radiation damage is due to the large system size used in the simulation. We also note that in our parTAD radiation damage simulation using eight processors, there are many low-barrier processes similar to the one shown in Fig. 6(b), which arise due to the presence of a large number of defects such as interstitials and vacancies, and which tend to introduce intermediate states for any moves with moderate or higher barriers. As a result, about 40% of the total TAD simulation time was spent resolving intermediate states and finding saddle-points and calculating their barriers. Thus, in this case the LSAD method can significantly improve the performance of TAD, while maintaining a negligibly small error in the energy barrier, as shown in Fig. 6.

We have also investigated the scaling of nudged elastic band (NEB) and Vineyard timings with system size (both with and without LSAD) for the case of Cu monomer diffusion on a Cu(100) substrate. In this case the Sheng EAM potential[39] was used while a cylindrical cutoff with $r_m = 10.725$ Å ($d_m = 1.65$) yields an error of 3×10^{-4} eV in the monomer diffusion barrier (0.5353 eV). For a fair comparison, the NEB timing without the LSAD method comes only from NEB calculations including barrier calculations. On the other hand, the NEB timing with the LSAD method also includes the computational cost of constructing chunks for the initial and final states along with their minimization. As can be seen in Fig. 7(a), while the NEB timing without LSAD increases as $t \sim N^{1.8}$ (where N is the total number of moving atoms in the system), the calculation time using LSAD exhibits no N dependence. Similarly, as shown in Fig. 7(b), the Vineyard timing without LSAD shows the expected N^3 behavior, while the Vineyard timing with LSAD method has no N dependence. Thus, for a large value of N, the performance (defined as the ratio of the calculation time without LSAD to that with LSAD) increases as $N^{1.8}$ and $N^{2.9}$ for NEB and Vineyard calculations, respectively, while for $N = 2401$, the performance factor is 26 (NEB) and 550 (Vineyard).

We now discuss the performance of TAD simulations carried out using our LSAD method. As a first example, we consider TAD simulations of Cu/Cu(100) growth at $T_{low} = 80$ K with deposition rate $F = 1$ ML/s on a $12a \times 12a$ substrate using the VC EAM potential. In our simulations, we used both cylindrical and spherical cutoffs with $r_m = 1.65 \times n^{0.2} r_{cut}$ and $T_{\text{high}} = 550$ K, which is the optimal fixed temperature in this case.[45] As can be seen in Fig. 8, the TAD simulations with LSAD are approximately 1.4 (2.2) times faster for cylindrical (spherical) cutoffs than those without LSAD. We have also

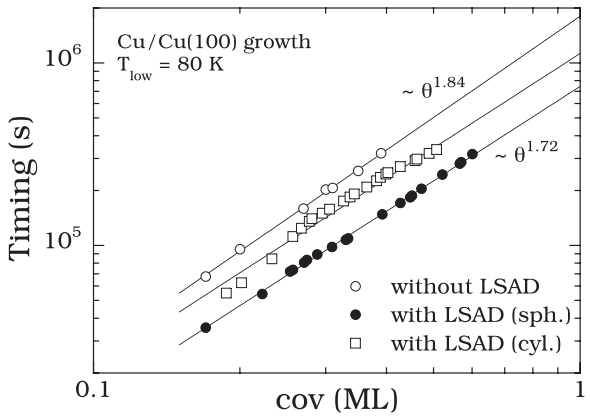

FIG. 8. Timing comparison as function of coverage in TAD simulations of Cu/Cu(100) growth at $T_{\text{low}} = 80$ K with and without LSAD, where a substrate size of $12a \times 12a$ and $d_m = 1.65 \times n^{0.2}$ with $r_{\text{cut}} = 4.961$ Å were used for both cylindrical and spherical chunks.

carried out adaptive TAD[45] simulations of Ag/Ag(100) annealing at $T_{low} = 80$ K using the VC EAM potential. In this case (see the supplementary material[46]), the LSAD method improves the performance of adaptive TAD by an additional 12%. This somewhat modest improvement is due to the small substrate size ($5a \times 5a$) used in the simulation.

IV. SUMMARY

We have presented an efficient method to speed up the calculation of activation barriers and Vineyard prefactors in large systems, by localizing the saddle-point calculation so that only atoms which are within a specified distance r_m from those atoms participating in the transition are included. In addition, by carefully examining the dependence of the error on cutoff r_m, we have shown that the errors introduced by such a cutoff can be minimized while still obtaining significant speed-up.

In particular, to understand the effects of the chunk radius r_m on the error in the activation energy, we have examined a variety of activated events found in Cu/Cu(100) and Ag/Ag(100) homoepitaxial and Ni/Cu(100) and Cu/Ni(100) heteroepitaxial simulations. For the homoepitaxial case, we have found that at least at low temperatures (for which any events deviating from such a power-law are unlikely to be accepted due to their high energy barriers) the cutoff radius satisfies the power-law form $r_m \sim n^{0.2} r_{cut}$. In contrast, we found that for Ni/Cu(100) and Cu/Ni(100), it is better to use a fixed chunk size in order to take into account the effects of long-range strain field created by lattice mismatch. In addition, we have also examined the effects of chunk size on the Vineyard prefactor. While the convergence is relatively slow and exhibits a strong dependence on the type of activated process, the LSAD method can still be used to efficiently estimate the prefactor.

While the results presented here were obtained using EAM potentials, we expect that in general the LSAD method may also be applied to systems with more complex interactions and may also be useful in *ab initio* calculations. In this connection we note that for the EAM potentials considered here, a cutoff radius which corresponds roughly to the fourth or fifth neighbor distance (for Ag and Cu, respectively) was found to be sufficient to ensure that the error in the activation energy is less than 10^{-3} eV. This suggests that a similar criterion may be used in other cases. Nevertheless, in general it is still desirable to perform some preliminary calculations for some representative cases as was done in Fig. 2, in order to determine r_m, before actually carrying out simulations with the LSAD method. On the other hand, if a somewhat lower accuracy (e.g., 0.005 eV) is required, which may be appropriate for simulations at higher temperature and is also consistent with the accuracy of most interaction potentials, then in general we expect that a cutoff corresponding to the fourth neighbor distance may be adequate.

We note that for computational efficiency we have restricted our definition of atoms "participating" in the transition to include only those atoms whose displacements in the final configuration are larger than a critical value d_c of the order of a few tenths of Angstroms. However, in general it may be desirable to include the saddle-point configuration as well as the initial and final configurations in defining the number of atoms n involved in a transition. For example, an initial LSAD calculation may be carried out using a value of r_m based on the displacements of the initial and final configurations, followed by a more accurate LSAD calculation if additional atoms are displaced at the saddle-point.

Finally, we have demonstrated that the performance of TAD simulations can be significantly improved by the use of LSAD calculations by considering two examples, Ag/Ag(100) annealing and Cu/Cu(100) growth. In both cases, we have found that the speed-up increases with increasing system size. As a result, one can achieve a significant speed-up in TAD simulations with larger system sizes, and also significantly improve the scaling of TAD with system size, while maintaining a negligibly small error in the energy barrier.

ACKNOWLEDGMENTS

This work was supported by National Science Foundation (NSF) grant DMR-0907399 as well as by a grant of computer time from the Ohio Supercomputer Center. N. B. Callahan gratefully acknowledges NSF funding from the Research Experience for Undergraduates (REU) program in the Department of Physics & Astronomy at the University of Toledo.

[1] A. F. Voter, F. Montalenti, and T. C. Germann, Annu. Rev. Mater. Res. **32**, 321 (2002).
[2] D. Perez, B. P. Uberuaga, Y. Shim, J. G. Amar, and A. F. Voter, Annu. Rep. Comp. Chem. **5**, 79 (2009).
[3] A. F. Voter, J. Chem. Phys. **106**, 4665 (1997).
[4] A. F. Voter, Phys. Rev. Lett. **78**, 3908 (1997).
[5] R. A. Miron and K. A. Fichthorn, Phys. Rev. Lett. **93**, 128301 (2004).
[6] K. A. Fichthorn, R. A. Miron, Y. Wang, and Y. Tiwary, J. Phys. Condens. Matter **21**, 084212 (2009).
[7] A. F. Voter, Phys. Rev. B **57**, R13985 (1998).
[8] M. R. Sørensen and A. F. Voter, J. Chem. Phys. **112**, 9599 (2000).
[9] F. Montalenti and A. F. Voter, J. Chem. Phys. **116**, 4819 (2002).
[10] F. Montalenti, M. R. Sørensen, and A. F. Voter, Phys. Rev. Lett. **87**, 126101 (2001).
[11] Y. Fan, A. Kushima, S. Yip, and B. Yildiz, Phys. Rev. Lett. **106**, 125501 (2011).
[12] S. Hara and J. Li, Phys. Rev. B **82**, 184114 (2010).

[13]A. Ishii, S. Ogata, H. Kimizuka, and J. Li, Phy. Rev. B **85**, 064303 (2012).
[14]G. Henkelman and H. Jónsson, J. Chem. Phys. **115**, 9657 (2001).
[15]L. Xu and G. Henkelman, J. Chem. Phys. **129**, 114104 (2008).
[16]O. Trushin, A. Karim, A. Kara, and T. S. Rahman, Phys. Rev. B **72**, 115401 (2005).
[17]F. El-Mellouhi, N. Mousseau, and L. J. Lewis, Phys. Rev. B **78**, 153202 (2008).
[18]D. Konwar, V. J. Bhute, and A. Chatterjee, J. Chem. Phys. **135**, 174103 (2011).
[19]R. Terrell, M. Welborn, S. T. Chill, and G. Henkelman, J. Chem. Phys. **137**, 014105 (2012).
[20]G. Nandipati, A. Kara, S. Islamuddin, and T. S. Rahman, J. Comput. Phys. **231**, 3548 (2012).
[21]G. A. Tribello, M. Ceriotti, and M. Parrinello, Proc. Natl. Acad. Sci. USA **107**, 17509 (2010).
[22]M. Eckert, E. Neyts, and A. Bogaerts, Cryst. Eng. Comm. **11**, 1597 (2009).
[23]A. Kushima, X. Lin, J. Li, J. Eapen, J. C. Mauro, X. Qian, P. Diep, and S. Yip, J. Chem. Phys. **130**, 224504 (2009).
[24]A. Kushima, J. Eapen, J. Li, S. Yip, and T. Zhu, Eur. Phys. J. B **82**, 271 (2011).
[25]A. Laio and F. L. Gervasio, Proc. Natl. Acad. Sci. USA **99**, 12562 (2002); Rep. Prog. Phys. **71**, 126601 (2008).
[26]B. M. Dickson, Phys. Rev. E **84**, 037701 (2011).
[27]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), p. 385.
[28]G. Henkelman, B. P. Uberuaga, and H. Jónsson, J. Chem. Phys. **113**, 9901 (2000).
[29]D. Sheppard, P. Xiao, W. Chemelewski, D. D. Johnson, and G. Henkelman, J. Chem. Phys. **136**, 074103 (2012).
[30]G. Henkelman and H. Jónsson, J. Chem. Phys. **111**, 7010 (1999).
[31]R. Malek and N. Mousseau, Phys. Rev. E **62**, 7723 (2000).
[32]W. E, W. Ren, and E. Vanden-Eijnden, Phys. Rev. B **66**, 052301 (2002); J. Chem. Phys. **126**, 164103 (2007); W. E and E. Vanden-Eijnden, Annu. Rev. Phys. Chem. **61**, 391 (2010); M. Cameron, R. V. Kohn, and E. Vanden-Eijnden, J. Nonlinear Sci. **21**, 193 (2011).
[33]M. S. Daw and M. I. Baskes, Phys. Rev. B **29**, 6443 (1984).
[34]Y. Shim, J. G. Amar, B. P. Uberuaga, and A. F. Voter, Phys. Rev. B **76**, 205439 (2007).
[35]Y. Shim, V. Borovikov, B. P. Uberuaga, A. F. Voter, and J. G. Amar, Phys. Rev. Lett. **101**, 116101 (2008).
[36]Y. Shim and J. G. Amar, Phys. Rev. B **71**, 125432 (2005).
[37]A. F. Voter, and S. P. Chen, Mater. Res. Soc. Symp. Proc. **82**, 175 (1987); A. F. Voter, in *Intermetallic Compounds: Principles and Practice*, edited by J. H. Westbrook and R. L. Fleischer (Wiley, New York, 1995), Vol. 1, p. 77.
[38]Y. Mishin, M. J. Mehl, D. A. Papaconstantopoulos, A. F. Voter, and J. D. Kress, Phys. Rev. B **63**, 224106 (2001).
[39]H. W. Sheng, M. J. Kramer, A. Cadien, T. Fujita, and M. W. Chen, Phys. Rev. B **83**, 134118 (2011).
[40]J. B. Adams, S. M. Foiles, and W. G. Wolfer, J. Mater. Res. **4**, 102 (1989).
[41]For the Ni (Cu) EAM potential, see Ref. 37 (Ref. 38). For the Cu-Ni cross potential, see G. Bonny *et al*., Philos. Mag. **89**, 3531 (2009).
[42]M. P. Allen and D. J. Tildesley, *Computer Simulations of Liquids* (Oxford, New York, 1987), p. 263.
[43]G. H. Vineyard, J. Phys. Chem. Solids **3**, 121 (1957).
[44]Y. Shim and J. G. Amar, Phys. Rev. B **83**, 245419 (2011).
[45]Y. Shim and J. G. Amar, J. Chem. Phys. **134**, 054127 (2011).
[46]See supplementary material at http://dx.doi.org/10.1063/1.4793218 for the performance of adaptive TAD simulations.