

Accelerated molecular dynamics through stochastic iterations and collective variable based basin identification

Pratyush Tiwary* and Axel van de Walle

School of Engineering, Brown University, Providence, Rhode Island 02912, USA

(Received 27 December 2012; published 18 March 2013)

We present a method called SISYPHUS (stochastic iterations to strengthen yield of path hopping over upper states) for extending accessible time scales in atomistic simulations. The method proceeds by separating phase space into basins and transition regions between the basins based on a general collective variable (CV) criterion. The transition regions are treated via traditional molecular dynamics (MD) while Monte Carlo (MC) methods are used to (i) estimate the expected time spent in each basin and (ii) thermalize the system between two MD episodes. In particular, an efficient adiabatic switching based scheme is used to estimate the time spent inside the basins. The method offers various advantages over existing approaches in terms of (i) providing an accurate real time scale, (ii) avoiding reliance on harmonic transition state theory, and (iii) avoiding the need to enumerate all possible transition events. Applications of SISYPHUS to low-temperature vacancy diffusion in bcc Ta and adatom island ripening in fcc Al are presented. A CV appropriate for such condensed phases, especially for transitions involving collective motions of several atoms, is also introduced.

DOI: [10.1103/PhysRevB.87.094304](https://doi.org/10.1103/PhysRevB.87.094304)

PACS number(s): 71.15.Pd, 02.70.Ns, 05.10.Ln, 68.43.Hn

Achieving usefully long timescales (seconds or longer) in atomistic simulations of materials is a problem of great interest, and the search for a practical and general solution has generated intense activity in the field over the last several decades^{1–9} (see Refs. 10–12 for excellent reviews of these and other efforts). The problem arises because, as the system moves from one energy basin to another through infrequent rare events, it stays trapped in some energy basin for extended periods of time. Along with the small time steps (on the order of femtoseconds) needed for the total energy staying conserved, this severely restricts the time scales accessible in molecular dynamics (MD) simulations and also leads to limited phase-space exploration.

Although many methods exist to increase the rate of rare events and efficiently explore the rugged free energy landscape, a frequent limitation is the inability to efficiently obtain an accurate estimate of the “real” time scale of the simulation in general physical systems. Existing schemes to achieve this either

- (1) require cataloguing all possible transitions paths out of a given basin,^{4–6} which can be computationally prohibitive in low-symmetry systems, or
- (2) rely on harmonic approximations to the system’s energy surface,¹³ or
- (3) involve computing averages that converge slowly, especially for large system sizes, because they involve exponentials of the system’s total energy.^{2,14}

In this paper, we propose a mixed Monte Carlo–molecular dynamics method that uses a collective variable (CV) χ solely to discriminate between basins and transition regions, thus placing very weak requirements on the choice of CV, less stringent than what is required in metadynamics methods.³² The method still provides the system’s dynamics via conventional atomic coordinates and thus provides more detailed information than the coarse grained dynamics typically provided by metadynamics methods.^{1,15} The method also provides an accurate real time scale that does not deteriorate with system size and that does not rely on harmonicity assumptions. There

is no need to construct *a priori* or *in situ* a catalog of possible transition mechanisms. The method is specially suited for exploiting massive parallel computing.

The proposed algorithm generalizes our previous work³ along multiple dimensions. First, we use a general CV χ (instead of the system’s potential energy) to discriminate between basins and transition regions and propose a type of CV suitable for this purpose in condensed phases. (Recently proposed dimensionality reduction algorithms^{16,17} that discover CVs automatically could be used as well.) Second, we introduce an adiabatic switching scheme to efficiently calculate the real time spent inside wells. Finally, we use a more robust criterion to determine when the system has been trapped in a basin for sufficiently long time to have equilibrated therein.

Let the N -particle system be characterized by position $x = (r_1, \dots, r_N)$ and velocity $v = (v_1, \dots, v_N)$. The CV χ is a function of x (we give an example of such a function later in this paper). For a user-specified cutoff value χ_{cut} , we define the basins, or wells, \mathbf{W} in this χ space as a set of connected states for which $\chi < \chi_{\text{cut}}$. In the wells, the method does not follow the system’s exact trajectory in phase space, but instead provides the expected amount of time spent in each well. In contrast, the $\chi \geq \chi_{\text{cut}}$ region of the phase space contains the interesting but infrequently occurring events whose dynamics is fully described. The choice of χ_{cut} thus specifies the level of details one wishes to retain in the simulation. Large values of χ_{cut} may cause wells to merge and limit the ability to resolve the precise dynamics of some events. The method is still formally correct but the definition of the wells \mathbf{W} may not have an obvious physical meaning. Nevertheless, the method is very robust to changes in χ_{cut} that do not change the topology of the well structure (we will provide numerical evidence of this fact).

When the CV χ of the system is above the threshold χ_{cut} , the system evolves via MD according to its true Hamiltonian with constant-pressure (or volume) or constant-temperature (or energy) ensemble as entailed by the simulation. When $\chi < \chi_{\text{cut}}$, the algorithm continues performing MD until either

(i) $\chi \geq \chi_{\text{cut}}$ (in which case the system is considered to have exited the well and standard MD continues as described above) or (ii) a time equal to the system's decorrelation time τ_c has elapsed. In that latter case, the system is considered to have been trapped in the well for long enough that it has reached a local thermodynamic equilibrium. This criterion is similar to the one used in other methods^{8,18} to define a transition event. At this time, we launch a Monte Carlo (MC) simulation (called *MCA*) whose aim is to generate a new random starting position at the well's boundary to initialize the next MD episode. *MCA* is run for long enough that the system loses memory of how it entered the well and visits the boundary of the well a few times. MD is restarted with the position x where the system last visited the well's boundary and with velocities v drawn from a truncated Maxwell-Boltzmann distribution conditional on $v \cdot \nabla \chi(x) > 0$ (i.e., we only consider velocities in the half space pointing outwards of the well).

There is an inherent assumption here that the system is thermalized in whichever basin it entered before leaving it. This assumption is well justified because we have already verified that regular MD stays in this well for longer than the system's correlation time [around τ_c ; see (ii) above]. Hence, by definition, we assume thermalization only when thermalization did take place.

In parallel to the first MC (*MCA*) we perform a second MC run (called *MCB*) that calculates the expected time t_W the system would have spent inside the well. This separation of two MC runs makes our algorithm extremely parallelizable and especially amenable to be used on loosely coupled clusters of computers. We can launch as many *MCB* runs as we have processors available. These runs do not need to communicate with each other, and because of the system's ergodicity, we can make a quick estimate of the quantity t_W by averaging over these independent runs. This parallelization is even simpler than for the parallel replica method.¹⁸

Before we describe how we calculate the time the system would have spent in whichever well it visited, we describe specifically how *MCA* is implemented. We define the boundary of the well \mathbf{W} with thickness w :

$$S_W = \{x = (r_1, \dots, r_N) : |\chi(x) - \chi_{\text{cut}}| < w\}. \quad (1)$$

Trying to visit S_W with an unbiased potential would be of no avail, since we will rarely visit states in S_W . We thus do MC with a biased potential $V^*(x)$ defined as follows:

$$V^*(x) = V(x) + \begin{cases} \infty & \chi(x) \geq \chi_{\text{cut}}, \\ V_0 \left(\frac{\chi_{\text{cut}} - \chi}{\chi_{\text{cut}}} \right)^m & \chi(x) < \chi_{\text{cut}}. \end{cases} \quad (2)$$

Per Eq. (2), *MCA* never visits the outside of the well \mathbf{W} and biases states with a penalty function that increases with their depth inside the well. Note that the bias is zero at the well boundary, which is important to obtain the correct sampling distribution of the boundary.² The bias inside the well changes the fraction of time spent at the boundary but not the ratio of the times spent at any two points of the boundary. The parameter m determines how sharp the boundary of the well is (a value around 0.5 is found to perform well in practice). V_0 is kept around the standard deviation in potential energy of the system. As we demonstrate numerically later, the

algorithm is very robust with respect to choice of parameters in Eq. (2).

Having described a way to accelerate the exploration of various wells, we now turn to the question of calculating (via a Monte Carlo labeled *Mcb*) the expected time the system would have spent inside well \mathbf{W} if there had been no acceleration of the dynamics. This time, denoted t_W , can be calculated as the reciprocal of the flux of states exiting the well:

$$t_W = \lim_{w \rightarrow 0} \left(\left\langle \frac{\bar{v}}{w} 1(x \in S_W) \right\rangle \right)^{-1}, \quad (3)$$

where the average $\langle \dots \rangle$ is taken over x drawn from the well \mathbf{W} with a probability density proportional to $e^{-V(x)/(k_B T)}$. k_B is Boltzmann's constant, T is the temperature, and $1(A)$ equals 1 if the event A is true and 0 otherwise. \bar{v} denotes the mean projection of a Maxwell-Boltzmann-distributed velocity along the unit vector u parallel to $\nabla \chi(x)$, conditional on $v \cdot \nabla \chi(x) > 0$. When all atoms have the same mass m , $\bar{v} = \sqrt{k_B T / 2\pi m}$ (a general expression can be found in Ref. 3).

A straightforward implementation of Eq. (3) will again suffer from a rare event problem. Even an importance sampling scheme, as suggested in Ref. 3, is not very efficient. Consider what happens if we use the biased potential as defined in Eq. (2) to calculate t_W :

$$t_W = \lim_{w \rightarrow 0} \frac{\langle e^{-\beta(V(x) - V^*(x))} \rangle^*}{\left(\frac{\bar{v}}{w} e^{-\beta(V(x) - V^*(x))} 1(x \in S_W) \right)^*}, \quad (4)$$

where $\langle \dots \rangle^*$ denote expectations taken under a density proportional to $e^{-\beta V^*(x)}$, in which β is $1/(k_B T)$. This approach is exact in the limit of an ensemble average, but there is a fundamental trade-off that limits its usefulness: A large bias $V^*(x) - V(x)$ leads to a more rapid convergence of the denominator (due to an increased sampling rate of the boundary) but a slower convergence of the numerator [due to an increase in $V^*(x) - V(x)$].

To avoid this problem, we now propose a technique that bears some resemblance to adiabatic switching methods,^{19,20} in which the system is continuously, *adiabatically* switched from $V(x)$ (the true potential) to $V^*(x)$ (identical to the potential used in *MCA*). Let $\hat{V}(x, \alpha)$ smoothly interpolate between $\hat{V}(x, 0) \equiv V(x)$ and $\hat{V}(x, 1) \equiv V^*(x)$. Then we can express the ensemble average in Eq. (4) as below (working in terms of rate t_W^{-1}):

$$\begin{aligned} t_W^{-1} &= \lim_{w \rightarrow 0} \frac{\int \frac{\bar{v}}{w} 1(x \in S_w) e^{-\beta \hat{V}(x, 0)} dx}{\int e^{-\beta \hat{V}(x, 0)} dx} \\ &\equiv \lim_{w \rightarrow 0} \left\langle \frac{\bar{v} 1(x \in S_w)}{w} e^{-\beta(\hat{V}(x, 0) - \hat{V}(x, 1))} \right\rangle_1 R, \end{aligned} \quad (5)$$

where dx denotes a differential volume in $3N$ dimensional configuration space for N particles, the integration being performed over the entire configuration space within the well \mathbf{W} and the expected value $\langle \dots \rangle_\alpha$ in Eq. (5) defined by

$$\langle \dots \rangle_\alpha = \frac{\int (\dots) e^{-\beta \hat{V}(x, \alpha)} dx}{\int e^{-\beta \hat{V}(x, \alpha)} dx}. \quad (6)$$

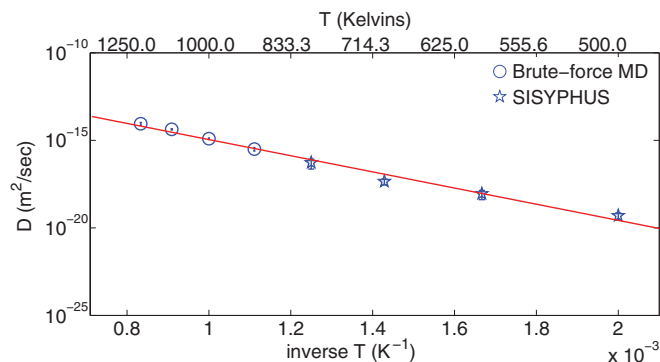


FIG. 1. (Color online) Diffusion constant for vacancy diffusion in Ta at various temperatures as through brute-force MD (circles) and SISYPHUS (stars). Error bars (roughly same as marker size) are also provided as obtained over 16 independent runs.

Below we define the term R in Eq. (5) and reexpress it in a computationally tractable form:

$$\begin{aligned}
 R &= \frac{\int e^{-\beta \hat{V}(x,1)} dx}{\int e^{-\beta \hat{V}(x,0)} dx} \\
 &= \exp \left(\ln \int e^{-\beta \hat{V}(x,1)} dx - \ln \int e^{-\beta \hat{V}(x,0)} dx \right) \\
 &= \exp \left(\int_0^1 \left(\frac{\partial}{\partial \alpha} \ln \int e^{-\beta \hat{V}(x,\alpha)} dx \right) d\alpha \right) \\
 &= \exp \left(-\beta \int_0^1 \frac{\int \frac{\partial \hat{V}(x,\alpha)}{\partial \alpha} e^{-\beta \hat{V}(x,\alpha)} dx}{\int e^{-\beta \hat{V}(x,\alpha)} dx} d\alpha \right) \\
 &= \exp \left(-\beta \int_0^1 \left\langle \frac{\partial \hat{V}(x,\alpha)}{\partial \alpha} \right\rangle_{\alpha} d\alpha \right). \quad (7)
 \end{aligned}$$

We pick a linear switching scheme for $\hat{V}(x,\alpha)$, i.e., an interpolation scheme between $\hat{V}(x,0)$ and $\hat{V}(x,1)$:

$$\hat{V}(x,\alpha) = (1 - \alpha)V(x) + \alpha V^*(x). \quad (8)$$

We now make a few observations regarding Eq. (5). It involves 2 parts. The first is $\lim_{w \rightarrow 0} \langle \frac{\mathbb{1}(x \in S_w)}{w} e^{-\beta(\hat{V}(x,0) - \hat{V}(x,1))} \rangle_1$. This is nonzero only when $x \in S_w$, and whenever it is nonzero, the difference $\hat{V}(x,0) - \hat{V}(x,1)$ is very small [see Eqs. (1) and (2)]. Since this average is calculated with the maximally biased potential $\hat{V}(x,1)$, the boundary $x \in S_w$ is visited frequently, and thus the first term in Eq. (5) can be evaluated very quickly. The second part in Eq. (5) is R , where the average $\langle \frac{\partial \hat{V}(x,\alpha)}{\partial \alpha} \rangle_{\alpha} = \langle V^*(x) - V(x) \rangle_{\alpha}$ does not contain any exponentials that could cause a slow convergence. In the Supplemental Material,³¹ we prove rigorously that for this switching scheme and for the choice of biasing potential in Eq. (2) we can use a nonuniform grid to evaluate R which can be made finer as $\alpha \rightarrow 0$ but kept coarse for larger α , leading to further computational efficiency.

For solid-state systems where bond breaking is the dominant mechanism of interest, we take χ to be the bond distortion function (BDF), defined below for an N -particle system:

$$\chi(x) = a_0 \left(\sum_{\langle i,j \rangle} \frac{|r_{ij} - r_{ij}^{eq}|^p}{r_{ij}^{eq}} \right)^{1/p}, \quad (9)$$

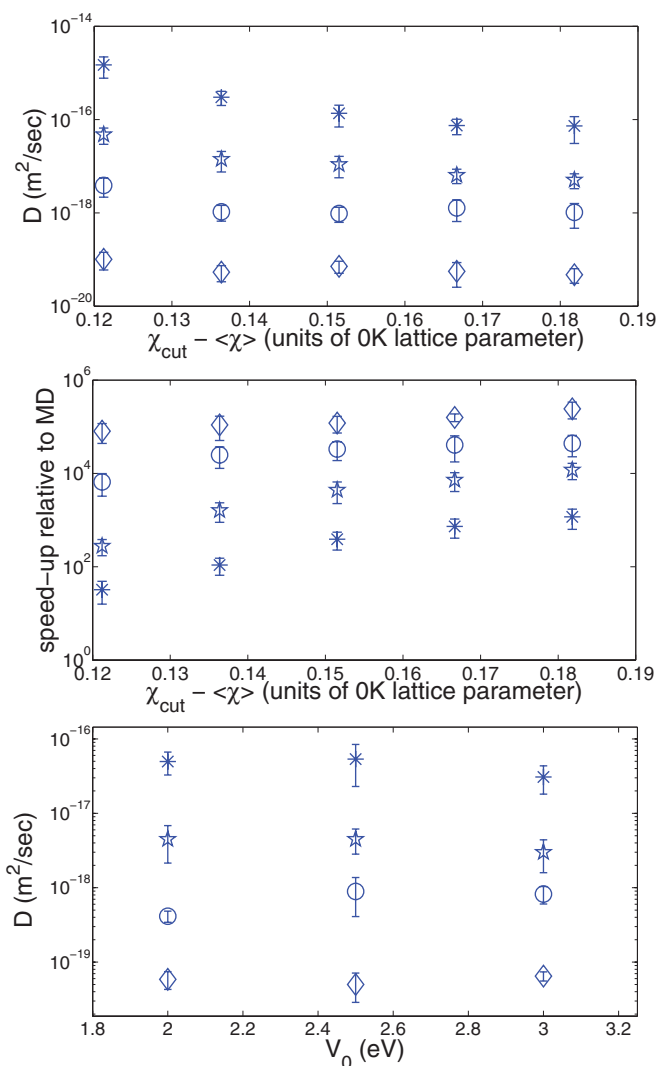


FIG. 2. (Color online) Top: Insensitivity of dynamics to choice of χ_{cut} (relative to average χ at that temperature) for vacancy diffusion in Ta across temperatures. \diamond , \circ , \star , and $*$ denote 500, 600, 700, and 800 K temperatures, respectively. Middle: Corresponding speed-ups relative to physical time achieved in brute-force MD run in the same wall-clock time. Bottom: Insensitivity of dynamics to choice of V_0 . Error bars over 16 independent runs for each data point.

where $p > 1$ and a_0 is the 0 K lattice parameter. For each bond, r_{ij}^{eq} is the equilibrium bond length that can be obtained by a few conjugate gradient steps each time a transition is detected. In the limit that $p \rightarrow \infty$, the BDF, which is a p norm over fractional bond distortions, approaches the maximum norm, i.e., the strain (times a_0) in the maximally strained bond. For $p \rightarrow \infty$ we thus recover the so-called bond-boost function.²¹ We pick p around 8–12 for the systems studied in this paper. Not taking $p = \infty$ allows us to treat on a similar footing cases where (a) one bond is distorted by a large amount, or (b) several bonds are collectively distorted by a significant amount which is however less than the amount in (a). As soon as either (a) or (b) happens, the BDF detects it through a spike and thus we can then switch back to doing completely unbiased MD. This way we can treat transition mechanisms involving small but concerted and collective motion of several atoms [see

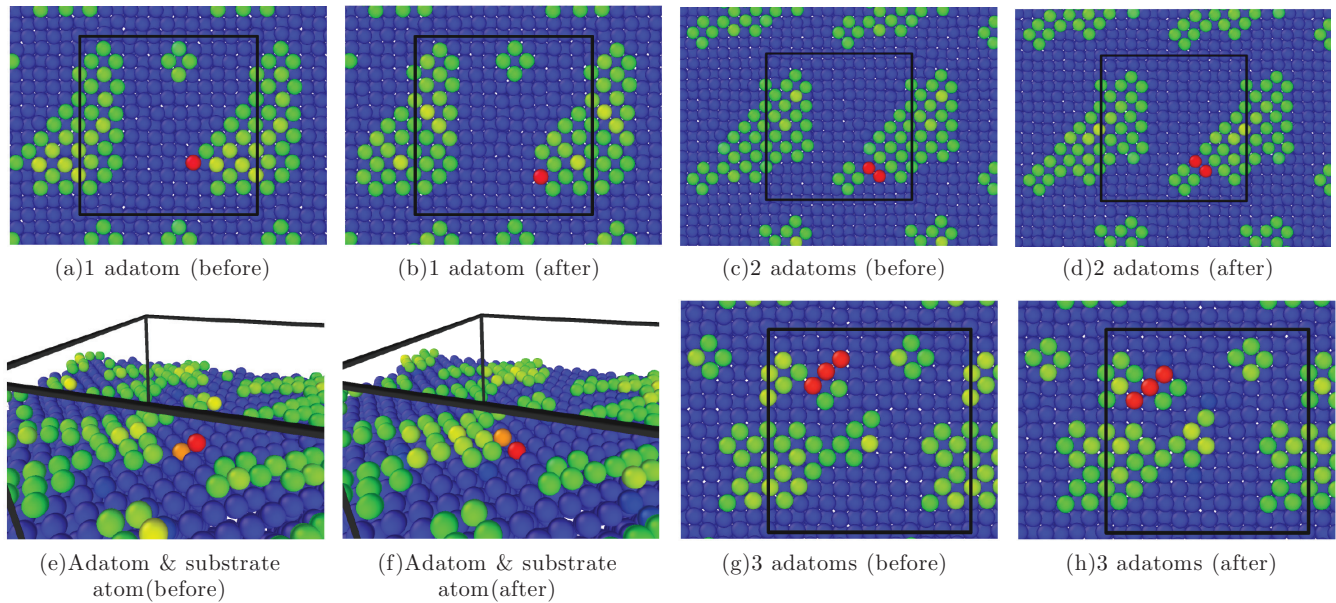


FIG. 3. (Color online) Mechanisms seen by SISYPHUS for adatom movement on Al (001) surface. Atoms colored per z coordinate. Blue/orange = substrate atom, green/red = adatom, red/orange = atom with maximum movement. Solid lines are periodic boundaries.

Figs. 3(g) and 3(h)]; for example in glasses, shear transformation zones²² involve several atoms displacing together by a small amount.

We now describe applications of the algorithm that validate it and demonstrate its insensitivity to choice of parameters. The first example is vacancy-assisted lattice diffusion at low temperatures in bcc tantalum. Lattice diffusion at low temperatures is a problem important in a spectrum of sciences from materials science to geology,^{23,24} but is beyond the time scales one can access in current MD simulations, requiring times longer than milliseconds.^{3,25} We describe the parameters for the MD part²⁶ of this simulation in the Supplemental Material.³¹ In Fig. 1 we demonstrate how the diffusion constants through brute-force MD and SISYPHUS (stochastic iterations to strengthen yield of path hopping over upper states) for vacancy-assisted diffusion lie on the same Arrhenius plot, giving an Arrhenius-type activation energy of $0.9(\pm 0.1)$ eV [in rough agreement with the harmonic transition state theory (HTST) calculation of 1.1 eV]. Figure 2 (top) demonstrates the lack of sensitivity of the dynamics to which χ_{cut} and V_0 values we pick. As expected the speed-up is higher for higher χ_{cut} [Fig. 2 (middle)]. At higher temperature we find slightly higher sensitivity to χ_{cut} (still within order of magnitude accuracy). This is because for a low χ_{cut} , the t_W values at high temperatures become closer to time τ_c the system takes to equilibrate in a well. Insensitivity to V_0 values can be seen from Fig. 2 (bottom). A smaller V_0 leads to slower sampling in Eq. (4), and MCa also takes longer to converge. With too high a V_0 , however, the periphery of the well \mathbf{W} can be too steep and one might again face sampling issues since the system can be trapped in some regions of the well boundary. In the Supplemental Material,³¹ we provide a back-of-the-envelope estimate of how to pick an optimum V_0 for a given system.

For our second application, we studied the room-temperature dynamics of Al adatoms on Al (001) thin film (625 atoms, roughly 3 times larger than the Ta vacancy

diffusion example). We picked this problem because first, from a technological perspective, it is of immense importance for fabrication processes in nanoscale devices involving growth of thin films from deposited adatoms.^{27,28} This problem is very interesting from a theoretical perspective too, given that it is an inherently nonequilibrium phenomenon dictated by the interplay between kinetics and thermodynamics. Being able to model and control the growth and properties of such films is hugely desirable; the time scales needed are however far beyond MD. Accurate 0 K saddle point search methods^{5,29} have shown the existence of a large number of transition mechanisms with low and similar activation energies (smaller than 0.4 eV). Such low lying barriers can be hard to deal with in most accelerated MD methods.¹⁰ On-the-fly kinetic Monte Carlo⁵ calculations have been used previously to get rough estimates of the time scale for adatom island ripening that we can compare SISYPHUS with. In the Supplemental Material³¹ we provide details of the simulation parameters for this system.

Figure 3 and the movies in the Supplemental Material³¹ illustrate the most common mechanisms seen through SISYPHUS. We have the single-adatom hop [(a) and (b)], the concerted two-adatom hop [(c) and (d)], and the concerted event involving an adatom and a substrate atom as the latter moves to the adatom layer [(e) and (f)]. Occasionally we see more complicated mechanisms like the 3-atom hop [(g) and (h)], and events with creation of a vacancy in the top substrate layer (see Supplemental Material³¹). The first three mechanisms are the most common and are in fact the lowest energy transitions found using the dimer method for saddle point search. In Fig. 4 and the movies in the Supplemental Material³¹ we show the typical evolution of the island ripening at room temperature over several milliseconds of physical time (the exact time can vary from run to run well but still within an order of magnitude). Shown along with are corresponding energies obtained by quenching each structure to its local minima, illustrating further the effectiveness of the algorithm

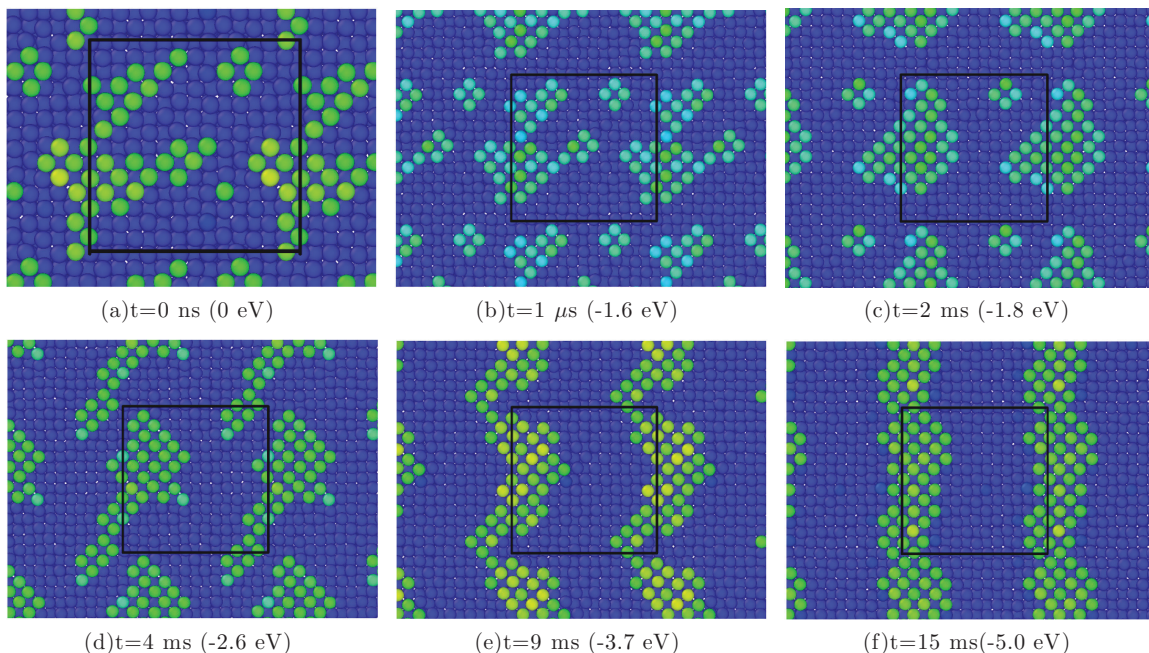


FIG. 4. (Color online) Room-temperature adatom ripening on Al (001) surface as a function of time. Each structure was quenched to find its corresponding local minima, and the energies after quenching [relative to (a)] are provided. Panel (a) shows the starting geometry. After $1 \mu\text{s}$, we have two disconnected clusters [(b) and (c)]. Corresponding brute-force MD runs were found trapped in similar configuration in the fraction of microsecond time they could achieve. At around 4 ms these two clusters join (d). At 9 ms (e) there is further joining and we have effectively one long chain of adatoms. At 15 ms (f) the chain has coarsened into one entity across simulation cells. Color scheme same as in Fig. 3.

in escaping and exploring various energy basins. The overall time can be compared with the analogous on-the-fly KMC work⁵ for the same system and interatomic potential³⁰ where 20 adatoms (half as many as current work) formed one compact cluster around 1 ms. As a verification that our proposed BDF is effective at decomposing phase space into disconnected wells, we show, in a movie accompanying the Supplemental Material,³¹ a superposition of snapshots of the system during a MCA run, illustrating that the system does not jump from one well to another within one MCA run (since it rejects all moves with $\chi \geq \chi_{\text{cut}}$).

In conclusion, we have shown SISYPHUS to be an extremely parallelizable and robust set of algorithms that help achieve a fraction of a second time scale for thousands of atoms. The method works well irrespective of system size and can be applied in the general setting of any collective

variable. We also introduced a CV appropriate for solid state systems especially for transitions involving collective motions of several atoms.

We would like to thank Arthur Voter and Graeme Henkelman for helpful discussions and stimulating the trajectory of this research, Christopher Weinberger for the HTST calculation for Ta vacancy diffusion, and Julia Greer for hospitality during P.T.'s visit to Caltech. This research was supported by the US National Science Foundation through XSEDE computational resources provided by NCSA under Grants No. DMR050013N and No. DMR120056, NSF Condensed Matter and Materials Theory Program No. DMR0907669, and the Office of Naval Research through Grant No. N00014-11-1-0886. The open-source and easy-to-use software package SISYPHUS is available for download.³³

*Current address: Computational Science, Department of Chemistry and Applied Biosciences, ETH Zurich, c/o USI Campus, via Buffi 13, CH-6900 Lugano, Switzerland.

¹A. Laio and M. Parrinello, *Proc. Natl. Acad. Sci. USA* **99**, 12562 (2002).

²A. F. Voter, *Phys. Rev. Lett.* **78**, 3908 (1997).

³P. Tiwary and A. van de Walle, *Phys. Rev. B* **84**, 100301 (2011).

⁴Y. Fan, A. Kushima, S. Yip, and B. Yildiz, *Phys. Rev. Lett.* **106**, 125501 (2011).

⁵G. Henkelman and H. Jonsson, *J. Chem. Phys.* **115**, 9657 (2001).

⁶G. T. Barkema and N. Mousseau, *Phys. Rev. Lett.* **77**, 4358 (1996).

⁷K. Lindorff-Larsen, S. Piana, R. O. Dror, and D. E. Shaw, *Science* **334**, 517 (2011).

⁸C.-Y. Lu, D. E. Makarov, and G. Henkelman, *J. Chem. Phys.* **133**, 201101 (2010).

⁹X.-M. Bai, A. F. Voter, R. G. Hoagland, M. Nastasi, and B. P. Uberuaga, *Science* **327**, 1631 (2010).

- ¹⁰A. Voter, F. Montalenti, and T. Germann, *Annu. Rev. Mater. Res.* **32**, 321 (2002).
- ¹¹A. Laio and F. L. Gervasio, *Rep. Prog. Phys.* **71**, 126601 (2008).
- ¹²D. Perez, B. Uberuaga, Y. Shim, J. Amar, and A. Voter, *Annu. Rep. Comput. Chem.* **5**, 79 (2009).
- ¹³M. Sørensen and A. Voter, *J. Chem. Phys.* **112**, 9599 (2000).
- ¹⁴C. Jarzynski, *Phys. Rev. E* **73**, 046105 (2006).
- ¹⁵G. Bussi, A. Laio, and M. Parrinello, *Phys. Rev. Lett.* **96**, 090601 (2006).
- ¹⁶G. A. Tribello, M. Ceriotti, and M. Parrinello, *Proc. Natl. Acad. Sci. USA* **107**, 17509 (2010).
- ¹⁷G. A. Tribello, M. Ceriotti, and M. Parrinello, *Proc. Natl. Acad. Sci. USA* **109**, 5196 (2012).
- ¹⁸A. F. Voter, *Phys. Rev. B* **57**, R13985 (1998).
- ¹⁹D. Frenkel and B. Smit, *Understanding Molecular Simulation: From Algorithms to Applications* (Academic Press, 2002), 2nd ed.
- ²⁰C. Jarzynski, *Phys. Rev. Lett.* **78**, 2690 (1997).
- ²¹K. A. Fichthorn, R. A. Miron, Y. Wang, and Y. Tiwary, *J. Phys.: Condens. Matter* **21**, 084212 (2009).
- ²²M. L. Falk and J. S. Langer, *Phys. Rev. E* **57**, 7192 (1998).
- ²³S. Mrowec and S. Marcinkiewicz, *Defects and Diffusion in Solids: An Introduction* (Elsevier, 1980).
- ²⁴E. B. Watson and E. F. Baxter, *Earth Planet. Sci. Lett.* **253**, 307 (2007).
- ²⁵M. I. Mendeleev and Y. Mishin, *Phys. Rev. B* **80**, 144111 (2009).
- ²⁶G. Ackland and R. Thetford, *Philos. Mag. A* **56**, 15 (1987).
- ²⁷Z. Zhang and M. G. Lagally, *Science* **276**, 377 (1997).
- ²⁸K. E. Becker, M. H. Mignogna, and K. A. Fichthorn, *Phys. Rev. Lett.* **102**, 046101 (2009).
- ²⁹G. Henkelman and H. Jonsson, *J. Chem. Phys.* **111**, 7010 (1999).
- ³⁰A. Voter and S. Chen, *Proc. Mater. Res. Soc.* **82**, 175 (1987).
- ³¹See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevB.87.094304> for details.
- ³²In metadynamics, the CV should distinguish between initial basin, final basin, and transition regions. Our CV however is allowed (but not required) to take the same value in two basins.
- ³³See <http://alum.mit.edu/www/avdw/sisyphus>.