

Hybrid deterministic and stochastic approach for efficient atomistic simulations at long time scales

Pratyush Tiwary*

Department of Applied Physics and Materials Science, California Institute of Technology, Pasadena, California 91125, USA

Axel van de Walle

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

(Received 30 August 2011; published 23 September 2011)

We propose a hybrid deterministic and stochastic approach to achieve extended time scales in atomistic simulations that combines the strengths of molecular dynamics (MD) and Monte Carlo (MC) simulations in an easy-to-implement way. The method exploits the rare event nature of the dynamics similar to most current accelerated MD approaches but goes beyond them by providing, without any further computational overhead, (a) rapid thermalization between infrequent events, thereby minimizing spurious correlations, and (b) control over accuracy of time-scale correction, while still providing similar or higher boosts in computational efficiency. We present two applications of the method: (a) Vacancy-mediated diffusion in Fe yields correct diffusivities over a wide range of temperatures and (b) source-controlled plasticity and deformation behavior in Au nanopillars at realistic strain rates ($10^4/s$ and lower), with excellent agreement with previous theoretical predictions and *in situ* high-resolution transmission electron microscopy observations. The method gives several orders-of-magnitude improvements in computational efficiency relative to standard MD and good scalability with the size of the system.

DOI: [10.1103/PhysRevB.84.100301](https://doi.org/10.1103/PhysRevB.84.100301)

PACS number(s): 71.15.Pd, 62.25.-g, 02.70.Ns, 87.10.Tf

With vast improvements in the quality of available interatomic force fields and computer power, the classical molecular dynamics (MD) simulation has seen a dramatic increase in its use across a variety of fields over the past few decades.¹⁻⁴ One of the features that makes MD so appealing is its ability to actually follow the dynamical evolution of the system, thus giving insight into the microscopic behavior of the material. However, this is where the major limitation of MD comes into light, too: Most of the interesting dynamics occurs as the system moves from one energy basin to another through infrequent rare events, while the system remains stuck in some energy basin for extended periods of time. This non-ergodicity, coupled with the small time steps (on the order of femtoseconds) needed for the total energy staying conserved, severely restricts the time scales accessible in MD simulations and also leads to limited phase-space exploration. There have been many attempts at addressing this time-scale problem in MD—examples include metadynamics,⁵ the activation relaxation technique,⁶ parallel replica dynamics, temperature-accelerated dynamics, and hyperdynamics.^{7,8} There are several excellent reviews such as Ref. 7 available on the subject. The hyperdynamics method⁸ offers an elegant and practical way to increase the rate of infrequent events. It consists of adding a potential energy bias that makes the potential wells, in which the system normally remains trapped for extended periods, less deep. A time-scale correction is also evaluated in terms of the bias potential. The hyperdynamics method, especially with the advent of a variety of easy-to-implement biasing forms,⁹ has seen several compelling applications over the past years.⁹⁻¹³ Our approach in this Rapid Communication builds upon the crucial insights of Voter and co-workers while seeking improvements along two important dimensions. First, it bypasses a fundamental trade-off present in hyperdynamics: A shallower potential well provides faster dynamics but, at the same time, reduces the ability of the modified potential

to properly thermalize the system in between the infrequent events, resulting in artificial correlation between these events. Second, our method provides better independent control over the accuracy of time-scale correction, while hyperdynamics time-scale estimates can remain noisy up to long simulation times, especially for large system sizes (see Ref. 14 for a discussion).

Let the state of the system be characterized by position x and velocity v , each being a $3N$ -dimensional vector for a system of N atoms. When the potential energy $V(x)$ of the system is above a threshold V_0 , the system evolves via constant-energy (or constant-temperature) MD according to its true Hamiltonian (Fig. 1). This high-energy region of the phase space contains the interesting but infrequently occurring events. The method is formally correct for any choice of V_0 ; a higher choice of V_0 merely limits our ability to monitor the detailed dynamics of some events. When the system's potential energy falls below V_0 , two MC simulations are initiated (denoted a and b). Simulation a runs MC with a perfectly uniform potential inside the potential well \mathbf{W} , consisting of the points x where the true potential energy $V(x)$ lies below V_0 (i.e., all moves are accepted as long as they do not go outside the well). Simulation a is run until the system is well thermalized and has lost memory of how it entered the well (this takes a few MC passes, an insignificant amount of wall-clock time). MD then resumes with positions drawn from the last MC state that visited the boundary of the potential well. The vector v of the velocities of all atoms for restarting MD is drawn from a Maxwell-Boltzmann distribution corresponding to the temperature T of interest, conditional on $v \cdot \nabla V(x) > 0$ (i.e., we only consider velocities in the half space pointing outward of the well). MC simulation a is first of the crucial differences between our approach and hyperdynamics: It ensures proper thermalization of the system between rare events even when using a completely flat potential in the well. Even though it is




FIG. 1. (Color online) Schematic and flowchart of the algorithm. Shell S_w of constant potential energy and energy well W as described in the text are shown here. See Eq. (6) for a definition of \bar{t}_W .

done with a uniform potential, it does not lead to the molecular structure being completely lost since we rule out all moves that lead to energy higher than V_0 .

In parallel to simulation a , another MC simulation b is launched to estimate the mean time the system should have spent in the well W . Akin to simulation a , b also rejects all moves that land outside the well W . The mean time spent in W is given by the reciprocal of the flux exiting¹⁵ the well W :

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

where the average $\langle \dots \rangle$ is taken over x drawn from the well W with a probability density proportional to $e^{-V(x)/(k_B T)}$, where k_B is Boltzmann's constant and the following definitions hold: $1(A)$ equals 1 if the event A is true and 0 otherwise, S_w is a shell of width w at the boundary of the well W , which can be defined in the limit of small w as

$$S_w = \{x : |V(x) - V_0| \leq w |\nabla V(x)|/2\}, \quad (2)$$

and \bar{v} denotes the mean projection of a Maxwell-Boltzmann-distributed velocity along the unit vector u parallel to $\nabla V(x)$, conditional on $v \cdot u > 0$. The latter is given by

$$\bar{v} = \sqrt{\frac{k_B T}{2\pi} \sum_{i=1}^N \frac{|u_i|^2}{m_i}}, \quad (3)$$

where m_i is the mass of atom i and $|u_i|$ denotes length of the three-dimensional subvector of u associated with atom i . Note that Eq. (3) reduces to the familiar expression¹⁵ $\bar{v} = \sqrt{k_B T/2\pi m}$ when all atoms have the same mass, in which case \bar{v} factors out of the average in (1). Since Eq. (1) involves an average, it can be approximated by MC simulations. However, the most straightforward implementation of this approach would be very inefficient because x would rarely visit the boundary S_w of the well. The efficiency can be considerably improved by using a biased potential $V^*(x)$ which is the same as the real potential in the high-energy regions [i.e., regions outside well W with $V(x) \geq V_0$], but lifted up in the deep energy basins. With this Eq. (1) becomes

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

where the averages $\langle \dots \rangle^*$ are taken over x drawn from the well W with a probability density proportional to $e^{-V^*(x)/(k_B T)}$ and $\beta = 1/k_B T$. MC simulation b is the second main difference with hyperdynamics: It provides separate control over the

accuracy of the speed-up factor since the length of the MC simulation b can be adjusted independently of the length of the whole simulation.

The form of biasing we use is a well-established and easy-to-implement biasing potential used in several implementations of Voter's hyperdynamics method, proposed by Hamelberg *et al.*:⁹

$$V^*(x) = V(x) + \begin{cases} 0, & V(x) \geq V_0, \\ \frac{(V_0 - V(x))^2}{\alpha + V_0 - V(x)}, & V(x) < V_0. \end{cases} \quad (5)$$

The times t_W obtained via MC simulations b can be directly added to the physical time spent doing MD simulations to yield the overall physical time of the simulation. However, refinements of the method can yield further improvement in efficiency. Instead of computing t_W for each well W , one may keep a running average

$$\bar{t}_W = \frac{1}{n_b} \sum_W t_W \quad (6)$$

of the time spent in the n_b wells sampled via MC simulation b (n_a , the number of wells actually visited, may well far exceed n_b). Once this average is converged, there is no need to initiate MC simulation b anymore. The overall time spent in all the wells will simply be $\bar{t}_W * n_a/n_b$. Note that there is no need to keep separate averages for different types of wells, which would have been difficult to implement. Although MC simulations a still need to be performed for all wells, these converge much more rapidly. Other efficiency improvements can be obtained by not performing fully converged MC simulations b and exploiting the fact that errors will average out over wells in Eq. (6). Note that this scheme must be used while ensuring that the biasing potential is sufficiently strong so that most of the random errors in Eq. (4) are concentrated in the numerator, to avoid a systematic bias due to nonlinearity of the ratio. We would like to point out that only the parameter w is additional to those in any typical hyperdynamics scheme (Hamelberg *et al.*'s⁹ in this case), the choice of which does not effect the result since we extrapolate t_W to the limit of small w .¹⁵ Our approach compares favorably with hyperdynamics⁸ where one does not have control over the accuracy of the accelerated time (hyperdynamics relies on this error canceling out over time but does not provide an estimate of how much it is^{8,14}), and one is obliged to keep performing dynamics with the biased potential at all stages of the calculation. Thus our method offers boosts as high as those that one could get from setting $\alpha = 0$ in Eq. (5) (akin to the flooding scheme^{16,17}), but still avoiding the slow convergence in time and problems with discontinuous forces that one encounters in implementing flooding based hyperdynamics. In addition, we avoid errors from sampling the system in the state when it is not thermalized between two rare events—once MD is relaunched in our scheme, the system is well thermalized by virtue of simulation a . To minimize the wall-clock time needed for calculation of time in Eq. (4) via simulation b , we use an optimal extent of biasing as suggested in Ref. 9. This involves setting $\alpha \simeq V_0 - V_{\min}$, which allows the biased potential to capture the shape of the potential wells.⁹ α smaller than this would improve sampling of the numerator in Eq. (4) but deteriorate that of the denominator. We picked two problems




FIG. 2. (Color online) $\log_{10}(\text{diffusivity})$ vs inverse temperature. (a) The straight lines denote Mendeleev *et al.*'s (Ref. 18) MD calculations. These are valid only until 700 K. (b) Asterisks denote diffusivity measurements per our approach. (c) The dashed line shows experimental measurements (Ref. 19) valid between 1000 and 1200 K.

to demonstrate that our method yields correct dynamics: (a) vacancy diffusion in bcc Fe at room temperature, and (b) deformation behavior in Au nanopillars at realistic strain rates.

Lattice diffusion at low temperatures is beyond the time scales one can access in current MD simulations, with most investigations¹⁸ going only beyond 700 K. The system we consider is 249 Fe atoms ($5 \times 5 \times 5$ bcc supercell with one vacancy) interacting through the embedded atom method (EAM) potential.¹⁸ For the MD part here and in deformation behavior problem, we performed *NVT* simulations using a time step of 2×10^{-15} s and a Langevin thermostat with a coupling constant 1×10^{-11} s⁻¹. The biasing parameter α was 50 eV. The V_0 values we used at 500 and 300 K were -975.5 and -984 eV, respectively (4 and 2.5 eV more than the mean energy at 500 and 300 K, respectively). We took the equilibrium concentration of defects¹⁸ to convert our effective diffusivity into equilibrium diffusivity. In Fig. 2 we plot the equilibrium diffusivity as obtained from (a) MD simulations,¹⁸ (b) our proposed approach, and (c) experimental measurements¹⁹ that include contributions from interstitial migrations also and hence are only slightly higher than both ours and MD values. We obtain an approximately five orders-of-magnitude boost, with similar speedup factors for system sizes up to 30 000 atoms.

For our second problem (see Fig. 3), we looked at the deformation behavior of Au nanopillars. With the advent of excellent *in situ* transmission electron microscopy (TEM) and high-resolution TEM (HRTEM) tools, there are many elegant experiments on sub-10-nm-sized crystals.^{20,23,24} Deformation in such small sizes is controlled by dislocation nucleation, and has been phenomenologically predicted²⁵ and experimentally found^{20,23,24} to have small activation volumes and strong strain-rate sensitivity. However, to the best of our knowledge, there is no direct MD-based confirmation of this strong strain-rate sensitivity due to the inability of MD to reach strain rates lower than 10^7 /s.




FIG. 3. (Color online) Simulation cell for stress-strain calculations. (a) Prior to application of any strain, (b) after yielding (strain = 12%) with strain rate = 5×10^4 /s. In (b) the leading partial has nucleated on a $\{111\}$ slip plane, leaving behind the two-layer-thick hcp region denoting an intrinsic stacking fault. Failure is thus through slip and not twinning, in agreement with Ref. 23. Atoms are identified as per bond order parameter Q_6 (Refs. 21 and 22). Perfect hcp atoms have been removed for clarity.

Using our method we were able to reach a 10^4 /s strain-rate regime with only ~ 48 h of computer time. We could also obtain several correct qualitative and quantitative aspects of the deformation dynamics, without assuming anything about the nature of deformation. The system we consider is 2016 Au atoms (a cylinder with height 7.4 nm and diameter 2.5 nm) interacting through an EAM potential.²⁶ The biasing parameter α was 1500 eV, while the starting V_0 value used was -7266 eV (8 eV more than the mean energy at 300 K), adjusted every 1000 MD steps to take into account the pressure-volume work contribution from the stress. The cylinder was initially carved out from perfect fcc lattice [Fig. 3(a)]. Periodic boundary conditions were imposed only along the cylinder axis z , which is the same as the compression axis $\langle 001 \rangle$. The cylinder was first equilibrated for 500 ps before beginning the compression carried out by uniformly rescaling the z coordinates of all atoms. The atomic virial stress²⁷ was used to obtain the Cauchy stress. Four different strain rates $\dot{\epsilon}$ were considered: 5×10^6 /s, 2.5×10^6 /s, 5×10^5 /s, and 5×10^4 /s. We present the resulting stress (σ)-strain (ϵ) plots in Fig. 4(a). Several conclusions can be drawn from Figs. 3 and 4 that prove our algorithm is capable of predicting correct dynamics in




FIG. 4. (Color online) (a) Stress-strain plots at three different strain rates: 2.5×10^6 /s (open circles), 5×10^5 /s (asterisks), 5×10^4 /s (pluses). The initial stress corresponding to zero strain is nonzero due to surface effects (Ref. 27). (b) $\log_{10}(\text{stress})$ at 11% strain (relative to surface stress at zero strain) vs \log_{10} (strain rate).

complicated systems. The yielding occurs at $\sim 10\%$ strain, and is through slip and not twinning or elastic instabilities: A leading partial nucleates on a $\{111\}$ slip plane at lower stresses than a trailing partial. This can be seen in Fig. 3(b), where the leading partial nucleated from the surface and left behind a two-layer-thick hcp region which again changes back to fcc after the trailing partial also nucleates at higher stresses and recombines with the leading partial. Figure 3(b) is identical to HRTEM images for the $\langle 001 \rangle$ loading of Au nanowires.²³ The strain rate sensitivity m in the relation $\sigma = \sigma_0 \dot{\epsilon}^m$ (derived by looking at stress at 11% strain) is approximately 0.14 ± 0.07 [see Fig. 4(b)], while Ref. 20 reports it to be ~ 0.11 for 75-nm-diam pillars. The activation volume Ω for the deformation as calculated through²⁰ $\Omega = \sqrt{3}k_B T \partial(lg\dot{\epsilon})/\partial\sigma$ is $\sim 1b^3$ ($b =$ Burgers vector), in excellent agreement with experimental observations.^{20,25}

To summarize, we have proposed an approach that combines the strengths of MC and MD, thus offering boosts of several orders of magnitudes with good system size scaling. We have applied the method to study lattice diffusion in bcc Fe at low temperatures and the deformation of Au nanopillars at low strain rates, and found it to work really well in both cases, predicting correct dynamics and exhibiting good scaling with an increase in the system size from 249 to 2016 atoms. We thus expect the method to be useful in a variety of situations.

This research was supported by the US National Science Foundation through TeraGrid resources provided by NCSA under Grant No. DMR050013N, through the US Department of Energy, National Energy Research Initiative for Consortia (NERI-C) Grant No. DE-FG07-07ID14893.

*pt@caltech.edu

¹S. J. Zhou, D. L. Preston, P. S. Lomdahl, and D. M. Beazley, *Science* **279**, 1525 (1998).

²M. Karplus and G. A. Petsko, *Nature (London)* **347**, 631 (1990).

³S. Piana and A. Laio, *J. Phys. Chem. B* **111**, 4553 (2007).

⁴T. T. Lau, A. Kushima, and S. Yip, *Phys. Rev. Lett.* **104**, 175501 (2010).

⁵A. Laio and M. Parrinello, *Proc. Natl. Acad. Sci. USA* **99**, 12562 (2002).

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

⁷A. F. Voter, F. Montalenti, and T. C. Germann, *Annu. Rev. Mater. Res.* **32**, 321 (2002).

⁸A. F. Voter, *Phys. Rev. Lett.* **78**, 3908 (1997).

⁹D. Hamelberg, J. Mongan, and J. Andrew McCammon, *J. Chem. Phys.* **120**, 11919 (2004).

¹⁰K. E. Becker, M. H. Mignogna, and K. A. Fichthorn, *Phys. Rev. Lett.* **102**, 046101 (2009).

¹¹S. Hara and J. Li, *Phys. Rev. B* **82**, 184114 (2010).

¹²K. A. Fichthorn, R. A. Miron, Y. Wang, and Y. Tiwary, *J. Phys. Condens. Matter* **21**, 084212 (2009).

¹³K. E. Becker, M. H. Mignogna, and K. A. Fichthorn, *Phys. Rev. Lett.* **102**, 046101 (2009).

¹⁴R. A. Miron and K. A. Fichthorn, *J. Chem. Phys.* **119**, 6210 (2003).

¹⁵A. F. Voter and J. D. Doll, *J. Chem. Phys.* **80**, 5832 (1984).

¹⁶H. Grubmüller, *Phys. Rev. E* **52**, 2893 (1995).

¹⁷M. M. Steiner, P.-A. Genilloud, and J. W. Wilkins, *Phys. Rev. B* **57**, 10236 (1998).

¹⁸M. I. Mendeleev and Y. Mishin, *Phys. Rev. B* **80**, 144111 (2009).

¹⁹M. Lubbehusen and H. Mehrer, *Acta Metall. Mater.* **38**, 283 (1990).

²⁰A. Jennings, J. Li, and J. R. Greer, *Acta Mater.* **59**, 5627 (2011).

²¹P. J. Steinhardt, D. R. Nelson, and M. Ronchetti, *Phys. Rev. B* **28**, 784 (1983).

²²J. Li, *Modell. Simul. Mater. Sci. Eng.* **11**, 173 (2000).

²³H. Zheng, A. Cao, C. R. Weinberger, J. Y. Huang, K. Du, J. Wang, Y. Ma, Y. Xia, and S. Mao, *Nat. Commun.* **1**, 244 (2010).

²⁴J. Y. Kim, D. Jang, and J. R. Greer, *Acta Mater.* **58**, 2355 (2010).

²⁵T. Zhu, J. Li, A. Samanta, A. Leach, and K. Gall, *Phys. Rev. Lett.* **100**, 025502 (2008).

²⁶G. Grochola, S. P. Russo, and I. K. Snook, *J. Chem. Phys.* **123**, 204719 (2005).

²⁷E. Rabkin, H. S. Nam, and D. J. Srolovitz, *Acta Mater.* **55**, 2085 (2007).