Kinetic activation-relaxation technique: An off-lattice self-learning kinetic Monte Carlo algorithm

Fedwa El-Mellouhi,* Normand Mousseau,† and Laurent J. Lewis‡
Département de Physique and Regroupement Québécois sur les Matériaux de Pointe (RQMP), Université de Montréal, Case Postale 6128, Succursale Centre-Ville, Montréal, Québec, Canada H3C 3J7
(Received 3 September 2008; published 7 October 2008)

Many materials science phenomena are dominated by activated diffusion processes and occur on time scales that are well beyond the reach of standard molecular-dynamics simulations. Kinetic Monte Carlo (KMC) schemes make it possible to overcome this limitation and achieve experimental time scales. However, most KMC approaches proceed by discretizing the problem in space in order to identify, from the outset, a fixed set of barriers that are used throughout the simulations, limiting the range of problems that can be addressed. Here, we propose a flexible approach—the kinetic activation-relaxation technique (*k*-ART)—which lifts these constraints. Our method is based on an off-lattice, self-learning, on-the-fly identification and evaluation of activation barriers using ART and a topological description of events. Using this method, we demonstrate that elastic deformations are determinant to the diffusion kinetics of vacancies in Si and are responsible for their trapping.

DOI: 10.1103/PhysRevB.78.153202 PACS number(s): 61.72.Cc, 05.10.Ln, 61.72.jd, 81.10.Aj

Many problems in condensed matter and materials science involve stochastic processes associated with the diffusion of atoms over barriers that are high with respect to temperature and therefore inherently slow under "normal" conditions. Because the associated rates are small, these processes may be considered independent; neglecting the thermal motion of atoms, it is thus possible to deal with them using the kinetic Monte Carlo (KMC) algorithm, a stochastic approach proposed by Bortz *et al.*[1–4] and based on transition state theory, whereby the evolution of a system is determined by a set of prespecified diffusion mechanisms, i.e., whose energy barriers are known beforehand. In KMC simulations, the time scale is determined by the fastest activated processes and, in practice, time scales of ms or longer can be reached—much longer than accessible in traditional molecular-dynamics (MD) simulations.

While KMC has been extensively and successfully used over the past 20 years, it suffers from a number of drawbacks. In particular, the systems investigated must be discretized and mapped onto a fixed lattice in order to define the various diffusion mechanisms that need to be considered at a given moment.[3] Once all processes on the lattice have been identified (and their barriers evaluated) *a priori*, the simulations simply consist in operating a diffusion event picked at random, updating the list of possible moves in the new configuration, and iterating this procedure long enough to cover the relevant physical time scales. This approach works very well for simple problems (e.g., surface diffusion, metal-on-metal growth) but fails when the systems undergo significant lattice deformations or when long-range elastic effects are important. There have been numerous efforts to lift these limitations, most solutions falling into one of two categories: introduction of continuum approximations for the long-range strain deformations, and on-the-fly evaluation of the energy barriers. The first category retains the lattice formulation but adds long-range contributions—which can be computed through various extrapolation schemes—to the barriers.[5,6] With the second class of solutions, there is no need to set up a catalog of all possible activation mechanisms. In a recently proposed self-learning KMC approach, Trushin *et al.*[7] introduced an on-the-fly search for barriers but displacements were restricted to be on-lattice. In other cases, a limited number of activated events using the activation-relaxation technique (ART)-like dimer[8,9] or eigenvector-following methods[10] are generated at each step in order to construct a small catalog which serves to determine the next move. Thus, these methods[5–7] are still limited by the lattice description of the problem and the approximate character of the elastic energies. On-the-fly/off-lattice approaches,[8–10] on the other hand, while more flexible, are currently inefficient as they do not take advantage of the knowledge of previously encountered events, and are therefore only useful for small systems with very few barriers.

In this Brief Report, we introduce a powerful on-the-fly/off-lattice KMC method which achieves speed-ups as large as 4000 over standard MD for complex systems, while retaining a complete description of the relevant physics, including long-range elastic interactions. Our approach is based on the activation-relaxation technique (ART nouveau)[11,12] for generating events and calculating barriers; the gain in efficiency is achieved through a topological classification of atomic environments, which allows configurations and events to be recognized and stored efficiently, and used again as the simulation proceeds, i.e., the method is self-learning. We demonstrate the validity and efficiency of this kinetic ART (*k*-ART) approach by applying it to the problem of vacancy diffusion in crystalline silicon, demonstrating the importance of elastic deformations for diffusion and trapping.

Before describing *k*-ART, it is useful to discuss the topological characterization. For each configuration, a connectivity graph formed by the network of local neighbors is first constructed. These may correspond to covalently bonded atoms in semiconductors, or faces in the Voronoï tessellation of compact materials. It is important that the configuration be uniquely defined through this network; i.e., the connectivity graph must lead to a unique structure once relaxed with a given interatomic potential. In order to classify the activated processes, a truncated graph is constructed around each atom, as illustrated in Fig. 1, the size of which depends on the physics of the system under study. In the case of Si, for example, we define the local environment around an atom by

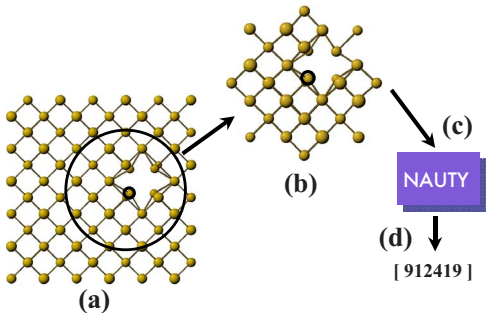

FIG. 1. (Color online) Local topology analysis in *k*-ART: A truncated graph (b) is extracted from the complete lattice around the highlighted central atom (a), and analyzed using NAUTY (c) which returns a unique key associated with the given topology (d).

a sphere of radius 5.0 Å, which includes about 40 atoms; two neighbors are bonded if their distance is less than 2.8 Å. An event is defined as a change in the topology of the local graph. This classification is performed using the freely available topological software NAUTY, developed by McKay,[13] which provides the topology index and all information necessary for uniquely identifying each environment, including the permutation key needed to reconstruct a specific geometry from the generic topology and a set of reference positions.

Events which have been learned are stored for subsequent use; in practice, the atomic positions of the initial state as well as the associated topologies for the initial, transition, and final truncated graphs are saved in memory. If needed, the transition and final state configurations may be reconstructed from the reference geometry through a series of symmetry operations extracted from the topological analysis. This results in a considerable reduction in the amount of data that need to be generated and manipulated. For a single vacancy in *c*-Si, for example, only 20 different topologies are necessary to describe all possible local environments, irrespective of the system size. Moreover, as the system evolves and previously encountered topologies are recognized, it is only necessary to update the table of active events, the cost of which is negligible as we will see below.

We now turn to a detailed description of the *k*-ART algorithm. Starting from an initial relaxed configuration, the various local topologies are characterized with NAUTY and, for each topology, possible events are constructed with ART nouveau,[11,12] which has been shown to efficiently identify the relevant diffusion mechanism in a wide range of systems, either crystalline or amorphous, with both empirical and *ab initio* methods.[14–17] Within this approach, the configuration is slowly pushed along a randomly selected direction until an unstable direction appears in the Hessian; this is followed while minimizing the energy in the perpendicular hyperplane until the system converges onto a saddle point and the system is then pushed over the barrier and relaxed into a new minimum. Since activated processes involve only a finite number of atoms, each event is initiated by displacing a given atom and its neighbors within a small, local region in a random direction. The exact size of the displacement regions depends on the system under study; in semiconductors, they typically involve first and second nearest neighbors. The initial convergence criterion for the saddle point search is set to 1.0 eV/Å in order to accelerate convergence (but see below).

To simplify labeling, each event is assigned to the topology centered on the atom that moved the most during the event, irrespective of the initial trial assignment. The events are stored as displacement vectors from the reference state to the transition and the final states; these are used to reconstruct all specific events associated with a given topology throughout the lattice. Once the list of topologies and associated barriers is set (or has been updated), all low-energy events (which we define for Si as having barriers of 15 k_BT or less) are reconstructed from the topology and re-relaxed with a stricter convergence criterion of 0.1 eV/Å in order to accurately take into account the local environment and the long-range interactions, leading to a precision of about 0.01 eV on the barrier height. At this point, two types of events are in the catalog: (i) "generic" events, which include all high-energy barriers, and (ii) "specific" events, where all low-energy barriers, dominating the kinetics, are relaxed individually. We associate a transition rate $r_i=\tau_0 \exp(\Delta E_i/k_BT)$ to each event, where τ_0 is fixed at the outset and, for simplicity, assumed to be the same ($=10^{13}$ s^{-1}) for all events. From this list, and following Bortz *et al.*,[1] the elapsed time to the next event is computed as $\Delta t=-\ln \mu/\Sigma_i r_i$, with μ as a random number in the [0,1] interval. Finally, an event is selected with a weight proportional to its rate and is operated; the clock is pushed forward and the process starts again. The topology of all atoms belonging to the local environment around the new state is constructed. If a new topology is found, a series of ART nouveau searches are launched; otherwise, we proceed to the next step. After all events are updated, the low-lying barriers are, once again, relaxed before calculating the time increment and selecting the next move.

As the system evolves, it may get trapped in a set of local configurations separated by very low energy barriers that dominate the dynamics without yielding diffusion. An exact solution to dealing with such "flickers" has been proposed by Athenes *et al.*,[18] but we elect here to use a simpler limited-memory Tabu-like approach[19] which proceeds by banning transitions rather than states.[20] In brief, at any given moment, we keep in memory (the "memory kernel") the n previous transitions. If a planned transition is already in memory, it is blocked and the initial or the final configuration of this move is adopted with the appropriate Boltzmann probability; the transition is also blocked for the next n jumps and removed from the list of possible events. As was shown in Ref. 20, this approach is thermodynamically exact and is kinetically valid as long as the memory is short compared to the timeline of evolution of the system.

We now demonstrate the validity and efficiency of our method by studying the diffusion of systems of two and six vacancies in a 1000-atom Stillinger-Weber *c*-Si sample. For the two-vacancy system, we start by removing two second-neighbor Si atoms, then perform a *k*-ART run for 200 CPU hours on a single 1.5 GHz Itanium 2 processor. During this time, the vacancies perform about 1000 jumps, corresponding to a diffusion time of about 100 μs. Figure 2(a) shows the total squared displacement of the atoms as a function of

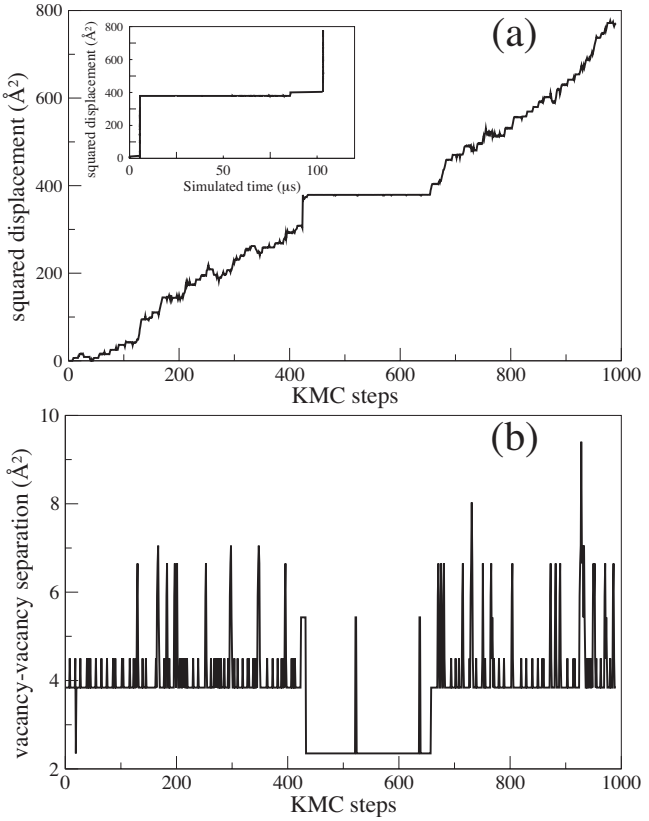

FIG. 2. (a) Total squared displacement as function of KMC steps (or time, in the inset). (b) Distance between the two vacancies as a function of KMC steps.

KMC steps (i.e., events) and, in the inset, effective time. Two types of behavior are clearly visible. During the first 400 steps (~4 μs), diffusion takes place through correlated single vacancy hops over barriers of 0.20–0.24 eV, the two vacancies maintaining a separation oscillating between 3.85 and 4.5 Å. At about step 400, the two vacancies become trapped as a single divacancy, characterized by small local rearrangements, and remain so for about 80 μs before partially breaking apart by going through a 0.7 eV barrier, and resuming its two-vacancy correlated walk. The correlated motion is best seen in Fig. 2(b), where we plot the distance between the two vacancies as a function of KMC steps: The two vacancies remain bound in a first- or second-neighbor state for the whole simulation, except for occasional excursions to larger distances. This striking result illustrates perfectly the impact of elastic deformations on diffusion: While an isolated vacancy diffuses with a barrier of 0.53 eV, the elastic field caused by the presence of a second vacancy in its vicinity lowers the barrier by about 0.3 eV, thus increasing the diffusion rate by almost 3 orders of magnitude. On the other hand, diffusion is strongly hindered when the two vacancies form a stable divacancy because of a dissociation barrier of 0.7 eV.

For the six-vacancy problem, now, we start with a configuration containing two three-vacancy clusters placed far away from each other, as shown in the $t=0$ snapshot in Fig. 3. This configuration is challenging because the dynamics is dominated by a series of local rearrangements and reorientations associated with low-energy barriers that preserve the compactness of the cluster; i.e., breaking it apart is very difficult. To test this, we first ran a 30 ns MD simulation at 500 K; no dissociation or diffusion events were observed. Likewise, nothing happened in a 5000-step k-ART simulation *without* a memory kernel, which covered 8 μs. These two calculations required roughly the same computational effort. We thus already conclude that k-ART is at least 250 times faster than MD; this is fast, but we can do much better by invoking the memory kernel to eliminate the flicker problem which is inherent to such complex materials.

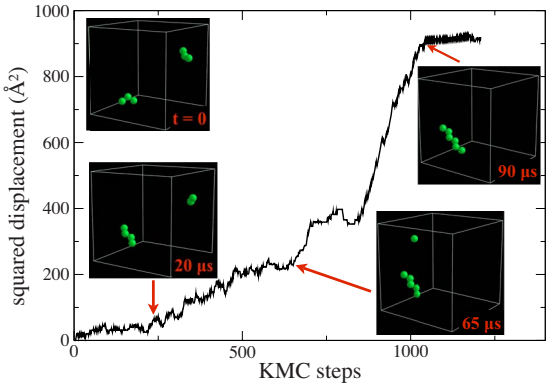

FIG. 3. (Color online) Total squared displacement as function of KMC steps for the six-vacancy system.

Thus, we carried out a third simulation of the six-vacancy problem using k-ART *and* the memory kernel. Figure 3 shows the squared displacement as a function of KMC steps. We observe, in agreement with the previous two simulations, that the initial state is fairly stable: The system flickers during the first 20 μs (160 KMC steps), in agreement with the MD and the k-ART simulation without memory kernel. At 20 μs, one vacancy breaks away from the top right cluster and quickly moves to the other cluster, forming a four-vacancy chain and leaving a divacancy behind. As in the two-vacancy simulation, this divacancy splits and diffuses through the box for about 45 μs (525 KMC steps) as a correlated pair. Finally, at event 685 (65 μs), the divacancy breaks apart and one vacancy rapidly joins the larger cluster. The remaining lone monovacancy diffuses through the lattice during 25 μs and eventually joins the five-vacancy cluster, forming a stable hexavacancy chain with a total energy 2 eV lower than that of the initial configuration. During the following 150 KMC steps (20 μs), the dynamics is dominated by rearrangements and reorientations that do not cause the dissociation of the stable hexavacancy cluster. In fact, the system requires much more time to overcome the 0.7 eV energy barrier needed to dissociate the hexavacancy than it does to dissociate the divacancy since the elastic deformations affecting the lattice are much more important in the former case than in the latter.

In terms of efficiency, k-ART with the memory kernel is about 4000 times faster than MD, with 110 μs simulated in 220 CPU hours. In k-ART most of the computational time is spent in identifying events associated with new topologies. This is clear in Fig. 4 where we plot CPU time versus simulated time for k-ART, with the learning phases indicated by

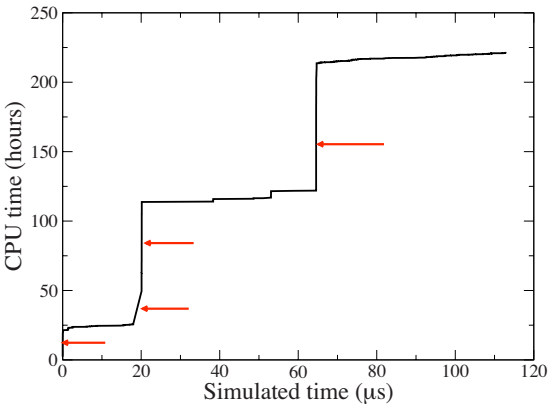

FIG. 4. (Color online) CPU time versus simulated time for the hexavacancy aggregation problems. The red arrows indicate extensive self-learning phases in *k*-ART.

arrows. Sampling is considerable: At the end of this run, 17 237 different events were generated, associated with 1964 initial topologies, for an average of almost 9 events per topology; at each step during the simulation, the system presents about 80–120 different topologies. Since each atom is associated with a topology, about 9000 different barriers are considered at each KMC step. Because *k*-ART is inherently local, a number of improvements can be envisaged that will yield considerable acceleration to the code. Parallelizing the management of events and barriers, for example, should speed up the calculations by a factor of 10 or 20. Moreover, because of the topological classification, the catalog of events may be stored and reused at a later time, thus accelerating new simulations.

Kinetic ART is an exciting self-learning, off-lattice kinetic Monte-Carlo algorithm that opens the door to the numerical study of problems such as semiconductor growth, self-organization, defect diffusion, and interface mixing, which have until now been out of the reach of simulations. While the cost of *k*-ART is significantly higher than that of ordinary lattice KMC, it has already provided essential information regarding the importance of elastic deformations in controlling the kinetics of vacancy diffusion and trapping in Si.

This work was supported by the Canada Research Chairs program and by grants from the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Fonds Québécois de la Recherche sur la Nature et les Technologies (FQRNT). We are grateful to the Réseau Québécois de Calcul de Haute Performance (RQCHP) for generous allocations of computer resources.

*f.el.mellouhi@umontreal.ca
†normand.mousseau@umontreal.ca
‡laurent.lewis@umontreal.ca

[1] A. B. Bortz, M. H. Kalos, and J. L. Lebowitz, J. Comput. Phys. **17**, 10 (1975).
[2] K. Fichthorn and W. Weinberg, J. Chem. Phys. **95**, 1090 (1991).
[3] A. F. Voter, F. Montalenti, and T. C. Germann, Annu. Rev. Mater. Res. **32**, 321 (2002).
[4] A. F. Voter, in *Radiation Effects in Solids*, edited by K. E. Sickafus, E. A. Kotomin, and B. P. Uberuaga (Springer, NATO, Dordrecht, The Netherlands, 2007), Vol. 235, pp. 1–24.
[5] D. Mason, R. Rudd, and A. Sutton, J. Phys.: Condens. Matter **16**, S2679 (2004).
[6] T. Sinno, J. Cryst. Growth **303**, 5 (2007).
[7] O. Trushin, A. Karim, A. Kara, and T. S. Rahman, Phys. Rev. B **72**, 115401 (2005).
[8] G. Henkelman and J. Jónsson, J. Chem. Phys. **111**, 7010 (1999).
[9] F. Hontinfinde, A. Rapallo, and R. Ferrando, Surf. Sci. **600**, 995 (2006).
[10] T. Middleton and D. Wales, J. Chem. Phys. **120**, 8134 (2004).
[11] G. T. Barkema and N. Mousseau, Phys. Rev. Lett. **77**, 4358 (1996).
[12] R. Malek and N. Mousseau, Phys. Rev. E **62**, 7723 (2000).
[13] B. D. McKay, Congr. Numer. **30**, 45 (1981).
[14] Y. Song, R. Malek, and N. Mousseau, Phys. Rev. B **62**, 15680 (2000).
[15] F. Valiquette and N. Mousseau, Phys. Rev. B **68**, 125209 (2003).
[16] F. El-Mellouhi, N. Mousseau, and P. Ordejón, Phys. Rev. B **70**, 205202 (2004).
[17] M.-A. Malouin, F. El-Mellouhi, and N. Mousseau, Phys. Rev. B **76**, 045211 (2007).
[18] M. Athenes, P. Bellon, and G. Martin, Philos. Mag. A **76**, 565 (1997).
[19] F. Glover and M. Laguna, *Tabu Search* (Kluwer, Dordrecht, 1997).
[20] M. V. Chubynsky, H. Vocks, G. T. Barkema, and N. Mousseau, J. Non-Cryst. Solids **352**, 4424 (2006).