

Kinetic activation-relaxation techniqueLaurent Karim Béland,^{1,*} Peter Brommer,^{1,†} Fedwa El-Mellouhi,^{2,‡} Jean-François Joly,^{1,§} and Normand Mousseau^{1,||}¹*Département de Physique and Regroupement Québécois sur les Matériaux de Pointe (RQMP), Université de Montréal, C.P. 6128, Succursale Centre-Ville, Montréal, Québec, Canada H3C 3J7*²*Science Program, Texas A&M at Qatar, Texas A&M Engineering Building, Education City, Doha, Qatar*

(Received 23 June 2011; published 17 October 2011)

We present a detailed description of the kinetic activation-relaxation technique (k-ART), an off-lattice, self-learning kinetic Monte Carlo (KMC) algorithm with on-the-fly event search. Combining a topological classification for local environments and event generation with ART nouveau, an efficient unbiased sampling method for finding transition states, k-ART can be applied to complex materials with atoms in off-lattice positions or with elastic deformations that cannot be handled with standard KMC approaches. In addition to presenting the various elements of the algorithm, we demonstrate the general character of k-ART by applying the algorithm to three challenging systems: self-defect annihilation in c-Si (crystalline silicon), self-interstitial diffusion in Fe, and structural relaxation in a-Si (amorphous silicon).

DOI: [10.1103/PhysRevE.84.046704](https://doi.org/10.1103/PhysRevE.84.046704)

PACS number(s): 02.70.-c, 61.72.Cc, 81.10.Aj, 61.72.J-

I. INTRODUCTION

Solid-phase diffusion in materials science and condensed matter is dominated by rare atomic diffusion events associated with high-energy barriers as measured with respect to temperature. These stochastic processes take place on an extended time scale that makes them very difficult to reproduce using linear simulation schemes such as molecular dynamics. Low rates, however, allow us to consider this series of processes as independent Markov chains. In this case, it is possible to apply the kinetic Monte Carlo (KMC) algorithm proposed by Bortz *et al.* [1–4]. Based on transition state theory, KMC uses a catalog of prespecified diffusion mechanisms to compute at every stage the exit rate from a local minimum. The clock is then advanced using Poisson's law, a move is selected with the appropriate rate, and a dynamical trajectory is constructed. While time steps in KMC are dominated by the lowest energy barriers, it is possible, under the right conditions, to simulate on experimental time scales.

Since it was proposed, KMC has been used extensively in materials science, condensed matter, and many other fields. Its well-known limitations have nevertheless prevented the method from being applied widely to complex systems. In particular, since KMC depends on a predefined event catalog, systems under study have to be discretized and atomic motion limited to fixed lattice positions [3]. In this way, it is possible to evaluate from the onset all possible moves that will be included in the catalog. For relatively simple kinetics, such as metal-on-metal growth, these limits are not major hurdles. They become a problem when trying to take full account of the lattice associated with long-range elastic effects and, more importantly, off-lattice and disordered conformations.

Over the years, a number of algorithms have been proposed to lift, at least partially, these limitations. Most can be

classified in one of two categories: the addition of a continuum approximation for computing the effects of long-range strain deformations on the energy barriers and on-the-fly evaluation of energy barriers. The first class of methods adds long-range contributions, computed from a number of extrapolation schemes, to the energy barriers extracted from the predefined catalog of events [5,6]. The second class relaxes the need for a predefined catalog and, in some cases, moves away from a lattice-based description. This is the case, for example, of the self-learning KMC approach by Trushin *et al.*, which keeps lattice-based displacement, but introduces an on-the-fly search for barriers [7]. To remove the constraints of a lattice-based description, other methods construct a new catalog at each step to determine the next step, using various climbing methods such as the dimer [8,9], eigenvector-following [10], or the autonomous basin climbing methods [11].

The first class of methods [5–7] remains limited to on-lattice positions in addition to providing often ill-controlled corrections to energy barriers modified by elastic deformation. The second class, with an on-the-fly, off-lattice approach, is much more flexible. However, it is inefficient as a catalog must be rebuilt at each step, making it costly to study complex systems with a large number of possible diffusion mechanisms [8–11].

In 2008 we introduced the kinetic activation-relaxation technique (k-ART), an on-the-fly, off-lattice KMC method that lifts these limitations [12]. In this initial work we showed that for a system of vacancies in Si at 500 K it achieves significant speed-ups over standard MD, while retaining a complete description of the relevant physics, including long-range elastic interactions. More recently, Kara *et al.* proposed a similar self-learning kinetic Monte Carlo method based, however, on less controlled procedures for finding barriers and classifying off-lattice configurations [13].

Here we present in detail the k-ART algorithm, which couples the activation-relaxation technique (ART nouveau) [14,15] for generating events and calculating barriers with NAUTY [16] for the topological classification of events. We also present improvements introduced for handling low-energy barriers and large systems. Finally, we demonstrate

*laurent.karim.beland@umontreal.ca

†peter.brommer@umontreal.ca

‡fadwa.el_mellouhi@qatar.tamu.edu

§jean-francois.joly.1@umontreal.ca

||normand.mousseau@umontreal.ca

the efficiency and versatility of the method by applying it to three systems: self-defect annihilation in crystalline Si, interstitial diffusion in Fe, and relaxation in amorphous silicon.

II. OVERVIEW OF KINETIC ART

Following standard KMC, the k-ART method uses an event catalog to compute the rate of escape from a local minimum and bring forward the simulation clock. There are three fundamental differences with respect to standard KMC, however. First, discretization of the local environment is done through topology instead of geometry, allowing atoms to adopt freely any spatial arrangement instead of being constrained to predefined lattice positions. Second, the catalog is not fixed at the simulation onset, but grows as new local environments are visited, allowing the study of very complex systems. Third, event energy barriers are fully relaxed at each step to take into account all geometrical rearrangements due to short- and long-range elastic deformations.

A simulation starts from an initial configuration relaxed into a local energy minimum. The local topology associated with each atom is first characterized with NAUTY [16]. All atoms sharing a specific topology are presumed to be associated with the same list of activated mechanisms. This is the basic assumption of k-ART and, as we will see below, it can be made to hold on a per-atom basis. This approach results in a considerably reduced amount of generated and handled data. For a single vacancy in crystalline silicon (c-Si), for example, one only needs to consider 20 different topologies to describe all local environments, irrespective of the system size.

A search for activated pathways is launched for each topology using ART nouveau [14,15]. ART nouveau was shown to identify efficiently relevant diffusion mechanisms in systems described with both empirical and *ab initio* methods [17–20]. This method has been extensively characterized in Ref. [21]. A number of technical improvements are also reported in Ref. [22] and new events can now be generated with as little as 300 force evaluations with either empirical or *ab initio* potentials. Each event is classified according to the initial minimum, the saddle configuration and the final state topologies and stored in the catalog.

Once the extensive search for events on all topologies is finished, relevant events for the current configurations are collected. All low-energy barrier events are relaxed for specific atoms, to include not only topological but also geometrical effects. At this point, following Bortz *et al.* [1], the elapsed time to the next event is computed as $\Delta t = -\ln \mu / \sum_i r_i$ where μ is a random number in the $[0, 1]$ interval and r_i is the rate associated with event i . The clock is pushed forward, an event is selected with the proper weight, and the atoms are moved accordingly, after a geometrical reconstruction.

Once in this new configuration, the process starts again: the topology of all atoms belonging to the local environment around the new state is constructed; if an unknown 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 applying the KMC algorithm.

III. ALGORITHMIC AND IMPLEMENTATION DETAILS

A. ART nouveau

The search for activated mechanisms is performed using ART nouveau [14,15], an open-ended climbing method for finding first-order saddle points surrounding a local minimum. As the most recent version of the algorithm is described in Refs. [21] and [22], we give here only a brief overview of the method. Event search with ART nouveau proceeds in three steps: (1) starting from an energy minimum, the system is deformed locally in a random direction until the lowest curvature of the Hessian matrix becomes negative, indicating an instability; (2) the configuration is then pushed along this direction of negative curvature while the energy is minimized in the hyperplane orthogonal to this direction until the total force falls below a set threshold, indicating that a first-order saddle has been reached; (3) the configuration is pushed over this point and is relaxed into a new minimum. This set of three configurations—initial minimum, saddle point, and final minimum—forms an event.

Since activated processes are local in nature, each event is initiated by displacing a given atom and its neighbors in a random direction. The exact size for this displaced region depends on the system studied. In semiconductors, it involves typically first and second nearest neighbors. In the case of Si vacancies, for example, the displacements are applied to the central atom and all atoms within a 3.0 Å radius from it. The initial convergence criterion for the saddle point is typically set to 0.5 eV/Å. This is associated with the generation of the generic event catalog (see below). Further relaxation associated with the calculation of specific events uses a 0.1 eV/Å threshold.

To decrease computational cost, ART nouveau never computes the Hessian directly, but rather uses a mixture of Lanczós [23] and DIIS [24] methods for converging to the saddle point. As discussed in Ref. [22], less than 300 force evaluations are generally needed to converge to a first-order saddle point. Taking into account all processes, including nonconverging steps and relaxation into a new minimum, about 600–800 force evaluations are required, on average, per successful event search. Relaxation of specific events, which starts near a reconstructed saddle point, are normally much faster, necessitating typically 1 to 80 force evaluations.

B. Topological classification and event generation

Topological classification of the local atomic environment is a crucial step in k-ART, as it provides a means of discretizing and cataloging local configurations, while taking into account all possible atomic arrangements and elastic deformations.

Atomic topologies for a given local configuration are computed as follows. We define a local environment consisting of all atoms within a sphere of a predefined radius centered around each atom, as illustrated in Fig. 1. These are then connected following a neighboring prescription, such as first neighbor distance cutoff or a Voronoï tessellation, forming a truncated connectivity graph, that is, a set of bonds connecting vertices, without geometrical information. This graph is then analyzed and classified using the freely available topological software NAUTY, developed by McKay [16]. This software

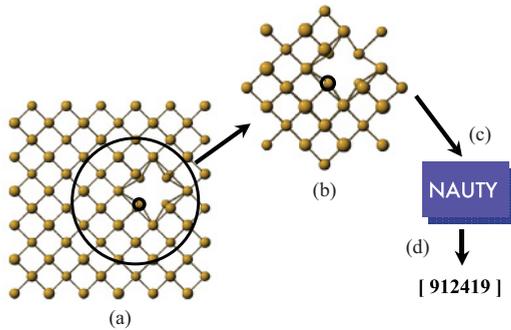


FIG. 1. (Color online) Local topology analysis procedure in k-ART. A truncated graph (b) is extracted from the complete lattice (a). This graph is analyzed through NAUTY (c), which returns a unique key and the associated topology (d).

package 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. The topology index from NAUTY has the form of an ordered set of three integers, which we use as input to a hash function to generate a unique topology label (hash key).

The geometrical reconstruction from a purely topological graph is made possible because we know the atomic positions of all atoms surrounding the local configuration described by this graph. This introduces sufficient constraints to ensure that most of time, as discussed below, a given graph corresponds to a unique fully relaxed geometry.

The neighboring prescription and the size of the truncated region are selected to ensure that, in most cases, the configuration is uniquely defined through this network, that is, the connectivity graph must lead to a unique structure once relaxed with a given interatomic potential. In the case of crystalline Si, for example, we define the local environment around an atom by a sphere of radius 5.0 Å, which includes about 40 atoms; two atoms are linked if their distance is less than 2.8 Å.

Once all new topologies are identified, a succession of ART nouveau searches are launched on each of them. The optimal number of searches done per topology depends on the nature of the system and can be adjusted to ensure that all low-energy barrier events, which dominate the dynamics, are found. In our current implementation, it is increased also with the number of times a given topology is seen, ensuring that the most common topologies are explored more often.

Every new event found with ART nouveau is analyzed and compared to the list of already known events for that initial topology. If an event with the same activation energy is already in memory, a series of tests are performed to assess whether or not it is the same event. When the generated event is judged to be new, it is then assigned to the topology centered on the atom that moved the most during the event, irrespective of the initial activation. The event label is obtained by combining the topology labels at the initial, saddle, and final states. We also keep in memory the position of the cluster of atoms for the initial, saddle, and final configurations. These are necessary (a) as a reference for the geometrical reconstruction from the saddle or the final configuration topological mapping for atoms belonging to the same topological class and (b) to compare the

event in the list with newly generated events. Once a generic event is added to the database, it is also added to the binary tree of events and to the histogram (see Sec. III C).

Finally, based on these data, a first generic rate is associated with the event by setting

$$r_i = \tau_0 \exp(-\Delta E_i/k_B T), \quad (1)$$

where τ_0 , the attempt frequency, is fixed at the onset and, for simplification, assumed to be the same for all events (10^{13} s^{-1}). ΔE_i is the barrier height, that is, the energy difference between saddle point and initial minimum. k_B and T are, as usual, the Boltzmann constant and the temperature, respectively.

While using a fixed attempt frequency is a simplification, it has been shown in several systems that τ_0 varies only weakly with the chosen pathway [25]. The value of 10^{13} s^{-1} is compatible with pre-exponential factors in iron derived from experiment [26] and in simulation [27]. For silicon, this value is also compatible with *ab initio* computations of neutral vacancy diffusion [19].

In certain symmetric configurations, it is possible that distinct events associated with a given atom have identical topologies of initial state, saddle point, and final state, and the same barrier height and absolute displacement. To distinguish between those and fully account for the symmetries, we also include the direction of atomic motion in the description of events. Checks on a number of highly symmetric states in Si and Fe show that this classification recovers all pathways and accurately discriminates between them.

Hash keys provide a fast way of storing and retrieving event or topological information. If the key of a new event or topology is already in use (a so-called hash collision), the key value is incremented by one until a free key is found. The arrays for events and topologies account for most of the memory used in k-ART simulations, and so a balance has to be struck between size and speed. Too large arrays waste RAM, while too small ones provoke frequent hash collisions and fill up earlier.

One of the major advantages of this approach is that this constantly updated catalog of generic events and topologies can be saved to the disk, made available to others, and reused on the same and similar systems. Multiple catalogs from different k-ART simulations can also be merged to create a larger database that can be used to start new simulations on the same system with a significant speed increase.

C. Adaptive generic event relaxation into specific events

Events generated for a given *topology* are known as *generic events*. It is assumed that all atoms sharing the same topology will have access to these events with, however, a small adjustment to the energy barrier due to local variations in position or long ranged elastic interactions. To take these changes into account, generic events with low barriers are reconverged for each realization, resulting in *specific events*, each linked to a particular atom.

Starting from a common topological generic event, a specific event is generated for each atom having the same topology by taking advantage of the ordered list of cluster atoms around the central atom obtained using NAUTY (permutation key). With this information, it is possible to reconstruct the geometry of the specific saddle and final

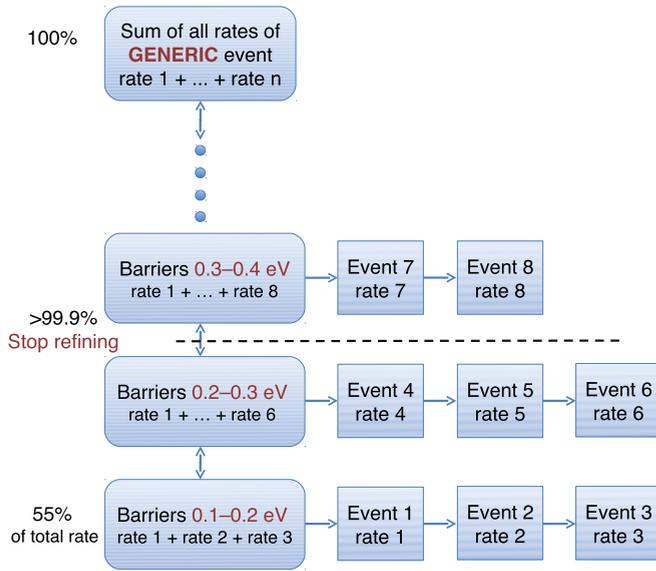


FIG. 2. (Color online) Generic events stored in a histogram used for an adaptive event relaxation procedure.

minimum conformations by mapping the displacement vector and transforming geometrically a given region in the system with respect to the generic configuration.

In an earlier version of the algorithm [12], specific events were identified as those belonging to generic classes with an energy barrier of $15 k_B T$ or less. Here, we adopt rather an adaptive algorithm based on the kinetics of the system. All generic events are ordered and stored into slices of 0.1 eV according to their energy barrier in a histogram (see Fig. 2). For each slice, the total rate is computed and the cumulative rate up to each energy barrier range is calculated.

The relaxation into specific events starts from the bottom of the histogram and then proceeds to higher energy barriers until events accounting for a large fraction of the total rate are relaxed (we use 99.9% here). The remaining generic events are copied to unrefined specific events. In this way, we ensure that almost all selected moves will be picked from the list of fully relaxed individual barriers. Geometric and elastic effects on the energies of local minima and transition states are therefore fully included in the rates.

D. Crucial aspects

1. Uniqueness in the correspondence between a geometry and topology

One of the main k-ART advantages resides in the topology-based discretization. To be valid, this classification requires a unique correspondence between a local geometry and a topology. With the right building rules for the truncated graph, it turns out that this approximation almost always works for covalently bonded systems, but also for metals. Since k-ART relies on the delicate reconstruction of transition states for computing specific events, a failure of the relation between topology and geometry leads systematically to the disappearance of this first-order saddle point and can be easily detected. In other words, event reconstruction automatically

fails when a given topology corresponds to more than one geometry.

When an ill-defined topology corresponding to more than one geometry is encountered for the first time, k-ART automatically adjusts the atom-atom connectivity cutoff used within the cluster of atoms forming this graph, until the geometries initially associated with the same topology are now placed into two different classes. The unique correspondence between local geometry and topology is now re-established.

For computational ease, a marker identifies already encountered difficult topologies so that k-ART can recognize them on the fly and test them to ensure that the correct event catalog is used. To do so, k-ART maps one event from each of the associated subtologies, until a successful map is achieved. If no event can be mapped using the current subtologies, a new label is again issued. Each difficult topology can have as many subtologies associated with it as necessary. In practice, topology collisions are extremely rare, if reasonable cluster and neighbor cutoffs are used. If there is a collision, it can usually be resolved with only one subtology.

2. Handling low-energy barriers

In KMC, dynamics is dominated by a system's lowest barrier energy. When the energy landscape consists of basins with numerous states connected by very low-energy barriers compared to those needed to leave these basins, the algorithm becomes trapped into computing nondiffusive events, decreasing significantly its efficiency in two ways. First, it limits the attainable simulated time, as the low-energy internal barriers produce a high total rate sum and thus a short average KMC time increment. Second, computational resources are bound to explore the states within a basin without yielding much information, as effective diffusion takes place typically outside the energy basins.

We had previously implemented a TABU-like approach [28], that bans transitions rather than states [29]. This algorithm is simple to implement and provides a thermodynamical solution when there is a clear energy separation; it fails, however, when few pathways are available or the energy spectrum is continuous. To account for these situations, we developed the basin-auto-constructing Mean Rate Method (bac-MRM), a basin-based acceleration scheme inspired by the mean rate method (MRM) of Puchala *et al.* [30]. A description of MRM can be found in the Appendix. In summary, MRM separates the trajectory into transient states and absorbing states, and accelerates the simulation by averaging over all possible jumps between transient states, yielding the correct probability to exit a basin to a certain absorbing state.

In kinetic ART, the relevant entities are not states, but *events*, characterized by an energy barrier between the initial and final states. In an event with energies E_i , E_s , and E_f at the initial state, the saddle point, and the final state, respectively, we define the forward energy barrier as $b_f = E_s - E_i$, and the inverse barrier by $b_i = E_s - E_f$. In both cases, rates going forward or backward are determined by Eq. (1). Basins are then identified on the fly by the barrier heights separating the basin states: Both the forward and the inverse barrier must be smaller than a user-defined threshold.

Starting in a local minimum, low-energy barrier events are marked as potential basin events, and the atomic displacement associated with these events is stored. If such an event is picked for execution, it is added to the current network of basin events (i.e., for the first basin event, this is at that time the only event), removed from the tree of available events, and executed as normal. The system is then in state two, and the k-ART event finding algorithm is started from this state. All events from previous basin states are kept in the tree and could be picked as KMC move.

After state two has been searched for possible events, and before the next event is picked, the MRM is applied to the basin consisting of two states: The rates leaving the basin are modified following Eq. (A4), and the total rate is adjusted accordingly. The next step is then selected. It can either lead to a new basin state, to be added according to the procedure described above, or out of the basin, in which case a standard k-ART move is applied. If an event is found to lead to an already explored basin state, it is rejected, removed from the tree, and added to the basin, adjusting the rates as needed.

Bac-MRM explores basins on the fly, and only as far as necessary. Simultaneously, no state is intentionally visited twice. While the internal dynamics within a basin is lost, the basin mean rate method in our implementation yields the correct distribution of exit states depending on the basin internal rates and the point of entry into the basin.

The computational overhead of bac-MRM is small and the CPU resources needed for all basin related operations are negligible compared to the time required to explore a single topology.

3. Optimizing k-ART for large scale systems

Because activated mechanisms are local in nature, it is relatively straightforward to optimize k-ART to handle systems with several tens of thousands of atoms. Indeed, diffusive motion typically involves regions composed of a few tens to a few hundreds of atoms, and the forces induced by this displacement typically propagate up to a few nanometers. Therefore, by coupling standard cell lists [31] and the Verlet algorithm [32] for constructing neighbor lists with a local force calculation, the computational effort of generating an event becomes almost system-size independent.

Local forces are first computed on all atoms involved in the event plus their first and second neighbors. As the system evolves, atoms on which the force exceeds a set threshold (of 0.01 eV/Å, in the case of c-Si) become labeled as *involved* and their first and second neighbors are added to the list. This process ensures that forces are computed only on the relevant atoms. The generation of new events and the relaxation of specific events therefore is entirely local, with only a global minimization performed after each KMC move to take into account all elastic effects.

We simulated several systems composed of 2744 to 27 000 c-Si atoms using an empirical Stillinger-Weber potential, periodic boundary conditions, and roughly the same density of vacancies and interstitials in equal numbers. The scaling is given as a function of system size in Fig. 3 for exploring a new topology on a single processor. Because of rare global calculations during the analysis of events, the algorithm has

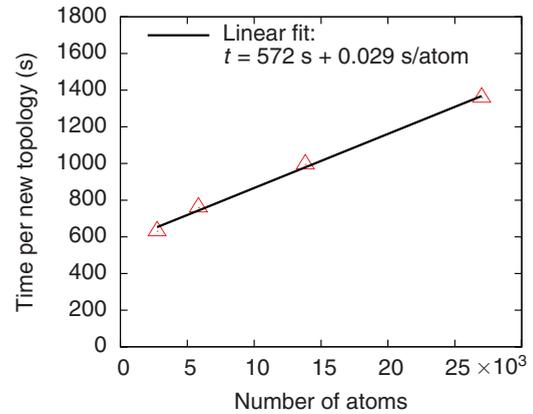


FIG. 3. (Color online) The average computational time required to explore a new topology as a function of system size. We show results with 2744 atoms, 5832 atoms, 13 824, and 27 000 atoms with 1, 2, 4, and 8 vacancy and interstitial defects, respectively.

a weak linear dependance of 0.03 s per atom. However, the sublinear terms, responsible for about 600 s, are dominant for system size of several tens of thousands of atoms. These sublinear terms are associated to the time required to attempt 15 saddle point searches on one topology and to analyze the results as computed on a single core of a 2.66 GHz Intel Xeon X5550 CPU (all CPU times reported in this paper are computed on the same CPU).

IV. APPLICATION OF k-ART

We now apply k-ART to three different systems, in order to demonstrate the flexibility of the method: (a) Vacancies and interstitials in c-Si. This work expands on the vacancy diffusion study presented in Ref. [12]. (b) Interstitials in Fe. Diffusion mechanisms for self-interstitial in iron are surprisingly complex; this system represents a good test of k-ART's ability to sample such landscape. (c) Relaxation of amorphous silicon (a-Si). By construction, KMC methods have been mostly limited to lattice-based problems; here we show that the topological approach of k-ART is sufficiently flexible to handle disordered materials.

A. Vacancies and interstitials in Si

We studied the annealing of eight pairs of vacancies and interstitials in an 8000 atoms c-Si box at 500 K using the Stillinger-Weber potential [33] and periodic boundary conditions. A snapshot of the initial state of the simulation is shown in Fig. 4. With this interatomic potential, the diffusion activation barrier for an individual vacancy is about 0.43 eV [34]. Interstitials can be stabilized in three states, two of which are almost degenerate in energy at about 0.75 eV above the first one. Diffusion for the single self-interstitial is dominated by a barrier at 0.94 eV [35].

Figure 5 (bottom) shows the evolution of the total energy, measured with respect to the crystalline state, and the squared displacement as a function of simulated time for the system above, representing a total of 2000 k-ART steps. Vacancies dominate the diffusion with significantly lower energy barriers. Each of the five large drops in energy corresponds to the

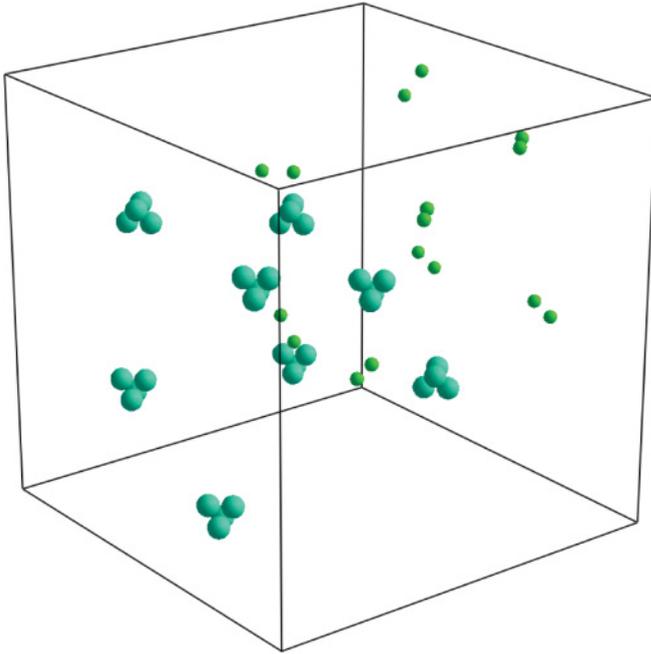


FIG. 4. (Color online) The initial state of our 8000 atoms c-Si box containing eight vacancies and eight interstitials. We only show over-coordinated (small spheres) and under-coordinated atoms (large spheres).

annihilation of an interstitial-vacancy (IV) pair, while from roughly $1 \mu\text{s}$ on, bound defects reorganize themselves without any recombination. We show in Fig. 6 an example of a typical isolated recombination. We observe several metastable states for the bound pair, before recombination, in general agreement with the findings of Tang *et al.* [36] and Marqués *et al.* [37].

Energy fluctuations in Fig. 5 are also associated with the formation of self-defect aggregates, such as a bi-interstitial and a vacancy complex, as well as elastic deformations. These can cause important differences between IV-pair recombinations. This is the case of the second IV recombination (at $t = 0.08 \mu\text{s}$), which differs markedly from the four others. Indeed, the presence of nearby vacancies (at a distance of about 1 nm from the main interstitial and vacancy) significantly modifies the relaxation mechanism. They introduce an intermediate oscillatory state which flickers for a duration of nearly 45 ns (see Fig. 7) before IV recombination. This metastable state underlines the importance of correctly handling long-range elastic and topological effects for the defect kinetics in semiconductors.

Even small changes in the environment, up to 20 \AA away from the defect, can change the kinetics of defect diffusion and recombination. For example, IV-pair recombination, for a nearly, but not completely isolated pair can follow at least two pathways, with barriers differing significantly: 0.45, 0.39, 0.19, and 0.20 eV in the first case, which is in good agreement with Gilmer *et al.* [38], and 0.36, 0.19, 0.22, and 0.51 eV in the second.

The advantage of k-ART's approach to catalog building can be seen in Fig. 5 (top). Since only previously seen topologies are included, the catalog grows as new regions of the configurational space are visited, avoiding the need to

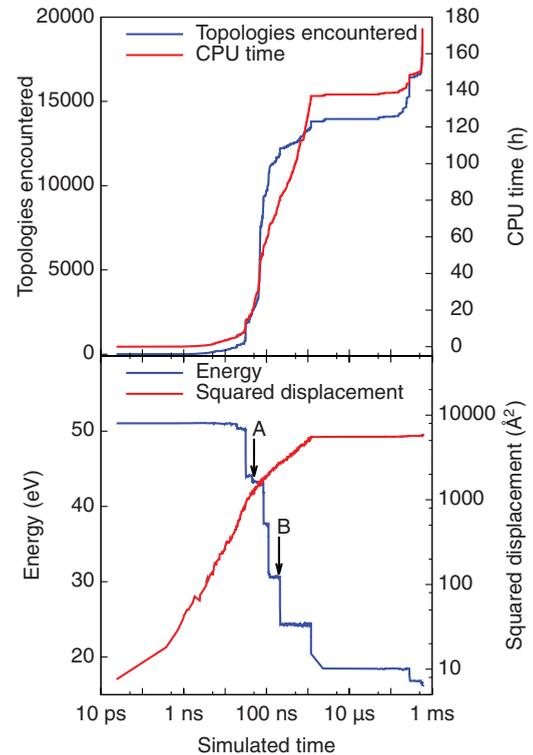


FIG. 5. (Color online) Simulation of eight IV pairs in an 800-atom box at 500 K. Top: The number of encountered topologies and the computation time as a function of the simulation time. Bottom: The evolution of the total energy, measured from the perfect crystal, and of the squared total displacement as a function of the simulation time. The zero on the energy scale corresponds to a box with no defect. Arrows indicate important interstitial-vacancy annihilation states shown as snapshots in Figs. 6 and 7.

construct all possible conformations from the onset, a task that would rapidly be impossible, even for a system with 16 defects. We see that as the defects first diffuse in isolation for the first 20 ns, very few new topologies are encountered. It is only as the recombination processes take place, between $t = 20 \text{ ns}$ and $1 \mu\text{s}$, that many new environments are visited, increasing rapidly the number of sampled conformations from a few hundreds to almost 20 000 topologies. While this number is large, all events are stored in a catalog and serve as such in any new simulation, reducing considerably the CPU cost over time.

B. Self-interstitial cluster diffusion in iron

Iron is widely used in nuclear power plants, which makes the modeling of irradiation-induced defects interesting, in particular single and clustered self-interstitial atoms (SIA) [39], from the microscopic to large scale, an important topic in computational materials science [40,41].

Simulating SIAs on an atomistic level is challenging for both molecular dynamics (MD) and standard KMC methods. On the one hand, the activation energies of defect migrations are rather high, so that MD simulations must be performed at comparatively high temperatures (up to 1200 K, cf. Refs. [42,43]) to explore the energy landscape within the accessible simulation times. This makes it difficult to identify

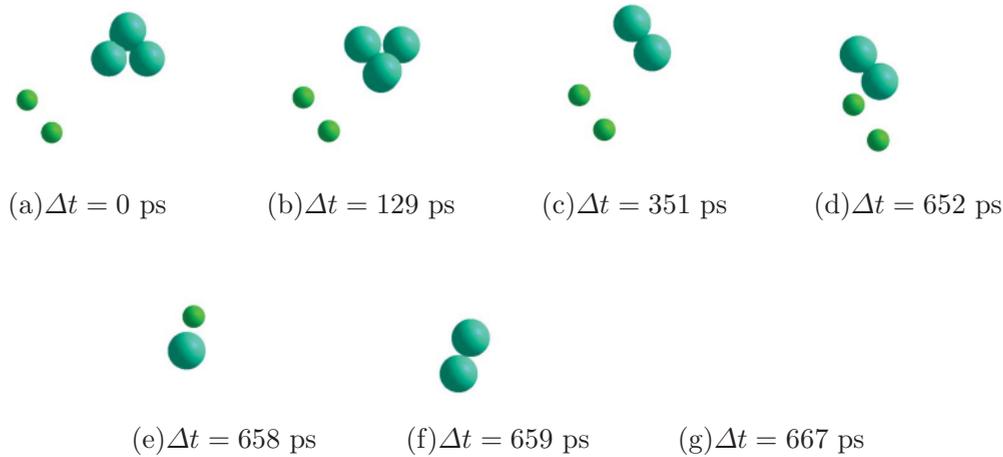


FIG. 6. (Color online) A typical interstitial-vacancy recombination. Time is measured from the initial snapshot, taken at $0.2 \mu\text{s}$ and indicated by arrow B in Fig. 5. Over-coordinated atoms are shown as small spheres and under-coordinated atoms are shown as large spheres. The empty figure in (g) indicates that there are no more topological defects in the local environment.

structures and mechanisms important at the significantly lower operating temperatures. On the other hand, due to the wealth of arrangements of interstitial atoms and other defects, it is a formidable task to build a catalog of possible transitions *a priori* (to use in a KMC simulation), without accidentally neglecting important migration paths. Often, a very reduced catalog of transition pathways is used [44]. All this is complicated by the fact that in iron clusters of interstitial atoms can glide in one dimension with very small migration energies (tens of milli-electron volts) [42].

K-ART is an ideal tool to explore the migration pathways of SIAs without any assumptions needed to assemble a catalog of events and the associated barriers *a priori*. Events are searched for on the fly for each topology in the system, with corrections applied at each KMC step to account for elastic distortions by surrounding defects. As outlined in Sec. III D 2, low barriers can also be treated efficiently with the basin-autoconstructing mean rate method.

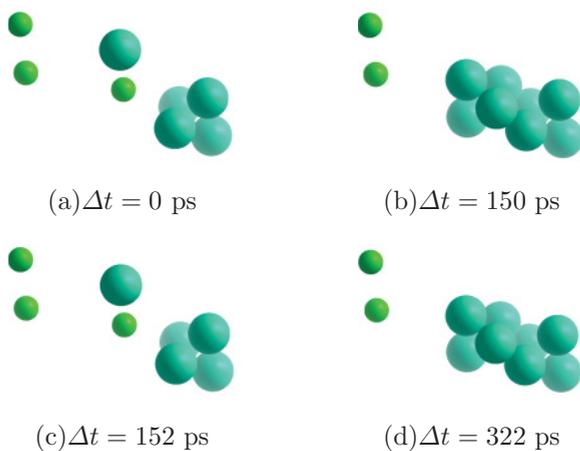


FIG. 7. (Color online) Oscillation of an interstitial-divacancy complex during an interstitial-vacancy recombination in the presence of nearby vacancies. Time are measured from the initial snapshot, taken at 50 ns and indicated by arrow A in Fig. 5. Over-coordinated atoms are shown as small spheres and under-coordinated atoms are shown as large spheres.

For our iron simulations, we used the Ackland-Mendelev potentials [45], an improved version of the potential developed by Mendelev *et al.* [46]. This potential describes defects accurately [47] and was used in MD [42,43] and ART nouveau [21] simulations of iron SIA systems. The EAM energy and force calculation routine was adapted from IMD [48] and can use any tabulated potential in IMD form.

To test the implementation, a single self-interstitial atom was embedded in a 1024-atom supercell. In the ground state, the interstitial forms a dumbbell in $[110]$ direction in agreement with earlier static simulations [49] and ART studies [21]. The transition with the lowest energy barrier follows the nearest neighbor (NN) translation-rotation mechanism proposed by Johnson [50] with an activation barrier of 0.3 eV. Higher barrier events include a transformation to the $[11\bar{1}\xi]$ dumbbell, followed by an on-site rotation, then pure translations to first and second NN sites. Other higher-energy events were found, but were not picked to be executed during our simulations.

In simulations with two clustered interstitials, k-ART recovers the mechanism for interstitial migration suggested by Johnson [50]: Both interstitials each perform a nearest-neighbor translation-rotation jump. This can happen in a single step, or in two sequential moves. The states and barriers found in our simulation agree with the results from Marinica *et al.* [21]. A sample trajectory of the di-interstitial system over 2000 KMC steps is shown in Fig. 8.

For a single self-interstitial, at the temperature of interest, barriers are well above $k_B T$, and there are no flickers. For SIA clusters, however, there are sequences of states separated by low barriers. In the di-interstitial case, Marinica *et al.* [21] find a basin of 0.25 eV above the ground state, with barriers below 0.1 eV separating the states. We reproduce that basin as demonstrated in a detail of the trajectory in Fig. 9. While the system explores the basin, low-energy barriers keep the system clock almost at a standstill.

In a system with a 4-SIA cluster, a basin is found around the ground state: Small reorientations of the four dumbbells lead to about 20 unique configurations separated by extremely low barriers ($<0.1 \text{ eV}$). Since the dynamics is dominated by

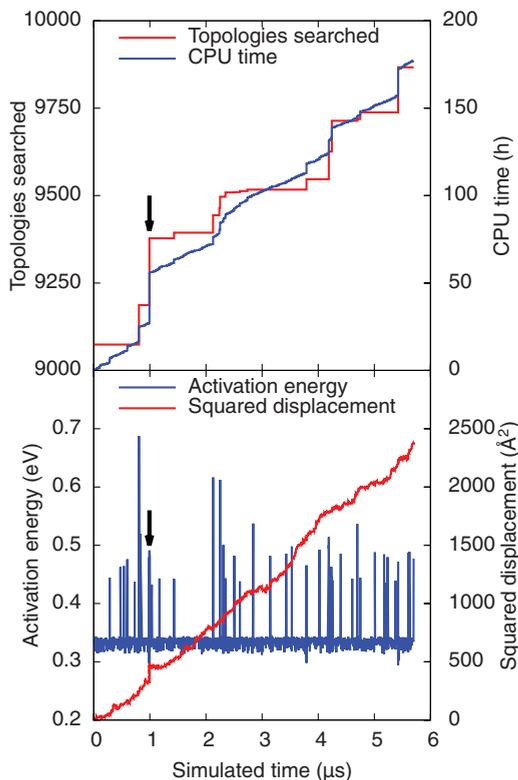


FIG. 8. (Color online) Simulation of an iron di-interstitial cluster in a 1024-atom box at 300 K. Top: The number of encountered topologies and the computation time as a function of the simulation time. Bottom: The evolution of the total energy, measured from the ground state, and of the the squared total displacement as a function of the simulation time. The system propagates mainly by the two-step Johnson process with an activation energy around 0.33 eV. The simulation was performed using a prebuilt catalog constructed from an earlier simulation, initially containing information about 9074 topologies. The simulation stalls when many new topologies must be explored. The arrow marks the basin shown in detail in Fig. 9. In that basin, the activation energy is under 0.3 eV. Two more basins are encountered at 3.8 and 5.5 μ s.

low-barrier events, the system manages to exit the basin only when all those states have been explored. A sample trajectory over 290 KMC steps is displayed in Fig. 10.

The wealth of structures and transition pathways found in iron systems with SIA clusters is virtually impossible to include in a catalog assembled *a priori* for standard KMC simulations. In contrast, the self-learning k-ART program will over the course of a simulation build a database of these configurations and events, thus saving time once the system revisits previous states. The presence of basins (i.e., groups of states separated by low barriers) in these systems dictates the use of an acceleration scheme. With bac-MRM, even such a rich system becomes a tractable problem in KMC simulation.

C. Relaxation dynamics in amorphous silicon

As kinetic ART is an off-lattice method with self-learning catalog building capabilities, it can also be used to study relaxation of disordered materials on long time scales. There have been previous applications of similar techniques to these

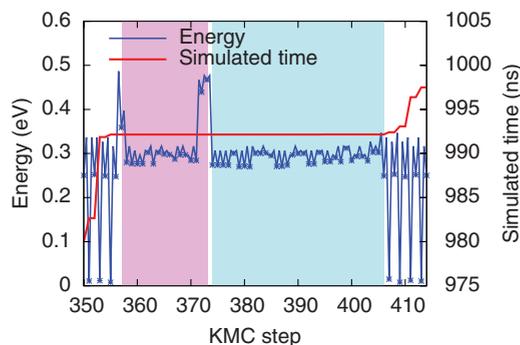


FIG. 9. (Color online) Detail of Fig. 8 (KMC steps 350–413, 0.99 μ s): After a series of two-step Johnson jumps, the system crosses into an excited state at KMC step 235. Then a number of states forming a basin are traversed. Shaded background indicates a basin (oscillatory) motion. After an exit event, the system resumes its basin trajectory, until another exit event leads it back to the ground state, from where it resumes its two-step motions. The energy trajectory passes through minima and saddle points alternately (minima marked by crosses). During the basin motion, the system clock is hardly moving.

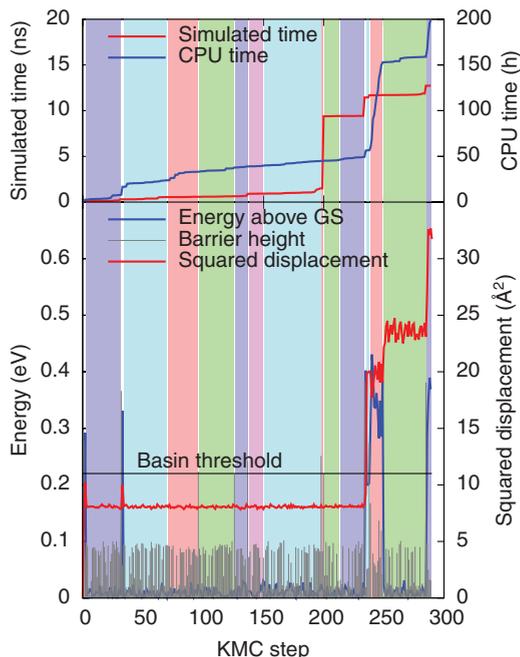


FIG. 10. (Color online) Trajectory of a tetra-interstitial cluster at 300 K in a 1024-atom cubic box over 290 KMC steps. Top: The simulation and the computation time as a function of KMC steps. Bottom: The evolution of the total energy, measured from the perfect crystal, and of the the squared total displacement as a function of KMC steps. The system makes several attempts to leave the ground state basin, but falls back until, at KMC step 234, it succeeds. It then moves through a sequence of excited states, before dropping back to a different ground state basin, with the whole cluster diffusing to the nearest neighbor site. Different background colors represent different basins (white: outside of basin). As the barriers (shown as impulses in the lower plot) are comparatively low, the clock advancement is rather small. Only if a barrier exceeding the basin threshold is picked, the KMC time step is noticeable. A significant share of the CPU time is spent exploring the sequence of excited states between steps 234 and 250.

systems, but these have suffered from limited sampling of events [10,51].

As a test case, we looked at amorphous silicon. Like crystalline silicon, this allotrope of silicon is fourfold coordinated with randomly oriented tetrahedra causing the loss of medium and long-range order in the system. This model system has been extensively studied with ART [14,52] and ART nouveau [18,53] and constitutes therefore a well-controlled model. Moreover, many fundamental questions remain regarding its dynamical properties. For example, in spite of considerable experimental efforts, the exact nature of defects responsible for structural relaxation is still a matter of debate [54–57]. As for many other disordered systems, only methods able to reach experimental time scales will be able to offer a satisfactory answer to these questions.

We start here with a well-relaxed 1000-atom a-Si configuration with periodic boundary conditions generated with the modified Wooten-Winer-Weaire procedure [58,59] and a reparametrized version of the popular Stillinger-Weber potential by Vink *et al.* [60] adjusted to describe appropriately this allotrope. All atoms in the generated a-Si sample are perfectly coordinated with a clean electronic gap [61] and a good agreement with the experimental radial distribution function [59].

For a disordered system, the advantage of recycling events based on the local atomic topologies takes a lot of time before becoming noticeable. For a well-relaxed 1000-atom model of a-Si, for example, no two atoms share the same topology and even after many thousands of events, topologies encountered more than once are rare. A meaningful catalog requires therefore the combination of many independent KMC trajectories started from various initial configurations.

At first, since each atom has its own topology, the number of initial events to be generated is very large. Successive steps tend to be much less expensive and the number of new topologies per step depends strongly on the amplitude of the displacement during the previous KMC step. Small displacements observed during flickers usually result in less than 10 new topologies while diffusive events can generate up to 140 new topologies in a single step.

The distribution of activation barriers is similar to the one found in previous ART studies [18,53]. Although it is continuous over a wide range of activation energies, the kinetics is dominated by nondiffusive low energy barriers. Contrary to crystalline systems, where a clear energy threshold separates diffusive from nondiffusive events, it is necessary to fix the basin threshold somewhat arbitrarily in the mean-rate method. Here we chose a cutoff of 0.3 eV for a simulation temperature of 300 K. Therefore, events associated with time scale of 16 ns or less are averaged over and the internal dynamics of these events is ignored. This is acceptable, as we are interested in simulations on the time scale of microseconds or more.

Figure 11 (bottom) shows the evolution of the total configurational energy as a function simulated time for a simulation of 3360 k-ART steps. Since we started from a very well-relaxed configuration, the proportion of flickers is important but the system still manages to relax by more than 6 eV over a 12 μ s simulation. While the initial relaxation with the modified Stillinger-Weber potential leaves the system

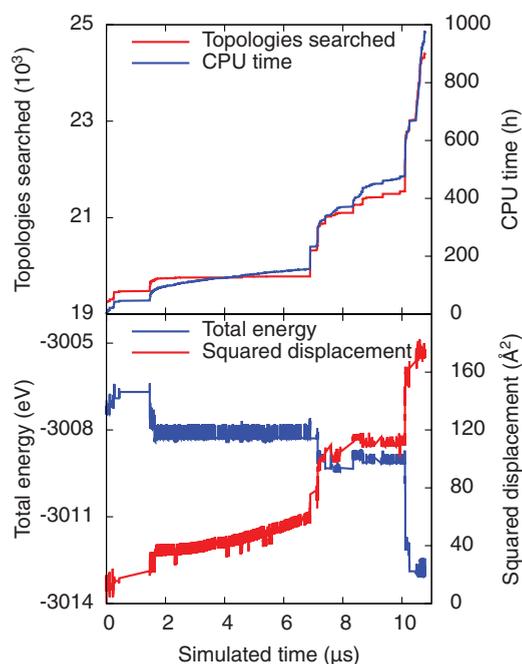


FIG. 11. (Color online) Simulation of a-Si in a 1000-atom box at 300 K. Top: The number of encountered topologies and the computation time as a function of the simulation time. Bottom: The evolution of the total energy and of the the squared total displacement as a function of the simulation time. The simulation was started with a catalog from an earlier simulation. The system flickers between two neighboring states until it finds a way to relax further. This leads to a sequence of configurations never seen before and the CPU time needed per step increases with the number of new topologies to explore.

perfectly coordinated, an average of 0.8 at.% defects are created in a few KMC steps. This concentration is relatively constant throughout the simulation. Moreover, almost all the low barrier events involve coordination defects. Defect migration events are hard to characterize since local atomic motions can affect the existence of low energy defects up to the third nearest neighbor distance. This can cause some defects to disappear while creating new ones.

The 4 eV drop at 10 μ s is initiated by a bond switching event of two fourfold coordinated neighboring atoms with a barrier of 0.28 eV. This allows for one atom to get rid of a highly strained bond, resulting in a energy drop of 0.84 eV. This event is then followed by a succession of 84 smaller relaxation events involving mostly spontaneous creation or destruction of low-energy coordination defects. The configuration eventually ends up in a lower energy basin where flickers again dominate. The average defect population goes from 0.8 to 1.4 at.% during the entire process.

K-ART in a-Si can be compared, on short times, with molecular dynamics. Using the MD software LAMMPS [62] with our a-Si model, we launch a 10-ns simulation at 300 K with a time step of 2 fs. At regular intervals, a configuration is frozen and relaxed into the nearest local minimum using steepest descent in order to compare with our k-ART simulation. Results (not shown) confirm that almost no deformation takes place on this short time scale and both MD and k-ART display atomic displacement of the same amplitude.

Total simulation time for this system is significant and Fig. 11 (top) shows the evolution of simulation time as a function of computer time for a code running on a single 2.66 GHz Intel Xeon X5550 CPU starting from a preconstructed catalog. The use of a parallel version of the algorithm coupled with a more extensive catalog is expected to reduce considerably the computational efforts for this simulation. Already, however, we see that k-ART can be a useful tool for these complex systems.

V. CONCLUSION

In this paper we present in detail the kinetic ART algorithm, a versatile self-learning on-the-fly, off-lattice kinetic Monte Carlo method. This method couples ART nouveau [15], a very efficient nonbiased open-ended algorithm for finding transition states [21,22], with a topological classification of events based on NAUTY, a powerful packaged developed by McKay [16].

Kinetic ART constructs a reusable event catalog that improves the efficiency of the algorithm over time. Events are stored as generic events coupled to a given topology. To fully include elastic deformations, the lowest-energy barriers are separately relaxed to specific events to fully account for geometrical and elastic deformations. By construction, the algorithm also automatically identifies cases when the topology does not correspond to a single geometry, ensuring that the basic approximations are valid for all events. For efficiency, k-ART also includes local force calculations, allowing sublinear scaling with system size, and an exact handling of flickers extended from the mean-rate method [30]. Other acceleration techniques, such as parallel handling of event relaxation and generation, are also implemented in the current version of the k-ART package.

To demonstrate k-ART's versatility, we applied the algorithm to three problems: vacancy-interstitial annihilation in c-Si, interstitial diffusion in Fe, and relaxation of a-Si. Clearly the algorithm, although slower than standard KMC, can handle accurately complex systems with many tens of thousands of topologies much faster than MD, opening the possibility of studying problems that have long remained out of reach of simulation.

ACKNOWLEDGMENTS

This work has been 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 *Calcul Québec* (CQ) for generous allocations of computer resources.

APPENDIX: THE MEAN RATE METHOD

Following Puchala *et al.* [30], the system is separated in transient states and absorbing states. To determine the probability to exit the basin to state x , we calculate the transition probability matrix $\underline{\mathbf{T}}$, with components

$$T_{ji} = \frac{R_{i \rightarrow j}}{\sum_k R_{i \rightarrow k}} = \tau_i^{-1} R_{i \rightarrow j}, \quad (\text{A1})$$

where $R_{i \rightarrow j}$ is the rate going from basin state i to basin state j , and the summation is over all basin and exit states k . τ_i^{-1} , the reciprocal of the sum of all rates leaving state i , is the mean residence time in state i each time it is visited. The occupation probability vector of all basin states after in-basin jump m (and before $m + 1$), $\Theta(m)$ is thus given by repeated application of $\underline{\mathbf{T}}$ to the initial occupation probability $\Theta_i(0) = \delta_{is}$, where s is the starting state. The sum of the occupation probabilities over all possible number of jumps gives the average number each basin state is visited:

$$\Theta^{\text{sum}} = \sum_{m=0}^{\infty} \underline{\mathbf{T}}^m \Theta(0) = (\mathbb{1} - \underline{\mathbf{T}})^{-1} \Theta(0), \quad (\text{A2})$$

from which the mean residence time in basin state i before leaving the basin can be calculated:

$$\tau_i = \tau_i^{-1} \Theta_i^{\text{sum}}. \quad (\text{A3})$$

These residence times are then used to accelerate the basin exit rates from basin state i to exit state j according to

$$\langle R_{i \rightarrow j} \rangle = \frac{\tau_i}{\sum_k \tau_k} R_{i \rightarrow j}, \quad (\text{A4})$$

with k summing over all basin states. The next KMC step is then determined using standard KMC rules, using these accelerated rates.

In contrast to the first passage time analysis (FPTA) [30,63], the mean rate method is computationally much simpler, as it requires a single matrix inversion to calculate the modified rates, after which the ordinary KMC rules apply. This comes at a cost: There is no correlation between the randomly determined residence time and the selected exit state. Puchala *et al.* find [30] that in measuring average quantities after many steps, both MRM and FPTA yield the same results.

-
- [1] A. B. Bortz, M. H. Kalos, and J. L. Lebowitz, *J. Comput. Phys.* **17**, 10 (1975).
 [2] K. A. Fichthorn and W. H. 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*, NATO Science Series, Series II: Mathematics, Physics and Chemistry, edited by

- K. Sickafus, E. Kotomin, and B. Uberuaga, Vol. 235 (Springer, Dordrecht, Netherlands, 2007), pp. 1–23.
 [5] D. R. Mason, R. E. Rudd, and A. P. 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 H. Jónsson, *J. Chem. Phys.* **111**, 7010 (1999).

- [9] F. Hontinfinde, A. Rapallo, and R. Ferrando, *Surf. Sci.* **600**, 995 (2006).
- [10] T. F. Middleton and D. J. Wales, *J. Chem. Phys.* **120**, 8134 (2004).
- [11] Y. Fan, A. Kushima, S. Yip, and B. Yildiz, *Phys. Rev. Lett.* **106**, 125501 (2011).
- [12] F. El-Mellouhi, N. Mousseau, and L. J. Lewis, *Phys. Rev. B* **78**, 153202 (2008).
- [13] A. Kara, O. Trushin, H. Yildirim, and T. S. Rahman, *J. Phys. Condens. Matter* **21**, 084213 (2009).
- [14] G. T. Barkema and N. Mousseau, *Phys. Rev. Lett.* **77**, 4358 (1996).
- [15] R. Malek and N. Mousseau, *Phys. Rev. E* **62**, 7723 (2000).
- [16] B. D. McKay, *Congressus Numerantium* **30**, 45 (1981).
- [17] Y. Song, R. Malek, and N. Mousseau, *Phys. Rev. B* **62**, 15680 (2000).
- [18] F. Valiquette and N. Mousseau, *Phys. Rev. B* **68**, 125209 (2003).
- [19] F. El-Mellouhi, N. Mousseau, and P. Ordejón, *Phys. Rev. B* **70**, 205202 (2004).
- [20] M.-A. Malouin, F. El-Mellouhi, and N. Mousseau, *Phys. Rev. B* **76**, 045211 (2007).
- [21] M.-C. Marinica, F. Willaime, and N. Mousseau, *Phys. Rev. B* **83**, 094119 (2011).
- [22] E. Machado-Charry, L. K. Béland, D. Caliste, L. Genovese, T. Deutsch, N. Mousseau, and P. Pochet, *J. Chem. Phys.* **135**, 034102 (2011).
- [23] C. Lanczos, *Applied Analysis* (Dover, New York, 1988).
- [24] P. Pulay, *Chem. Phys. Lett.* **73**, 393 (1980); R. Shepard and M. Minkoff, *Mol. Phys.* **105**, 2839 (2007).
- [25] H. Yildirim, A. Kara, and T. S. Rahman, *Phys. Rev. B* **76**, 165421 (2007).
- [26] J. A. Stroschio and D. T. Pierce, *Phys. Rev. B* **49**, 8522 (1994).
- [27] N. I. Papanicolaou and H. Chamati, *Comput. Mater. Sci.* **44**, 1366 (2009).
- [28] F. Glover and M. Laguna, *Tabu Search* (Kluwer Academic, Norwell, MA, 1997).
- [29] M. V. Chubynsky, H. Vocks, G. T. Barkema, and N. Mousseau, *J. Non-Cryst. Solids* **352**, 4424 (2006).
- [30] B. Puchala, M. L. Falk, and K. Garikipati, *J. Chem. Phys.* **132**, 134104 (2010).
- [31] M. P. Allen and D. J. Tildesley, *Computer Simulation of Liquids* (Clarendon, New York, 1987).
- [32] L. Verlet, *Phys. Rev.* **159**, 98 (1967).
- [33] F. H. Stillinger and T. A. Weber, *Phys. Rev. B* **31**, 5262 (1985).
- [34] D. Maroudas and R. A. Brown, *Phys. Rev. B* **47**, 15562 (1993).
- [35] T. Sinno, Z. K. Jiang, and R. A. Brown, *Appl. Phys. Lett.* **68**, 3028 (1996).
- [36] M. Tang, L. Colombo, J. Zhu, and T. Diaz de la Rubia, *Phys. Rev. B* **55**, 14279 (1997).
- [37] L. A. Marqués, L. Pelaz, J. Hernández, J. Barbolla, and G. H. Gilmer, *Phys. Rev. B* **64**, 045214 (2001).
- [38] G. H. Gilmer, T. Diaz de la Rubia, D. M. Stock, and M. Jaraiz, *Nucl. Instrum. Meth. Phys. Res. Sec. B: Beam Interactions with Materials and Atoms* **102**, 247 (1995).
- [39] C.-C. Fu, J. Dalla Torre, F. Willaime, J.-L. Bocquet, and A. Barbu, *Nat. Mater.* **4**, 68 (2005).
- [40] J.-P. Massoud, S. Bugat, B. Marini, D. Lidbury, and S. Van Dyck, *J. Nucl. Mater.* **406**, 2 (2010).
- [41] L. Malerba *et al.*, *J. Nucl. Mater.* **406**, 7 (2010).
- [42] D. A. Terentyev, L. Malerba, and M. Hou, *Phys. Rev. B* **75**, 104108 (2007).
- [43] N. Anento, A. Serra, and Y. N. Osetsky, *Modell. Simul. Mater. Sci. Eng.* **18**, 025008 (2010).
- [44] C. S. Becquart and C. Domain, *Phys. Status Solidi B* **247**, 9 (2010).
- [45] G. J. Ackland, M. I. Mendeleev, D. J. Srolovitz, S. Han, and A. V. Barashev, *J. Phys. Condens. Matter* **16**, S2629 (2004).
- [46] M. I. Mendeleev, S. Han, D. J. Srolovitz, G. J. Ackland, D. Y. Sun, and M. Asta, *Philos. Mag.* **83**, 3977 (2003).
- [47] L. Malerba *et al.*, *J. Nucl. Mater.* **406**, 19 (2010).
- [48] J. Stadler, R. Mikulla, and H.-R. Trebin, *Int. J. Mod. Phys. C* **8**, 1131 (1997).
- [49] C.-C. Fu, F. Willaime, and P. Ordejón, *Phys. Rev. Lett.* **92**, 175503 (2004).
- [50] R. A. Johnson, *Phys. Rev.* **134**, A1329 (1964).
- [51] N. Mousseau, *Defect Diffus. Forum* **194–199**, 775 (2001).
- [52] N. Mousseau and G. T. Barkema, *Phys. Rev. B* **61**, 1898 (2000).
- [53] H. Kallel, N. Mousseau, and F. Schiettekatte, *Phys. Rev. Lett.* **105**, 045503 (2010).
- [54] S. Roorda, W. C. Sinke, J. M. Poate, D. C. Jacobson, S. Dierker, B. S. Dennis, D. J. Eaglesham, F. Spaepen, and P. Fuoss, *Phys. Rev. B* **44**, 3702 (1991).
- [55] S. Coffa, J. M. Poate, D. C. Jacobson, and A. Polman, *Appl. Phys. Lett.* **58**, 2916 (1991).
- [56] S. Coffa, J. M. Poate, D. C. Jacobson, W. Frank, and W. Gustin, *Phys. Rev. B* **45**, 8355 (1992).
- [57] J. F. Mercure, R. Karmouch, Y. Anahory, S. Roorda, and F. Schiettekatte, *Phys. Rev. B* **71**, 134205 (2005).
- [58] F. Wooten, K. Winer, and D. Weaire, *Phys. Rev. Lett.* **54**, 1392 (1985).
- [59] G. T. Barkema and N. Mousseau, *Phys. Rev. B* **62**, 4985 (2000).
- [60] R. L. C. Vink, G. T. Barkema, W. F. van der Weg, and N. Mousseau, *J. Non-Cryst. Solids* **282**, 248 (2001).
- [61] M. Durandurdu, D. A. Drabold, and N. Mousseau, *Phys. Rev. B* **62**, 15307 (2000).
- [62] S. Plimpton, *J. Comput. Phys.* **117**, 1 (1995); [<http://lammps.sandia.gov>].
- [63] M. A. Novotny, *Phys. Rev. Lett.* **74**, 1 (1995).