Understanding long-time vacancy aggregation in iron: A kinetic activation-relaxation technique study

Peter Brommer,[1,2,*] Laurent Karim Béland,[1,†] 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]*Department of Physics and Centre for Scientific Computing, University of Warwick, Gibbet Hill Road, Coventry CV4 7AL, United Kingdom*

(Received 16 June 2014; revised manuscript received 9 October 2014; published 30 October 2014)

Vacancy diffusion and clustering processes in body-centered-cubic (bcc) Fe are studied using the kinetic activation-relaxation technique (k-ART), an off-lattice kinetic Monte Carlo method with on-the-fly catalog building capabilities. For monovacancies and divacancies, k-ART recovers previously published results while clustering in a 50-vacancy simulation box agrees with experimental estimates. Applying k-ART to the study of clustering pathways for systems containing from one to six vacancies, we find a rich set of diffusion mechanisms. In particular, we show that the path followed to reach a hexavacancy cluster influences greatly the associated mean-square displacement. Aggregation in a 50-vacancy box also shows a notable dispersion in relaxation time associated with effective barriers varying from 0.84 to 1.1 eV depending on the exact pathway selected. We isolate the effects of long-range elastic interactions between defects by comparing to simulations where those effects are deliberately suppressed. This allows us to demonstrate that in bcc Fe, suppressing long-range interactions mainly influences kinetics in the first 0.3 ms, slowing down quick energy release cascades seen more frequently in full simulations, whereas long-term behavior and final state are not significantly affected.

DOI: 10.1103/PhysRevB.90.134109 PACS number(s): 61.80.Az, 61.72.jd, 61.82.Bg, 02.70.−c

I. INTRODUCTION

The study of irradiation damage is central to understanding materials kinetics. Ion bombardment leaves behind defect cascades of self-interstitial atoms (SIA) and vacancies that can migrate and affect the mechanical properties of the materials, forming cracks or swelling, for example [1,2]. Because of the microscopic nature of these processes, numerical simulations are needed to provide crucial understanding about atomistic details that is difficult to obtain from experiments. While simulations have contributed significant information on this topic, they have also been limited due to the extended time scale over which many of these processes take place: a time scale that is out of reach from standard molecular dynamical simulations. Kinetic Monte Carlo (KMC) [3] provides a solution to reaching long-time dynamics, however, the standard implementation requires an upfront knowledge of the relevant barriers and cannot take into account crucial elastic deformations [2]. While results from such simulations are enlightening, their quantitative validity is limited since the full details of local atomic configurations can affect greatly diffusion kinetics [4]. Over the last few years, numerous improvements have been proposed to standard KMC simulations, in order to overcome these limitations [5–7] (see Ref. [8] for additional references), leading to a new interest in the fundamental mechanisms associated with defect diffusion in nuclear materials.

Radiation-damage recovery in Fe is generally described in five stages [2,9]. At low temperature, the first two stages are associated with interstitial-vacancy (I–V) pair recombination; the third one, below 200 K, with di-interstitial diffusion, the fourth one, near room temperature, with vacancy diffusion and the final one, at 520–550 K, with defect-cluster dissociation. In spite of this apparent clean energy differentiation between the various stages, a number of recent results suggest that damage recovery is a more complex process, where the various defects play an intertwined role [10,11].

Using the kinetic activation-relaxation technique (k-ART), an off-lattice kinetic Monte Carlo method with on-the-fly catalog building capabilities, we revisit this issue in order to clarify these mechanisms. K-ART allows us to reach experimental time scales, many orders of magnitude longer than reachable by molecular dynamics while incorporating exact elastic effects and identifying the atomistic details of diffusion mechanisms at every step. In this paper, we focus on vacancies, which dominate diffusion above room temperature. This is also the temperature range where KMC techniques, which move the system across a potential energy surface determined at 0 K, are the method of choice. For simulations at higher temperatures, when the time scales accessible to KMC decrease due to increasing reaction rates, alternative methods such as molecular dynamics (MD) may be better suited to the task. At the temperatures under study here, following diffusion mechanisms and pathways allows us to show that we can recover previous simulation results and provide a number of new insights regarding both the diffusion process of small vacancy clusters and the aggregation phase.

II. METHODS

Simulations presented here are based on kinetic ART [5,12], a powerful and versatile algorithm with on-the-fly event catalog building capacity and exact treatment of elasticity.

*Current address: Centre for Predictive Modelling, School of Engineering, Library Road, Coventry CV4 7AL, United Kingdom; p.brommer@warwick.ac.uk

†Current address: Materials Science and Technology Division, Oak Ridge National Laboratory, Oak Ridge, Tennessee 37831-6138, USA.

‡Current address: Department of Mechanical and Aerospace Engineering, Carleton University, 1125 Colonel By Drive, Ottawa, Ontario, Canada K1S 5B6.

§normand.mousseau@umontreal.ca

 ©2014 American Physical Society

K-ART uses the activation-relaxation technique (ART nouveau) for event searching [10,13–15], an efficient saddle-point searching method that was shown to be particularly effective for finding events in iron [10]. Combining ART nouveau with a topological analysis package [16], k-ART is not limited to crystalline environments and it can classify configurations and events with any level of disorder. Over the last few years, k-ART was applied with success to complex systems such as ion-bombarded [17] and amorphous silicon [18]. A description of the method can be found in Ref. [12]. We present here only the basic steps associated with a k-ART run:

(1) Starting from a local minimum, we first determine the topology associated with the local environment of each atom in the cell.

(2) We then verify in the event catalog whether each topology currently present has been already observed and has been sufficiently searched for events.

(3) If not, ART nouveau searches are launched on the topologies to complete the catalog; if yes, we move directly to the next step.

(4) All events associated with the topologies characterizing the configuration are placed in the list of available events for this local minimum.

(5) Events that account for at least 99.99% of the total rate are fully reconstructed and their saddle point fully relaxed to include all elastic deformations associated with the specific configuration.

(6) The total rate is computed again, the clock is moved forward according to a Poisson distribution, an event is selected at random with the appropriate weight.

(7) The event is applied and the configuration is relaxed into a new minimum. We can go back to step 1.

We use a 2000-atom body-centered-cubic (bcc) Fe crystal ($10 \times 10 \times 10$ cubic unit cells) with periodic boundary conditions and the Ackland-Mendelev effective interaction potential A04 [19]. A number of simulations are also performed with the potential M07 by Malerba *et al.* [20] in order to assess the influence of the effective potential used. We remove between one and six atoms to focus on the vacancy diffusion associated with the fourth stage in radiation-damaged recovery. We also remove 50 atoms to compare vacancy-aggregation results with other simulations [21–24] and positron-annihilation experiments [25].

All simulations are performed with a constant prefactor of 5×10^{12} s^{-1}, which was shown to be a good approximation [21,26–28]. We use a radius of 5.6 Å for the topological classification, with an initial 50 searches for a new topology. Additional searches are launched based on logarithmic increase in the number of times this topology is observed, i.e., a topology appearing n times is searched $50\ (1 + \log n)$ times in total. We ignore events associated with perfectly crystalline environments as barriers are too high to be selected on the simulation time scale. For the initial configuration of the 50-vacancy system, we identify between 616 and 766 different generic events. By the time the simulation reaches 1 ms, we will have generated tens of thousands of events (between 38 056 and 82 924 events for the four runs). Comparing with other off-lattice KMC simulations performed on the same system but with a different algorithm (SEAKMC [7]), we observe that both approaches lead to similar time scales for vacancy aggregation, suggesting that both methods manage to generate physically relevant pathways, as discussed in Sec. III C [24,29].

The first simulations are run at 573 K to allow an essential comparison with earlier lattice KMC results and analytical estimations [30]. Aggregation of large clusters is performed near room temperature, in the vacancy-dominated temperature regime.

A. Handling small barriers

A major obstacle to overcome in achieving efficient KMC simulations is the problem of small barriers, which plagues many KMC simulation techniques. If a system has a basin, i.e., a group of states separated by energy barriers that are significantly lower than those connecting the basin to other states, then standard KMC will show frequent jumps within the basin (so-called flickers). Such flickers do not progress the simulation and provide little insight beyond the first tour through those states. To avoid these fruitless KMC steps, k-ART uses the basin-autoconstructing mean rate method (bac-MRM) [12,31]. With bac-MRM, final states of events with both forward and backward barriers below a user-specified basin threshold are dynamically added to the current basin. K-ART then averages over all possible transitions within a basin to pick the next step. This ensures that each k-ART step advances the simulation by either expanding a basin (until its complete exploration) or by leaving the basin behind. MRM-like acceleration methods have successfully been used recently in other KMC codes [32].

Because the low-barrier treatment is statistically exact, the choice of the low-barrier threshold does not affect the system's kinetics. Since particular trajectories are lost within the basin, however, it is preferable to use the lowest threshold possible for specific systems in order to preserve atomistic details about these. Typically, the barrier threshold needs to be increased over the course of a vacancy clustering simulation, as flickers on growing energy scales become significant obstacles. For example, in the 50-vacancy clustering simulations, we start with initial threshold of 0.152 eV, increasing it to 0.75 eV as the system evolves (see also Sec. III C for details).

B. Generic embedding: Eliminating long-range effects

To study the significance of long-range interactions between defects, we created a modified version of k-ART that intentionally suppresses all structural information in barrier determination beyond the local environment used for topological identification. In a normal k-ART run, these elastic interactions are accounted for during event reconstruction (step 5 in the sequence described in Sec. II above), when the most important events for a certain topology are refined for each atom that shares that topology. These atoms have the same local environment up to the topology cutoff radius, but may differ beyond that. A crude approach to ignore these differences would be to just skip the refinement step, but this would introduce a bias, as every single catalog event was found in a certain configuration of all atoms.

To counter this problem, we introduce *generic embedding* (GE): We launch the ART nouveau searches (step 3 in the

k-ART method) in a modified structure, where the atoms within the sphere used for topological classification (5.6 Å radius) are embedded in a defect-free crystal of bcc iron. The events found in this way are free from any hidden bias caused by atomic configurations outside the environment classification. These events are then used without further refinement (which would defy the purpose of GE). The final configuration of a GE k-ART event is relaxed into the new minimum as in the standard k-ART recipe.

GE simulations are computationally less demanding than full k-ART simulations. The fact that no event reconstruction is necessary far outweighs the overhead of creating the embedding environment, which is negligible compared to performing an ART nouveau search in that structure. During the early simulation phases of a system far from equilibrium, when there are many new local environments introduced after almost every step, both methods use a comparable amount of CPU time per KMC step. But, as soon as the system at least temporarily reaches a stationary state with no new environments to search, GE can perform a KMC step in a few CPU minutes, whereas the barrier refinement limits full k-ART to about 1 CPU hour per KMC step for the simulations described in Sec. III C. Overall, those GE simulations were a factor 4–5 faster than their full k-ART counterparts for a comparable number of KMC steps.

C. Force field

As k-ART simulations, like other improved or adaptive KMC simulations, require a large number of force evaluations to build a catalog (on the order of 300 to 600 per successful saddle-point search [8]), they are only feasible with a classical effective potential or force field, where the interactions between electrons and nuclei are reduced to an effective interaction between atoms. Such simplification comes at a cost: While *ab initio* calculations using density functional theory (DFT) methods do not in principle require any additional input, the force-field parameters need to be determined beforehand. The parameters are usually obtained by optimizing certain physical properties (like activation barriers) within the limit of the selected model. However, this implies that there is a loss of generality; properties *not* included in the optimization process may not be equally well reproduced.

Iron is typically described by embedded atom method (EAM) potentials [33]

$$V = \sum_i U\left[\sum_j \rho(r_{ij})\right] + \sum_{ij} \phi(r_{ij}), \qquad (1)$$

where $U[n]$ is the embedding energy of an atom in a local density $n = \sum_j \rho(r_{ij})$, defined as the sum of the contributions from neighbors at distance r_{ij} via a transfer function $\rho(r_{ij})$. $\phi(r_{ij})$ is a pair potential function. The properties of the material in simulation are determined by the choices for the three functions $U(n)$, $\rho(r_{ij})$, and $\phi(r_{ij})$.

One of the first Fe potentials created specifically for the simulation of defects was by Ackland *et al.* (referenced as A04) [19], based on an earlier version by the same authors [34]. While these potentials did not include the vacancy migration barrier energy in the fit, they can reproduce the activation energy for self-diffusion, which is the sum of vacancy formation and migration energies, to within 10% of the experimental value [34]. For this potential, the energy profile for vacancy migration is double humped, i.e., there is a minimum at the midpoint of the trajectory, and the maximum is slightly displaced towards either end. This is an artifact of the potential; DFT calculations do not show this intermediate minimum [20]. This is why we also use here a potential developed by Malerba *et al.* (referenced as M07) [20], which uses vacancy formation and migration energies directly in the fitting procedure, removing the double-hump problem. M07 also yields a slightly higher value for the migration barrier (0.68 eV compared to 0.63 eV with A04), in agreement with DFT calculations. This difference directly influences the time scales for vacancy kinetics and will be addressed in Sec. III A.

III. RESULTS

A. Small vacancy clusters

To establish a reference point, we first examine the monovacancy using the A04 potential. At low temperature, in simulations over 5000 KMC steps corresponding to a time scale on the order of 1 s at 300 K, the vacancy first-nearest-neighbor (1NN) jump (0.64 eV) takes place directly or via a split vacancy position (0.12 eV below saddle-point barrier), with 100% probability. The barrier heights agree with literature values [20]. As the temperature rises above 700 K, the next highest barrier, corresponding to the third-nearest-neighbor (3NN) split vacancy (1.07 eV), starts occurring (0.1% at 700 K and 1.3% at 1200 K); given the energy difference between the lowest barrier and the next one, the diffusion remains controlled by the 0.64-eV barrier even at high temperature.

In a recent comparative study of effective potentials for iron, Malerba *et al.* [20] report an artificial local minimum in the vacancy 1NN diffusion barrier that is not found in first-principles simulations (cf. Sec. II C). This local minimum corresponds to the split vacancy position found in our simulations with the A04 potential. The authors suggest a modified potential referenced as M07 that does not show this artifact, however, its 1NN vacancy diffusion barrier is slightly higher at 0.68 eV. At 573 K, this difference in barrier height corresponds to a rate that is lower by a factor of 2.7. Additionally, the A04 k-ART simulation finds twice the number of events with the 1NN barrier for a single vacancy, as it identifies the direct process and the artificial one via an intermediate split vacancy as physically distinct events. This accounts for a total rate difference of factor 5.4 between M07 and A04. As the use of the A04 potential is prevalent in point-defect simulations in iron [4,7,21–24,35–37], we will also use it for reasons of comparability, but will compare with results obtained with M07 in certain cases.

For a divacancy, we find that both A04 and M07 reproduce three distinct bound states at 4NN, 1NN, and 2NN position (in order of increasing binding energy). The energies of divacancy states and connecting saddle points are shown in Fig. 1. DFT calculations predict an additional bound state at the 5NN position, but this is usually not reproduced by an effective interaction potential [4]. Our results for A04 agree with a previous comparison between various force fields and

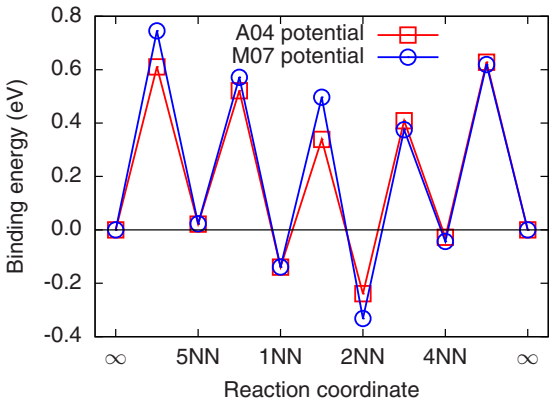

FIG. 1. (Color online) Energy landscape of the divacancy states and saddle points according to the A04 and M07 force fields. The energy zero corresponds to two isolated vacancies (∞).

first-principles results in Ref. [4] (where the A04 potential is called AM); the M07 was not part of this study. The most significant difference between M07 and A04 is that the 2NN state is more strongly bound by 0.1 eV for M07 with barriers to the 1NN (4NN) state higher by 0.25 (0.06) eV, respectively. This implies that the M07 favors a divacancy propagation via an intermediate 4NN state rather than a 1NN state, while still being more strongly bound overall. At a simulation temperature of 573 K, the divacancy dissociates within a few tens of KMC steps.

Next, we study lifetime and diffusivity of small vacancy clusters (one–six vacancies). These simulations are performed in a 2000-atom box at 573 K. The vacancies are initially placed in the lowest-energy configuration for a vacancy cluster of that size. A cluster is deemed decayed, if at least one of its vacancies moves further away than the 4NN site from each of the remaining cluster vacancies. If a cluster reforms during a simulation, the lifetime is sampled another time. The diffusion coefficient quantifies the diffusion of the cluster as a whole during its lifetime. The results can be found in Table I. For monovacancies and trivacancies, the diffusion coefficients for A04 are higher than those of M07 by a factor of 5, which agrees with the rate difference between the two potentials. In larger clusters, the results from the two potentials diverge: M07 yields significantly longer lifetimes and lower diffusion coefficients than A04. This may be due to the artificial

TABLE I. Lifetimes and diffusion constants of vacancy clusters at 573 K for the A04 [19] and M07 [20] potentials. The A04 divacancy dissociates before a diffusion coefficient can be measured.

	Thermal lifetime (s)		Diffusion coefficient ($nm^2\ s^{-1}$)	
Cluster	A04	M07	A04	M07
v	N/A		1×10^7	2×10^6
v_2	3×10^{-8}	5×10^{-7}	–	2×10^5
v_3	6×10^{-8}	2×10^{-6}	2×10^7	4×10^6
v_4	1×10^{-6}	1×10^{-4}	2×10^6	4×10^4
v_5	5×10^{-6}	2×10^{-3}	7×10^4	1×10^3
v_6	4×10^{-5}	4×10^{-2}	5×10^3	3×10^1

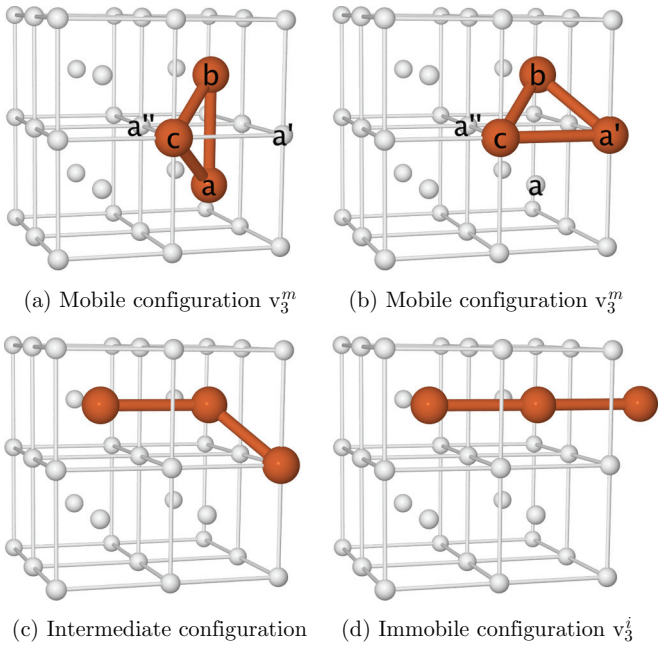

(a) Mobile configuration v_3^m (b) Mobile configuration v_3^m

(c) Intermediate configuration (d) Immobile configuration v_3^i

FIG. 2. (Color online) Three-vacancy cluster diffusion. Small light spheres are occupied lattice sites, large dark spheres are vacancies.

minimum near the saddle point between two nearest-neighbor vacancy positions, and thus be an artifact of the potential. In agreement with Ref. [30], we find that the 3-vacancy has an exceptionally high diffusion coefficient. This is due to an extremely effective migration mechanism with a low activation barrier (cf. Ref. [2]).

In the mobile ground-state configuration v_3^m [bound by -0.54 eV (-0.67 eV) for A04 (M07)], vacancies a and b at 2NN separation have a third vacancy c on a common 1NN site [see Fig. 2(a)]. The cluster diffuses by a 1NN vacancy hop of either a or b to one of the two sites that are 1NN to b or a and 2NN to c. In Fig. 2(a), a or b could hop to either sites a' or a''. Figure 2(b) shows one of the two possible results. For the A04 (M07) potential, the barrier for this event is 0.46 eV (0.52 eV). This mobile v_3 configuration can transition to an immobile v_3^i configuration, which is located 0.06 eV (0.02 eV) above the ground state, by a multistep process [cf. Figs. 2(c) and 2(d)]. However, the barriers to leave this metastable state are significantly higher; at 573 K this results in a mean residence time that is longer by a factor of about 7 (5) with respect to the mobile state. Figure 3 presents the energies and barrier heights along the trajectory shown in Fig. 2 for the two potentials.

For larger clusters, our simulations yield significantly longer lifetimes and lower diffusion coefficients than the lattice KMC simulations by Borodin *et al.* [30]. This seems to imply that those simulations are not able to fully capture the complete dynamics of the system, leaving aside important diffusion pathways or neglecting elastic effects that are taken into account using k-ART.

B. Six-vacancy aggregation

In this section and the following one, we assign a vacancy to a cluster, if it is at nearest or next-nearest atom site

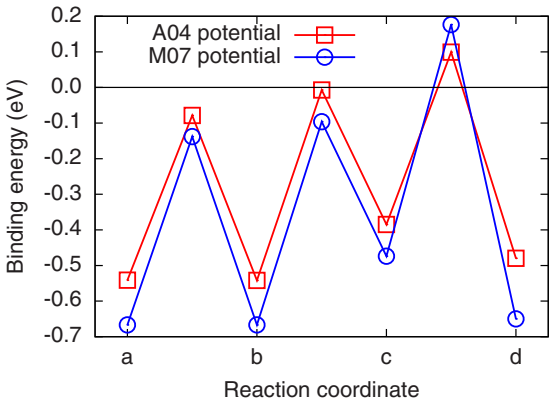

FIG. 3. (Color online) Binding energy of minima and saddle points along the *a*-*b*-*c*-*d* trajectory (cf. Fig. 2). Note that k-ART does not provide any energy information for intermediate points.

separation from at least one other vacancy of that cluster. At the lower simulation temperature, the 4NN vacancy, which is technically still bound, is rarely sampled, and it does not affect the overall aggregation evaluation whether or not that site is included in a cluster definition. We now look at vacancy aggregation by removing six atoms at random positions from the 2000-atom box. The ground state for this system, at 2.63 eV below six isolated vacancies, is a single cluster, with four vacancies sitting at the corners of one face of the cubic unit cell, and the other two in the body centers of the two cubic unit cells that share this face. Four independent simulations with the A04 potential are launched at 50 °C, in the vacancy-dominated diffusion regime, until vacancies are completely agglomerated (1600–5400 KMC steps not counting the intrabasin flickering moves) for about 0.1 s. The time evolution for these clusters is presented in Fig. 4 and shows the evolution of the cluster-size distribution, the total square displacement, and the energy. We note that even though some dimers are formed very early on, as soon as 50 μs, diffusion becomes important only after almost 0.5 ms. It is completed typically within 20 ms, when the full six-vacancy cluster is formed.

As is seen in Fig. 5, all simulations show a similar aggregation pathway, with the formation of a first divacancy, followed by a second (skipped in run 3) and then the formation of more complex trivacancy, a divacancy with a tetravacancy, a pentavacancy, and, finally, the ground state. Overall, the diffusivity is about 10^6 Å^2/s but certain clusters stand out. This is the case for the mobile trivacancy v_3^m, which shows the fastest diffusivity, while the pentavancy is extremely slow (see Fig. 6, which plots the iron self-diffusion constant for various cluster configurations). This agrees with the findings from the three-vacancy simulations presented in the previous section and can be observed in Fig. 4, where the fastest increase in total square displacement occurs in the presence of trivacancy. Run 4, which does not sample a trivacancy, shows a final square displacement that is between three and four times smaller than the three other runs, and confirms the rich pathway diversity to be found even for such a small cluster.

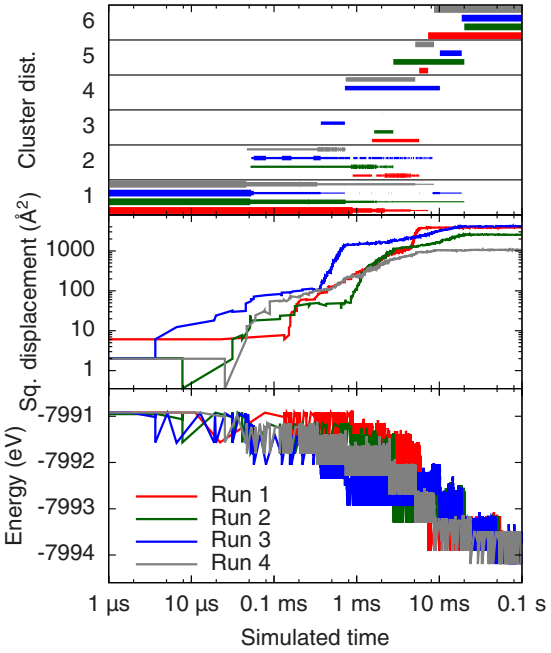

FIG. 4. (Color online) Trajectory of four independent runs with six randomly placed vacancies. The bottom plot shows the cohesive energy of the system over time. The central plot shows the total squared displacement. The top plot shows the cluster-size distribution, where the linewidth is proportional to the number of vacancies in clusters of the respective size. High atomic mobility appears to be correlated with clusters of size three.

C. 50-vacancy aggregation

Having established the diffusion mechanisms for small aggregates, we can now see what role they play in an environment with a larger concentration of defects. We prepare four independent models of 2000-atom Fe cell with 50 vacancies each inserted at random. Figure 7 shows the trajectory for these four runs, with a similarly energy profile associated to three regimes: 0 to 10 μs, 10 μs to 1 ms, and from 1 ms to the end of the simulation at 1 s. To understand

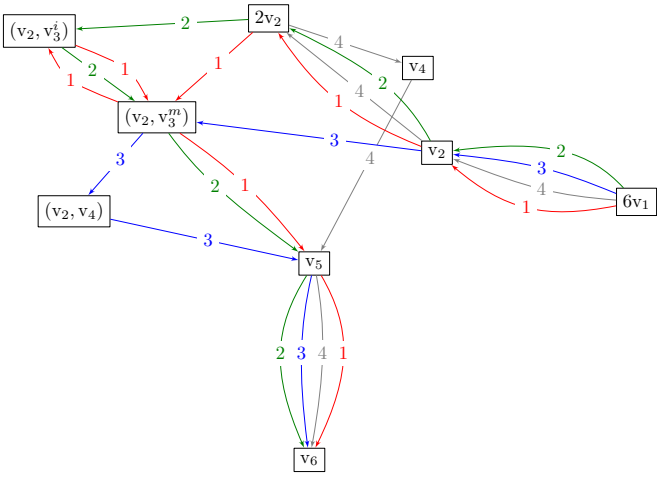

FIG. 5. (Color online) Flow chart of vacancy cluster formation in the six-vacancy system. Clusters with n vacancies are labeled v_n, monovacancies are not considered except for the initial state.

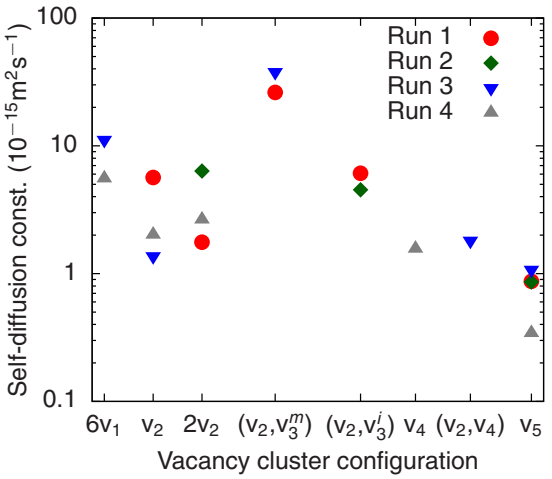

FIG. 6. (Color online) Top: diffusion rate of various cluster combinations in six-vacancy simulations. Clusters of n vacancies are denoted as v_n. Remaining vacancies are monovacancies. There are two distinct configurations of three-vacancy clusters: mobile (v_3^m) and immobile (v_3^i) (see text).

these three regimes, it is useful to look at the evolution of the average cluster size and the monovacancy fraction shown in Fig. 9 for Run 2.

The first regime is associated with internal cluster relaxation, with an almost constant cluster size and number of vacancies. The second regime shows clustering, with the fraction of monovacancies dropping to almost zero (except for transient defects) followed by sharp increase in average cluster size. This leads to a significant energy relaxation. The final regime shows the slow internal evolution of larger clusters, that vary in size between 5 and 14, with an average of 8 vacancies, in agreement with positron annihilation spectroscopy measurements on iron samples [25]. After almost one second, the relaxation is still 20 eV above the lowest-energy structure, −7789 eV, found for a single 50-vacancy cluster relaxed using k-ART, suggesting that further aggregation takes place on a longer time scale at 300 K.

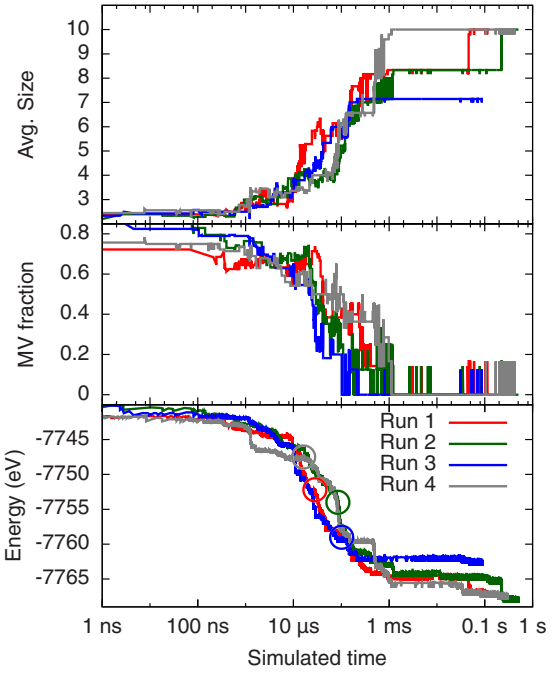

FIG. 7. (Color online) K-ART simulation results of four independent runs with randomly placed vacancies. Bottom plot: energy of the four runs. Most energy is released within the first 100 μs. Circles mark the highest effective barrier of the runs (see Fig. 8 for magnifications). Center plot: fraction of monovacancies. After about 1 ms, all vacancies are clustered. Top plot: average cluster size. The simulations are terminated when seven (run 3) or five (other runs) clusters remain and no further progress is made.

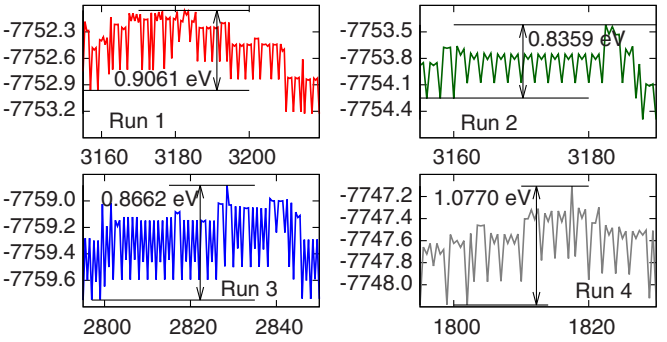

FIG. 8. (Color online) Magnifications from the four runs showing the trajectory through minima and saddle points. The highest effective barrier (i.e., difference between saddle-point energy and previous lowest energy) is highlighted. In all four plots, the abscissas are KMC steps, while the ordinates are energy in eV. For runs 1 and 2, an energetically favorable configuration is obtained within a few KMC steps. For runs 3 and 4, the energy does not drop below the former optimal energy right away. This is also an example of the "replenish and relax" mechanism described in Ref. [17].

The time scale of vacancy relaxation is almost logarithmic. For run 2, for example, 215 KMC steps are needed to reach 1 μs, 650 KMC steps to reach 10 μs, 3360 for 0.1 ms, 5000 for 1 ms, 6600 for 10 ms, and 7500 to reach 0.1 s (total 10 600 KMC steps for a 0.5-s simulation, other simulations are comparable). This is only possible because the effective barrier, i.e., the one that leads to a structural evolution, increases roughly logarithmically with time. The effective

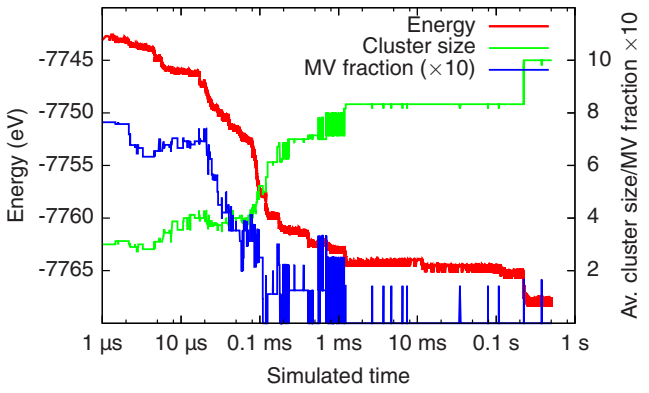

FIG. 9. (Color online) Details of run 2: Energy, average cluster size, and monovacancy fraction are plotted over simulation time. For further details, see text.

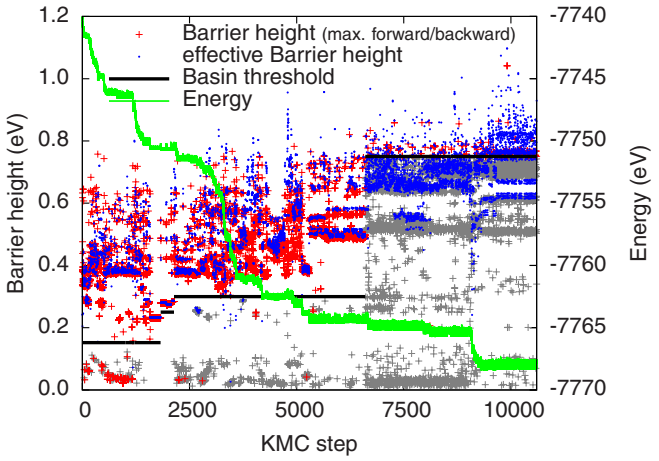

FIG. 10. (Color online) Basin threshold control and relaxation progression of run 2. The crosses show the height of the saddle point above the lower of the two states that it links. Red symbols are regular barriers, gray symbols represent barriers included in the basin. The smaller blue symbols show the effective barrier, i.e. the height of the saddle point above the minimal state obtained so far. The basin threshold is shown by the solid black line.

barrier is defined as the energy difference between the local saddle point and the lowest energy obtained previously (cf. Fig. 8, which shows the maximal effective barriers of the four runs in the relaxation phase). For relaxation simulations, this is a good measure of the time scale, as this determines the time scale to escape from the previous optimal configuration.

The efficiency of bac-MRM allows us to follow the overall relaxation over a long time scale. An example, taken from run 2, is given in Fig. 10. Crosses indicate barriers selected: gray crosses are basin exploration events that will not directly advance a simulation; red crosses are events outside of basins and may lead to progress or return to previous state. Around step 2000, we observe a flicker associated with barrier of 0.23 eV (and no more basin events). We increase the basin threshold to 0.25 eV, just to find a flicker at 0.28 eV. We must therefore increase the threshold again to allow the simulation to progress. The new threshold, set at 0.3 eV remains effective until step 6600, when it is increased to 0.75 eV.

The blue dots in Fig. 10 show the effective barrier height. This energy dictates on which time scale a relaxation takes place, while the local barrier height is closely related to the time increment per KMC step. This demonstrates the need for an effective basin acceleration scheme, as without bac-MRM the simulation would get locked in on the significantly shorter time scale associated with the lower local barriers associated with nondiffusive events.

During the course of this run, low effective barriers die out, except when there are large relaxations. In these situations, we are finding new minimal energy states frequently, including skewed events, with small forward and large backward barriers. Overall, the effective barrier increases roughly linearly with number of KMC steps or logarithmically with time.

The 50-vacancy simulations were stopped when the vacancies aggregated in 5 clusters (7 clusters for run 3), as no further progress could be seen. The smallest cluster in this stage was size 6 (run 3: 5), whose effective migration barrier is significantly above the basin threshold of 0.75 eV at the end of the simulations. To go further in time, beyond the 1-s time scale, it would therefore be necessary to further increase the basin threshold, to handle properly the high-barrier flickers that appear in this time regime. In the current bac-MRM implementation, all events for all atoms in all basin states need to be kept in memory, and the basin would comprise several hundreds, if not thousands, of states for higher basin thresholds.

Aggregation of vacancies in this system has been used to compare the efficiency of two other off-lattice KMC methods looking, for example, at the time needed to reach an energy level of −7760 eV. Using the autonomous basin climbing (ABC) method combined with nudged elastic band (NEB) and KMC methods, Fan *et al.* obtain a relaxation time scale of more than 20 000 s [36] compared with 10^{-4} to 6×10^{-5} s with k-ART. This indicates that ABC is unable to identify the most important relaxation pathways, significantly overestimating the relaxation time scale by missing relevant mechanisms [22,24]. With the self-evolving atomistic kinetic Monte Carlo (SEAKMC) method [29,38] and running 12 simulations, Xu *et al.* find a relaxation time between 3×10^{-6} and 10^{-4} s [24]. The relaxation time scales overlap with our results, but are within an interval about an order of magnitude faster than ours, a small difference given that the time scale depends exponentially on the determination of barrier height. Nevertheless, since Xu *et al.* replicate carefully our simulation setup, it is important to identify the origin of this small discrepancy and efforts are currently underway to address this issue. From our current and joint assessment with Xu and Stoller, both SEAKMC and k-ART provide the correct kinetics and the difference is due to small details in the implementation and choice of parameters that do not affect the overall results.

D. Long-range elastic effects

We repeated the four runs with 50 vacancies using the generic embedding k-ART code, thus purposefully ignoring any elastic effects on defects beyond the topological classification radius of 5.6 Å. The runs were stopped if no further progress was made. The resulting trajectories are shown in Fig. 11. Overall, the differences to a full k-ART run appear to be minor: Two runs finish up with six clusters and a final energy above −7765 eV, and two with five clusters and below −7765 eV (run 4 briefly samples four clusters, with two close clusters temporarily connected via a second-nearest neighbor site, before settling with two separate clusters; this might be an effect of GE). This is in agreement with simulations performed with SEAKMC [24], a method that involves cutting off elastic interactions at roughly 8 Å in the case of Fe vacancies, which exhibit similar trajectories.

There are real differences, however, mostly in the intermediate time regime. First, due to the absence of long-range elastic contributions, monovacancies appear to remain present until after the first ms, whereas they do no longer factor in full k-ART simulations at that time. Final structures show comparable distribution of cluster sizes in both cases.

Second, we observe differences in the trajectories in Fig. 12. During the clustering regime (cf. Sec. III C) after the initial relaxation, all full k-ART simulations show a pronounced drop

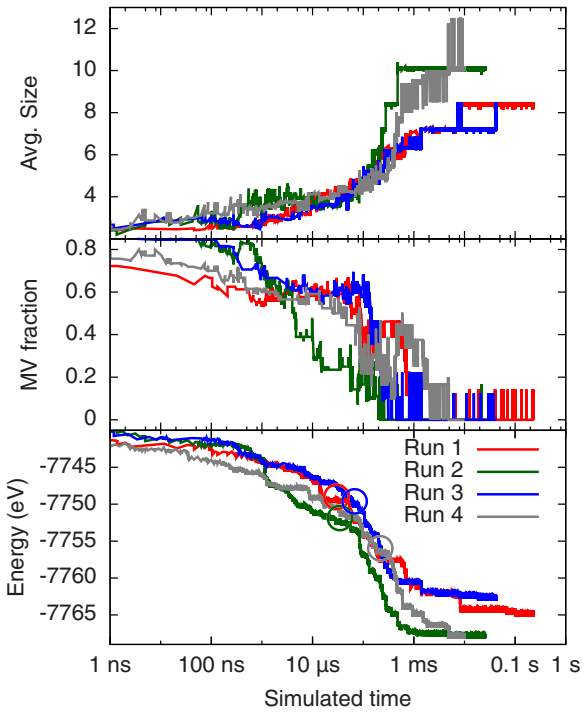

FIG. 11. (Color online) GE k-ART simulation results of four independent runs with randomly placed vacancies. Bottom plot: energy. Circles highlight the highest effective barrier of each run (0.858/0.897/0.823/0.903 eV). Center plot: fraction of monovacancies. Top plot: average cluster size.

in energy of about 4 eV within a few microseconds, where the energy release rate accelerates after first slowing down. In contrast, this behavior is found only in run 2 of the GE simulations. The other three runs yield an almost linear time-energy relationship during this regime of vacancy clustering, with a continuous and gradual release of energy. This effect could be explained by the inherent symmetry of GE simulation runs: Without elastic effects, the quick cascades seen in full

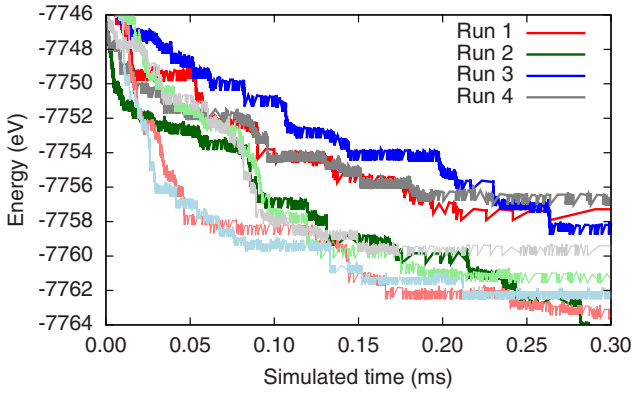

FIG. 12. (Color online) Comparison between standard and generic embedding k-ART runs from same initial configurations. GE runs are shown in dark colors and bold lines, whereas the corresponding standard runs use thinner lines and lighter colors. The time axis is nonlogarithmic.

simulations are no longer favored, and the energy decreases more uniformly.

This does not seem to lead to significant effects on the long-term evolution, where these differences on the μs–ms scale become irrelevant, as the final vacancy aggregation limits long-range elastic effects and the GE approach still treats exactly the short range (5.6 Å) elastic deformations. As we compute the energy correctly in both methods, the final states reached are more determined by thermodynamics than by kinetics. It should also be remarked that even with generic embedding the simulations are still considerably more sophisticated than with a standard, short-range, on-lattice KMC method, as even the artificially simplified GE simulations still allow complicated and off-lattice atomic defect structures.

IV. DISCUSSION AND CONCLUSION

Using the kinetic activation-relaxation technique, an off-lattice kinetic Monte Carlo method with on-the-fly catalog generation, we study vacancy diffusion and aggregation in bcc Fe providing detailed information regarding mechanisms and pathways that is difficult to obtain through standard simulation methods. Our results show a richness in the diffusion mechanisms as well as a complex balance between elastic and chemical effects.

Comparison with previously published works of the monovacancy and divacancy diffusion, using two different force fields, allows us to show that k-ART samples correctly diffusion mechanisms, including artificial ones induced by one of the force fields [4,20].

First, considering small vacancy cluster diffusion, with one to six defects, we can compare with MD simulations as well as provide a clear description of diffusion and assembly mechanisms. For example, we observe that trivacancies diffuse an order of magnitude faster than monovacancies, divacancies, and tetravacancies, as had been predicted by *ab initio* calculations of dominant cluster diffusion mechanisms for one to four vacancies [2]. Here, however, we go beyond these results and show, for example, that pentavacancy clusters are between 2 and 10 times slower and larger clusters, counting 6 to 14 vacancies, are almost immobile on the simulation time scale. The identification of pathway and diffusion mechanisms for these structures is therefore essential for understanding the long-time kinetics of materials. Methods such as k-ART, which includes both elastic and kinetic effects, greatly facilitate this characterization as compared to standard approaches where these pathways must be deduced one by one, an approach that rapidly becomes impossible when dealing with long-range elastic effects or complex systems.

Simulations of the kinetics of 50 vacancies in Fe were performed at 50 °C in the vacancy-migration regime discussed by Fu *et al.* [2]. Here, we observe the aggregation in on average 10-vacancy clusters, on a time scale between 1 ms and 1 s, depending on the specific initial vacancy distribution. Once these clusters are formed, we do not observe further aggregation on the simulation time scale. This observation is in agreement with position-annihilation spectroscopy experiments that show that room-temperature annealed Fe-implanted samples still present a signal compatible with clusters of 9 to 14 vacancies in size at room temperature. On the other hand,

larger voids with 40 to 50 vacancies are observed only after annealing at temperatures between 100 °C and 280 °C [25].

Comparative simulation runs, where long-range elastic interactions were artificially excluded, show that elastic effects are particularly significant during the clustering phase up to about 0.3 ms. There, these effects trigger quick energy release cascades that are not regularly seen otherwise. On longer time scales, long-range elastic interactions do not appear to influence either defect evolution or final states reached. As the elastic strain associated with vacancies in a compact material such as bcc iron decays rapidly with distance, it is to be expected that elastic effects play an even larger role for interstitials or in more open (e.g., covalently bound) systems. This shows the importance of including these interactions in the determination of energy barriers and thus simulation kinetics.

Beyond providing detailed information regarding the long-time kinetics of vacancies in Fe, this work opens up a new time regime to simulations, expanding the overlap between simulations and experiments and technological applications. Kinetic ART, which can treat off-lattice positions and disordered materials while including exactly long-range elastic effects, will allow us to further our study of defects kinetics in Fe, but also in a number of other technologically relevant materials, such as alloys, semiconductors, and cements. In these materials, mechanical and electronic properties are determined by microscopic kinetics and structure, and we can now answer those questions, which have remained out of reach until today.

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 de recherche du Québec – Nature et technologies (FRQ-NT). We are grateful to Calcul Québec (CQ) for generous allocations of computer resources.

[1] T. Diaz de la Rubia, H. M. Zbib, T. A. Khraishi, B. D. Wirth, M. Victoria, and M. Caturla, Nature (London) **406**, 871 (2000).
[2] C.-C. Fu, J. Dalla Torre, F. Willaime, J.-L. Bocquet, and A. Barbu, Nat. Mater. **4**, 68 (2005).
[3] A. Bortz, M. Kalos, and J. Lebowitz, J. Comput. Phys. **17**, 10 (1975).
[4] F. Djurabekova, L. Malerba, R. C. Pasianot, P. Olsson, and K. Nordlund, Philos. Mag. **90**, 2585 (2010).
[5] F. El-Mellouhi, N. Mousseau, and L. J. Lewis, Phys. Rev. B **78**, 153202 (2008).
[6] A. Kushima, X. Lin, J. Li, J. Eapen, J. C. Mauro, X. Qian, P. Diep, and S. Yip, J. Chem. Phys. **130**, 224504 (2009).
[7] H. Xu, Y. N. Osetsky, and R. E. Stoller, J. Phys.: Condens. Matter **24**, 375402 (2012).
[8] N. Mousseau, L. K. Béland, P. Brommer, J.-F. Joly, F. El-Mellouhi, E. Machado-Charry, M.-C. Marinica, and P. Pochet, J. At. Mol. Opt. Phys. **2012**, 925278 (2012).
[9] S. Takaki, J. Fuss, H. Kuglers, U. Dedek, and H. Schultz, Radiat. Eff. **79**, 87 (1983).
[10] M.-C. Marinica, F. Willaime, and N. Mousseau, Phys. Rev. B **83**, 094119 (2011).
[11] M.-C. Marinica, F. Willaime, and J.-P. Crocombette, Phys. Rev. Lett. **108**, 025501 (2012).
[12] L. K. Béland, P. Brommer, F. El-Mellouhi, J.-F. Joly, and N. Mousseau, Phys. Rev. E **84**, 046704 (2011).
[13] G. T. Barkema and N. Mousseau, Phys. Rev. Lett. **77**, 4358 (1996).
[14] R. Malek and N. Mousseau, Phys. Rev. E **62**, 7723 (2000).
[15] E. Machado-Charry, L. K. Béland, D. Caliste, L. Genovese, T. Deutsch, N. Mousseau, and P. Pochet, J. Chem. Phys. **135**, 034102 (2011).
[16] B. D. McKay, Congressus Numerantium **30**, 45 (1981).
[17] L. K. Béland, Y. Anahory, D. Smeets, M. Guihard, P. Brommer, J.-F. Joly, J.-C. Pothier, L. J. Lewis, N. Mousseau, and F. Schiettekatte, Phys. Rev. Lett. **111**, 105502 (2013).
[18] J.-F. Joly, L. K. Béland, P. Brommer, and N. Mousseau, Phys. Rev. B **87**, 144204 (2013).
[19] G. J. Ackland, M. I. Mendelev, D. J. Srolovitz, S. Han, and A. V. Barashev, J. Phys.: Condens. Matter **16**, S2629 (2004).
[20] L. Malerba, M.-C. Marinica, N. Anento, C. Bjorkas, H. Nguyen, C. Domain, F. Djurabekova, P. Olsson, K. Nordlund, A. Serra, D. A. Terentyev, F. Willaime, and C. S. Becquart, J. Nucl. Mater. **406**, 19 (2010).
[21] Y. Fan, A. Kushima, S. Yip, and B. Yildiz, Phys. Rev. Lett. **106**, 125501 (2011).
[22] P. Brommer and N. Mousseau, Phys. Rev. Lett. **108**, 219601 (2012).
[23] Y. Fan, A. Kushima, S. Yip, and B. Yildiz, Phys. Rev. Lett. **108**, 219602 (2012).
[24] H. Xu, R. E. Stoller, and Y. N. Osetsky, J. Nucl. Mater. **443**, 66 (2013).
[25] M. Eldrup and B. Singh, J. Nucl. Mater. **323**, 346 (2003).
[26] H. Yildirim, A. Kara, and T. S. Rahman, Phys. Rev. B **76**, 165421 (2007).
[27] J. A. Stroscio and D. T. Pierce, Phys. Rev. B **49**, 8522 (1994).
[28] N. I. Papanicolaou and H. Chamati, Comput. Mater. Sci. **44**, 1366 (2009).
[29] H. Xu, Y. N. Osetsky, and R. E. Stoller, J. Nucl. Mater. **423**, 102 (2012).
[30] V. A. Borodin and P. V. Vladimirov, J. Nucl. Mater. **362**, 161 (2007).
[31] B. Puchala, M. L. Falk, and K. Garikipati, J. Chem. Phys. **132**, 134104 (2010).
[32] K. A. Fichthorn and Y. Lin, J. Chem. Phys. **138**, 164104 (2013).
[33] M. S. Daw and M. I. Baskes, Phys. Rev. B **29**, 6443 (1984).
[34] M. I. Mendelev, S. Han, D. J. Srolovitz, G. J. Ackland, D. Y. Sun, and M. Asta, Philos. Mag. **83**, 3977 (2003).
[35] S. M. J. Gordon, S. D. Kenny, and R. Smith, Phys. Rev. B **72**, 214104 (2005).
[36] Y. Fan, A. Kushima, and B. Yildiz, Phys. Rev. B **81**, 104102 (2010).
[37] C. Björkas, K. Nordlund, and M. J. Caturla, Phys. Rev. B **85**, 024105 (2012).
[38] H. Xu, Y. N. Osetsky, and R. E. Stoller, Phys. Rev. B **84**, 132103 (2011).