Synchronous parallel spatially resolved stochastic cluster dynamics

Aaron Dunn^{a,b,*}, Rémi Dingreville^a, Enrique Martínez^c, Laurent Capolungo^b

^aSandia National Laboratories, Albuquerque, 87185 NM, USA

^bGeorge W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology, UMI 2958 Georgia Tech CNRS, 57070 Metz, France

^cMaterial Science and Technology Division, MST-8, Los Alamos National Laboratory, Los Alamos, 87545 NM, USA

Abstract

In this study, a spatially resolved stochastic cluster dynamics (SRSCD) model for radiation damage accumulation in metals is implemented using a synchronous parallel kinetic Monte Carlo algorithm. The parallel algorithm is shown to significantly increase the size of representative volumes achievable in SRSCD simulations of radiation damage accumulation. Weak scaling performance of the method is tested in two cases: (1) an idealized case of Frenkel pair diffusion and annihilation, and (2) a characteristic example problem including defect cluster formation and growth in α -Fe. For the latter case, weak scaling is tested using both Frenkel pair and displacement cascade damage. To improve scaling of simulations with cascade damage, an explicit cascade implantation scheme is developed for cases in which fast-moving defects are created in displacement cascades. For the first time, simulation of radiation damage accumulation in nanopolycrystals can be achieved with a three dimensional rendition of the microstructure, allowing demonstration of the effect of grain size on defect accumulation in Frenkel pair-irradiated α -Fe.

Keywords: cluster dynamics, kinetic monte carlo, parallel

1. Introduction

Structural materials in nuclear reactors are subjected to extreme environments including radiation damage and high temperatures [1, 2]. In addition, materials to be used in future nuclear reactors need to withstand even greater radiation doses and temperature ranges than those in current reactors [1, 3–6]. Recently, several types of radiation-resistant materials have been designed with microstructures that allow increased recombination of radiation damage [7–9]. Materials with microstructure length scales on the order of tens to hundreds of nanometers have shown increased resistance to radiation damage. Examples such as nano-grained and nano-laminate materials show higher radiation tolerance with increasing interface-to-volume ratio [8, 10–13]. Several studies have investigated the role of interfaces in mitigating radiation damage accumulation in metals [13–15]. However, models that can explicitly represent nanoscale microstructures and can simulate radiation damage accumulation to doses and timescales that are directly comparable to experiments have not yet been developed.

Techniques such as cluster dynamics (CD) [16–24] and kinetic Monte Carlo (KMC) [16, 17, 25–30] based methods have led to a greater understanding of the effect of irradiation on microstructure evolution in a variety of materials and irradiation conditions. In CD-based methods, populations of radiation-induced defects are tracked

via rate equations that represent the concentration evolution of defects of a given type in the material. These rate equations are solved deterministically. By contrast, KMC-based methods such as object kinetic Monte Carlo (OKMC) represent a stochastic approach for simulating defect evolution. These methods track the diffusion and interaction of individual defects inside simulated volumes, according to Arrhenius laws governing their diffusion and binding behavior. Both CD and KMC methods derive their input parameters such as binding and migration energies from a combination of experimental, atomistic, and ab-initio results and have been shown to produce similar results for defect evolution in simplified cases [16, 17, 31, 32]. However, the benefits and limitations of each method are quite different. CD models are typically very computationally efficient and can model large doses and timescales, and grouping schemes such as that of Golubov et al. [33] allow simulation of very large clusters. However, these methods are typically limited to modeling systems with only a few mobile defects and chemical species, limiting the number of reactions that can occur between defects [16, 17, 20–22, 32]. KMC models can model very complex defect behaviors and can capture such effects as spatial correlations between defects [34], but are typically limited by computation time and simulation volumes used for complex studies such as cascade damage accumulation are typically limited to approximately 100 nm in length [25, 28, 29].

Spatially resolved stochastic cluster dynamics (SRSCD)

*email: adunn32@gatech.edu

[35] has recently been developed as an alternate method for simulating defect accumulation in irradiated materials. In SRSCD, defects are assumed to be homogeneously distributed within volume elements and the reaction rates for allowed reactions such as clustering, dissociation, and diffusion between volume elements are derived from classical CD methods. However, reactions and time steps are chosen according to a classical kinetic Monte Carlo algorithm. It therefore avoids limitations of CD in the complexity of the allowed defects and behaviors while reducing the computational demands compared to OKMC, which tracks all defect positions within a volume at every step.

To reach large simulation volumes one necessarily needs to implement these tools in parallel. In a deterministic model, such implementation is straightforward as each parallel domain will explicitly march in time with the same time step. However, in the case of kinetic Monte Carlo approaches the problem is far more complex because the time increment at each step is a weighted random choice based on the allowed reactions in the system. Therefore, domains running in parallel may become asynchronous, making reactions such as defect diffusion from one to the other difficult to properly carry out. In order to address this problem, a synchronous parallel kinetic Monte Carlo algorithm has been developed by Martínez et al. [36] in which a single time increment is chosen for all processors and the possibility of choosing null events is added to the KMC algorithm. Although this algorithm has been implemented for OKMC and lattice kinetic Monte Carlo (LKMC) systems [36–38], it has not been applied to the SRSCD methodology. The objective of the present study is to introduce a parallel implementation of the SRSCD scheme and to assess the numerical performance of the approach as well as its ability to simulate, for the first time, irradiation in polycrystals.

The manuscript is organized as follows. In Sections 2 and 3, SRSCD and the parallel kinetic Monte Carlo algorithm are described. In Section 4, the parallel scaling of the SRSCD code is investigated in an idealized system with only single vacancies and interstitials as well as in characteristic simulations of Frenkel pair and cascade implantation in α -Fe. Here, an explicit cascade implantation scheme is developed to address problems of poor weak scaling due to fast-moving defects that occur in simulations of cascade damage accumulation. In Section 5, synchronous parallel SRSCD is shown to enable simulations of damage accumulation in nano-grained polycrystalline α -Fe and the effect of grain size on damage accumulation is studied as a model case. The choice of kinetic Monte Carlo domain size is discussed in Section 6. In Section 7 conclusions are discussed including the limitations of the present methodology.

2. Description of SRSCD

2.1. Formulation of stochastic cluster dynamics (SCD)

Stochastic cluster dynamics (SCD) was developed by Marian and Bulatov as an intermediate method between OKMC and CD methods for simulating defect accumulation in irradiated metals [39]. It is based on the work of Gillespie [40], in which an algorithm is presented to simulate the evolution of chemical rate equations in a kinetic Monte Carlo setting. To describe this algorithm in the context of radiation damage, we consider a system containing N defect species S_i , $i \in \{1, \dots, N\}$, with the concentration c_i (in units of atomic fraction, or defects per lattice atom) of species S_i . In typical deterministic rate theory, the first-order rate equation for c_i is in the form:

$$\frac{dc_i}{dt} = \sum_{\mu=1}^{M_i} \pm a_{\mu}^i(c_1, \dots, c_N) \quad (1)$$

where a_{μ}^i is the reaction rate of reaction R_{μ}^i that impacts the population of species S_i (in units of reactions per lattice atom per second) and the $\pm$ symbol indicates whether the reaction adds or subtracts defects of this type from the system. In general, a_{μ}^i is dependent on the concentration of defects of various types in the system. There are a total of M_i reactions involving species i in this example. To treat a system of coupled rate equations $\frac{dc_i}{dt}$ in a stochastic setting, the concentrations of each species $c_i(t)$ are converted into populations $X_i(t)$ inside a finite volume V by multiplying by the ratio of the chosen volume element size to the atomic volume Ω :

$$\frac{dX_i}{dt} = \sum_{\mu=1}^{M_i} a_{\mu}^i(X_1, \dots, X_N) \cdot \left(\frac{V}{\Omega}\right) = \sum_{\mu=1}^{M_i} A_{\mu}^i(X_1, \dots, X_N) \quad (2)$$

Here, reaction rates A_{μ}^i are in units of reactions per second and represent the rate for the specified reaction R_{μ}^i to occur inside the chosen volume element V .

At this point, instead of measuring the evolution of the various defect populations X_i through a deterministic solution of the rate equations presented in equation (2), we can instead choose a stochastic approach to population evolution. Given a system at time t with volume V and defect populations $\{X_1, \dots, X_N\}$, we can create a list of reaction rates A_{μ} for all allowed reactions R_{μ} in the system, $\mu \in \{1, \dots, M\}$. If the populations X_i are known, the rates A_{μ} are analytically calculable via the terms of equation (2). We now restrict ourselves to considering integer populations of each defect species S_i , reducing the number of nonzero reaction rates that must be calculated.

Given such a system, the classical kinetic Monte Carlo algorithm can be applied to choose a time step τ and reaction μ to carry out in the system by choosing two random numbers r_1 and $r_2 \in (0, 1)$ and applying the formula [40]:

$$\tau = \frac{1}{A} \ln\left(\frac{1}{r_1}\right), \quad \sum_{\nu=1}^{\mu-1} A_{\nu} < r_2 A \leq \sum_{\nu=1}^{\mu} A_{\nu}. \quad (3)$$

Once reaction R_μ is chosen, the populations of defects X_i in the system are subsequently updated according to the reaction chosen and the changed reaction rates A_μ are recalculated. Therefore, all defects are accounted for at each step and any imbalances in defect populations are due to preferential absorption of defects at sinks, for example due to the greater diffusivity of self-interstitials compared to vacancies. This formalism assumes that when reactions occur, they are carried out instantaneously so that two reactions cannot occur simultaneously. It has been proven [41, 42] that the stochastic description of events presented here is equivalent to the deterministic rate theory formulation given in equation (1). Note that for a given reaction such as two single vacancies joining to form a cluster, $V+V \rightarrow 2V$, the reverse reaction $2V \rightarrow V+V$ may or may not be allowed, and if it is allowed it may have a different reaction rate than the forward reaction. This reflects the different mechanisms associated with defect clustering and dissociation, and can lead to a system with no equilibrium state in which clusters grow indefinitely.

2.2. Spatially resolved stochastic cluster dynamics

The development of SCD allows simulations of radiation damage accumulation to model more complex systems and defect behaviors than CD because reaction rates A_μ are only calculated for defects present in the volume element V [39]. However, both methods remain spatially homogeneous and therefore cannot resolve complex microstructures or heterogeneous effects such as displacement cascades. In CD-based methods, these barriers are overcome by solving the deterministic rate equations presented in equation (1) at spatial nodes, allowing diffusion of defects between nodes and allowing simulation of defect evolution in thin films and nano-laminate microstructures [18, 23, 24]. Displacement cascade implantation has similarly been included in CD-based methods using cascades that are pre-annealed using spatially resolved techniques such as OKMC and EKMC [21, 22].

Spatially resolved stochastic cluster dynamics (SRSCD) has been developed in order to use SCD to simulate non-homogeneous microstructures and spatially correlated reactions between defects [35]. Here, instead of carrying out the kinetic Monte Carlo algorithm described in section 2.1 inside a single volume element, this methodology is applied to several smaller volume elements each with volume V . Defects are only allowed to interact within the same volume element, and diffusion of defects between volume elements is given by a finite-volume solution of Fick's law. The reaction rate A_{jk} for diffusion of a defect between volume elements j and k is given by:

$$A_{jk} = \frac{D\Omega_{jk}(N_j - N_k)}{VL_{jk}} \quad (4)$$

where D is the diffusion rate of that defect type, Ω_{jk} is the shared area of the two (rectangular) volume elements, N_j and N_k are the numbers of defects of that type in volume

elements j and k , respectively, and L_{jk} is the distance between the centers of the volume elements.

Cascade damage simulations in SRSCD require volume element sizes less than or equal to the size of the cascade in order to achieve correct initial densities of the various defects in the cascade [43]. However, the use of small volume elements decreases the computational efficiency of the simulation by increasing the number of diffusion events between volume elements. To increase the computational efficiency of cascade implantation simulations in bulk iron, a pseudo-adaptive meshing scheme has been implemented in which the bulk material is treated with larger volume elements for the majority of the simulation and separate refined meshes are generated when cascades are implanted [44]. Despite the computational gains made by this method, the total simulation volume and doses in simulations carried out to date are still limited to cubes with side length 160 – 300 nm and 10^{-1} dpa, respectively [44]. Non-adaptive meshing simulations of defect evolution in thin films have exhibited similar limitations on volume and total dose [43]. Simulations of damage accumulation in explicitly represented polycrystalline nano-grained microstructures require large volumes as well as fine spatial resolution. The large number of volume elements required by such simulations prompts the use of a parallel implementation of SRSCD.

3. Synchronous parallel kinetic Monte Carlo algorithm

The synchronous parallel kinetic Monte Carlo method of Martínez et al. [36] was developed to avoid asynchronous time evolution in multiple domains which can occur in parallel kinetic Monte Carlo methods. In this method, the total volume is divided between D domains. Therefore, X_i^d and A_μ^d represent the population of species X_i in domain d and the reaction rate of reaction R_μ^d in domain d , respectively, with $\mu \in \{1, \dots, M\}$, $i \in \{1, \dots, N\}$, and $d \in \{1, \dots, D\}$. Similarly to the kinetic Monte Carlo algorithm presented in Section 2.1, the total reaction rate A^d is computed in domain d :

$$A^d = \sum_{\mu=1}^M A_\mu^d \quad (5)$$

The maximum total reaction rate is then found among all domains:

$$A_{\max} = \max_d \{A^d\} \quad (6)$$

In order to implement out the synchronous parallel kinetic Monte Carlo algorithm, a null event R_0^d is then added to the list of reactions in each domain. This reaction represents the possibility of no reaction being carried out in a given domain during a global time step. The reaction rate of the null event A_0^d is defined in each domain as the

1 difference between the total reaction rate in domain d and
 2 the maximum reaction rate among all domains:

$$3 \quad A_0^d = A_{\max} - A^d \quad (7)$$

5 Therefore the rate of R_0^d is nonzero in all domains with
 6 total reaction rate $A^d < A_{\max}$. By adding null events to
 7 each domain in this way, each domain has a new total
 8 reaction rate A^{*d} with value given by:

$$10 \quad A^{*d} = \sum_{\mu=0}^M A_{\mu}^d = A_{\max} \quad (8)$$

13 where here $\mu = 0$ is included in the sum in order to in-
 14 clude the null event in the list of reactions. Since the total
 15 reaction rate A^{*d} is the same in each domain, the kinetic
 16 Monte Carlo algorithm can choose a single time step for
 17 the global system while choosing an independent reaction
 18 in each system. This is done by choosing independent ran-
 19 dom numbers r_1 and $r_2^d \in (0, 1)$. These are used to choose
 20 a single time step τ for all domains and a separate reaction
 21 R_{μ}^d in each domain d :

$$24 \quad \tau = \frac{1}{A_{\max}} \log \left(\frac{1}{r_1} \right), \quad \sum_{\nu=0}^{\mu-1} A_{\nu}^d < r_2^d A_{\max} \leq \sum_{\nu=0}^{\mu} A_{\nu}^d \quad (9)$$

27 The main difference between equations (3) and (9) is
 28 the fact that τ is treated as a global time step and that in
 29 each domain the possibility of choosing a null event now
 30 exists. The addition of null events to the kinetic Monte
 31 Carlo algorithm in this way has been proven not to alter
 32 the overall time evolution of the system [36], while allowing
 33 the choice of a single time step and multiple reactions per
 34 step in the entire system.

36 The term ‘domain’ used above describes a sub-volume
 37 of the simulation volume within which one choice of re-
 38 action is made during each step. The smallest size this
 39 domain can take in SRSCD is a single volume element,
 40 such that a separate reaction is chosen inside each volume
 41 element during each step, and the largest size that this
 42 domain can take is all of the volume elements ascribed
 43 to a single processor. Choosing smaller domains increases
 44 the number of reactions at each step but makes null events
 45 more likely, and the optimal choice of domain size can vary
 46 depending on the simulation. All simulations in Section 4
 47 are carried out with one domain per processor. The im-
 48 pact of making this choice on scaling is dependent on the
 49 simulation parameters and is discussed in Section 6.

51 The parallel kinetic Monte Carlo algorithm presented
 52 here can cause boundary errors in OKMC methods when
 53 two reactions are chosen during the same step in neigh-
 54 boring domains such that two defects move within a pre-
 55 defined interaction distance of each other [36]. In SRSCD,
 56 boundary errors are caused when defects in neighboring
 57 volume elements in different domains exchange positions
 58 during the same step. In doing so, the defects are unable
 59 to interact because they are never located in the same vol-
 60 ume element. However, these errors are extremely unlikely

except for unusual cases, for example in a system with
 two volume elements and two domains only. Therefore,
 these errors are considered negligible in the simulations
 performed in this work. Indeed, no difference between de-
 fect populations were found in serial and parallel imple-
 mentations of SRSCD, even in cases with very small total
 volume per domain.

4. Scaling of synchronous parallel SRSCD

If a kinetic Monte Carlo simulation with total reaction
 rate A is divided among D domains using the synchronous
 parallel algorithm presented in Section 3, the best-case sce-
 nario for scaling (assuming zero probability of null events
 in any domain) gives $A_{\max} = \frac{A}{D}$. Therefore, the best-case
 increase in time step is $\tau_{\text{parallel}} = D \cdot \tau_{\text{serial}}$. Any deviation
 from this is due to uneven distribution of total reaction
 rates A^d among the domains, leading to non-zero proba-
 bilities of null events being chosen inside some domains at
 each step. When simulating very stiff problems, e.g. prob-
 lems in which reaction rates A_{μ} vary by orders of mag-
 nitude for different reactions, total reaction rates A^d are
 likely to vary as well due to heterogeneous distributions of
 fast-reacting defects in the system. Because of the large
 number of defect types modeled in simulations of damage
 accumulation in metals such as α -Fe, as well as the fact
 that their diffusivities can vary by many orders of magni-
 tude, this problem occurs frequently in such simulations.
 Therefore, when investigating the performance of SRSCD,
 it is key to include both ideal cases that demonstrate the
 capabilities of the architecture itself as well as cases that
 demonstrate parallel performance in typical simulations of
 damage accumulation.

The following scaling simulations were performed on
 Sandia’s High Performance Computing Sky Bridge cluster
 from Cray, Inc. The cluster has 16 processors and 64GB
 of RAM per node with 2.6 GHz Intel Sandy Bridge:2S:8C
 processors. Processor communication is handled through
 a fully connected QDR InfiniBand interconnect.

4.1. Damage accumulation in characteristic simulations of irradiated α -Fe: allowed defects and reactions

In Sections 4.2 and 4.3, the parallel performance of
 SRSCD is investigated for two cases: an idealized case in
 which clustering is not allowed and the only reactions are
 Frenkel pair formation, diffusion, and recombination, and
 a case representing realistic defect behaviors. In the lat-
 ter case, vacancy and self-interstitial clusters are allowed
 to form, and displacement cascades are introduced in Sec-
 tion 4.3. In this section, a brief description of the allowed
 defects and reactions included in these SRSCD simulations
 is presented.

The defect types included in the simulations in this
 work are the following:

- 1 1. *Single vacancies and self-interstitials*: These defects
2 are treated as spherical and migrate in three dimen-
3 sions with diffusion rates given by a standard Arrhe-
4 nius law, $D = D_0 e^{-\frac{E_m}{k_b T}}$. They are allowed to form
5 clusters and annihilate with defects of the opposite
6 type.
- 7 2. *Small mobile vacancy and self-interstitial clusters*:
8 Vacancy and self-interstitial clusters size 2-4 are treated
9 as spherical clusters that diffuse in 3D. Pairs of clus-
10 ters are allowed to combine, annihilate, and undergo
11 thermal emission of point defects.
- 12 3. *Larger immobile vacancy clusters*: Vacancy clusters
13 size 5 or greater are assumed to be immobile, spher-
14 ical defects. These defects can interact with mobile
15 defects via clustering or annihilation, and can emit
16 single vacancies.
- 17 4. *Interstitial dislocation loops*: Self-interstitial clusters
18 size 5 or greater are assumed to form circular dis-
19 location loops [45]. These dislocation loops can be
20 mobile and diffuse in 1D. Mobile dislocation loops
21 can become immobile if two dislocation loops meet
22 and form a junction [46] or if they encounter a trap.

23 Migration and binding energies E_m and E_b as well as
24 diffusion prefactors D_0 required to compute reaction rates
25 are taken from atomistic simulations in α -Fe [34, 45] and
26 are shown in Table 1. The reaction rates A_μ for the various
27 allowed reactions between defects depend on the diffusivity
28 of the various defects, their geometry, and the character of
29 their diffusion. Reaction rates for all allowed reactions in
30 this work are found in [44].

31 In sections 4.2 and 4.3, scaling is investigated using the
32 case of one KMC domain per processor only. Simulations
33 with one KMC domain per volume element are less effi-
34 cient when defect clustering is allowed due to the large
35 number of null events chosen when the parameters pre-
36 sented above are applied. Further discussion of the scaling
37 of simulations with one KMC domain per volume element
38 can be found in Section 6.

39 4.2. Weak scaling during Frenkel pair implantation

40 In this work, weak scaling is used as the principal met-
41 ric for parallel performance of SRSCD. In weak scaling, the
42 change in computation time is measured as the simulation
43 volume per processor is held constant while the total sim-
44 ulation volume is varied. This metric indicates how effec-
45 tively the parallel algorithm enables simulations of larger
46 volumes than are feasible using the serial algorithm. The
47 metric used in this work for parallel performance associ-
48 ated with weak scaling, $\eta_w(p)$, is defined as:

$$49 \eta_w(p) = \frac{t_s}{t_p}, \quad (10)$$

50 where p is the number of processors, t_s is the computation
51 time for a standard serial simulation, and t_p is the compu-
52 tation time for a parallel simulation with p processors and
53

total volume p times the volume of the serial simulation.
When $\eta_w = 1$, the system exhibits perfect weak scaling,
indicating that the addition of simulation volume and pro-
cessors does not increase computation time. If $\eta_w < 1/p$,
the parallel simulation is slower than the equivalent simu-
lation in serial, indicating that the parallel algorithm does
not allow efficient simulation of large volumes.

In this section, the weak scaling performance of syn-
chronous parallel SRSCD is characterized for two cate-
gories of simulations: (1) simulations in which the only
allowed reactions in the system are Frenkel pair implanta-
tion, vacancy and self-interstitial diffusion, and vacancy-
interstitial recombination and (2) simulations in which de-
fect clustering is also allowed, as described in Section 4.1.
All simulations include volume elements with 20 nm length
and periodic boundary conditions. Weak scaling simula-
tions are performed with 1000 volume elements per pro-
cessor. In order to compare scaling in simulations with
different allowed reactions, simulations with and without
clustering are started from the same initial defect state.
The initial defect state chosen corresponds to the steady-
state population of single vacancies and interstitials found
when clustering is not allowed for a given dose rate (in
dpa·s⁻¹). Steady-state populations of defects are reached
before 10⁻⁴ dpa for all dose rates studied here, and there-
fore all simulations are started from an initial dose of 10⁻⁴
dpa. The final dose in all weak scaling simulations is 10⁻³
dpa. Simulations with no clustering remain at this steady
state population while simulations in which clustering is al-
low deviate from steady-state as clusters form. Although
the latter case does not represent a true physical problem,
it was chosen in order to better compare scaling between
problems with different allowed reactions. The steady-
state population of vacancies and interstitials is shown in
Figure 1(c) for the various dose rates used in this study.

Weak scaling results are shown in Figures 1(a) and 1(b)
for the cases of no clustering and clustering, respectively.
Dose rates from 1 – 10⁻⁶ and 1 – 10⁻⁸ dpa·s⁻¹ are shown
in Figures 1(a) and 1(b), respectively. The lowest dose
rates are not shown in Figure 1(a) due to computational
limitations. In most cases, scaling remains above the 1/ p
limit, indicating that significant computational gains are
achieved by using the synchronous parallel SRSCD algo-
rithm. When no clustering is allowed (1(a)), η_w is indepen-
dent of dose rate for dose rates greater than 10⁻⁴ dpa·s⁻¹
indicating that very few null events were chosen in these
simulations and the distribution of reaction rates is fairly
homogeneous between KMC domains. Therefore, at these
dose rates all decreases in η_w with increasing numbers of
processors are due to increased communication time be-
tween processors rather than the occurrence of null events.
Below 10⁻⁴ dpa·s⁻¹, the low concentration of defects con-
tributes to spatial inhomogeneities in defect distributions,
making null events more likely and decreasing weak scal-
ing.

By contrast, when clustering is allowed (1(b)), the re-
lationship between parallel scaling and dose rate is more

(a) Weak scaling: no clustering

(b) Weak scaling: clustering

(c) Steady-state concentrations (no clustering)

(d) Fraction of null events chosen

Figure 1: Parallel performance associated with weak scaling using several different dose rates corresponding to different initial concentrations of Frenkel pairs. Results are compared for the cases of no defect clustering (1(a)) and clustering (1(b)). The initial concentration of Frenkel pairs in each simulation is given by the steady-state concentration with no clustering allowed for various dose rates (1(c)). The fraction of null events chosen at each dose rate for the case of 64 processors is shown in 1(d). Results are not shown for dose rates below 10^{-6} dpa·s $^{-1}$ for simulations in which clustering is not allowed due to computational limitations.

Migration parameters	$D = D_0 e^{-\frac{E_m}{k_b T}}$	
Single vacancy	$E_m = 0.67$ eV, $D_0 = 8.2 \times 10^{11}$ nm ² /s	[34]
2- <i>v</i> cluster	$E_m = 0.62$ eV, $D_0 = 8.2 \times 10^{11}$ nm ² /s	[34]
3- <i>v</i> cluster	$E_m = 0.35$ eV, $D_0 = 8.2 \times 10^{11}$ nm ² /s	[34]
4- <i>v</i> cluster	$E_m = 0.48$ eV, $D_0 = 8.2 \times 10^{11}$ nm ² /s	[34]
Larger <i>v</i> clusters	(immobile)	
Single interstitial	$E_m = 0.34$ eV, $D_0 = 8.2 \times 10^{11}$ nm ² /s	[34]
2- <i>i</i> cluster	$E_m = 0.42$ eV, $D_0 = 8.2 \times 10^{11}$ nm ² /s	[34]
3- <i>i</i> cluster	$E_m = 0.43$ eV, $D_0 = 8.2 \times 10^{11}$ nm ² /s	[34]
4- <i>i</i> cluster	$E_m = 0.43$ eV, $D_0 = 8.2 \times 10^{11}$ nm ² /s	[34]
<i>n</i> - <i>i</i> cluster (1D)	$E_m = .06 + .11(n^{-1.6})$, $D_0 = (3.5 \times 10^{10} + 1.7 \times 10^{11} n^{-1.7})$ nm ² /s	[45]
Binding energies		
V ₂ → V+V	$E_b = 0.3$ eV	[34]
V ₃ → V ₂ +V	$E_b = 0.37$ eV	[34]
V ₄ → V ₃ +V	$E_b = 0.62$ eV	[34]
Vacancy clusters ($n > 4$)	$E_b^v(n) = 2.07 - 3.01(n^{\frac{2}{3}} - (n-1)^{\frac{2}{3}})$	[34]
I ₂ → I+I	$E_b = 0.8$ eV	[34]
I ₃ → I ₂ +I	$E_b = 0.92$ eV	[34]
Interstitial clusters ($n > 3$)	$E_b^i(n) = 3.77 - 5.05(n^{\frac{2}{3}} - (n-1)^{\frac{2}{3}})$	[34]

Table 1: Migration and binding parameters for vacancies and self-interstitials in bulk α -Fe.

complex. Weak scaling η_w decreases with decreasing dose rate from $1 - 10^{-3}$ dpa·s⁻¹ and then increases slightly between 10^{-3} and 10^{-4} dpa·s⁻¹. The initial decrease in η_w is due to the increased heterogeneity of reaction rates between processors due to the lower number of defects and the presence of small numbers of fast-diffusing defects like dislocation loops. Below 10^{-3} dpa·s⁻¹, fast-moving loops become scarce again and the population of immobile loops grows, leading to improved scaling as reaction rates become more homogeneous between processors once again. As the dose rate continues to decrease from $10^{-4} - 10^{-8}$ dpa·s⁻¹, η_w once again decreases. At these low dose rates, the population of interstitial loops decreases and the mobility of vacancies becomes significant relative to the irradiation time, leading to poor scaling as diffusion of vacancies and vacancy clusters between volume elements dominates the reactions chosen. At the lowest dose rate studied here, 10^{-8} dpa·s⁻¹, η_w is below the $1/p$ scaling limit, indicating that computational gains from the parallel algorithm are not attained for this case.

To investigate the relationship between weak scaling and null events in more detail, Figure 1(d) shows the fraction of kinetic Monte Carlo steps that result in the choice of a null event for a given processor in simulations using 64 processors. These results are in agreement with the trends discussed above, including (1) the relatively constant scaling of simulations when no clustering is allowed above 10^{-4} dpa·s⁻¹, followed by a decrease in weak scaling at lower dose rates; and (2) the non-monotonic increase in null events as dose rate decreases from $1 - 10^{-8}$ dpa·s⁻¹ when clustering is allowed. Note that even when very few null events are chosen, such as in the case of no clustering at 1 dpa·s⁻¹, Figure 1(a) shows that weak scaling η_w is sig-

nificantly less than 1, indicating that the majority of the increase in computation time in this case is due to communication between processors rather than null events.

In summary, weak scaling simulations show that for the simplified system studied here, fast computation of defect evolution is made possible by the synchronous parallel algorithm even for very large total simulation volumes except for simulations at low dose rates that include many defect types and reactions. The scaling of any given SRSCD simulation is dependent on the heterogeneity of reaction rates among processors as well as factors such as the amount of computation time expended choosing reactions and the amount of time spent communicating information between processors.

4.3. Scaling during displacement cascade implantation

Simulations of cascade damage in SRSCD use a pseudo-adaptive meshing scheme [44] in which a fine mesh is introduced when cascade implantation events are chosen in the kinetic Monte Carlo algorithm, allowing spatially correlated reactions between defects in the cascade to occur. When all such spatially correlated reactions have occurred, the fine mesh is discarded and the simulation continues with a coarser mesh. This scheme was developed in order to enable simulations of larger volumes and greater radiation doses than feasible in past SRSCD simulations.

This methodology causes simulations of cascade damage to exhibit poor parallel performance, as the total reaction rate when a cascade is present can be as much as seven orders of magnitude greater than after the spatially correlated reactions in the cascade have been carried out. This is shown in Figure 2(a), which depicts the total reaction rates inside each processor in a two-processor simula-

tion of cascade implantation in α -Fe at room temperature. The total reaction rate inside each processor in the simulation oscillates between 10^9 s^{-1} while a cascade is present and 10^2 s^{-1} otherwise. Due to the Monte Carlo nature of cascade implantation, this oscillation leads to a situation where a cascade is typically present in only one processor at a time, and during that time the other processor chooses null events. Thus weak scaling of this method leads to values of η_w close to the $1/p$ limit, as only a single processor typically chooses a reaction at each step.

In order to mitigate this problem, an explicit cascade implantation scheme has been implemented. In this case, cascade implantation is not treated as a reaction that can be chosen with the Monte Carlo algorithm. Instead, cascades are regularly implanted into all processors simultaneously at regular time intervals. The frequency of cascade implantation is chosen such that the overall dose rate remains the same. Thus, spatially correlated reactions between defects in the cascades can be carried out simultaneously for several cascades in multiple processors, decreasing the probability that a null event will be chosen. This is shown in Figure 2(b), in which the total reaction rates for each processor inside a two-processor system are once again shown. Since cascades are implanted simultaneously in both processors, the amount of overlap between the reaction rates in each processor increases, allowing for better weak scaling.

To demonstrate the computational gain provided by this explicit cascade implantation scheme, the speedup gained by using the explicit method has been calculated for both computation time and number of kinetic Monte Carlo steps:

$$\text{Speedup} = \frac{t_{\text{MC}} - t_{\text{exp}}}{t_{\text{MC}}}, \quad (11)$$

where t_{exp} and t_{MC} are the computation time in explicit or Monte Carlo cascade implantation mode, respectively. These simulations are performed with 8 coarse volume elements per processor due to computational constraints with length 80 nm and periodic boundary conditions, and displacement cascades are implanted up to 10^{-3} dpa at $7 \times 10^{-7} \text{ dpa}\cdot\text{s}^{-1}$ and room temperature. Figure 3(a) shows the speedup calculated in this way as a function of the size of the simulated system (keeping volume per processor constant as in weak scaling simulations).

Implementation of this method could produce damage accumulation results that are different from the case of Monte Carlo cascade implantation if defects implanted simultaneously into neighboring processors undergo correlated reactions that are caused by the explicit implantation scheme and are therefore non-physical. However, several factors mitigate this risk. First, each processor typically contains several volume elements, and only a single displacement cascade is inserted into each processor at a time. This makes the probability that two cascades will be inserted into neighboring volume elements low, decreasing the chances that correlated reactions will occur between

defects in neighboring elements. Next, the use of adaptive meshing [44] prevents two displacement cascades from occupying adjacent volume elements. In order to demonstrate the reliability of this method, defect populations and average sizes in the simulations described above are compared for the case of 64 processors using the Monte Carlo and explicit methods Figure 3(b). Both defect populations and average size show good agreement over two orders of magnitude of displacement damage.

Although the explicit implantation scheme does provide a decrease in computation time and number of steps as the system becomes larger, the improvement is not as strong as seen in Sections 4.2. This is due to the fact that fast-moving interstitial clusters are more likely to form due to the spatial correlation between defects in the cascades, and these defects cause the system to exhibit poor scaling due to their extremely high diffusivities even when the explicit cascade implantation scheme is used. A different choice of material parameters in which dislocation loops are assumed to be immobile, such as in the work of Fu et al. [34], or a first-passage Monte Carlo algorithm [47] could be used to remove these fast-diffusing clusters and would significantly improve the scalability of the explicit implantation method.

5. Application: Frenkel pair implantation in polycrystalline α -Fe

The synchronous parallel SRSCD algorithm presented here allows simulation of explicitly represented polycrystalline microstructures with grain sizes on the order of tens of nanometers. To demonstrate this capability, Frenkel pair implantation in nano-grained polycrystalline α -Fe is simulated using two representative polycrystalline structures with 15 grains and 30 grains generated using Sandia National Laboratories' Material Point Method (MPM) code [48]. This polycrystalline microstructure is obtained by means of a 3D Potts model of grain evolution [49]. The simulated grain evolution used to generate the representative microstructure employs a Monte Carlo Potts model using a non-physical kinetic Monte Carlo temperature with isotropic grain boundary energy and mobility. The dimensions of each representative volume are $400 \times 100 \times 100 \text{ nm}$, and spatial resolution is achieved using 5 nm cubic volume elements in SRSCD. The polycrystalline grain structures are shown in Figures 4(a) and 4(b).

Radiation damage in these polycrystals is simulated with uniform Frenkel pair implantation at a dose rate of $7 \times 10^{-7} \text{ dpa}\cdot\text{s}^{-1}$. As a first approximation, grain boundaries are treated as perfect sinks for defects and more complex behaviors such as defect diffusion along grain boundaries and grain boundaries acting as imperfect sinks are not accounted for. Damage accumulation is simulated up to a total dose of 10^{-1} dpa , using 32 processors per simulation. Simulation time for irradiation to 10^{-1} dpa is approximately four hours under these conditions. By contrast, carrying out these simulations in serial using SRSCD

Figure 2: Total reaction rates inside each processor during the first 3000 Monte Carlo steps in a two-processor cascade implantation simulation, using the Monte Carlo algorithm for cascade implantation (top) or using an explicit implantation scheme at the same dose rate (bottom). Increasing overlap of reaction rates allows for better scaling of parallel SRSCD simulations.

Figure 3: (3(a)) Speedup in computation time using the explicit cascade implantation method vs Monte Carlo cascade implantation for simulations of cascade damage in α -Fe. (3(b)) Comparison of vacancy (V) and self-interstitial (SIA) type defect populations and average sizes in a simulation of cascade damage to 10^{-3} dpa using 64 processors with Monte Carlo (MC) and explicit cascade implantation schemes.

(a) 15 Grains

(b) 30 Grains

Figure 4: Grain structure and vacancy clusters (spheres) in simulated polycrystals implanted with Frenkel pairs to 10^{-2} dpa. Fewer total defects are found in the case of 30 grains due to the smaller average grain size and corresponding higher grain boundary sink strength. The total simulation volume is $400 \times 100 \times 100$ nm.

requires approximately twice as much computation time. Therefore, while the parallel SRSCD algorithm does provide significant computational benefits compared to serial SRSCD for this simulation, these gains are limited by poor scaling at low dose rates, as discussed in Section 4.2.

The final population of vacancy clusters inside the grains at 10^{-1} dpa is shown for the 15-grain and 30-grain representative volume elements in Figures 4(a) and 4(b). Note that in these figures, spheres represent the total population of vacancy clusters within each volume element rather than individual defects, and the sphere size has been magnified in order to provide a better guide to the eye. Individual vacancy cluster sizes remain well below the volume element size in these simulations. Self-interstitial populations are not shown because their high diffusivity causes the majority of self-interstitials to be absorbed by grain boundaries, which are treated as perfect sinks. The difference in damage content in terms of vacancy cluster concentration and average cluster size, in units of vacancies per cluster, between the two grain sizes is shown as a function of dose in Figure 5. The effect of increasing the number of grains (and therefore decreasing the average grain size) is shown to be decreased vacancy cluster concentration and average size, in agreement with experimental trends [10–12].

Figure 5: Average vacancy cluster concentration (in m^{-3}) and average number of vacancies per cluster as a function of dose for a polycrystals implanted with Frenkel pairs to 10^{-1} dpa. Two polycrystals were simulated, with 15 and 30 grains, respectively. The total simulation volume is $400 \times 100 \times 100$ nm.

6. Discussion: Limiting cases of kinetic Monte Carlo domains

In parallel SRSCD simulations, two competing factors can influence computation time: the amount of communication carried out between processors and the time spent searching for which reaction to choose in each KMC domain. Both of these factors are influenced by choosing the size of the KMC domain. Here we discuss the limiting cases of KMC domain sizes: one KMC domain per processor (KMC_P, the choice made for all simulations in Section 4) and one KMC domain per volume element (KMC_E). These two choices of KMC domain will be investigated under two limiting cases: (1) homogeneously distributed reaction rates, and (2) strongly heterogeneously distributed reaction rates, in which only one fast-moving defect is active in the entire simulation at a time. For simplicity, we assume an ideal case for this analysis: p processors with m volume elements per processor and n possible reactions per volume element. In case (1), we assume that n is equal for all volume elements and therefore the probability of null events is zero. We also assume that the same amount of time T_c is spent in communication between processors each step.

We first investigate case (1). When using the KMC_P domain, the implementation of synchronous parallel SRSCD used in this study uses an optimized search algorithm for choosing reactions that requires $\mathcal{O}(m)$ computation time to choose a volume element and $\mathcal{O}(n)$ computation time to choose a reaction. By contrast, when using the KMC_E domain, the search time for a reaction in that domain is $\mathcal{O}(n)$ but m such searches are made.

In addition to differences in computation required to search for reactions, the average timestep taken by each method is different. In case (1), the timestep chosen when using the KMC_P domain scales with $\frac{1}{m}$ while the timestep when using the KMC_E domain is independent of the size of the domain. Therefore, the computation times $T(\tau)_P$ and $T(\tau)_E$ required in order to allow the same amount of simulated time τ to pass in the cases of KMC_P and KMC_E, respectively, are:

$$\begin{aligned} T(\tau)_P &= m [T_c + \mathcal{O}(m + n)] \\ T(\tau)_E &= T_c + \mathcal{O}(mn) \end{aligned} \quad (12)$$

where the first equation has been multiplied by m to indicate that m steps are required for KMC_P to reach the same simulated time τ reached in by KMC_E in a single step. In the limiting case of negligible communication time and $n \gg m$, equation (12) gives the same computation time using either method, while in the limit that $n \ll T_c$ or $n \ll m$, KMC_E is faster. In a typical simulation, $n \ll m$, and therefore KMC_E is the preferred method in case (1).

This advantage disappears in the case of very stiff problems, or problems in which a few fast-moving defects are heterogeneously distributed throughout the volume (case (2)). This is frequently realized in simulations of defect

accumulation in α -Fe, in which a fast-moving SIA loop with diffusivity several orders of magnitude higher than single vacancies and interstitials can form. In this case, for KMC_P simulations, the reaction chosen at each step will most likely include the fast-moving defect. By contrast, in KMC_E simulations, the volume element containing the fast-moving defect will choose a reaction involving that defect and all other volume elements will choose null events. Therefore, the average timestep of the two methods is the same as it is dominated by the reaction rates associated with the fast-moving defect. As the probability of the system including such a fast-moving defect is proportional to m , the timestep of each method scales with $\frac{1}{m}$ in this case. Therefore, assuming equal communication time per step, the total computation times $T(\tau)_P$ and $T(\tau)_E$ required in order to allow simulated time τ to pass are:

$$\begin{aligned} T(\tau)_P &= m [T_c + \mathcal{O}(m + n)] \\ T(\tau)_E &= m [T_c + \mathcal{O}(mn)] \end{aligned} \quad (13)$$

where both equations have been multiplied by m due to the $\frac{1}{m}$ scaling of the timestep. In this case, the choice of KMC_P is superior.

To demonstrate the different scaling situations described in equations (12) and (13), the impact of the choice of KMC domain on computation is compared using weak scaling simulations of Frenkel implantation and clustering shown in Figure 1(b). Only two implantation rates are chosen for the sake of comparison, $1 \text{ dpa}\cdot\text{s}^{-1}$ and $10^{-4} \text{ dpa}\cdot\text{s}^{-1}$. At $1 \text{ dpa}\cdot\text{s}^{-1}$, the high dose rate causes the reaction rates in the system to remain relatively homogeneous, and the system more closely resembles the one described by equation (12). At $10^{-4} \text{ dpa}\cdot\text{s}^{-1}$, reaction rates are more heterogeneously distributed and the system more closely resembles the one described by equation (13). Although the scaling behavior described in equations (12) and (13) represents limiting cases not achieved in these simulations, results still agree with the trends predicted by these equations. This can be seen in Figures 6(a) and 6(b) as the weak scaling is higher and computation time is lower for KMC_E than for KMC_P at $1 \text{ dpa}\cdot\text{s}^{-1}$, while the opposite is true at $10^{-4} \text{ dpa}\cdot\text{s}^{-1}$. Figure 6(c) shows the ratio N_E/N_P of the number of steps in KMC_E simulations (N_E) to the number of steps in KMC_P simulations (N_P). At $1 \text{ dpa}\cdot\text{s}^{-1}$, KMC_E simulations take two orders of magnitude fewer steps, allowing a significant reduction in communication time. By contrast, at $10^{-4} \text{ dpa}\cdot\text{s}^{-1}$, KMC_E take less than an order of magnitude fewer steps than KMC_P simulations while expending significantly more computation time searching for reactions during each step. This is due to the high fraction of null events chosen at $10^{-4} \text{ dpa}\cdot\text{s}^{-1}$ for KMC_E simulations.

Due to the fact that SRSCD simulations typically include significant heterogeneity of reaction rates among volume elements, the scaling described in equation (13) more closely describes the typical performance of the algorithm than the scaling described in equation (12). Therefore,

Figure 6: Weak scaling (6(a)), computation time (6(b)), and number of steps (6(c)) compared between SRSCD simulations with one KMC domain per volume element (KMC_E) and one KMC domain per processor (KMC_P). Results are shown for two dose rates: 1 dpa·s⁻¹ and 10⁻⁴ dpa·s⁻¹, corresponding to cases of homogeneous and heterogeneous reaction rate distributions, respectively. In Figure 6(c), the ratio between the number of steps for KMC_E and KMC_P simulations is shown.

all results presented in Section 4 reflect the choice of the KMC_P domain as this is expected to be the most frequently used scheme for implementing synchronous parallel SRSCD. However, this analysis is dependent on many factors such as allowed reactions, temperature, and dose rate. The choice of optimal KMC domain size may vary depending on these factors, as seen in Figure 6. Overall, a combination of increased communication time, decreased time spent searching for reactions, and increased timestep all influence the total computation time when increasing the number of processors for a given simulation of fixed size. Therefore, the best choice of simulation volume per processor and size of KMC domains depends on the particular simulation parameters. In an ideal case, volume elements are subdivided among KMC domains and processors such that the number of null events is minimized while keeping the number of volume elements per KMC domain low enough to avoid computation-intensive searches for reactions.

7. Conclusions

In this work, implementation of a synchronous parallel kinetic Monte Carlo algorithm in a spatially resolved stochastic cluster dynamics (SRSCD) model for radiation damage accumulation has enabled simulations of damage accumulation in larger volumes and to greater doses than previously achievable by SRSCD simulations. The parallel scaling performance of this methodology has been tested in α -Fe for cases of Frenkel pair implantation at several dose rates and cascade implantation at a dose rate of 7×10^{-7} dpa·s⁻¹. Computational efficiency results obtained from weak scaling simulations are shown to be strongly dependent on the allowed defects and reactions as well as the simulation parameters and size of the kinetic Monte Carlo domain. Overall, significant gains in computation time are achieved when running large simulations with multiple processors rather than a single processor without changing the simulation results.

For the case of cascade implantation, the frequent occurrence of fast-diffusing self-interstitial loops makes the scaling of the synchronous parallel algorithm poor, as orders of magnitude of difference can occur between the total reaction rates A^p in different processors p . This causes the synchronous parallel algorithm to choose null events in all processors except for the one that has the fast-diffusing defect, decreasing the quality of the scaling. An explicit cascade implantation method has been implemented to ameliorate this problem, and it has been showed to decrease computation time moderately for the material properties and implantation conditions considered here.

The synchronous parallel SRSCD algorithm developed in this study allows simulation of radiation damage accumulation in volumes large enough to contain explicitly represented nanocrystalline microstructures. This enables the study of effects such as the impact of grain size and grain boundary sink strength on damage accumulation. These

additional capabilities represent a step towards predictive modeling of the radiation-resistant behavior of new materials in damage environments that are difficult to achieve in experiments, such as fast fission and fusion neutron irradiation.

Acknowledgments

Supported by the Laboratory Directed Research and Development program at Sandia National Laboratories, a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

This work was also funded by the Sandia National Laboratories/Georgia Tech Excellence in Engineering Research Program.

References

References

- [1] R. Klueh, A. Nelson, Ferritic/martensitic steels for next-generation reactors, *Journal of Nuclear Materials* 371 (1) (2007) 37–52.
- [2] G. Odette, G. Lucas, Embrittlement of nuclear reactor pressure vessels, *Jom* 53 (7) (2001) 18–22.
- [3] B. Frost, Radiation damage in fast-breeder and fusion reactors, *Nuclear Science, IEEE Transactions on* 19 (6) (1972) 230–236.
- [4] E. E. Bloom, The challenge of developing structural materials for fusion power systems, *Journal of nuclear materials* 258 (1998) 7–17.
- [5] S. Zinkle, N. Ghoniem, Operating temperature windows for fusion reactor structural materials, *Fusion Engineering and design* 51 (2000) 55–71.
- [6] T. Muroga, M. Gasparotto, S. Zinkle, Overview of materials research for fusion reactors, *Fusion engineering and design* 61 (2002) 13–25.
- [7] S. Ukai, M. Fujiwara, Perspective of ODS alloys application in nuclear environments, *Journal of Nuclear Materials* 307 (2002) 749–757.
- [8] A. Misra, M. Demkowicz, X. Zhang, R. Hoagland, The radiation damage tolerance of ultra-high strength nanolayered composites, *Jom* 59 (9) (2007) 62–65.
- [9] S. J. Zinkle, Advanced materials for fusion technology, *Fusion Engineering and Design* 74 (1) (2005) 31–40.
- [10] C. Sun, K. Yu, J. Lee, Y. Liu, H. Wang, L. Shao, S. Maloy, K. Hartwig, X. Zhang, Enhanced radiation tolerance of ultrafine grained Fe–Cr–Ni alloy, *Journal of Nuclear Materials* 420 (1) (2012) 235–240.
- [11] H. Kurishita, S. Kobayashi, K. Nakai, T. Ogawa, A. Hasegawa, K. Abe, H. Arakawa, S. Matsuo, T. Takida, K. Takebe, et al., Development of ultra-fine grained W–(0.25–0.8) wt% TiC and its superior resistance to neutron and 3MeV He-ion irradiations, *Journal of Nuclear Materials* 377 (1) (2008) 34–40.
- [12] K. Yu, Y. Liu, C. Sun, H. Wang, L. Shao, E. Fu, X. Zhang, Radiation damage in helium ion irradiated nanocrystalline Fe, *Journal of Nuclear Materials* 425 (1) (2012) 140–146.
- [13] M. Demkowicz, R. Hoagland, J. Hirth, Interface structure and radiation damage resistance in Cu–Nb multilayer nanocomposites, *Physical review letters* 100 (13) (2008) 136102.
- [14] X.-M. Bai, A. F. Voter, R. G. Hoagland, M. Nastasi, B. P. Uberuaga, Efficient annealing of radiation damage near grain boundaries via interstitial emission, *Science* 327 (5973) (2010) 1631–1634.

- [15] M. A. Tschopp, K. Solanki, F. Gao, X. Sun, M. A. Khaleel, M. Horstemeyer, Probing grain boundary sink strength at the nanoscale: Energetics and length scales of vacancy and interstitial absorption by grain boundaries in α -Fe, *Physical Review B* 85 (6) (2012) 064108.
- [16] R. Stoller, S. Golubov, C. Domain, C. Becquart, Mean field rate theory and object kinetic Monte Carlo: A comparison of kinetic models, *Journal of Nuclear Materials* 382 (2008) 77–90.
- [17] A. Barbu, C. S. Becquart, J. Bocquet, J. D. Torre, C. Domain, Comparison between three complementary approaches to simulate large fluence irradiation: application to electron irradiation of thin foils, *Philosophical Magazine* 85 (2005) 541–547.
- [18] D. Xu, B. D. Wirth, Modeling spatially dependent kinetics of helium desorption in BCC iron following He ion implantation, *Journal of Nuclear Materials* 403 (2010) 184–190.
- [19] C. J. Ortiz, M. J. Caturla, Simulation of defect evolution in irradiated materials: Role of intracascade clustering and correlated recombination, *Physical Review B* 75 (2007) 184101.
- [20] C. J. Ortiz, M. J. Caturla, C. C. Fu, F. Willaime, He diffusion in irradiated α -Fe: An ab-initio-based rate theory model, *Physical Review B* 75 (2007) 100102.
- [21] T. Jourdan, J.-P. Crocombette, Rate theory cluster dynamics simulations including spatial correlations within displacement cascades, *Physical Review B* 86 (2012) 054113.
- [22] E. Meslin, A. Barbu, L. Boulanger, B. Radiguet, P. Pareige, K. Arakawa, C. Fu, Cluster-dynamics modelling of defects in α -iron under cascade damage conditions, *Journal of Nuclear Materials* 382 (2008) 190–196.
- [23] A. Dunn, M. McPhie, L. Capolungo, E. Martinez, M. Cherkaoui, A rate theory study of helium bubble formation and retention in Cu–Nb nanocomposites, *Journal of Nuclear Materials* 435 (1) (2013) 141–152.
- [24] M. Li, M. Kirk, P. Baldo, D. Xu, B. Wirth, Study of defect evolution by TEM with in situ ion irradiation and coordinated modeling, *Philosophical Magazine* 92 (2012) 2048–2078.
- [25] N. Soneda, S. Ishino, A. Takahashi, K. Dohi, Modeling the microstructural evolution in bcc-Fe during irradiation using kinetic Monte Carlo computer simulation, *Journal of Nuclear Materials* 323 (2003) 169–180.
- [26] J. D. Torre, J.-L. Bocquet, N. V. Doan, E. Adam, A. Barbu, JERK, an event-based Kinetic Monte Carlo model to predict microstructure evolution of materials under irradiation, *Philosophical Magazine* 85 (2005) 549–558.
- [27] C. Domain, C. Becquart, L. Malerba, Simulation of radiation damage in Fe alloys: an object kinetic Monte Carlo approach, *Journal of Nuclear Materials* 335 (2004) 121–145.
- [28] C. Becquart, A. Barbu, J. Bocquet, M. Caturla, C. Domain, C.-C. Fu, S. Golubov, M. Hou, L. Malerba, C. Ortiz, A. Souidi, R. Stoller, Modeling the long-term evolution of the primary damage in ferritic alloys using coarse-grained methods, *Journal of Nuclear Materials* 406 (2010) 39–54.
- [29] V. Jansson, L. Malerba, Simulation of the nanostructure evolution under irradiation in Fe–C alloys, *Journal of Nuclear Materials* 443 (2013) 274–285.
- [30] I. Martin-Bragado, A. Rivera, G. Valles, J. L. Gomez-Selles, M. J. Caturla, MMonCa: An Object Kinetic Monte Carlo simulator for damage irradiation evolution and defect diffusion, *Computer Physics Communications* 184 (12) (2013) 2703–2710.
- [31] J. Dalla Torre, C.-C. Fu, F. Willaime, A. Barbu, J.-L. Bocquet, Resistivity recovery simulations of electron-irradiated iron: Kinetic Monte Carlo versus cluster dynamics, *Journal of nuclear materials* 352 (1) (2006) 42–49.
- [32] T. Jourdan, C. C. Fu, L. Joly, J. Bocquet, M. Caturla, F. Willaime, Direct simulation of resistivity recovery experiments in carbon-doped-iron, *Physica Scripta* 2011 (T145) (2011) 014049.
- [33] S. Golubov, A. Ovcharenko, A. Barashev, B. Singh, Grouping method for the approximate solution of a kinetic equation describing the evolution of point-defect clusters, *Philosophical Magazine A* 81 (3) (2001) 643–658.
- [34] C.-C. Fu, J. D. Torre, F. Willaime, J.-L. Bocquet, A. Barbu,

- Multiscale modelling of defect kinetics in irradiated iron, *Nature Materials* 4 (2005) 68–74.
- [35] A. Y. Dunn, L. Capolungo, E. Martinez, M. Cherkaoui, Spatially resolved stochastic cluster dynamics for radiation damage evolution in nanostructured metals, *Journal of Nuclear Materials* 443 (2013) 128–139.
- [36] E. Martínez, J. Marian, M. H. Kalos, J. M. Perlado, Synchronous parallel kinetic Monte Carlo for continuum diffusion-reaction systems, *Journal of Computational Physics* 227 (8) (2008) 3804–3823.
- [37] E. Martínez, P. R. Monasterio, J. Marian, Billion-atom synchronous parallel kinetic Monte Carlo simulations of critical 3D Ising systems, *Journal of Computational Physics* 230 (4) (2011) 1359–1369.
- [38] I. Martin-Bragado, J. Abujas, P. Galindo, J. Pizarro, Synchronous parallel Kinetic Monte Carlo: Implementation and results for object and lattice approaches, *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 352 (2015) 27–30.
- [39] J. Marian, V. V. Bulatov, Stochastic cluster dynamics method for simulations of multispecies irradiation damage accumulation, *Journal of Nuclear Materials* 415 (1) (2011) 84–95.
- [40] D. T. Gillespie, A general method for numerically simulating the stochastic time evolution of coupled chemical reactions, *Journal of computational physics* 22 (4) (1976) 403–434.
- [41] I. Oppenheim, K. Shuler, G. Weiss, Stochastic and deterministic formulation of chemical rate equations, *The Journal of Chemical Physics* 50 (1) (1969) 460–466.
- [42] T. G. Kurtz, The relationship between stochastic and deterministic models for chemical reactions, *The Journal of Chemical Physics* 57 (7) (1972) 2976–2978.
- [43] A. Dunn, L. Agudo-Merida, I. Martin-Bragado, M. McPhie, M. Cherkaoui, L. Capolungo, A novel method for computing effective diffusivity: Application to helium implanted α -Fe thin films, *Journal of Nuclear Materials* 448 (2014) 195–205.
- [44] A. Dunn, L. Capolungo, Simulating radiation damage accumulation in α -Fe: a spatially resolved stochastic cluster dynamics approach, *Computational Materials Science* (2015) (In press).
- [45] N. Soneda, T. D. de La Rubia, Migration kinetics of the self-interstitial atom and its clusters in bcc Fe, *Philosophical Magazine A* 81 (2001) 331–343.
- [46] D. Terentyev, L. Malerba, P. Klaver, P. Olsson, Formation of stable sessile interstitial complexes in reactions between glissile dislocation loops in bcc Fe, *Journal of Nuclear Materials* 382 (2) (2008) 126–133.
- [47] T. Opplestrup, V. V. Bulatov, G. H. Gilmer, M. H. Kalos, B. Sadigh, First-passage Monte Carlo algorithm: diffusion without all the hops, *Physical review letters* 97 (23) (2006) 230602.
- [48] T. Bartel, R. Dingreville, J. Robbins, M. Lusk, L. Zhang, L. Brewer, iMPALE - Material Point Arbitrary Lagrangian Eulerian code: A time implicit/explicit integration 3-D Lagrangian material point method code, Tech. Rep. SAND2012-0957, Sandia National Laboratories, 2012.
- [49] L. Zhang, R. Dingreville, T. Bartel, M. T. Lusk, Hybrid Monte Carlo simulation of stress-induced texture evolution with inelastic effects, *Metallurgical and Materials Transactions A* 42 (3) (2011) 575–581.

Graphical Abstract

$$\tau = \frac{1}{A_{\max}} \log \left(\frac{1}{r_1} \right)$$

$$\sum_{\nu=0}^{\mu-1} A_{\nu}^d < r_2^d A_{\max} \leq \sum_{\nu=0}^{\mu} A_{\nu}^d$$