Daniel Mason, Andrea Sand, Sergei Dudarev

Atomistic-Object Kinetic Monte Carlo simulations of irradiation damage in tungsten

Atomistic-Object Kinetic Monte Carlo simulations of irradiation damage in tungsten

Daniel Mason, Andrea Sand, Sergei Dudarev

Atomistic-Object Kinetic Monte Carlo simulations of irradiation damage in tungsten

D.R. Mason[*1], A.E. Sand[2], and S.L. Dudarev[1]

[1]CCFE, Culham Centre for Fusion Energy, Abingdon, Oxfordshire OX14 3DB, United Kingdom
[2]Department of Physics, University of Helsinki, P.O. Box 43, FI-00014, Helsinki, Finland

December 13, 2018

Abstract

We describe the development of a new object kinetic Monte Carlo code where the elementary defect objects are off-lattice atomistic configurations. Atomic-level transitions are used to transform and translate objects, to split objects and to merge them together. This gradually constructs a database of atomic configurations- a set of relevant defect objects and their possible events generated on-the-fly. Elastic interactions are handled within objects with empirical potentials at short distances, and between spatially distinct objects using the dipole tensor formalism. The model is shown to evolve mobile interstitial clusters in tungsten faster than an equivalent molecular dynamics simulation, even at elevated temperatures. We apply the model to the evolution of complex defects generated using molecular dynamics simulations of primary radiation damage in tungsten. We show that we can evolve defect structures formed in cascade simulations to experimentally observable timescales of seconds while retaining atomistic detail. We conclude that the first few nanoseconds of simulation would be better performed using molecular dynamics, and that, in the 20keV PKA cascades annealing simulations we consider here, the first few microseconds establish the conditions for long timescale evolution.

1 Introduction

Nuclear materials for Gen IV fission and fusion power stations will be required to operate under unprecedented irradiation fluence. The successful retention

---

[*]daniel.mason@ukaea.uk

of good thermal and mechanical properties depends on the balance between damage generation due to neutron irradiation and restoration through defect recombination. Predicting the evolution of defect clusters is therefore of utmost importance.

Two popular tools for characterizing defects, transmission electron microscopy (TEM)[1, 2], atom-probe tomography (APT) [3] are very successful techniques for identifying (respectively) nanoscale dislocation loops and precipitates, but both struggle to identify the smallest point defects and clusters. This is a significant omission, as state-of-the-art STEM microscopy[4], Field-Ion Microscopy [5], and molecular dynamics (MD)[6, 7] suggest that the majority of damage produced in primary damage cascades should be invisibly small defects.

But MD has also shown us that some defects generated during the heat spike phase of cascades are not small - a power-law size-frequency distribution of defects has been found in simulations of irradiated tungsten and iron[8, 9], and confirmed in experimental observations of self-ion irradiated tungsten[10]. Recently the spatial distribution[11] and count per incident ion[12] of larger defect clusters has also been measured in MD simulations and *in situ* irradiations at cryogenic temperature, giving a characteristic length-scale of a single primary damage cascade for 150keV PKA ions in W of order one nanometre. This proves that large defects are generated sufficiently close together that the interaction between them is a significant driving force on their subsequent evolution[13, 14].

To model the evolution of nanoscale defects to experimentally observable timescales at elevated temperature, we often turn to cluster dynamics (CD)[15], or object kinetic Monte Carlo (okMC)[16, 17, 18]. In the limit of no spatial correlation, ie defects produced homogeneously and with no elastic interactions, CD and okMC give essentially the same results[19]. However, if defect objects are produced in a spatially correlated manner, then this could lead to some rapid recombination in okMC [20, 21]. A challenge therefore is how to handle the spatial correlation in defect production- MD will give this information in the form of atomic positions, but okMC typically simplifies complex configurations to a small handful of parameters describing size, orientation and position of the defect. It is not obvious that the complex atomic configurations seen in MD, particularly in overlapping cascades, should resolve quickly to prismatic loops and clusters[22, 23]. Furthermore it may be the case that if this resolution to simpler forms does occur, it is because of the mutual proximity of defects. okMC often relies on simple capture radii to determine whether clusters react, which may not correctly describe self-climb[24] or correlated atom motion in a dislocation core[25]. Finally we note that including elastic interactions into okMC shows that dislocation loops can be trapped to experimentally observable times by drawing each other into favourable positions in their stress fields[13]. The elastic interaction is most pronounced when loops are close together, which is exactly when the assumptions of okMC starts to break down.

On-the-fly kinetic Monte Carlo methods[26, 27, 28] are sometimes used to evolve arbitrarily complex defects, as they work by searching the configuration space of atomic positions for saddle points between atomic configurations[29]. Accelerated MD can also be used to seek transitions at high temperatures and

then map back to transition rates at the desired low temperature[30, 31]. To find a barrier requires a considerable calculation effort, but when seeking infrequent events over thermally activated barriers, these methods are extremely powerful. For very low barriers, of the order 0.1eV or less, such as might be seen for interstitial cluster movement, they struggle to compete with MD: put simply kMC methods are stochastic and so sample a range of local minima, whereas MD is deterministic and so moves from one minimum to the next. Irradiation damage cascades give us very low barriers with quasi-independent interstitials, mid-range barriers for strongly interacting clusters and high barriers for vacancy clusters. All must be treated consistently.

In this paper we develop okMC with arbitrarily complex off-lattice atomic configurations. The okMC framework constructs a database of previously visited objects and transitions, minimising recalculation. However, rather than search for a comprehensive set of saddle points on-the-fly[32], we use a prescribed set of correlated atomic moves based on those well-known in MD and coarse-grained simulations, and well-parameterized by DFT. We discuss the advantages and disadvantages of this simplified scheme. We also demonstrate how to incorporate elastic interactions efficiently in an off-lattice kMC code.

In section 3.1 we apply the model to some simple test cases to demonstrate correct and efficient running, and in section 3.2 we apply the model directly to radiation damage cascades generated using MD. We conclude that very mobile interstitial clusters may be rapidly lost to the boundaries of the simulation cell, or may collide. If they do collide, then a complex sessile interstitial cluster may be formed first, which then relaxes to a low-energy mobile cluster. This is consistent both with the proposed mechanism of Marian et al [33, 34] that two $\frac{1}{2}\langle 111 \rangle$ loops can collide to form a single $\langle 100 \rangle$ loop, and the experimental *in situ* observation that when loops collide, the resultant loop takes the Burgers vector of the larger[35] - our simulations are exploring both the short collision timescale and the longer relaxation timescale and only the latter is seen in experiment.

2 Atomistic-Object Kinetic Monte Carlo

The code we describe here could be described either as atomistic or object kinetic Monte Carlo, and either as having transitions computed on-the-fly or predetermined. As these terms are common in the kMC literature, and can appear to be mutually exclusive, we will first clarify our meaning.

In materials science, particularly when discussing nano-scale defects in metals, it is common to talk about point defects, clusters, dislocation loops, voids etc. In doing so we are implicitly stating that such defects have a spatial localisation (they exist as quasi-independent objects within a crystal), and a temporal persistence (they are metastable, rather than ephemeral atomic configurations). This is the theoretical basis of object Kinetic Monte Carlo and cluster dynamics - we define the nano-scale defects as objects, and rules determine their dynamic evolution. An okMC state can be completely specified by the types of object, and their positions. An okMC model consists of the current

state, together with rules for evolving the objects. A clear exposition of okMC is given in [18].

Often defect objects will be defined with a very small number of parameters - for example a void may be given a position and number of vacancies contained. The small number of parameters used is for convenience only, there is no *a priori* reason why a void should not be given additional parameters defining its shape, if these could be meaningfully employed in the dynamics. The logical limit of object Kinetic Monte Carlo is to describe each object with atomic resolution: provided each atomic configuration has a padding boundary of perfect crystal, the defect regions are spatially localised, and provided the atoms are elastically relaxed then the configuration is metastable. We describe our okMC state as the types, and positions of objects that it contains, where the type of each object is stored as a hash key to an atomic configuration stored in a database.

In okMC the rules for evolving the state are similarly usually simplified to rigid body translations; a transformation from one object to another; and splitting one object into two (or equivalently its reverse combining two objects into one, as required by detailed balance). These rules often are given as simple rates which vary depending on object type, but with atomistic detail it is possible to find these rates explicitly. The rates for evolving the atomistic defect objects are constructed on-the-fly by considering possible atomistic processes which change the configuration.

Finally we note that in okMC the host crystal is hardly referenced, except possibly as an homogeneous elastic medium, in contrast to MD or standard atomistic KMC. Similarly in our simulation the atoms in the "perfect crystal" far from the objects is not stored. In common with existing okMC codes such as MMONCA[18], the computational cost using our code for a simulation of one object contained in a one thousand atom system is therefore similar to the cost of one object in one million atom system. In contrast with other okMC codes, there is only one type of object in our work- an atomic configuration, with one set of rules for its dynamics as described below.

2.1 An atomistic object

The problem of identifying an atomic configuration or transition are related and have been tackled previously with graph theory [26] and bitmaps [36]. We will use the bit-twiddling Zobrist hashing method [37]. We start with the observation that atoms repel each other at short range and so are never too close together. Consider a simple cubic crystal mesh with side $a_0/4$. The furthest apart two atoms could get while mapping to the same node is $a_0\sqrt{3}/4$. This is 50% the nearest neighbour separation on a bcc lattice side a_0, or 61% the nearest neighbour separation on an fcc lattice. If atoms would prefer to be on bcc or fcc lattice sites, then except during high energy collisions they will never get so close together they will map to the same $a_0/4$ simple cubic node. For practical purposes any complex atomic configuration may be represented uniquely by having zero or one atoms slightly displaced from $a_0/4$ simple cubic nodes. If interstitial alloys are considered the validity of the finer lattice may need to be

verified, but the principle of non-multiply occupied sites on a sufficiently fine lattice remains clear: if we cannot guarantee zero or one atom per node spaced by $a_0/4$, it may yet be possible if the node spacing is $a_0/8$. The extension to non-cubic lattices is similarly trivial.

It is less clear that the mapping of the real space positions in $\mathcal{R}^{3N}$ to fine-scale nodes is one-to-one rather than many-to-one. For the former to be true, then only one relaxed configuration of atoms maps to a particular set of fine-scale nodes, and a small displacement of the atoms would relax back to the same point. This is impossible to guarantee for all potentials, as the watershed hypersurfaces between local minima can be arbitrarily complex [38]. If during a kMC simulation, two slightly perturbed configurations of atoms were associated with a single local minimum, then the damage done to the dynamics would be small, as all the evolution is quasi-static and the barrier between such close minima is most likely negligible anyway. We therefore assert that each configuration of atoms on a fine mesh is associated with a unique minimum, and use the fine mesh occupations to define the configuration. This reduces the description of the configuration to the atomic-chess-board problem, and so we use the extremely efficient Zobrist key[37] to find a 128 bit hash key for an atomic configuration. These hashes index stored objects - that is the atomic configurations plus their associated transitions - on a large database.

We define an object by first identifying all defected unit cells - meaning those which do not have the same occupation of the same fine-scale nodes as the perfect crystal. A buffer region of m unit cells in each direction (ie the cube of $(2m+1)^3$ cells including the defected unit cell) is added, and then centered in a cubic minimal bounding box. This region of atoms is ascribed to an object. An object can therefore be rather large, spanning 32^3 unit cells in some of our cascade evolution simulations, and contain both interstitials and vacancies. We have found that insisting on a buffer region of $m=2$ unit cells around defects is adequate to reproduce the elastic interaction energy between pairs of defects to within 10%, as determined by comparing the energy computed using the dipole tensor formalism (see section 2.3) to a full atomistic relaxation.

In this work we do not include alloying species. We would expect adding an interstitial or substitutional atom like hydrogen or helium to work using this same model for objects (note that bcc tetrahedral [$\frac{1}{2}$ $\frac{1}{4}$0] and octahedral [$\frac{1}{2}$ $\frac{1}{2}$0] interstices are perfectly resolved). As an aside, we note that alloying elements have previously been tackled using a grey-alloy approach [39], but this is left for future work for our model.

2.2 Finding a transition

In a lattice-based kMC simulation, a transition consists of a mobile atom species moving from one lattice site to another, or for a pair of atoms to exchange sites. The number of transitions is determined by the number of mobile atoms and the number of moves each can make. The rate of each transition may be found by consulting a table using the before and after local configurations.

In an off-lattice kMC simulation, transitions can be sought on-the-fly using

the dimer method[32], kART[26, 40], or AMD[30, 41]. The number of transitions is limited by the length of computational time available to search, there will always be more high energy (and therefore low probability) transitions. Though it offers a more exhaustive list of possible transitions in principle, on-the-fly searching is not without its difficulties. Some transitions may be more readily found than others, and care must be taken not to double count them. It is also difficult to ensure that all relevant barriers are added. As an example consider the eight $\langle \frac{1}{2} \frac{1}{2} \frac{1}{2} \rangle$ nearest-neighbour vacancy-atom exchanges in a perfect bcc lattice (we ignore second neighbour $\langle 100 \rangle$ exchanges for clarity of exposition). Each saddle has the same probability of discovery, so the number of times each saddle is found should be Poisson distributed. The probability of not finding a particular saddle is $\bar{p} = exp(-m/8)$, where $m/8$ is the mean number of hits per saddle given m trials. The probability of finding all saddles is $p_{all} = (1 - \bar{p})^8$. To have a $p_{all} > 50\%$ chance of finding all 8 saddles, we need $\bar{p} < 0.083$, and so $m > 20$. To have a $p_{all} > 99\%$ chance of finding all 8 saddles, we need $m > 53$. If objects are to be reused, we need a good coverage of saddles, and so this is potentially a lot of wasted work. Using the symmetry of the system could help reduce wasted work, but in general cases the local surrounding environment will show little symmetry and so make it difficult to achieve great gains. SEAKMC[27] avoids this problem by sampling the saddle points rather than attempting an exhaustive survey. This is efficient, but requires careful tuning before starting a simulation to get correct residence times. It would not be appropriate for a model where barriers are reused. Recently progress has been made to minimise the impact of unsearched transitions using the AMD framework to find transitions[31].

An alternative route for computing the transition energy for this off-lattice atom-vacancy exchange example is to recognise that the move remains identifiable even when atoms are displaced slightly from lattice sites. For $\langle 111 \rangle$ crowdions the moves are also well-known - there exists a Frenkel-Kontorova string pull in the $\langle 111 \rangle$ direction[42, 43], and a rotation through the $\langle 110 \rangle$ dumbbell to $\langle 11\bar{1} \rangle$[43]. We can code these moves by searching for the correct atomic environment to permit the move, then moving atoms appropriately. The number of possible transitions is easily established, and moves which have the same final configuration clearly have different routes, so double- counting them is not a problem. We explicitly sacrifice the chance of finding unusual or unexpected transition paths in favour of a reusable and smaller, but complete, set of transitions.

With the fine mesh defined, we can define three prototype atomic transformations for the bcc lattice illustrated in figure 1.

Vacancy-atom exchange: Find a high energy atom, and look for an unoccupied 'vacancy' lattice site separated by $\langle \frac{1}{2} \frac{1}{2} \frac{1}{2} \rangle$. We permit the transition attempt if the unoccupied site itself is surrounded by 26 unoccupied neighbours on the fine mesh lattice (the region $-\frac{1}{4} : \frac{1}{4}$).

Rotation: Look for a pair of high energy atoms, and displace in opposite $\langle \frac{1}{4} \frac{1}{4} \frac{1}{4} \rangle$ directions.
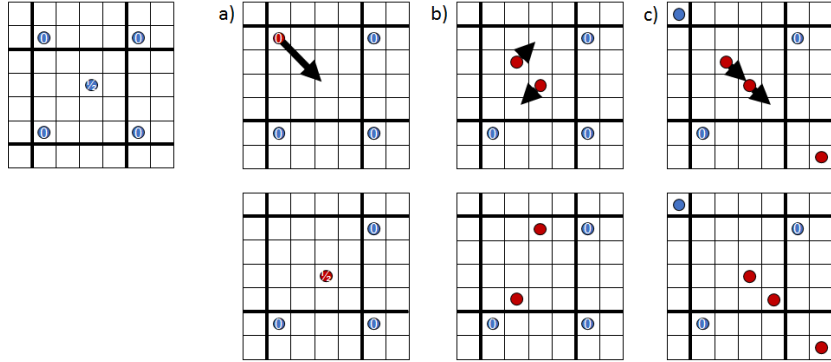
Figure 1: A cartoon illustrating transitions defined on a fine lattice. From left to right, the bcc lattice, a) an atom-vacancy exchange, b) a rotation move, and c) a string pull move. Note that inside the code the atom positions are fully off-lattice, and the fine $a_0/4$ simple cubic lattice is present to distinguish quickly between similar states.

String pull: Look for a high energy atom, and then check in successive $\langle \frac{1}{4}\ \frac{1}{4}\ \frac{1}{4} \rangle$ cells for two more high energy atoms. If they are present, then the original and its first neighbour are displaced along the string.

We also need a fourth prototype transformation for the cooperative motion of clusters of atoms. The multi-string F-K model[25] is a simple extension for understanding the motion of clusters and loops. Here many $\langle 111 \rangle$ strings move over the potential surface, but are additionally coupled together with elastic springs. The springs represent the energy stored in the dislocation core, and acts to prevent one string getting too far ahead of its neighbours. If a single string of atoms is pulled over the PES, the restoring force of its neighbours will tend to bring it back again. There may not even be a relaxed metastable configuration of a dislocation line with just one string of atoms pulled. This is the physical origin of the Ohsawa transition[44] between motion of an interstitial cluster as a single rigid body and by the double-kink mechanism for a large dislocation loop.

For now we can ensure there is a macroscopic diffusion of the centre of mass of a cluster by exploiting the okMC nature of the model. Within an object, we search for a contiguous subcluster of defected unit cells (excluding padding) containing only interstitials. Then for this subcluster we check the principle eigenstresses, as computed by diagonalizing the dipole tensor. An interstitial cluster will have one large eigenstress along its Burgers vector. If these conditions are met, then the rate of transition to the same subcluster translated by $[\frac{1}{2}\ \frac{1}{2}\ \frac{1}{2}]$ along the Burgers vector is fixed to $\nu/\sqrt{n_I}$, where ν is the attempt frequency of a single string pull (see table 1) and n_I the number of interstitials. This square root dependence of the prefactor has been observed in simulation[45, 46] and experiment[47].

A better model might consider simultaneous movement of multiple strings of atoms. But it would be difficult to justify moving two strings together rather than three, or more. The multiplicity of final states to investigate would, at present, lead to an unsupportable slowdown of the code. We therefore leave investigation of the cooperative motion of multiple strings to future work, and acknowledge that our model is currently limited to smaller interstitial clusters.

With these sets of moves and with the atoms off-lattice, a great range of atomic configurations can be explored. It is interesting to note, however, that all transition events that an object can make can be simply categorised.

- Events can define a transformation of the atomic configuration, taking one object to another.

- A transformation may require/permit a resizing of the object to preserve its buffer of perfect crystal.

- A transformation may translate the origin of the object.

- Events can split the object into two or more. In our code we do not consider splitting into more than two objects- which can temporarily leave an object with parts far enough apart that they do not overlap. This is a rare case, and there is no damage done by this to the dynamic evolution anyway, only a minor suboptimal evolution until the full decomposition has been recognised.

There is no requirement to include a transition event combining pairs of objects. This can be automatically included by testing for the overlap of objects at the beginning of every KMC step.

2.3 Parameterization of kMC simulations

All kMC simulations reported here use the same parameterization, with no additional tuning. If the transition takes an object from configuration (A) to (A'), the transition rate $(A \to A')$ is taken to be

$$r_{(A \to A')} = \nu_{(A \to A')} \exp\left[-\frac{\Delta E_m^{(A \to A')}}{k_B T}\right],$$ (1)

where $\nu_{(A \to A')}$ is a rate prefactor, and $\Delta E_m^{(A \to A')}$ a migration barrier. As shown in ref[48], for pure tungsten at least, the simple Kang-Weinberg model[49] is a good approximation for the migration barrier for an atom-vacancy exchange. If we consider a transition which takes an object configuration (A) to (A'), then the migration barrier is

$$\Delta E_m^{(A \to A')} = E_{saddle}^{(A \to A')} - E^{(A)} = \max\left\{\frac{E^{(A')} - E^{(A)}}{2} + \Delta E_m, \frac{\Delta E_m}{2}\right\},$$ (2)

8

where ΔE_m is a constant for the transition category, and $E^{(A)}$ is an appropriate "relaxed" energy for object configuration (A). ΔE_m barriers for all moves are given in table 1. The maximum function is used to ensure that the migration energy is always positive; $E^{(A)}$ may take any value but (A) is known to be metastable. We take the minimum value for the barrier as $\frac{\Delta E_m}{2}$, which is an empirical choice taken to be sufficiently low to be rarely needed in practice, and not unduly affect the dynamics. For fast-moving interstitial clusters the barrier is close to zero in any case; for vacancy defects this limit will only be applied where there is a strong driving force to select one move in preference to another. We define the appropriate relaxed energy below.

As previously noted, we store atomic configurations of defects as localised objects surrounded by a buffer region of perfect crystal, and an object can be large. We compute the energy barrier for a transition by first finding the region where atoms are expected to move, then adding a buffer region to this (see table 1). This cut-out region of 8-10 unit cells can be significantly smaller than the original object which might be 20-30 unit cells.. We fix atoms in the outermost unit cell boundary of the cut-out, and relax. This takes the energy of the cutout region from $E_{cut}^{(A)}$ to $E_{cut}^{(A')}$.

Fixing the atoms ensures that the transition region after the event fits back into its parent object, and takes some account of the strain field in which the transition is located, but note that this is *not* treating the elastic interaction fully. Firstly, we need to account for the response of the remainder of the object to the transition. If $\vec{f}_{embed}$ is the force on the atoms in the object introduced by reintroducing the cutout after the event, and $D_{ab} = \partial^2 E / \partial \vec{r}_a \partial \vec{r}_b$ is the Hessian computed with atoms in their 'after' positions, then the linear elastic response in the object is

$$E_{embed} = -\tfrac{1}{2}\vec{f}_{embed} \cdot D^{-1} \vec{f}_{embed}. \tag{3}$$

This can be evaluated using a Lanczos recursion technique[50, 51], and typically takes order ten milliseconds on a single core, a few percent of the time taken to fully relax the atom positions in the cut-out region. Secondly, as the object is defined by atomic positions, we need to account for the self-energy of the object due to its periodic images. This can be done using the method of Varvenne et al[52]. As we are using embedded atom potentials, the dipole tensor is easily computed using

$$P_{ij} = -\sum_a \sum_{b \in \mathcal{N}_a} \left(\tfrac{1}{2} \left.\frac{\partial V}{\partial r}\right|_{r_{ab}} + \left.\frac{\partial F}{\partial \rho}\right|_{\rho_a} \left.\frac{\partial \phi}{\partial r}\right|_{r_{ab}} \right) \frac{r_{ab,i} r_{ab,j}}{r_{ab}}, \tag{4}$$

where $r_{ab,i}$ is the i^{th} Cartesian coordinate of the separation between atoms a and b, $V(r)$ is a pairwise potential and $F[\rho]$ a many body contribution from the embedding density function $\rho = \sum \phi(r)$. This is a simple sum over atoms and their neighbours, using the same first derivatives as a force, and so is of negligible computational cost. The elastic energy between an object with dipole tensor $P_{ij}^{(A)}$ and a second with dipole tensor $P_{ij}^{(B)}$ located at separation $\vec{R}_{AB}$

is[53]

$$E_{elas}^{(A,B)}\left(\vec{R}_{AB}\right) = P_{ij}^{(A)}\frac{\partial^2 G_{ik}\left(\vec{R}_{AB}\right)}{\partial x_j \partial x_l}P_{kl}^{(B)}, \tag{5}$$

where G_{ij} is the elastic Greens function. In this work we use the isotropic elastic Green's function for convenience, as tungsten is nearly elastically isotropic. This requires the bulk elastic Lamé parameters λ and μ, which we compute at the beginning of the simulation using the empirical potential supplied. The centre-of-position of an object is taken to be [54]

$$\mathbf{R} = \frac{\sum_{a=1}^{n}\|\mathbf{P}^a\|\,\mathbf{R}^a}{\sum_{a=1}^{n}\|\mathbf{P}^a\|} \tag{6}$$

where $\|\mathbf{P}^a\| = \sqrt{\mathrm{Tr}((\mathbf{P}^a)^2)}$ is the Frobenius norm- a measure of the strength of the stress field generated by the a^{th} atom.

The self energy of the object is found in principle by the conditionally convergent sum of equation 5 over defect periodic images[52, 55]:

$$E_{self}^{(A)} = \sum_{uvw} E_{elas}^{(A,A)}\left(\vec{R}_{uvw}\right), \tag{7}$$

where $\vec{R}_{uvw}$ is the position of the image translated $\{u, v, w\}$ periodic repeats in each Cartesian direction. To avoid convergence issues we simplify using $\{u, v, w\} \in [-1 : 1]$.

Thirdly we need to account for the changing elastic response of the system as a whole. This is again done with the dipole tensor formalism, an approach previously exploited by Subramanian et al[56], and in a simplified form in ref [13]. If the transition takes the object from configuration (A) to (A'), then the energy difference is

$$
\begin{aligned}
E^{(A')} - E^{(A)} &= E_{cut}^{(A')} - E_{cut}^{(A)} \\
&+ E_{embed}^{(A')} - E_{embed}^{(A)} \\
&+ E_{self}^{(A')} - E_{self}^{(A)} \\
&+ \sum_{B \neq A}\left(P_{ij}^{(A')}\frac{\partial^2 G_{ik}\left(\vec{R}_{A'B}\right)}{\partial x_j \partial x_l} - P_{ij}^{(A)}\frac{\partial^2 G_{ik}\left(\vec{R}_{AB}\right)}{\partial x_j \partial x_l}\right)P_{kl}^{(B)}
\end{aligned} \tag{8}
$$

The spatial regions for each level of relaxation are illustrated in figure 2.

The migration barriers, characteristic frequencies and padding using for the cutout region are given in table 1. As a side-benefit of having categories of transition defined, the rate prefactor may itself be a function of temperature or characteristic of the defect structure. We set the rate prefactor of string-pull type low-energy barrier moves to be linearly dependent on temperature, corresponding to the friction limited regime [57].

Note that the embedding energies can be stored as data for the transition and reused , and the self-energy is a property of the object. The long-range
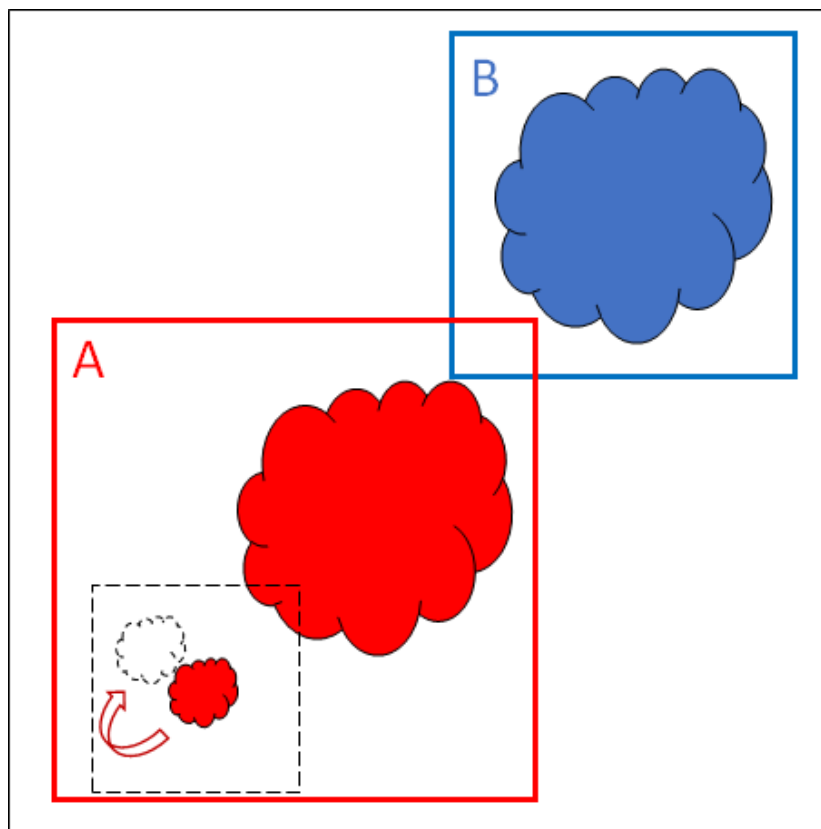
Figure 2: A cartoon illustrating a transition in object $(A) \to (A')$. The transition is computed in a small cutout region (indicated by the hashed lines), with the atoms on the boundary fixed. The energy of the atoms in the cutout region is computed using empirical potentials. The relaxation of object (A') due to re-embedding the cutout is computed at the harmonic level using the Hessian over the whole object (A'). The self-energy of object (A') due to interactions with periodic copies of itself is computed using its dipole tensor. The interaction of object (A') with object (B) is computed with dipole tensors.

move	high energy atom (eV)	barrier ΔE_m (eV)	rate prefactor (THz)	padding (u.c.)
atom-vacancy exchange	0.25	1.75	6.45	2
atom pair rotation	1.10	0.40	1.0	3
string pull	0.50	0.013	0.75 k_BT	3
cooperative	cluster	0.013	0.75 k_BT / $\sqrt{N}$	2

Table 1: Parameterization of the three akMC and one okMC transition events considered in our simulations. High energy atoms with potential energy over the threshold are considered as candidates for the move. The barriers are computed weighted using the before- and after- energies, including full elastic correction, according to equation 8, with an additional barrier given in the table. The rate 6.45 THz is the Debye frequency for tungsten [58]. The energy barrier for vacancy migration is taken from DFT [59], which is in good agreement with experiment [60]. The energy barrier for string pull and crowdion rotation are taken from DFT [43]. The string pull and translation moves are assumed to be friction limited, with rate constant linearly proportional to $k_B T$.

inter-object energy of interaction needs to be recomputed at every okMC step. If the dipole tensor is stored on the object, this takes negligible time compared to relaxing the atoms.

We compute and store all the transition rates (according to our search rules) and so can use rejection-free kinetic Monte Carlo[61] using the basin-autoclimbing Mean Rate Method[62].

3   Results

3.1   Individual defects

Before moving to complex systems with multiple moving defects, we present some validation work with isolated individual defects. This offers a comparison with literature, and demonstrates the working of the code.

We consider the dynamics of a single monovacancy, a quad-vacancy cluster, a monointerstitial, and interstitial clusters of size 2,7,13,19. Here, all are simulated at 600K, using an empirical potential known to be good for vacancy defects ref[48]. In section 3.3 we show this empirical potential is also good for interstitial defects.

In section 2.3 we explained how the migration barrier $\Delta E^{(A \to A')}$ was computed. In figure 3 we present histograms of barriers computed over a selection of 1000 step simulations. The monovacancy always has equal-energy before- and after- states, so every barrier is exactly $\Delta E^{(A \to A')} = \Delta E_m = 1.75$ eV. By contrast the quadvacancy is a low energy cluster, so to evolve it must gain energy. It therefore shows a range of barriers above 1.75 eV, and a few transitions
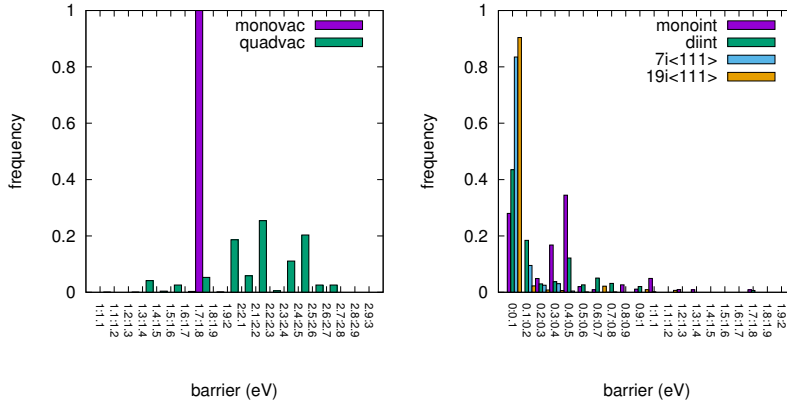
Figure 3: Histograms of relative frequencies of migration barriers computed over 1000 kMC steps at 600K, using the MNB potential in tungsten. Left: vacancy-type objects. A single monovacancy has a barriers at exactly 1.75eV by construction, but a quad-vacancy explores a wider range. As the quad-vacancy itself is a low-energy cluster, many are above the isolated vacancy barrier. Right: interstitial-type objects. A single crowdion shows barriers near 0eV and 0.4eV, corresponding to translation and rotation modes, and a few higher energy transitions to other single-interstitial formations, octahedral, $\langle 100 \rangle$ dumbbell etc. The diinterstitial has a wider range of translation modes where one string pulls past the other, and a few rotations are found. The 7 interstitial cluster and 19-interstitial loop show mostly translation modes.

back to the low-energy cluster below 1.75 eV. The crowdion shows translations and transformations between crowdion and dumbbell near 0.0 eV, and some rotations near 0.4 eV. It also shows some high energy barriers to higher energy single-interstitial formations. The interstitial clusters show an increasing fraction of low-energy string-pull and cooperative cluster transitions either to different configurations or translations.

Our code finds a range of possible atomistic moves using expensive atomistic relaxations, and stores them. When all low energy barriers describing translations and transformations have been computed, the object KMC model is fully determined, and so the code can evolve them at the same speed as any other okMC. This is illustrated in figure 4. On this plot is also indicated the performance of a good MD code such as LAMMPS[63], which currently takes about one second wall time to perform an update step for 1 million atoms, the update step being a simulated time of one femtosecond. An important point to note in figure 4 is that our KMC code starts slower, but becomes faster than this canonical benchmark for MD because of the efficient reuse of object information.

We have computed the diffusion constant for interstitial defects using the method of ref[43]. The result is shown in figure 5. The string-pull attempt fre-
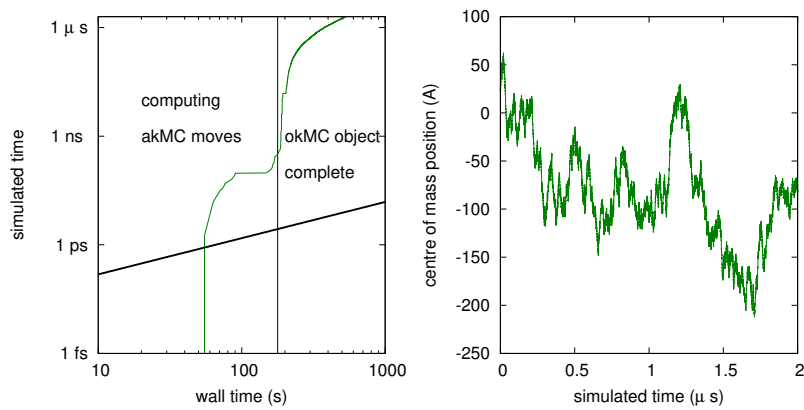
Figure 4: Data from a single run computing the diffusion constant of a 13 interstitial cluster at 600K in a $32 \times 32 \times 32$ unit cell box. Left: for three minutes wall time (using a single core on a desktop PC), string pulling moves are considered and stored. After this they can just be recalled, and so the code accelerates. The solid black line is an indication of MD speed, assuming a good MD code can perform one million atom update steps per second (per core), and each timestep is 1fs. Right: the (x-) position of the centre of mass of the cluster, computed using equation 6, showing expected Brownian motion with no significant drift. This is a trivial exercise using standard okMC where the rules for translations are predetermined, but less so for an on-the-fly code where the rules for translation are computed.
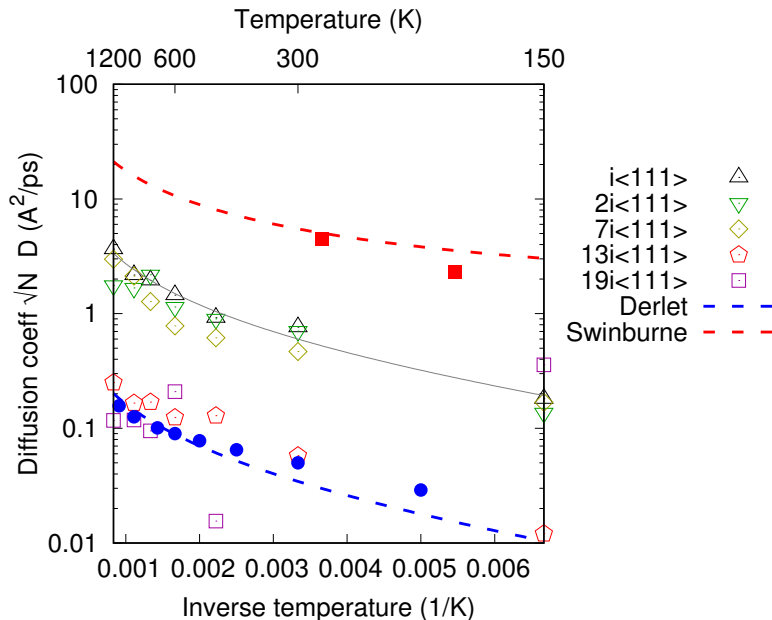
14

Figure 5: The diffusion constant (points) computed using our model, multiplied by the square root of the interstitial count. The dashed lines are the theoretical prediction from Swinburne et al[64], and Derlet et al [43], and solid symbols data points from MD simulations reported in those papers. We conclude that the model presented here has order-of-magnitude correct diffusion coefficients for isolated interstitial clusters. Note that non-Arrhenius behaviour can be seen in this plot, a consequence of the temperature-dependent rate prefactor (see table 1.)

quency and energy barrier are fixed to give a good crowdion diffusion constant. The computed diffusion constant for clusters is within an order of magnitude of literature results. This error is not expected to have a significant impact - clusters undergo a fast macroscopic displacement of the centre of mass, so this will not be a rate limiting step in longer term microstructural evolution. We would prefer not to overfit the model by introducing additional rules.

3.2 kMC annealing of primary radiation damage cascades

In this section we anneal the complex defects produced in MD simulations of irradiation damage cascades.

Irradiation cascade simulations were performed using the classical molecular dynamics code PARCAS [65]. The simulations were performed in bulk tungsten, using the interatomic potential for tungsten by Derlet et al. [43], stiffened at short range[66]. Simulation cells were $68 \times 68 \times 68$ unit cells (629k atoms), initially at 0 K, with the atoms on the periodic boundaries thermostatted
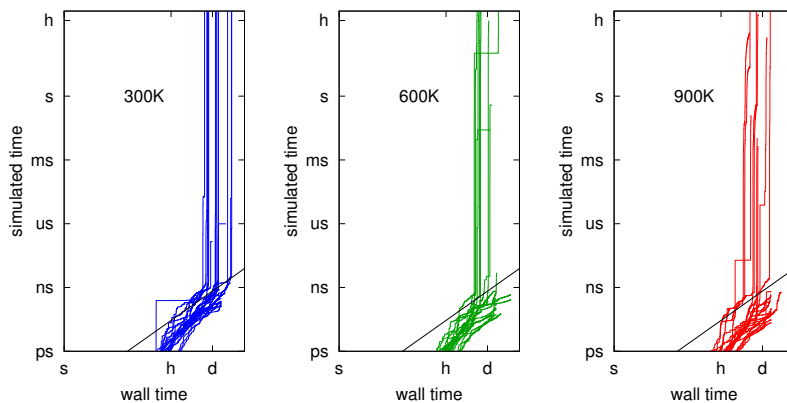
Figure 6: Performance of the code, plotted as simulated vs wall time. The thick black line corresponds to a good MD code, running at 1 fs/1M atoms/s. After the first nanosecond of simulated time the KMC code overtakes MD. Note that we are considering multiple independent simulations on a single processor, and that we acknowledge parallelizing MD is significantly easier than parallelizing kMC.

to 0 K[67]. Electronic stopping was modelled using a frictional force applied to atoms with kinetic energy over 10 eV, with the magnitude of the electronic friction determined by SRIM [68]. As discussed in section 2.3, our model is currently only parameterized to correctly simulate the dynamics of small interstitial clusters, so we use low energy cascades, which have a low probability of producing a large loop [69, 9]. One atom was given an intitial kinetic energy of 20 keV, and the simulations followed until cool (< 10 K) at 40 ps. Details of the simulation method are given in [70]. Five cascade configurations so formed were used as the starting point for further simulations described below. We cut out a cubic region of $40 \times 40 \times 40$ unit cells (128000 atoms) containing the cascade with ample buffer. We ran the simulations at 300K,600K and 900K multiple times with different random seeds.

One additional simulation technique used for the cascade simulations is to employ absorbing boundary conditions. This is simple to implement in atomistic KMC: at the beginning of each move all atoms in the unit cells on the boundary of the simulation are set to be perfect crystal. If this affects an object, then the object is relaxed before continuing, but now it may contain a different atom count. Using this method there is no strong bias towards the boundary, but neither will the defect return from the boundary.

First we consider the timing results of our KMC code. It is important that an off-lattice KMC code is actually faster than MD, and this is a difficult milestone to reach. We see in figure 6 that our code overtakes MD after about 1ns of simulated time, which corresponds to about 24 hours wall time using a single processor.

16

Next we consider the evolution of the energy of the system. In figure 8 we see that the greatest portion of the energy reduction is in the first nanosecond of simulated time. This again we attribute to crowdion diffusion, absorption and annihilation. Figure 7 shows this behaviour for a cascade which starts with a diffuse collection of small defect clusters. Three crowdions are lost to the boundaries within 100 ps and two more have annihilated by recombination with vacancies. Two further crowdions are elastically trapped by each other and a small cluster of vacancies nearby. After 150 ps only two interstitial clusters remain, and both are glissile, so that after 120 ns both have left the simulation cell.

After one nanosecond simulated time, the system evolves through a set of descending energy plateaux. This is common behaviour in kinetic Monte Carlo, associated with escaping trapping basins of energy states. In our case we find sessile interstitial clusters form through collisions of smaller clusters. These then eventually relax to mobile clusters.

In figure 9 we show snapshots from a simulation showing the rapid formation and slow relaxation of a sessile interstitial cluster. We first see the collision $2i\langle 111\rangle + 6i\langle 111\rangle \rightarrow 8i\langle 100\rangle$. This is a stable sessile defect, which would not then further evolve if simulated in isolation. In our case the sessile cluster is formed close to the centre of a dense cascade. This simulation was run at a temperature of 600K, where vacancies have a small mobility. The cluster absorbs a vacancy, in the reaction $8i\langle 100\rangle + v \rightarrow 7i(\text{complex})$. This then initiates a transformation into a glissile cluster $7i(\text{complex}) \rightarrow 7i\langle 111\rangle$. In other simulations we have seen small sessile 3- and 4- interstitial clusters transform into a glissile cluster by absorption of an interstitial.

The last phase of annealing that we observe occurs after 1ms at 900K, or 1s at 600K. This phase is the onset of vacancy mobility. In the simulation shown in 7 the vacancies start well separated, and after one hour simulated time only a trivacancy and a quad-vacancy remain.

3.3 Discussion

The most physically interesting phase of the annealing is in the nanosecond to microsecond timescale, where complex sessile interstitial clusters are formed and subsequently relax. To investigate the occurrence of sessile defect clusters, we have systematically explored the energy landscape of interstitial clusters in tungsten. We placed N interstitials into a $2 \times 2 \times 2$ unit cell box, randomly placed at the fine-mesh quarter lattice positions, then embedded this into a $16 \times 16 \times 16$ unit cell box and relaxed at constant volume. We took $10N^2$ initial configurations. The binding energy is defined as the difference in formation energy between the lowest energy mono-interstitial and the N-interstitial cluster:

$$E_b \equiv NE_1^f - E_N^f. \tag{9}$$

Figure 10 shows the result as a histogram over binding energies. The height of the peaks is proportional to the number of times the bin is hit, without
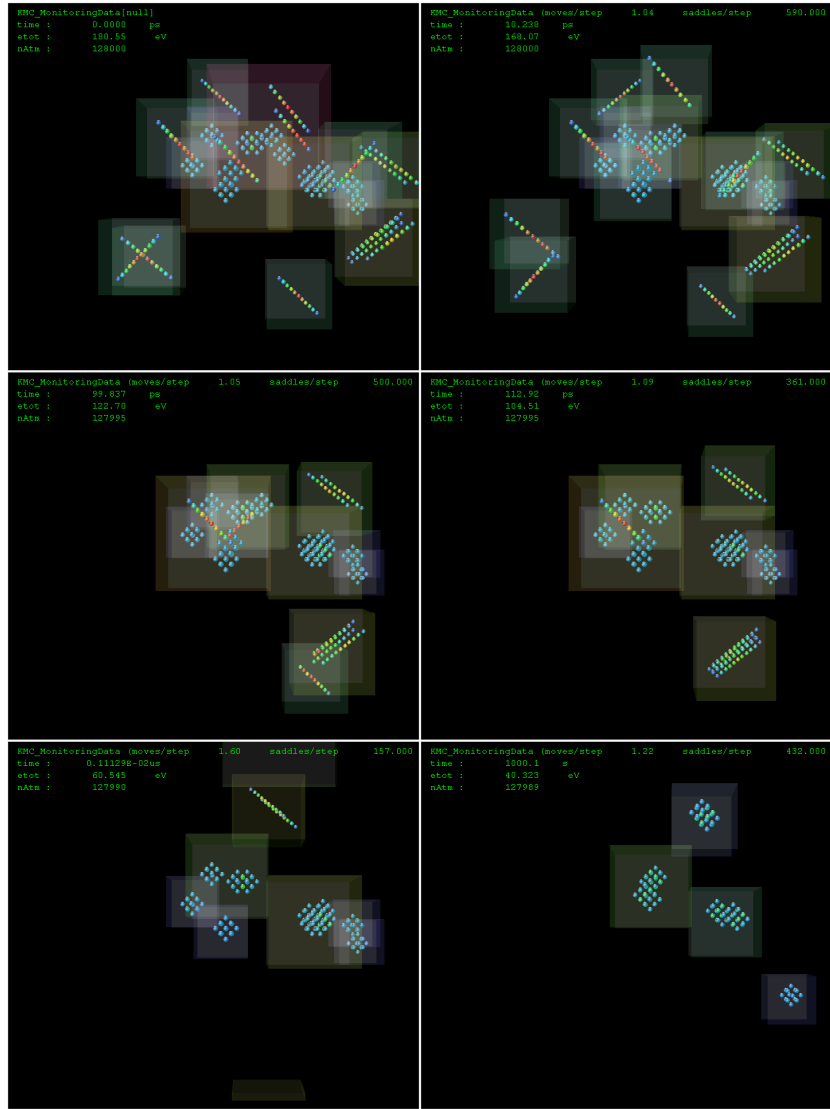
Figure 7: Snapshots from the evolution of a 20keV irradiation damage cascade in tungsten evolving at 600K. Reading left-to-right and top-to-bottom, the snapshots are taken at time $t = 0, 10ps, 100ps, 113ps, 100ns, 1000s$. Starting from the top left image, the cascade starts with a number of crowdions. After 100ps three have diffused out of the simulation cell, and two more have annihilated on vacancies. At the bottom right we see a loose 4-interstitial cluster gather a crowdion to become a 5-interstitial cluster. After 150ps only two small interstitial clusters remain, and by 120ns both have left the simulation cell. Brownian motion of the vacancies ultimately leads to accumulation into small 3- and 4- vacancy clusters. Only high energy atoms are shown, coloured from blue (+0.1eV) through green to red (+1.5eV). This calculation took 26 h on a single processor.
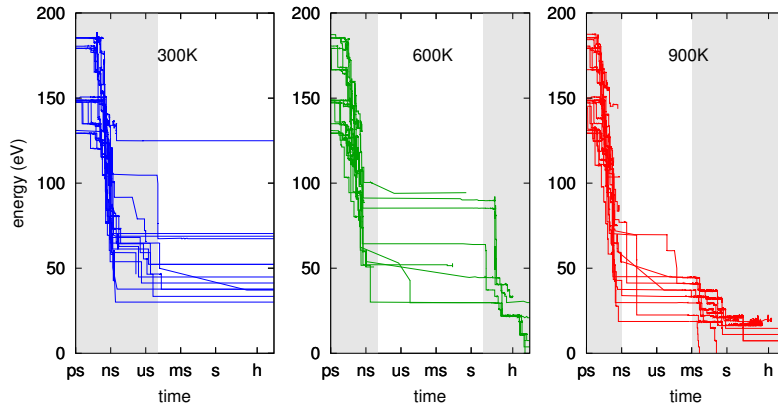
Figure 8: Evolution of the energy in the cascades considered as a function of simulated time. We can broadly split the evolution into the three parts, indicated by the vertical bands. The greatest reduction in energy comes in the first few ns of simulated time. After this the energy evolves through a number of descending plateaux. At the longest times simulated vacancy motion becomes possible, and further relaxation is possible.

accounting for the same structure being found multiple times. We find that the lowest energy structures for each cluster size correspond to all interstitials close together and orientated along $\langle 111 \rangle$, but there are also a large number of higher energy clusters, many of which are difficult to categorize. This same conclusion about distinct low energy clusters and a semi-continuum band of metastable states was previously drawn by Marinica et al for the Fe-interstitial landscape [72]. A full study of these interstitial defect clusters can be found in ref [73]. We conclude that the MNB empirical potential [48] gives good formation energies and relaxation volumes for these clusters, compared to DFT calculations.

The complex energy landscape of uncategorisable interstitial cluster configurations is significant for the microstructural evolution found in this study. This result suggests that if interstitial clusters generated in a cascade collide they can first form a single high-energy interstitial cluster, before subsequently relaxing further, ultimately to a low energy mobile defect cluster.

4 Conclusions

We have developed a new okMC code which takes atomistic configurations as the elementary objects, rather than using simple idealised defects. The code operates by searching for, and storing atomistic transitions, based on vacancy-atom exchange, and correlated atom rotation and string pulling moves. These moves are stored, and recalled, so that an object once discovered needs never be recomputed. Typically a large number of configurations of each defect cluster
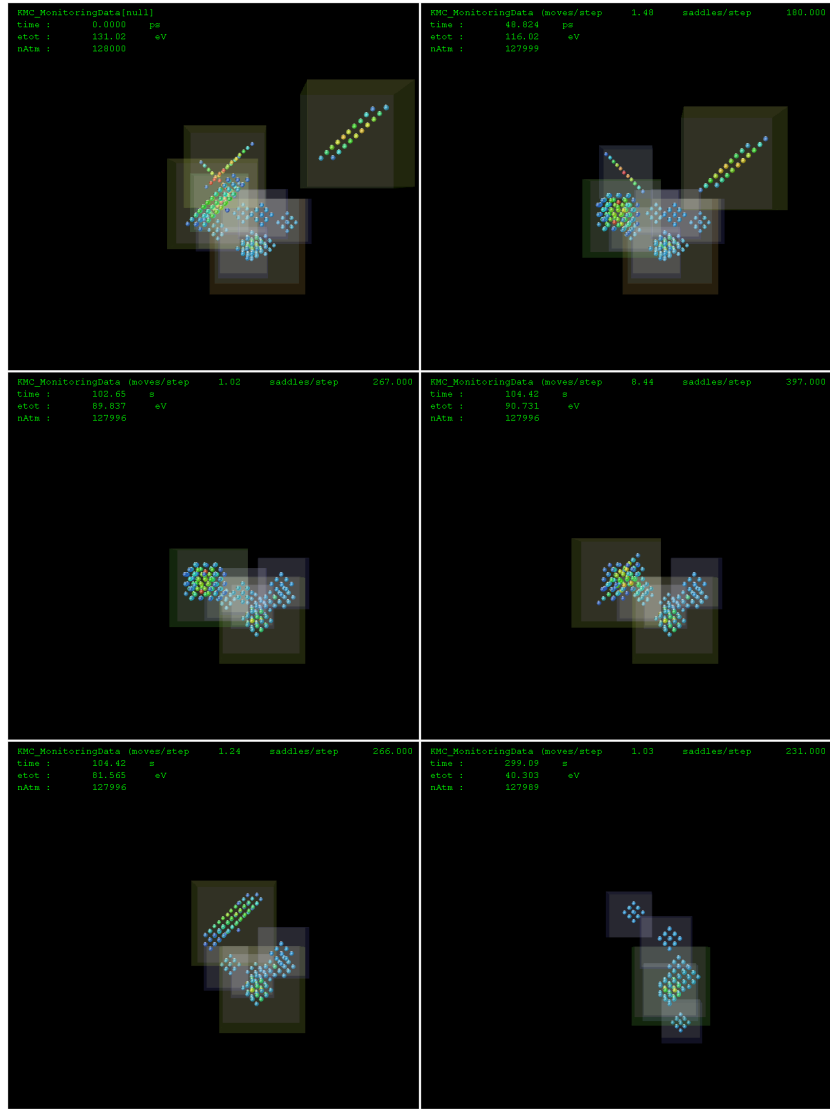
Figure 9: Snapshots from the evolution of a 20keV irradiation damage cascade in tungsten evolving at 600K. Reading left-to-right and top-to-bottom, the snapshots are taken at time $t = 0, 50ps, 103s, 104s, 104s, 300s$. Starting from the top left image, the diinterstitial at the top left combines with the larger 6-interstitial cluster, forming a sessile 8-interstitial cluster. The remaining monointerstitial and diinterstitial rapidly diffuse out of the simulation cell leaving only sessile objects. After 100s the 8-interstitial cluster absorbs a monovacancy, starting a rapid transformation into a mobile 7-intersitial $\langle 111 \rangle$ cluster. Brownian motion of the vacancies leads to accumulation into a single 9-vacancy void at about 1000s. Only high energy atoms are shown, coloured from blue (+0.1eV) through green to red (+1.5eV). This calculation took 58 h on a single processor.
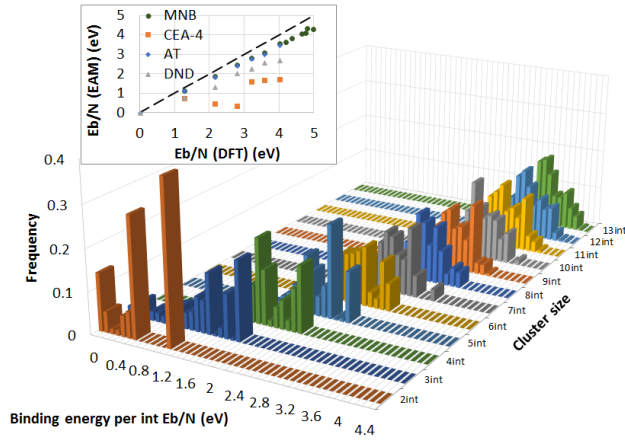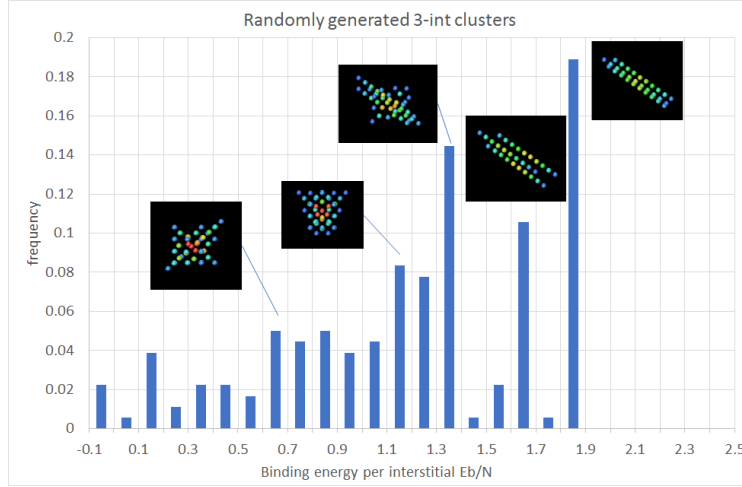
20

Figure 10: If interstitials are randomly placed in close proximity and relaxed, a wide range of different structures are found. The top figure shows a histogram of binding energies found for three-interstitial clusters, with some of the structures drawn. The bottom figure shows similar histograms for cluster sizes 2-13 interstitials. Inset: a comparison of the binding energy for the minimum energy structures computed with DFT[71] and different empirical potentials. The lowest energy structures found for each cluster size are mobile ⟨111⟩-type defect clusters, but the majority are sessile and difficult to simply categorize. We conclude that if defects produced in cascades collide, they are likely to form a sessile metastable configuration before transforming to a mobile configuration.

will be computed, but when all relevant configurations are found no further expensive atomic searches or relaxations are required and the code speed is comparable to a regular okMC code, albeit with a higher memory footprint. If defects collide, or transform into a previously unexplored configuration, the object can not be recalled from the database and so more atomistic transitions are added.

Elastic interactions are treated differently at four different length-scales. Within an object, the atoms are kept relaxed using interatomic potentials, allowing for arbitrarily complex atomic configurations. When a local transition is considered, this is done by clamping atoms on the boundary of the transition active volume. The energy impact of this clamping is removed by computing the harmonic elastic relaxation energy, computed using the Hessian. The self-energy due to interactions with periodic replicas is subtracted using the dipole tensor (computed with interatomic potentials) and the isotropic elastic Greens function, in the manner of Varvenne et al[52]. Finally the interaction between objects is treated using dipole tensors and isotropic elasticity. The long-range elastic energy difference between the before- and after- states is added to each transition.

We have found that the first nanosecond simulated time is by far the most computationally expensive part of the simulation. In future work it would be advantageous therefore to exploit MD to run the first nanosecond of simulation, rather than stopping the simulation shortly after the PKA was set in motion. Our cascade annealing results were produced using a serial code, taking tens of hours to complete each run. While it may be possible to parallelize the code in the future, this is not a trivial task due to the difficulty in load-balancing: each kMC event requires atomic relaxations with very different times required for convergence. As our computational burden is light, we can of course benefit from the trivial parallelization offered by running multiple independent simulations for generating statistics. It is also important to note that we have deliberately simulated cascade annealing in the dilute (low-fluence) limit, exploiting the loss of mobile defects to distant sinks to accelerate the dynamics. In the dense limit we would expect cascade overlaps, which would be better modelled with MD rather than kMC alone.

Our results demonstrate that there may exist interesting modes of cascade relaxation, where glissile defects collide to form a sessile defect. This sessile defect may then spontaneously transform to a glissile form, or may have such a transformation initiated by absorption of a mobile defect. At present we do not have sufficient statistical information to be able to draw strong conclusions about the relative rate of these relaxation modes, but this does indicate that it may be necessary to take into account the spatial correlation of multiple defects generated in a small cascade.

5 Acknowledgements

References

[1] M.L. Jenkins and M.A. Kirk. *Characterization of Radiation Damage by Transmission Electron Microscopy*. Series in Microscopy in Materials Science. IOP, Bristol, 2001.

[2] Marquis Kirk, Xiaoou Yi, and Michael Jenkins. Characterization of irradiation defect structures and densities by transmission electron microscopy. *Journal of Materials Research*, 30(9):1195?1201, 2015.

[3] Philip D Edmondson, Alan Xu, Luke R Hanna, Michal Dagan, Steve G Roberts, and Lance L Snead. Atomic scale, 3-dimensional characterization of radiation effects in tungsten for fusion applications. *Microscopy and Microanalysis*, 21(S3):579–580, 2015.

[4] C. Liu, L. He, Y. Zhai, B. Tyburska-Püschel, P.M. Voyles, K. Sridharan, D. Morgan, and I. Szlufarska. Evolution of small defect clusters in ion-irradiated 3C-SiC: Combined cluster dynamics modeling and experimental study. *Acta Materialia*, 125:377 – 389, 2017.

[5] Michal Dagan, Luke R. Hanna, Alan Xu, Steve G. Roberts, George D.W. Smith, Baptiste Gault, Philip D. Edmondson, Paul A.J. Bagot, and Michael P. Moody. Imaging of radiation damage using complementary field ion microscopy and atom probe tomography. *Ultramicroscopy*, 159:387 – 394, 2015. 1st International Conference on Atom Probe Tomography & Microscopy.

[6] Roger E Stoller and Lawrence R Greenwood. Subcascade formation in displacement cascade simulations: Implications for fusion reactor materials1. *Journal of Nuclear Materials*, 271–272:57 – 62, 1999.

[7] Wahyu Setyawan, Giridhar Nandipati, Kenneth J. Roche, Howard L. Heinisch, Brian D. Wirth, and Richard J. Kurtz. Displacement cascades and defects annealing in tungsten, part i: Defect database from molecular dynamics simulations. *Journal of Nuclear Materials*, 462:329 – 337, 2015.

[8] A. E. Sand, S. L. Dudarev, and K. Nordlund. High-energy collision cascades in tungsten: Dislocation loops structure and clustering scaling laws. *EPL (Europhysics Letters)*, 103(4):46003, 2013.

[9] A. E. Sand, M. J. Aliaga, M. J. Caturla, and K. Nordlund. Surface effects and statistical laws of defects in primary radiation damage: Tungsten vs. iron. *EPL (Europhysics Letters)*, 115(3):36001, 2016.

[10] X. Yi, A. E. Sand, D. R. Mason, M. A. Kirk, S. G. Roberts, K. Nordlund, and S. L. Dudarev. Direct observation of size scaling and elastic interaction between nano-scale defects in collision cascades. *EPL (Europhysics Letters)*, 110(3):36001, 2015.

[11] D.R. Mason, A.E. Sand, X. Yi, and S.L. Dudarev. Direct observation of the spatial distribution of primary cascade damage in tungsten. *Acta Materialia*, 144:905 – 917, 2018.

[12] D. R. Mason, X. Yi, A. E. Sand, and S. L. Dudarev. Experimental observation of the number of visible defects produced in individual primary damage cascades in irradiated tungsten. *EPL (Europhysics Letters)*, 122(6):66001, 2018.

[13] D.R. Mason, X. Yi, M.A. Kirk, and S.L. Dudarev. Elastic trapping of dislocation loops in cascades in ion-irradiated tungsten foils. *J. Phys.: Condens. Matter*, 26:375701, 2014.

[14] S. L. Dudarev, M. R. Gilbert, K. Arakawa, H. Mori, Z. Yao, M. L. Jenkins, and P. M. Derlet. Langevin model for real-time Brownian dynamics of interacting nanodefects in irradiated metals. *Phys. Rev. B*, 81:224107, Jun 2010.

[15] Jaime Marian and Vasily V. Bulatov. Stochastic cluster dynamics method for simulations of multispecies irradiation damage accumulation. *Journal of Nuclear Materials*, 415(1):84 – 95, 2011.

[16] C. Domain, C.S. Becquart, and L. Malerba. Simulation of radiation damage in fe alloys: an object kinetic Monte Carlo approach. *Journal of Nuclear Materials*, 335(1):121 – 145, 2004.

[17] C.S. Becquart and C. Domain. An object kinetic Monte Carlo simulation of the dynamics of helium and point defects in tungsten. *Journal of Nuclear Materials*, 385(2):223 – 227, 2009. Nuclear Materials {IIIProceedings} of the E-MRS 2008 Spring Meeting: Third Symposium N on Nuclear Materials.

[18] Ignacio Martin-Bragado, Antonio Rivera, Gonzalo Valles, Jose Luis Gomez-Selles, and María J. Caturla. MMonCa: An object kinetic Monte Carlo simulator for damage irradiation evolution and defect diffusion. *Computer Physics Communications*, 184(12):2703 – 2710, 2013.

[19] R.E. Stoller, S.I. Golubov, C. Domain, and C.S. Becquart. Mean field rate theory and object kinetic Monte Carlo: A comparison of kinetic models. *Journal of Nuclear Materials*, 382(2):77 – 90, 2008. Microstructural Processes in Irradiated Materials.

[20] T. Jourdan and J.-P. Crocombette. Rate theory cluster dynamics simulations including spatial correlations within displacement cascades. *Phys. Rev. B*, 86:054113, Aug 2012.

[21] Giridhar Nandipati, Wahyu Setyawan, Howard L. Heinisch, Kenneth J. Roche, Richard J. Kurtz, and Brian D. Wirth. Displacement cascades and defect annealing in tungsten, part ii: Object kinetic Monte Carlo simulation of tungsten cascade aging. *Journal of Nuclear Materials*, 462:338 – 344, 2015.

[22] A.E. Sand, J. Byggmstar, A. Zitting, and K. Nordlund. Defect structures and statistics in overlapping cascade damage in fusion-relevant bcc metals. *Journal of Nuclear Materials*, 511:64 – 74, 2018.

[23] F. Granberg, J. Byggmästar, A. E. Sand, and K. Nordlund. Cascade debris overlap mechanism of $< 100 >$ dislocation loop formation in Fe and FeCr. *EPL (Europhysics Letters)*, 119(5):56003, 2017.

[24] T.D. Swinburne, K. Arakawa, H. Mori, H. Yasuda, M. Isshiki, K. Mimura, M. Uchikoshi, and S.L. Dudarev. Fast, vacancy-free climb of prismatic dislocation loops in BCC metals. *Sci. Rep.*, 6:30596, 2016.

[25] S. L. Dudarev. Coherent motion of interstitial defects in a crystalline material. *Philosophical Magazine*, 83(31-34):3577–3597, 2003.

[26] L.K. Béland, P. Brommer, F. El-Mellouhi, J.-F. Joly, and N. Mousseau. Kinetic activation-relaxation technique. *Phys. Rev. E*, 84:046704, Oct 2011.

[27] Haixuan Xu, Yuri N Osetsky, and Roger E Stoller. Self-evolving atomistic kinetic Monte Carlo: fundamentals and applications. *Journal of Physics: Condensed Matter*, 24(37):375402, 2012.

[28] Laurent Karim Béland, Yuri N. Osetsky, Roger E. Stoller, and Haixuan Xu. Kinetic activation-relaxation technique and self-evolving atomistic kinetic Monte Carlo: Comparison of on-the-fly kinetic Monte Carlo algorithms. *Computational Materials Science*, 100:124 – 134, 2015. Special Issue on Advanced Simulation Methods.

[29] Eduardo Machado-Charry, Laurent Karim Béland, Damien Caliste, Luigi Genovese, Thierry Deutsch, Normand Mousseau, and Pascal Pochet. Optimized energy landscape exploration using the ab initio based activation-relaxation technique. *The Journal of Chemical Physics*, 135(3):034102, 2011.

[30] Danny Perez, Blas P. Uberuaga, Yunsic Shim, Jacques G. Amar, and Arthur F. Voter. Chapter 4 accelerated molecular dynamics methods: Introduction and recent developments. volume 5 of *Annual Reports in Computational Chemistry*, pages 79 – 98. Elsevier, 2009.

[31] Thomas D. Swinburne and Danny Perez. Self-optimized construction of transition rate matrices from accelerated atomistic simulations with Bayesian uncertainty quantification. *Phys. Rev. Materials*, 2:053802, May 2018.

[32] Graeme Henkelman, Gísli Jóhannesson, and Hannes Jónsson. *Methods for Finding Saddle Points and Minimum Energy Paths*, pages 269–302. Springer Netherlands, Dordrecht, 2002.

[33] J. Marian and B.D. Wirth. Mechanism of formation and growth of $< 100 >$ interstitial loops in ferritic materials. *Phys. Rev. Lett.*, 88(25):255507, 2002.

[34] H. Xu, R.E. Stoller, Y.N. Osetsky, and D. Terentyev. Solving the puzzle of h100i interstitial loop formation in bcc iron. *Phys. Rev. Lett.*, 110:265503, 2013.

[35] K. Arakawa, T. Amino, and H. Mori. Direct observation of the coalescence process between nanoscale dislocation loops with different burgers vectors. *Acta Materialia*, 59(1):141 – 145, 2011.

[36] Giridhar Nandipati, Niranjan Govind, Amity Andersen, and Aashish Rohatgi. Self-learning kinetic Monte Carlo simulations of Al diffusion in Mg. *Journal of Physics: Condensed Matter*, 28(15):155001, 2016.

[37] D.R. Mason, T.S. Hudson, and A.P. Sutton. Fast recall of state-history in kinetic Monte Carlo simulations utilizing the Zobrist key. *Computer Physics Communications*, 165(1):37 – 48, 2005.

[38] Chris J Pickard and R J Needs. Ab initio random structure searching. *Journal of Physics: Condensed Matter*, 23(5):053201, 2011.

[39] M. Chiapetto, L. Malerba, and C.S. Becquart. Effect of Cr content on the nanostructural evolution of irradiated ferritic/martensitic alloys: An object kinetic Monte Carlo model. *Journal of Nuclear Materials*, 465:326 – 336, 2015.

[40] Fedwa El-Mellouhi, Normand Mousseau, and Laurent J. Lewis. Kinetic activation-relaxation technique: An off-lattice self-learning kinetic monte carlo algorithm. *Phys. Rev. B*, 78:153202, Oct 2008.

[41] B.P. Uberuaga and Voter A.F. *Accelerated Molecular Dynamics Methods*, volume 235, pages 25–43. Springer, Dordrecht, 2007.

[42] Oleg M. Braun and Yuri S. Kivshar. Nonlinear dynamics of the Frenkel-Kontorova model. *Physics Reports*, 306(1):1 – 108, 1998.

[43] P. M. Derlet, D. Nguyen-Manh, and S. L. Dudarev. Multiscale modeling of crowdion and vacancy defects in body-centered-cubic transition metals. *Phys. Rev. B*, 76:054107, 2007.

[44] Kazuhito Ohsawa and Eiichi Kuramoto. Thermally activated transport of a dislocation loop within an elastic model. *Journal of Nuclear Materials*, 367-370:327 – 331, 2007. Proceedings of the Twelfth International Conference on Fusion Reactor Materials (ICFRM-12).

[45] Yu. N. Osetsky, A. Serra, B. N. Singh, and S. I. Golubov. Structure and properties of clusters of self-interstitial atoms in fcc copper and bcc iron. *Philosophical Magazine A*, 80(9):2131–2157, 2000.

[46] T. D. Swinburne, S. L. Dudarev, and A. P. Sutton. Classical mobility of highly mobile crystal defects. *Phys. Rev. Lett.*, 113:215501, Nov 2014.

[47] K. Arakawa, K. Ono, M. Isshiki, K. Mimura, M. Uchikoshi, and H. Mori. Observation of the one-dimensional diffusion of nanometer-sized dislocation loops. *Science*, 318(5852):956–959, 2007.

[48] D R Mason, D Nguyen-Manh, and C S Becquart. An empirical potential for simulating vacancy clusters in tungsten. *Journal of Physics: Condensed Matter*, 29(50):505501, 2017.

[49] H. C. Kang and W. H. Weinberg. Dynamic Monte Carlo with a proper energy barrier: Surface diffusion and two-dimensional domain ordering. *The Journal of Chemical Physics*, 90(5):2824–2830, 1989.

[50] D R Mason, R E Rudd, and A P Sutton. Atomistic modelling of diffusional phase transformations with elastic strain. *Journal of Physics: Condensed Matter*, 16(27):S2679, 2004.

[51] Robert E. Rudd, D.R. Mason, and A.P. Sutton. Lanczos and recursion techniques for multiscale kinetic Monte Carlo simulations. *Progress in Materials Science*, 52(2):319 – 332, 2007. Modelling electrons and atoms for materials science.

[52] Céline Varvenne, Fabien Bruneval, Mihai-Cosmin Marinica, and Emmanuel Clouet. Point defect modeling in materials: Coupling ab initio and elasticity approaches. *Phys. Rev. B*, 88:134102, Oct 2013.

[53] S.L. Dudarev and A.P. Sutton. Elastic interactions between nano-scale defects in irradiated materials. *Acta Materialia*, 125:425 – 430, 2017.

[54] S.L. Dudarev, D.R. Mason, E. Tarleton, P.-W. Ma, and A.E. Sand. Multiscale model for stresses, strains and swelling of reactor components under irradiation. *Nuclear Fusion*, 2018.

[55] S. L. Dudarev and Pui-Wai Ma. Elastic fields, dipole tensors, and interaction between self-interstitial atom defects in bcc transition metals. *Phys. Rev. Materials*, 2:033602, Mar 2018.

[56] Gopinath Subramanian, Danny Perez, Blas P. Uberuaga, Carlos N. Tomé, and Arthur F. Voter. Method to account for arbitrary strains in kinetic Monte Carlo simulations. *Phys. Rev. B*, 87:144107, Apr 2013.

[57] S.P. Fitzgerald and D. Nguyen-Manh. Peierls potential for crowdions in the bcc transition metals. *Phys. Rev. Lett.*, 101:115504, 2008.

[58] C.S. Becquart, C. Domain, U. Sarkar, A. De Backer, and M. Hou. Microstructural evolution of irradiated tungsten: *Ab initio* parameterisation of an OKMC model. *J. Nucl. Mater.*, 403:75–88, 2010.

[59] D. Nguyen-Manh, S.L. Dudarev, and A.P. Horsfield. Systematic group-specific trends for point defects in bcc transition metals: An ab initio study. *Journal of Nuclear Materials*, 367-370:257 – 262, 2007. Proceedings of the Twelfth International Conference on Fusion Reactor Materials (ICFRM-12).

[60] J.N. Mundy, S.J. Rothman, H.A. Lam, N.Q.and Hoff, and L.J. Nowicki. Self-diffusion in tungsten. *Phys. Rev. B*, 18:6566–6575, 1978.

[61] A.B. Bortz, M.H. Kalos, and J.L. Lebowitz. A new algorithm for Monte Carlo simulation of Ising spin systems. *Journal of Computational Physics*, 17(1):10 – 18, 1975.

[62] M. A. Novotny. Monte Carlo algorithms with absorbing Markov chains: Fast local algorithms for slow dynamics. *Phys. Rev. Lett.*, 75:1424–1424, Aug 1995.

[63] Steve Plimpton. Fast parallel algorithms for short-range molecular dynamics. *Journal of Computational Physics*, 117(1):1 – 19, 1995.

[64] Thomas D Swinburne, Pui-Wai Ma, and Sergei L Dudarev. Low temperature diffusivity of self-interstitial defects in tungsten. *New Journal of Physics*, 19(7):073024, 2017.

[65] K. Nordlund. PARCAS computer code. The main principles of the molecular dynamics algorithms are presented in Nordlund et al, PRB 57:7556 (1998) and Ghaly et al, Phil Mag A79:795 (1999). The adaptive time step is the same as in Nordlund, Comp Mat Sci 3:448 (1995), 2006.

[66] C. Björkas and K. Nordlund. Comparative study of cascade damage in Fe simulated with recent potentials. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 259(2):853 – 860, 2007.

[67] H.J.C. Berendsen, J.P.M. Postma, W.F. van Gunsteren, A. DiNola, and J.R. Haak. Molecular dynamics with coupling to an external bath. *J. Chem. Phys.*, 81(8):3684–3690, 1984.

[68] J.F. Ziegler, J.P. Biersack, and U. Littmark. *The stopping and range of ions in solids.* Pergamon, 1982.

[69] A. De Backer, A. E. Sand, K. Nordlund, L. Luneville, D. Simeone, and S. L. Dudarev. Subcascade formation and defect cluster size scaling in high-energy collision events in metals. *EPL (Europhysics Letters)*, 115(2):26001, 2016.

[70] A.E. Sand, K. Nordlund, and S.L. Dudarev. Radiation damage production in massive cascades initiated by fusion neutrons in tungsten. *Journal of Nuclear Materials*, 455(1–3):207–211, 2014. Proceedings of the 16th International Conference on Fusion Reactor Materials (ICFRM-16).

[71] R. Alexander, M.-C. Marinica, L. Proville, F. Willaime, K. Arakawa, M. R. Gilbert, and S. L. Dudarev. Ab initio scaling laws for the formation energy of nanosized interstitial defect clusters in iron, tungsten, and vanadium. *Phys. Rev. B*, 94:024103, Jul 2016.

[72] M.-C. Marinica, F. Willaime, and N. Mousseau. Energy landscape of small clusters of self-interstitial dumbbells in iron. *Phys. Rev. B*, 83:094119, Mar 2011.

[73] Relaxation volumes of microscopic and mesoscopic irradiation-induced defects in tungsten. *submitted*, 2019.