



Cite this: *Mater. Adv.*, 2025, 6, 7379

Atomistic simulations of irradiation damage on the engineering timescale: examining the dose rate effect in tungsten

Max Boleininger,^{ID *a} Daniel R. Mason,^a Thomas Schwarz-Selinger^b and Pui-Wai Ma^a

The change in materials properties subjected to irradiation by highly energetic particles strongly depends on the irradiation dose rate. Atomistic simulations can in principle be used to predict microstructural evolution where experimental data is sparse or unavailable, however, fundamental limitations of the method make it infeasible to replicate the experimental timescale spanning from seconds to hours. Here, we present an atomistic simulation method where the motion of vacancies is accelerated, while the fast degrees of freedom are propagated with standard molecular dynamics. The resulting method is free of adjustable parameters and can predict microstructural evolution under irradiation at elevated temperatures. Simulating the microstructural evolution of tungsten under irradiation at dose rates of 10^{-5} , 10^{-4} , and 10^{-3} dpa s⁻¹, we find that increasing the temperature or reducing the dose rate primarily results in a reduction of the steady-state defect concentration, in qualitative agreement with deuterium retention and post-irradiation resistivity recovery experiments. The formation of a nanoscale void is observed if a system initially containing a large dislocation loop is irradiated. We present a minimally simple rate theory model which reproduces the time-dependent defect concentration and volume swelling behaviour obtained from the simulations.

Received 25th June 2025,
Accepted 22nd August 2025

DOI: 10.1039/d5ma00677e

rsc.li/materials-advances

1. Introduction

In recent years, atomistic simulations were used to great effect to predict changes in materials properties under irradiation.^{1–4} The simulation method is conceptually simple: atoms are assigned an initial position and momentum (for example, they may be arranged in a base-centred cubic lattice with velocities drawn from the Maxwell-Boltzmann distribution for a given temperature), and just like for any classical mechanical problem, the trajectories of atoms are individually propagated by solving Newton's equation of motion. In such molecular dynamics (MD) simulations, radiation damage is introduced to following a simple schema; when high-energy particles collide with atoms of the material, they may impart sufficient kinetic energy to displace them from their lattice site. If sufficient energy is transferred to the recoil atom, this leads to a self-similar cascade of ballistic secondary, tertiary, and so on, atomic collisions. The average number of times each atom in the system has been ballistically displaced is a common

measure for irradiation dose in the dpa unit (“displacements per atom”),⁵ and similarly, the dose rate is quantified by the dpa s⁻¹ unit.

Introducing the generation of radiation damage into atomistic simulations simply involves assigning individual atoms a large amount of kinetic energy and propagating the resultant trajectories in time.⁶ The various phenomena observed in the context of radiation damage simulations, such as the distinct phases of the collision cascade,⁷ the saturation of radiation damage at high dose,⁸ and the rapid relaxation of external stresses⁴ emerge from the collective behaviour of atoms over the course of propagation. Using this approach, the recent years saw the generation of irradiation microstructures in qualitative and quantitative agreement with experimental observations of materials properties such as thermal resistivity,² lattice strains,¹ vacancy concentration,⁸ and low-temperature irradiation creep.⁴ These were genuine predictions, in the sense that no adjustable parameters or fitting were required.

Some external input may be required, for example, SPECTRA-PKA⁹ can be used to generate a recoil spectrum so that atoms can be assigned recoil energies consistent with ion, electron, or neutron irradiation, while SRIM¹⁰ can be used to generate input for the electronic stopping model. These inputs affect the quantitative outcome of the simulations. The most

^a United Kingdom Atomic Energy Authority, Culham Campus, Abingdon, Oxfordshire, OX14 3DB, UK. E-mail: max.boleininger@uksaea.uk

^b Max Planck Institute for Plasma Physics, Boltzmannstr. 2, 85748, Garching, Germany



critical external input is the choice of the model for the interatomic forces. In large-scale atomistic simulations, these forces are typically described by empirical potentials, parameterised such that relevant physics emerging from atomic mechanisms are qualitatively reproduced. Recent years also saw the emergence of machine-learned interatomic potentials, which can offer forces with density-functional theory accuracy at a fraction of the computing time.¹¹

However, what gives atomistic simulations such predictive power—the explicit propagation of atomic trajectories—simultaneously leaves the method incapable of simulating the experimentally relevant timescales of seconds, minutes, days, and longer. Atoms in a metal vibrate at a frequency of around 10^{13} Hz, and so to resolve this oscillation, a propagation timestep order a hundredth of the oscillation period is needed. With contemporary high-performance computing, it is possible to propagate atomic trajectories on the order of a few billion steps, but even that results in a total simulated time of only a few microseconds; this is still far removed from the timescale of a tensile test, or the service lifetime of a fusion reactor component. Experimentally accessible irradiation dose rates span from 10^{-3} dpa s⁻¹ in ion-irradiation experiments to 10^{-7} dpa s⁻¹ for plasma-facing materials in a fusion reactor, necessitating simulated times on the order of hours to years. Under certain conditions, such as high dose rates or low irradiation temperatures, this presents no concern as microstructural evolution is primarily driven by the irradiation damage generation process, which is accessible on the simulated time-scale.⁸ However, simulations at elevated irradiation temperatures, where a significant amount of thermal relaxation is expected to occur, have conventionally been out of reach.

Coarse-grained methods, such as rate theory^{12–18} or object kinetic Monte Carlo^{19–29} have been employed to simulate the time-evolution of defect concentrations under various conditions. In these methods, the evolution of selected defect entities, such as point defects, dislocation loops, and voids, are simulated following a predefined set of rules. Although these coarse-grained methods can simulate microstructural evolution spanning the timescale from microseconds to hours, they require prior knowledge of the defect types emerging during evolution and an appropriate selection of input parameters describing properties of these defects. While quantitative predictions derived from these methods are highly sensitive to the choice of the aforementioned defect types and properties, they offer valuable insight into whether phenomena observed in experiments are consistent with simple dynamic models.

Here, we present a hybrid atomistic/coarse-grained method where faster degrees of freedom are handled by explicit time-propagation of atomic trajectories, while the motion of vacancies is accelerated using a kinetic Monte Carlo (kMC) scheme. The rationale behind this approach is that microstructural evolution of irradiated materials exhibits a separation of time-scales,³⁰ with interstitial-type defects diffusing many orders of magnitude faster than vacancies. While the fast degrees of freedom are evolved in a standard MD simulation, we intermittently accelerate the next rate-limiting mode of

microstructural evolution, which is vacancy migration. Complex, correlated atomic motion over short times is therefore handled with a high accuracy dynamic scheme (MD), and only the simpler vacancy migration is treated with the appropriate dynamics of kMC, as recommended by Mason *et al.*³¹ The resulting method can be used to predict microstructural evolution under irradiation at moderately elevated temperatures, where thermal diffusion of vacancies is sufficiently active to lead to thermal recombination of defects and formation of void clusters. While the method relies on specific assumptions on the timescales governing microstructural evolution, it is free of adjustable parameters. We apply the method to quantify the influence of temperature on the irradiation microstructure in tungsten at dose rates of 10^{-5} , 10^{-4} , and 10^{-3} dpa s⁻¹. The data is matched to a minimally simple rate theory model, allowing to predict the vacancy concentration and volume swelling as a function of temperature and dose rate valid up to doses of order 1 dpa, before the occurrence of break-away swelling. Current limitations of the method are also described.

We conclude with a comparison to experiment, which shows a qualitatively correct reduction in defect content at elevated temperature, with simulation slightly overestimating the degree of defect recombination. We attribute this to a lack of impurities in our model.

2. Theory

2.1. Rationale of the model

The accelerated simulation method is built on the assumption that irradiated microstructures evolve on multiple starkly different timescales. Consider the following illustrative simulation: an irradiated single-crystal tungsten, simulated by applying successive recoils up to 0.32 dpa at effectively cryogenic conditions, see ref. 4, is heated up. As the crystal is annealed, the irradiation microstructure coarsens and recovers, releasing its stored potential energy. The anneal begins with a temperature at absolute zero and proceeds with a rapid heating rate of 2.5×10^{11} K s⁻¹ until a temperature of 2500 K is reached. In Fig. 1 we plot the stored energy per atom relative to a perfect tungsten crystal at identical temperature over the course of the anneal, as well as the energy release rate over the anneal. The temperature scale is converted into an effective activation energy E_A using the expression:

$$E_A = k_B T W_0 \left(\frac{16\pi R c_0 T \nu_D}{a\alpha} \right), \quad (1)$$

where W_0 is the Lambert W function, $k_B = 8.61733 \times 10^{-5}$ eV K⁻¹ the Boltzmann constant, R the distance at which defects react, $c_0 = 0.31$ at% the initial vacancy concentration, $\nu_D = 8 \times 10^{12}$ Hz the Debye frequency for tungsten, and $\alpha = 2.5 \times 10^{11}$ K s⁻¹ the heating rate. We refer to Appendix A for a derivation of eqn (1).

While some recovery occurs at a low activation energy peaked at 0.1 eV, most of it occurs at a much higher activation energy peaked at 1.4 eV, consistent with the vacancy migration barrier of 1.53 eV for this potential, *c.f.* 1.73 eV for density-functional theory (DFT).³² As vacancies become mobile at



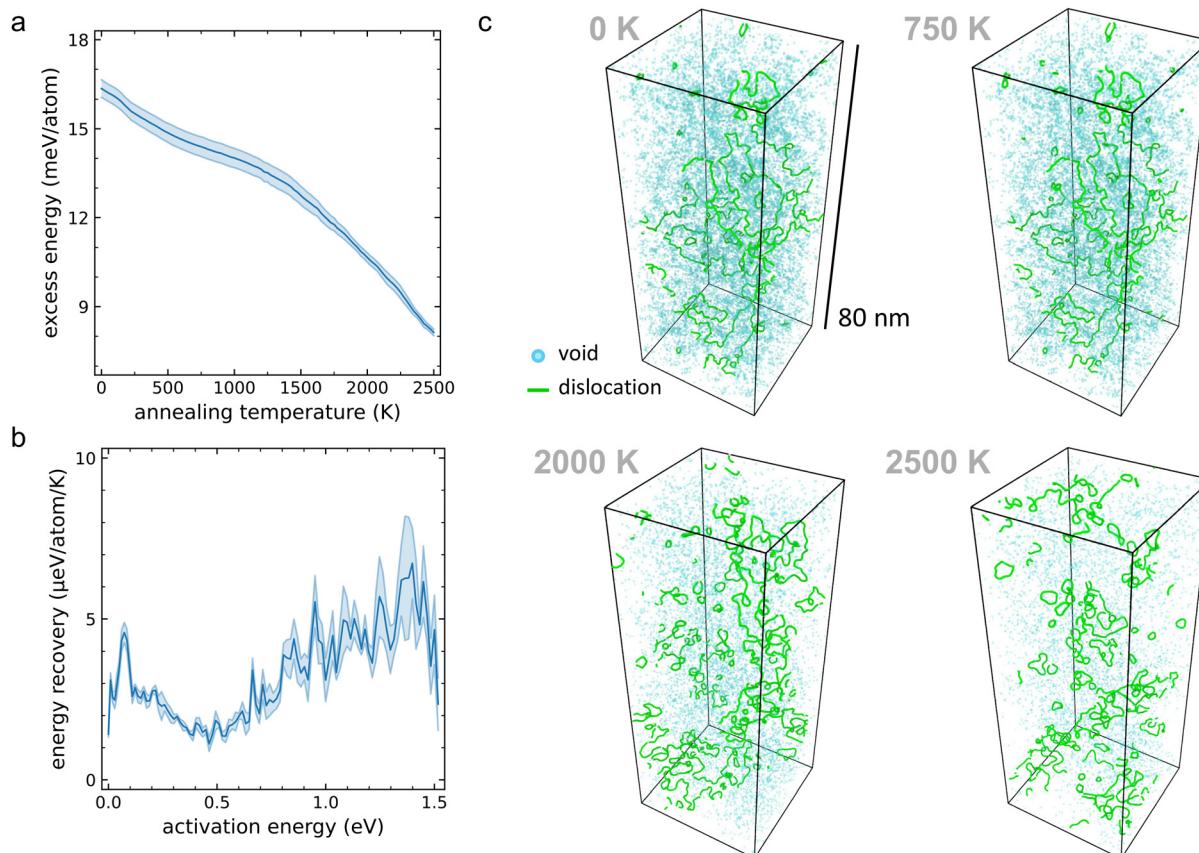


Fig. 1 Basic concepts of the recovery of irradiated tungsten illustrated by a rapid annealing simulation. (a) The microstructure of a tungsten crystal, here irradiated to 0.32 dpa, stores potential energy. As the crystal is heated, the microstructure evolves and coarsens, releasing the excess energy, shown here on top of the thermal energy. (b) The rate of energy release (bottom) allows inference on the activation energy by which the recovery proceeds. The simulated heating rate is unrealistically high with $2.5 \times 10^{11} \text{ K s}^{-1}$; in experiment, recovery is expected to occur at lower temperature. Shown here is the mean recovery over five simulations, each starting with a different 0.32 dpa microstructure. Shaded areas indicate the standard error. (c) Microstructural view of tungsten during post-irradiation annealing. Pictured are dislocation lines (green: $\frac{1}{2}(111)$) and void-type defects (blue). As monovacancies become mobile at around 1500 K at this heating rate, stage-III recovery commences: vacancies coalesce into vacancy-type dislocation loops and nanoscale voids, clearing up the microstructure. Shaded areas indicate the standard error over 5 independent annealing simulations.

higher temperature, they recombine with interstitials, thereby lowering the defect population in what is considered a second order reaction kinetics process.³³ In resistivity recovery experiments, this peak is known as stage-III recovery, typically occurring at a temperature of around 200 °C to 300 °C for tungsten, depending on the heating rate (*c.f.* 2000 °C in the simulation due to the rapid annealing rate used here).^{30,34}

Given that there is quite little recovery between the two identified recovery stages, we do not necessarily need to accelerate the time propagation of the microstructure for all thermally activated processes equally. At elevated temperature, for example 300 °C in tungsten, a monovacancy is expected to readily hop from lattice to lattice site with a rate of about 0.01 Hz. Over the timescale of hours, the vacancy degrees of freedom are expected to undergo significant evolution at this temperature. Degrees of freedom with a very low activation energy, for example 0.10 eV, would evolve with a comparatively enormous rate of 10¹² Hz at this temperature, and hence any meaningful evolution of these degrees of freedom concludes within a few picoseconds; this timescale is

easily accessible in a standard atomistic simulation, and hence does not need acceleration.

In what follows, we only consider simulations of tungsten due to its importance in fusion as the primary candidate material for plasma-facing components owing to its high melting temperature and high thermal conductivity;³⁵ but we note that the above discussion extends to most base-centred cubic (BCC) metals. DFT calculations find self-interstitial atom migration energies mostly in the order of meV,³⁶ or 10 meV for Cr, Mo, and W,³⁷ with notable exception of Fe with 0.34 eV^{36,38} due to the effect of magnetism. While the migration energy of vacancies in different materials can vary substantially, for the BCC transition metals it ranges from 0.5 to 2.0 eV,³² which confirms the general notion that self-interstitial atoms diffuse thermally at much lower homologous temperature than vacancies. This separation of timescales appears to be a general feature of BCC metals. While there may be processes active within the time-scale gap, these cannot be simply addressed within this single split picture. We present a method that is suited for describing the recovery of elemental metals with

irradiated microstructure, rather than a general time acceleration method.

2.2. Acceleration method

In order to propagate the vacancy system by a simulated time of t_{sim} , we follow the accelerated atomistic simulation method illustrated in Fig. 2. It proceeds as follows: First, the system is propagated with standard molecular dynamics for a short amount of time Δt_{MD} , a few picoseconds, to relax the degrees of freedom with low activation energy. Next, the system is analysed for monovacancies (as divacancies are not binding in tungsten,³⁹ they are treated as two neighbouring monovacancies). Finally, the vacancy system is propagated by a time Δt_{vac} by means of the rejection-sampling kinetic Monte Carlo (rkMC) algorithm,⁴⁰ with vacancy hopping rates depending on the local

stress to account for stress-biased diffusion. The vacancy propagation time is chosen such that the rejection rate averages 20%, as $\Delta t_{\text{vac}} \sim (5Z\nu)^{-1}$, where

$$\nu = \nu_D \exp\left(-\frac{E_m}{k_B T}\right) \quad (2)$$

is the hopping rate of a vacancy to one of its adjacent atomic sites ($E_m = 1.73$ eV), as estimated by Arrhenius' law, and $Z = 8$ is the number of vacancy hopping sites in a BCC crystal. These three steps are repeated for N_{acc} cycles until the total simulated time $t_{\text{sim}} = N_{\text{acc}}\Delta t_{\text{vac}}$ is reached. The rkMC algorithm is presented in more detail in the following section.

2.3. Implementation

The first step is the detection of vacancies in the microstructure. We note that it is crucial that the algorithm never misidentifies a larger vacancy cluster, *i.e.* size trivacancy or larger, as a mono or divacancy. Since divacancies in tungsten are not strongly bound, the nucleation of voids requires a rare coincidental meeting between three vacancies to form a trivacancy. If such a larger cluster is then misidentified as one or more monovacancies, the acceleration algorithm could dissociate the cluster by driving the vacancies apart again, thereby irrevocably disturbing the delicate clustering process. In addition, the vacancy detection algorithm needs to conclude within a runtime comparable to, or faster than, the runtime of the MD propagation step Δt_{MD} in order to keep the overhead of the algorithm small.

The void detection algorithm is based on reusing the (full) neighbour list from the MD simulation to detect empty space around each atom. The algorithm is outlined in Fig. 2. If an atom is adjacent to a void in a BCC metal, then there should be free space in the crystal around the first neighbour distance of $\sqrt{3}/2a$ to the atom, where $a = 3.165$ Å is the tungsten lattice constant. In practice, we sample points on a geodesic sphere of radius $\sqrt{3}/2a$ centred on each atom. If a point on the geodesic is further than $0.75 \times \sqrt{3}/2a$ to any other atom in the neighbour list of the central atom, then the central atom is adjacent to a void. Next, all void-adjacent atoms are clustered, and connectivity graphs for each cluster are built with a cut-off of $1.1a$. The cut-off is chosen as larger than a to improve the accuracy when thermal fluctuations are present. Next, the connectivity graphs are compared to precomputed connectivity graphs of the monovacancy and the two unique divacancy structures. If an isomeric match is detected, then the cluster of void-adjacent atoms is surrounding a monovacancy or a divacancy. In the case of a monovacancy, the monovacancy coordinate is placed in the centre of the cluster, and in the case of a divacancy, the vacancies are placed along the long axis determined by the second central moments of the cluster atom coordinates. This method can be applied to detect larger vacancy clusters by including more precomputed connectivity graphs—in practice, we found that enumeration of the unique number of vacancy cluster configurations up to size 5 was

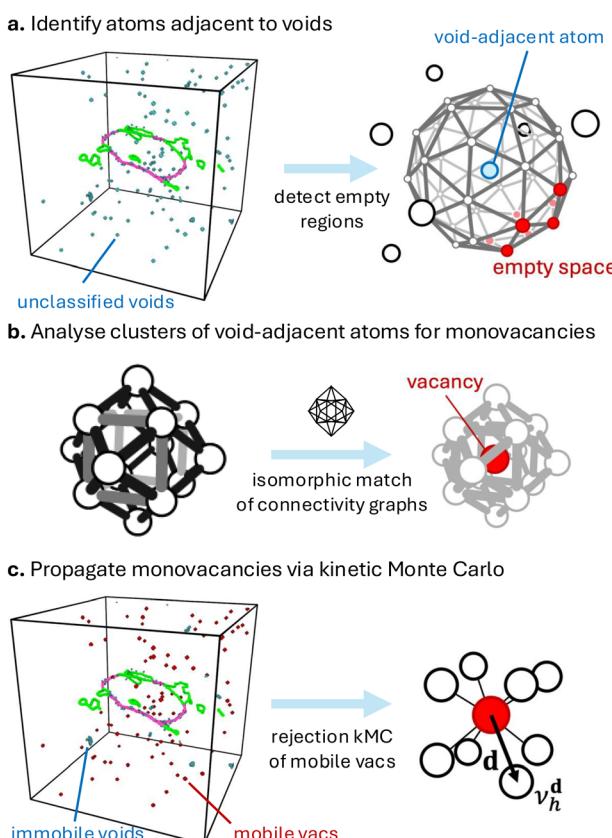


Fig. 2 Vacancy-acceleration algorithm. (a) Left: From the microstructure, shown are voids and dislocation lines, we identify atoms adjacent to void regions. Right: An atom is classified as adjacent to a void if any point (red) in a radius of $|\frac{1}{2}(111)|$ from the atom is further than $0.75 \times |\frac{1}{2}(111)|$ from any other atom. (b) Atoms adjacent to voids are clustered by proximity. For each cluster, a connectivity graph is constructed. If the graph is isomorphic to the precomputed graphs of the monovacancy or the two unique divacancy configurations, then the corresponding defect is identified. (c) As only monovacancies and divacancies are accelerated here, larger void structures are effectively immobile. For each vacancy, the hopping rates ν_h^d to the positions of its first neighbours \vec{d} are computed, accounting for the virial atomic stress at its current and target positions. The outcome of the random vacancy migration event is evaluated with a rejection kinetic Monte Carlo algorithm. Vacancy migration is considered a stochastically independent and spatially local process.



feasible. The specificity and detection accuracy is benchmarked in Appendix B.

Once the vacancies are detected, the next step is to propagate their positions with the rkMC algorithm. From the previous vacancy clustering step, we already have the list of potential hopping sites available. We now evaluate the list of hopping rates of each individual vacancy to their respective hopping sites. To account for the effect of elastic stress on the vacancy hopping rate, consider the hopping rate of a vacancy at position \vec{r}_0 to a neighbouring site $\vec{r}_0 + \vec{d}$. The hopping rate is estimated by the Kang–Weinberg model,^{31,41} with elastic stresses accounted for in the elastic dipole approximation.^{42–44}

$$\nu_h^{\vec{d}} = \nu_D \exp \left(-\frac{E_m + \text{tr} \left(\sigma_0 \cdot \Omega_0 - \frac{1}{2} (\sigma_{\vec{d}} + \sigma_{\vec{d}}) \cdot \Omega_{\text{sd}}^{\vec{d}} \right)}{k_B T} \right), \quad (3)$$

where σ_0 and $\sigma_{\vec{d}}$ are the virial atomic stress tensors at the vacancy and exchanging atom sites, respectively, and Ω_0 and $\Omega_{\text{sd}}^{\vec{d}}$ are the relaxation volume tensors of the vacancy at the ground state and saddle point, respectively. Here, the stress tensor at the saddle point is approximated by the average of the stress tensors at the vacancy and exchanging atom sites. The dipole approximation is precise provided the spatial scale of heterogeneity in the stress field is comparable to the size of the vacancy. Large stress gradients occur for example in dislocation cores, however, this is not a concern as a vacancy in a core will be absorbed during MD propagation.

In the following, we use the relaxation volume tensors of the vacancy as computed with DFT,³² such as Ω_0 :

$$\Omega_0 = \Omega_{\text{iso}}^{(\text{eq})} I_3 \quad (4)$$

where I_3 is the identity matrix of size 3 and $\Omega_{\text{iso}}^{(\text{eq})} = -0.115$ atomic volumes. The general expression for the relaxation tensor at the saddle-point for a hop in an arbitrary direction is:

$$\begin{aligned} \Omega_{\text{sd}}^{\vec{d}} = & \left(\Omega_{\text{iso}}^{(\text{sd})} + 2\Omega_{\text{aniso}}^{(\text{sd})} \right) \left(\hat{\vec{d}} \otimes \hat{\vec{d}} \right) \\ & + \left(\Omega_{\text{iso}}^{(\text{sd})} - \Omega_{\text{aniso}}^{(\text{sd})} \right) \left(\hat{\vec{u}} \otimes \hat{\vec{u}} + \hat{\vec{v}} \otimes \hat{\vec{v}} \right) \end{aligned} \quad (5)$$

where $\Omega_{\text{iso}}^{(\text{sd})} = -0.099$ and $\Omega_{\text{aniso}}^{(\text{sd})} = -0.094$ in atomic volumes, and vectors \vec{u} and \vec{v} are given by

$$\begin{aligned} \vec{u} = & \left(-\frac{d_y + d_x}{d_z}, 1, 1 \right) \\ \vec{v} = & \left(1, \frac{d_x^2 + d_z^2 + d_y d_z}{d_x(d_z - d_y)}, -\frac{d_x^2 + d_y^2 - d_y d_z}{d_x(d_z - d_y)} \right). \end{aligned} \quad (6)$$

The above expression is derived by rotating the relaxation volume tensor for a hop in [111] direction to an arbitrary direction \vec{d} . This formulation does not require prior identification of the crystal orientation, and can hence be applied to polycrystalline systems. To avoid numerical issues in evaluating

\vec{v} when $d_y = d_z$, we note that the singularity is removable in $\hat{\vec{v}}$:

$$\lim_{d_y \rightarrow d_z} \hat{\vec{v}} = \left(0, \mp \frac{d_x^2 + 2d_z^2}{\sqrt{2d_x^4 + 4d_x^2 d_z^2 + 4d_z^4}}, \pm \frac{d_x^2}{\sqrt{2d_x^4 + 4d_x^2 d_z^2 + 4d_z^4}} \right), \quad (7)$$

and that the tensor product $\hat{\vec{v}} \otimes \hat{\vec{v}}$ is invariant to the direction of the limit. The relaxation volume tensor is therefore well defined for any hopping direction.

In summary, evaluating the rates only requires the vacancy migration energy E_m and the relaxation volume tensors Ω_0 and $\Omega_{\text{sd}}^{\vec{d}}$ respectively in the ground state and at the saddle-point. These can be obtained from first principles methods, such as DFT, and represent constants fundamental to a given material. No adjustable parameters are featured in this method.

We note that eqn (3) for the vacancy hopping rate does not include any quasiharmonic or anharmonic effects. To account for these, we require first-principles calculations of the temperature-dependent migration barrier and relaxation volume tensors. Recent progress in simulating the anharmonic vacancy migration in tungsten with a machine-learned potential indicates that such data may be available in the near future.⁴⁵

Before each rkMC step, we extract the coordinates $\{\vec{r}_i\}$ and virial atomic stresses $\{\sigma_i\}$ of all atoms within a range of vacancies, from which we obtain the virial stress at an arbitrary position \vec{R} , which may for example be the vacancy or hopping site positions, by kernel density estimation:

$$\sigma_{\vec{R}} = \frac{\sum_i f(|\vec{R} - \vec{r}_i|) \sigma_i}{\sum_j f(|\vec{R} - \vec{r}_j|)}, \quad (8)$$

where we chose the kernel function

$$f(r) = \cos\left(\frac{\pi r}{2r_0}\right)^2 \Theta(r_0 - r), \quad (9)$$

where $\Theta(x) = 1$ if $x > 0$, otherwise 0, is the Heaviside function, and $r_0 = 5.8$ Å is a cutoff radius chosen between the fifth and sixth nearest neighbours. We note that the virial atomic stress is often reported in dimensions of pressure times volume, thereby requiring division by the atomic volume. Here, this is avoided by representing the relaxation volume tensors in units of atomic volume, resulting in the product of virial stress and relaxation volume returning the correct dimensions. The vacancy acceleration method is tested by comparison with standard molecular dynamics for the diffusion of a monovacancy with and without external stress in Appendix C.

The rkMC iteration is performed after the hopping rates $\{\nu_h^{\vec{d}}\}$ for all vacancies have been computed. In a given time interval t_{vac} , the expectation value for the number of hopping events of a given vacancy is $\langle n \rangle = t_{\text{vac}} \sum_h \nu_h^{\vec{d}}$. In a compromise between keeping the number of propagation steps low and reducing the probability of multiple hopping events occurring

within the rkMC propagation time step, we choose the time $t_{\text{vac}} = (5Z\nu)^{-1}$ such that $\langle n \rangle \approx 0.2$, where the net vacancy hopping rate is estimated by the hopping rate in a stress-free crystal. As we do not account for stress, the actual acceptance rate recorded in our simulations varies between 15% and 30%. For each vacancy, we generate a uniform random number $0 < r \leq 1$. If $r \leq 1 - \exp\left(-t_{\text{vac}} \sum_h v_h^d\right)$, then the vacancy undergoes a migration event, chosen from the set of possible hops, with probability weighted by the hopping rate. Otherwise, the vacancy remains at its site for this iteration. The migration event is implemented by creating an atom at the position of the vacancy and deleting the atom at the vacancy hopping site. If two vacancies are selected to move into the same atomic site, then only one of the two events is randomly chosen.

The method for simulating overlap cascades is described in detail in ref. 4. In order to apply the acceleration method to an irradiation damage simulation, we note that an atomic recoil increments the dose by $\Delta\phi = 0.8 T_d/(2E_d N_{\text{atoms}})$ in units of NRT dpa,⁵ where T_d is the damage energy evaluated from the recoil energy using the Lindhard model⁵ and $E_d = 90$ eV the threshold displacement energy for tungsten. Hence, for a given target dose rate of $\dot{\phi}$ in units of dpa s⁻¹, in between each atomic recoil, the vacancy subsystem needs to be propagated for a total time of $t_{\text{sim}} = \Delta\phi/\dot{\phi}$, or a total number of $N_{\text{acc}} = t_{\text{sim}}/\Delta t_{\text{vac}}$ rkMC steps.

The MD simulation are performed in LAMMPS,⁴⁶ which is compiled as a shared library and driven from Python⁴⁷ along with the vacancy detection and rkMC algorithms.

3. Results

We performed irradiation damage simulations to quantify the effect of dose rate on the defect concentration. We simulated the irradiation of an initially perfect, pure tungsten single crystal consisting of $80 \times 80 \times 80$ body-centred crystal unit cells ($25 \times 25 \times 25$ nm³), in total containing 1 024 000 atoms, with periodic boundary conditions applied in all three directions. We performed one simulation each per the following conditions of dose rate $\dot{\phi}$ and temperature T :

- $\dot{\phi} = 10^{-3}$ dpa s⁻¹, $T = 600, 650, 700, 750$, and 800 K
- $\dot{\phi} = 10^{-4}$ dpa s⁻¹, $T = 550, 600, 650, 700$, and 750 K
- $\dot{\phi} = 10^{-5}$ dpa s⁻¹, $T = 500, 550, 600, 650$, and 700 K

The temperature ranges were chosen to sample cascade-driven microstructural evolution at lower temperatures, and diffusion-driven evolution at higher temperatures, for a given dose rate. The dose rates were chosen to span multiple orders of magnitude whilst remaining accessible to ion-irradiation experiments for validation. To test for the effect of a sink, we also repeated the higher-temperature 10^{-4} dpa s⁻¹ dose rate simulations with a tungsten crystal initially containing a interstitial-type dislocation loop of 13 nm diameter with $\langle 100 \rangle$ burgers vector, corresponding to a dislocation density of 2.5×10^{15} m⁻².

In all simulations, a mono-energetic recoil spectrum with 10 keV was used, resulting in a dose increment of

3.5×10^{-5} dpa per recoil. A barostat was applied to keep the system-averaged stress to zero, using the approach by Shinoda *et al.*⁴⁸ implemented in Lammmps. The system was propagated for 10 ps after each atomic recoil to ensure that the collision cascade would conclude. Next, the fast degrees of freedom were relaxed for $t_{\text{MD}} = 1$ ps in between each of the N_{acc} rkMC steps. This process was repeated until the simulations reached a maximum dose of 0.3 dpa. The higher temperature simulations however were unable to reach this dose; at high temperature, the vacancies undergo thousands of migration events between each cascade, slowing down the simulation substantially.

The MD propagation time of 1 ps was chosen as a compromise between allowing fast degrees of freedom sufficient time to relax whilst keeping computational runtimes feasible. To test whether the propagation time was sufficient, we performed a simulation at 750 K and 10^{-4} dpa s⁻¹ (without loop) with a shorter propagation time of 0.1 ps. If the propagation time is sufficient, we expect the fastest degrees of freedom to conclude during the MD propagation phases between cascades. As a measure for whether the propagation time is sufficient, we consider the fastest defects in the system, *i.e.* self-interstitial crowdions:^{36,49} if the propagation time is sufficient, the crowdion concentration should remain vanishingly low throughout the simulation. Comparing the crowdion concentration at the highest simulated dose for this irradiation condition of 0.03 dpa, we found crowdion concentrations of 40 appm and 5 appm for 0.1 ps and 1 ps, respectively. A propagation time of 1 ps appears sufficient for the simulations considered here. In practice, the time required for the fast degrees of freedom to reach steady-state is expected to depend on features of the elastic energy landscape generated by the microstructure.

In Fig. 3, we plot the vacancy concentration as a function of dose and temperature for the three dose rates considered here, and show the resultant microstructures at the highest simulated doses for the dose rate of 10^{-4} dpa s⁻¹. We define 'vacancy concentration' as the void volume fraction, here shown in units of atomic percentage, or at%. The void content is quantified using the method developed by Mason *et al.*² Note that the x-axis is broken into two different linear scales to aid resolving the graphs at low dose. In Fig. 4, we compare the vacancy concentration with and without the initial $\langle 100 \rangle$ dislocation loop for the dose rate of 10^{-4} dpa s⁻¹ and show the final microstructures.

We make the following observations: At low temperature or high dose rate, the vacancy concentration saturates to a value of around 0.31 at%. At athermal conditions, when vacancies are effectively immobile in between cascade impacts, the vacancy concentration is driven to a recoil energy dependent steady-state.^{2,8} The saturation concentration of 0.31 at% observed here is lower than that found in athermal cascade simulations for the same tungsten potential, *c.f.* 0.34 at%⁵⁰ and (0.40 ± 0.03) at%,⁸ possibly due to the enhanced mobility of interstitial defects at higher temperature. We note that the aforementioned athermal simulations were performed in systems sized 340 000 and 21 296 000 atoms, respectively, demonstrating that

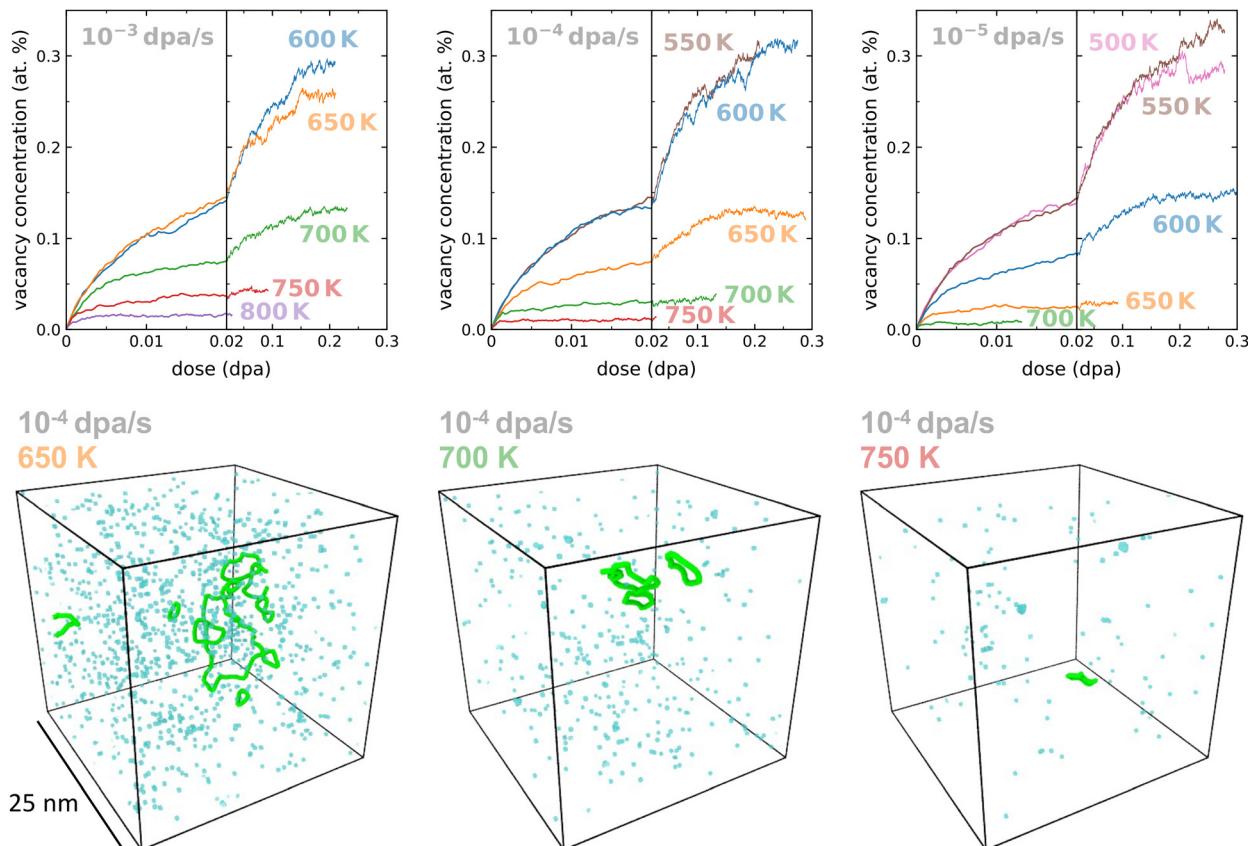


Fig. 3 Simulated irradiation of tungsten as a function of dose, dose rate and temperature. Top row: the vacancy concentration is found to decrease with increasing temperature and decreasing dose rate. Note the broken x-axis scale. Bottom row: simulation snapshots of the highest simulated doses for the $10^{-4} \text{ dpa s}^{-1}$ dose rate, 0.30, 0.13 and 0.03 dpa, at a temperature of 650, 700, and 750 K, respectively. Pictured are dislocation lines (green: $\frac{1}{2}\langle 111 \rangle$), here all of interstitial type, and void-type defects (blue). Some void clustering is observed at 750 K.

the system size chosen here is sufficient for a systematic study of changes in vacancy concentration.

The primary effect of increasing the temperature or decreasing the dose rate is a reduction in the steady-state defect concentration of the material. The steady-state concentration emerges from the balance between the continuous generation of radiation defects by high-energy atomic recoil and the continuous recombination between mobile vacancies and interstitials. The secondary effect is that vacancies cluster together, eventually forming small void clusters that remain stable under irradiation. The final effect is that the presence of the sink, here in the form of a $\langle 100 \rangle$ loop, leads not only to a higher defect concentration than in an otherwise identical simulation without a sink, but also to the formation of larger, nanoscale voids, see the 750 K microstructure in Fig. 4.

From the simulations with an initial dislocation loop, we observe that the loop acts as a sink for interstitial-type defects, thereby reducing the chances of vacancy-interstitial recombination, and giving each vacancy more time to find a void cluster to attach to. At the same time, the nature of the sink itself transforms over the course of irradiation, with the dislocation loop Burgers vector changing from $\langle 100 \rangle$ to $\frac{1}{2}\langle 111 \rangle$ as interstitial defects are absorbed.

4. Discussion

The primary effect, the increased defect recombination, can be rationalised by a minimal rate theory model. Consider first that the highly mobile interstitials generated by cascades cluster together into $\frac{1}{2}\langle 111 \rangle$ -type interstitial loops. As these defects only diffuse one-dimensionally in direction collinear to the Burgers vector, they are unlikely to find many vacancies to recombine with. Vacancies, on the other hand, diffuse three-dimensionally, and are hence more likely to encounter interstitial clusters to recombine with. Assuming that interstitial clusters are effectively sessile in the context of defect recombination, the differential equation describing time-evolution of the vacancy concentration $c(t)$ in units of atomic fraction, which is nominally equal to the Frenkel pair concentration, is written as a second order reaction kinetics expression:

$$\frac{dc(t)}{dt} = -4\pi D\eta R c^2(t) + \left(1 - \frac{c(t)}{c_{\text{sat}}}\right) g \dot{\phi}, \quad (10)$$

where $D = a^2\nu$ is the vacancy diffusion constant, $\eta = 2/a^3$ is the bcc atomic number density, R is the distance at which Frenkel pairs recombine, g is the vacancy generation efficiency in units of atomic fraction per dpa, and c_{sat} is the saturated vacancy

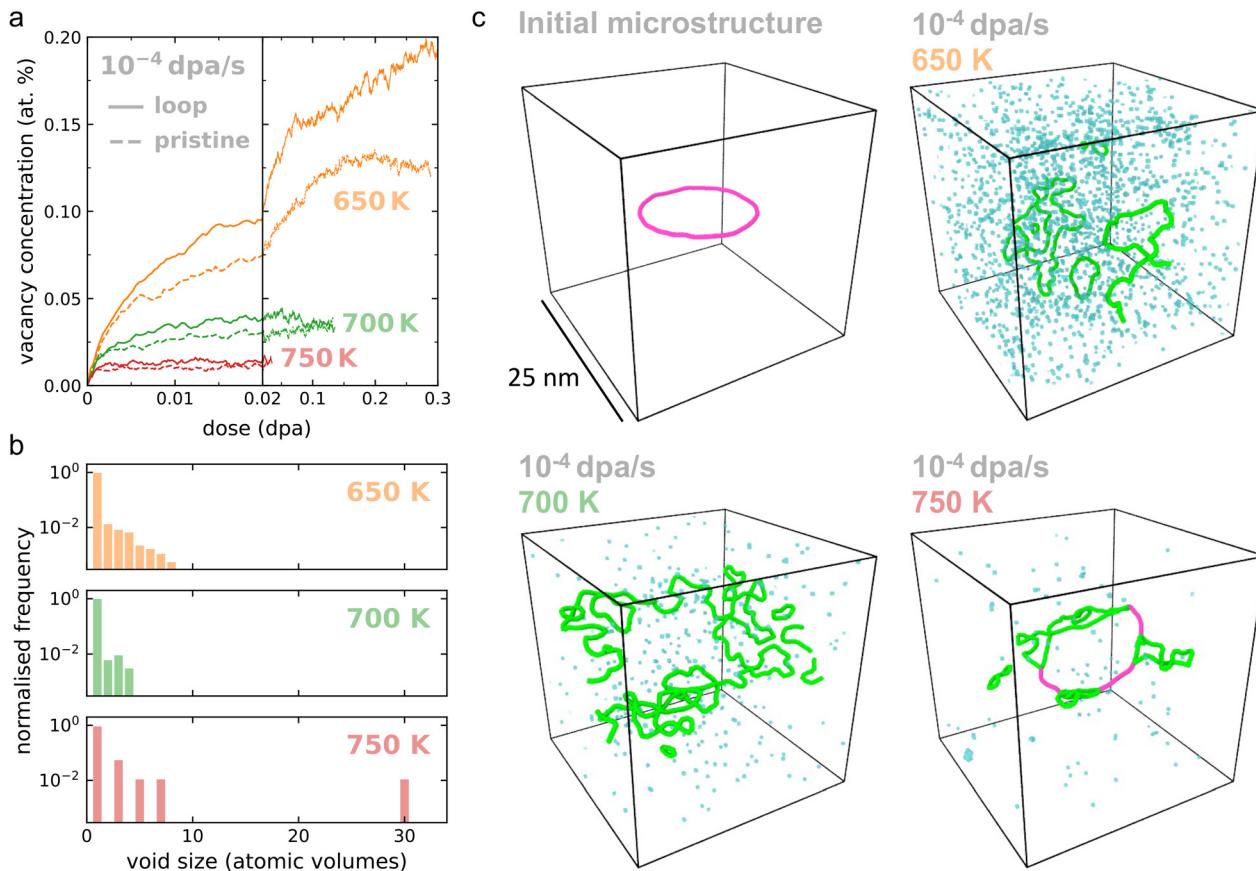


Fig. 4 Simulated irradiation of tungsten containing an initial $\langle 100 \rangle$ loop of 13 nm diameter as a function of dose and temperature. (a) The initial dislocation loop appears to act as a sink, absorbing interstitials preferentially, which suppresses Frenkel pair recombination and consequently leads to a higher vacancy concentration. (b) No void formation is observed in the 650 K and 700 K structures, while a void of diameter 1.3 nm forms in the 750 K structure. (c) Simulation snapshots of the initial microstructure and highest simulated doses at 0.30, 0.14 and 0.03 dpa, at a temperature of 650, 700, and 750 K, respectively. Pictured are dislocation lines (green: $\frac{1}{2}\langle 111 \rangle$, purple: $\langle 100 \rangle$) and void-type defects (blue). We find that the $\langle 100 \rangle$ loop is transformed to a $\frac{1}{2}\langle 111 \rangle$ loop as interstitials are absorbed.

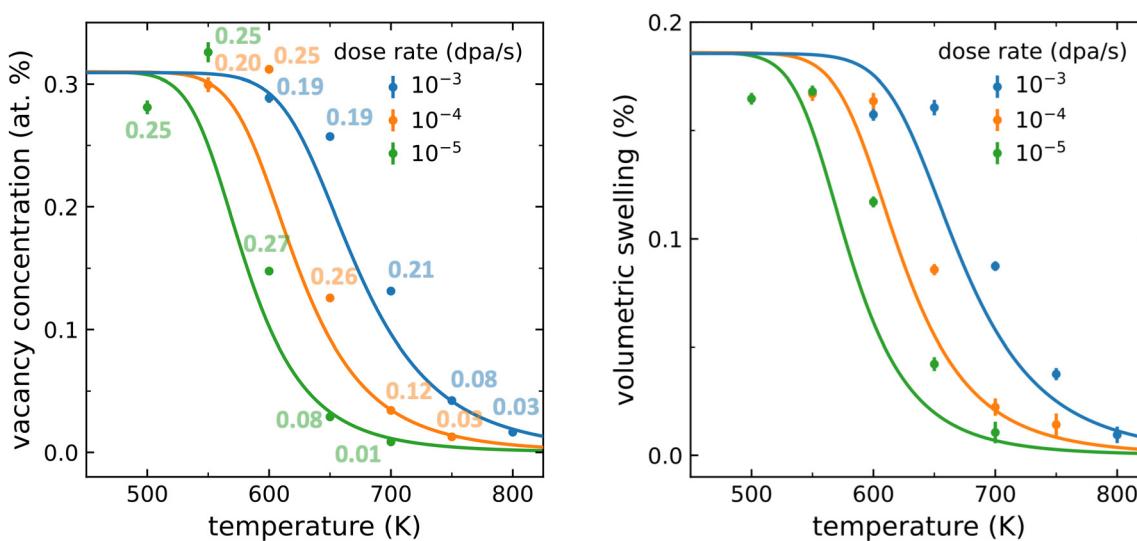


Fig. 5 Comparing prediction to simulation of irradiation swelling in tungsten. A minimally simple and analytical model of vacancy recombination (line), see text, is consistent with the accelerated irradiation damage simulation (dots), especially under consideration that the model contains no adjustable parameters. The data points and error bars represent the mean and standard deviation of the vacancy or swelling values taken over the latest 20% of the simulation snapshots, with points labelled by the respective mean dose (dpa). Here, a Debye frequency for tungsten of $\nu_D = 8 \times 10^{12}$ Hz was used. The thermal expansion strain was removed from the volumetric swelling strain shown here, see main text. The labels on each data point indicate the mean dose in dpa. While the doses appear low, the transient component of vacancy evolution appears already concluded, see Fig. 3.



concentration in the athermal limit in units of atomic fraction. The first term on the right-hand side drives recombination, and was derived assuming that interstitial and vacancy concentrations are equal. The second term is phenomenological, with the purpose of driving the vacancy concentration towards the athermal steady-state, where it saturates to c_{sat} .⁸ Parameters D and R are fundamental materials constants, obtainable from first-principles methods such as DFT, while g and c_{sat} depend on the material and cascade energy and can either be obtained from overlap cascade simulations or estimated using the arc-dpa model.⁵¹ For this tungsten potential, with 10 keV recoil energy cascades, we find $g = 11.25$ at% dpa⁻¹, $c_{\text{sat}} = 0.31$ at%, and choose $R = a$. For a constant dose rate and initial condition $c(0) = 0$, the above differential equation has the analytical solution:

$$c(t) = \frac{2c_{\text{sat}}}{1 + \sqrt{1 + \chi} \coth\left(\frac{g\dot{\phi}t}{2c_{\text{sat}}}\sqrt{1 + \chi}\right)}, \quad (11)$$

where

$$\chi = 16\pi D\eta R \frac{c_{\text{sat}}^2}{g\dot{\phi}} \quad (12)$$

For large times, the vacancy concentration asymptotically approaches a steady-state:

$$\lim_{t \rightarrow \infty} c(t) = \frac{2c_{\text{sat}}}{1 + \sqrt{1 + \chi}}. \quad (13)$$

In Fig. 5, we show a plot comparing the vacancy concentrations from the accelerated atomistic simulation and the rate theory model. The model is quantitatively consistent with simulation, predicting the extent by which the defect concentration is reduced as vacancies become more mobile at higher temperatures or lower dose rates. The simulated vacancy concentrations and volume swelling appear to be shifted up by about 20 K relative to the model at higher vacancy concentrations, suggesting that the recombination rate in the model is overestimated compared to simulation. This is expected, as the model is formulated under the assumption that the concentration of interstitials available for recombination is equal to the vacancy concentration. However, Vacancies can only recombine with interstitials that lie at the perimeter of defect clusters or dislocation loops, and therefore the recombination rate is overestimated especially for higher defect densities, where large interstitial-type dislocation networks are able to form.

Further, we note that the model can also be used as a minimally simple description of irradiation swelling by converting the defect concentrations into volumetric swelling ε_v via the relaxation volumes of the defects. Assuming that interstitials cluster together or join larger pre-existing defect structures, where their relaxation volume asymptotically approaches one atomic volume as the structure grows, and that vacancies appear in the form of either mono-vacancies or small voids, where their relaxation volume is approximately -0.4 atomic volumes,⁵² we find for the volumetric swelling strain:

$$\varepsilon_v \approx (1 - 0.4)c(t) = 0.6c(t) \quad (14)$$

The swelling model is compared to the volumetric swelling directly measured on the MD simulation box, see Fig. 5. The thermal expansion strain has been removed from the atomistic data for this comparison. This was done by performing MD of pristine tungsten in the isobaric-isoenthalpic ensemble to determine the temperature-dependent lattice constant $a(T)$. To obtain the volumetric swelling, the (volumetric) thermal expansion strain $\varepsilon^*(T) = 3(a(T)/a - 1)$ is subtracted from the total volumetric strain of the simulation cell, defined relative to the zero dose and zero temperature state.

We note that the above model predicts a saturation in vacancy concentration and swelling. Void swelling, that is, a measure for swelling obtained by counting nano-sized and larger voids under a transmission electron microscopy (TEM), is not a feature of the model because it is not accounted for in its formulation; the simulations have not progressed to a dose high enough to accurately quantify this phenomenon. Therefore, the model is expected to offer a useful description of volumetric swelling only up to intermediate dose of order 1 dpa.

The secondary effect, the occurrence of void clustering, is in broad agreement with TEM characterisation of tungsten self-ion irradiated at comparable temperature and identical dose rate.⁵³ Void nucleation requires temperatures sufficient for monovacancy diffusion to occur, and hence nanometre-sized voids do not form in standard athermal collision cascade simulations. Similarly, they are not observed in tungsten self-ion irradiated at room temperature under the TEM. The number concentration of voids appears lower than in experiment, however some degree of underestimation is expected as the simulations here do not contain any impurities; tungsten vacancies bind strongly to impurity elements such as C, N, and O, and could thereby catalyse the void nucleation process.⁵⁴

Some coarse-grained models, such as rate theory, introduce the notion of sinks,^{13,14} microstructural features acting as entities with endless capacity to remove point defects from the system. Here, we see that the presence of a large $\langle 100 \rangle$ dislocation loop causes an increase in the steady-state vacancy concentration. Dislocation loops have considerably larger elastic interactions with self-interstitials than with monovacancies, thereby preferentially absorbing self-interstitials over vacancies. The remaining vacancies find fewer interstitials to recombine with in their immediate vicinity, leading to a higher steady-state vacancy concentration. This observation is consistent with the sink bias⁵⁵ phenomenon in rate theory, which is considered the driving mechanism for void swelling.¹⁴ The transformation of the loop illustrates that sinks represented by dislocation networks are not immutable and bottomless, as the network topology and size changes over the course of irradiation. A detailed analysis of this phenomenon requires a dedicated study.

We conclude with a comment on the computational performance of this simulation method. While we succeeded in developing a simple model accounting for the effect of irradiation temperature on microstructural evolution, it is still



Table 1 Rate of migration events for a tungsten vacancy, as estimated with Arrhenius' law ($E_A = 1.73$ eV)

Temperature (K)	Vacancy hops per hour
300	5×10^{-13}
500	2×10^{-1}
700	2×10^4
900	8×10^6
1100	5×10^8

challenging to perform simulations up to high dose at high temperature. In the same vein that standard atomistic simulations cannot reach experimental timescales because the intrinsic frequency of atomic vibration needs to be numerically resolved, in the vacancy-accelerated simulation method, we find that at high temperature, the vacancies become so mobile that in a given time interval, they are required to undergo an enormous amount of individual migration events. To illustrate the point, we list the expected number of vacancy hops in one hour of real time as a function of temperature, as estimated with Arrhenius' law, see Table 1. As the hopping rate increases exponentially with temperature, any acceleration method that resolves individual migration events becomes computationally infeasible at high temperature.

For the acceleration scheme presented here, each rkMC cycle is accompanied by 5 MD propagation cycles of 1 ps duration each (note the factor 5 is due to the 20% rkMC acceptance rate). To propagate a single simulation up to one hour at 900 K would hence be equivalent to running on the order of 2×10^{10} MD time steps. Assuming no overhead, with current computing capabilities, it would require on the order of 10 000 GPU-hours[†] to propagate a 1 million atom system to one hour at 900 K. For comparison, this represents about 0.2% of the annual capacity of the GPU partition of the PITAGORA cluster (27 PFLOPS).⁵⁶ While it is feasible to run a few such simulations, in order to reach longer simulation times at higher temperatures, it is necessary to develop a method where vacancy migration is accelerated in a way that does not require resolution of individual migration events.

5. Comparison to experiments

For a number of reasons, it is not straightforward to draw quantitative comparisons between experimental measurements and predictions derived from these simulations. First of all, the simulations predict that the majority of voids appear in the form of small vacancy clusters, with exception of a single large void forming at 750 K when a sink is present, see Fig. 4. There are few microstructural characterisation techniques that are sensitive to small vacancy clusters. One such technique being positron-annihilation spectroscopy (PAS), which has recently been used in conjunction with predictions from first-principles simulations to construct quantitative vacancy cluster size distributions.⁵⁷ Instead, we shall focus on experimental studies

reporting on quantities that can act as indirect measures of defect concentration. We consider specifically ion-irradiation experiments, as neutron irradiation introduces additional complexity due to transmutation effects.⁵⁸

For our first comparison, we consider depth-resolved deuterium concentration measurements in irradiated tungsten^{59,60} as a proxy for vacancy concentration. In these experiments, tungsten was first subjected to self-ion-irradiation at a given temperature, and then subsequently exposed to a deuterium plasma at elevated temperature over multiple days, allowing the surface-implanted deuterium to diffuse through the damaged layer and bind to the irradiation defects. The temperature is elevated to promote deuterium diffusion, but still kept below stage-III recovery to avoid further microstructural evolution. The deuterium concentration can then be resolved as a function of depth using 3-He nuclear reaction analysis.

Experimental data was obtained for recrystallized, polycrystalline tungsten irradiated at 290 K and 800 K with 20.3 MeV tungsten ions at an average dose rate of 10^{-5} dpa s⁻¹, following the procedure outlined in Schwarz-Selinger *et al.*⁶¹ Here, however, deuterium loading was performed at 370 K instead of 450 K to maximise deuterium retention. Deuterium retention data is also available for the same two irradiation temperatures for single-crystalline samples with $\langle 100 \rangle$ ⁶² and $\langle 111 \rangle$ ⁶³ surface orientations, here irradiated with 10.8 MeV tungsten ions at an average dose rate of 2×10^{-4} dpa s⁻¹. We do not consider the difference in ion energy to affect the result, as the recoil spectra are nearly identical below the cascade fragmentation energy.⁶⁴ The data for the two surface orientations is plotted together, as the retained deuterium concentration does not appear to be significantly affected by the surface orientation.

In our simulation, we estimate the number of retained deuterium atoms directly from the void surface area Σ :

$$N_D = 5 \frac{\Sigma}{\Sigma_v}, \quad (15)$$

assuming that deuterium predominantly binds to internal surfaces, with the surface of monovacancy Σ_v trapping up to 5 deuterium atoms at room temperature,⁶⁵ and that the introduction of deuterium does not significantly alter the concentration of microstructural defects. The surface areas are computed using the void detection method presented by Mason *et al.*⁶⁶

The measured and simulated deuterium concentrations are shown in Fig. 6 as a function of dose for different temperatures. To allow for a direct comparison with the 800 K experimental irradiation temperature, we use simulation data for a 10^{-3} dpa s⁻¹ dose rate, noting that we would expect more defect recombination in a 800 K and 10^{-4} dpa s⁻¹ simulation. As we did not explicitly simulate irradiation at room temperature, we used here the 600 K data-set for comparison, labelled as 290 K with an asterisk; at this temperature and dose rate, vacancies undergo no thermal diffusion, and therefore the system evolves athermally.

[†] 5.26×10^8 atom step/s for an EAM potential on a H200 nvidia card, see <https://developer.nvidia.com/hpc-application-performance>, accessed 01/17/2025.



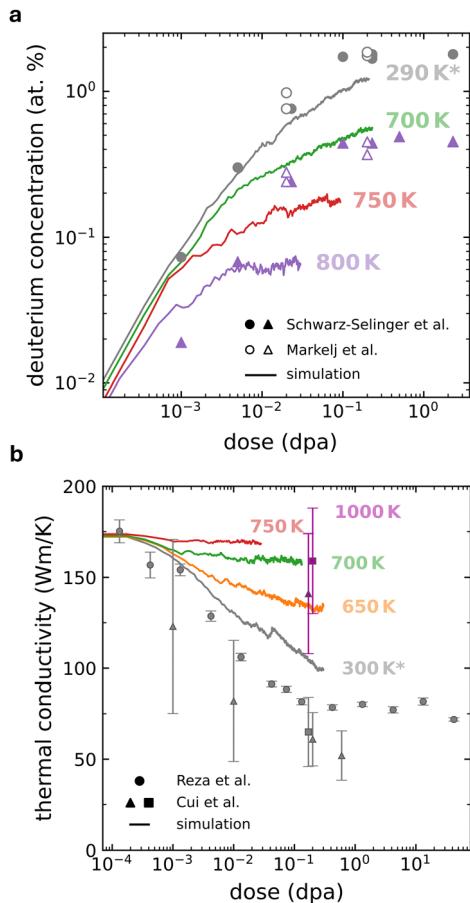


Fig. 6 Comparison to experiment. (a) Deuterium retention of tungsten irradiated at elevated temperature. Experimental data is shown for 290 K and 800 K, in polycrystalline (Schwarz-Selinger *et al.*⁶¹) and single-crystalline tungsten (Markelj *et al.*^{62,63}). At 290 K, experiment and simulation are in agreement. At 800 K, simulation underestimates the retained amount. (b) Thermal conductivity of irradiated tungsten. Experimental data is shown for 300 K (Reza *et al.*,⁶⁶ Cui *et al.*⁶⁷) and for 1000 K (Cui *et al.*⁶⁷). As before, experiment and simulation are in broad agreement for 300 K. At higher irradiation temperature, simulations underestimate defect concentration, resulting in an overestimation of thermal conductivity. Asterisks indicate that the simulated temperature is higher than stated, but representative for room temperature irradiation, see main text for details.

For room temperature irradiation, we find that the predicted deuterium concentration is in good qualitative and quantitative agreement with experiment, which corroborates the findings of a previous study,² where a hybrid method for simulating microstructural damage called *cascade annealing* was used, combining cascades and Frenkel pair insertion to reach higher doses. For irradiation at 800 K, simulation and experiment both show a reduction in deuterium concentration, however, simulation underestimates the retained amount by at least a factor of three. Another curious point is that similar deuterium retention experiments^{59,60} were performed on tungsten samples irradiated at average dose rates spanning three orders of magnitude, ranging from 5.8×10^{-3} to 3.9×10^{-6} dpa s⁻¹. However, each sample was found to retain a saturated deuterium concentration of about 0.5 at% irrespective of dose rate.

This effect is not reproduced in our simulations; our predicted steady-state vacancy concentration is dependent on both temperature and dose rate. The notable exception is at low temperature, where the vacancy concentration approaches the athermal steady-state of about 0.3 at%. We further note that irradiations were not performed with a continuous beam as recommended by ASTM standards⁶⁸ but beams were rastered with 1 kHz over the sample surface to achieve a laterally homogenous irradiation. As a consequence, the peak damage rate is two orders of magnitude larger than the average rate stated here. For the 290 K irradiation we expect rastering not to influence the defect density in this athermal regime but eventually, for a given dose rate at higher temperature and hence higher vacancy mobility, the increased peak dose rate may have an influence on the vacancy density. For our next comparison, we consider thermal conductivity measurements of ion-irradiated tungsten. The nanoscale defects generated by irradiation act as electron scattering sites, thereby lowering the electrical and thermal conductivities. As ion-irradiation damage is typically localised to within the first few microns of the sample, an accurate measurement of the change in thermal conductivity requires the use of carefully calibrated, surface-sensitive techniques, such as the 3- ω method⁶⁹ or transient grating spectroscopy (TGS).⁷⁰ In our simulations, we compute the thermal conductivity using the method developed by Mason *et al.*,⁶⁶ where atoms with a high potential energy are considered as scattering sites. We report thermal conductivities at room temperature.

Measurements of thermal conductivity in tungsten irradiated at higher temperature is sparse; data is available for 300 K^{67,71} and 1000 K.⁶⁷ We computed thermal conductivities for the 10^{-4} dpa s⁻¹ dataset, matching the dose rate of the experiments by Reza *et al.*⁷¹ Similar to the previous comparison, we find qualitative agreement between simulation and experiment at room temperature, with simulation overestimating the conductivity by about 30% at higher doses. A previous simulation study² using the cascade annealing method found closer quantitative agreement. At high temperature, we find that simulation and experiment both show an increase in thermal conductivity, which can be attributed to a reduction in scattering sites as fewer crystal defects are present at high temperature. The simulations predict a near full recovery of thermal conductivity by 750 K, while the experiments indicate that recovery might not be quite as complete at 1000 K—however, the reported uncertainties are large, impeding a quantitative comparison.

Finally, we note that the choice of the recoil energy in our simulations represents a source of uncertainty. At athermal conditions, it was shown that the recoil energy directly affects the steady-state vacancy concentration.⁸ Cascade simulations with recoil energies drawn from the 20.3 MeV self-ion-irradiation recoil spectrum, up to the fragmentation energy of 100 keV,⁷² generate a steady-state vacancy concentration of (0.31 ± 0.01) at%, which is close to the vacancy concentration resulting from 10 keV recoils. While the overall defect concentration appears consistent between these two irradiation

conditions, it remains an open question whether the direct generation of voids by high-energy recoils strongly affects the void clustering statistics at higher temperature.

To summarise, these comparisons suggest that the vacancy acceleration method overestimates the rate of vacancy and interstitial recombination compared to experiments in real materials. Furthermore, TEM characterisation of samples irradiated at comparable temperature and dose rate⁵³ suggest that the simulation method underestimates the number density of nanoscale-sized voids. A possible explanation for these two discrepancies might lie in the neglect of impurities in the simulated microstructure. Light element impurities, such as C, N, and O, can trap vacancies, and thereby suppress defect recombination whilst catalysing the void nucleation process.⁵⁴ Vacancies immobilised by impurities do not contribute to vacancy-driven defect recombination, which would increase the defect concentration compared to a pure material. For comparison, consider that even the extra-high purity tungsten samples (99.9999 wt%) used by Hu *et al.*⁵³ could contain up to ~ 10 appm of light element impurities, which would translate to ~ 10 impurity atoms in a 1 million atom system. In conclusion, while the effect of impurities on microstructural evolution is missing in the simulations performed here, it is feasible to investigate this in a future study.

6. Conclusion

We developed a model that enables simulation of microstructural evolution under irradiation and simultaneous thermally driven recovery. This is achieved by accelerating specifically vacancy migration with a kinetic Monte Carlo algorithm, while the rest of microstructural evolution, including collision cascades and defect clustering, proceeds implicitly by standard molecular dynamics simulation. We discussed the rationale behind the methodology and outlined the algorithm. We studied the effect of the dose rate on the irradiation defect concentration in tungsten and presented a simple predictive rate theory model that matches the simulation result. We made three observations: First, the defect concentration is reduced as either the temperature increases or the dose rate decreases. Second, at higher temperatures, vacancies begin clustering into nanoscale voids. Third, the presence of a large dislocation loop in the initial microstructure leads to a higher defect concentration as it acts as a sink for primarily interstitial defects. These predictions are in qualitative agreement with experimental observations, however, an assessment of the quantitative accuracy is impeded by low volume of simulation data and sparsity of quantitative experimental data. The model appears to overestimate vacancy-defect recombination; a possible explanation is the present neglect of impurities in the model.

We conclude that a hybrid molecular dynamics/Monte Carlo method is a viable option for exploiting the gap in diffusion time-scales between interstitial-type defects and vacancies. The results of the model are consistent with the assertion that vacancies are the primary drivers of defect recombination,

rather than interstitials. The method allows for an explicit and calibration-free description of sinks, as demonstrated here by inclusion of a large dislocation loop, and the simulation of individual vacancy migration events allows for the inclusion of impurity effects in a future study. However, we note that the simulation method becomes computationally impractical at higher temperatures due to the exponentially increasing number of vacancy migration events that need to be simulated in a given period of time.

Conflicts of interest

There are no conflicts to declare.

Data availability

Atomistic simulation snapshots and the data constituting the figures are available at DOI: <https://doi.org/10.5281/zenodo.17091685>.

Appendix

A Activation energy

In order to determine the activation energy E_A of a recovery process, typically a dedicated set of isothermal annealing experiments would be undertaken.⁷³ Here, we assume that the recovery process follows second order reaction kinetics, specifically through the reaction of a particular type of mobile defect with itself. In that case, we expect the defect concentration $c(t)$ in units of atomic fractions in an annealing experiment with temperature ramping to follow the differential equation:

$$\frac{dc(t)}{dt} = -8\pi D\eta R c^2(t) \quad (\text{A.1})$$

with initial condition $c(0) = c_0$, where a is the tungsten lattice constant, D is the diffusion constant, $\eta = 2/a^3$ is the bcc atomic number density, and R is the distance at which defects react. We consider here the case where the defects undergo 3-dimensional diffusion by moving *via* $1/2\langle 111 \rangle$ crystal sites, leading to $D = a^2 \nu_D \exp[-E_A/(k_B T)]$. If the temperature starts at absolute zero and increases with time with constant rate, $T = \alpha t$, eqn (A.1) has the solution:

$$c(T) = \frac{c_0}{1 + 8\pi a^2 R \eta c_0 \nu_D \left[\frac{T}{\alpha} \exp\left(-\frac{E_A}{k_B T}\right) + \frac{E_A}{k_B \alpha} \text{Ei}\left(-\frac{E_A}{k_B T}\right) \right]}, \quad (\text{A.2})$$

where Ei is the exponential integral function. To determine the temperature T^* of the recovery peak, we need to find the root

$$\left. \frac{\partial^2 c(T)}{\partial^2 T} \right|_{T=T^*} = 0. \quad (\text{A.3})$$

As activation energies of interest are typically high relative to $k_B T$, we can find the approximate root by series expansion of the second order derivative with respect to x , where $x = k_B T/E_A$, to



second order at $x = 0$. Solving T^* for E_A , we find an expression relating the activation energy to the temperature of the recovery peak:

$$E_A = k_B T^* W_0 (8\pi a^2 R \eta c_0 T^* \nu_D \alpha^{-1}). \quad (\text{A.4})$$

B. Vacancy detection validation

Here, we test the specificity and accuracy of the vacancy detection algorithm. We apply the algorithm to a tungsten simulation cell containing a void volume fraction of 0.01%, 0.1%, and 1%. The void structures are generated by iteratively deleting atoms in randomly placed spheres with radii R drawn from the exponential distribution $R \sim \text{Exp}(\lambda)$, here $\lambda = 1$ until the target void volume fraction is reached. The structures are analysed for monovacancies and divacancies using Wigner–Seitz Analysis, representing the ground truth data. Example void structures are shown in Fig. 7.

Next, to introduce thermal noise consistent with a system temperature T , each atomic coordinate is perturbed by a random displacement vector \vec{u} , where each component is drawn from a normal distribution with zero mean and variance of

$$\langle u_i \rangle^2 = aT + bT^2, \quad (\text{B.1})$$

where $a = 6.638 \times 10^{-6} \text{ \AA}^2 \text{ K}^{-1}$ and $b = 2.710 \times 10^{-9} \text{ \AA}^2 \text{ K}^{-2}$. Parameters a and b were fitted to reproduce the mean-squared

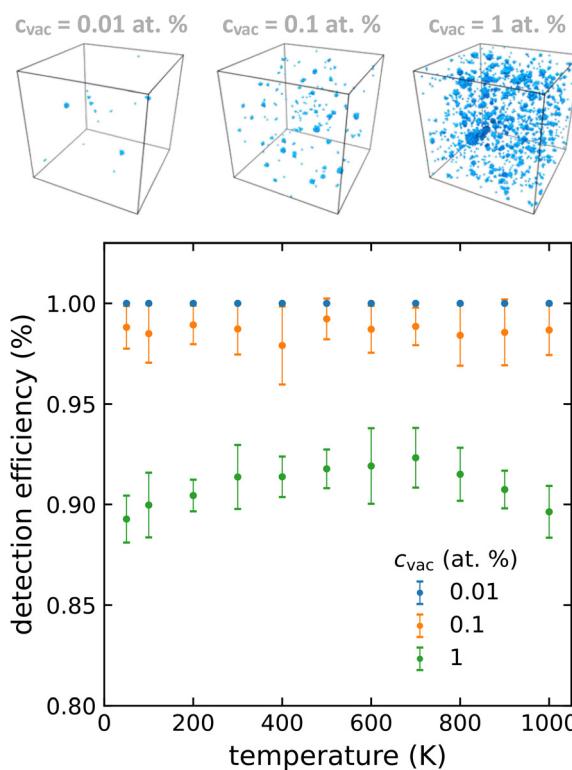


Fig. 7 Vacancy detection validation. Top: Microstructures with varying void volume fractions are analysed for monovacancies and divacancies as a function of temperature. Bottom: The fraction of detected over nominal vacancies, the detection efficiency, is found to be insensitive to thermal noise.

atomic displacement of BCC tungsten atoms at finite temperature but can also, in the linear approximation, be estimated from the Debye temperature $T_D = 400 \text{ K}$ ⁷⁴ via the Debye–Waller factor:⁷⁵

$$\langle u_i \rangle^2 = \frac{3\hbar^2}{k_B M} \frac{T}{T_D^2} = 8.249 \times 10^{-6} \frac{\text{\AA}^2}{\text{K}} T \quad (\text{B.2})$$

where $M = 183.84 \text{ u}$ is the atomic mass of tungsten.

Next, we run the vacancy detection algorithm on the perturbed structures, and test if the detected vacancies are consistent with the reference vacancies. We perform this analysis for temperatures spanning 50 K to 1000 K, for 10 independent perturbed structures per temperature, with results shown in Fig. 7. We find that the vacancy detection algorithm is unaffected by thermal noise for the temperatures tested here. The vacancy detection algorithm becomes less accurate in detecting monovacancies and divacancies as the void volume fraction increases.

Further, we find that no vacancies were misclassified; there were no false positive detections of vacancies, and all of the identified vacancy clusters were identified correctly compared to the Wigner–Seitz ground truth. We note that at the highest tested void volume fraction, about 10% of the vacancies are not detected, suggesting that vacancies might be diffusing about 10% slower than they should. Solving Arrhenius' rate for the effective temperature, this slow-down effectively corresponds to a reduction in the simulated temperature on the order of 5 K. This error is comparatively small compared to the error due to the neglect of anharmonic effects in the vacancy hopping rate, see Appendix 8, and hence we shall not discuss it further.

C. Vacancy acceleration validation

The rkMC algorithm is validated by simulating the diffusion of a monovacancy without external stress, and with an external shear stress of $\sigma_{xy} = 3 \text{ GPa}$. Using standard MD, we sample the diffusion tensor over temperatures spanning 1000 K to 1700 K. Using the vacancy-accelerated method, we sample temperatures spanning 300 K to 900 K. To reconstruct the diffusion tensor in the MD simulation, we track the vacancy hops using an on-the-fly Wigner–Seitz Analysis method. In the accelerated simulation, we track the successful rkMC hopping events.

The simulated diffusion tensor is also compared to the analytical diffusion tensor. For a monovacancy at position \vec{r}_0 in a BCC crystal, it is given by

$$D = \frac{1}{6} \sum_{h=1}^8 \left(\vec{d}_h \otimes \vec{d}_h \right) \nu_h^{\vec{d}}, \quad (\text{C.1})$$

where the summation runs over the eight unique $\frac{1}{2}\langle 111 \rangle$ adjacent hopping sites placed at $\vec{r}_0 + \vec{d}_h$, and $\nu_h^{\vec{d}}$ is the stress-dependent hopping rate, see eqn (3). For the comparison of the analytical model with molecular dynamics, we note that expression eqn (3) does not account for vibrational entropy, which, in the harmonic approximation, acts to effectively lower the migration barrier $E'_m = E_m - TS$.⁷⁶ For vacancy migration with



this interatomic potential,⁷⁷ we find $S = 1.8 k_B$. We evaluate the model diffusion tensor including and excluding the entropic contribution.

Because we are validating the analytical model and accelerated simulations in reference to explicit molecular dynamics, we use the vacancy migration barrier and relaxation volume tensors consistent with the chosen interatomic potential.⁷⁷ We determined the migration barrier using the nudged elastic band method,⁷⁸ and the relaxation volume tensors following the method outlined in ref. 79: $E_m = 1.523$ eV, $\Omega_{\text{iso}}^{(\text{eq})} = -0.1221$, $\Omega_{\text{iso}}^{(\text{sd})} = -0.1897$, and $\Omega_{\text{aniso}}^{(\text{sd})} = -0.0800$ in atomic volumes.

The simulation results are shown in Fig. 8. The results from the molecular dynamics and accelerated methods are in close agreement to the model prediction with and without the entropic contribution, respectively. In the absence of external

shear stress, the dipole tensor is isotropic. When a shear stress of $\sigma_{xy} = -3$ GPa is included, the diffusion tensor gains off-diagonal elements $D_{xy} = D_{yz}$ with magnitude comparable to the diagonal elements, demonstrating that elastic stress is accounted for in the vacancy acceleration method.

Acknowledgements

We acknowledge useful discussions with Felix Hofmann, Thomas Jourdan, and Sabina Markelj. This work has been carried out within the framework of the EUROfusion Consortium, funded by the European Union *via* the Euratom Research and Training Programme (Grant Agreement No 101052200—EUROfusion) and from the EPSRC [grant number EP/W006839/1]. To obtain further information on the data and models underlying this paper please contact PublicationsManager@ukaea.uk. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them. The authors acknowledge the use of the Cambridge Service for Data Driven Discovery (CSD3) and associated support services provided by the University of Cambridge Research Computing Services (<https://www.cs3.cam.ac.uk>) in the completion of this work.

References

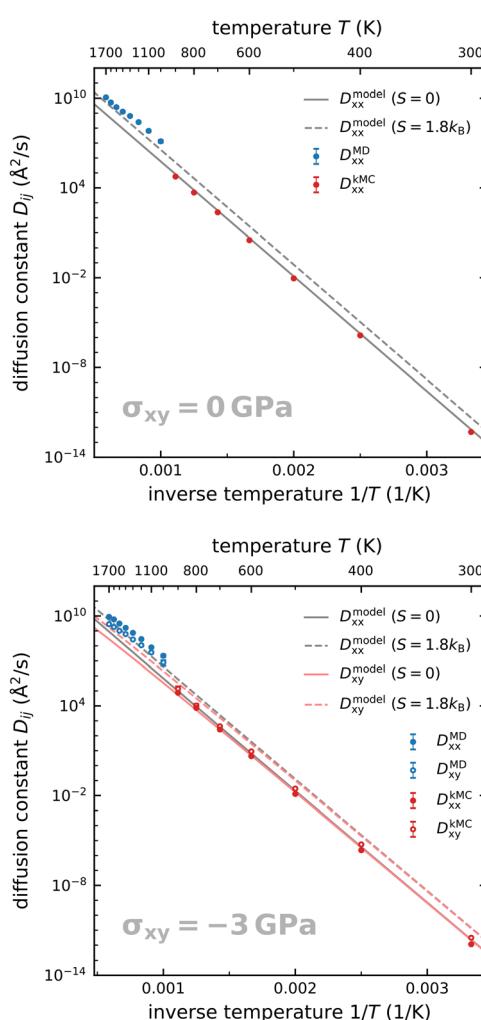


Fig. 8 Kinetic Monte Carlo validation. Shown are the diffusion tensors of a tungsten monovacancy, with and without external stress, obtained from standard molecular dynamics (D_{xx}^{MD}), from the vacancy acceleration method (D_{xx}^{KMC}), and from the analytical model (D_{xx}^{model}). The molecular dynamics and accelerated method results are consistent with the model predictions with ($S = 1.8 k_B$) and without ($S = 0$) the entropic contribution to the migration barrier, respectively. Error bars indicating the standard error are too small to be visible on this axis.

- 1 D. R. Mason, S. Das, P. M. Derlet, S. L. Dudarev, A. J. London, H. Yu, N. W. Phillips, D. Yang, K. Mizohata, R. Xu and F. Hofmann, *Phys. Rev. Lett.*, 2020, **125**, 225503.
- 2 D. R. Mason, F. Granberg, M. Boleininger, T. Schwarz-Selinger, K. Nordlund and S. L. Dudarev, *Phys. Rev. Mater.*, 2021, **5**, 095403.
- 3 C. A. Hirst, F. Granberg, B. Kombaiah, P. Cao, S. Middlemas, R. S. Kemp, J. Li, K. Nordlund and M. P. Short, *Sci. Adv.*, 2022, **8**, eabn2733.
- 4 A. Feichtmayer, M. Boleininger, J. Riesch, D. R. Mason, L. Reali, T. Höschen, M. Fuhr, T. Schwarz-Selinger, R. Neu and S. L. Dudarev, *Commun. Mater.*, 2024, **5**, 218.
- 5 M. J. Norgett, M. T. Robinson and I. M. Torrens, *Nucl. Eng. Des.*, 1975, **33**, 50–54.
- 6 F. Granberg, K. Nordlund, M. W. Ullah, K. Jin, C. Lu, H. Bei, L. M. Wang, F. Djurabekova, W. J. Weber and Y. Zhang, *Phys. Rev. Lett.*, 2016, **116**, 135504.
- 7 A. F. Calder, D. J. Bacon, A. V. Barashev and Y. N. Osetsky, *Philos. Mag.*, 2010, **90**, 863–884.
- 8 M. Boleininger, D. R. Mason, A. E. Sand and S. L. Dudarev, *Sci. Rep.*, 2023, **13**, 1684.
- 9 M. R. Gilbert, J. Marian and J.-C. Sublet, *J. Nucl. Mater.*, 2015, **467**, 121–134.
- 10 J. F. Ziegler, *Nucl. Instrum. Methods Phys. Res., Sect. B*, 2004, **219**, 1027–1036.



11 A. M. Goryaeva, J. Dérès, C. Lapointe, P. Grigorev, T. D. Swinburne, J. R. Kermode, L. Ventelon, J. Baima and M.-C. Marinica, *Phys. Rev. Mater.*, 2021, **5**, 103803.

12 L. Glowinski, C. Fiche and M. Lott, *J. Nucl. Mater.*, 1973, **47**, 295–310.

13 L. Mansur, *J. Nucl. Mater.*, 1978, **78**, 156–160.

14 A. Brailsford and R. Bullough, *J. Nucl. Mater.*, 1972, **44**, 121–135.

15 J. E. Westmoreland, J. A. Sprague, F. A. Smidt Jr. and P. R. Malmberg, *Radiat. Eff.*, 1975, **26**, 1–16.

16 D. Xu, G. VanCoevering and B. D. Wirth, *Comput. Mater. Sci.*, 2016, **114**, 47–53.

17 F. Li, C. Chen, D. Guo, X. Zhou, Y. Chen, Y. Long, L. Guo, L. Li, Q. Ren, Y. Liao, T. Liu and R. Shu, *J. Alloys Compd.*, 2021, **874**, 159751.

18 P. Saidi, M. Topping, C. Dai, F. Long, L. K. Béland and M. R. Daymond, *J. Nucl. Mater.*, 2021, **543**, 152478.

19 Q. Xu, H. L. Heinisch and T. Yoshiie, *J. Comput.-Aided Mater. Des.*, 1999, **6**, 215–223.

20 M. Chiapetto, C. Becquart and L. Malerba, *Nucl. Mater. Energy*, 2016, **9**, 565–570.

21 J. Li, C. Zhang, Y. Yang, T. Wang and I. Martin-Bragado, *J. Nucl. Mater.*, 2022, **561**, 153529.

22 C. Domain, C. Becquart and L. Malerba, *J. Nucl. Mater.*, 2004, **335**, 121–145.

23 I. Martin-Bragado, A. Rivera, G. Valles, J. L. Gomez-Selles and M. J. Caturla, *Comput. Phys. Commun.*, 2013, **184**, 2703–2710.

24 N. Castin, A. Bakaev, G. Bonny, A. Sand, L. Malerba and D. Terentyev, *J. Nucl. Mater.*, 2017, **493**, 280–293.

25 A. B. Sivak, V. A. Romanov and V. M. Chernov, *Crystallogr. Rep.*, 2010, **55**, 97–108.

26 A. Sivak, V. Chernov, V. Romanov and P. Sivak, *J. Nucl. Mater.*, 2011, **417**, 1067–1070.

27 M. J. Caturla, *Comput. Mater. Sci.*, 2019, **156**, 452–459.

28 F. Jiménez and C. Ortiz, *Comput. Mater. Sci.*, 2016, **113**, 178–186.

29 M. Mock, P. Stein, C. Hin and K. Albe, *J. Nucl. Mater.*, 2019, **527**, 151807.

30 L. Keys, J. Smith and J. Motteff, *Phys. Rev.*, 1968, **176**, 851.

31 D. R. Mason, A. E. Sand and S. L. Dudarev, *Modell. Simul. Mater. Sci. Eng.*, 2019, **27**, 055003.

32 P.-W. Ma and S. L. Dudarev, *Phys. Rev. Mater.*, 2019, **3**, 063601.

33 M. Anand, B. Pande and R. Agarwala, *Radiat. Eff.*, 1978, **39**, 149–155.

34 A. Reza, G. He, C. A. Dennett, H. Yu, K. Mizohata and F. Hofmann, *Acta Mater.*, 2022, **232**, 117926.

35 C. Luo, L. Xu, L. Zong, H. Shen and S. Wei, *Fusion Eng. Des.*, 2023, **190**, 113487.

36 P.-W. Ma and S. L. Dudarev, *Phys. Rev. Mater.*, 2019, **3**, 013605.

37 P.-W. Ma and S. L. Dudarev, *Phys. Rev. Mater.*, 2019, **3**, 043606.

38 C.-C. Fu, F. Willaime and P. Ordejón, *Phys. Rev. Lett.*, 2004, **92**, 175503.

39 K. Heinola, F. Djurabekova and T. Ahlgren, *Nucl. Fusion*, 2017, **58**, 026004.

40 A. F. Voter, *Radiation effects in solids*, Springer, 2007, pp. 1–23.

41 H. Kang and W. Weinberg, *J. Chem. Phys.*, 1989, **90**, 2824–2830.

42 C. Varvenne, F. Bruneval, M.-C. Marinica and E. Clouet, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 2013, **88**, 134102.

43 C. Varvenne and E. Clouet, *Phys. Rev. B*, 2017, **96**, 224103.

44 S. L. Dudarev and P.-W. Ma, *Phys. Rev. Mater.*, 2018, **2**, 033602.

45 X. Zhang, S. V. Divinski and B. Grabowski, *Nat. Commun.*, 2025, **16**, 394.

46 S. Plimpton, *J. Comput. Phys.*, 1995, **117**, 1–19.

47 G. Van Rossum and F. L. Drake, *Python 3 Reference Manual*, CreateSpace, Scotts Valley, CA, 2009.

48 W. Shinoda, M. Shiga and M. Mikami, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 2004, **69**, 134103.

49 P. M. Derlet, D. Nguyen-Manh and S. Dudarev, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 2007, **76**, 054107.

50 M. Boleininger, S. L. Dudarev, D. R. Mason and E. Martnez, *Phys. Rev. Mater.*, 2022, **6**, 063601.

51 K. Nordlund, S. J. Zinkle, A. E. Sand, F. Granberg, R. S. Averback, R. Stoller, T. Suzudo, L. Malerba, F. Banhart and W. J. Weber, *et al.*, *Nat. Commun.*, 2018, **9**, 1–8.

52 D. R. Mason, D. Nguyen-Manh, M.-C. Marinica, R. Alexander, A. E. Sand and S. L. Dudarev, *J. Appl. Phys.*, 2019, **126**, 075112.

53 Z. Hu, P. Desgardin, C. Genevois, J. Joseph, B. Décamps, R. Schäublin and M. F. Barthe, *J. Nucl. Mater.*, 2021, **556**, 153175.

54 C. Song, J. Hou, L. Chen, C. Liu and X.-S. Kong, *Acta Mater.*, 2024, **263**, 119516.

55 A. A. Kohnert and L. Capolungo, *Phys. Rev. Mater.*, 2019, **3**, 053608.

56 Cineca, ENEA and EUROfusion choose Lenovo Supercomputer to Develop Fusion Energy Research in Italy, <https://news.lenovo.com/pressroom/press-releases/cineca-enea-eurofusion-lenovo-supercomputer-fusion-energy-research/>, Accessed: 2025-02-28.

57 Z. Hu, J. Wu, Q. Yang, F. Jomard, F. Granberg and M.-F. Barthe, *Nano Lett.*, 2025, **25**, 10787–10793.

58 M. Gilbert and J.-C. Sublet, *Nucl. Fusion*, 2011, **51**, 043005.

59 W. Chrominski, L. Ciupinski, P. Bazarnik, S. Markelj and T. Schwarz-Selinger, *Mater. Charact.*, 2019, **154**, 1–6.

60 T. Schwarz-Selinger, *Nucl. Mater. Energy*, 2017, **12**, 683–688.

61 T. Schwarz-Selinger, *Mater. Res. Express*, 2023, **10**, 102002.

62 S. Markelj, E. Punzón-Quijorna, M. Kelemen, T. Schwarz-Selinger, R. Heller, X. Jin, F. Djurabekova, E. Lu and J. Predrag, *Nucl. Mater. Energy*, 2024, **39**, 101630.

63 J. Zavašnik, A. Šestan, T. Schwarz-Selinger, K. Hunger, E. Lu, F. Tuomisto, K. Nordlund, E. Punzón-Quijorna, M. Kelemen and J. Predrag, *et al.*, *Mater. Charact.*, 2025, **224**, 115050.

64 B. Wielunská, M. Mayer, T. Schwarz-Selinger, A. E. Sand and W. Jacob, *Nucl. Fusion*, 2020, **60**, 096002.

65 K. Heinola, T. Ahlgren, K. Nordlund and J. Keinonen, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 2010, **82**, 094102.



66 D. R. Mason, A. Reza, F. Granberg and F. Hofmann, *Phys. Rev. Mater.*, 2021, **5**, 125407.

67 S. Cui, R. P. Doerner, M. J. Simmonds, C. Xu, Y. Wang, E. Dechaumphai, E. Fu, G. R. Tynan and R. Chen, *J. Nucl. Mater.*, 2018, **511**, 141–147.

68 L. Shao, J. Gigax, D. Chen, H. Kim, F. A. Garner, J. Wang and M. B. Toloczko, *Nucl. Instrum. Methods Phys. Res., Sect. B*, 2017, **409**, 251–254.

69 S. Cui, M. Simmonds, W. Qin, F. Ren, G. R. Tynan, R. P. Doerner and R. Chen, *J. Nucl. Mater.*, 2017, **486**, 267–273.

70 F. Hofmann, M. P. Short and C. A. Dennett, *MRS Bull.*, 2019, **44**, 392–402.

71 A. Reza, H. Yu, K. Mizohata and F. Hofmann, *Acta Mater.*, 2020, **193**, 270–279.

72 A. De Backer, A. E. Sand, K. Nordlund, L. Luneville, D. Simeone and S. L. Dudarev, *EPL*, 2016, **115**, 26001.

73 H. Schultz, *Acta Metall.*, 1964, **12**, 649–664.

74 C. Kittel, *Introduction to Solid State Physics*, John Wiley & Sons, Inc., New York, 8th edn, 2004.

75 G. Paradezhenko, N. Melnikov and B. Reser, *arXiv*, 2017, preprint, arXiv:1711.04302, DOI: [10.48550/arXiv.1711.04302](https://doi.org/10.48550/arXiv.1711.04302).

76 G. H. Vineyard, *J. Phys. Chem. Solids*, 1957, **3**, 121–127.

77 D. R. Mason, D. Nguyen-Manh and C. S. Becquart, *J. Phys.: Condens. Matter*, 2017, **29**, 505501.

78 G. Henkelman, B. P. Uberuaga and H. Jónsson, *J. Chem. Phys.*, 2000, **113**, 9901–9904.

79 P.-W. Ma and S. L. Dudarev, *Phys. Rev. Mater.*, 2019, **3**, 013605.

