Adaptive moving environment for efficient molecular dynamics simulations of high-fluence ion irradiation

Fridlund, Christoffer

2022-11

Fridlund , C , Toijala , R , Flinck , O , Laakso , J , Nordlund , K & Djurabekova , F 2022 , '
Adaptive moving environment for efficient molecular dynamics simulations of high-fluence
ion irradiation ' , Computational Materials Science , vol. 214 , 111708 . https://doi.org/10.1016/j.commatsci.2022.111

http://hdl.handle.net/10138/348426
https://doi.org/10.1016/j.commatsci.2022.111708

Computational Materials Science 214 (2022) 111708

Contents lists available at ScienceDirect

Computational Materials Science

journal homepage: www.elsevier.com/locate/commatsci

Full length article

Adaptive moving environment for efficient molecular dynamics simulations of high-fluence ion irradiation

C. Fridlund [a,*], R. Toijala [a], O. Flinck [a], J. Laakso [a], K. Nordlund [a], F. Djurabekova [a,b]

[a] Department of Physics, University of Helsinki, P.O. Box 43, FIN 00014, Finland
[b] Helsinki Institute of Physics, University of Helsinki, P.O. Box 43, FIN 00014, Finland

ARTICLE INFO

Keywords:
Molecular dynamics
High-performance computing
Moving atom algorithm
High-fluence irradiation

ABSTRACT

Ion-beam processing of materials is widely supported by atomistic simulations by means of molecular dynamics. Although the approach has given several valuable insights, it has a limited operational window in both time and length scales. In particular, for high-fluence ion irradiation with multiple consecutive cascades, the direct molecular dynamics method becomes prohibitively time consuming.

In this work, we propose a speed-up algorithm for molecular-dynamics simulations of multiple consecutive collision cascades employing an *adaptive moving environment* model. In the model, the computational power is primarily focused on calculating the atomic movement in the propagating cascade regions, while thermally equilibrated regions outside the cascades are excluded. Up to a five-fold efficiency increase was seen with the *adaptive moving environment* compared to classical molecular dynamics, without any significant statistical difference in the results of multiple individual ion-cascade simulations in a heterostructure of alternating Si and SiO_2 layers. Simulations of temperature-driven dynamic annealing during high-fluence ion irradiation of Si nanopillars at elevated temperatures using the *adaptive moving environment* showed similar trends as experiments with respect to temperature dependence. The model is included in the atomistic simulator toolkit, COSIRMA (COmputer Simulator for IRradiation of MAterials), and can easily be enabled through the user-friendly graphical interface.

1. Introduction

Ion beams have been used for several decades both in materials research for exploring the properties of materials as well as in industrial fabrication of integrated circuits. Computer models for simulating the ion-beam effects of materials have been around for almost as long [1].

Molecular dynamics (MD), based on classical Newtonian mechanics, has actively been used for the last six decades to simulate the atomistic behavior of the target materials in everything from static systems with only a few hundred atoms, to high-dose ion irradiation – over the entire energy regime from a few eV to several GeV – on systems constituting millions of atoms, as well as biological systems with complex molecules [2–6]. With potential models expressed through optimized functional formalism and proper multi-threaded computing (i.e. OpenMPI [7], OpenMP [8], etc.), it is possible to simulate atomic systems of even billions of atoms with modern computers [9]. Nevertheless, the process can be unbearably slow, especially if one wishes to reach long time scales.

Simulating high-fluence ion irradiation with MD is even slower, as these simulations often include thousands of cascades and fairly large systems to hinder the cascades from overlapping with themselves when periodic boundary conditions (PBC) are in place. To decrease the computational time and resources needed to simulate relatively high fluences ($>1\times10^{10}$ ions/cm^2), one has to start looking into ways of optimizing the models. Moreover, optimized algorithms for atomistic simulations may increase the overall interest for the simulation models outside academia.

In this paper, we present the *adaptive moving environment* (AME, pronounced [ɑ:mɛ]), an extension to the classical MD algorithm. The model is inspired by existing *moving atom* algorithms [10,11], where the lattice atoms are present only around the projectile (i.e., moving atom) currently propagating through the matrix. No information about lattice atoms further away from the projectile is updated at every MD step [5,12]. Traditionally, these models are very fast, as they only need to know about the current neighborhood of the ion. However, this effectively limits the structures that can be simulated, and it is very difficult to simulate materials that are rendered disordered or amorphous during irradiation.

In AME, the atomic coordinates are available whether or not the atoms are included in any calculations. This allows for dynamic track-

* Corresponding author.
E-mail address: christoffer.fridlund@helsinki.fi (C. Fridlund).

https://doi.org/10.1016/j.commatsci.2022.111708
Received 23 May 2022; Received in revised form 10 July 2022; Accepted 3 August 2022
Available online 18 August 2022

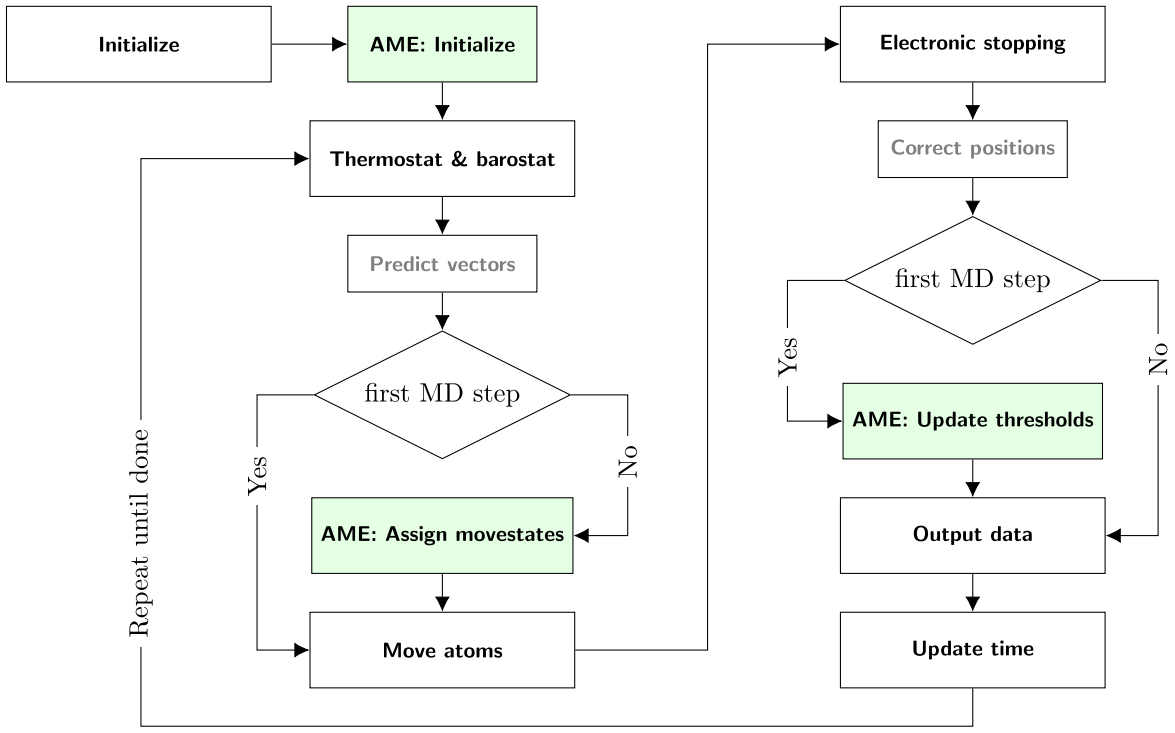

Fig. 1. AME included in the basic execution sequence of a classical predictor–corrector MD model. AME is initialized during the first MD step and all atoms are activated to get the initial force and energy distributions. During subsequent MD steps the threshold parameters are already defined and ready to use when assigning the `movestates`. Temperature and pressure controls are applied before entering the predictor, where the velocities and accelerations are computed based on the previous state. AME `movestate` assignment happens right after the predictor, before the atom positions are updated using the predicted vectors and the current forces. The new vectors are slightly adjusted based on the electronic stopping power before they are finally corrected with finite adjustments to compensate drifting in the numerical integrator. During the first MD step the threshold tables are updated for future use right before the physical quantities are printed. Finally the time step is updated and the model iterates the simulator loop, beginning from the thermostat and barostat again using the updated AME thresholds.

ing of all the structural modifications that occur during the irradiation, allowing inhomogeneous amorphization and various atomic compositions and concentrations to be present in different regions of the material. By carefully excluding atoms not involved in any cascades from the force calculation, it is possible to increase the simulation efficiency significantly in systems comprising hundreds of thousands of atoms. The interaction between the thermally equilibrated regions and the propagating region of the developing cascade is crucial for a physical model in all kinds of atomic systems and environments.

Currently we use the AME model to simulate irradiation of a Si/SiO_2 heterostructure, similar to the ones described in Refs. [13,14], with individual Si^+ ions and Si^+-ion clusters of various sizes to test the limitations of the model. As a production test of the model we investigated the feasibility of MD as a method for simulating temperature-dependent recrystallization of Si pillars through dynamic annealing caused by prolonged irradiation. The phenomenon has been observed experimentally in Si pillars at ambient temperatures above room temperature [15]. While MD has been previously used to simulate amorphous-to-crystalline phase transitions in Si structures [16,17], to our knowledge, there are no reports on simulations of high-fluence irradiation at elevated ambient temperatures.

2. Methodology

2.1. Adaptive moving environment

Our implementation of AME can be found in PARCAS [18–20] and in the newly released simulator toolkit *COmputer Simulator for IRrradiation of MAterials* (COSIRMA) [21]. AME aims to increase the efficiency of atomistic simulations of ion-induced collision cascades by avoiding time-consuming calculations of forces on atoms in thermally equilibrated regions, unaffected by the energetic impacts. AME defines atoms with kinetic energies exceeding the ambient temperature or atoms that were involved directly in collision cascades as "active". The activated regions of AME thus follows the propagating collision cascade. Fig. 1 shows a flow chart of the classical MD model with emphasis on the parts where the AME algorithm is included. AME is currently implemented in most of the potential models available in PARCAS. Nevertheless, AME can be implemented for any existing analytical potential model, as long as the specifics of the potential formalism are taken into account. The parameters of the model need to be carefully selected for every supported material, paying special attention to the calculation of energy losses when parts of the bi-directional force calculation are omitted. It is easy to turn the model on and off With the `movemode` parameter (see Table A.4).

2.1.1. Initialization

It is important to keep track of the atom species in a uniquely identifiable manner throughout the simulations with AME. Even atoms of the same element type may behave differently in different environments, e.g. c-Si, a-Si, and Si in a-SiO_2 all have different bonding energies, and hence different equilibrium forces.

A dynamic activation threshold is either computed from the equilibrated system or provided as an input parameter (`Fcrit` and `Ecrit`) during the initialization of the model, see Table A.4 for how it is used as an input parameter. The dynamically computed thresholds are estimated from the kinetic energy (proportional to the temperature according to the equipartition theorem) and the force distributions in the relaxed system. To obtain the required activation distributions, all atoms are assigned the HOT `movestate` (see Table 1) during the very first MD step to let the atoms move in accordance with classical MD. At the end of this step all atoms will have force values and kinetic energy computed by the numerical integrator of the MD model. These values can be used to initialize the activation thresholds.

Table 1
The `movestates` used by AME to define regions around a propagating collision cascade. Each `movestate` is presented with a description of how it is treated by the MD simulator. The `movestates` are presented in descending priority order (e.g. HOT will have priority over COLD), used when communicating atom data between neighboring computational nodes in a parallel multi-core setup.

Movestate	Description
ACTIVE	Atoms are simulated with classical MD without thermostat and have the ability to activate more atoms when the ion cascade propagates.
HOT	Atoms are simulated with classical MD without thermostat.
COLD	Atoms are simulated with classical MD with a Berendsen thermostat [22].
FIXED	Atoms are simulated with classical MD without thermostat but cannot move. This functions as a damping layer between the COLD and the STATIC layer.
STATIC	These atoms are considered to be in thermal equilibrium and can be skipped in the force calculation.

Since the atoms at surfaces and interfaces inherently have lower binding energy than their bulk counterparts, it is important that the atoms of these regions are not included in the definition of the activation thresholds, as that would skew the sensitiveness of the activation criteria. To effectively exclude the atoms at these regions, we generate the distributions only from atoms with the number of neighbors close to the average value in the structure, representing the bulk parts of the material. These regions are still activated when a cascade overlaps with them due to the lower binding energy.

Thermally equilibrated atoms have a small probability of gaining high kinetic energy according to the Maxwell–Boltzmann distribution. Some of these atoms might have very high values in the histograms. An effective way of preprocessing the distributions, before extracting the activation thresholds from them, is to clear all single-count bins at high values to find the activation thresholds more effectively. Exemplary histograms of the preprocessed force and kinetic energy distributions used for defining the activation thresholds (the vertical dashed lines) can be seen in Figs. 2(a) and 2(b) for a structure relaxed at 300 K using the Watanabe–Samela potential. Procedure 2 in the Appendix shows the initialization in pseudo code.

We note that it is highly advisable to relax the structure completely before computing the activation histograms to enable an accurate definition of the activation thresholds. A low activation threshold will reduce the efficiency of AME, while too high thresholds result in more aggressive deactivation, affecting the atoms that are still involved in a cascade.

2.1.2. Activation

At the beginning of every MD step (with the exception of the very first as described in Section 2.1.1), the simulation continues by setting the `movestate` of every atom to STATIC to start from a clean sheet every iteration (see Fig. 1). This is directly followed by the search of atoms satisfying any of the activation criteria presented in Eqs. (1)–(3), where the index i denotes ith atom and the subscript "a" denotes activation thresholds: F_a the activation force magnitude and E_a the activation energy and assigning them with the ACTIVE `movestate`. The forces and energies of the previous MD step are used in the succeeding step when using the activation thresholds to find which atoms should be activated.

$$F_i \geq F_\mathrm{a} \wedge E_i \geq E_\mathrm{a} \tag{1}$$

$$F_i \geq 2F_\mathrm{a} \tag{2}$$

$$E_i \geq 2E_\mathrm{a} \tag{3}$$

Eq. (1) is based on the distributions seen in Fig. 2 and can be shifted along the x-axis using a combination of the `FcritA` & `FcritFact` and `EcritA` & `EcritFact` input parameters (see Table A.4). Eqs. (2) and (3) were arbitrarily chosen to be clearly larger than the first criteria in case some atoms move fast without experiencing much force from other atoms or vice versa. In our case the chosen values worked quite well.

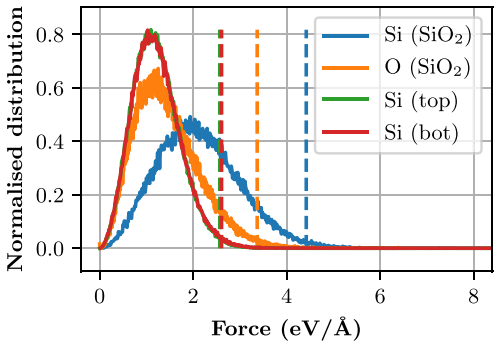

(a) Threshold values: Si (SiO_2) = 4.4 eV/Å, O (SiO_2) = 3.4 eV/Å, Si (top) = 2.6 eV/Å, and Si (bot) = 2.6 eV/Å

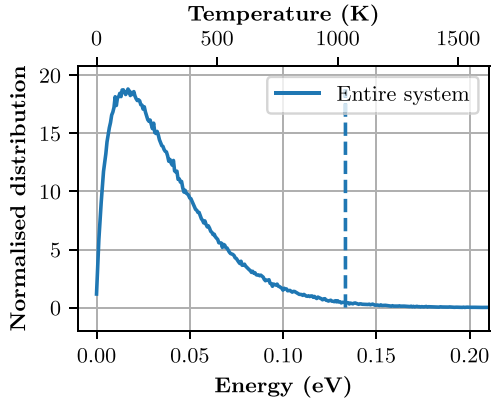

(b) Threshold value: 0.13 eV

Fig. 2. The activation thresholds are computed in real time from the force and energy distributions in the system after the first MD step. (a) shows the force histograms and (b) the energy/temperature histogram of the relaxed test structure seen in Fig. 4. The force distributions of the top and bottom c-Si layers are overlapping in the histogram as they both consist of similar crystalline Si. The energy distribution is the same for all atomic species as they have the same temperatures when properly relaxed. The dashed lines indicate the threshold limits (default at 99 %).

The spatial reach of the default Verlet list [23], setup with the link-cell method [24], in PARCAS was not efficient enough for AME to set the `movestates` to the atoms surrounding the ACTIVE atoms. Increasing the potential model's cutoff distance caused a general, unwanted slowdown of the simulation, and lead us to investigate other ways of assigning the `movestates`. A suitable compromise was reached by recreating a secondary Verlet list, with extended reach comprising the user-specified cutoffs for the boundaries between HOT/COLD, COLD/FIXED, and FIXED/STATIC (`movecuth`, `movecutc`, and `movecutf`, respectively in Table A.4), only for atoms in the ACTIVE `movestate`. By optimizing the activation criteria we could limit the number of simultaneous ACTIVE atoms, to further minimize the efficiency loss and the extra requirement for memory allocation in the secondary Verlet list.

Fig. 3 shows an example of the assigned `movestates` around ACTIVE atoms at 300 K using the default parameters in Fig. 3(a) a c-Si lattice and Fig. 3(b) an a-SiO_2 structure and Table 1 explains the behavior of the `movestates` in detail. Every atom outside the extent of the FIXED-cutoff is considered thermally equilibrated and left in the STATIC `movestate`. As every mobile atom is activated from an

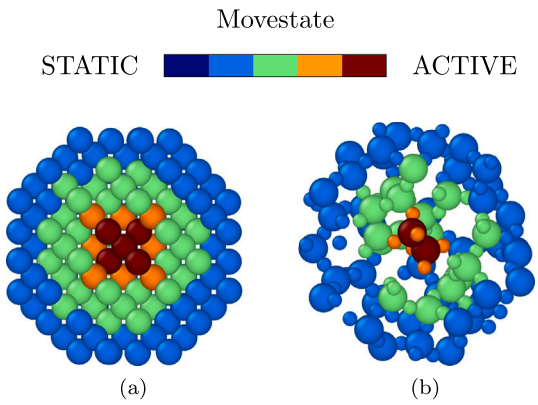

Fig. 3. During the activation stage, atoms exceeding the activation criteria are assigned the ACTIVE `movestate` (dark red in the picture), the surrounding atoms are assigned the four other `movestates` based on distance from the ACTIVE atom. In this example we have (a) crystalline Si and (b) amorphous SiO_2 at 300 K ambient temperature, with default parameters (Table A.4) and cutoff distances R_{hot} =3 Å, R_{cold} =6 Å, and R_{fixed} =9 Å. The structures are illustrated with a cross-sectional cut through the spherical activated region. The O atoms in the amorphous SiO_2 are shown with smaller radius and the atoms in both figures are surrounded by STATIC atoms, intentionally not shown for visual clarity of the `movestate` shells (for an illustration of the STATIC atoms, please see Fig. 7).

ACTIVE atom, it is necessary to find these as accurately as possible. Procedure 3 in the Appendix illustrates the activation algorithm with pseudo code.

In some cases, a new atomic species (uniquely identifying atom types in the structure) – that was not present initially – may appear during the simulation. E.g. new atomic species might be created when dynamically introducing ions or the user introduces new atomic species manually in-between collision cascades. AME cannot set the `movestate` to ACTIVE of atoms with atomic species not present during the initial MD step. To circumvent this issue, we have implemented the `mcritt` array that maps "unknown" atom identifiers to "known" ones. This basically means that if AME tries to set any of these atoms to ACTIVE, it will use the histograms of the known species instead.

When `allmovet` is greater than zero and there is `allmovet` fs left of the simulation, AME will activate all atoms by forcing them to have the ACTIVE `movestate` and behave just like in classical MD. This is useful if e.g. a longer classical relaxation of the structure is needed at the end of a cascade.

2.1.3. Heat transport in simulations

Since the majority of the atoms are assigned with the STATIC `movestate` (thus, immobilized), they cannot participate in the heat transport. To compensate for this we introduced the COLD region, where the temperature is controlled by a thermostat if it is applied in the simulation. The COLD atoms control the temperature of the HOT and ACTIVE regions. All three of them are simulated according to the classical MD algorithm. The COLD atoms interact with the FIXED atoms in the same fashion as the fixed atoms in classical MD, i.e. they experience forces, but do not transfer the momentum. The atoms in the STATIC `movestate` are not included in the force calculation and, hence, their contribution must be included explicitly. Currently we implement this contribution using Newton's third law if a STATIC atom is encountered in the neighbor loop of any atom. This results in a small offset in the force calculation caused by the asymmetric components in the three-body term. The impact of this offset is minimal since it appears in the region with the atoms in a state of relaxed equilibrium. A schematic of the algorithm, implemented into the Stillinger–Weber potential model can be seen in Procedure 1 in the Appendix.

2.1.4. Further optimization of the model

Usually, the best speedup for MD calculations is achieved through parallelization. In classical MD, the parallelization is done by spatially dividing the simulation cell among all available computation nodes. The results of the interatomic-interaction calculations of every node are communicated and synchronized between neighboring nodes every MD step.

There is usually a layer of "ghost atoms" at the boundaries between neighboring nodes, representing the same atoms. These "ghost atoms" enables smooth communication between neighboring units. AME fully supports this parallelization scheme. When the cascade propagates through the atomic structure, it is necessary to regularly balance the work distribution between the computation nodes. This is done by assigning the spatial extents of every computation node dynamically to ensure that every node has approximately the same number of actively simulated atoms, reducing possible computational bottlenecks.

The `movestates` have strict mathematical priority over each other, ensuring that all atoms always have the highest applicable `movestate` after the atom data has been communicated between neighboring computation nodes in a multi-core setup. All `movestates` are listed in descending priority order with descriptions in Table 1.

2.2. General setup of MD simulations

In all our simulations, unless mentioned otherwise, we used the Berendsen thermostat [22] with τ_T =100 fs^{-1} to control the temperature by scaling the velocity of either all atoms in the system (emulating the NVT ensemble) or within 5 Å from the sides with PBC or fixed atoms (quasi-NVE ensemble, since the temperature control affects only a part of the simulation cell). The former control was used during the relaxation runs and the latter during the cascade runs.

The cascade simulations in the quasi-NVE ensemble were followed by relaxation runs in the NVT ensemble to dissipate the remnants of heat brought in to the simulation cell by the impacting ions. This additional simulation phase helped the system to find a local equilibrium state at the given ambient temperature.

2.3. Setup of AME test simulations

To verify the performance of the newly developed AME model, we performed a set of simulations with individual 5 keV Si^+ ion impacts (N = 300) at 300 K and 700 K as well as cluster ion impacts on the $Si/SiO_2/Si$ semiconductor heterostructure shown in Fig. 4 at 300 K with three clusters sizes: 7, 26, and 69 atoms. The smallest cluster was simulated with 1, 3, 5, 7, and 9 keV; the middle cluster with 1, 3, 5, and 7 keV; and the largest cluster with 1 and 3 keV kinetic energy per cluster atom. All the simulations were repeated N = 100 times for statistics.

The simulation system was constructed as described in Ref. [13] by merging a 21.7 nm × 21.7 nm wide and 7.2 nm thick a-SiO_2 layer with two 10.9 nm thick c-Si layers with matching widths. The Si–Si, Si–O, and O–O interactions were described in these simulations by the Stillinger–Weber-like Watanabe–Samela interatomic potential [25,26].

The system was relaxed for 25 ps in the NVT ensemble at both 300 K and 700 K, keeping the Z-surface open to allow for relaxation of stresses caused by thermal expansion. We used PBCs in the X- and Y-directions to emulate an infinite surface. The relatively low ion energy ensured that the entire cascade developed within the cell sufficiently far from the sides and the bottom to prevent the cascade from overlapping with itself through the PBC or with the temperature-controlled bottom.

In the cascade simulations with AME, the Berendsen thermostat was applied only to the COLD atoms, independently of their position in the cell. As described in Section 2.1.3, the interaction of the ACTIVE and HOT atoms with the surrounding COLD atoms enables heat dissipation from the cascade region without involving the disabled STATIC atoms in the simulation.

Table 2
The AME settings used in the test simulations. The *Default* settings are the ones provided by the software and the *No move all* has disabled the classical MD simulation of all atoms at the end of the simulation. See the main text for more information about the parameters.

Parameters	Unit	Default	No move all	Short cutoff	Long cutoff
`movemode`	flag	1 (on)	1 (on)	1 (on)	1 (on)
`allmovet`	fs	1000.0	0.0	1000.0	1000.0
`movecuth`	Å	3.0	3.0	2.0	5.0
`movecutc`	Å	6.0	6.0	4.0	10.0
`movecutf`	Å	9.0	9.0	6.0	15.0
`FcritA`		0.99	0.99	0.99	0.99
`EcritA`		0.99	0.99	0.99	0.99
`FcritFact`		1.0	1.0	1.0	1.0
`EcritFact`		1.0	1.0	1.0	1.0
`mcritt`		−1	−1	−1	−1
`Fcrit`	eV/Å	−∞	−∞	−∞	−∞
`Ecrit`	eV	−∞	−∞	−∞	−∞

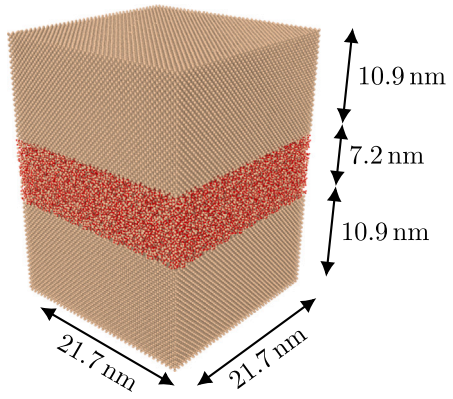

Fig. 4. The semiconductor heterostructure with periodic boundaries in X- and Y-directions used in the benchmark simulations. Red (darker) spheres indicate positions of O atoms and light brown ones Si.

The entry point of the ion was selected randomly within a unit cell, 5 Å above the middle of the surface. To avoid channeling effects, the incident θ angle was tilted with respect to the surface normal by 7° [27], while the azimuthal angle ϕ was selected at random. The spherical ion clusters were set to impact the surface at normal incidence, since the probability of channeling in these dense cascade events is very low. The clusters were rotated randomly around their own cardinal axes in all three directions for more variation. The lateral positioning of the center of the ion cluster was randomized within one unit cell and the vertical position was determined by placing the center of the cluster 15 Å above the surface of the target, making sure there was empty space between them. All cascades were allowed to develop for up to 5 ps. AME was set to enter the relaxation phase earlier if all atoms were deactivated before the time limit of 5 ps (see Section 2.1.2).

We ran one individual ion-impact simulation using AME with the default parameter set, two slightly varying the HOT-, COLD-, and FIXED-cutoff distances (referred to as "Short cutoff" and "Long cutoff"), and one turning off the `allmovet` (called "No move all") to assess the sensitivity of AME and how it responded to variations in the input parameters (see Table 2 for a complete list of parameters used). We further inspected possible model inconsistencies resulting from poor choice of AME parameters. We compare the performance of AME in all four tests with the performance of classical MD in Fig. 8.

Finally, we compared the quantitative results of all test simulations obtained with and without AME, taking into account the stochastic nature of the ion irradiation events [28,29]. Note that we do not compare individual cascades since small differences in the simulation setups may affect the exact path of energetic particle trajectory. To verify the reliability of AME calculations, we compared the ion ranges and the defect depth distributions for point defects produced on average in the 300 independent classical MD and AME simulations. The defects were analyzed using the Wigner–Seitz defect analysis modifier in OVITO [30].

Although the SiO_2 layer in our simulations is amorphous to begin with and the c-Si layers may render amorphous during irradiation, we apply the same definition of "interstitials" and "vacancies" for these amorphous structures as for a crystal lattice. With this we aim for a quantitative comparison of performances of the two algorithms (classical MD and AME) in a consistent manner rather than reporting the actual defect concentration as a result of the studied irradiation process.

2.4. Setup of the Si recrystallization simulations

To analyze the applicability of AME in more demanding situations, we used the model to simulate temperature-dependent recrystallization through dynamic annealing of radiation damage in Si structures. In these simulations we used the Stillinger–Weber (SW) [31] and the Tersoff (T) [32] interatomic potentials, which have been extensively used for Si simulations for decades [31–34]. We first ensured that our MD setup would be able to capture the thermally activated recrystallization with a flat Si surface before moving on to a more complex Si nanopillar structure.

We analyzed the degree of crystallinity in the structure after every ion impact by determining the number of atoms in diamond lattice configuration compared to any other configuration using the *identify diamond structure* modifier found in OVITO [30]. Furthermore, we analyzed the increase in amorphization resistance of Si with increasing temperature by performing the same simulation series of 500 ions at four different ambient temperatures: 300 K, 1000 K, 1200 K, and 1500 K for the SW and 700 K, 1000 K, 1200 K, and 1500 K for the T potentials.

2.4.1. Temperature scaling

Since recrystallization is a thermally activated process, we have to consider how well the potentials can reproduce the experimental melting point of Si. The SW potential predicts the melting point of Si very accurately ($T_{melt,\ SW} \approx 1690$ K [35,36] versus the experimental value $T_{melt,\ exp} \approx 1683$ K [35]), since the potential was fitted to the melting point [31]. The T potential, however, overestimates the value significantly ($T_{melt,\ T} \approx 2570$ K [36]). To setup MD irradiation conditions comparable to experiments, we used homologous temperature, i.e. rescaled the target ambient temperature of the simulation to be the same fraction of the melting point as in the experiment. In other words, we find the ratio $b = T_{ambient,\ exp}/T_{melt,\ exp}$ for the experiment and use it to find the comparable ambient temperature for the simulation. E.g. Xu et al. in Ref. [15] used the experimental $T_{ambient,\ exp} = 673.15$ K, which gives us $b \approx 0.4$. The comparable value in the SW potential is $T_{ambient,SW} = 0.4 \times 1690$ K ≈ 676 K and in the T potential $T_{ambient,T} = 0.4 \times 2570$ K ≈ 1028 K.

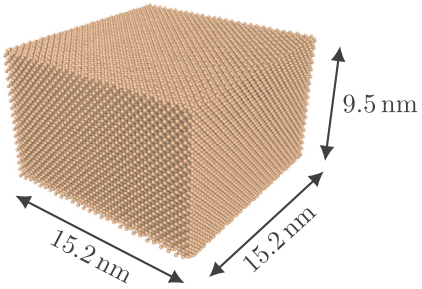

Fig. 5. The c-Si block structure used in the recrystallization bulk test for verifying the recrystallization process using MD.

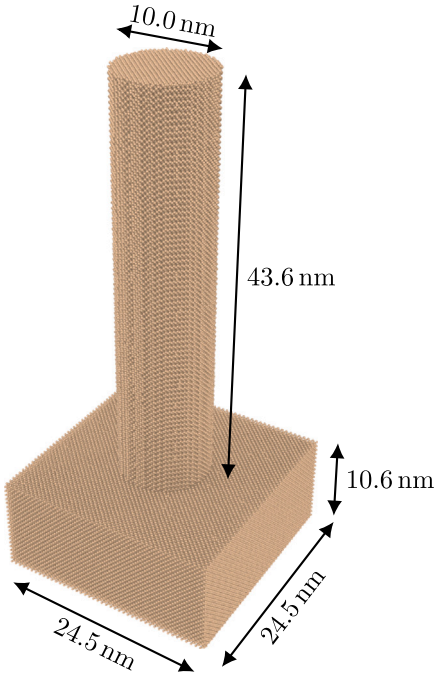

Fig. 6. The c-Si nanopillar used for simulating recrystallization with AME enabled. The dimensions are marked in the figure.

2.4.2. Flat surface simulation setup

To prepare the flat surface, we relaxed a 15 nm × 15 nm × 10 nm Si block (see Fig. 5) for 25 ps in the NPT ensemble at P =0 bar and the desired ambient temperature with the PBC enabled in all directions. After that, we opened the surface in the Z-direction (fixed the three bottom-most layers) and relaxed for 25 ps more in the NVT ensemble at the same ambient temperature to create a stable surface.

The irradiation was performed at room temperature in the quasi-NVE/NVT ensembles (see details in Section 2.2) with 500 consecutive 1 keV Si^+ ion impacts at slightly tilted ($\theta = 7$ ° and $\phi = 27$ °) incidence to avoid the channeling directions.

2.4.3. Nanopillar simulation setup

The Si nanopillar was cut out from a rectangular Si block relaxed in the NPT ensemble (at room temperature and 0 bar pressure) for 25 ps. The 43.6 nm long nanopillar with diameter 10 nm was left on a pedestal with dimensions 24.5 nm × 24.5 nm × 10.6 nm, as shown in Fig. 6. The prepared pillar was finalized for all the chosen ambient temperatures through a 25 ps relaxation in the NVT ensemble at the desired ambient temperature (300 K, 700 K, and 1200 K).

We used 25 keV broad-beam Si^+ ion irradiation up to $\sim 1.25 \times 10^{14}$ cm^{-1} fluence on the Si nanopillar at the chosen ambient temperatures – all lower than the melting point of the potential model – using both the T and the SW potentials in the quasi-NVE/NVT ensembles (see details in Section 2.2). In these simulations the heat was dissipated by a few atom layers at the bottom of the pedestal. Every single collision cascade was run for 25 ps with `allmovet` (see Table A.4 for more information about the parameters) activating all atoms for the final 5 ps of the simulation to cool the entire system to the ambient temperature in the NVT ensemble.

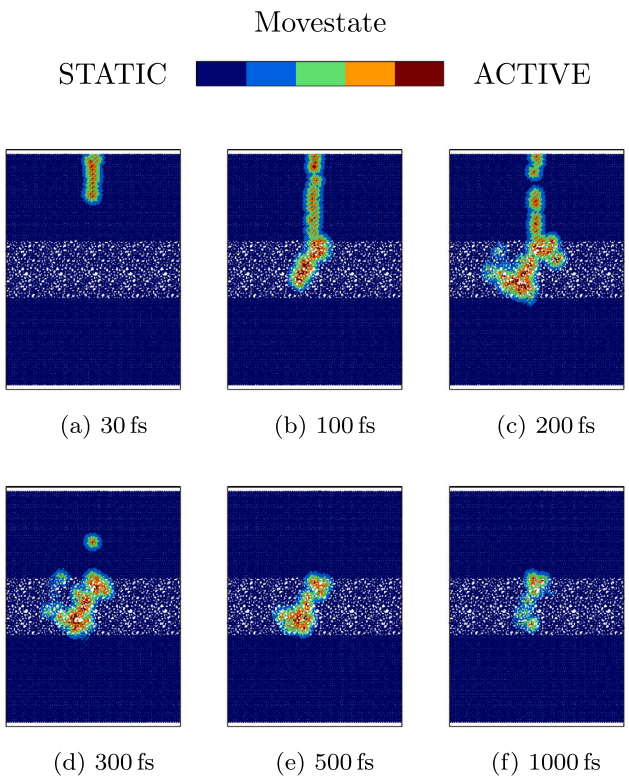

Fig. 7. The first 1000 fs of one ion cascade at 300 K, seen through a 1.0 nm cross section of the X–Z-plane, with the *Default* AME parameter set in the heterostructure shown in Fig. 4. The ballistic phase is over before 500 fs and the ion cascade subdues after the ion has stopped. The `movestates` are color coded.

3. Simulation results

3.1. Performance of the Adaptive Moving Environment

We demonstrate the use of AME by illustrating the cascade propagation triggered by an individual ion impact in Fig. 7. The figure clearly shows two significant points: Using AME does not affect trajectories of energetic atoms in the ballistic phase (the cascade and its branches are followed as expected); the algorithm successfully deactivates atoms in the regions where the thermal equilibrium has been reached (see the upper part of the simulation cell shown in Figs. 7(a)–7(f)).

Additionally we analyze the computational efficiency of MD simulations with AME. In Fig. 8 we show the computational time distribution of 300 simulations performed both with classical MD (without AME) and with AME using the four parameter sets described in Section 2.3. We see that AME increased the time efficiency multi-fold. However, we note that longer cutoffs perform with lower time efficiency than shorter cutoffs, but with statistical results closer to the ones obtained with classical MD. Nevertheless, the suggested default parameters provided the most optimal efficiency with respect to reliability of the obtained results and time spent on calculation per cascade.

Table 3
Mean depths of formation of vacancies and interstitials as well as the mean ion range averaged over $N = 300$ simulations performed at 300 K and 700 K in the structure shown in Fig. 4. The error bars are the standard error of the mean. The shape of depth profiles can be seen in Fig. 9.

Simulation	Int. depths (Å)	Vac. depths (Å)	Ion depths (Å)
Distributions at 300 K			
Classical	115.9 ± 2.6	113.0 ± 2.7	125.9 ± 2.9
Default	121.5 ± 3.2	108.8 ± 2.8	125.0 ± 2.8
No move all	122.2 ± 3.1	110.5 ± 2.8	125.0 ± 2.8
Short cutoff	125.8 ± 3.3	112.6 ± 2.9	127.6 ± 2.8
Long cutoff	130.6 ± 3.0	123.1 ± 2.7	125.8 ± 2.8
Distributions at 700 K			
classical	112.1 ± 2.7	109.9 ± 2.7	123.5 ± 3.1
Default	103.3 ± 2.8	104.0 ± 2.8	121.7 ± 2.9

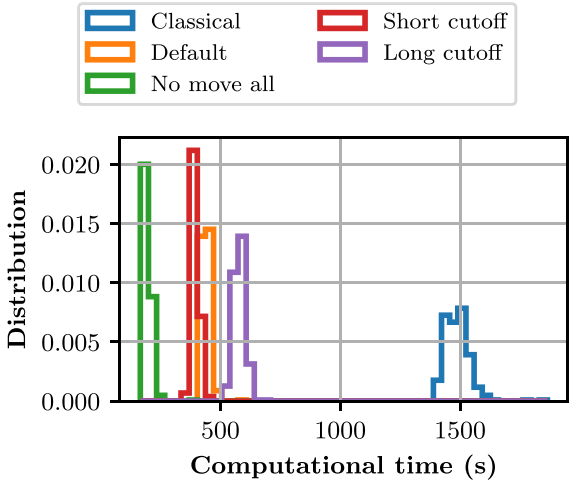

Fig. 8. Comparison of the computational time used by classical MD and AME with the four different parameter sets in Table 2 for simulations of individual ion impacts on the $Si/SiO_2/Si$ heterostructure system at 300 K. The statistical results of the tests are shown in Table 3.

3.1.1. Simulations of individual ion impacts

In Table 3 and Fig. 9 we compare the statistical results obtained from the four test simulations averaged over 300 individual ion impacts at 300 K ambient temperature. At the bottom of Table 3 we present the results obtained for 700 K (~400 °C) with classical MD and with AME using the default parameters. This temperature is known to increase the Si resistance to radiation-induced amorphization [15,37,38], this has also been confirmed in MD simulations previously. For individual cascades, there is no significant difference between the results obtained at 300 K and 700 K, hence we focus the following discussion only on the 300 K case.

We see that all test simulations at 300 K resulted in very similar distributions with the mean values within the standard error of the means. Since the initial damage is produced in the ballistic phase of collision cascades, such an agreement is expected. The most aggressive regime with *No move all* parameter still resulted in the similar defect distributions to other simulations, which indicates that even in this case, the heat dissipation, which was the most efficient in these simulations, was sufficient and it did not affect the results significantly. Additional relaxation run, which allows for relaxation of small displacements and recombination of shallow defects, did not change the results remarkably. Hence, we conclude that the initial defect production in collision cascades can be fully simulated with the use of AME and no additional computational resources for the relaxation runs with the `moveallt` parameter are needed. However, this may not be true for multiple consecutive collision cascades and a regular activation of all atoms is desirable to enable interaction of previously formed defects with the newly produced ones in all regions of the simulations box.

3.1.2. Simulations of cluster ion impacts

In the following, we analyze the AME efficiency for the cluster ion impact simulations. In Fig. 11 we show an almost three-fold speedup at the lowest incident energy ($E_0 = 7$ keV with $N_{at} = 7$), with decreasing efficiency for larger energy deposition either for the same or for the larger cluster sizes, until the simulation performs at the same speed as classical MD (or even slower, due to the additional neighbor list).

In Fig. 10 we compare the atomic displacement distributions, depth profiles of the mean atomic displacements as well as interstitials and vacancies produced in the cluster ion impact simulations. We excluded all sputtered atoms from the distributions since they would only add to the long-range tail. Moreover, the sputtering results are identical both with and without AME, since the model does not affect this process. In all three cases, the stronger vacancy peak near the surface compared to the interstitial one in the same region, indicates the loss of surface atoms to sputtering. The corresponding quantities in the distributions were averaged over 100 simulations and divided in 200 bins along the Z-direction (the height) of the structure.

We note that the shock waves caused by thermal expansion of the thermal spikes in the cluster ion simulations were damped less efficiently with AME compared to classical MD. This resulted in a higher short-range peak with much fewer long-range displacements in the displacement distributions in the AME simulations (compare results with AME "on" and "off" in the far-left panels of Figs. 10(a)–10(c)).

In the simulations with the smallest cluster size ($N_{at} = 7$, see Fig. 10(a)), the strict activation thresholds resulted in relatively narrow cascades limiting movement of less active atoms with AME "on" as compared to classical MD (AME "off"). This is evident from the atomic displacement distributions (center-left panel in Fig. 10(a)), which shows a strong peak at the surface for the simulations with AME "on", while already below 50 Å from the top surface, this quantity is much smaller than in the simulations with AME "off". At the upper Si/SiO_2 interface, the atoms are much less active resulting in much fewer defects in this region (compare blue peak – AME "off" – at the upper interface is much higher than the orange one — AME "on").

The largest cluster size ($N_{at} = 69$, $E_0 = 69$ keV, see Fig. 10(c)) yielded the distributions closest to those obtained with classical MD. The displacement distribution (far-left panel in Fig. 10(c)) still exhibits the same artificial short-range peak, however, the distribution of mean displacements is in much better agreement between classical MD and AME (the center-left panel in Fig. 10(c)). The peaks of the number of interstitials and vacancies in the SiO_2 layer near the interface were somewhat broader in the simulations with AME on, which may also relate to the less efficient dynamic annealing, which can be expected in thermal spikes that are stronger in these high-energy cluster ion simulations. However, the defects are generated at the same depths and the total displacements are very similar.

At the interfaces, the potential energy is the highest, which explains the point-defect distribution peaks around the interface locations when the energetic cascade reach and cross them in the classical MD case. This is expected as some atoms – due to imperfect bonding between

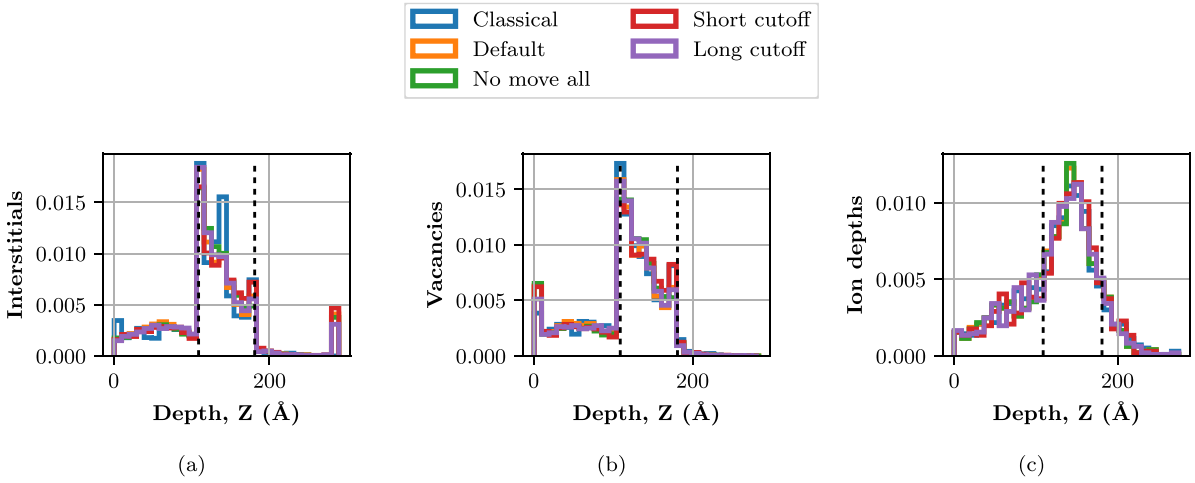

Fig. 9. Defect depth distributions averaged over $N = 300$ individual ion-impact simulations, using classical MD and AME with the different threshold parameters shown in Tables 2 and 3. In these plots, ions enter from the left in the distribution plots. (a) shows the depth profile of interstitials, (b) shows the depth profile of vacancies, and (c) shows the ion range profile. The vertical dashed lines indicate where the Si/SiO_2 interfaces are located.

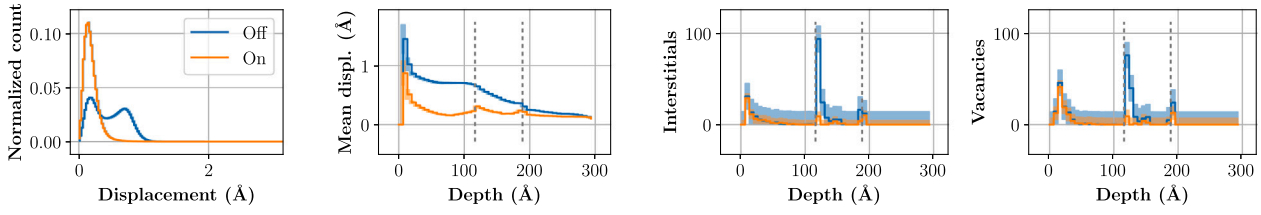

(a) Cluster $N_{at} = 7$ with total incident energy $E_0 = 7\,\text{keV}$.

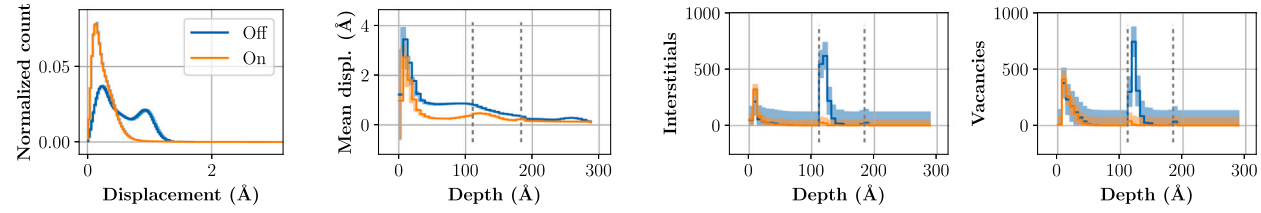

(b) Cluster $N_{at} = 26$ with total incident energy $E_0 = 26\,\text{keV}$.

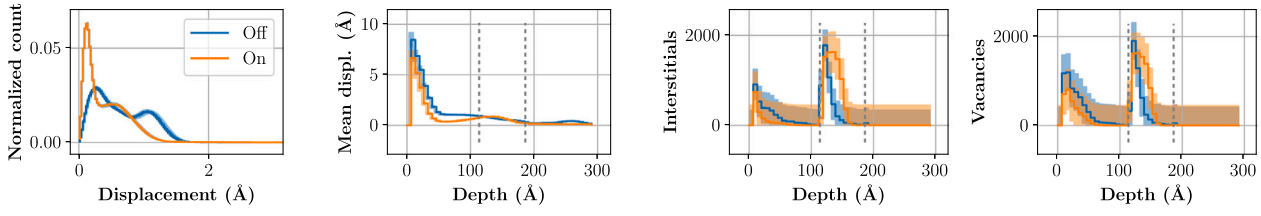

(c) Cluster $N_{at} = 69$ with total incident energy $E_0 = 69\,\text{keV}$.

Fig. 10. Comparison of different distributions of several quantities obtained with AME ("On" in legend) and classical MD ("Off" in legend) simulations of the cluster ion impacts on the $Si/SiO_2/Si$ heterostructure for three different cluster sizes: (a) 7 atoms, (b) 26 atoms and (c) 69 atoms, with 1 keV per atom in each of the clusters. For all cluster energies, the atomic displacement distributions (far left), the depth profiles of the mean atomic displacements (center left), interstitials (center right) and vacancies (far right) were averaged over 100 simulations. The shadowed areas in the distributions indicate standard deviation. The sputtered atoms are not included in the analysis. The vertical dashed lines indicate Si/SiO_2 interface at ~110 Å and ~180 Å. The surface where the ions enter is at 0 Å.

the two substances – have higher potential energy and hence, are more easily displaced from their initial positions.

When the velocity of an atom approaches $(r_{HOT} + r_{COLD})/\Delta t$, the atom may reach the FIXED region before the surrounding atoms are activated, which ultimately might lead to a sudden large local increase in the potential energy. This situation is particularly prominent at unphysically high incident energies. Overlapping, dense cascades initiated by high-energy cluster ion impacts lead to the development of thermal spikes with strong non-linear effects. This in turn leads to a significant increase in the velocities of moving atoms, and as a result, the energetic

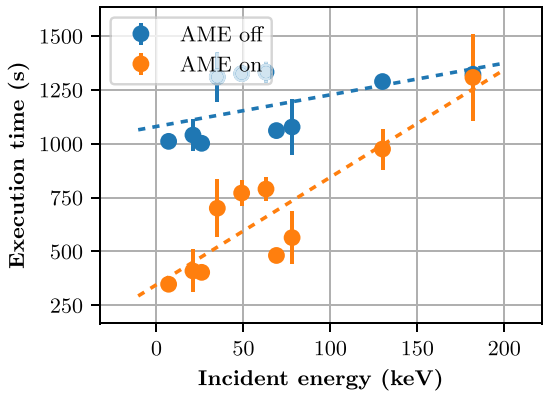

Fig. 11. Computational efficiency of AME ("AME on") compared to classical MD ("AME off") for cluster ion impacts of various incident energy. The almost three-fold speedup at the lowest incident energy was achieved with the smallest ion cluster size and energy ($N_{at} = 7$, $E_0 = 7$ keV). The higher the energy, the closer the time-consumption is to classical MD. At the very high energies, the time-consumption might even surpass classical MD due to the extended neighbor list with many active atoms.

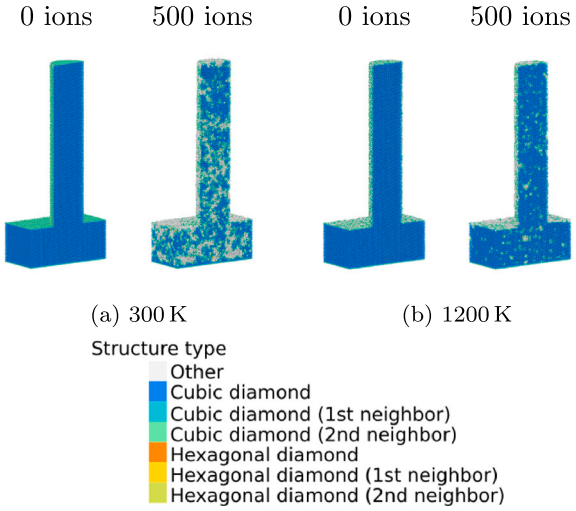

Fig. 12. Cross-sectional images of the initial and final states of the Si nanopillar after the 25 keV broad-beam Si^+ irradiation at (a) 300 K and (b) 1200 K. Blue color indicates diamond lattice (unamorphized c-Si structure).

atoms might physically overlap with frozen COLD atoms during said high-energy event. Hence to benefit fully from AME at borderline extreme high-energy and dense-cascade events, it is necessary to properly optimize the activation criteria in Fig. 2, for instance, the atoms can be activated at even lower force and temperature.

3.2. Simulations of dynamic annealing at elevated temperatures

The recrystallization simulation was done only with AME and addressed the dynamic annealing observed in c-Si irradiated to high fluences at elevated temperatures. In recent experiments of prolonged ion irradiation of Si nanopillars, the shape of the nanopillars was observed to not change drastically if the ambient temperature was above 823 K to 873 K [15].

Fig. 13 shows the crystallinity trends of $N = 500$ consecutive ion cascades for both the T and the SW potentials. The analysis was done with the "identify diamond structure" modifier in OVITO [30]. The presented plots can be interpreted as amorphization rates, that clearly decrease when the ambient temperature increases (see Fig. 12). The low amorphization rates observed with the SW potential at all tested temperatures is explained by the explicit bond-angle term in the

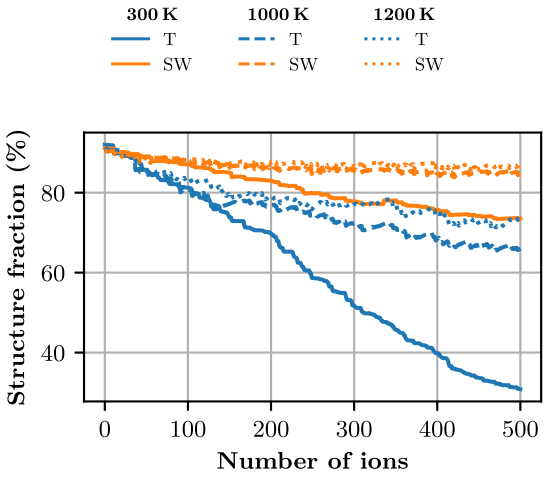

Fig. 13. Evolution of the structure fraction (the fraction of atoms in diamond environment as given by the "Identify diamond structure" modifier in OVITO) in the Si nanopillar with number of incoming consecutive ions simulated at 300 K, 1000 K, and 1200 K with the Tersoff (T) and the Stillinger–Weber (SW) potentials. The effect of temperature is stronger in the SW potential.

potential's functional form, while the atoms in the T potential require longer time to stabilize at the correct bond-angle distribution.

The bulk irradiation simulations showed a clear difference in crystallinity at different ambient temperatures already after a few ion impacts. This result was further improved after we used the ambient temperatures rescaled with respect to the melting point of the potential model (see Section 2.1.3). While we clearly see a thermal dependence of the recrystallization in the Si nanopillars with both potential models, the MD timescales cannot fully capture the enhanced dynamic annealing observed in the experiments by Xu. et al. [15]. Since the annealing is partly driven by diffusion mechanisms, that takes place at much longer timescales beyond the MD timespan, it will not be fully accessible by the means of MD. While AME allows for faster execution of the hot ion cascades (leading to higher fluences), it will not be able to simulate longer MD scales without returning to classical MD for relaxation. This means that simulation of the long-term relaxation processes – including recrystallization – will not be improved with the method, although it can capture the process on very similar level as the classical MD.

4. Discussion

This far we have demonstrated the abilities of the proposed extension to the classical MD algorithm, AME. AME aims to speed up high-fluence ion-irradiation simulations keeping the accuracy of MD. The key feature of the algorithm is to treat atoms that are not involved directly in collision cascades as non-interacting. These atoms are not included in the numerical integration of equations of motion, which greatly reduces the computational resources needed when processing a cascade event.

We observed a significant increase in computational efficiency for high-fluence ion-irradiation simulations of individual ion impacts with AME and shorter `movestate` cutoff radii. There was, however, no notable difference in the number of produced defects nor in the defect depth profiles when altering the cutoff radii. The cluster ions simulations at lower energies resulted in a slightly poorer statistical agreement between the two models, however, almost a three-fold speedup was achieved. The higher energetic clusters had better statistical agreement with classical MD, but lower speedup efficiency. This result is expected as a higher concentration of energetic events leads to a significant increase in energetic collisions that require a larger volume of atoms to dissipate the impact energy. It is important to keep in mind that the

parameters must be properly optimized before applying the method to new materials. While it is possible to perform the entire AME simulation without any thermal equilibration at the end of the simulation (further speeding up the simulation), we strongly recommend it, especially if the structure is being exposed to multiple consecutive ion cascades.

While simulating recrystallization due to radiation annealing with AME, we noticed the same timespan limitation as observed with classical MD. Since the speedup does not deal with acceleration of the algorithm, but more efficient use of computational resources, AME is not recommended for simulations of near-equilibrium processes, such as diffusion. The advantage of AME is in its ability to simulate more ion-bombardment events in larger systems compared to classical MD using the same CPU resources.

Although the Si nanopillar rendered amorphous in our simulations with elevated temperature, AME was able to produce a clear temperature dependence for the amorphization and recrystallization rates at different ambient temperatures. Hence, the MD timescale is sufficient to deduce a positive effect of higher ambient temperatures on the recrystallization rate and in general, the elevated temperatures can be used to increase the resistance to amorphization of crystalline Si structures under prolonged ion irradiation even in confined nanosystems.

5. Conclusions

One of the major limitations of the widely used molecular dynamics (MD) method is the short time scale that it can handle. While the limitation is intrinsic, it is interesting to search for application-specific speedup schemes. In this work, we present a speedup extension, *adaptive moving environment* (AME), to the classical MD algorithm suitable for studies of ion-irradiation effects in materials. The approach is based on excluding the atoms that are not directly involved in collision cascades from the numerical integration of the MD algorithm with an adaptive scheme that regulates the states of the atoms between "mobile" and "immobile" depending on whether they are in the vicinity of the collision cascades. The AME parameters are chosen from analysis of force and energy distributions in the equilibrated system prior to any irradiation. The presented approach is validated in simulations of ion irradiation of nanoscale Si heterostructures. The simulations showed significant speedup compared to similar classical MD simulations, while obtained results agreed very well, especially, for individual ion-impact simulations.

CRediT authorship contribution statement

C. Fridlund: Simulation work and analysis, Speedup model, Software, Conceptualization, Implementing speedup model, Draft and writing. **R. Toijala:** Implementing speedup model. **O. Flinck:** Simulation work and analysis. **J. Laakso:** Software design, Implementation of COSIRMA. **K. Nordlund:** Supervision. **F. Djurabekova:** Supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The source code is referenced in the article.

Acknowledgments

This project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 688072 and from the Vilho, Yrjö and Kalle Väisälä Foundation. The views and opinions expressed herein do not necessarily reflect those of the European Commission. We thank the Finnish Grid and Cloud Infrastructure (FGCI), urn:nbn:fi:research-infras-2016072533, and CSC–IT Center for Science, Finland, for computational resources.

Appendix. Pseudo-code of AME implementations

A.1. Interatomic potential example

Procedure 1 The adaptive moving environment implemented in the Stillinger–Weber potential. The implementations in other potential models differ slightly, depending on the three-body k-loop.

Input: Positions, types, `movestates`, and neighbors of atoms.
Output: E_p and F

```
Initialize E_p and F to zero
for i ← 1, N_atoms do
    if movestate[i] is STATIC then
        Skip to next atom
    end if
    for j ← 1, neighbor_count[i] do
        j_n ← neighbor_list[i, j]
        Compute and save dist. between atoms i and j_n
        if movestate[j_n] is STATIC then
            Compute 2-body E_p and F for i–j_n
            Assign to both i and j_n using Newton III
        end if
        for k ← 1, j − 1 do
            k_n ← neighbor_list[i, k]
            Load distance between atoms i and k_n
            Compute 3-body E_p and F for j_n–i–k_n
            Assign parts of contrib. to each atom i, j_n, k_n
        end for
    end for
end for
```

A.2. Initialization

Procedure 2 Initialization of the adaptive moving environment during the very first MD step.

Input: F, E_k, $F_\mathrm{crit}[]$, E_crit, types, box, and neighbors.
Output: $F_\mathrm{a}[]$, E_a

```
for i ← 1, N_atoms do
    movestate[i] ← HOT
end for
Generate E_k histogram
Generate separate F histogram for each atom type
Compute activation energy, E_a, based on E_k histogram and E_crit parameters
for t in atom types do
    Compute activation force, F_a[t], based on F[t] histogram and F_crit[t] parameters
end for
```

A.3. Activation

Procedure 3 Deactivation/Activation method of the adaptive moving environment. Turn everything off at the beginning, and only reactivate affected atoms.

Input: positions, F, E_k, $F_\mathrm{a}[]$, E_a types, `movestates`, neighbor list cutoff, and neighbors of atoms.
Output: positions, `movestates`

```
movestate[1,N_atoms] ← STATIC
for i ← 1, N_atoms do
    movestate[i] ← ACTIVE based on E_a and F_a[i]
end for
Update speedup neighbor list of ACTIVE atoms
for i ← 1, N_active atoms do
    for n in neighbors of active do
        dr ← distance between neighbor and ACTIVE
        Compare dr with HOT, COLD, and FIXED cutoffs and update the movestate of n
    end for
end for
Communicate the movestates between MPI nodes
```

Table A.4
A list of the parameters controlling AME. The model is set up to have working default values for all parameters, while providing the user with flexible solutions to modify the execution. Some of the parameters depend on the potential model, e.g. R_{cutoff} is the extent of the potential model in use. If the simulation is restarted from an earlier stage, it is possible to read in the activation thresholds from the previous run with the last two parameters presented in this table. The $-\infty$ in `Fcrit[]` and `Ecrit` are interpreted as "not given" and used to compute the values as described in Section 2.1.1.

Parameter	Default	Comment
`movemode`	0	0 or 1; Enables/Disables the algorithm. Disabled by default.
`allmovet`	0.0	>= 0.0 [fs]; Sets the time that should be reserved out of the total simulation time for relaxing all atoms at the end of the simulation. No relaxation by default.
`movecuth`	−1.0	−1.0 or > 0.0 [Å]; The extent of the HOT region. $1 \times R_{\text{cutoff}}$ by default.
`movecutc`	−1.0	−1.0 or > 0.0 [Å]; The extent of the COLD region. $2 \times R_{\text{cutoff}}$ by default.
`movecutf`	−1.0	−1.0 or > 0.0 [Å]; The extent of the FIXED region. $3 \times R_{\text{cutoff}}$ by default.
`FcritA`	0.99	0.0 to 1.0; Percentage of the force activation criteria. 99 % by default.
`EcritA`	0.99	0.0 to 1.0; Percentage of the energy activation criteria. 99 % by default.
`FcritFact`	1.0	>= 0.0; Scaling factor of the force activation criteria. 1.0 by default.
`EcritFact`	1.0	>= 0.0; Scaling factor of the energy activation criteria. 1.0 by default.
`mcritt[]`	[−1, …, -1]	−1 or Atom type ID; Array mapping unknown atom types to known types (e.g. those introduced by the simulator vs those read from the lattice file).
`Fcrit[]`	[−∞,…, −∞]	[eV/Å]; Array of previously computed force activation thresholds for each atom type.
`Ecrit`	−∞	[eV]; Previously computed energy activation threshold.

A.4. Parameters

All parameters have a default value if they are not set in the input file of PARCAS. Some of the default values are negative to inform AME to automatically compute a positive value. E.g. the default extents of the HOT, COLD, and FIXED `movestates` are automatically determined from potential cutoff defined in the potential model.

References

[1] K. Nordlund, Historical review of computer simulation of radiation effects in materials, J. Nucl. Mater. 520 (2019) 273–295, http://dx.doi.org/10.1016/j.jnucmat.2019.04.028, URL https://linkinghub.elsevier.com/retrieve/pii/S0022311518314703.

[2] B.J. Alder, T.E. Wainwright, Studies in molecular dynamics. General method, J. Chem. Phys. 31 (1959) 459, http://dx.doi.org/10.1063/1.1730376, URL http://link.aip.org/link/?JCPSA6/31/459/1.

[3] R. Car, M. Parrinello, Unified approach for molecular dynamics and density-functional theory, Phys. Rev. Lett. 55 (1985) 2471, http://dx.doi.org/10.1103/PhysRevLett.55.2471, URL https://link.aps.org/doi/10.1103/PhysRevLett.55.2471.

[4] D. Fenyö, B.U.R. Sundqvist, B.R. Karlsson, R.E. Johnson, Molecular-dynamics study of electronic sputtering of large organic molecules, Phys. Rev. B 42 (1990) 1895–1902, http://dx.doi.org/10.1103/PhysRevB.42.1895, URL http://link.aps.org/doi/10.1103/PhysRevB.42.1895.

[5] K. Nordlund, Molecular dynamics simulation of ion ranges in the 1–100 kev energy range, Comput. Mater. Sci. 3 (1995) 448–456, http://dx.doi.org/10.1016/0927-0256(94)00085-Q, URL http://linkinghub.elsevier.com/retrieve/pii/092702569400085Q.

[6] M. Hou, Z.-Y. Pan, Cascade statistics in the binary collision approximation and in full molecular dynamics, Nucl. Instrum. Methods Phys. Res. B 102 (1995) 93–102, http://dx.doi.org/10.1016/0168-583X(95)80123-4, URL http://linkinghub.elsevier.com/retrieve/pii/0168583X95801234.

[7] E. Gabriel, G.E. Fagg, G. Bosilca, T. Angskun, J.J. Dongarra, J.M. Squyres, V. Sahay, P. Kambadur, B. Barrett, A. Lumsdaine, R.H. Castain, D.J. Daniel, R.L. Graham, T.S. Woodall, Open MPI: Goals, concept, and design of a next generation MPI implementation, in: Proceedings, 11th European PVM/MPI Users' Group Meeting, Budapest, Hungary, 2004, pp. 97–104.

[8] L. Dagum, R. Menon, Openmp: an industry standard api for shared-memory programming, IEEE Comput. Sci. Eng. 5 (1998) 46–55, http://dx.doi.org/10.1109/99.660313, URL http://ieeexplore.ieee.org/document/660313.

[9] K. Kadau, T.C. Germann, P.S. Lomdahl, Molecular dynamics comes of age: 320 billion atom simulation on bluegene/l, Internat. J. Modern Phys. C 17 (2006) 1755–1761, http://dx.doi.org/10.1142/S0129183106010182.

[10] D.E. Harrison, Application of molecular dynamics simulations to the study of ion-bombarded metal surfaces, Crit. Rev. Solid State Mater. Sci. 14 (1988) s1–s78, http://dx.doi.org/10.1080/10408438808244782.

[11] R. Smith, D.E. Harrison, B.J. Garrison, Simulation of kev particle bombardment of covalent materials: An investigation of the yield dependence on incidence angle, Nucl. Instrum. Methods Phys. Res. B 46 (1990) 1–11, http://dx.doi.org/10.1016/0168-583X(90)90661-D.

[12] K.M. Beardmore, N. Grø nbech Jensen, Efficient molecular dynamics scheme for the calculation of dopant profiles due to ion implantation, Phys. Rev. E 57 (1998) 7278–7287, http://dx.doi.org/10.1103/PhysRevE.57.7278, URL http://link.aps.org/doi/10.1103/PhysRevE.57.7278.

[13] C. Fridlund, J. Laakso, K. Nordlund, F. Djurabekova, Atomistic simulation of ion irradiation of semiconductor heterostructures, Nucl. Instrum. Methods Phys. Res. B 409 (2017) 14–18, http://dx.doi.org/10.1016/j.nimb.2017.04.034, URL https://linkinghub.elsevier.com/retrieve/pii/S0168583X17304627.

[14] F. Djurabekova, C. Fridlund, K. Nordlund, Defect and density evolution under high-fluence ion irradiation of Si/SiO_2 heterostructures, Phys. Rev. Mater. 4 (2020) 013601, http://dx.doi.org/10.1103/PhysRevMaterials.4.013601, URL https://link.aps.org/doi/10.1103/PhysRevMaterials.4.013601.

[15] X. Xu, K.-H. Heinig, W. Möller, H.-J. Engelmann, N. Klingner, A. Gharbi, R. Tiron, J. von Borany, G. Hlawacek, Morphology modification of Si nanopillars under ion irradiation at elevated temperatures: plastic deformation and controlled thinning to 10 nm, Semicond. Sci. Technol. 35 (2020) 015021, http://dx.doi.org/10.1088/1361-6641/ab57ba, URL https://iopscience.iop.org/article/10.1088/1361-6641/ab57ba.

[16] L. Marques, M.-J. Caturla, H. Huang, T.D.D.L. Rubia, Molecular dynamics studies of the ion beam induced crystallization in silicon, MRS Proc. 396 (1995) 201, http://dx.doi.org/10.1557/PROC-396-201, URL http://journals.cambridge.org/abstract_S1946427400371433.

[17] C. Krzeminski, Q. Brulin, V. Cuny, E. Lecat, E. Lampin, F. Cleri, Molecular dynamics simulation of the recrystallization of amorphous Si layers: Comprehensive study of the dependence of the recrystallization velocity on the interatomic potential, J. Appl. Phys. 101 (2007) 123506, http://dx.doi.org/10.1063/1.2743089, URL http://aip.scitation.org/doi/10.1063/1.2743089.

[18] K. Nordlund, M. Ghaly, R.S. Averback, M. Caturla, T.D. de la Rubia, J. Tarus, Defect production in collision cascades in elemental semiconductors and fcc metals, Phys. Rev. B 57 (1998) 7556–7570, http://dx.doi.org/10.1103/PhysRevB.57.7556, URL http://link.aps.org/doi/10.1103/PhysRevB.57.7556.

[19] K. Nordlund, J. Keinonen, M. Ghaly, R.S. Averback, Coherent displacement of atoms during ion irradiation, Nature 398 (1999) 49–51, http://dx.doi.org/10.1038/17983, URL http://www.nature.com/doifinder/10.1038/17983.

[20] PARCAS Source, 2021, URL https://gitlab.com/acclab/parcas/.

[21] C. Fridlund, J. Laakso, COSIRMA source, 2021, URL https://gitlab.com/acclab/cosirma.

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

[23] L. Verlet, Computer experiments on classical fluids, I. Thermodynamical properties of lennard-jones molecules, Phys. Rev. 159 (1967) 98–103, http://dx.doi.org/10.1103/PhysRev.159.98, URL http://link.aps.org/doi/10.1103/PhysRev.159.98.

[24] M.P. Allen, D.J. Tildesley, Computer Simulations of Liquids, Clarendon Press, 1987.

[25] T. Watanabe, H. Fujiwara, H. Noguchi, T. Hoshino, I. Ohdomari, Novel interatomic potential energy function for Si, O mixed systems, Japan. J. Appl. Phys. 38 (1999) L366.

[26] J. Samela, K. Nordlund, Origin of nonlinear sputtering during nanocluster bombardment of metals, Phys. Rev. B 76 (2007) 125434, http://dx.doi.org/10.1103/PhysRevB.76.125434, URL https://link.aps.org/doi/10.1103/PhysRevB.76.125434.

[27] K. Nordlund, F. Djurabekova, G. Hobler, Large fraction of crystal directions leads to ion channeling, Phys. Rev. B 94 (2016) 214109, http://dx.doi.org/10.1103/PhysRevB.94.214109, URL https://link.aps.org/doi/10.1103/PhysRevB.94.214109.

[28] J.C. Moreno-Marin, U. Conrad, H.M. Urbassek, A. Gras-Marti, Fractal structure of collision cascades, Nucl. Instrum. Methods Phys. Res. B 48 (1990) 404–407, http://dx.doi.org/10.1016/0168-583X(90)90149-O.

[29] M.T. Robinson, The statistics of sputtering, Nucl. Instrum. Methods Phys. Res. B 90 (1994) 509–512, http://dx.doi.org/10.1016/0168-583X(94)95604-9.

[30] A. Stukowski, OVITO, 2020, URL http://ovito.org/.

[31] F.H. Stillinger, T.A. Weber, Computer simulation of local order in condensed phases of silicon, Phys. Rev. B 31 (1985) 5262–5271, http://dx.doi.org/10.1103/PhysRevB.31.5262, URL https://link.aps.org/doi/10.1103/PhysRevB.31.5262.

[32] J. Tersoff, New empirical model for the structural properties of silicon, Phys. Rev. Lett. 56 (1986) 632–635, http://dx.doi.org/10.1103/PhysRevLett.56.632, URL http://link.aps.org/doi/10.1103/PhysRevLett.56.632.

[33] J. Tersoff, Empirical interatomic potential for silicon with improved elastic properties, Phys. Rev. B 38 (1988) 9902–9905, http://dx.doi.org/10.1103/PhysRevB.38.9902, URL http://link.aps.org/doi/10.1103/PhysRevB.38.9902.

[34] S.J. Cook, P. Clancy, Comparison of semi-empirical potential functions for silicon and germanium, Phys. Rev. B 47 (1993) 7686–7699, http://dx.doi.org/10.1103/PhysRevB.47.7686, URL https://link.aps.org/doi/10.1103/PhysRevB.47.7686.

[35] J.Q. Broughton, X.P. Li, Phase diagram of silicon by molecular dynamics, Phys. Rev. B 35 (1987) 9120–9127, http://dx.doi.org/10.1103/PhysRevB.35.9120, URL https://link.aps.org/doi/10.1103/PhysRevB.35.9120.

[36] S. Yoo, X.C. Zeng, J.R. Morris, The melting lines of model silicon calculated from coexisting solid–liquid phases, J. Chem. Phys. 120 (2004) 1654–1656, http://dx.doi.org/10.1063/1.1633754, URL http://aip.scitation.org/doi/10.1063/1.1633754.

[37] J. Linnros, R.G. Elliman, W.L. Brown, Divacancy control of the balance between ion-beam-induced epitaxial cyrstallization and amorphization in silicon, J. Mater. Res. 3 (6) (1988) 1208–1211, http://dx.doi.org/10.1557/JMR.1988.1208, URL http://link.springer.com/10.1557/JMR.1988.1208.

[38] T. Henkel, V. Heera, R. Kögler, W. Skorupa, M. Seibt, The temperature dependence of the ion beam induced interfacial amorphization in silicon, Appl. Phys. Lett. 68 (24) (1996) 3425–3427, http://dx.doi.org/10.1063/1.115781, URL http://aip.scitation.org/doi/10.1063/1.115781.