Journal of Nuclear Materials 402 (2010) 116–123

Contents lists available at ScienceDirect

Journal of Nuclear Materials

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

Intragranular Xe bubble population evolution in UO_2: A first passage Monte Carlo simulation approach

D. Schwen *, R.S. Averback

Department of Materials Science and Engineering, University of Illinois at Urbana-Champaign, IL 61801, USA

ARTICLE INFO

Article history:
Received 12 March 2010
Accepted 6 May 2010

ABSTRACT

We introduce a first passage based Monte Carlo [1] code to investigate the population evolution of Xe fission gas bubbles in UO_2 fuels. Growth laws are obtained for homogeneous and heterogeneous re-solution models for a wide range of gas and bubble diffusivities. Under certain irradiation conditions, bubble populations find dynamic steady states. Homogeneous re-solution is included using a Monte Carlo binary collision model, while heterogeneous re-solution is modeled as the *ad hoc* destruction of bubbles.

1. Introduction

Understanding and possibly controlling the evolution of the population of intragranular fission gas (fg) bubbles is a key challenge in the development of high burn-up generation IV reactor technologies. The size and spatial distributions of fission gas bubbles in uranium dioxide and mixed oxide fuel elements are key parameters in determining the gas transport properties and the fuel performance. Fission gas retention can adversely influence the thermal and mechanical properties of reactor fuels, depending on the size of the bubbles, while release of gaseous species from the fuel can lead to cladding failure at high burn-up.

In previous publications we looked at the re-solution of fission gas from bubbles, through homogeneous [2] and heterogeneous [3] mechanisms. We employed molecular dynamics and Monte Carlo (binary collision) simulations, since they are feasible for the short timescales of re-solution processes, of the order of picoseconds. The aim of the present work is to extend the modeling capabilities to nucleation and growth of fission gas bubbles, processes with timescales of the order of hours, days or even years.

We developed a code based on first passage Green's functions to track the diffusive motion of individual atoms and bubbles over long timescales. The idea was presented by Opplestrup et al. [1] to treat diffusing objects undergoing two-body reactions. The core idea is that the migrations of objects undergoing successive identical diffusive jumps, with only rare encounters with other walkers, can be treated by allowing many jumps to be computed as a single super-jump. The method guarantees that no walkers interact during each super-jump. We present an overview of the algorithm, but focusing on our implementation of it for calculations of bubble evolution.

Numerous efforts at predicting the release rates and evolution of the size distribution of fission gas bubbles as functions of temperature and burn-up rate are documented in the literature; however, they are all conceptually different from our approach. Existing bubble population models can be partitioned into two groups, single-size and size-distribution models. The former assume that all bubbles have the same or at least very narrowly distributed sizes. Turnbull [4] proposed the first single-size bubble population model for UO_2 motivated by Cornell's [5] electron microscope observation of fission gas bubbles. Both bubbles and gas atoms in solution are characterized by a mean field concentration. Both nucleation and re-solution are assumed to be heterogeneous in the wake of high energetic fission fragments. Speight [6] expanded Turnbull's model to include a size distribution of gas bubbles. Notably absent from Speight's model is bubble coalescence. Bubble growth is entirely driven by re-solution and absorption of fission gas atoms. Lösönen's locally accurate model (LOAM) [7–9] offers another single bubble size model, geared at gas-release calculations. It refines the spherical grain approximation through spatial partitioning of the grains, thus allowing for varying bubble sizes as a function of position inside the grain. However, it similarly neglects bubble diffusion and coalescence.

The second group of models allows for a broad distribution of bubble sizes. Wood [10] uses a rate-theory approach to model the evolution of the size distribution of the fission gas bubbles. The mathematical basis for his approach is a system of differential conservation equations – one for each possible bubble size. A solution requires setting an upper bound for the bubble size to keep the number of equations in the system finite. Wood chooses homogeneous nucleation and a heterogeneous re-solution mechanism, which ignores any possible dependence of the re-solution cross section on bubble size. The model does not take coalescence into account and the parameterization features an unusually high gas atom diffusion coefficient, several orders of magnitude above experimental findings. The size space approach by Rusell [11] fea-

* Corresponding author. Tel.: +1 217 333 1096; fax: +1 217 244 2946.
E-mail address: daniel@schwen.de (D. Schwen).

0022-3115/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.jnucmat.2010.05.006

tures both heterogeneous re-solution and nucleation. The size j_0 of nucleated bubbles is a sensitive input parameter in the model, strongly influencing the final size distribution. The size distribution is assumed to be only slowly changing in time. Coalescence, local fluctuations, and bubbles below the nucleation size j_0 are not considered in this model, leading to bubble sizes that are underestimated.

The above models all have in common some kind of mean field approximation, either over entire grains, or in case of Lösönen's model, over subregions of a grain. Assumptions about mechanisms for diffusion, nucleation, growth, and re-solution are hard-wired into the mathematical formulations. Our model offers far more flexibility to test the interplay of different mechanisms, as it is geared towards parametric studies that help extract specific information from complex experiments and guide the design of new experiments. We do not employ mean field approximations, which allows us to analyze spatial correlations and local concentration fluctuations, which, as we will show, could have a significant influence on the bubble population.

For practical time and length scales relevant to typical fuel operating life, lattice based simulation methods are not feasible. Even movement of a small gas bubble over a distance of just one nanometer can require hundreds or thousands of individual atomic jumps. Instead we treat the gas atoms and bubbles as Brownian random walkers. Typical fission gas concentrations are less than 1%. The small concentration of walkers in this system results in a sparseness of walker interactions. For significant amounts of time the diffusing gas atoms and bubbles occupy disjoint regions of space. First passage Monte Carlo offers a way to replace most of the eventless diffusion jumps by a vastly reduced number of super-jumps. It does so by partitioning the sample volume into non-overlapping sub volumes (protective zones), with each containing only one diffusing walker. The super-jump is given by the direct move of the walker from its original location to the edge of its protective zone. The time associated with this super-jump is the time it takes the walker to reach the edge of the protective zone for the first time, hence *first* passage. This first passage time can be calculated analytically for certain geometries of protective zones. In particular the sphere is a natural choice, as the probability density for the super-jump to go to any point on its surface is uniform in the case of isotropic diffusion. It should be noted that the random walk processes of individual diffusion jumps and the first passage super-jumps are statistically equivalent.

First passage Monte Carlo can offer speed-ups of several orders of magnitude in dilute systems, and thus it enables simulation times up to several years. In contrast to mean field theories, local fluctuations and spatial correlations are included. The model provides simple access to a multitude of statistical data on the simulated systems. The size evolution of individual bubbles can be tracked, average lifetimes computed, individual knockouts can be followed and return probabilities calculated as a function of precipitate size. Our approach offers flexibility to change or switch out underlying mechanisms allowing direct access to the influence of each input parameter and every sub-process in the population evolution.

2. Methodology

Both single fission gas atoms and fission gas bubbles are described as one type of object in our simulation, the so-called *walkers*. Every walker has a set of properties, the most important being the number of atoms contained in the walker, n, and its diffusion constant, D. The value n can range from one for single fission gas atoms, to hundreds or thousands for larger bubbles. For the purpose of this study we take the radius to be a simple function of the number of atoms in the bubble, by assuming a constant fission gas density and spherical shape of the walker. Different functional dependencies for the diffusion constant of each walker can be selected. The models used in this work are discussed below.

While we examine different models for Xe re-solution and diffusion, we do not try to be inclusive. For example, we do not treat thermal re-solution of Xe, although Veshchunov [12] comments that it becomes significant in UO_2 above 1800 K. We also do not include the effects of bubble pressure on diffusion, an effect that is likely to become important at low temperatures and low irradiation fluxes. In this regime, moreover, faceted bubbles, which are believed to move by step nucleation on the bubble surface, have been observed under certain conditions [13,14]. We leave these extensions of our model for a future study.

Periodic boundary conditions are applied to our simulation domain to approximate a bulk material. Absorbing or reflecting boundaries are currently being implemented to treat, for example, thin films or grain boundaries.

The walkers in our code are surrounded by a spherical region, a so-called *protective zone*. The size of each protective zone is constructed using a simple geometrical process to be as large as possible without overlapping any other zone. Per default, each zone contains a single walker. As long as the walker moves inside the zone no walker–walker interactions can occur. To propagate the system it is therefore sufficient to calculate the time it will take the walker to reach the border of its protective zone for the first time, the *first passage time*.

Every protective zone has a set of properties, its radius R, its time of creation t_c, its time of expiration t_e when the contained walker will hit the boundary, and a pointer to its neighboring protective zone.

At every time step the list of protective zones is searched for the nearest expiration time t_e. The nearest expiration time t_e is then compared to the times of the next recoil t_r and the next fission gas atom production t_p; whichever event is soonest is then executed. In case of the event being the expiration of a protective zone, the walker contained in the zone is moved to a random location on the boundary of the zone, and the protective zone is then redrawn around the walker. Each advancement step in time for a protective zone is called *promotion*. The algorithm demands the ability to promote protective zones ahead of their expiration time, which we call *premature promotion*. For example, consider a small zone with a short expiration time; it can get trapped next to a large zone with long expiration time. Chances for the particle in the small zone to move closer to the big zone are about the same as moving away from the big zone; however, every step away is smaller than the preceding step that brought the walker closer, owing to the decreasing distance of the radius of its protective zone. This can lead to an asymptotic behavior with a single walker moving in a shrinking protective zone and consuming all of the CPU time with high frequency updates. A solution was hinted at in Ref. [1]. We implemented it by checking the ratio of expiration times of the last promoted protective zone and its neighbor zone. If this ratio exceeds a pre-determined threshold, the larger zone is promoted prematurely, and its new protective zone is shrunk by 50% from its maximum possible radius, thereby freeing the trapped neighboring walker.

With only one particle per protective zone and no two protective zones ever overlapping, particle interactions are prohibited. We solve this dilemma by performing a simple test after each simulation step. The distance to the neighboring protective zone and the radius of the minimal sphere containing all of the neighboring zone's walkers are calculated. If the difference of these values is smaller than a threshold R_{bunch} both protective zones are prematurely promoted and a new protective zone, termed cooperative protective zone (CPZ), is constructed around the mean of the coor-

dinates of their combined walkers. The radius of this CPZ is chosen so that the walkers from the two zones are on the inside of the new zone and have a distance of at least R_{pad} to the new zone boundary.

This new zone might intersect uninvolved nearby protective zones. In this case all intersected zones must be prematurely promoted and walkers then ending up inside the CPZ will be added to the CPZ. Walkers lying entirely outside are provided with a new, smaller PZ that does not overlap the CPZ. Walkers intersecting the boundary of the CPZ trigger a radius increase of the CPZ, until they are entirely on the inside and have a distance of at least R_{pad} to the zone boundary. This radius increase prompts for a rechecking of intersection, making it an iterative process, which can lead to cascading growth of the CPZ.

The walkers inside the CPZ are then advanced by conventional step-by-step MC until either one of the walkers touches the boundary or two walkers touch each other. The expiration clock of the CPZ is then set to the total time of this conventional MC run. Each of these MC steps is recorded by the software to allow premature propagation of the CPZ by rolling back the walker coordinates to any point in time. In practical situations only a few hundred conventional MC steps are needed for each CPZ and out of hundreds of protective zones only a few are CPZs. In fact, the fraction of CPZ's varies approximately as the square of the concentration of walkers. By varying the bunching threshold R_{bunch} and the CPZ padding length R_{pad} the execution times spent on analytical first passage steps and the times spent on conventional MC steps can be shifted back and forth, enabling an optimum balance to be found for every sample configuration. The mixture of conventional and accelerated first passage Monte Carlo provides added flexibility. For example short ranged ($<R_{bunch}/2$) interactions between walkers can be easily implemented, such as an effective potential barrier a walker has to overcome when entering an over-pressured bubble. Any interactions beyond hard spheres would be very hard to accommodate in a purely first passage formalism.

Prematurely promoting a protective zone means obtaining the coordinates of the contained walkers ahead of its expiration time. For multi-walker zones this is achieved by storing a history of all classical Monte Carlo steps, allowing for a replay to any intermediate time step. For single walker zones this is achieved by a simple scaling of the zone radius according to the square root of the passed fraction of the expiration time.

For improving the scalability of our simulation code three tasks deserve mention. The first one involves the temporal sorting of the protective zones in the simulation. At each regular time step the protective zone with the smallest (i.e. closest) expiration time has to be determined. This is best achieved through the implementation of the abstract data type *priority queue*. We store pointers to all protective zones in a binary tree, whereby the pointer in the root node of the tree always points to the protective zone that expires next.

The second task consists of selecting a random candidate walker for a re-solution process. Depending on the mechanism, the walkers are assigned different relative probabilities for participation in a re-solution event. We augment the aforementioned binary tree by storing in each node the sum of all weights in the sub-tree rooted at that node. A walker is then randomly chosen by obtaining a random number p between zero and the sum of all walker weights, and walking the tree from the top. At each node a tri-valued choice is made to determine whether to select the current node, or continue down the tree either to the left or right child node, decrementing p accordingly. Every change in weight to a walker triggers an update of the tree which is of $O(\log(N))$. The advantage is the reduction in complexity for choosing a walker from $O(N)$ to $O(\log(N))$. For systems where the number of walker weight changing reactions, e.g., coalescence, does not drastically exceed the number of re-solution events, a substantial speed-up can be achieved.

The third task is finding the maximum non-overlapping radii for the protective zones, and determining if a newly inserted particle appears in an existing protective zone. A spatially sorted index of all protective zones is kept in a modified link-cell data structure. This index is consulted whenever a proximity search for protective zones is performed to obtain the distance to the closest neighboring protective zone surface. As opposed to the common application of link cells in molecular dynamics simulations, no fixed cut-off distance exists for the neighbor search. Link cells must be searched in an outward spiraling pattern until at least one neighbor is found and even then further cells must be searched until the existence of further potentially closer neighbor zones can be safely excluded. Protective zones up to a maximum size r_m are sorted into the link cells. Zones exceeding the threshold are sorted into a separate large zone list. By capping the maximum radius inside the link cells, an upper bound can be put on the search radius.

To avoid rejection sampling and frequent evaluation of the first passage time distribution (see Eq. (1) and Table 1 for definition of terms) we chose inverse transform sampling to obtain the random first passage times. From the derivation above it is evident that the cumulative distribution function (CDF) of the first passage probability is the complement to the survival probability. Since an explicit inversion of this function is not feasible, we tabulate its values by evenly sampling its codomain [0;1), the interval from zero to one, excluding one. As the CDF is monotonous an efficient binary search can be performed. We tabulate several thousand data points. This tabulation need only be done for the first run of our code; the results can be stored and used in later runs. The tabulated sampling function is linearly interpolated, yielding a fast source of correctly distributed random first passage times.

$$P(t) = 1 - \frac{R}{\sqrt{\frac{D}{\pi}t}} \sum_{n=-\infty}^{\infty} e^{-R^2\frac{\left(\frac{1}{2}-n\right)^2}{Dt}} \tag{1}$$

In our calculations, the diffusion constant of each walker is taken strictly as a function of its radius, with the radius in turn being a function of the number of contained atoms, only. In this work we focus on a single functional dependence, namely that D is proportional to r^{-4}. This dependence is representative of surface diffusion, with matrix atoms performing random jumps along the surface of the fission gas bubbles. This mechanism is assumed to be dominant for fission gas bubbles in UO_2 at temperatures above 1500 K [15]. In our simulations we set the diffusion coefficient of single-atom walkers, which automatically determines the diffusion coefficient of bubbles of all sizes. It should be noted that this constitutes a simplification of the real physics, as single fission gas atoms, unlike bubbles, are likely migrating through vacancy jumps rather than surface diffusion. Furthermore, the surface diffusion in small bub-

Table 1
Quantities used in the first passage Monte Carlo fission gas bubble population evolution simulation algorithm.

Quantity	Symbol
Walker diffusion constant	D
Number of atoms per walker	n
Fission gas atomic volume	Ω_{fg}
First passage time probability distribution	$P(t)$
Protective zone radius	R
Threshold distance for protective zone combination	R_{bunch}
Minimum protective zone radius needed to fully enclose all contained walkers	R_{min}
Padding added to new multi-walker protective zones	R_{pad}
Walker radius	r
Protective zone creation time	t_c
Protective zone expiration time	t_e

bles is believed to be suppressed by the presence of gas atoms, which block jump target sites along the bubble surface [16]. Preliminary simulation results show that the single gas atom diffusivity has an influence on the concentration of fission gas in solution. The reduced retention time in solution, caused by a higher diffusivity, will reduce the nucleation rate of new bubbles. Suppression of small bubble mobility, however, will hamper bubble growth by coalescence, which we identify as the dominant coarsening mechanism. The analysis of this interplay will be subject of a later work. A more refined treatment of the gas atom and bubble diffusion behavior can easily be added to the simulation code. Trapping of Xe atoms, which is not thought to be relevant for in-pile conditions [17], is also not presently implemented in the code.

We included two basic re-solution mechanisms in the simulations, homogeneous and heterogeneous re-solution. Homogeneous re-solution is implemented as a single-atom knockout process; multiple knockouts are negligible for homogeneous re-solution [2]. Using the binary collision Monte Carlo Software 3dTRIM [2] we calculated the fission gas atom recoil energy spectrum per fission event, averaged over the energy and mass distribution of ^{235}U fission fragments. The recoil rate was then taken as the number of recoils per fission event multiplied with the fission rate. Previous work [2] using binary collision Monte Carlo and molecular dynamics simulations showed that recoils with energies less than 100 eV can be excluded since they do not escape the bubbles. The times between recoil events are sampled from a Poisson distribution and the recoil energies are then sampled from the normalized recoil energy distribution. For every recoil event, a walker is chosen randomly, applying the number of atoms in each walker as a probability weight. A random location inside the walker is then selected as the starting point of the fission gas recoil. A cascade is then launched using the 3dTRIM algorithm using a random recoil direction. As soon as the remaining energy of the fission gas recoil falls below a threshold value (100 eV) the cascade stops and the final position of the recoil atom is returned to the main program. A new single-atom walker is inserted at this position and the walker of origin is shrunk by one atom. Notably this procedure is also applied to handle recoils of gas atoms in solution. While this does not alter the bubble size distribution directly, it does add to the diffusion of gas atoms.

The heterogeneous re-solution is implemented as the *ad hoc* destruction of entire bubbles. We assume that re-solution is caused by the electronic energy loss S_e of fission fragments passing close to bubbles. Thus the re-solution rate depends on the fission rate and the average path length with which fission fragments are traveling with an electronic energy loss above a certain threshold value. Again we use the binary collision Monte Carlo Software 3dTRIM, which also treats the electronic stopping component, to average over all path lengths with S_e above a threshold value of ions sampled from the energy and mass distribution of ^{235}U fission fragments. To locate a bubble for re-solution we assume the re-solution cross section to be proportional to the square of the bubble radius [6]. The selected walker is again prematurely propagated to obtain its specific location. A sphere with a diameter of 10 nm is then drawn around the walker location and the atoms in the walker are redistributed randomly inside this re-solution volume. The

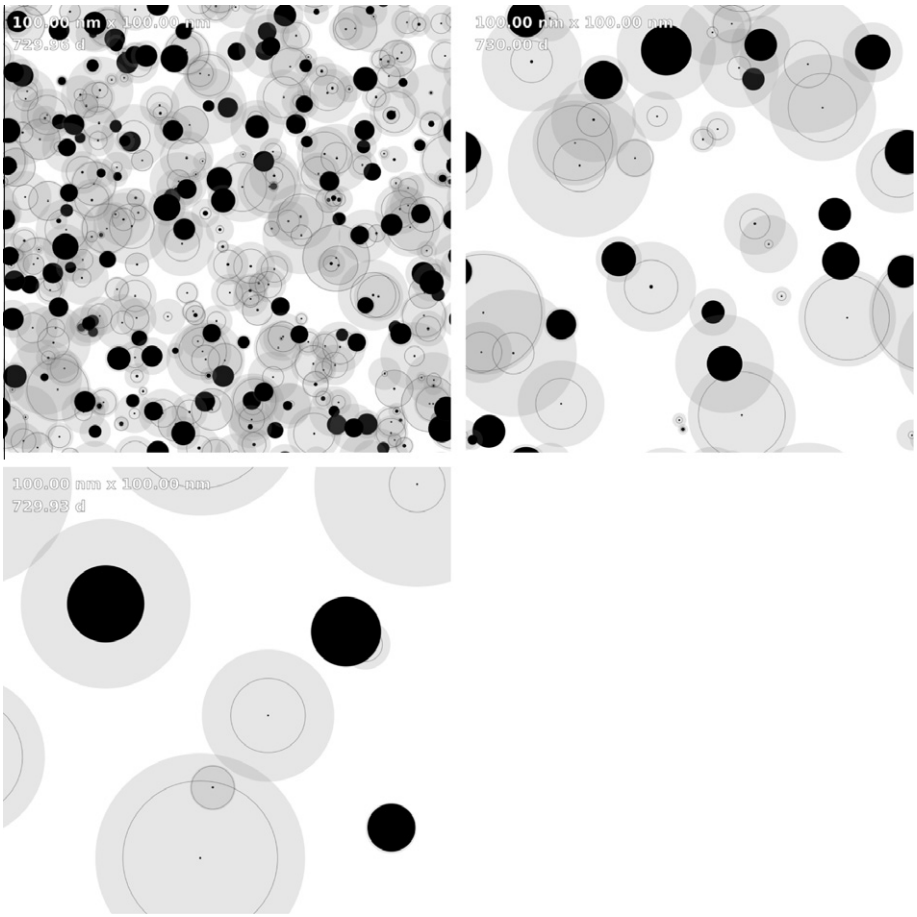

Fig. 1. Snapshots of three $(100\,\mathrm{nm})^3$ simulation cells taken at 1–2 years of simulated time at constant gas accumulation of $2.75 \times 10^{-9}\,\mathrm{Xe\,nm^{-3}\,s^{-1}}$ and homogeneous re-solution for varying diffusion constants D_0 (from left to right $10^{-19}\,\mathrm{m^2\,s^{-1}}$, $10^{-18}\,\mathrm{m^2\,s^{-1}}$, and $10^{-17}\,\mathrm{m^2\,s^{-1}}$) Black disks are fission gas bubbles and atoms, light gray disks indicate the protective zones, and the dark circles indicate the progress of each protective zone as the distance each particle would be displaced in case of a premature update of the protective zone. The average bubble radius clearly grows with increasing diffusivity.

size of the re-solution volume is guided by the size of the molten zone along the fission fragment track as calculated using a two-temperature model [18,19].

Care must be taken to recognize and potentially avoid finite size effects in the simulation runs. An obvious limitation for the maximum attainable bubble size is given by the total number of fission gas atoms in the simulation volume. Should growth dominate over re-solution and gas generation, eventually nearly all gas atoms would coalesce into one large bubble, which then would then have no further means to grow. This can also be re-interpreted as a limitation on the maximum reachable simulation time. One way of overcoming such a limitation is by dynamically growing the simulation volume and thus increasing the amount of gas in the simulation cell. We implemented a simple algorithm into our software which increased the dimensions of the simulation cell in integer multiples and fills the added space with replicas of the original simulation volume. The artificial periodicity thereby introduced quickly disappears as the replicas continue to evolve randomly and separate from each other. This allows us to keep the number of walkers manageable by only growing the cell when the number of walkers drops below a fixed threshold.

3. Simulation results

Selected simulation results are now presented to illustrate the quite different behaviors of bubble evolution that result from homogeneous and heterogeneous dissolution, while at the same time demonstrating the broad scope of the method. We ran simulations with continuous fission gas production, as well as runs with a fixed fission gas concentration. Simulation cell sizes range between $(100\,\mathrm{nm})^3$ and $(200\,\mathrm{nm})^3$, periodic boundary conditions were applied in every run. The diffusion coefficient D for walkers containing n atoms was assumed to be $D_0 n^{-4/3}$, which is proportional to r^{-4}. The single Xe atom diffusion coefficient, D_0, in this model, therefore, determines the diffusion coefficients of all bubbles. The walker radius was calculated, assuming constant density, as $(3\Omega_{fg} n/4\pi)^{1/3} + r_{rec}$. For simplicity the recombination distance r_{rec} was chosen to be zero.

The three viewgraphs in Fig. 1 illustrates the general features of the code; they show snapshots of three different $(100\,\mathrm{nm})^3$ simulation cells taken at up to 2 years of simulated time at the constant gas accumulation rate of $2.75 \times 10^{-9}\,\mathrm{Xe\,nm^{-3}\,s^{-1}}$ and homogeneous re-solution. The diffusion constant D_0 is increasing from left to right with values of $10^{-19}\,\mathrm{m^2\,s^{-1}}$, $10^{-18}\,\mathrm{m^2\,s^{-1}}$, and $10^{-17}\,\mathrm{m^2\,s^{-1}}$. Fission gas bubbles and atoms are drawn as black disks, while light gray disks indicate the protective zones. Shown are two dimensional projections of the three dimensional simulation cell. No two protective zones are overlapping in space. The dark circles, concentric to the protective zones, indicate the spherical surface to which a given particle would jump, if the protective zone were updated ahead of its expiration time. Its radius is given by the size of a hypothetical protective zone that would yield an expiration time coinciding with the time of the snapshot, taking the same random sample from the first passage probability distribution as is used for the actual protective zone. The average bubble radius clearly grows with increasing diffusivity.

We first investigated the stability of various bubble populations under irradiation by preparing several initial bubble populations. These samples were then subjected to fission fragment irradiation, but without generating additional fission gas. All states employed the same gas concentration of $5.6 \times 10^{25}\,\mathrm{m^{-3}}$ (0.23 at.%) at different levels of precipitation, ranging from a random solution of single gas atoms, to bubbles with a radius of 6 nm. The irradiation condi-

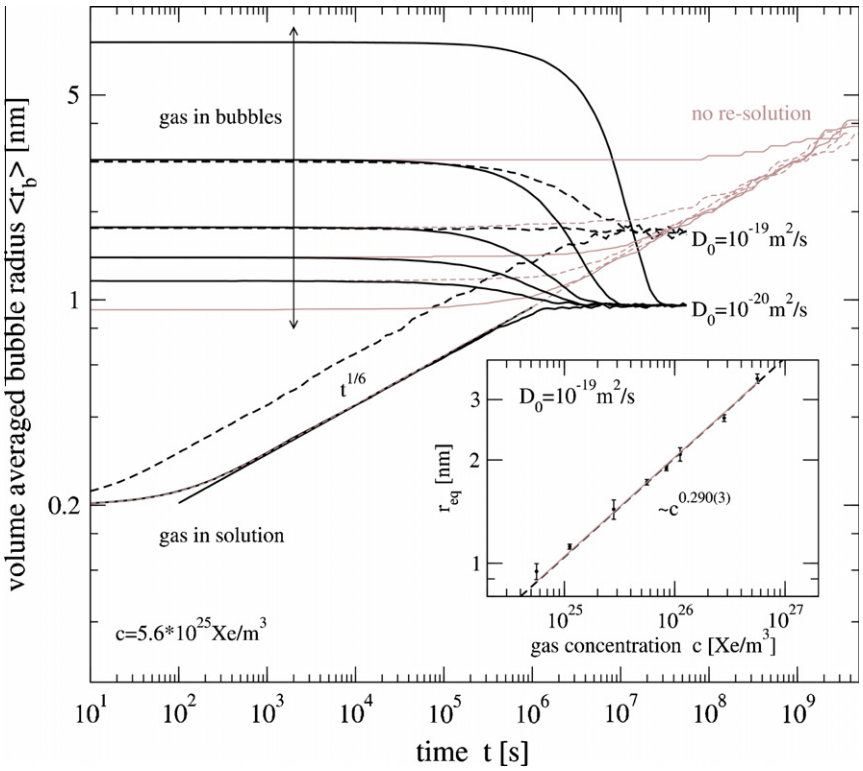

Fig. 2. Evolution of the volume-averaged bubble radius as a function of time for various initial configurations and two different diffusivities. Keeping the total Xe concentration constant at $5.6 \times 10^{25}\,\mathrm{Xe/m^3}$ (0.23 at.%) the gas was randomly distributed as bubbles with radii between 1 nm and 6 nm, and as single gas atoms in solution. Black curves are obtained from simulation runs with homogeneous re-solution, while the gray curves show results from runs with the re-solution turned off. All initial configurations evolve towards an equilibrium configuration with the same mean radius depending on the diffusion constant D_0. The inset shows equilibrium bubble radii as a function of the gas concentration for a diffusion constant D_0 of $10^{-19}\,\mathrm{m^2\,s^{-1}}$. The data can be fitted with the equation $r_{eq} = r_0 * c^{\alpha}$ with the exponent $\alpha = 0.290(3)$.

tions correspond to a fission rate of 10^{-8} nm^{-3} s^{-1}. Control simulations were run for each initial state with the re-solution mechanism deactivated. The evolutions of the volume-averaged bubble radius as a function of time for various initial configurations and two different diffusivities are plotted in Fig. 2. The black curves are obtained from simulation runs with homogeneous re-solution, while the gray curves show results from the control runs with re-solution turned off. With homogeneous re-solution included, all initial configurations evolve toward a steady state configuration with the same volume-averaged radius depending on the diffusion constant D_0. The inset shows these steady state bubble radii as a function of the gas concentration for a diffusion constant D_0 of 10^{-19} m^2 s^{-1}. The data can be fitted with the equation $r_{eq} = r_0 * c^{\alpha}$ with the exponent $\alpha = 0.290(3)$. A similar steady state patterning for immiscible alloys under irradiation has been theoretically predicted by Enrique and Bellon [20,21] and simulated using kinetic lattice Monte Carlo by Heinig et al. [22]. Our model, however, is considerably different, since we do include particle coalescence as a growth mechanism, but do not treat thermal re-solution.

We next look at growth laws for fission gas bubbles under constant fission gas production, assuming typical reactor conditions with Xe being generated at a rate of 2.75×10^{-9} nm^{-3} s^{-1}. The volume-averaged bubble radius as a function of time for two different re-solution mechanisms and various diffusion constants is shown in Fig. 3. The scaling of the diffusion constants, which is related to the sample temperature, is reflected in a scaling of the volume-averaged bubble radius. The growth laws clearly depend on the re-solution mechanism. For heterogeneous re-solution we observe a $t^{1/6}$ dependence, independent of diffusivity. The homogeneous re-solution case exhibits bubble growth proportional to t^{β}, with β ranging from 1/3 to 1/4. At constant gas production, and thus proportionality between time t and gas concentration c, a value of 1/3 for β requires a constant bubble number density, while lower values require an increase in bubble number density, and thus the nucleation of new bubbles. For a single-atom diffusion constant D_0 of 10^{-19} m^2 s^{-1}, a fit of the slope yields an exponent β of 0.29, which, notably, is identical to the exponent α in the relation of equilibrium bubble radius and gas concentration, shown inset in Fig. 2. The bubble population, therefore, must be constantly in steady state. The curve labeled r_{max} denotes the maximum possible bubble size attainable in the simulation volume by aggregating all available Xe atoms into a single bubble. Care must be taken

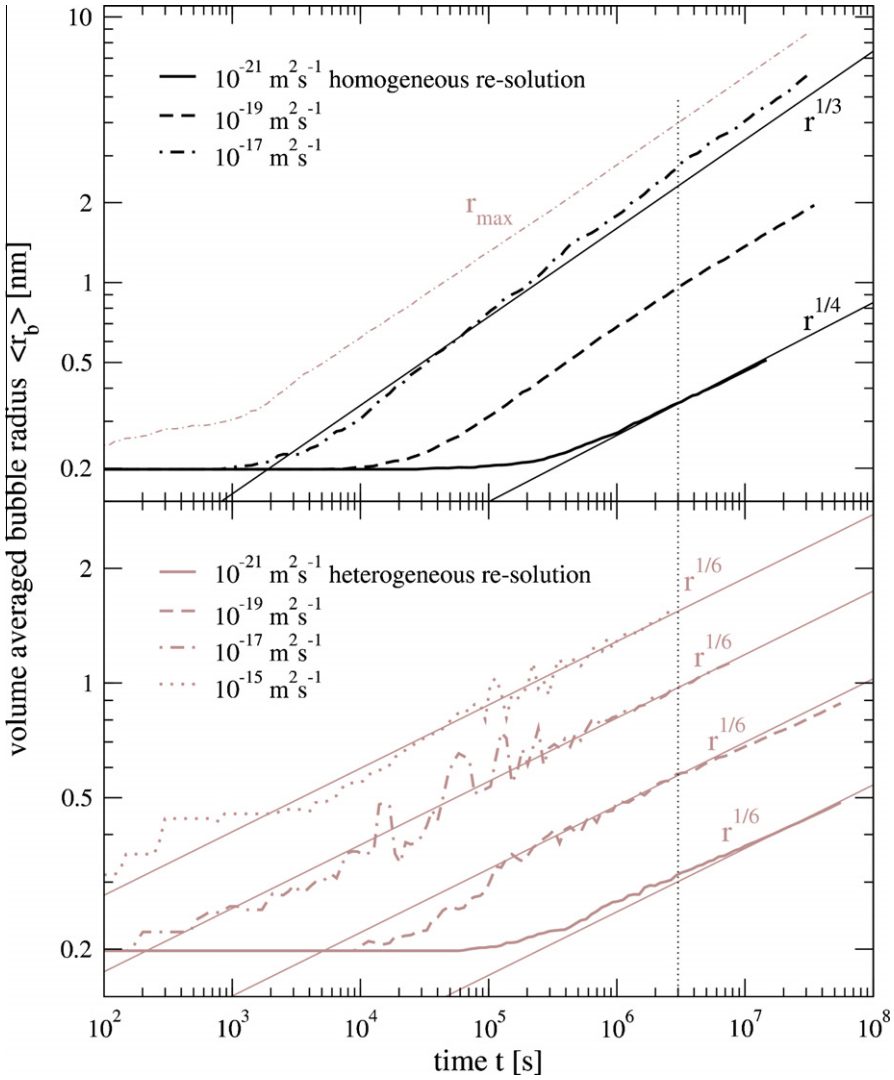

Fig. 3. Volume averaged bubble radius as a function of time, for two different re-solution mechanisms and various diffusion constants. Fission gas is added at a constant rate of Xe/s. The growth law depends strongly on the re-solution mechanism, yielding a $t^{1/6}$ dependence for the heterogeneous re-solution case and $t^{1/3}$ as well as $t^{1/4}$ dependencies for the homogeneous re-solution case. The curve labeled r_{max} denotes the maximum possible bubble size attainable in the simulation volume by aggregating all available Xe atoms into a single bubble. The scaling of the diffusion constants, which is related to the sample temperature, is reflected in a scaling of the volume-averaged bubble radius.

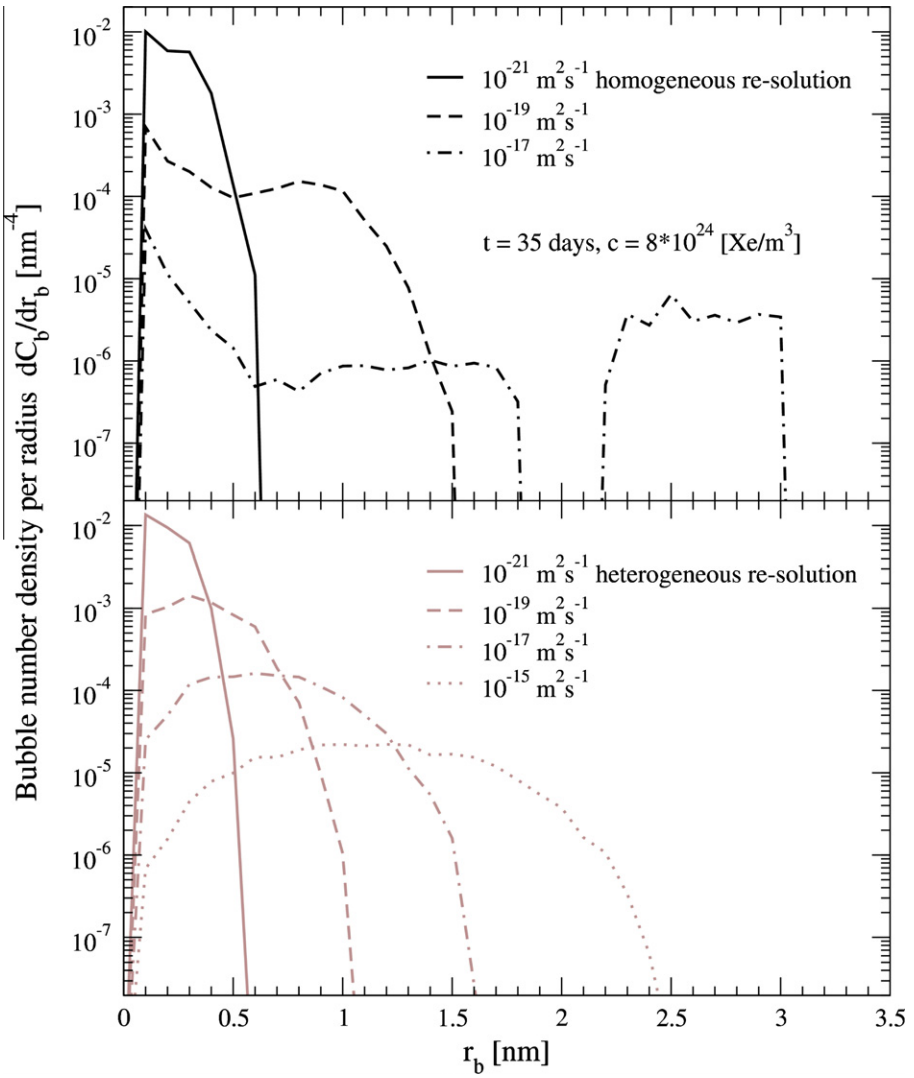

Fig. 4. Bubble size distributions for the homogeneous and heterogeneous re-solution cases as a function diffusion constant. All histograms were taken at $t = 3 \times 10^6$ s. The size distributions for the heterogeneous re-solution runs (gray curves) are clearly peaked around an average radius, as opposed to the histograms of homogeneous re-solution runs (black curves) which exhibit a maximum at the single gas atom radius. The amount of single gas atoms in solution is comparable for both re-solution mechanisms at the same diffusivities.

to ensure that final size effects do not influence the results and that the average bubble size stays well below r_{max}. We verified that the results are indeed identical with an enlarged simulation box size.

As the volume-averaged bubble size provides only a limited description of the system, we examined the bubble size distributions more closely. Bubble size distributions for the homogeneous and heterogeneous re-solution cases as a function of diffusion constant are plotted in Fig. 4. All size distributions were taken at $t = 3 \times 10^6$ s, $\sim$35 days. The size distributions for the heterogeneous re-solution runs (gray curves) are peaked around an average radius, in contrast to the histograms of homogeneous re-solution runs (black curves) which exhibit a maximum at the single gas atom radius. The concentration of single gas atoms in solution, however, is comparable for both re-solution mechanisms at the same diffusivities. The different cross sections in the homogeneous and heterogeneous re-solution mechanisms respectively favor and suppress the development of large bubbles.

4. Summary and conclusions

We demonstrated the applicability of a first passage method to the problem of fission gas bubble evolution. Rather than describing a specific model for the evolution of the bubble population, we have provided a framework that enables a systematic analysis of the impact of various model aspects. In this work we explored bubble population evolution as a function of gas concentration and diffusivity for two different re-solution mechanisms. We find the interesting results that for homogeneous re-solution, and a fixed quantity of Xe gas, the average bubbles size acquires a steady state value that depends on the bubble diffusion coefficient but not the initial size distribution, i.e., the steady state is independent of the initial conditions. We also find that for cases including Xe gas generation, the bubble size growth laws differ significantly for homogeneous and heterogeneous re-solution. These findings should guide appropriate experiments and suggest suitable observables. In this regard, an experimental determination of the bubble growth law should lead to the selection of a re-solution mechanism which is consistent with the observations. The principal advantages of this model over the rate-theory approach are the treatment of coalescence and fluctuations, and no upper boundary for the bubble size has to be assumed. The return probability for re-solved atoms, moreover, is calculated explicitly; it does not need to be approximated with the usual r/R isolated particle approximation, and it includes the possibility for gas atoms to recoil back into the bubble.

The simulation method is in principle applicable to other systems like helium in metal, or dilute immiscible alloys such as CuMo, CuW or oxide dispersed steels.

Acknowledgments

The research was supported by the US Department of Energy, under NERI-C Award No. DE-FG07-14891, and by the US Department of Energy, Basic Energy Sciences under Grant DE-FG02-05ER46217. The authors would like to thank Prof. M. Junge of the Department of Mathematics at UIUC for fruitful discussions.

References

[1] T. Opplestrup, V. Bulatov, G. Gilmer, M. Kalos, B. Sadigh, Phys. Rev. Lett. 97 (2006) 10–13.
[2] D. Schwen, M. Huang, P. Bellon, R. Averback, J. Nucl. Mater. 392 (2009) 35–39.
[3] M. Huang, D. Schwen, R. Averback, J. Nucl. Mater. 399 (2010) 175–180.
[4] J. Turnbull, J. Nucl. Mater. 38 (1971) 203–212.
[5] R.M. Cornell, J. Nucl. Mater. 38 (1971) 319–328.
[6] M. Speight, J. Nucl. Mater. 38 (1971) 236–238.
[7] P. Lösönen, J. Nucl. Mater. 304 (2002) 29–49.
[8] P. Lösönen, Nucl. Eng. Des. 196 (2000) 161.
[9] P. Lösönen, Nucl. Eng. Des. 201 (2000) 139.
[10] M. Wood, J. Nucl. Mater. 82 (1979) 264.
[11] K. Russell, Acta Metall. 19 (1971) 753.
[12] M. Veshchunov, J. Nucl. Mater. 277 (2000) 67–81.
[13] L.F. Willertz, P.G. Shewmon, Metall. Trans. 1 (1970) 2217.
[14] C. Baker, J. Nucl. Mater. 1 (1977) 117–123.
[15] M. Veshchunov, V. Shestak, J. Nucl. Mater. 376 (2008) 174–180.
[16] E.Ya. Mikhlin, J. Nucl. Mater. 4 (1979) 405–408.
[17] H.J. Matzke, Radiat. Eff. Defects Solids 53 (1980) 219–242.
[18] M. Toulemonde, Ch. Dufour, A. Meftah, E. Paumier, Nucl. Instrum. Meth. Phys. Res. B 166–167 (2000) 903.
[19] M. Huang et al., J. Nucl. Mater. (2010). doi:10.1016/j.jnucmat.2010.01.015.
[20] R.A. Enrique, P. Bellon, Phys. Rev. Lett. 84 (2000) 2885.
[21] R.A. Enrique, P. Bellon, Phys. Rev. B 63 (2001) 134111.
[22] K. Heinig, T. Müller, B. Schmidt, M. Strobel, W. Möller, Appl. Phys. A 77 (2003) 17–25.