Prediction of Bubble and Cavity Nucleation in High Damage Rate Irradiation of Simulated Fe-Cr Alloys Using a Hybrid Cluster Dynamics Model

By

Gerrit L VanCoevering

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Nuclear Engineering and Radiological Sciences)
in The University of Michigan
2020

Doctoral Committee:

Professor Gary S. Was, Chair
Professor Fei Gao
Professor Amit Misra
Professor Brian D. Wirth, University of Tennessee

Gerrit L VanCoevering

gvanc@umich.edu

ORCID iD: 0000-0003-1959-1253

Acknowledgments

Foremost among those I must recognize is my advisor Professor Gary Was, who kindled my interest in this field in this classroom and nurtured it by taking me in when I was something of a graduate school orphan. Most importantly, he never gave up on me, even when I had begun to give up on myself. I would also like to distinguish Professor Brian Wirth whose advice, along with his research group, made it possible to be a modeling student in an experimental group. I would also like to the thank the remainder of my committee, Professor Fei Gao and Professor Amit Misra whose input has significantly contributed to the quality of this thesis.

Over my somewhat long tenure in the Was research group, I had the privilege of working alongside a bevy of fantastic scientists. I would like to thank Dr. Stephen Taller in particular, who, as my cohort and project-mate, has provided countless valuable conversations and feedback over the years. He is the hardest working person I know, and his passion for this field is contagious. I would also like to thank Dr. Drew Johnson, who has given me not only a couch to crash on and many a beer to drink but has been consistently able to lift my spirits through commiseration and pedantic arguments. I didn't spend much time in the laboratory, but I was always to happy to see the rotating cast of Was group students and postdocs around the office over the years: Dr. Shyam Dwaraknath, Dr. Elizabeth Getto, Dr. Anthony Monterrosa, Dr. Tyler Moss, Dr. Kale Stephenson, Dr. Anne Campbell, Dr. Stephen Raiman, Dr. Justin Hesterberg, Dr. Michael Higgins, Rigel Hanbury, Dr. David Woodley, Samara Levine, Katey Thomas, Dr. Peng Wang, Dr. Miao Song, Dr. Mi Wang, Dr. Calvin Lear, Dr. Wenjun Kuang, Dr. Donghai Du, Dr. Kai Chen, Valentin Pauly, and Logan Clowers.

I was also fortunate to have several extremely valuable conversations with collaborators from other institutions thanks to my involvement with the Simulating Neutrons with Accelerated Particles (SNAP) project and a summer internship for the IUP fellowship at Oak Ridge National Labortory. I owe a great deal to Dr. Aaron Kohnert, Dr. Roger Stoller, and Dr. Alex Barashev in particular, whose guidance and feedback over the years were critical to the quality of this work.

I would like to thank my parents, Jim and Gail Van Coevering, for boosting my confidence to almost irresponsible levels and instilling me with lust for lifelong learning. Not to mention that I owe them my debt-free status. I am truly spoiled, and I could not be more grateful.

Finally, I must acknowledge my favorite ladies Shannon and Matilda. Most this work was done with Matilda laying against my legs and sometimes licking my toes. That sounds strange, but it is important to note that she is a dog and that is normal for them. The human woman Shannon, on the other hand, has provided an anchor of emotional intelligence, companionship, zesty tattoos, and adorable looks of contempt that I could not do without. She may not have contributed scientific merit to this thesis, but she kept me in the right mindset to complete it, and I have every intention of making her my wife someday.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Abstract

A new generation of safer, more efficient nuclear reactors is perhaps the most viable way to combat climate change by weening the power grid from carbon-laden fossil fuels. However, many of these advanced designs require high temperatures, corrosive environments, and large irradiation doses which degrade structural components. Development of radiation-tolerant materials is crucial, but conventionally requires long, expensive test reactor irradiations which activate the sample, making examination difficult and costly as well. Ion irradiation experiments present an extremely appealing alternative, as they can introduce decades of radiation damage in a matter of days with minimal activation. Unfortunately, these accelerated testing techniques come with a cost of several experimental artefacts which must be accounted for, the most fundamental of which is that damage rate itself has a tremendous impact on its accumulation. Material degradation on an engineering scale is the result of an extraordinary number of individual atomic-scale picosecond damage events generating defects that migrate, coalesce, and dissociate in just nanoseconds continuously over decades of reactor operation. Clever kinetic models that bridge this massive gap in length and time scales are necessary to gain physical insight into radiation damage which will aid in the development of candidate materials designed on a microscopic scale to minimize defect accumulation. This gap is particularly pronounced for the slow, but often life-limiting phenomenon of swelling, where vacant lattice sites eventually accumulate into sizable cavities which uniformly increase the volume of a component beyond its tolerance.

The objective of this work is to use cluster dynamics modeling to better understand the temperature, helium generation, and dose rate dependencies of swelling behavior in simulated Ferritic-Martensitic (FM) steels. An initial study using a conventional modeling approach based on reaction rate theory was fundamentally incapable of reproducing several key experimental swelling observations, but most of these were recovered by iteratively adding novel physical mechanisms. Foremost among these discrepancies was an overprediction of swelling rates, a complete lack of incubation periods, and an almost negligible effect of co-generated helium. Introducing a bias toward interstitials to cavities informed by molecular statics calculations,

presented a new barrier to cavity nucleation that could only be overcome through the accumulation of stabilizing helium. Since FM steels have low transmutation rates, helium generation rates are relatively low, resulting in long incubation periods and low swelling rates. Second, the temperature shift—an offset in optimal swelling temperature between high and low dose rate conditions—was overpredicted. A heterogeneous nucleation mechanism, promoting helium self-clustering, weakened dose rate dependence by disproportionately facilitating cavity nucleation at low dose rate, decreasing temperature shift. Finally, observations of swelling without helium were rationalized with bias suppression catalyzed by segregation of chemical impurities to the cavity surface.

These mechanisms ultimately combine to constitute a more complete, nuanced understanding of cavity nucleation as a function of temperature, dose rate, and helium generation rate. This creates a more complete account of how microstructures evolve under irradiation and strengthens the predictive link between accelerated ion irradiation experiments and the reactor conditions they seek to emulate. Therefore, in addition to contributions to the fundamental science of irradiation damage, this model could inform the design of future experiments to optimize their investigative potential, hastening the qualification of candidate materials.

Chapter 1: Introduction

Nuclear technology is the most reliable source of carbon-free power capable of combatting climate change by supplanting dirty fossil fuels. However, the current fleet of reactors may be approaching the end of their already extended lifespans. Between a strict regulatory environment in the wake of high-profile accidents, and economic competition from a boom in natural gas, safety, and efficiency requirements for the next generation of reactors have never been higher. From the standpoint of neutronics many adequate designs exist, but they often require higher operating temperatures, more intense radiation exposure, and corrosive environments for structural components. Understanding how these extreme conditions affect the engineering properties of materials is a crucial step toward designing a reactor which can tolerate them.

As high energy particles strike a material, they dislodge atoms from their lattice sites, which strike and dislodge neighboring atoms in a chaotic event called a damage cascade. The result is that some lattice sites are left empty (vacancies), while the atom it previously held is wedged between other sites (interstitials). As damage events continuously occur, the point defects produced begin to aggregate into larger defects, such as dislocation loops, stacking fault tetrahedra, voids, and bubbles. Furthermore, the resulting disruption of thermodynamic equilibrium can cause redistribution of chemical species as phases form or dissolve and impurities segregate to grain boundaries and extended defects. Depending on the material and its operating temperature, these processes can lead to failure when mechanical properties are degraded by embrittlement or creep, the dimensions of the material are changed by swelling or growth, or corrosion is accelerated by radiolysis and segregation.

Therefore, material failure on an engineering scale is the result of an astronomical number of individual atomic-scale picosecond damage events generating defects that migrate, coalesce, and dissociate in just nanoseconds continuously over decades of reactor operation. To mitigate these deleterious changes to physical and mechanical properties, materials must be designed on a microscopic scale to minimize defect accumulation. Doing so will require a thorough understanding of defect behavior to bridge the enormous length and time scale gap between atomic

level solid-state physics and engineering experience with candidate materials. Insofar as these processes cannot be controlled, reactor designs need to thoroughly understand them to tolerate negative effects.

Most experimental insight into radiation damage comes from post-irradiation examination of samples inserted into a test reactor. These studies are exceptionally long, expensive, and induce hazardous radioactivity. Furthermore, the reactor environment is difficult to precisely control, resulting in variable and uncertain temperature and dose histories. Irradiation experiments conducted using accelerated protons or heavy ions present an appealing alternative, offering extremely fast damage production in a controlled laboratory setting with minimal sample activation. These myriad benefits come at the cost of several experimental artefacts which must be mitigated. Foremost among these, from the perspective of defect kinetics, is that accelerated damage production causes major changes to its accumulation.

One strategy, proposed by Mansur, attempts to counteract the effect of dose rate by systematically changing another irradiation variable, temperature. Despite being built upon several assumptions which reduce it to a simple analytic form, this approach has proven extremely successful in experiments conducted on pure metals and austenitic steels. It is extremely appealing to use a similar irradiation variable shift, ideally in all metals, to establish a predictive link between ion irradiations and the reactor conditions they seek to emulate. If the prediction of microstructure is sufficiently rigorous for regulators, it could dramatically hasten qualification of candidate materials, and thus the regulatory approval, and eventually construction, of the next generation of reactors.

Unfortunately, the simplifying assumptions implicit to the Mansur temperature shift miss several key physical components of swelling behavior. To start, it is only valid for cavity growth, which neglects the much more complex physical process of nucleation. Particularly for alloys which have sizable incubation periods, in which the nucleation regime extends significantly into the lifetime of the component, this significantly hampers its applicability. Furthermore, the time-evolution of irradiation damage is limited only to the production, recombination, and absorption of single point defects. Modeling on the atomic scale suggests that a more realistic picture is complicated by production and mobility of defect clusters, the co-evolution of competing

microstructural features, and complicated interactions with atomic impurities that stabilize or immobilize defects and their clusters.

Indeed, recent experiments suggest that this temperature shift strategy may be particularly insufficient for Ferritic-Martensitic (FM) steels, which have emerged as a preeminent candidate for high dose applications due to their swelling resistance and high-temperature stability. This belies the need for more rigorous models of defect kinetics that contain a more complete account of the microstructure and generate predictions of its evolution. While predictive capacity of these simulations alone may not be sufficient to qualify reactor materials, they can be coordinated with experiments to inform experimental design to maximize relevance to reactor conditions. Furthermore, this approach can assist in recontextualizing the existing bodies of accelerator and test reactor irradiation data, and through more fundamental understanding of damage effects, inform design of even more tolerant materials.

This work will use mesoscale defect modeling in pursuit of three ends. First, to enhance the fundamental understanding of defect and defect cluster behavior to better predict the evolution of damage phenomena. Second, to bridge fundamental knowledge of defect properties at an atomic scale with experimental observations to discern their implications on defect kinetics. Third, to quantify the effect of dose rate on damage accumulation to assess the viability of ion irradiation as a surrogate for reactor experiments in the FM system and provide suggestions to optimize experimental design.

Chapter 2: Background

The discovery of irradiation-induced void swelling in the 1960's presented an unforeseen, life-limiting form of degradation for core internals that has significantly set back the design of fast spectrum reactors [1]. Over the next several decades, a campaign of experimentation on pure metals and austenitic steels has provided material to significantly advance theoretical understanding of swelling – how vacancy-type irradiation damage accumulates into uniform volumetric increases which cause component failure. Swelling is an exceptionally difficult phenomenon to study in ferritic-martensitic (FM) alloys. Unlike the pure metals and austenitic alloys, which constitute most of the community's experience, FM steels exhibit long incubation periods and low linear swelling rates that require hundreds of displacements per atom (dpa), a standard unit of radiation damage, to study properly (See Figure 2.1). This makes standard experiments costly in both time and money, resulting in a scarcity of experimental data, particularly under neutron irradiation conditions.

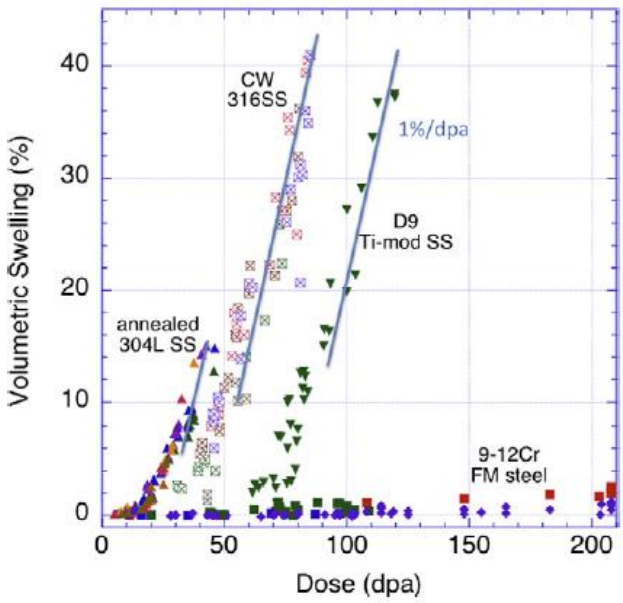

Figure 2.1: Comparison of Swelling Behavior in austenitic stainless steels (SS) and ferritic-martensitic steel (FM) [2]

2.1 Summary of Key Experimental Data in Ferritic-Martensitic Steels

2.1.1 Peak Swelling Temperature, Incubation, and Linear Swelling Rate

Test reactor facilities such as the Fast Flux Test Facility (FFTF), have been used to study radiation effects on ferritic alloys [3–12] at temperatures varying from around 380 to 440˚C. Figure 2.2 provides a summary of the swelling behavior observed in commercial ferritic martensitic steels containing chromium in the range from 8 to 12% from the FFTF database. It is important to note that this data set, in addition to covering a limited temperature range, also is subject to some degree of data uncertainty related to power changes during reactor operation that impact the accumulated dose, in addition to temperature and dose gradients. Thus, it is difficult to make systematic evaluation in the swelling behavior, and more importantly there is a lack of data at different doses in systematic increments can challenge the interpretation of incubation and growth regimes. Even when growth behavior is distinctly observed, the accuracy of determining the rate of swelling is limited by how large dose increments are, and the existence of a transient period of non-linear growth makes it difficult to assert whether a true 'steady-state' has been reached. A temperature at which swelling is maximized may slowly change as dose accumulates and therefore be fundamentally indeterminable. Finally, high sensitivity to initial heat treatment and microstructure also influence the uncertainty [6]. Indeed, Garner and co-workers [3] call into question the way dose was calculated in key results, which would significantly shorten incubation periods and increase linear swelling rates.

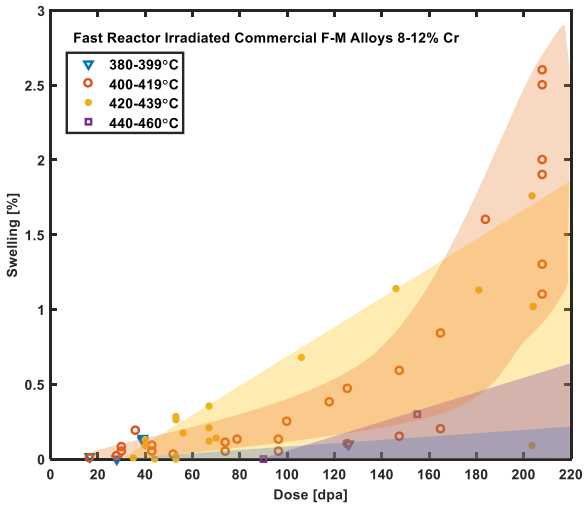

Figure 2.2: Collection of Relevant Test Reactor Data for Commercial Ferritic-Martensitic Steels from [3–12]

Since there is so little data to work with, incomplete data like those which contain no cavity size information because swelling was measured via density column was also considered. However, there are some experiments that contained a full TEM characterization of cavity size distributions [11–13]. Table 2.1 contains the range of observed or inferred outcomes for key swelling metrics in ferritic alloys under test reactor irradiation conditions.

Table 2.1: Range of Experimental Outcomes for Swelling under Neutron Irradiation (5×10^{-7}-5×10^{-6} dpa/s) [3–12]

Metric	Min	Max
Linear Swelling Rate [%/dpa]	0.01	0.5
Incubation Period [dpa]	0	100
Peak Swelling Temp. [C]	400	440

The high dose nature of the phenomenon makes ion irradiation particularly appealing, as it allows irradiations to be reach doses of interest in much shorter time frames, and with minimal sample activation. This has allowed significantly more systematic work to be done in pursuit of determining temperature peaks and separating incubation and growth periods [14–21]. While it can provide a more complete data set, ion beam techniques introduce several complications. The damage created is concentrated within a few microns of the sample surface, and an analysis region must be carefully selected to avoid surface sink effects and the ion implantation region [22]. Some ion beam vacuum systems have also been observed to permit some influx of carbon [23], which is expected to interact with and trap vacancies and create carbides which compete with other sinks for available defect flux. Finally, much of the legacy ion beam data uses beam rastering to ensure even profiles, which may alter cavity evolution compared to a steady state beam, which is more representative of the reactor case [24]. Recent work has done much to address this by simultaneously co-injecting helium with a secondary beam with its energy systematically degraded to match the spatial damage profile [25]. The degree to which experimenters account for these effects, and the high variation in dose rates available to ion beam facilities creates significant spread in this data as well.

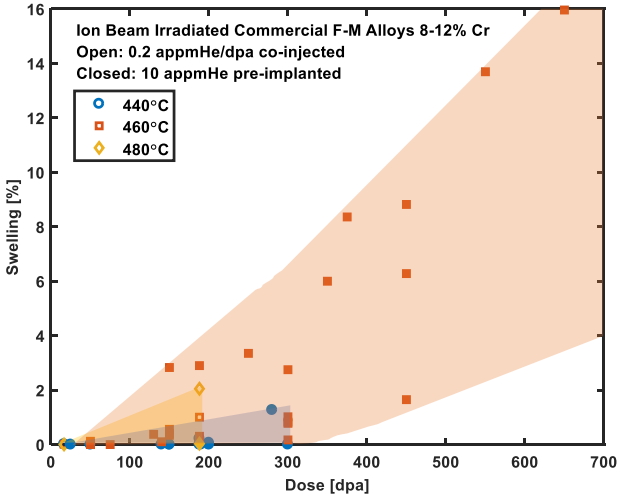

Figure 2.3: Collection of Relevant Ion Beam Data for Commercial Ferritic-Martensitic Steels from MIBL [19,26]

The introduction of inert helium simultaneously with damage, and at a ratio similar to transmutation rates in reactor is very difficult. Much ion irradiation data pre-implants a fixed amount of helium before the irradiation, or co-injects it with a second beam, but in ratios much larger than experienced in test reactor conditions. Focus will be placed on conditions with the most comparable helium levels and introduction methods. Table 2.2 contains the range of observed or inferred outcomes for key swelling metrics in ferritic alloys under ion irradiation.

Table 2.2: Range of Experimental Outcomes for Swelling under Ion Irradiation (5×10^{-4}-5×10^{-3} dpa/s)

Metric	Min	Max
Linear Swelling Rate [%/dpa]	0.02	0.2
Incubation Period [dpa]	100	400
Peak Swelling Temp. [C]	460	500

2.1.2 The Effect of Chemical Impurities on Cavity Microstructure

Helium generated by (n,α) transmutation reactions is an unavoidable bi-product of neutron irradiation, and its effect on cavity nucleation has been thoroughly studied. The extremely high mobility of this inert gas in the lattice (~0.07 eV) [27] allows the helium to rapidly accumulate on grain boundaries [28], dislocations [29,30], and cavities. The pressure He exerts on the surface of a cavity significantly suppresses thermal evaporation, particularly at small sizes where vacancies are most weakly bound [31]. One of the most prevalent models for void nucleation, the critical

bubble model [32], is predicated on the necessity of helium to overcome the barrier of thermal emission and nucleate stable, growing voids.

In copper (Figure 2.4), swelling has been observed to peak as a function of helium generation rate. At low He/dpa ratios, helium serves to increase stability of cavities which eventually become growing voids which contribute substantially to swelling. Eventually, around 10 appm/dpa in copper, helium enhancing nucleation begins to suppress swelling as more available vacancies partition to the large density of small bubbles that serve as recombination sites. In a study performed on FM T91 with various levels of pre-implanted helium (Figure 2.5), the latter effect was observed, but the peak in helium-to-damage ratio was at 0 appm implanted. Since swelling not only occurs, but is sometimes maximized, when no helium is present, mechanisms which require helium to overcome the thermal barrier to nucleation need to be reconsidered.

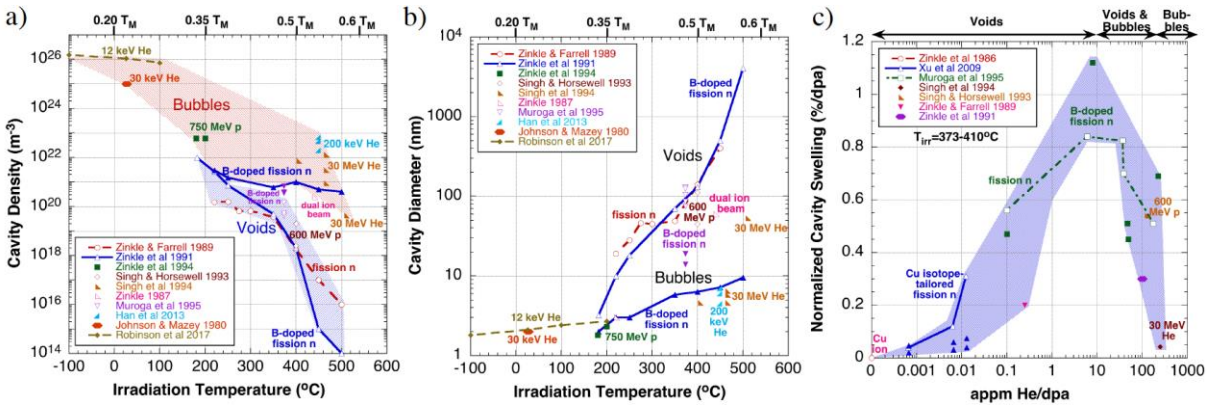

Figure 2.4: a) Cavity density as a function of irradiation temperature for Cu irradiated to ~1 dpa at He/dpa ratios from 0.1 to 300 appm He/dpa, b) cavity diameter as a function of irradiation temperature. The filled symbols denote bubbles, while the open symbols denote voids. c) Normalized cavity swelling as a function of He/dpa ratio for copper irradiated near 400°C at ~10^{-7} dpa/s. As reproduced from Ref.[33] compiling [34–39]

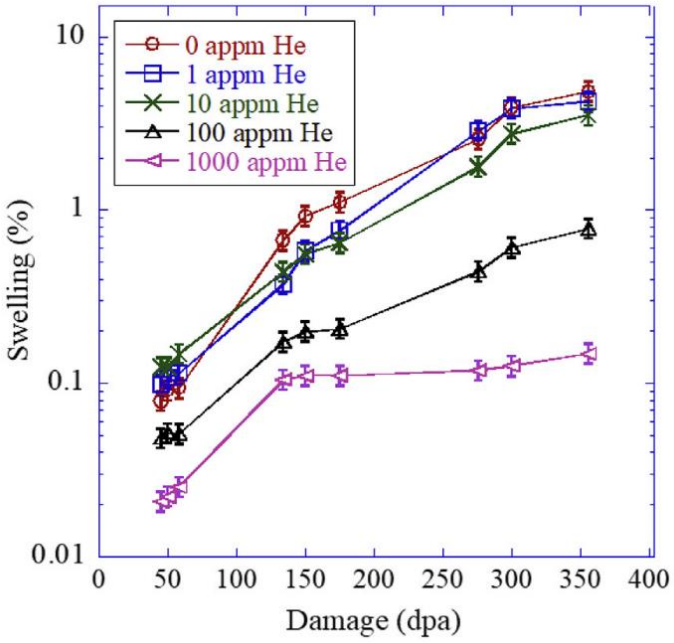

Figure 2.5: Swelling in ion-irradiated T91 with several levels of He cold pre-implantation [40]

In addition to helium, other trace impurities already present in commercial alloys have demonstrated an ability to influence swelling. Several DFT models have demonstrated significant trapping interactions between vacancies and impurities, particularly when the impurities are organic elements [41,42]. The effective reduction in vacancy motion retards nucleation, promoting recombination. Conversely, other impurities, such as phosphorus and silicon have (Figure 2.6) been shown to enhance swelling in trace amounts before reverting to impeding swelling as concentration increases. The reason for the initial increase is unclear and suggests a competition between a swelling-enhancing mechanism dominant at low densities and defect trapping dominant at high densities.

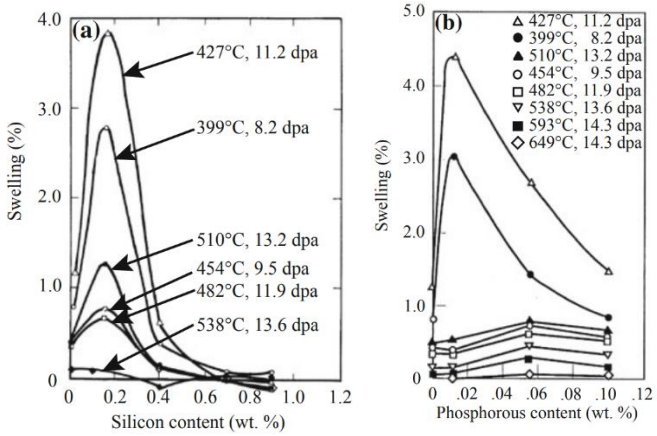

Figure 2.6: The effect of small substitutional impurities on swelling rates in austenitic steel [43]

2.2 Mean Field Rate Theory

The fundamental theoretical framework for understanding the evolution of cavity microstructure at experimental timescales is mean field rate theory (MFRT) and the void growth equation. In its simplest form, MFRT predicts the fate of non-equilibrium defects with only two differential equations that describe how vacancies and interstitials partition between recombination and competing defect sinks. These rates are described by solving the diffusion equation in an effective, lossy medium to generate concentration-independent rate constants [44]. These point defect balance equations are governed by

$$\frac{dC_{\mathrm{v}}}{dt} = G - R_{\mathrm{iv}} C_{\mathrm{v}} C_{\mathrm{i}} - K_{\mathrm{v}} C_{\mathrm{v}} , \tag{2.1}$$

and

$$\frac{dC_{\mathrm{i}}}{dt} = G - R_{\mathrm{iv}} C_{\mathrm{v}} C_{\mathrm{i}} - K_{\mathrm{i}} C_{\mathrm{i}} , \tag{2.2}$$

where v and i subscripts denote vacancies and self-interstitials (which will persist throughout this work), G denotes Frenkel pair generation rate (dose rate), and C's denote concentration of defects per unit volume. $R_{\mathrm{iv}} = 4\pi r_{\mathrm{iv}}(D_{\mathrm{i}} + D_{\mathrm{v}})$ is the recombination constant with D's denoting defect diffusion coefficients and r_{iv} is the recombination radius. K's are the sink strength for each defect type. Considering the principle sinks required for swelling, the total sink strength is the sum of dislocation and cavity absorption according to $K_{\mathrm{i}} = \rho_{\mathrm{d}} Z_{\mathrm{i}} D_{\mathrm{i}} + 4\pi \overline{r_{\mathrm{V}}} N_{\mathrm{V}} D_{\mathrm{i}}$ and its vacancy analogue $K_{\mathrm{v}} = \rho_{\mathrm{d}} Z_{\mathrm{v}} D_{\mathrm{v}} + 4\pi \overline{r_{\mathrm{V}}} N_{\mathrm{V}} D_{\mathrm{v}}$ where ρ_{d} is the dislocation density, N_{v} is the void (cavity) density, Z is dislocation bias for each defect type, and $\overline{r_{\mathrm{V}}}$ is the mean void radius. For a more robust picture, additional unbiased sinks such as grain boundaries and precipitates can be included as well.

Owing to the preference of dislocations for interstitials, the defect fluxes become unbalanced and result in a net flux of vacancies to cavity nuclei. The void growth equation translates this driving force for void growth into a radial growth rate, given by

$$\frac{dr_{\mathrm{V}}}{dt} = \frac{\Omega}{r_{\mathrm{V}}} \left[D_{\mathrm{v}} C_{\mathrm{v}} - D_{\mathrm{i}} C_{\mathrm{i}} - D_{\mathrm{v}} C_{\mathrm{v}}^{\mathrm{e}}(r_{\mathrm{V}}) \right] , \tag{2.3}$$

where Ω is the atomic volume and $C_v^e(r_V)$ is given by

$$C_{\mathrm{v}}^{\mathrm{e}}(r_{\mathrm{V}}) = C_{\mathrm{v}}^{\mathrm{e}} \exp\left[-(p - 2\gamma / r_{\mathrm{V}})\Omega / kT\right], \tag{2.4}$$

where $C_{\mathrm{v}}^{\mathrm{e}} = \exp(S_{\mathrm{f}}^{\mathrm{v}}) \exp(-E_{\mathrm{f}}^{\mathrm{v}} / kT) / \Omega$ is the bulk equilibrium vacancy concentration, γ is the cavity surface energy, p is the gas pressure on the cavity, and $S_{\mathrm{f}}^{\mathrm{v}}$ and $E_{\mathrm{f}}^{\mathrm{v}}$ are the entropy and energy of vacancy formation.

The temperature that maximizes swelling rate will be a focus of this effort, because it synthesizes a large amount of information on cavity nucleation into a single useful measure. Temperature exerts strong control over the density of stable cavities by effectively modifying the ratio of the sink strength of dislocations to that of cavities given by

$$Q_{\mathrm{i,v}} = \frac{\rho_{\mathrm{d}} Z_{\mathrm{i,v}} D_{\mathrm{i,v}}}{4\pi \overline{r}_{\mathrm{V}} N_{\mathrm{V}} D_{\mathrm{i,v}}}, \tag{2.5}$$

where D's denote defect diffusion coefficients, ρ_d is dislocation density per unit volume, Z's are capture efficiencies for each defect type, $\overline{r_V}$ is the mean cavity radius, and N_V is the number density of cavities [45]. At high temperatures, cavities are thermally unstable, and dissociation impedes the formation of a significant density, increasing Q. At low temperatures, cavities are so stable that they almost never dissociate, resulting in large densities of small clusters that act as powerful recombination sites, decreasing Q. Between these two extremes lies a window of stability where a density of cavities is sustained that can compete with biased sinks to drive a net flux of vacancies to the void. Swelling is generally maximized when $Q \simeq 1$, creating a parity which optimizes conditions for swelling and gives rise to a peak in the temperature distribution.

As suggested by the prominence of its term in the rate equations, dose rate has a significant impact on swelling rate. Faster production means less time between damage events for absorption at sinks, which results in increased recombination and decreased swelling. Mansur [46] used MFRT in two limiting regimes, sink-dominant and recombination-dominant, to derive expressions which predict an increase in temperature to compensate for high dose rates. In the sink dominant regime relevant to these complex FM alloys, the shift between two cases denoted by subscripts 1 and 2 is given by

$$T_2 - T_1 = \frac{(kT_1^2 / (E_m^v + E_f^v)) \ln(G_2 / G_1)}{1 - (kT_1 / (E_m^v + E_f^v)) \ln(G_2 / G_1)} \quad (2.6)$$

,

where E_m^v is the migration energy of a vacancy. The most noteworthy feature of these equations, is that they can generate predictions that can provide guidelines for the selection of a representative ion irradiation temperature using only the neutron irradiation temperature to be emulated, the ratio of damage rates, and the fundamental energies of vacancy formation and migration in the alloy.

While these equations have proved extremely useful, particularly in the thoroughly studied austenitic and pure metal cases, the very simplicity that makes them convenient is their main shortcoming. Temperature shifts for ferritic systems are significantly overpredicted by this simple approach. In the many years of experimental and theoretical advancement since Mansur's work have revealed several new physical phenomena that complicate this simple description. While theoretical constructs describing these features have been created in many cases, they complicate the math to a degree which precludes their incorporation into convenient analytical expressions that can generate predictions of temperature shifts.

Foremost among these deficiencies is the mean size approximation, which assumes a single static density of cavities exists at a single mean size and grow uniformly in tandem with one another. This completely removes any predictive power over nucleation phenomena. If differences in nucleation behavior between peaks are to be considered, their ratio must be used as an input to the more complete form of the temperature shift equation [46]. This means experimenters using these equations are implicitly assuming that cavity densities at the peak swelling temperature for different dose rates are identical. In reality, cavity evolution is a chaotic process with a complex bimodal size distribution in which some sizes are likely to grow rapidly while others shrink, dissociate, and recombine.

Treating time evolution of a size distribution with dose is well beyond the scope of any single expression. The large cavities at doses of experimental relevance are themselves as large as a typical simulation cell for atomistic techniques, and the long length and time scales involved in defect diffusion processes proscribe them techniques of this complexity. Object kinetic Monte Carlo (OKMC) techniques are well suited to treatment of the stochastic nature of defect accumulation, and are capable of accounting for spatial correlations between reacting species [47].

They are well suited to studying the early stages of nucleation, but high computational demands typically limit maximum dose to less than 1 dpa, which prohibits study of steady state void growth, particularly in systems where that is unlikely to occur before 100 dpa. They are also limited to species whose density is sufficiently high to be likely to appear in a simulation cell, which could be an issue if the size distribution is dominated by small densities of large voids. The only technique capable of providing redress to key assumptions on experimental length and timescales is cluster dynamics.

Although cluster dynamics (CD) models, a generalized implementation of MFRT which will be discussed in more detail in Chapter 3:, are the main tool available for simulating radiation damage at the length and timescales relevant to swelling, their application to date has been somewhat limited. The way the models are constructed causes the number of partial differential equations to increase dramatically with the size of cluster considered. As a result, many of the innovations in the last several decades have been focused on compressing the size of the problem to a computationally feasible size while preserving the integrity of the underlying physics [48,49]. The resultant complexity of these models, and the potential for artefacts they introduce has largely constrained swelling models to qualitative observations of the pathways to nucleation [50] and the behavior of helium vacancy clusters in tandem with desorption experiments [51,52].

Models of the interstitial components of the defect microstructure are more common. Relative to swelling, nucleation and growth of dislocation loops tends to saturate much faster, occurs over a much wider range of temperatures, and does not require the inclusion of a second defect species (helium). While the complex mobility of loops makes this problem far from trivial, these comparatively modest computation requirements have contributed to significant developments in this field. The very first loop models predicted a state where immobile clusters created directly in cascades were introduced at extremely high densities which were far too temperature dependent [53]. This ‘nucleation catastrophe’ was rectified with the implementation of a production bias model, which allowed these clusters to migrate rapidly through one-dimensional glide, annihilating at sinks such as grain boundaries. This rapid mobility itself conflicts with observations of low thermal mobility observed in in-situ electron microscopy studies. Further work directly comparing to in-situ TEM experiments and accounting for the limitations therein have demonstrated the

efficacy of trapping mechanisms which mediate this rapid diffusion in commercial F-M alloys [54–56].

Modeling with atomistic and stochastic methods have revealed several key deviations from MFRT with respect to fundamental generation and subsequent behavior of point defect. First, the only form of irradiation damage considered by Eqs. (2.1) and (2.2) are point defects, but a large body of molecular dynamics (MD) data has shown that high PKA energy damage events create damage in a state that already has some clustering of point defects [57]. In the case of interstitials this may affect their mobility, as larger loops can migrate one-dimensionally through the lattice [58]. Cascade-generated vacancy clusters can serve as preferred nucleation sites circumventing the smallest cluster sizes that are most prone to thermal emission. Second, the sink strengths, which govern the defect partitioning that allows bias-driven growth, are calculated using several underlying assumptions. Some, like the anisotropy of their absorptive power and the implication of even sink spacing, are unable to be relaxed without a more computationally costly, spatially resolved approach. It has also been widely assumed that cavities act as unbiased absorbers of point defects, but there has been work using both elasticity theory [59] and molecular statics [60] which suggests that, like dislocations, small cavities are biased toward interstitial absorption. This could constitute a considerable, and hitherto untreated, impediment to nucleation.

Chapter 3: Methods

3.1 Cluster Dynamics Model Structure

Cluster dynamics models are the intuitive logical extension of the classical MFRT. Where Mansur's analysis for instance, uses only two rate equations (one for each defect type), a cluster dynamics approach employs an intricate network of interconnected differential equations for each allowed cluster size. For a given species *j*, the concentration is governed by

$$\frac{dC_j}{dt} = G_j + R_j(C) - \sum_{m=1,n_{\text{sinks}}} k_{\text{sink}}^m D_j C_j \quad , \tag{3.1}$$

where $R_j(C_j)$ is a sum of all possible aggregation events

$$R_j^+(C) = \sum_{n+m\to j} k_{n,m}^+ C_n C_m - \sum_k k_{j,k}^+ C_j C_k \tag{3.2}$$

and dissociation events

$$R_j^-(C) = \sum_{n\to n-j,j} k_{n-j,j}^- C_n - \sum_k k_{j,k}^- C_j \tag{3.3}$$

which create or consume a species *j*, where *k*'s are unique rate constants encapsulating the kinetics of each reaction. This approach, while significantly more complex, allows the simplifying assumptions inherent to Mansur's correlation, most notably a static number density and the mean size approximation, to be relaxed.

The multiscale nature of the resultant model, with characteristic times ranging from the nanosecond life of a single defect to the years-long formation of large cavities in reactor, creates a particularly stiff system of ordinary differential equations ODEs (or PDEs in cases where spatial discretization is considered). This essentially mandates implicit solution strategies which generally require the generation of a Jacobian matrix for the system. For extremely large systems like that of cluster dynamics, a full Jacobian is prohibitively memory intensive, and its interconnected nature neuters the efficacy of sparse storage strategies. To combat this, a Jacobian-

free Newton-Krylov (JFNK) method is employed to solve for the unknown right endpoint of the backward Euler method. When well converged this prohibits negative concentrations, conserves defect mass, and keeps the problem computationally manageable.

Using cluster dynamics to draw the picture of defect kinetics into sharper focus will require many more rate constants than a simple MFRT approach. This section aims to describe and justify how each of these interactions is calculated.

3.1.1 Primary Damage Formation

Decades of molecular dynamics simulations of high PKA energy damage events have demonstrated their propensity to form defect clusters directly in cascades. This can be folded into a PKA distribution from SRIM [61,62] or SPECTER [63] calculations to generate a vector of defect cluster production terms following the methodology of Xu and co-workers [64]. Two such vectors were created (Figure 3.1), a 9-vacancy and 20-vacancy clusters as the largest size generated. The 9-vacancy maximum is representative of the direct result of an MD cascade in iron. The 20-vacancy maximum is based on a similar cascade in which the spatial correlation of the damage event promotes coalescence beyond that which would be adequately described by defects evenly distributed throughout a mean field.

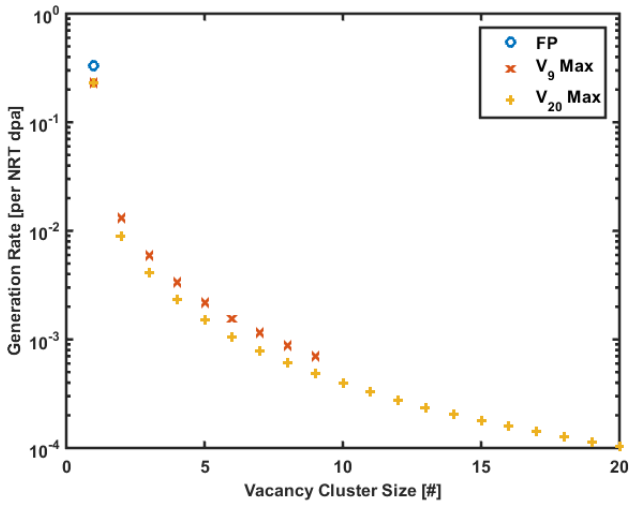

Figure 3.1: Vacancy cluster distributions used in cascade primary damage formulation as compared to Frenkel pair treatment

3.1.2 Sinks

Sinks, defined here as any feature distributed throughout the microstructure which absorb or promote recombination of defects, are perhaps the most crucial element of any MFRT or CD

model. If interstitial and vacancy-type defects were consumed equally by all sinks, there would be no driving force for any long-term evolution of extended defect clusters. The primary sinks considered in most prior mesoscale modeling studies are network dislocations, dislocation loops, grain boundaries, precipitates, and cavities. Dislocations have a known propensity to interact elastically with interstitials over long distances relative to vacancies resulting in a preferential absorption termed bias. The resulting supersaturation of vacancies in the matrix promotes void growth.

The conventional calculation of sink strength is performed by defining a cell of appropriate geometry around a sink (spherical for voids, cylindrical for dislocations, etc.) and assuming a steady state solution to the diffusion therein [44,65–68].

For the purposes of this work, cavity sink strengths are an intrinsic part of the system of cluster equations considered, and their absorption will be treated individually as part of the larger reaction network. The only external, static sink of major import is therefore the dislocation network which will serve as the consistent source of bias driving void growth. The analytical solution to the diffusion equation in a cylindrical cell and generalizing to all dislocations in the system yields the sink strength [69],

$$k_j^{\text{disl}} = \frac{2\pi\rho}{\log\left(\frac{1}{\sqrt{\pi\rho r_c^2}}\right)}, \tag{3.4}$$

where r_c is the dislocation's capture radius, and ρ is the dislocation density. A simplified form of Eq. (3.4) expresses the sink strength as a simple product of a bias factor and dislocation density. This significantly simplified the mathematics in pursuit of analytical solutions, but the approximation breaks down when large dislocation densities, like the ones in ferritic systems, are used.

Assuming that grain boundaries act as ideal sinks (that is, incident defects are eliminated rather than merely trapped), mean field formulations of sink strength are available [67]. Some initial results considered these treatments, but they were largely dispensed with, as the high sink density provided by dislocations within the grains dwarf the GB contribution in a bulk sense. The localized

effects of grain boundary sink strengths were considered in a more nuanced form via a spatially-dependent finite difference term to be described in section 5.1.2.

3.1.3 Aggregation Reactions

The nucleation and growth of visible cavity microstructure is the result of a series of individual defect capture events by smaller clusters, represented by the positive component of the interaction term R in Equation (3.2). The defining features of the kinetics in this system are the rate constants $k_{j,n}$, which define the rate at which clusters of size j and n coalesce to form a new cluster (j+n). Presuming spherical symmetry and isotropic diffusion, the classical sink strength formulation described in 3.1.2 can be generalized as

$$k_{j,n}^{+} = 4\pi r_{j,n}(D_j + D_n), \tag{3.5}$$

where $r_{j,n}$ is the reaction distance and D_j and D_n are the diffusivities of each species which are calculated as a function of migration energy E_m and an exponential prefactor D_0

$$D_j = D_{0,j} \exp\left(\frac{-E_{m,j}}{k_b T}\right). \tag{3.6}$$

3.1.4 Dissociation Reactions

The countervailing process to aggregation at sufficiently high temperatures is thermal emission of mobile monomers (primarily vacancies and interstitial helium atoms). This is especially critical to the understanding of cavity nucleation as it is responsible for the falling edge of peak swelling. For purposes of cluster accounting each event eliminates one parent cluster and produces one daughter and a mobile monomer according to Eq. (3.3).

The rate constant for these reactions is fundamentally a function of how strongly bound each monomer is. Due to the mean field nature of these models, determination is not quite as simple as a treatment mirroring simple diffusion, where an emitted monomer must overcome an activation barrier with some attempt frequency. However, the emitted defect is strongly spatially correlated to its parent cluster, and there is a high likelihood of rapid reabsorption. To account for this in a mean field sense requires a thermodynamic approach. Since each dissociation reaction has a reverse aggregation reaction, we can build on the assumption that these two are at an equilibrium

condition. If the total change in free energy between a bound and unbound state is presumed to be dominated by the binding energy, E_b, the resulting rate constant is

$$k^{-}_{(j+n)} = \frac{k^{+}_{j,n}}{\Omega} \exp\left(\frac{-E_{b,(j+n)}}{k_b T} \right) \tag{3.7}$$

These binding energies generally follow the capillary extrapolation law [70] modified by the partial pressure of helium according to a hard sphere equation of state [31]. In addition to the more common thermal emission of vacancies, interstitials can also be emitted via a loop punching mechanism, generating a new vacancy member within the cavity. The energetic barrier to interstitial formation is too great for this to occur thermally in the swelling temperature regime, but it is possible with sufficient helium overpressure.

3.1.5 Parameterization

While cluster dynamics promises more accurate quantitative knowledge of how defects (and impurities) diffuse and aggregate in a material, the rate constants governing the system behavior are still dependent on input parameter selection. In some cases, experimental techniques like resistivity recovery experiments can be used to provide a direct measurement [71]. For many others, atomistic modeling techniques like density functional theory (DFT) and molecular dynamics (MD) are the best options for providing accurate estimates. The parameters of most relevance to ferritic alloys have been curated and are given in Table 3.1. It is worth noting that CD cannot explicitly and distinctly model iron and chromium, and as such many of the fundamental material parameters (e.g. lattice parameter) are for bcc iron. Despite this, there are several reasons for the claim that this model simulates FM steel. First, the migration parameters for point defects, which are the most fundamental contributor to defect kinetic behavior, were taken from atomistic simulations in Fe-9Cr [70–75] when available. Second, the sink parameters (dislocations and grain boundaries) were gathered directly from experimental characterizations of FM alloys, providing description of features which would be far less prevalent in a pure iron sample. Experimentally-informed parameters considered here were from nominally 12Cr HT-9, but they are consistent with newer data for 9Cr T91. Finally, the development of this model at every stage is guided by comparison with FM experiments.

Table 3.1: Material Parameters used to describe defect evolution

Parameter	Value	Source	Alloy
Dislocation Density* [m^{-2}]	3×10^{14}	Characterization[5]	HT-9 (12Cr)
Grain Diameter* [m]	1×10^{-6}	Characterization[5]	HT-9 (12Cr)
Di-vacancy Binding Energy [eV]	0.3	DFT Modeling [70]	bcc Fe
Vacancy Formation Energy [eV]	1.79	DFT Modeling [75]	Fe-9Cr
Vacancy Migration Energy [eV]	0.63	DFT Modeling [72]	Fe-9Cr
Interstitial Formation Energy [eV]	4.76	MD Modeling [73]	bcc Fe
Interstitial Migration Energy [eV]	0.22	DFT Modeling [71]	Fe-9Cr
Helium Migration Energy [eV]	0.07	MD Modeling[74]	Fe-9Cr
Vac. Diffusion Pre-exponential [m^2/s]	8.2×10^{-7}	DFT Modeling [72]	Fe-9Cr
Int. Diffusion Pre-exponential [m^2/s]	1.1×10^{-7}	DFT Modeling [71]	Fe-9Cr
He Diffusion Pre-exponential [m^2/s]	4×10^{-7}	MD Modeling[74]	Fe-9Cr
Dislocation Bias for Interstitials	1%		

*Consistent with subsequent characterization of nominally 9Cr T91

The cluster dynamics framework accommodates the addition of a second dimension to this system of equations to describe each cavity's helium content. Each rate equation for a cluster is replete with capture and emission terms as function of size. Capture reactions are directly proportional to radius and diffusivity, but emission terms take on an additional Arrhenius dependence on some binding energy. Each binding energy is a sum of contributions to stability provided by vacancies and helium atoms. Vacancies enhance stability as size is increased, asymptotically approaching the formation energy according to capillarity. For very small cluster sizes where capillarity breaks down (entering a negative energy space), ab initio binding energies are available, starting with the least stable divacancy at 0.3 eV [70]. The stabilizing helium pressure inside the bubble is approximated using a hard sphere equation of state [31] and added to the vacancy contribution. Furthermore, when helium partial pressure is in excess of equilibrium by more than the formation of an interstitial, a 'loop-punching' event can occur described as an emission event where the monomer is an interstitial. These, along with other critical quantities were compiled from ab initio data for α-iron and experimental measurements into a reasonable parameter set (Table 3.1).

Stochastic and deterministic cluster dynamics methods have been used to analyze radiation phenomena such as bubble formation, loop nucleation and growth, and precipitation [54,55,76–78]. However, existing cluster dynamics work attempting to compensate for dose rate has been confined to very low temperatures and doses due to memory limitations imposed by the large number of equations necessitated by the treatment of three clustering species (interstitials, vacancies, and helium gas atoms). This work reduces these to a computationally manageable level using a moment-based scheme that groups hundreds or thousands of equations which would be associated with larger clusters into just three representative quantities such that the concentration of defects is conserved [48,79]. This permits assessment of high dose swelling behavior which would otherwise require the solution of billions of differential equations.

The grouping method functions by taking the first two moments of a power series approximation of the size distribution function:

$$f(n) = A_j + B_j(n - \bar{n}_j) \tag{3.8}$$

where n is the number of vacancies in a cluster, and a given group j is represented by the 0th moment A and 1st moment B of that group. Provided these two degrees of freedom, one can express both the number of clusters

$$C_j = A_j \Delta n_j \tag{3.9}$$

and the number of individual members residing in those clusters,

$$N_j = \left(A_j \bar{n}_j + B_j \sigma_j^2\right) \Delta n_j \tag{3.10}$$

as a function of the group width Δn_j, in order to ensure that they are both conserved. Differentiating these with respect to time and performing some substitutions yields the ODEs

$$\frac{dA_j}{dt} = \frac{1}{\Delta n_j}\left(R_j(\mathrm{n}_{j-1}) - R_j(n_j)\right) \tag{3.11}$$

$$\frac{dB_j}{dt} = \frac{\Delta n_j - 1}{2\sigma_j^2}\left(R_j(\mathrm{n}_{j-1}) + R_j(n_j) - 2\bar{R}_j\right) \tag{3.12}$$

where $\bar{R}_j$ is the reaction rates internal to the group. An additional 1st moment can then be applied to generalize this method to multidimensional phase spaces (in this case, the incorporation of helium).

Despite being the most promising option to conserve both total defect mass and the total number of clusters in which they reside, moment-based grouping schemes are not without artefacts. It has been demonstrated that these schemes are prone to numerical oscillations, particularly at the largest sizes, as rapid changes in density as a function of size cause the moment to under-sample the distribution [48,49]. Without great caution, these oscillations can themselves propagate, allowing this noise to dominate the underlying physical signal. This is clearly a problem, particularly as one attempts to model bi-modal size distributions which experience a rapid density change between the bubble and void peaks. For this reason, the model has been reconstructed to encompass a full, two-dimensional cluster dynamics phase space at small sizes where nucleation occurs which is interfaced to a coarse, mean size approximation when cluster grow beyond the upper bound. The density is controlled by the sum of all aggregation terms resulting in a cluster size larger than the upper bound and its growth is controlled by the void growth rate equation (3.13), presuming all clusters reside at the mean size. If the exchange between these regimes is sufficiently dynamic, and group sizes are kept small, these maintain defect conservation and produce results consistent with well-behaved cases employing a full phase space.

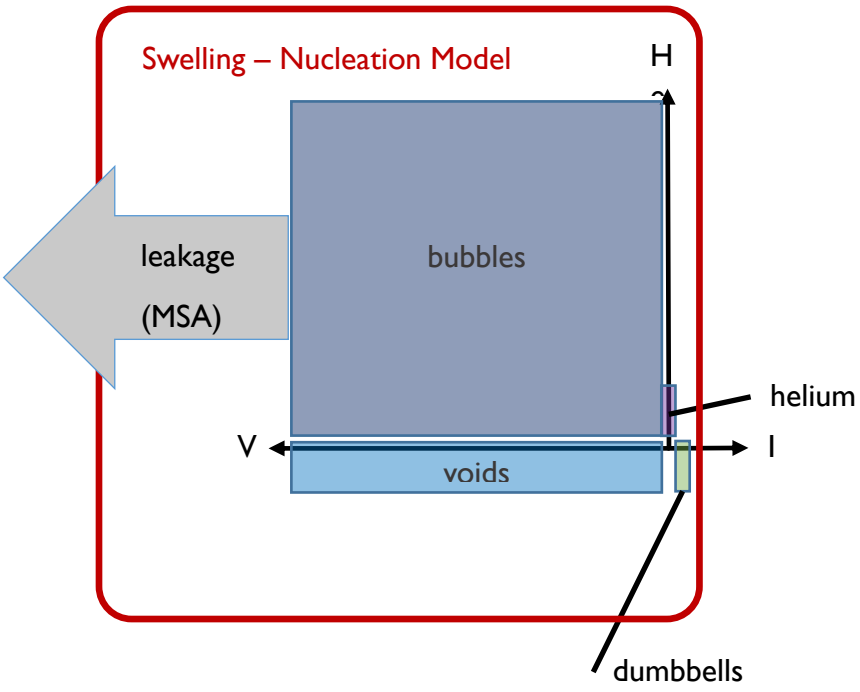

Figure 3.2: Visual representation of the phase space of cluster equations considered with respect to defect and helium component axes

At any stage of irradiation, the resultant swelling can be calculated by integrating over the phase space to determine the total number of vacancy members residing in cavities, using the grouped form Eq. (3.10) if applicable. Similar integrations can be performed to obtain other aggregate quantities such as total cavity density and mean cavity radius. Since experimental results are primarily obtained via TEM characterization, only clusters larger than 1.5nm are considered to account for the limitations of image resolution. Separate calculations were performed to ascertain the contribution of sub-resolution features to swelling. Since swelling scales with the cube of cavity radius, even at large densities, these small features amounted to a volumetric change a few orders of magnitude below their TEM-visible counterpart.

Even with the grouping scheme in place, the helium-vacancy space is very large such that a robust treatment of dislocation loops and their glide results in a model which is too memory intensive. Instead loop nucleation is neglected in favor a single three-dimensionally mobile interstitial instance of equation (3.1), and a fixed edge dislocation network with a reasonable density of 3×10^{14} m^{-2}, which falls within the range of experimental outcomes summarized in Section 2.1. Experimental experience suggests that even over the course of as much as 650 dpa, the total dislocation line length is unlikely to change by more than about a factor of two as a result of loop

nucleation [19], suggesting that this assumption is reasonable. Furthermore, preliminary results (to be presented in Figure 4.3) suggest that in the high density regime relevant to FM steels (10^{14}-10^{15} m^{-2}), swelling is significantly more sensitive to the input bias parameter than to dislocation density itself. With regard to the character of dislocations, screw dislocations, which are prevalent in FM steels, exhibit significantly lower bias than their edge counterpart [80]. This was initially considered as a justification for the very low 1% value in the initial parameterization. However, the loop microstructure, which is not explicitly treated here, is expected to introduce a comparable line length of edge dislocations within the first few dpa, which is too short to explain long incubation periods.

Ideally, the model would consider each of the many features of the complicated FM microstructure, such as lath/sub-grain boundaries, Ni/Si-rich clusters, and α' precipitates. Unfortunately, even this relatively coarse technique lacks the computational ability to earnestly model their time evolution, spatial dependencies, and the various temperature and dose rate dependencies thereof. Thus, two assumptions are made for homogeneously nucleating cavities:

- The only interaction between other extended microstructural features and the cavities is inter-sink competition for point defects.
- In swelling-relevant temperature and dose rate regimes, the sink strength of these features will be dominated by the dislocation network.

The former is a quite easily justified, since more explicit treatments of non-homogeneous interactions between helium or vacancies with these features will be explicitly discussed later. The latter corroborated by a body of characterization data for comparable ferritic alloys [26,81,82], which consistently show dislocations as the dominant sink.

3.2 Additional Mechanisms

3.2.1 Heterogeneous Nucleation

A heterogeneous nucleation mechanism was incorporated through the addition of a parallel phase space associated with a preferred nucleation site. A constant density of static sites was introduced to the system, which immobilize helium atoms according to a trapping/detrapping pair defined in Eqs. (3.2) and (3.3), providing a site for self-clustering. Eventually these helium clusters will achieve sufficient pressure to punch out an interstitial, or absorb a vacancy, transforming into a

bubble. The aggregation-emission behavior of larger helium clusters is presumed to have the same clustering energetics as free interstitial helium in the lattice according to Gao [83] with only mobility affected. In order to ensure the number of sites is conserved, there are effectively two interfaced copies of the phase space: one where none of the He-V clusters are associated with a site (homogeneous contribution), and one where all of them are (heterogeneous contribution).

The major unknowns of this model which define its effect are the density of static helium traps, and the binding energy defining the strength of this interaction. To gauge the strength of these effects on swelling, a suite of calculations of temperature shift was conducted at various values of these controlling parameters. There is a significant quantity of data from helium desorption experiments [84–87] and molecular dynamics [29,30,88] estimating the interaction energy of helium with assorted microstructural features. These, paired with microstructural characterization of irradiated F-Ms [13,19] can provide density estimates for potential sites (dislocations, small precipitates, loops). Comparison between this identifying information and parameters that provide reasonable temperature shifts will be used to posit the type of site most likely to be dominant.

3.2.2 Cavity Bias

Elastic interactions between point defects and nucleating voids were treated by applying a size-independent extension ($r_{\text{ext},j,n}$) to the geometric radii of each reacting species in the reaction radius used in Eq. (3.5),

$$r_{j,n} = r_{\text{geom},j} + r_{\text{geom},n} + r_{\text{ext},j,n} \tag{3.14}$$

Molecular statics show that these elastic extensions are larger for interstitials than for vacancies, as shown visually in Figure 3.3, creating a net bias on cavities. This creates an inhibition to nucleation which, unlike thermal emission, cannot be overcome by simply shifting to a lower temperature. The magnitude of these elastic extensions is approximately independent of cavity size, so as cavity size increases, the relative increase in capture volume decays as its contribution relative to the geometric volume is reduced according to

$$Z = \frac{k_i - k_v}{k_v} = \frac{r_{\text{ext,i}} - r_{\text{ext,v}}}{r_{\text{void}} + r_{\text{ext,v}} + r_{\text{defect}}} \tag{3.15}$$

Where Z is the net bias and k is the reaction rate constant given in Eq. (3.5). The result as a functional form, shown in Figure 3.4, creates a strong impediment to nucleation.

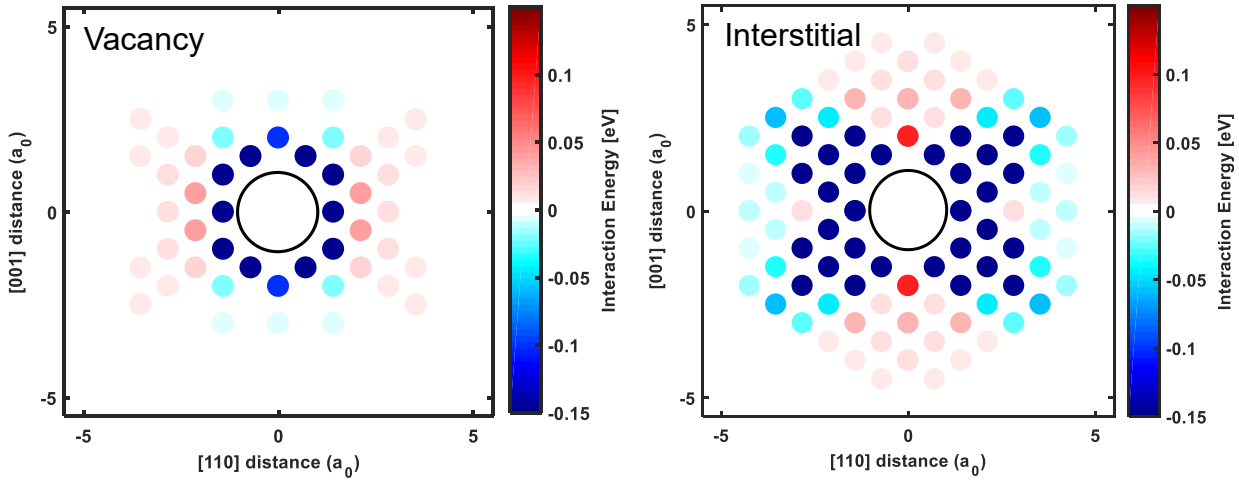

Figure 3.3: Maps of elastic interaction energy between a 59-vacancy cavity at the center of a molecular statics simulation box and point defects of vacancy and interstitial type placed at each nearby lattice position reproducing [60]

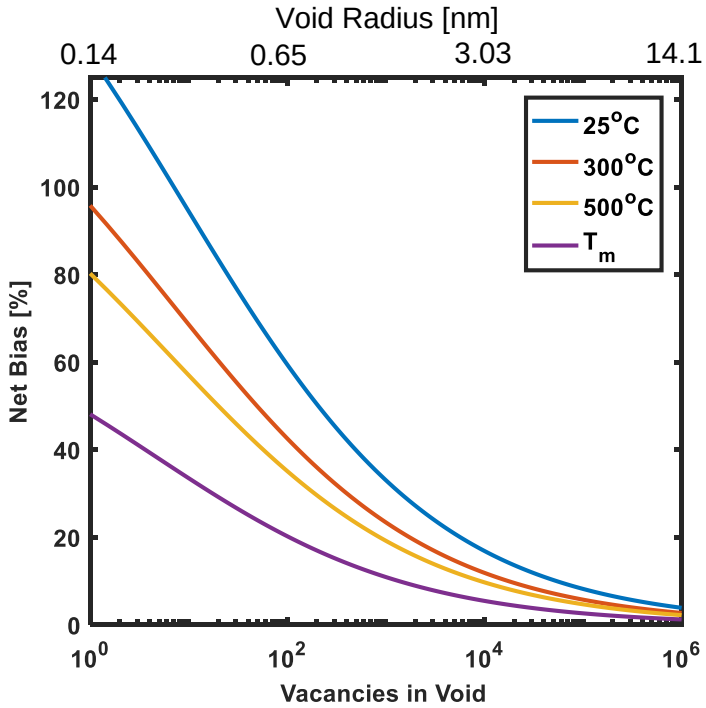

Figure 3.4: Net bias toward interstitial capture on a cavity as a function of its size at several temperatures reproduced using the methodology in [60]

Statics calculations in bcc iron reproducing the work of Kohnert [60] specify the size of these elastic reaction extensions as 0.32nm for interstitials and 0.08nm for vacancies. The resulting impact on the bias-balance of the system was mapped against the dislocation bias to determine

conditions where the most reasonable incubation dose and linear swelling rate occur at the temperature peak.

3.2.3 Impurity Trapping

Atomistic modeling efforts have recently computed the stabilities of these immobile complexes formed by vacancies and impurities (predominantly carbon) [41]. To remove the residual stability effects of the effective diffusion approach, these binding energies were utilized to determine the strength of this effect directly by including a new reaction pair for formation and dissociation of a vacancy-carbon complex. A static density of sessile carbon atoms was introduced with an input binding energy informing the stability of complexes.

For an even more rigorous treatment of these interactions, an equation was added to provide mobility to small vacancy clusters and carbon solutes [89], and by allowing accumulation of immobile V-C clusters up to size V_4C_2 using binding energies from ab initio calculations and simple extrapolation [41] (Table 3.2)

Table 3.2: Vacancy and Carbon Mobilities, as well as binding energies of carbon atoms to assorted species of V_xC_y complexes [41]

Species	E_m [eV]	D_0 [m^2/s]	$E_b^{V_xC}$ [eV]	$E_b^{V_xC_2}$ [eV]
V	0.63	8.2×10^{-7}	0.41	0.77
V_2	0.62	8.2×10^{-7} / 2	0.63	0.68
V_3	0.35	8.2×10^{-7} / 3	0.63	0.68
V_4	0.48	8.2×10^{-7} / 4	0.63	0.68
C	0.86	1.4×10^{-7}		

3.2.4 Vacancy-Driven Heterogeneous Nucleation

With helium-driven heterogeneous nucleation and cavity bias in place within the model, creating a vacancy driven form will be a matter of turning off or turning down vacancy bias when they are associated with a preferred nucleation site. The idea is that some chemical species is modifying the strain field about the vacancy such that it no longer prefers interstitials. Site densities, and extent of reaction extension radius reduction will be profiled to find a process which results in slower nucleation than cases with helium. Attempts at explicitly modeling the accumulation of impurities on the surface of the cavity through a trapping/de-trapping reaction pair caused the

implicit dilute limit assumptions used to compute reaction rate constants (described in 3.1) to break down due to the large densities relative to that typical of defect populations. Furthermore, the mechanism by which impurities would be expected to accumulate, radiation-induced segregation, is inadequately described by an Arrhenius binding energy.

An explicit model for segregation to a cavity surface would infeasibly expand the complexity of the model, so an effective treatment was pursued. To determine whether such a bias suppression could explain swelling behavior complete with incubation periods without co-injected helium, a smooth, linear reduction in bias was prescribed as a function of total accumulated defect flux. While this method will not capture all of the physics of the problem, parameters can be informed by both observations of segregation behavior at static sinks such as grain boundaries, and the results of coordinated atomistic modeling as described in the next section to ensure that results are justifiable, even if they cannot be considered predictive.

3.3 Molecular Statics

The origin of the renewed interest in cavity bias as a key impediment to swelling with the potential to explain incubation periods is primarily the result of molecular statics calculations [60]. To determine the effect of impurities on these bias factors, a comparable study was performed following the methodology therein.

Molecular statics simulations progress by simply finding the minimal potential energy state for an input configuration of atoms at 0K. Each atom is defined by its position in the simulation cell and an interatomic potential function which defines its energy state based on the other atoms in the system. In much the same way that a cluster dynamics model makes assumptions about the way atomic scale features that allow spatial fluctuations to be suppressed, interatomic potentials smooth fluctuations in electronic structure to simplify atoms to objects obeying Newtonian mechanics. These potentials are generated utilizing empirical fits to a variety of data measured experimentally or calculated using more rigorous atomistic modeling techniques. These fitting criteria can range wildly, from defect formation energies to elastic moduli. As a result, careful potential selection is key, as fits generated with respect to one specific subset of atomic behaviors may fail spectacularly when trying to reproduce another phenomenon.

Using the LAMMPS [90] package, a $(30\times30\times30)a_0$ simulation box was populated with a perfect bcc iron lattice. The total potential energy of the system is then minimized algorithmically. By

introducing various defects, their formation energy can be calculated by comparing the new energy state to the perfect lattice according to

$$E_f = U_{\text{tot}}^{\text{w/defect}} - \frac{N_{\text{atoms}}^{\text{w/defect}}}{N_{\text{atoms}}^{\text{perfect}}} U_{\text{tot}}^{\text{perfect}} \tag{3.16}$$

In the same way, placing two defects with known formation energies into the lattice allows the interaction energy to be computed according to

$$E_{\text{int}} = U_{\text{tot}}^{\text{w/defect+void}} - U_{\text{tot}}^{\text{perfect}} - E_{\text{f}}^{\text{defect}} - E_{\text{f}}^{\text{void}} \tag{3.17}$$

If these interaction energies are negative, there is an attraction between these features. This energy is presumed to be sufficient to capture a mobile defect if the interaction meets the criterion [91]

$$E_{\text{int}} \leq -k_b T \tag{3.18}$$

Repeating this method with a void in the cell center, and defects at all surrounding lattice positions creates a map of effective capture sites. Such maps were obtained for each defect type, averaging all six configurations for a (110) dumbbell in the interstitial case.

These anisotropic interaction maps must then be synthesized into a quantity which can be utilized as a CD sink strength. Each capture site satisfying (3.18) was treated as though it were a member vacancy in the void and its "capture radius" was calculated assuming a spherical geometry. The increase to capture efficiency resulting from elastic interactions is deemed the radius extension. Any mismatch between interstitial and vacancy can be converted directly into a net bias using Eq. (3.15). These bias calculations were repeated with several impurities, represented by appropriate embedded atom method pair potentials for iron with phosphorus [92], carbon [93], nickel [94], and vanadium [95].

Chapter 4: Initial Results & Objective

4.1 Simple Frenkel Pair Model

Beginning with the simplest form of such a CD model, and most comparable to the invariance relations, temperature distributions were generated for both fast reactor and heavy ion beam relevant dose rates (Figure 4.1). The resulting distributions exhibit peak swelling temperatures at 140˚C and 240˚C for neutron and ion dose rates respectively. Even though these temperature peaks manifest well below all experimental data, they coincide with Mansur's invariance equation (2.6) given the appropriate input energies (see Table 3.1), which results in an estimate of approximately 106˚C. Initially, upward shifts in these peaks toward the values from experiment were accomplished by increasing the vacancy migration energy. This strategy could produce the desired result with increases of only a few tenths of an eV. Such an increase requires a departure from ab initio and desorption data, which must be physically justified using an argument of impurity interactions. Since the migration energy also effects the rate of emission however, this would imply that every single nucleating cluster is associated with an impurity. Modifying this reduced diffusivity to only apply to capture reactions has a qualitatively dissimilar effect, reducing peak magnitude with minor change in peak location. This discrepancy indicates that this preliminary model severely underestimates thermal stability of nucleating voids. Size and density information (Figure 4.2) complicates this picture, as even though the low peak temperatures suggest low stability, void densities overpredict experiment by about an order of magnitude.

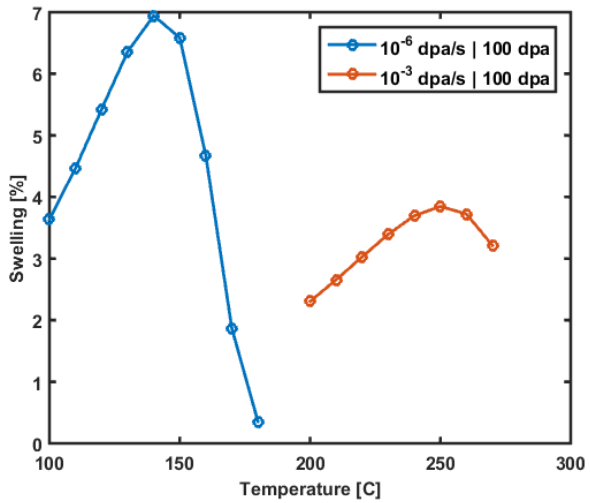

Figure 4.1: Initial temperature distributions of swelling at 100 dpa under reactor and ion beam relevant dose rates using Frenkel pair damage

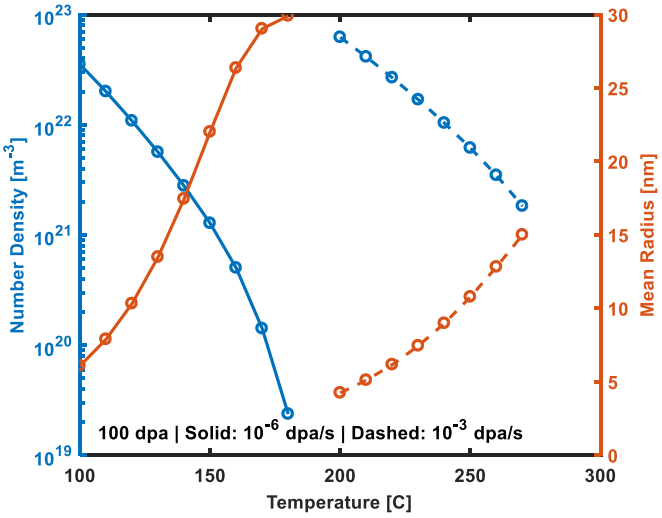

Figure 4.2: Cavity density and mean radius at reactor and ion beam relevant dose rates using Frenkel pair damage

In addition to the peaks presented in Figure 4.1, a suite of calculations varying the parameters most intuitively connected to swelling behavior were performed. Establishing bounds of reasonability and making justifiable adjustments to these parameters is the simplest way to improve agreement. The most notable parameters in consideration for changes are cavity binding energy, vacancy mobility, dislocation bias, and unbiased sink strengths. Temperature peaks generated from these slight changes are provided in Figure 4.3.

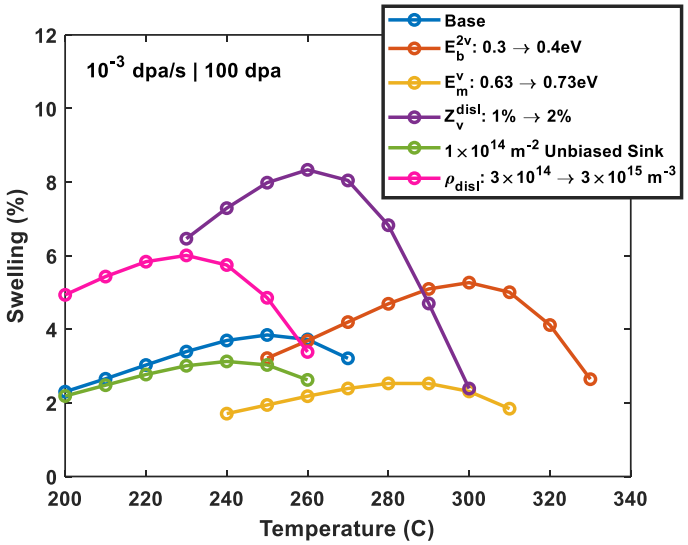

Figure 4.3: Ion-irradiated temperature distributions of swelling at 100 dpa with small variations in key material parameters

Both vacancy migration and surface energy modifications show the ability to shift the location of temperature peaks. Some may argue that a sufficient increase in vacancy migration could be justified by interactions with solute atoms, but it will be described in a later section that this treatment is flawed. Also, peaks tend to shift in tandem, providing no redress to the overestimate of temperature shift. Dislocation bias and unbiased sink additions exhibit a strong effect on swelling rates by controlling the number of incoming defects and what proportion are vacancies which contribute to cavity growth. While this creates an extremely strong control over the magnitude of swelling, it has little control over the temperature regime in which swelling occurs. An order of magnitude increase in dislocation density exhibits a significant effect on both swelling rate and optimal swelling temperature. This is attributed to an increase in total competing sink strength, which reduces peak temperature and swelling rate, offset by an increase to the net bias of the system which increases swelling rate. Unfortunately, none of these parametric changes are capable of addressing the large discrepancy with experiment without assuming values unreasonably out of line with measurement.

To summarize, there is outstanding disagreement between these initial results and experiment. Model results show that:

1. Dose rate effects are much stronger than observed, with test-reactor-relevant damage resulting in a peak swelling temperature around 300˚C, well below any experimental measurements.
2. Linear swelling rates are far higher than observed unless a very small dislocation bias on the order of 1% for interstitials is used.
3. Long incubations periods are effectively absent, with linear swelling being reached in only 1-2 dpa.

4.2 Objective

The objective of this work is to identify and include the physical processes describing cavity nucleation in a cluster dynamics model and compare model predictions against experimental results for ferritic-martensitic swelling behavior by iteratively adding self-consistent mechanisms to the model and mapping the temperature, helium, and dose rate dependencies thereof.

Chapter 5: Results and Discussion

This section aims to systematically account for the physical processes missing in existing descriptions of cavity nucleation in irradiated F-M steels: heterogeneous nucleation, cavity bias, and impurity-based bias suppression.

To address overestimation of temperature shift, a heterogeneous nucleation mechanism was implemented. The low temperature at which swelling is maximized in reactor-relevant dose rates implies an underestimation of cavity stability. Increases in the parameters that control stability have an identical effect on the ion irradiation peak, leaving the temperature shift unchanged. Since there is no reason to believe cavities are more thermally stable when neutrons are creating them, an athermal process must exert control over nucleate density. A heterogeneous nucleation mechanism provides such a device by introducing a density of preferential sites wholly independent of irradiation parameters. With less difference in the number of nucleation sites between dose rates, the compensatory temperature offset required can be expected to be much smaller. Justification for using these improvements will be predicated on the ability to identify their correlation to another feature of the microstructure, either dislocations or precipitates. The anticipated densities and interaction energies of these sites will be key to predicting the strength of their effect.

A bias of small cavities for interstitials will be studied to see if providing such an impediment to cavity growth can reduce swelling rates and create incubation periods. The overestimates of swelling rate and the absence of incubation dose suggest a problem with the conventional depiction of bias as a constant value applied only to dislocations. The concept of incubation is predicated on an irradiation-induced change to the effective bias of the entire sink system. While some have suggested that irradiation induced changes in network density could constitute such a change in bias [96], such changes typically occur much faster than the observed incubation periods of 100 dpa or more, and are not borne out by measurements. Other work using elasticity theory suggests that cavities, particularly small cavities, are biased toward interstitial absorption as well, effectively negating the dislocation bias early in irradiation. As some cavities overcome this

impediment through accumulation of helium, they become less biased, allowing the bias of dislocations to sustain linear growth. This compelling explanation of the incubation phenomenon then requires accurate estimates of its strength and dependencies.

The role of helium is of critical importance in both proposed mechanisms, as its aggregation drives heterogeneous nucleation, and provides much needed stability to biased cavities so they can achieve growth. While this rectifies the discrepancy of its miniscule impact on the conventional, homogeneously nucleating physics, it introduces a new one; Ion irradiation experiments with no helium additions still induce swelling in FM alloys. With these mechanisms in play, an irradiation without helium will completely suppress the nucleation of cavities. If these physical processes are indeed dominant, another process that is slower, but qualitatively similar, must be in play. Since no other inert gasses are anticipated, vacancies themselves must be accumulating in some heterogeneously distributed sites where the local energetics are suppressing the attractive forces on interstitials that create cavity bias. Much akin to heterogeneous nucleation via helium accumulation, understanding this mechanism will be predicated on knowledge of the chemical or structural nature of these sites and the strength of their effect on vacancy bias.

5.1 Initial Model Refinements

5.1.1 Cascade Damage Clustering and Helium Co-Generation

Introduction of cascade damage production as described in Section 3.1.1 results in larger clusters which climb up the capillarity curve, making them more stable against thermal emission. These act as preferred cavity nucleation sites, resulting in temperature distributions (Figure 5.1) with peaks much closer to the experimentally observed peaks at 420-460˚C, particularly in the ion-relevant case. The shift in temperature, however, remains much larger than the expected 50-70˚C. While the ion case produces a swelling peak near that measured experimentally, the reactor-relevant case manifests at temperatures well below where any neutron-irradiated ferritic steels have observed any swelling. A comparison was made to the temperature shift predicted by Eq. (2.6), modified such that instead of the simple sum of migration and formation energy, the effective dissociation energy, E,

$$E = E_m^v + E_f^v + E_b^v, \tag{5.1}$$

was used and includes vacancy binding energy, E_b^v=(2γ/r), of the largest, most stable, injected cluster (Table 5.1) given by the capillarity law described in section 3.1.4. The agreement in shifts demonstrates encouraging alignment between this model and classical rate theory. The discrepancy with experiment despite this theoretical agreement underscores the fact that there are still missing physics.

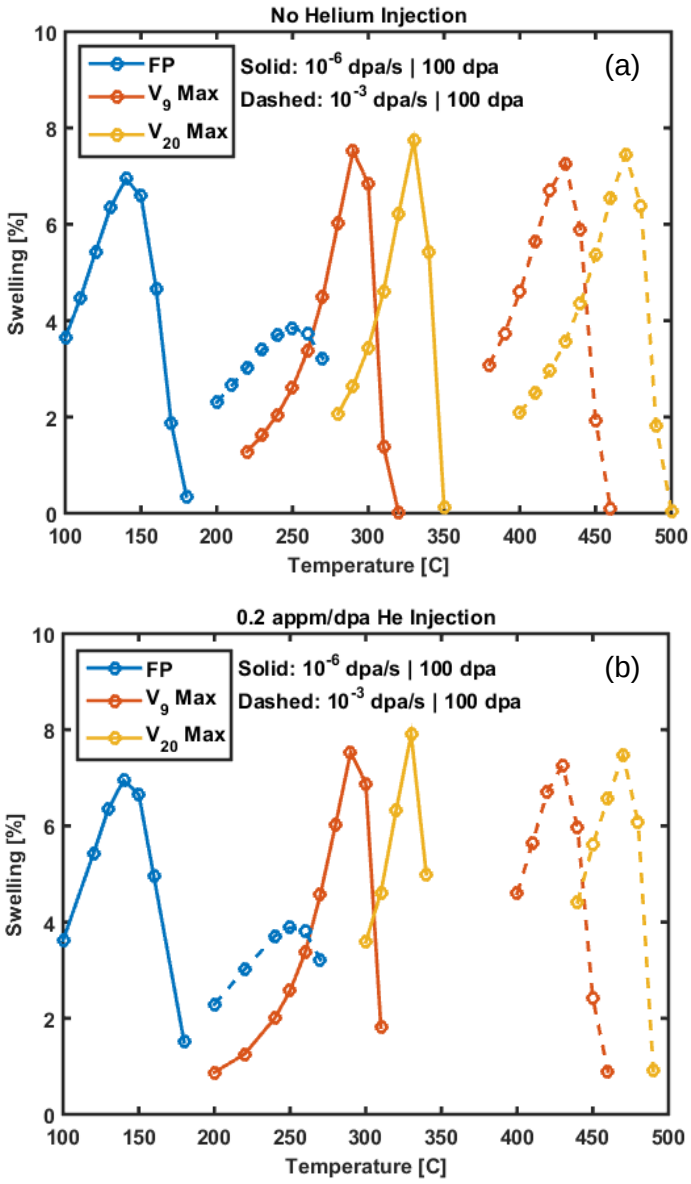

Figure 5.1: Temperature distributions of swelling with increasing cascade vacancy cluster production (a) without helium and (b) with cogenerated helium

Table 5.1: A comparison of model-generated peak swelling temperature shifts between ion and fast reactor dose rates with analytical predictions of Mansur

Damage Mode	Cluster Dynamics ΔT [K] (±10)	Computed Mansur ΔT [K]
Frenkel Pair	110	106.7
9-Vacancy Maximum Cascade Cluster Size	140	146.8
20-Vacancy Maximum Cascade Cluster Size	140	151.0

In a further effort to compensate for the apparent underprediction of cavity stability, the full two-dimensional helium-vacancy system was utilized with a constant helium generation rate of 0.2 appm He/dpa (a reasonable transmutation rate for FM steel in a fast reactor spectrum [10]). The result as shown in Figure 5.1b has a negligible effect on swelling. Only the least stable Frenkel Pair case showed any discernable change, and even then, only in swelling rates rather than peak locations. It seems that the stabilizing effect of vacancy clustering in primary damage overwhelms any effects of helium. This runs contrary to the wide body of experimental evidence for a significant stabilizing effect of helium pressure which can greatly enhance swelling, primarily by shortening incubation periods.

A more rigorous treatment of primary damage significantly improved agreement with experimental swelling, and cascades producing 9-member clusters will be incorporated in all forthcoming results. This is in part because the peak swelling temperature at high dose rate is closer to experimental observations. Furthermore, the 20-member cluster generation terms are based upon aging after the damage event, which, despite being able to consider spatial correlations, consider the same sort of aggregation reactions that cluster dynamics does. They are thus eschewed to avoid double counting these reactions in the production and reaction terms.

The inclusion of helium and its lack of an effect on swelling contradicts experimental evidence (e.g. Figure 2.5). To rectify this, all future results will include 0.2 appm/dpa of co-generated helium unless otherwise stated.

5.1.2 Spatial Dependence:

A key aspect of ion irradiation is that the damage is concentrated near the irradiated surface, with steep gradients in damage and a deeper region populated with injected SIAs. This could be having

nonintuitive effects on swelling and its temperature distribution. The spatial dependence of damage under ion irradiation is relatively well known, and its shape and depth can be reliably predicted with the Monte Carlo code SRIM [61]. Performing a polynomial fit to these stochastic profiles as shown in Figure 5.2 allows for determination of a unique damage generation term at each mesh point of a finite difference model. The finite difference calculation adds a diffusion term to each mobile defect's rate equation by approximating the second spatial derivative in Fick's second law according to

$$\left(\frac{dC}{dt}\right)_{x_i} = D\frac{\left(\frac{C(x_{i+1})-C(x_i)}{x_{i+1}-x_i}\right)-\left(\frac{C(x_i)-C(x_{i-1})}{x_i-x_{i-1}}\right)}{\frac{x_{i+1}-x_{i-1}}{2}} \tag{5.2}$$

Where D is the defect diffusion coefficient, x_i is the depth of the spatial mesh point under consideration.

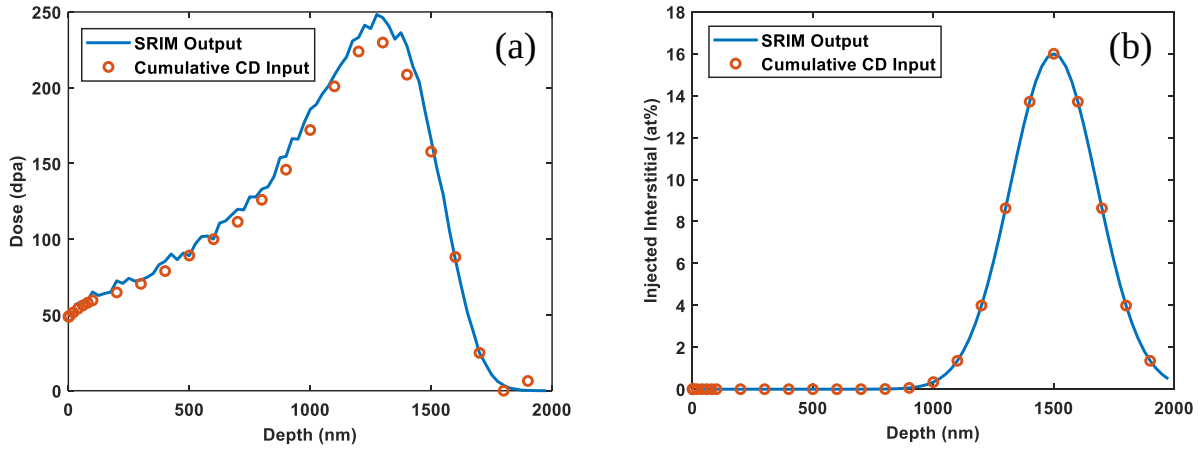

Figure 5.2: (a) Radiation damage and (b) ion implantation profiles generated by SRIM and the corresponding cluster dynamics input at each node of the spatial mesh when the analysis region (600nm) has reached 100 dpa.

Experimental work routinely avoids the regions most affected by the spatial distribution of damage utilizing FIB liftouts that allow analysis to be confined to an acceptable experimental region of interest. Figure 5.3 shows a comparison between swelling behavior at one such depth of interest, depths plagued by artefacts, and a bulk case. As expected, the near surface (100 nm) case, and the 1500nm case atop the interstitial injection peak are areas of significant swelling suppression. The 600nm depth which would be analyzed in a typical experiment shows no discernable difference in swelling behavior from that of a bulk case (such as a reactor irradiation), indicating that spatial

variation is not a concern if the surface and interstitial injection peak are avoided. In these conditions, the mean free path of a typical point defect is too short to sustain enough long-range mass transport along the damage gradient which allow the analysis region to be affected. Therefore, data taken in this depth regime is expected to be a good representation of bulk irradiation behavior in this case. This may no longer be the true in different irradiation conditions particularly at higher temperatures or lower damage rates when mean free paths are lengthened. Future results presented here will eschew spatial dependence in favor of a continuum model to ease computational demands.

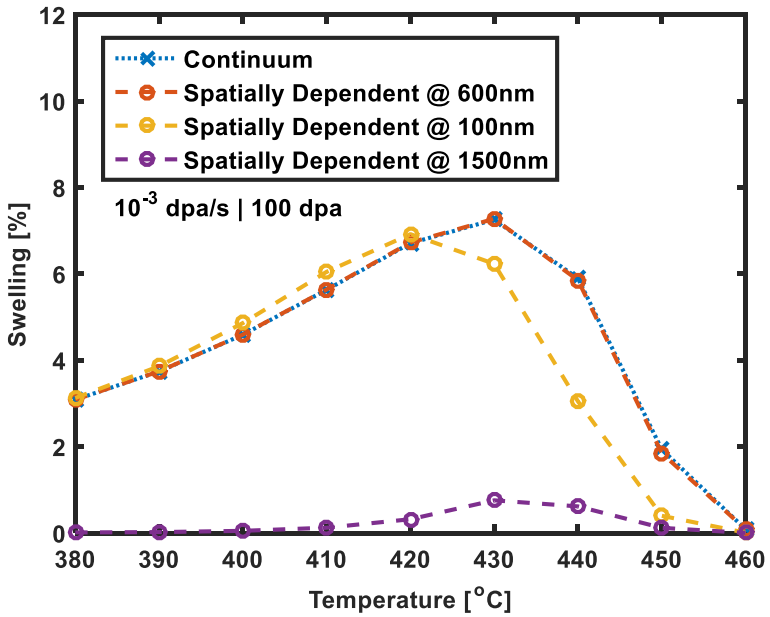

Figure 5.3: Swelling temperature distributions at various points throughout a spatially dependent damage profile. Dose rate is adjusted to be held constant at the point of interest.

5.1.3 Carbon-Vacancy Interactions

Initially, upward shifts in the swelling peak were accomplished by increasing the vacancy migration energy. This strategy produced the desired result with increases of only about a tenth of an eV in the ion case (Figure 5.4). A departure from ab initio and desorption data must be physically justified using an argument of impurity interactions. Since the emission rate constant is proportional to the capture rate, increases to the migration energy decrease both aggregation and dissociation events, thereby reducing peak swelling rates and causing upward shift in peak swelling temperature. This creates an artificial increase in cavity stability which falsely implies that each nucleating cluster is associated with an impurity. While it can be reasonable to assert

that a freely diffusing vacancy through the matrix is likely to encounter an impurity and be temporarily immobilized, there is no such reason to assume that an analogous vacancy escaping a cavity from which it recently dissociated will be reflected back to the cavity by such an encounter.

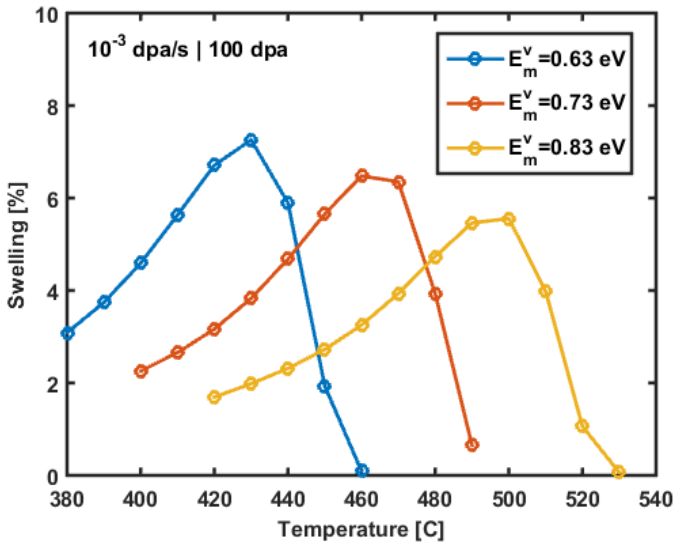

Figure 5.4: Swelling temperature distributions with small increases in vacancy migration energy.

Atomistic modeling efforts have recently computed the stabilities of these immobile complexes formed by vacancies and impurities (predominantly carbon) [41]. To remove the artificial stability increase implicit to the effective diffusion approach, these binding energies were utilized to determine the strength of this effect directly by including a new reaction pair for formation and dissociation of a vacancy-carbon complex. A static density of sessile carbon atoms was introduced with an input binding energy informing the stability of complexes. The resulting temperature distributions show that the primary effects are a reduction in void growth rate, and a small decrease in effective stability indicated by a small downward temperature shift (Figure 5.5). At a reasonable condition based on the solubility of carbon in iron and the composition, which assumes 200 appm is dissolved and uses the binding energy of a V-C complex, the effect is modest. However, pushing to higher-than-expected carbon solubility or binding energies at the magnitude of V-C_2 complexes increases the capability to retard swelling in a significant way.

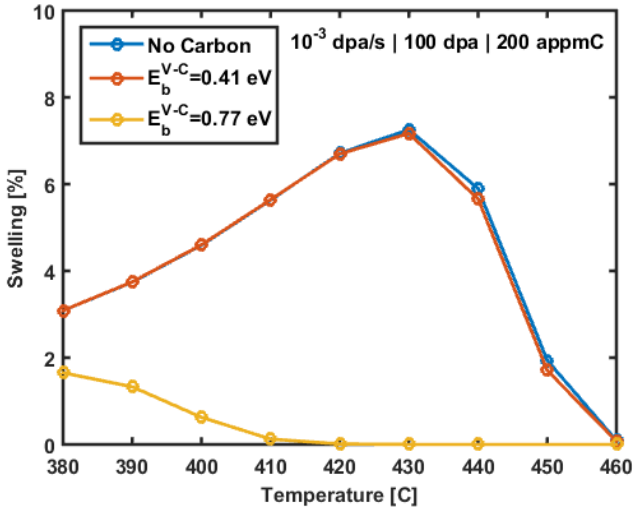

Figure 5.5: Swelling temperature distributions at the ion relevant case demonstrating the effect of carbon-vacancy complex formation with limiting monomer binding energies.

For a more rigorous treatment of these interactions, mobility terms were added to the rate equations of small vacancy clusters and carbon solutes [89], and by allowing accumulation of immobile V-C clusters up to size V_4C_2 using binding energies from ab initio calculations and simple extrapolation [41] (Table 3.2). Figure 5.6 demonstrates that when these more tightly bound complexes are formed explicitly via capture of two carbon atoms, they do not sustain densities sufficient to significantly impede cavity nucleation. The secondary nature of this effect clearly indicates that it is not the dominant physical phenomenon responsible for the incubation period. The effects of carbon on swelling are most likely relegated primarily to the formation and coarsening of carbides which act as unbiased sinks competing with cavities for defect flux [26].

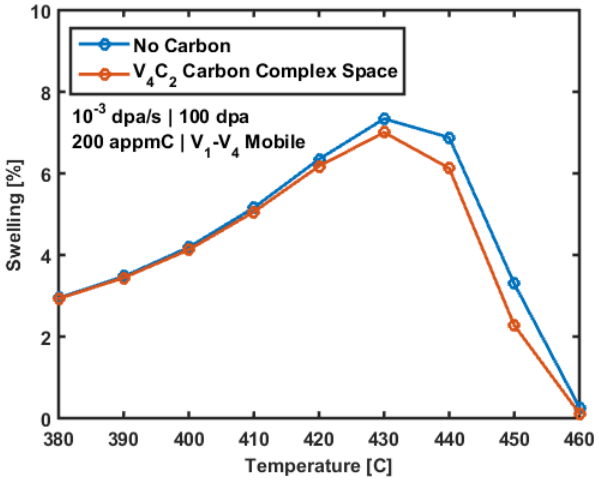

Figure 5.6: Swelling temperature distributions where clusters of up to 4 vacancies are mobile and can form complexes with up to 2 carbon solute atoms

While most of the mechanisms examined demonstrated the capacity to significantly affect either steady state swelling rates or the temperature regime where swelling occurs (ergo nucleation), they were, like the parameter variations in Figure 4.3, fundamentally unable to reproduce the key experimental observations. This is clearly seen in Figure 5.7, where reasonable swelling rates only occur with low dislocation bias, incubation is absent, temperature shift is too large (140˚C) and helium has almost no effect. Accounting for spatial effects and impurity trapping, which modify net vacancy flux can reduce swelling rates, but not as significantly as a modest change in net dislocation bias. Even with the superimposition of the most swelling-retardant parameters which could be considered reasonable, their impact is not substantial enough to counteract an increase in net dislocation bias to expected values in the tens of percent. Furthermore, there is no obvious way to rationalize an evolution of the strength of these mechanisms over the course of the irradiation, which prevents them from reproducing incubation behavior. The effect of clustered cascade damage on stability affects each dose rate condition equally, causing any changes in location of the temperature peak to occur by approximately equal amounts, with no significant impact on temperature shift. In order to reduce the temperature shift toward expectation, a mechanism that disproportionally affects one temperature distribution is required. In summary, while cascade damage, spatial dependence, and carbon-vacancy interactions create a more nuanced understanding of the fate of radiation-induced defects, their effects cannot be responsible for the outlying discrepancies between modeling and experiments.

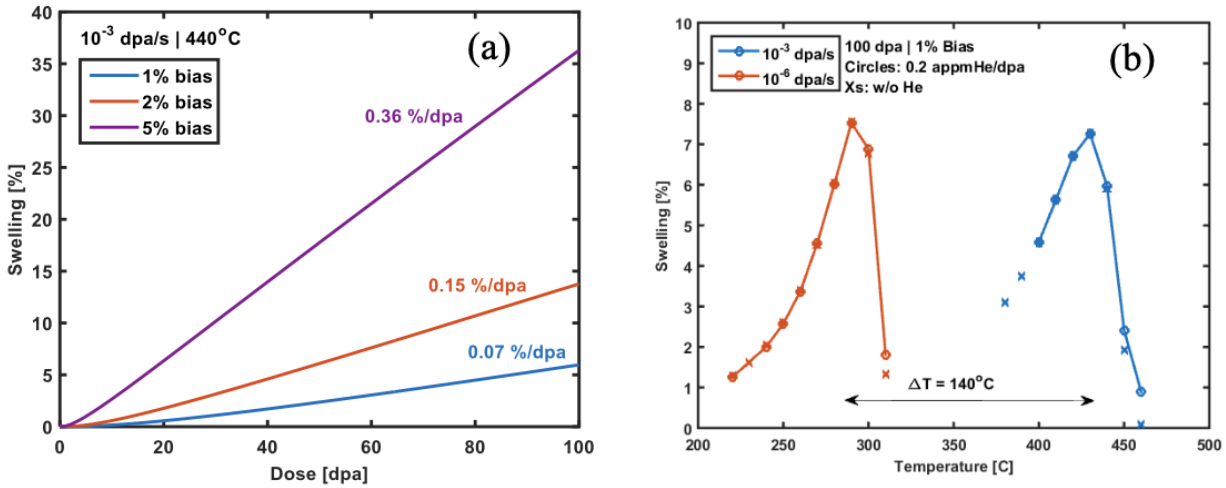

Figure 5.7: (a) Swelling as a function of dose at the peak swelling temperature for ion beam dose rates with varied net dislocation bias and (b) the temperature distributions of swelling at 100 dpa for ion beam and fast reactor relevant dose rate conditions, with and without 0.2 appm/dpa of cogenerated helium.

5.2 Cavity Bias

The incorporation of a net bias toward interstitials on small cavities produced a strong barrier to nucleation that fundamentally changes the pathway to cavity nucleation. Unlike thermal emission as the primary obstacle to nucleation, which can primarily be overcome by shifting irradiation conditions to a temperature where it is weakened, cavity bias uniformly impedes the smallest cavities at all temperatures. The only recourse is the slow accumulation of helium, which both combats thermal emission and bias directly by pushing back against the intrinsic strain field around the void as pressure meets or exceeds mechanical equilibrium. As co-generated helium accumulates in the embryonic cavity population, it slowly begins to grow as bias is suppressed and, in sufficiently over-pressurized bubbles, interstitials are punched out. When the bubbles are sufficiently large, the bias on them decays geometrically (as shown in Figure 3.4) until it reaches approximate parity with the opposing bias on dislocations. Bubbles exceeding this critical size then begin a bias-driven growth period. A representative cluster density map showing this phenomena is presented in Figure 5.8, where helium content is transformed into equivalent radius space according to $R = \left(3\Omega m / 4\pi\right)^{1/3}$.

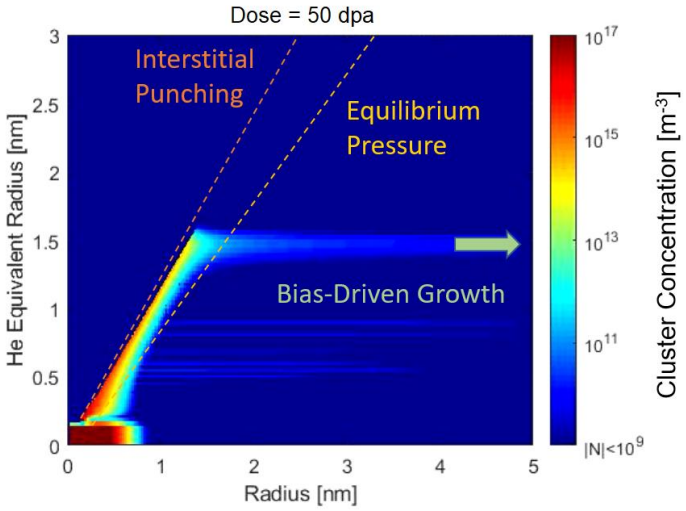

Figure 5.8: A representative two-dimensional (*n* vacancy × *m* helium) cavity size distribution showing the pathway to nucleation in a cavity-biased system.

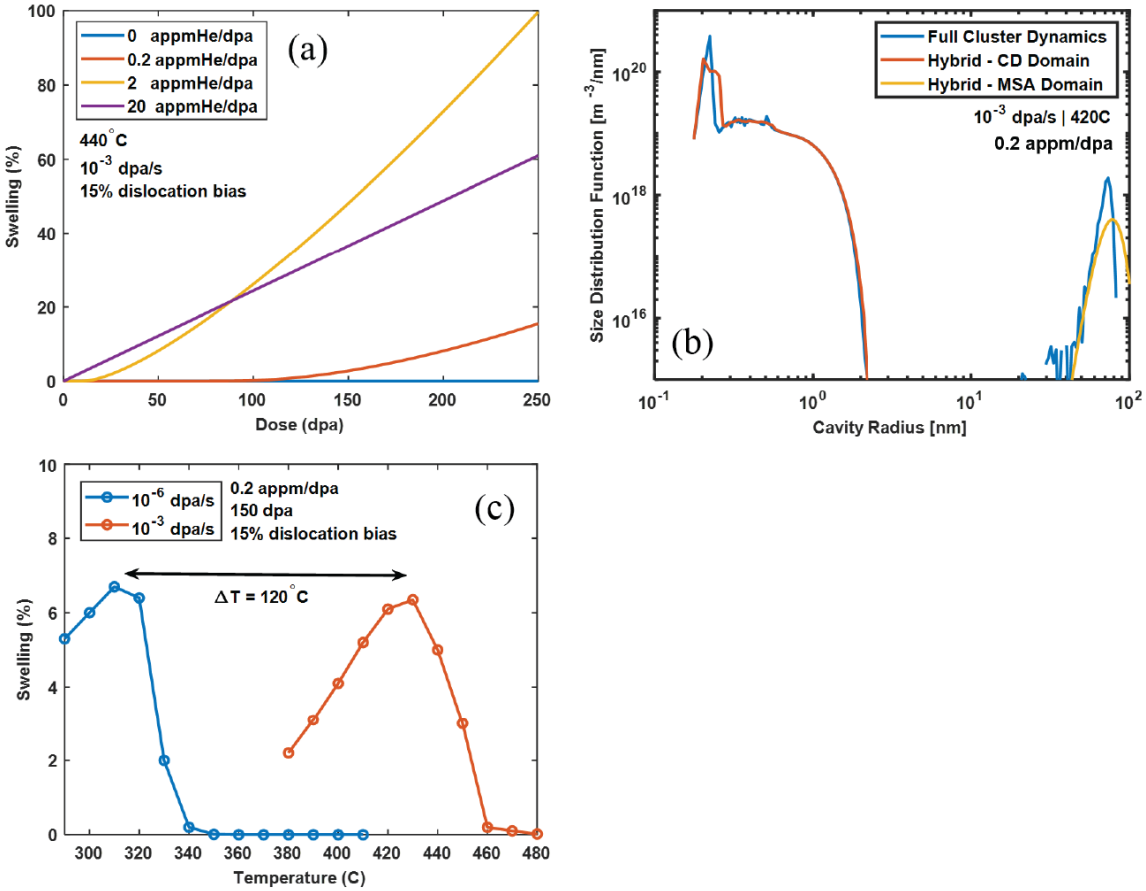

Figure 5.9: Behavior of swelling when accounting for cavity bias; (a) predicted swelling as a function of dose and helium co-injection rates at high dose rate. (b) representative bi-modal cavity size distribution at high dose rate for a full-phase-space and hybrid cluster dynamics/mean size approximation. (c) temperature distribution of swelling at fast reactor and ion irradiation dose rates and

This new nucleation behavior directly reproduces three key swelling behaviors absent in previous iterations of the model: incubation dose, bi-modal size distribution, and low steady state swelling rate as shown in Figure 5.9.

For studies conducted in a system where thermal evaporation is the only barrier to nucleation, a given irradiation was either in a temperature regime where it was weak enough for swelling to occur or one where it was not. Cavity bias is strong enough at all temperatures to retard nucleation, in fact, its magnitude increases with decreasing temperature (Figure 3.4). Therefore, one cannot simply find a lower temperature regime where emission is weak enough to sustain rapid nucleation upon the onset of radiation, necessitating the often slow process of helium accumulation. At low generation rates it can take as much as 100 dpa to accumulate the helium concentration necessary to reach the critical size and promote bias driven growth. During this period, negligible swelling occurs due to the very small sizes of these bubbles and the high degree of cavity-bias-driven recombination, leading to an incubation period as shown in Figure 5.9a. As helium generation rate increases, these periods get progressively shorter as it takes less time to reach the helium content necessary to transition to growth, although this is sometimes accompanied by a lower swelling rate once linear growth is reached. Furthermore, the persistent, stable density of subcritical bubbles gives way to a bi-modal size distribution, shown schematically in Figure 5.9b. The maximum steady state swelling rate, at least for the reactor-relevant helium generation rate of 0.2 appm/dpa is now a more reasonable 0.1 %/dpa even with a dislocation bias 15 times larger than the initial results in Chapter 4:.

Despite the profound changes to the physics of cavity nucleation introduced by cavity bias, the temperature dependence at each dose rate condition remains fundamentally unchanged. Figure 5.9c shows that swelling temperature peaks and the consequent temperature shift are effectively the same as the unbiased work at the same cascade conditions (Figure 5.7b). Thermal emission still fundamentally controls this temperature dependence in an indirect way, by governing helium partitioning. At low temperature, where very few vacancy clusters dissociate, helium content is divided among a large density of cavity nuclei. Not only does this high density limit the pressure available to a given bubble, it is also a significantly stronger sink, promoting recombination rather than growth. At high temperature helium has an extremely low density of vacancy clusters to partition to. The few that survive can certainly grow, but their density is too low to amount to

significant swelling. If helium can bind to other sites in the microstructure, such as grain boundaries and dislocations which have sink strengths dwarfing that of cavities in this regime, helium is even less likely to find and stabilize them, resulting in even lower densities. Once again, the optimal swelling conditions arise when an intermediate cavity density is realized, where the helium is not spread too thin to overcome cavity bias, but there is a sufficient density to create significant swelling. Figure 5.10 shows a direct comparison with experimental cavity size distributions and demonstrates encouraging qualitative agreement. The experimental and CD bubble size distributions predict a peak of around r = 1 nm which is relatively insensitive to irradiation temperature. Growing cavity distributions also exhibit consistent trends. As temperature increases, less cavities can transition from bubbles to voids, but the ones that do have less competition for defects and are able to grow faster. At high enough temperature (>500˚C), cavity density is too low to constitute significant swelling, and their low sink strength precludes growth. The main discrepancy between the size distributions is quantitative. While swelling is reasonable, the model size distributions are dominated by very low densities of very large cavities.

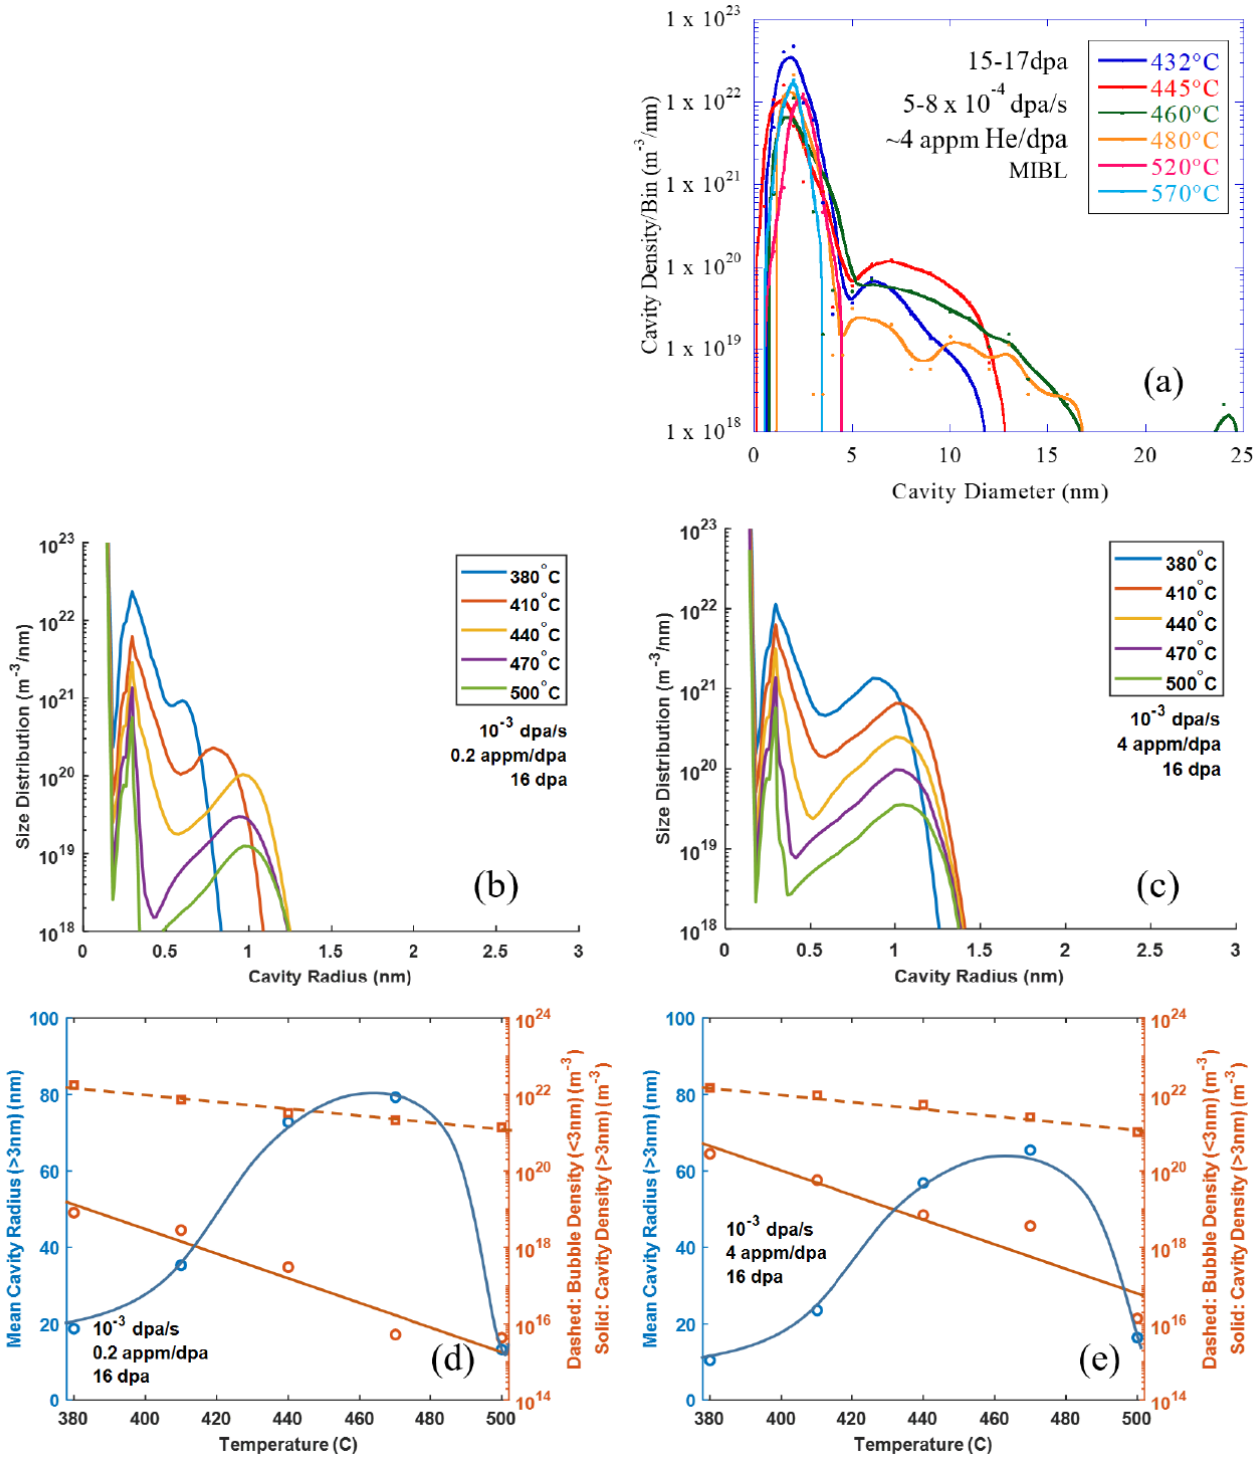

Figure 5.10: Cavity size distributions resulting from (a) 5 MeV Fe ion irradiation in the Michigan Ion Beam Laboratory characterized by S. Taller [82], compared against (b) cluster dynamics bubble size distributions at 0,2 appm He/dpa and (c) 4 appm He/dpa, as well as (d) mean radius and density of cavities for 0.2 appm He/dpa, and (e) 4 appm He/dpa for a series of temperatures.

Performing a systematic study on the effect of helium generation rate reveals a dependence qualitatively similar to the temperature peak, as shown in Figure 5.11. Without helium the cavity-biased model will never predict significant swelling. Very low helium generation rates (e.g. 0.02 appm/dpa) may eventually produce significant steady state swelling, but the incubation periods become so long they exceed the final dose of most ion irradiation experiments or the lifetime dose of a reactor component. Without a sufficient quantity of helium, no cavity nuclei can reach a large enough size to begin bias-driven growth. At very high helium generation rates (200+ appm/dpa), the density of bubbles becomes very large and competition for vacancy flux is fierce. The high sink strength this creates competes with the biased dislocation network, weakening the driving force for cavity growth. Furthermore, the available vacancy supersaturation is subdivided between so many embryos that any one of them is unlikely to reach the critical size. The intermediate region between these two extremes optimizes swelling conditions. This maximum occurs at approximately 20 appm/dpa at 50 dpa, which is consistent with the experimental data in Figure 2.4 (albeit for a different alloy system).

Cavity size distributions from cluster dynamics capture key behaviors of the experimental result (Figure 5.11a-c). Bubble and cavity density have an intuitively strong, positive correlation with helium generation rate in all cases. Lower helium generation rates, such as those used experimentally to replicate reactor transmutation, exhibit a relatively consistent cavity size. The main experimental disagreement with respect to helium generation rate is an opposite relationship with peak bubble size. The tendency toward size distributions which are dominated by low densities of large cavities also persists.

A comparable helium generation rate series conducted at fast reactor dose rate at its peak swelling temperature of 320˚C, suggests that trends in helium behavior are mostly intact (Figure 5.11d-e). Bubble size distributions tend to be slightly broader than their high dose rate counterpart, but the peak bubble size and density dependence are consistent. Growing cavities have qualitatively consistent behavior, but the peak in helium generation rate forms slightly lower. This is particularly encouraging considering the 4 appm/dpa ion irradiation data effectively emulated the BOR-60 fast reactor data at nominal 0.22 appm/dpa [97].

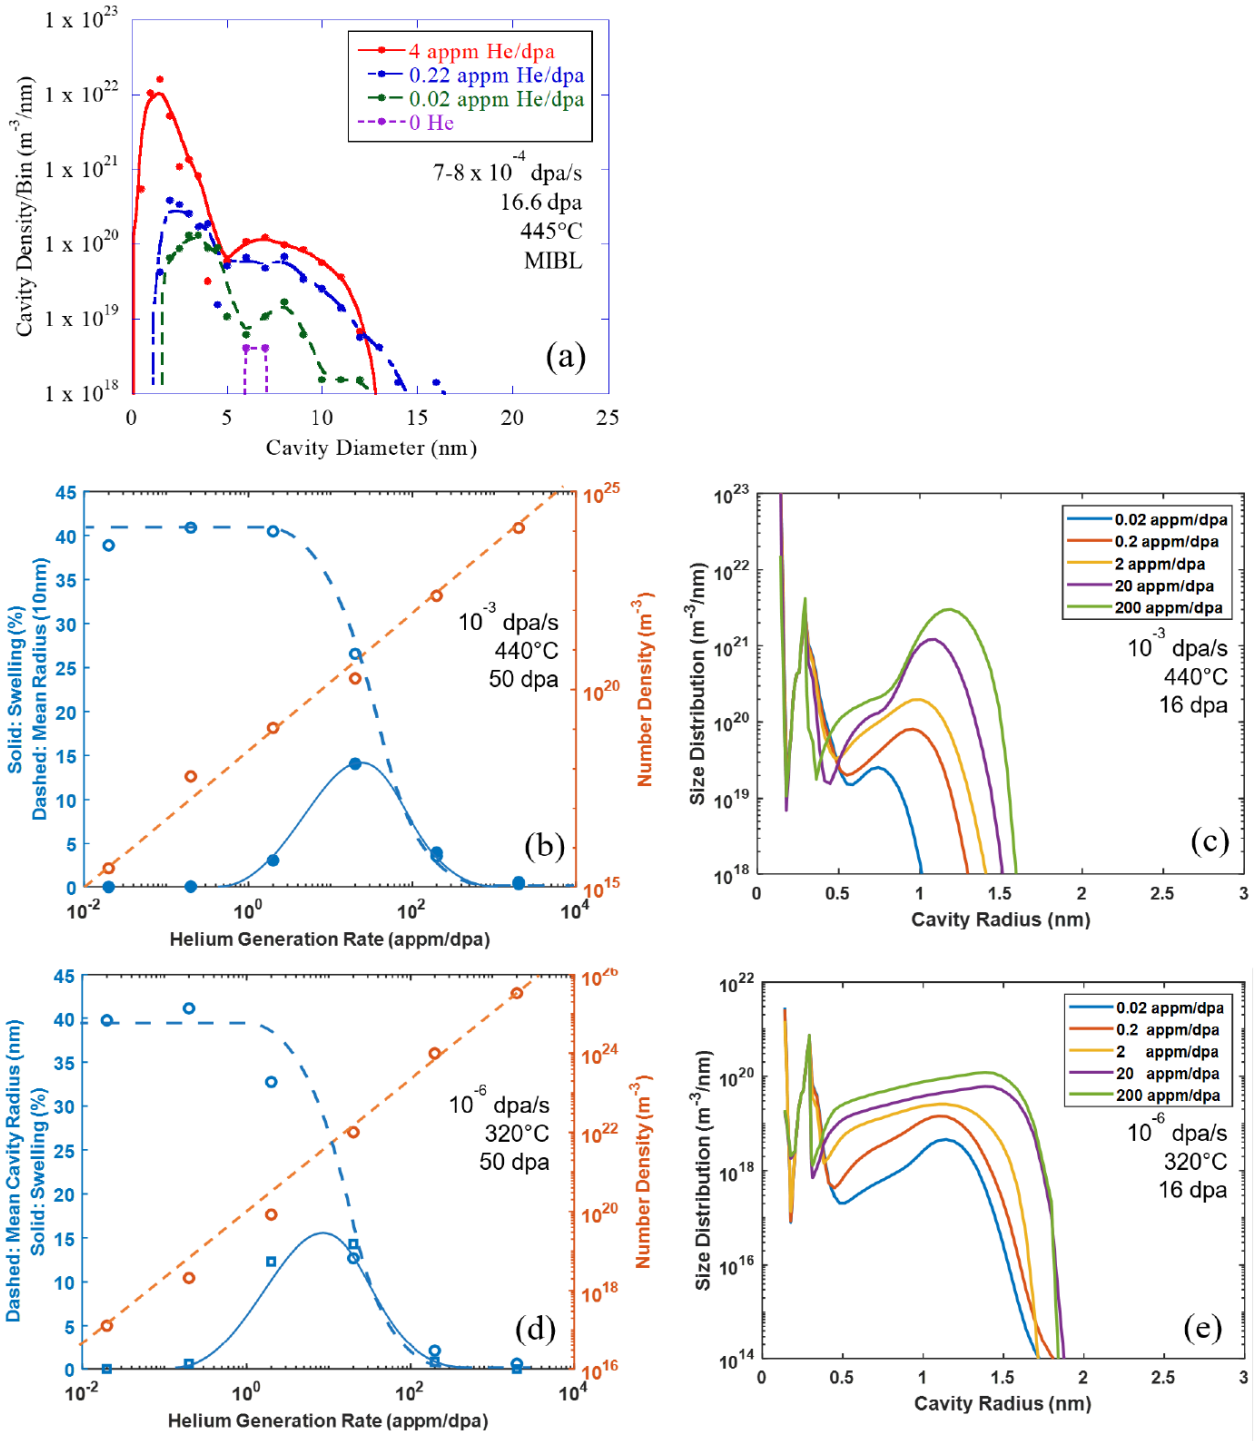

Figure 5.11: Cavity size distributions resulting from (a) 5 MeV Fe ion irradiation in the Michigan Ion Beam Laboratory [82], and compared against (b) CD model predicted swelling, mean cavity size and cavity density for ion, and (d) fast reactor irradiation dose rate at 50 dpa, and (c) bubble size distributions for ion and (e) fast reactor irradiation dose rate to 16 dpa for various helium co-generation rates.

Focusing on a single constant point in helium generation space should not be considered a complete descriptor of helium dependence. As an example, prediction of the swelling as a result of helium to dpa ratio at a dose of 50 versus 150 dpa is shown in Figure 5.12, and demonstrates that the higher dose produces a shift in the maximum down to 2 appm/dpa. Lower helium conditions exhibit longer incubation periods, but once linear growth is achieved, swelling rates tend to be much larger. This can be attributed to the significantly lower density of bubbles that compete with growing cavities to absorb the newly produced mobile vacancies. This will likely continue to slowly shift down as dose increases, which is consistent with experimental work [40,81].

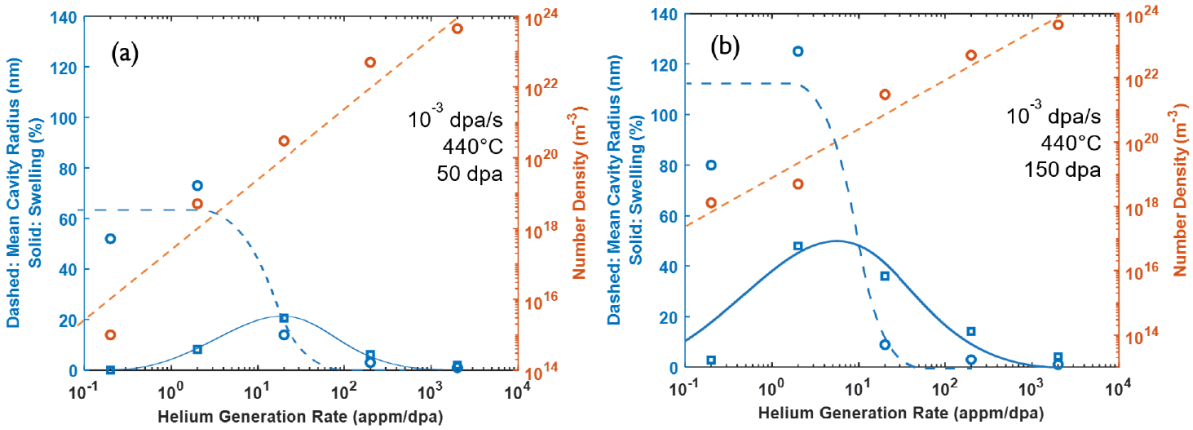

Figure 5.12: Swelling, mean cavity size and cavity density as a function of helium co-injection rate for ion irradiation dose rate to 50 and 150 dpa

In summary, cavity bias fundamentally changes the nucleation pathway of a cavity by providing a previously missing barrier, resulting in three key experimental behaviors being qualitatively reproduced:

- Helium promotes nucleation that is necessary for swelling to occur, but in abundance suppresses swelling via an over-nucleation of subcritical bubbles, resulting in a peak with respect to generation rate which is qualitatively similar to temperature.
- At low transmutation-relevant helium generation rates, the slow rate of helium accumulation creates an incubation period in which negligible swelling occurs, before eventually achieving a critical size where a transition to linear swelling begins.

- Cavity bias counteracts dislocation bias as a driving force for cavity growth, resulting in reasonable linear swelling rates even when net bias on dislocations is increased to a more reasonable 15-20%.

Because the nucleation phenomenon as a function of temperature is still dependent on the same underlying mechanism, thermal emission, the dose rate dependence of these temperature distributions is unchanged. This implies the existence of a missing nucleation mechanism which does not share the same temperature and dose rate dependencies.

5.3 Heterogeneous Nucleation

A heterogeneous nucleation mechanism provides a means for helium to self-cluster independently of the available vacancy population, which, given an appropriate density and interaction energy of available nucleation sites, can fundamentally change the temperature dependence of swelling. No longer are temperature peaks determined solely by thermal emission of vacancy clusters sustaining a cavity density that balances under-nucleation where most cavities evaporate, and over-nucleation where helium is spread between too many clusters for any to achieve growth. The effect of these sites is twofold. First, the density of nucleation sites is temperature independent, effectively providing a floor for the number of available locations for helium to partition. Second, a new temperature dependence is explicitly introduced by the thermal dissociation reaction defining the strength at which helium is bound to, and thus immobilized by, these sites.

The effect of heterogeneous nucleation on peak swelling temperature at each dose rate and several site configurations is shown in Figure 5.13 and Figure 5.14. Provided a sufficiently high site density and binding energy, a consistent upward shift in peak swelling temperature was realized in all cases. Site binding energy demonstrated particular sensitivity, with 0.5 eV-bound cases barely affecting the location of the temperature peak compared to the homogeneous nucleation case, and then 1.5 eV case shifting the peak over 100°C at fast reactor dose rate. Given the same binding energy, site density exhibited an almost log-linear relationship with peak swelling temperature over a threshold near 10^{19} m^{-3}. This mirrors trends in cavity density as a function of temperature, demonstrating that providing more sites raises the floor on cavity nucleate density.

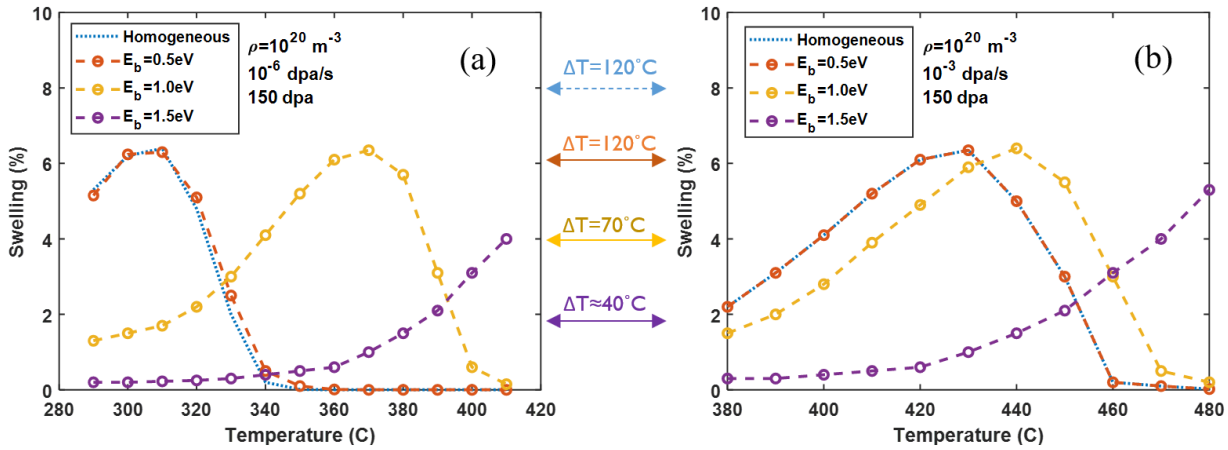

Figure 5.13: Swelling at 150 dpa as a function of temperature at (a) fast reactor and (b) ion irradiation dose rates generated by a cavity-biased, heterogeneously-nucleating cluster dynamics simulation for a range of helium-site binding energies.

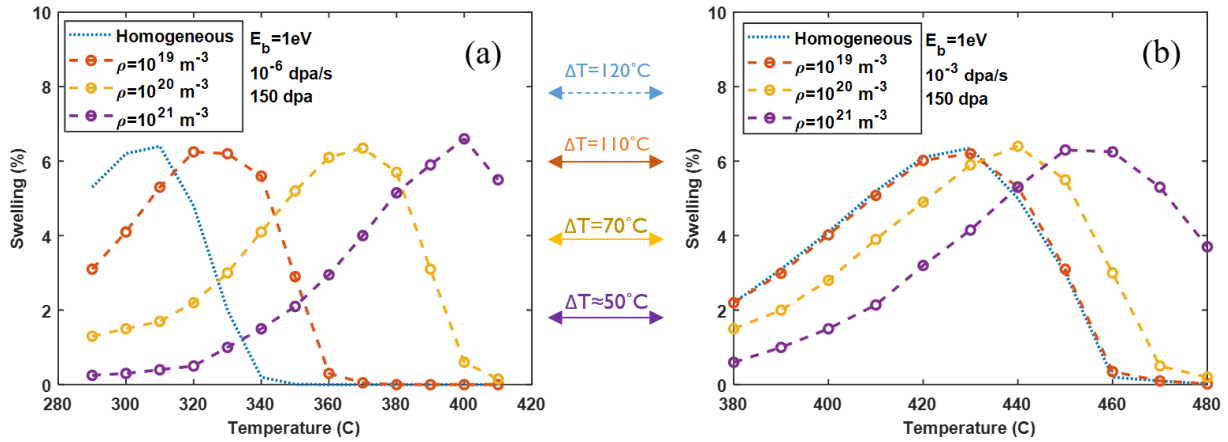

Figure 5.14: Swelling at 150 dpa as a function of temperature at (a) fast reactor and (b) ion irradiation dose rates generated by a cavity-biased, heterogeneously-nucleating cluster dynamics simulation for a range of nucleation site densities.

The enhancement to nucleation is disproportionately strong at low dose rate. At low dose rates, the characteristic time between damage events is much longer, allowing significantly more thermal emission, which is independent of incoming damage. This means that peak swelling at low dose rates occurs at lower temperatures where thermal emission rates are lower. This higher propensity for thermal emission increases the importance of an alternative nucleation site which is independent of damage production and defect microstructure. Furthermore, since swelling peaks at low dose rate would ordinarily occur at much lower temperatures than their high dose rate counterparts, the dissociation of helium with nucleation sites is less prevalent. The result is a significant reduction in temperature shift toward experimental prediction.

A reasonable value of temperature shift is achieved with a site binding energy of about 1 eV and site density at or above 10^{20} m^{-3}. Guided by predictions of helium interactions with various microstructural features derived from atomistic modeling (Table 5.2), one may begin to speculate on the identity of these nucleation sites. This work asserts that interactions on the order of 0.5 eV are insufficient to significantly affect swelling, so coherent precipitates as they are presented in atomistic literature [98] can be ruled out as too weakly interacting. It is worth noting that helium-precipitate interactions are the least represented in the atomistic literature, likely owing to the increase in model complexity that comes with incorporating more atomic species and the requisite interatomic potentials thereof. Grain boundaries, depending on the orientation, have suitable binding energies. However, in swelling relevant conditions, both ion beam and fast reactor experiments have observed void-denuded zones in their vicinity [6,19], implying that these explicitly retard swelling rather than aiding nucleation. This leaves dislocations, which strongly interact with helium and are distributed throughout the microstructure as loops and network dislocations, as the most likely candidate at present.

Table 5.2: Maximum binding energies between interstitial helium and key microstructural features in bcc iron

Helium Interaction	E_b
Edge Dislocation [29]	2.3 eV
Screw Dislocation [30]	1.0 eV
Grain Boundary [88]	0.55-2.3 eV
Coherent Cu Precipitate [98]	0.6 eV

A series of helium generation rates was also investigated for the configuration of the heterogeneous nucleation model that best emulated experimental temperature shift ($\rho_{sites}=10^{20}m^{-3}$, $E_b=1eV$). The resulting swelling distributions are compared directly against a homogeneous (control) condition at the heterogeneous peak swelling temperature in Figure 5.15. In the homogeneous control cases (Figure 5.15a), the optimum helium generation rate for ion dose rate conditions is about a factor of 5 larger than the fast reactor case, when compared at their peak swelling temperature in homogeneous conditions (320˚C and 440˚C respectively). The low dose rate conditions also exhibit a lower magnitude of swelling at optimal conditions than the high dose rate counterpart. These trends persist when comparing dose rates at their peak swelling temperature in heterogeneous nucleation conditions with consistently higher peak helium generation rates and larger swelling magnitudes at those peaks.

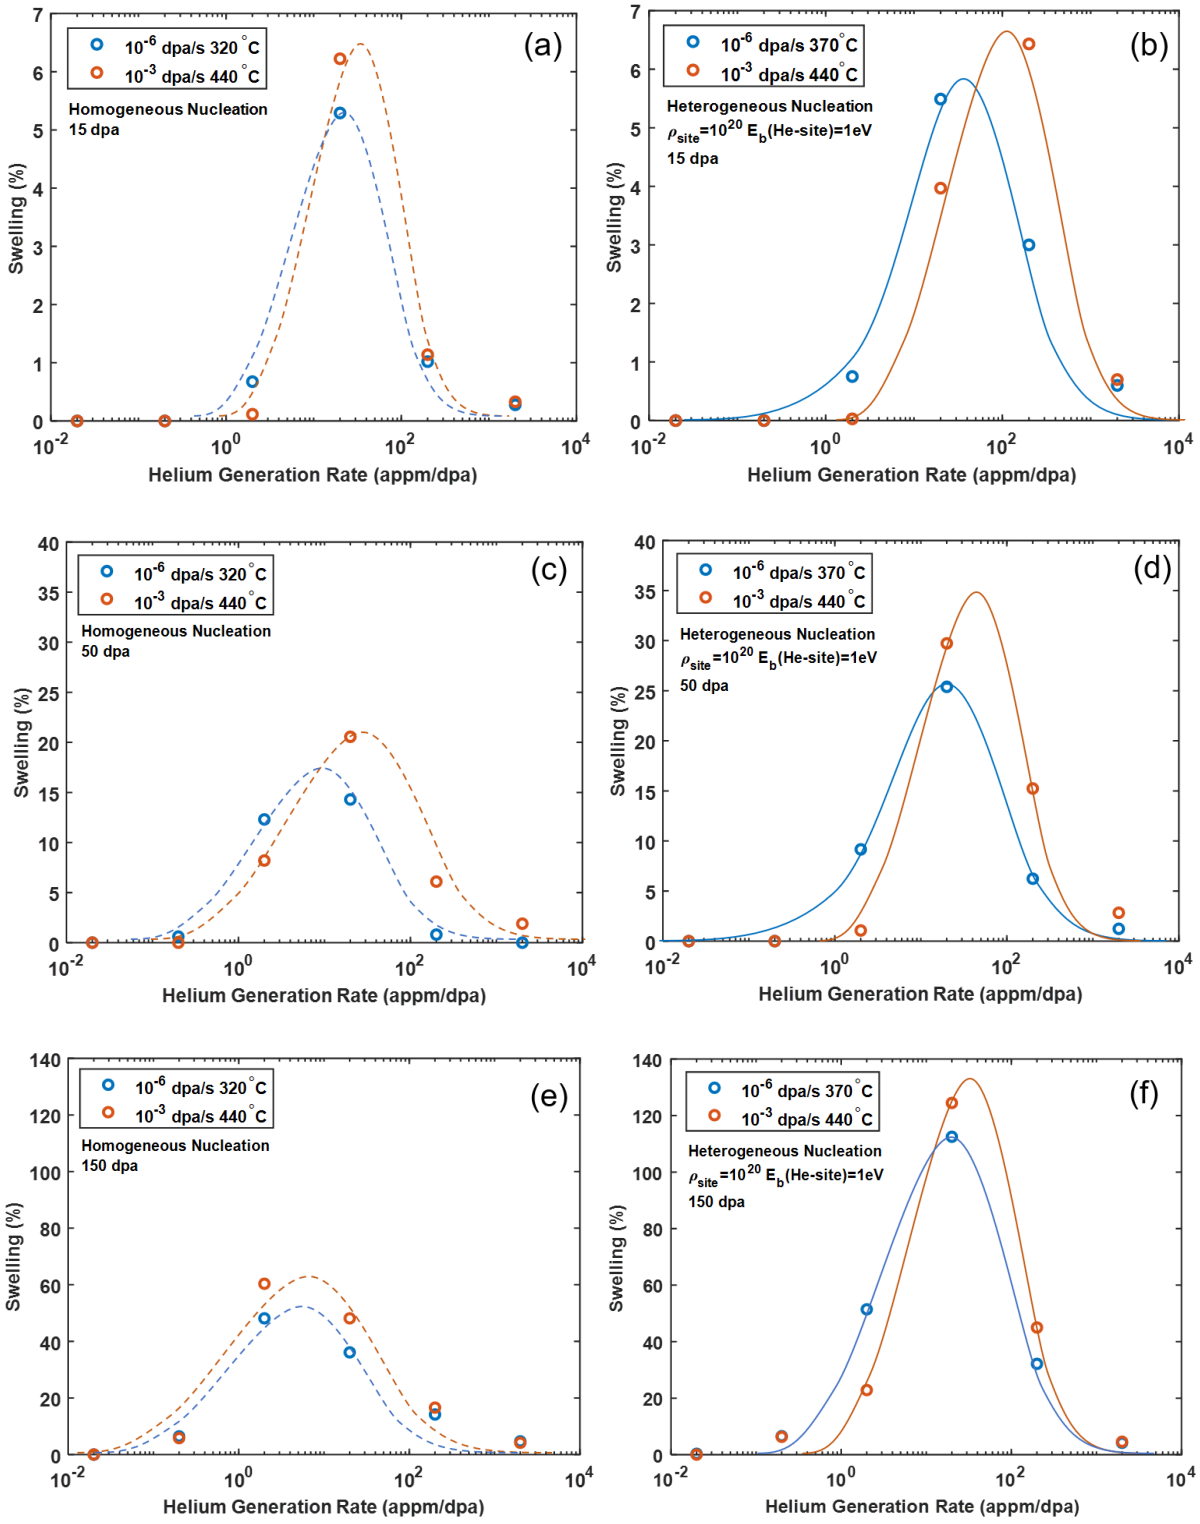

Figure 5.15 Swelling as a function of helium generation rate for ion and fast reactor dose rates to 15 (a-b), 50 (c-d), and 150 dpa (e-f) with homogeneous nucleation and heterogeneous nucleation with E_b=1eV and site density of 10^{20}m^{-3}

These observations can be rationalized by considering helium partitioning between different nucleation sites. Generally, high dose rates can sustain higher densities of vacancy clusters, owing to less time for thermal emission between damage events. Since helium is divided between more embryonic cavities, the pressure per cavity is reduced, lengthening incubation periods as more total helium inventory is required to sustain cavity growth. However, once a transition to growth occurs, a greater number of cavities are present to make that transition. Since the biased-cavity model is in a regime dominated by small densities of large cavities, this does not result in over-nucleation and much larger swelling magnitudes are observed. The same argument can be leveraged to explain the consistently higher peak in the helium generation rate distribution in heterogeneous nucleation conditions, which introduces even more partitioning sites.

Given the consistent upward shift in peak helium generation rate for high dose rate conditions, which is only sharpened by the inclusion of heterogeneous nucleation, it may be prudent for researchers to co-inject helium at rates 5-10x higher than calculated reactor transmutation rates in order to more closely simulate reactor swelling. The time-dependent nature of the phenomenon raises some concerns. At 15 dpa (Figure 5.15a-b), the ratio of helium generation rate maxima is approximately a full order of magnitude, particularly in the heterogeneous nucleation case. However, once 150 dpa is reached (Figure 5.15e-f), and most cases are firmly in the growth regime, peaks shift toward lower generation rates (consistently with Figure 5.12) and closer together. In the homogeneously nucleating case, the distinction between high and low dose rate all but disappears, and in heterogeneous conditions it is significantly reduced.

To summarize, a heterogeneous nucleation mechanism fundamentally changes the temperature dependence of swelling by partially decoupling nucleate density from damage rate. The role of characteristic dissociation event is shared by the evaporation of the smallest vacancy clusters and that of helium with nucleation sites independent of radiation-induced defects. Low dose rate cases are more prone to evaporation and tend toward lower steady-state concentrations of point defects, resulting in a swelling peak which occurs at lower temperature. Since sites trap helium more strongly at lower temperature combined with these lower defect concentrations increase the relative importance of sites to the nucleation process at low damage rates. They are disproportionately affected, causing a reduction in the magnitude of the temperature shift. Given an appropriately configured heterogeneous nucleation mechanism, in this case a binding energy of

about 1eV and a density around $10^{20}m^{-3}$, the peak swelling temperature falls in the experimentally observed regime. Based on atomistic modeling estimates of the energetics of helium interactions, the sites are likely dislocations. These same site-partitioning arguments were leveraged to assert that reactor emulation would benefit from another irradiation variable shift, a 5x-10x increase in helium cogeneration rate for dual ion irradiations, with a caveat that this may no longer be valid at high dose levels.

5.4 Impurity-Driven Nucleation

With significant improvements to experimental agreement in swelling rate, incubation, period, helium effects, and temperature shift, the main outstanding discrepancy is the inability for swelling to occur in the absence of helium. Cavity bias was responsible for many of these improvements, but it is simply too strong a barrier to nucleation to be overcome without a stabilizing presence to counterbalance it. Self-ion irradiations have consistently demonstrated the ability to introduce swelling with no co-injected or pre-implanted helium [18,19,40]. This implicates a process qualitatively similar to helium accumulation, which slowly facilitates changes to the stability or bias of the system in order to transition between an environment which is not prone to swelling and one that is. This section will introduce solute segregation as one such potential mechanism.

5.4.1 Cavity Bias Calculations

Molecular statics (MS) calculations were carried out following the procedures outlined in Section 3.3 to determine the effect of impurities introduced to the surface of a cavity. Phosphorus in particular was a focus of this work, because its interatomic potential is widely used and has been shown to accurately reproduce key properties of bcc iron [92]. The elastic interactions between point defects and a 59-vacancy void in the center of the simulation box are presented in Figure 5.16. Comparing against the pure iron control, the phosphorous has a nearly imperceptible effect on the interstitial map. However, the vacancy interaction map which is comparatively sparse and confined to the void surface, now exhibits a small plume of additional interaction sites. While this may appear small in this context, when generalized to an elastic extension applicable to CD, as presented in Figure 5.17a, the effect is an increase of nearly a factor of 2x.

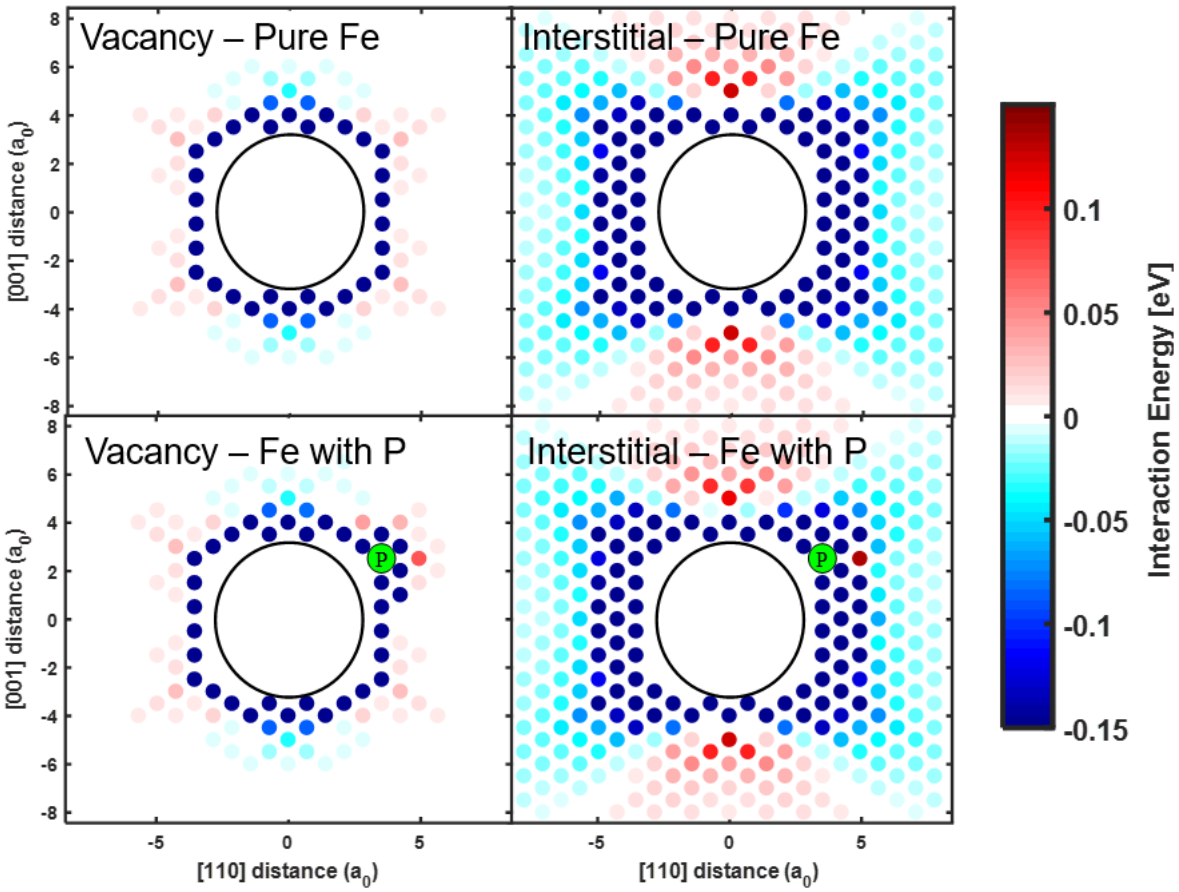

Figure 5.16: Interaction energy maps for each defect type oriented about at 59-member void with and without a phosphorous atom decorating the surface of the (111) direction.

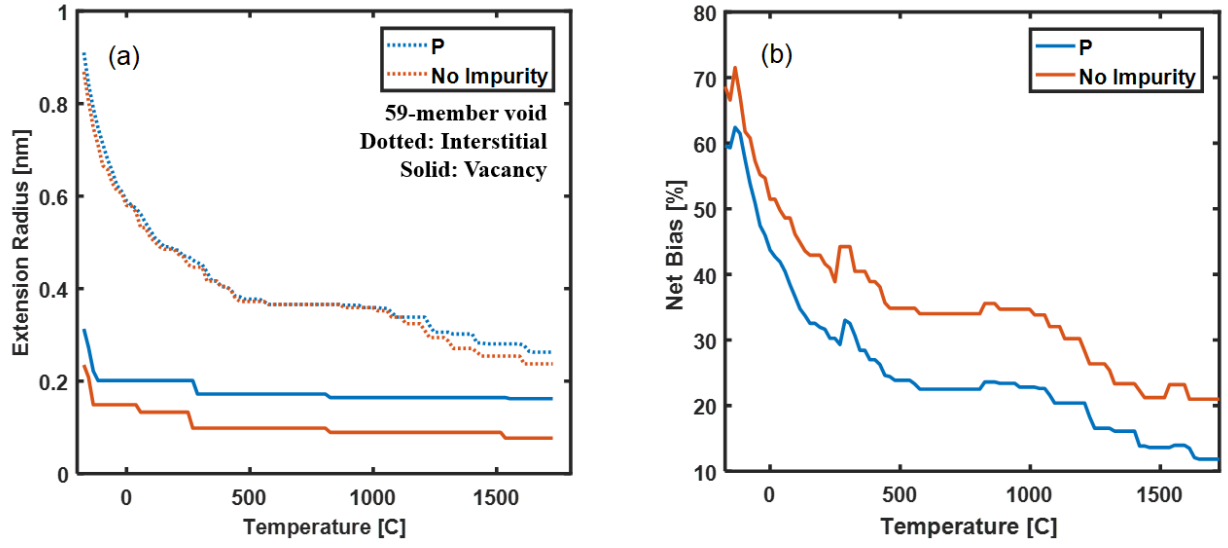

Figure 5.17: (a) Elastic extension to cavity defect capture radius and (b) net bias toward interstitial capture as a function of temperature with and without a phosphorous impurity

Figure 5.17b shows the resulting suppression to net bias as a result of this disproportionate increase in vacancy capture. All temperatures of relevance to swelling (300-500˚C) experience a net bias which is reduced by approximately one third. Figure 5.18 shows that this effect is not only strong, but cumulative, increasing to an almost 80% reduction relative to pure iron when additional phosphorus is added to the (111) directions. This effect is not uniform, however. Figure 5.18 also shows that impurities positioned in the least favorable orientation relative to the void center – (100) – have almost no effect on bias. A few other orientations were considered, but they all fell between these two limiting cases.

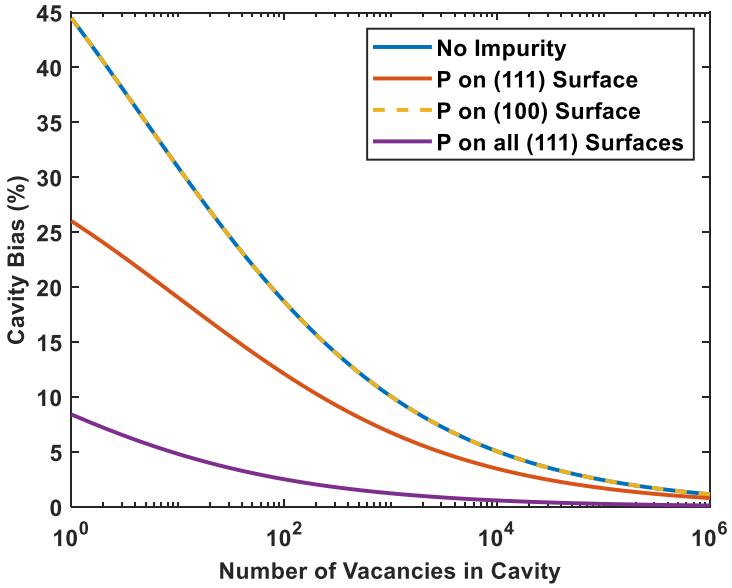

Figure 5.18: Cavity bias as function of size with a phosphorous atom oriented on the (111) surface, the (100) surface, and all six surfaces in the (111) family.

Generalizing this treatment to several other impurity types present in FM alloys illustrates qualitatively similar effects (Figure 5.19). All substitutional impurities studied reduced cavity bias by a magnitude proportional to their size mismatch with iron. Phosphorus, being the most undersized, had the most pronounced effect. Only carbon, which was placed on an octahedral interstitial site increased bias toward interstitial capture. The conclusion is that substitutional solutes facilitate vacancy capture while interstitial solutes have the inverse effect. In order to rule out conflated effects these impurities may have on the dissociation events, binding energies were also calculated with the impurities present on the surface (Figure 5.20). These results produced values in line with literature values [70] and collapse to the capillarity law used to approximate the binding energy of larger clusters in CD simulations.

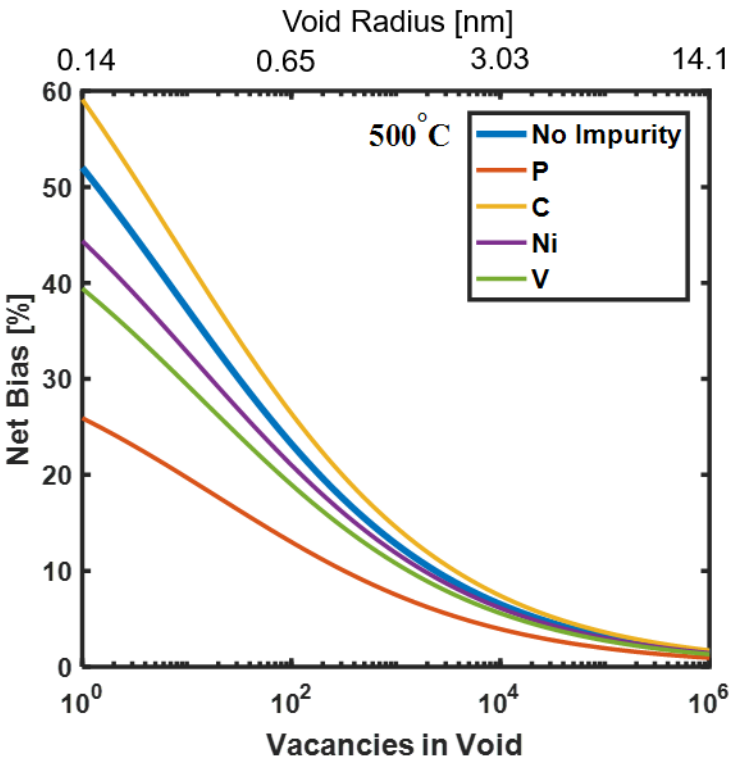

Figure 5.19: Net bias for interstitial capture as a function of void size with substitutional (P,C,V) and interstitial (C) impurities located on the (111) surface.

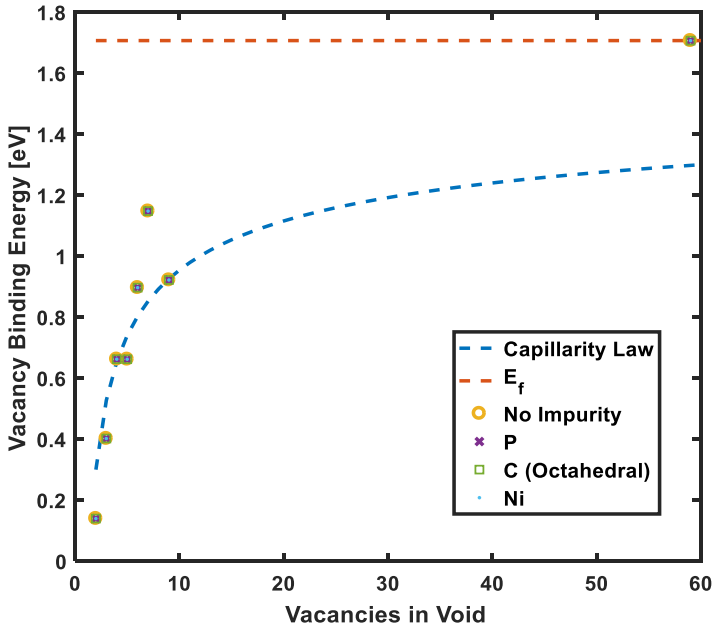

Figure 5.20: The binding energy of a single vacancy to void with various impurities on the (111) surface for void sizes ranging from 1 to 59 vacancy members.

5.4.2 Implications for Swelling

The statics work serves as encouraging evidence that these impurities indeed have the capability to locally change the bias, but in order to extend this to quantitative swelling information, a mechanism that predicts how likely they are to be situated at the appropriate location on the cavity surface. Unfortunately, these impurities are not necessarily expected to bind to the cavity, so treatments such as simply introducing them with a trapping-detrapping reaction may not be representative. Furthermore, the densities at which some of these impurities are present in ferritic alloys would cause the dilute limit assumptions implicit to rate constants to break down. Ideally, these would be treated as solute species in a radiation induced segregation model, but an accurate RIS model requires a fine spatial mesh oriented about a static sink. This is prohibitively difficult in the case of cavity nucleation, which is predicated on a rapidly evolving density of small embryonic cavities forming and evaporating. This section will aim to provide a reasonable facsimile of a full treatment which provides a preliminary map of what impurity-catalyzed swelling in the absence of helium could look like which is informed by experimental data.

Radiation-induced segregation has been extensively observed in FM steels at a variety of irradiation conditions [99,100]. Chromium, as the primary alloying element, has been most studied, with segregation behavior alternating between enrichment and depletion depending on irradiation temperature, dose rate, and alloy composition. Other minor solutes such as nickel and silicon have also been observed to enrich on grain boundaries [97]. Other researchers have done well to describe and model mechanisms controlling this behavior [101], and since they are too complex to directly model, they will not be discussed in great detail in this work. Regardless, the fundamental underpinning of segregation is a correlation between the migration of radiation-induced defects and various chemical species. As point defects eventually collect at sinks, the entanglement of these fluxes drives a redistribution of solutes at the sink-matrix interface. Grain boundaries are the archetypical site for segregation observations, since their nature as static, planar sinks makes them accessible to electron diffraction spectroscopy techniques. Since the underlying mechanism is not inherently dependent on sink geometry, the phenomenon may be generalized to other sinks like cavities and dislocations where characterization is more difficult. Indeed, there is some evidence of Cr segregation to dislocation loops in FM steels [102,103].

In lieu of an explicit segregation model, cavity bias was smoothly reduced in a prescribed manner shown schematically in Figure 5.21. The suppression follows a smooth cubic spline from the initial pure-iron cavity bias to an input ultimate suppressed value over the course of an input saturation dose. Since experimental characterization data on segregation to small cavities is unavailable, the values were selected to obtain swelling results reasonable for single ion irradiation experiments.

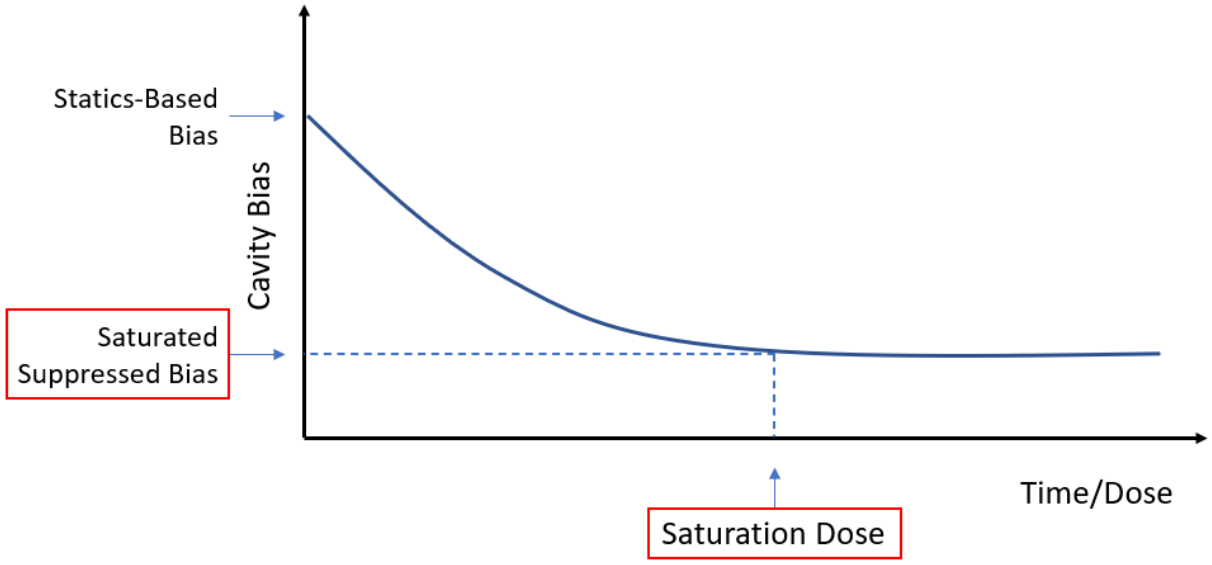

Figure 5.21: Schematic representation of the two-parameter prescribed bias suppression treatment.

The results of such a treatment, shown in Figure 5.22, demonstrate that this method is capable of quantitatively reproducing experimental swelling results without co-injected helium when properly parameterized. When the saturation dose for segregation is set sufficiently high, a significant incubation period is observed. Grain boundary segregation as observed experimentally typically saturates at much lower doses on the order of 10 dpa [97], but given the transient nature of nucleating cavities which are rapidly forming and evaporating in a manner which resets the segregation process, a much longer saturation dose seems rational. Once incubation gives way to growth, the value of ultimate suppressed bias controls cavity growth. When segregation suppresses cavity bias to 75% or more of its pure-iron value, it poses too strong a barrier and very little growth is observed. As it proceeds to lower values linear swelling rates increase rapidly to almost 1%/dpa. Suppression below around 50% of the initial pure-iron value yields less sensitivity

since the dislocation bias which drives growth has already dominated the defect kinetics of the system.

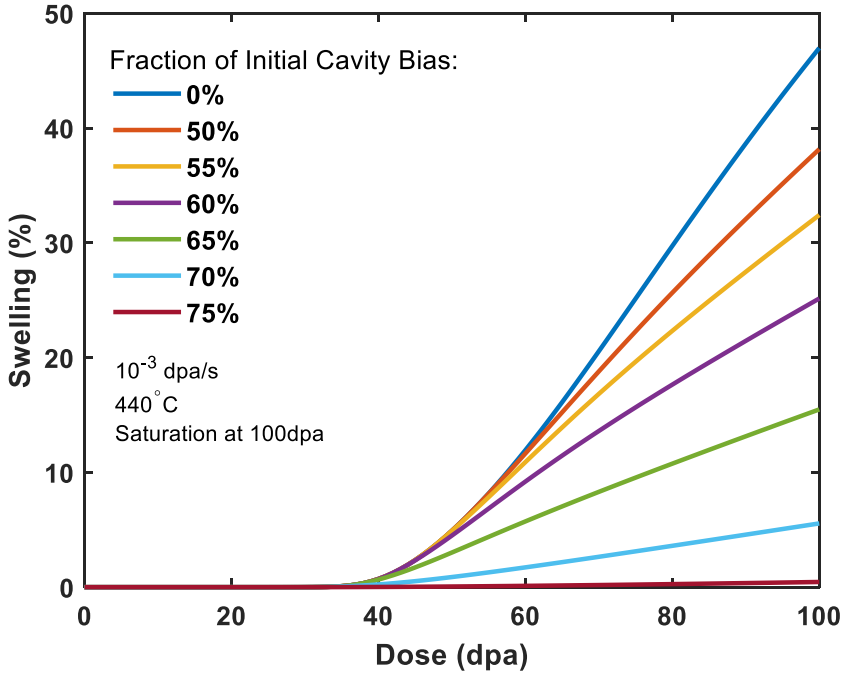

Figure 5.22: Swelling without co-injected helium as a function of dose using a prescribed bias suppression for various ultimate suppressed bias values

Temperature distributions generated by this model, presented in Figure 5.23, raise some concerns about this model in this simple form. While there is still a peak swelling temperature, it is significantly broader than any previously observed in this work. Swelling takes on a plateau-like distribution, extending to temperatures well below any experimental observations of swelling at high dose rate. It seems that since cavity bias, not thermal emission is acting as the primary barrier to nucleation, and this model reduces that bias uniformly at all temperatures, over-nucleation at low temperatures is averted. In reality, segregation is not uniform as a function of temperature. Two sets of ion-irradiated grain boundary segregation data for the FM alloy T91 are presented in Figure 5.24 [97,101,104]. Gaussian fits to these data which closely resemble the Modified Inverse Kirkendall modeling results presented by Wharry alongside their measurements [101], were used as an approximation of the RIS temperature dependence.

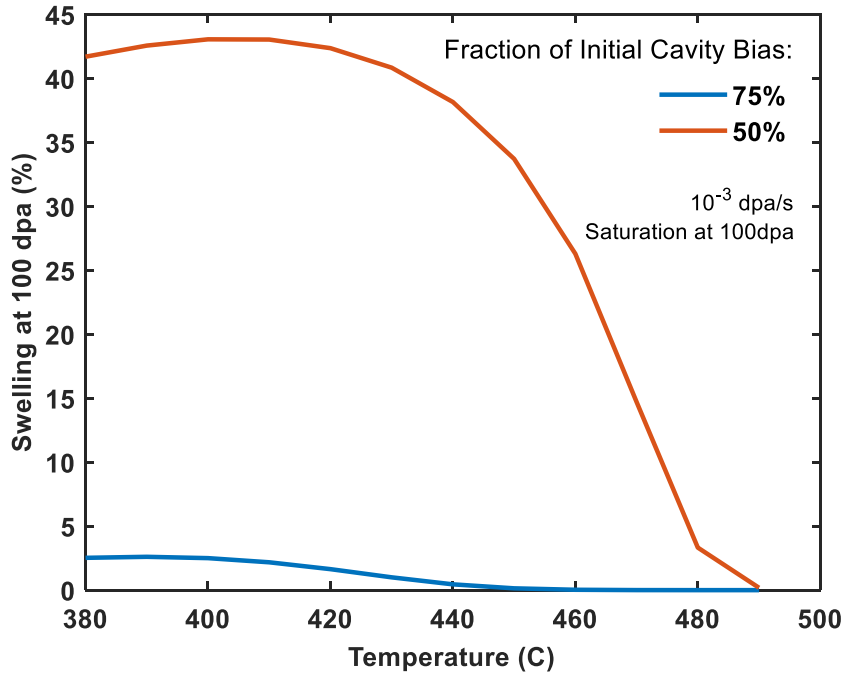

Figure 5.23: Temperature distributions of swelling without co-injected helium for cavity bias suppressed to 50% and 75% of its initial value over the course of 100 dpa

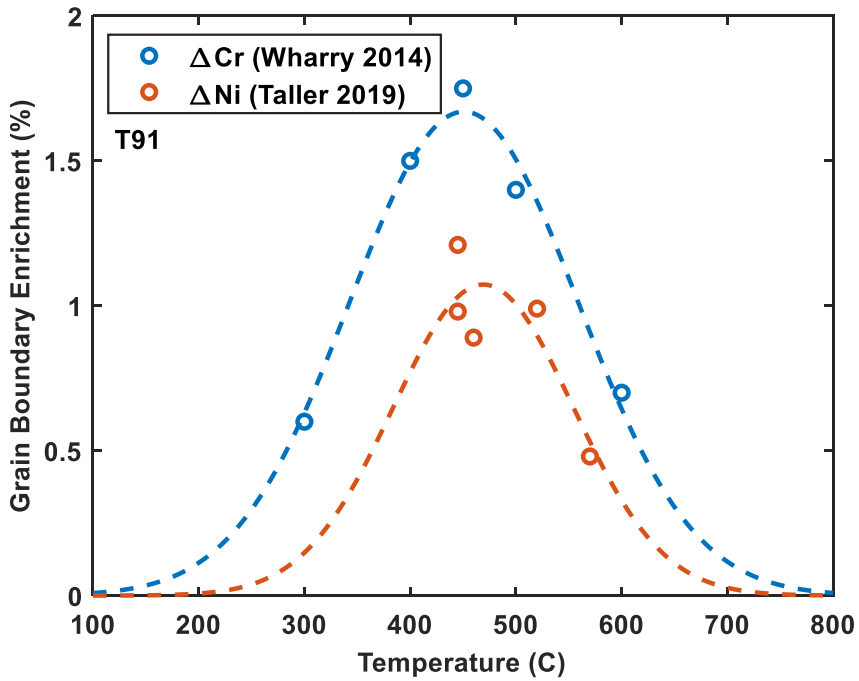

Figure 5.24: Temperature distributions of radiation-induced segregation of Cr and Ni to grain boundaries in T91 irradiated with protons at 10^{-5} dpa/s [101] and 5 MeV Fe at 10^{-4}-10^{-3} dpa/s [97] accompanied by gaussian fits

Accounting for the temperature distribution of segregation dramatically improves the result (Figure 5.25). When the ultimate suppressed bias is directly proportional to experimental RIS temperature distributions, the swelling temperature distributions return to a reasonable width. At high temperatures, swelling is still limited by thermal emission, but at low temperatures, swelling is instead suppressed by high cavity bias resulting from insufficient impurity concentrations on the cavity surface.

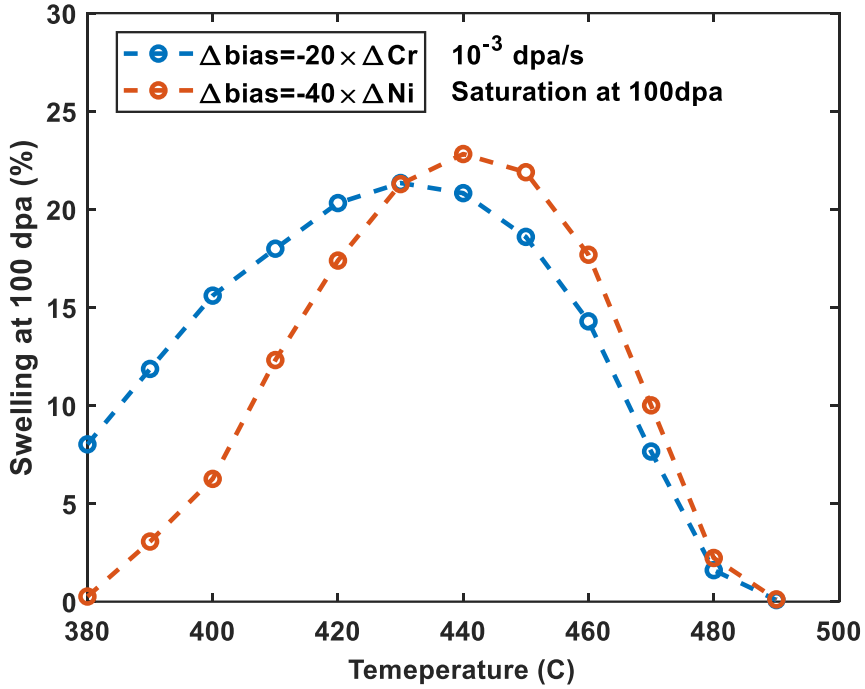

Figure 5.25: Temperature distributions of swelling without co-injected helium for cavity bias suppressed in proportion to gaussian fits to experimental GB segregation distributions over the course of 100 dpa

Accounting for surface microchemistry and its effects on cavity bias shows great promise for explaining observations of swelling without helium cogeneration. The statics calculations show quite clearly that cavity bias can be significantly suppressed with even a single undersized solute atom decorating the surface of a void. Furthermore, a cluster dynamics implementation of these results shows that incubation periods, reasonable swelling rates, and a consistent peak swelling temperature can be obtained without helium if cavity bias is suppressed. There are several underlying assumptions underpinning this work. Segregation is not explicitly modelled but informed indirectly by experimental observations at grain boundaries which are a fundamentally larger, more static sink. Segregation is presumed to follow a gaussian temperature distribution fit to grain boundary RIS in one FM alloy, T91. Furthermore, segregation is assumed to occur more

slowly than at grain boundaries due to the transient state of nucleating cavities. These cement the present status of this model as descriptive rather than predictive and highlight the need for more rigorous study of local microchemistry surrounding nucleating cavities.

Chapter 6: Conclusions

This work has introduced physical mechanisms necessary for the reproduction of several key aspects of cavity nucleation. A bias toward interstitials on cavities produced a barrier to cavity nucleation, which, unlike thermal emission alone could not be overcome by simply considering a lower temperature regime. The result was a system that required significant helium accumulation to achieve growth, and in cases where helium generation is slow, like low-activation FM steels, results in long incubation periods. An overly strong dependence on dose rate, reflected by an overestimation of temperature shift, was remedied through the addition of heterogeneous nucleation sites. Introduction of a site for static helium self-clustering which was itself independent of temperature and dose rate, served to disproportionately promote nucleation at low dose rate where thermal emission between damage events is more prevalent, reducing the dose rate dependence, and thus, temperature shift. Finally, a promising explanation for swelling that occurs absent helium was found in impurity-driven bias suppression. Backed by molecular statics calculations, impurities decorating the surface of a cavity could weaken its bias toward interstitials and facilitate growth. Individual contributions are further itemized below.

This work was focused on addressing 4 key discrepancies between conventional rate theory modeling of swelling with experimental observations:

- Dose rate effects were much stronger than observed, with test-reactor-relevant damage resulting in a peak swelling temperature around 300˚C, well below any experimental measurements, and overprediction of temperature shift between 110-140˚C.
- Linear swelling rates were far higher than observed unless an unreasonably small dislocation bias on the order of 1% for interstitials is used.
- Long incubations periods were effectively absent, with linear swelling being reached in only 1-2 dpa.

- Helium had almost no effect on cavity nucleation at generation rates of experimental relevance, unlike experiments where it is observed to greatly increase cavity density, shorten incubation periods, and create bi-modal size distributions.

Several mechanisms were incorporated, which did not fully reconcile these differences:

- Varying the many input parameters defining the kinetics of cavity nucleation and growth resulted in changes to swelling rates and peak swelling temperatures, but could not ultimately rectify the discrepancies, even using physically unreasonable values.
- Incorporation of vacancy clusters generated directly in cascades raised peak swelling temperatures, resulting in a reasonable high dose rate peak of 440˚C, but ultimately had little effect on temperature shift.
- Accounting for damage gradients and injected interstitials present in ion irradiations showed that in the appropriate analysis region, the results were equivalent to a continuum approach.
- Trapping interactions between vacancies and carbon impurities on the order of E_b=0.5eV were unable to significantly retard cavity nucleation and growth.

A cavity bias toward interstitials based on elastic interactions near the surface fundamentally changed the nucleation pathway, providing a previously missing barrier to nucleation. This results in three key experimental behaviors being qualitatively reproduced:

- Helium has a strong effect, promoting nucleation that is necessary for swelling to occur, but in abundance suppresses swelling via an over-nucleation of subcritical bubbles, resulting in a peak with respect to generation rate which is qualitatively similar to temperature.
- At low transmutation-relevant helium generation rates, the slow rate of helium accumulation creates an incubation period in which negligible swelling occurs, before eventually achieving a critical size where a transition to linear swelling begins.
- Cavity bias counteracts dislocation bias as a driving force for cavity growth, resulting in reasonable linear swelling rates even when net bias on dislocations is increased to a more reasonable 15-20%.

Cavity bias was unable to explain two key phenomena:

- Because the nucleation phenomenon as a function of temperature is still dependent on the same underlying mechanism, thermal emission, the dose rate dependence of these temperature distributions is unchanged.
- Helium is now absolutely necessary for swelling to occur, and without it vacancy clusters act only as centers for recombination.

A heterogeneous nucleation mechanism was introduced, which fundamentally changes the temperature dependence of swelling by partially decoupling nucleate density from damage rate.

- Providing sites for static helium self-clustering disproportionately enhanced cavity nucleation in low dose rate conditions, and when appropriately configured (a site binding energy of 1eV and a site density around $10^{20}m^{-3}$) it increased the reactor-relevant peak swelling temperature to 370˚C, and reduced the temperature shift to 70 ˚C.
- Based on the density and binding energy of these self-clustering sites, they are most consistent with dislocations.
- Assessing the helium generation rate dependence showed that agreement between dose rates may be improved with a 5x-10x increase in helium cogeneration rate for dual ion irradiations, however, this may no longer be valid at high dose levels.

The final discrepancy, a bi-product of cavity bias was addressed by accounting for the effects of chemical impurities at the cavity surface.

- Molecular statics calculations showed that undersized substitutional solutes oriented on the (111) surface of a cavity could reduce its bias toward interstitials by about a third.
- A CD implementation of these results shows that incubation periods, reasonable swelling rates, and a consistent peak swelling temperature can be obtained without helium if cavity bias is gradually suppressed to about 60% of its initial value over the course of 100 dpa.

Despite the significant improvements to experimental agreement, due to the intrinsic complexity of such a heterogeneous, non-equilibrium phenomenon, a few key discrepancies remain. The model still fails to quantitatively capture the cavity size distribution, with cavity microstructures dominated by small densities of large features. Microchemical arguments to justify swelling in the absence of helium are not rigorously derived directly from the underlying physics, existing as an interpretation rather than a prediction of experimental data.

Chapter 7: Future Work

Despite the clear improvement to agreement with experiment, several deficiencies remain. Finding effective ways to relax many of the assumptions implicit to a continuum model such as cluster dynamics is critical to rectifying them.

First, in the model as presented here, size distributions are consistently made up of small densities of very large cavities. This suggests that despite the significant improvements offered here and the insight they provide, there are still missing physics of cavity nucleation. This may be the result of a continuum model having fundamental inability to capture the intrinsic heterogeneity and spatial correlations of an evolving microstructure. Discovering ways to incorporate these without a treatment which is fully spatially resolved demands creative physical workarounds.

Second, models for microchemical effects on cavity evolution presented here are admittedly laden with assumptions. Finding self-consistent ways to model these interactions without introducing prohibitive computational complexity is an immense challenge. Each species explicitly incorporated requires many more equations to solve and interaction parameters which would require a parallel modeling effort on the atomic scale (if an appropriate interatomic potential exists). To the extent that these parameters exist in the modeling literature, some studies expanding this work could be conducted, for example varying the cluster mobility and stability parameters as a function of chromium content. More rigorous treatments directly considering more chemical species results in a "combinatorial explosion" of the phase space, exponentially increasing the size of an already very large problem. Focusing on one species at a time, with an emphasis on abundant species like chromium, and species with known interaction parameters with point defects like carbon presents a manageably incremental strategy.

Third, for similar reasons of computational demand, the present model essentially neglects the interstitial half of the defect microstructure. While these two halves have each provided valuable contributions in isolation, a complete account of the physics would have them fully coupled to one another. To incrementally proceed toward this ideal, one can propose a few preliminary treatments

of dislocation microstructural evolution. As a first approximation, the characterization of dislocation loops and networks from irradiation experiments could be directly used as sinks via an empirical fit. Refining this slightly, a very loosely interfaced pair of separate cavity and loop microstructure models could pass sink strengths between one another at some dose step. The assumption implicit to each of these techniques is that cavity and dislocation microstructure only interact as independent sinks in competition with one another. In reality, these clusters would be directly reacting with one another in many cases, particularly when dislocation loops are mobile, but neglecting these complicated reaction networks is what makes the problem manageable. The main benefit of a time-evolving dislocation sink strength is the gradual replacement of low-bias screw network dislocations with edge dislocation loops, resulting in a change in bias which could promote a short incubation period independent of helium generation rate.

As a final note, experiments and models are mutually more effective as tools of scientific discovery when they complement, corroborate, and inform one another. The main three mechanisms proposed in this work – cavity bias, heterogeneous nucleation, and segregation-driven bias suppression – are more valuable as descriptions of FM swelling behavior if experimenters can find independent evidence of them. This work aims to encourage electron microscopy in the FM system to look for cavities oriented co-linearly or exhibiting solute enrichment at their surface.

References

[1] C. Cawthorne, E.J. Fulton, Nature (1967).

[2] S.J. Zinkle, G.S. Was, Acta Mater. 61 (2013) 735–758.

[3] F. a. Garner, M.B. Toloczko, B.H. Sencer, J. Nucl. Mater. 276 (2000) 123–142.

[4] N. Packan, K. Farrell, J. Stiegler, J. Nucl. Mater. 78 (1978) 143–155.

[5] B.H. Sencer, J.R. Kennedy, J.I. Cole, S.A. Maloy, F.A. Garner, J. Nucl. Mater. 393 (2009) 235–241.

[6] D.S. Gelles, J. Nucl. Mater. 237 (1996) 293–298.

[7] D.S. Gelles, J. Nucl. Mater. 212–215 (1994) 714–719.

[8] E. a. Little, D. a. Stow, J. Nucl. Mater. 87 (1979) 25–39.

[9] M.B. Toloczko, F.A. Garner, C.R. Eiholzer, J. Nucl. Mater. 212–215 (1994) 604–607.

[10] J.J. Kai, R.L. Klueh, J. Nucl. Mater. 230 (1996) 116–123.

[11] J.M. Vitek, R.L. Klueh, J. Nucl. Mater. 122 (1984) 254–259.

[12] T. Morimura, A. Kimura, H. Matsui, J. Nucl. Mater. 239 (1996) 118–125.

[13] B.H. Sencer, J.R. Kennedy, J.I. Cole, S.A. Maloy, F.A. Garner, J. Nucl. Mater. 414 (2011) 237–242.

[14] J.J. Kai, G.L. Kulcinski, J. Nucl. Mater. 175 (1990) 227–236.

[15] P. a. Schmidt, P.R. Malmberg, J. a. Sprague, J.E. Westmoreland, Irradiat. Eff. Microstruct. Prop. Met. (1976) 227.

[16] Y.E. Kupriiyanova, V.V. Bryk, O.V. Borodin, A.S. Kalchenko, V.N. Voyevodin, G.D. Tolstolutskaya, F.A. Garner, J. Nucl. Mater. 468 (2016) 264–273.

[17] G. Ayrault, J. Nucl. Mater. 114 (1983) 34–40.

[18] E. Getto, Z. Jiao, a. M. Monterrosa, K. Sun, G.S. Was, J. Nucl. Mater. 462 (2015) 458–469.

[19] E. Getto, K. Sun, A.M. Monterrosa, Z. Jiao, M.J. Hackett, G.S. Was, J. Nucl. Mater. 480 (2016) 159–176.

[20] X. Wang, A.M. Monterrosa, F. Zhang, H. Huang, Q. Yan, Z. Jiao, G.S. Was, L. Wang, J. Nucl. Mater. 462 (2015) 119–125.

[21] X. Wang, Q. Yan, G.S. Was, L. Wang, Scr. Mater. 112 (2016) 9–14.

[22] M.P. Short, D.R. Gaston, M. Jin, L. Shao, F.A. Garner, J. Nucl. Mater. 471 (2016) 200–207.

[23] I.L. Singer, J. Vac. Sci. Technol. A Vacuum, Surfaces, Film. 1 (1983) 419.

[24] J.G. Gigax, E. Aydogan, T. Chen, D. Chen, L. Shao, Y. Wu, W.Y. Lo, Y. Yang, F.A. Garner, J. Nucl. Mater. 465 (2015) 343–348.

[25] S. Taller, D. Woodley, E. Getto, A.M. Monterrosa, Z. Jiao, O. Toader, F. Naab, T. Kubley, S. Dwaraknath, G.S. Was, Nucl. Instruments Methods Phys. Res. Sect. B Beam Interact. with Mater. Atoms 412 (2017) 1–10.

[26] E. Getto, G. Vancoevering, G.S. Was, J. Nucl. Mater. 484 (2016) submitted.

[27] D.M. Stewart, Y.N. Osetsky, R.E. Stoller, S.I. Golubov, T. Seletskaia, P.J. Kamenski, Philos. Mag. 90 (2010) 935–944.

[28] F. Gao, H.L. Heinisch, R.J. Kurtz, J. Nucl. Mater. 386–388 (2009) 390–394.

[29] H.L. Heinisch, F. Gao, R.J. Kurtz, E. a. Le, J. Nucl. Mater. 351 (2006) 141–148.

[30] H.L. Heinisch, F. Gao, R.J. Kurtz, J. Nucl. Mater. 367–370 (2007) 311–315.

[31] R.E. Stoller, G.R. Odette, J. Nucl. Mater. 131 (1985) 118–125.

[32] Y. Dai, G.R. Odette, T. Yamamoto, Comprehensive Nuclear Materials, Elsevier, 2012.

[33] S.J. Zinkle, Comprehensive Nuclear Materials, Elsevier, 2012.

[34] T. Muroga, H. Watanabe, N. Yoshida, H. Kurishita, M.L. Hamilton, J. Nucl. Mater. (1995).

[35] Q. Xu, T. Yoshiie, K. Sato, J. Nucl. Mater. (2009).

[36] S.J. Zinkle, K. Farrell, H. Kanazawa, J. Nucl. Mater. (1991).

[37] S.J. Zinkle, K. Farrell, J. Nucl. Mater. (1989).

[38] B. Singh, M. Eldrup, A. Möslang, in:, Eff. Radiat. Mater. Sixt. Int. Symp., 2009.

[39] B.N. Singh, A. Horsewell, J. Nucl. Mater. (1994).

[40] A.M. Monterrosa, Z. Jiao, G.S. Was, J. Nucl. Mater. 509 (2018) 707–721.

[41] C.J. Ortiz, M.J. Caturla, C.C. Fu, F. Willaime, Phys. Rev. B 80 (2009) 134109.

[42] C. Barouh, T. Schuler, C.C. Fu, M. Nastar, Phys. Rev. B - Condens. Matter Mater. Phys. 90 (2014) 1–20.

[43] F.A. Garner, A.S. Kumar, ASTM STP 955 (1987) 289–314.

[44] A.D. Brailsford, R. Bullough, M.R. Hayns, J. Nucl. Mater. 60 (1976) 246–256.

[45] L.K. Mansur, J. Nucl. Mater. 206 (1993) 306–323.

[46] L.K. Mansur, J. Nucl. Mater. 78 (1978) 156–160.

[47] R.E. Stoller, S.I. Golubov, C. Domain, C.S. Becquart, J. Nucl. Mater. 382 (2008) 77–90.

[48] A.A. Kohnert, B.D. Wirth, Model. Simul. Mater. Sci. Eng. 25 (2017).

[49] T. Jourdan, G. Bencteux, G. Adjanor, J. Nucl. Mater. 444 (2014) 298–313.

[50] S.I. Golubov, R.E. Stoller, S.J. Zinkle, a. M. Ovcharenko, J. Nucl. Mater. 361 (2007) 149–159.

[51] X. Hu, D. Xu, B.D. Wirth, J. Nucl. Mater. 442 (2013) S649–S654.

[52] D. Xu, B.D. Wirth, J. Nucl. Mater. 403 (2010) 184–190.

[53] a. V. Barashev, S.I. Golubov, J. Nucl. Mater. 389 (2009) 407–409.

[54] A.A. Kohnert, B.D. Wirth, J. Appl. Phys. 117 (2015) 154306.

[55] D. Xu, G. Vancoevering, B.D. Wirth, Comput. Mater. Sci. 114 (2016) 47–53.

[56] D. Xu, B.D. Wirth, M. Li, M. Kirk, Appl. Phys. Lett. 101 (2012) 101905.

[57] R.E. Stoller, G.R. Odette, B.D. Wirth, J. Nucl. Mater. 251 (1997) 49–60.

[58] Y.N. Osetsky, D.J. Bacon, a. Serra, B.N. Singh, S.I. Golubov, J. Nucl. Mater. 276 (2000) 65–77.

[59] M.P. Surh, W.G. Wolfer, J. Comput. Mater. Des. 14 (2007) 419–424.

[60] A.A. Kohnert, M.A. Cusentino, B.D. Wirth, J. Nucl. Mater. 499 (2018) 480–489.

[61] J.F. Ziegler, M.D. Ziegler, J.P. Biersack, Nucl. Instruments Methods Phys. Res. Sect. B Beam Interact. with Mater. Atoms 268 (2010) 1818–1823.

[62] R.E. Stoller, M.B. Toloczko, G.S. Was, A.G. Certain, S. Dwaraknath, F.A. Garner, Nucl. Instruments Methods Phys. Res. Sect. B Beam Interact. with Mater. Atoms 310 (2013) 75–80.

[63] R.E. Stoller, L.R. Greenwood, in:, ASTM Spec. Tech. Publ., 2000.

[64] D. Xu, B.D. Wirth, M. Li, M. Kirk, Acta Mater. 60 (2012) 4286–4302.

[65] A.D. Brailsford, R. Bullough, J. Nucl. Mater. 44 (1972) 121–135.

[66] R. Bullough, N.M. Ghoniem, J. Nucl. Mater. 127 (1985) 47–55.

[67] R. Bullough, M.R. Hayns, M.H. Wood, J. Nucl. Mater. 90 (1980) 44–59.

[68] R. Bullough, R.C. Newman, Reports Prog. Phys. 33 (2002) 101–148.

[69] H. Trinkaus, V. Naundorf, B.N. Singh, C.H. Woo, J. Nucl. Mater. (1994).

[70] C.C. Fu, J.D. Torre, F. Willaime, J.L. Bocquet, A. Barbu, Nat. Mater. 4 (2004) 68–74.

[71] D. Terentyev, P. Olsson, T.P.C. Klaver, L. Malerba, Comput. Mater. Sci. 43 (2008) 1183–1192.

[72] D. Nguyen-Manh, M.Y. Lavrentiev, S.L. Dudarev, Comptes Rendus Phys. (2008).

[73] B.D. Wirth, G.R. Odette, D. Maroudas, G.E. Lucas, J. Nucl. Mater. 244 (1997) 185–194.

[74] D. Terentyev, N. Juslin, K. Nordlund, N. Sandberg, J. Appl. Phys. 105 (2009) 103509.

[75] E. del Rio, J.M. Sampedro, H. Dogo, M.J. Caturla, M. Caro, A. Caro, J. Perlado, J. Nucl. Mater. 408 (2011) 18–24.

[76] A.A. Kohnert, B.D. Wirth, J. Appl. Phys. 117 (2015) 154305.

[77] A.Y. Dunn, L. Capolungo, E. Martinez, M. Cherkaoui, J. Nucl. Mater. 443 (2013) 128–139.

[78] A. Dunn, B. Muntifering, R. Dingreville, K. Hattar, L. Capolungo, J. Nucl. Mater. 480 (2016) 129–137.

[79] S.I. Golubov, A.M. Ovcharenko, a. V. Barashev, B.N. Singh, Philos. Mag. 81 (2001) 643–658.

[80] Z. Chang, D. Terentyev, N. Sandberg, K. Samuelsson, P. Olsson, J. Nucl. Mater. 461 (2015) 221–229.

[81] D. Woodley, The Role of Helium on Cavity Growth and Swelling at High Damage Levels in Ferritic-Martensitic Steels, 2020.

[82] S. Taller, G.S. Was, Acta Mater. 198 (2020) 47–60.

[83] F. Gao, H.Q. Deng, H.L. Heinisch, R.J. Kurtz, J. Nucl. Mater. 418 (2011) 115–120.

[84] H. Lefaix-Jeuland, S. Moll, T. Jourdan, F. Legendre, J. Nucl. Mater. 434 (2013) 152–157.

[85] R. Sugano, K. Morishita, H. Iwakiri, N. Yoshida, J. Nucl. Mater. 307–311 (2002) 941–945.

[86] I.I. Chernov, B.A. Kalin, M.S. Staltsov, K.Z. Oo, S.Y. Binyukova, O.S. Staltsova, A.A. Polyansky, V.S. Ageev, A.A. Nikitina, J. Nucl. Mater. 459 (2015) 259–264.

[87] X.Z. Cao, Q. Xu, K. Sato, T. Yoshiie, J. Nucl. Mater. 417 (2011) 1034–1037.

[88] H.L. Heinisch, F. Gao, R.J. Kurtz, J. ASTM Int. 4 (2007) 190–196.

[89] K. Tapasa, a. V. Barashev, D.J. Bacon, Y.N. Osetsky, Acta Mater. 55 (2007) 1–11.

[90] S. Plimpton, J. Comput. Phys. (1995).

[91] K. Schroeder, K. Dettmann, Zeitschrift Für Phys. B Condens. Matter Quanta 22 (1975) 343–350.

[92] G.J. Ackland, M.I. Mendelev, D. Srolovitz, S. Han, A.V. Barashev, J. Phys. Condens. Matter 16 (2004) S2629–S2642.

[93] D.J. Hepburn, G.J. Ackland, Phys. Rev. B - Condens. Matter Mater. Phys. 78 (2008) 1–6.

[94] G. Bonny, R.C. Pasianot, L. Malerba, Model. Simul. Mater. Sci. Eng. (2009).

[95] M.I. Mendelev, S. Han, W.J. Son, G.J. Ackland, D.J. Srolovitz, Phys. Rev. B - Condens. Matter Mater. Phys. 76 (2007) 1–11.

[96] M.P. Surh, J.B. Sturgeon, W.G. Wolfer, J. Nucl. Mater. 336 (2005) 217–224.

[97] S. Taller, Z. Jiao, K. Field, G.S. Was, J. Nucl. Mater. 527 (2019) 151831.

[98] R.J. Kurtz, H.L. Heinisch, F. Gao, J. Nucl. Mater. 382 (2008) 134–142.

[99] Z. Lu, R.G. Faulkner, G. Was, B.D. Wirth, Scr. Mater. 58 (2008) 878–881.

[100] Z. Jiao, G.S. Was, Acta Mater. 59 (2011) 4467–4481.

[101] J.P. Wharry, G.S. Was, Acta Mater. 65 (2014) 42–55.

[102] A. Bhattacharya, E. Meslin, J. Henry, A. Barbu, S. Poissonnet, B. Décamps, Acta Mater. 108 (2016) 241–251.

[103] C. Heintze, F. Bergner, A. Ulbricht, M. Hernández-Mayoral, U. Keiderling, R. Lindau, T. Weissgärber, J. Nucl. Mater. 416 (2011) 35–39.

[104] Z. Jiao, S. Taller, K. Field, G. Yeli, M.P. Moody, G.S. Was, J. Nucl. Mater. 504 (2018) 122–134.