5-2017

Numerical Simulation of Metallic Uranium Sintering

Bruce Berry
University of Arkansas, Fayetteville

Numerical Simulation of Metallic Uranium Sintering

A thesis submitted in partial fulfillment
of the requirements for the degree of
Master of Science in Mechanical Engineering

by

Bruce Berry
Georgia Institute of Technology
Bachelor of Science in Nuclear Engineering, 2013

May 2017
University of Arkansas

This thesis is approved for recommendation to the Graduate Council.

Dr. Paul Millett
Thesis Director

Dr. Rick Couvillion
Committee member

Dr. Arun Nair
Committee member

ABSTRACT

Conventional ceramic oxide nuclear fuels are limited in their thermal and life-cycle properties. The desire to operate at higher burnups as is required by current utility economics has proven a formidable challenge for oxide fuel designs. Metallic formulations have superior thermal performance but are plagued by volumetric swelling due to fission gas buildup. In this study, we consider a number of specific microstructure configurations that have been experimentally shown to exhibit considerable resistance to porosity loss. Specifically, a void sizing that is bimodally distributed was shown to resist early pore loss and could provide collection sites for fission gas buildup. We employ the phase field model of Cahn and Hilliard, solved via the finite element method using the open source Multi-User Object Oriented Simulation Environment (MOOSE) developed by INL.

ACKNOWLEDGEMENTS

This work would not have been possible without the help of my advisor, Paul Millett. In addition, the frequent discussions with my lab coworkers proved to be invaluable. Support from family and friends is also greatly appreciated.

TABLE OF CONTENTS

LIST OF FIGURES

NOMENCLATURE

k -	Thermal conductivity
TD -	Percent theoretical density
ϕ -	Phase field order parameter
M -	Material mobility constant
κ -	Interface constant
f_{loc} -	Local free energy of mixing
E_d -	Lumped forces term
F -	Free energy functional
Ω -	(Mathematical) hypothetical bounding volume
S -	(Mathematical) hypothetical bounding surface
χ -	Ratio of small to total pore space
D -	Diameter
O -	(Complexity) worst case execution cost
Ω -	(Complexity) best case execution cost
Φ -	(Complexity) tightly bounded execution cost

1 Introduction

1.1 Background

As humanity settles into the Information Age, we strive to apply new technologies to problems that were previously thought to be out of reach. These new challenges continually motivate us to reconsider how best to manage the resources available to us, including energy. We require more of our sources of energy, be it in the form of superior economics or through more sustainable operation. As our demands increase, infrastructure must be modernized to reflect these 21^{st} century needs. Improved nuclear power is expected to play an important role in tomorrow's energy mixture, and significant research and development is currently underway to develop systems that can satisfy future demands. The current research described here focuses on an improved fuel formulation intended for use in next generation reactors. The successful application of this fuel promises next generation reactors with greater safety, profitability, and sustainability, satisfying many of the projected requirements of a next-generation power source. Here we will give a brief overview of the concepts and systems being investigated in the current work and explain the technical motivations of the particular research area.

Nuclear power plants (NPPs) generate power via the splitting of fuel atoms in a self-sustaining chain reaction. The splitting generates significant heat within the fuel which is then removed and carried to steam generators via a coolant, such as water. The water is transformed into steam, and energy is extracted from the steam to produce electricity, in much the same way as coal and gas fired power plants generate power. The splitting of the

fuel atoms results in the production of lighter elements, many of which are radioactive over various time scales. Some of these elements are in the gaseous form, a fact that will be of some importance for this work.

Although very similar in overall operation to fossil fuel plants, NPPs must be additionally secured against the possibility of coolant loss to the fuel, as even when shut down the fuel produces significant residual heat (as a result of radioactive decay) that must be continually removed to avoid fuel melting and plant damage. Such an accident has been observed in operation in previous events - the Fukushima Disaster is a leading example - and significant effort is put into designing systems and operational procedures that minimize the risk of this occurring and mitigate damage should such an event occur. The leading designs for next-generation reactors prioritize passively safe behavior in accident scenarios, which currently employed designs lack in the most extreme situations (such as Fukushima).

Even considering this drawback, NPPs posses several advantages over fossil fuel power plants. They are among the most environmentally friendly power sources available, releasing far lower amounts of radiation than fossil fuel plants, and are rivaled only by hydroelectric power plants in their per capita carbon emissions. They have low sensitivity to fuel prices in comparison to fossil plants, and enjoy very high uptimes, resulting in a very predictable and reliable baseload power source.

Traditional reactors are powered almost exclusively by uranium oxide, UO_2. This formulation is classified as a ceramic material. It has good chemical compatibility with common structural and cooling materials, and like most ceramics possess a relatively high melting point. Not without drawbacks, oxide fuels have a fairly low thermal conductivity, and also exhibit conductivity reduction with increasing temperature for certain relevant temperature

ranges. This behavior imposes operational restrictions that reduce plant efficiency in order to avoid unsafe behavior should unexpected events occur. In addition, plant cycle times between refuellings are somewhat short, owing to the fact that oxide fuels have a rather low heavy metal density, the ratio of uranium (fuel) to other compounds (oxygen) in the fuel material. In somewhat simple terms, this limits the "potency" of the fuel after extended operation, and requires frequent refuellings to maintain a nuclear reaction.

Metallic uranium formulations exist that alleviate many of the problems of the oxide fuel forms. U-Zr alloys exhibit superior thermal conductivity at all temperature ranges, significantly relaxing operational restrictions in place to avoid fuel melting. While their melting point is lower than oxide formulations, their ability to remove heat much more efficiently allows for larger thermal safety margins while maintaining performance. This property lends a certain amount of intrinsic accident-tolerance to the fuel. In traditional reactors, the primary mechanism employed to avoid damage during an accident scenario is by restricting operation to ranges that fail gracefully, that is, by adopting operating procedures that make it difficult for the system to depart from nominal behavior. This technique has been widely successful, and contributes to nuclear power's impressive safety record. It is difficult for this method to secure against events that are highly unlikely, and more destructive than otherwise planned for. Accident tolerant fuels behave acceptably even during events beyond the traditionally designed-for range. Thermal behavior, while only one aspect of accident tolerance in fuels, is an important attribute, and metallic fuels perform admirably in this regard.

In addition to improved thermal performance and safety, metallic fuels posses high initial heavy metal density. This enables operation for longer duration than traditional oxide

fuels, and better total fuel consumption. This value, termed "burnup", is a measure of the utilization of the fissile material in the fuel. The burnup that can be attained is influenced by numerous factors, including the compositional and morphological changes that occur in the fuel as it is consumed. With longer operation between cycles, power plants tend to be more profitable, as refuelling periods are typically considered to be extremely expensive.

1.2 Motivation

The complexities of tomorrow's economy require advanced energy sources that promote growth. It is natural that as society progresses to tackle increasingly difficult problems, our infrastructure must evolve to support these efforts. Fossil fuels have proven invaluable in that they have enabled the cheap fueling of the industrial revolution, which has led to today's technological state. Not without their drawbacks, their continued use will become more expensive, as their extraction becomes more difficult. Additionally, externalities of their consumption will become more evident, both in their local and global effects on the environment.

It is clear that new sources of energy must power tomorrow's society. Next-generation nuclear power is expected to play a part in fulfilling this need. Among the goals in designing the next generation of NPPs, two focus areas emerge: improved safety and enhanced economics. We have already discussed some of the considerations when selecting a fuel for its safety performance. Any system in widespread deployment must be shown to be cost-competitive with alternatives. For nuclear power, upfront capital expenditure is a significant hurdle to adoption. Various solutions to this problem have been proposed, including the design of plants with significantly longer operational lives, enabling the recovery of capital

over longer periods. In addition to capital costs, utility companies are concerned with the frequency at which reactors must be shutdown for refueling and maintenance, as shutdowns are expensive. It is this area in which new fuel formulations become very attractive, as almost all current efforts seek to extend the duration for which fuel can be utilized.

Metallic fuel formulations, while beneficial from a safety standpoint, have behaved poorly during operation in that they experience rapid microstructure changes that inhibit their long term operation (where the microstructure is defined as structural details that are resolvable at the μm size). Specifically, they experience rapid swelling with high temperature operation that prematurely stresses the protective cladding surrounding the fuel pellet. This swelling occurs as a result of pore space elimination within the fuel, followed by fission gas creation as the fuel is consumed. The gases, which are initially contained within the fuel pellet, eventually migrate to the outer region of the fuel, and exert excessive pressure on the cladding. This requires early shutdown and prevents operation to high burnup. It is this particular problem that this research seeks to address, in particular, via microstructure configurations that resist pore loss, and provide space for the capture of fission gases. This modification is an important step in the design of metallic fuels that are suitable for use in NPPs.

1.3 Objectives

The first objective of this study is to develop an understanding of the changes that occur to metallic fuel microstructures as they are thermally annealed, that is, under conditions comparable to those experienced in power reactors. We examine the mechanisms by which the fuel densifies as void space within the structure is eliminated. Certain microstructure

configurations have been shown to behave differently in this regard, and we seek to understand the reasons behind this difference. We find direction for this study from experimental results showing a bimodally-distributed pore size configuration as exhibiting improved resistance to pore loss, and seek to develop computational models that duplicate this result. By developing an understanding into how metallic fuels respond to sintering conditions, it is hoped that future work will be able to design fuels that improve reactor economics and safely.

The operational environment that exists in an active fuel element is complex, and involves rapidly changing temperature, radiation, and stress components. In this work, we develop a model focusing on the behavior in response to thermal effects, but which is capable of extension to include other effects as well. The goal of such a model is to allow follow-on studies to add physical effects incrementally, at each step focusing on the fundamental mechanisms and their affects on the behavior of the fuel system.

1.4 Thesis Organization

This thesis is divided into five chapters. In the introduction, the necessary background knowledge for understanding the scope of this work has been given. The motivations are also presented. Chapter two, the literature review, summarizes the search of existing knowledge that was performed. This formed a starting point in developing a suitable model for the problem, and helped narrow the scope of the work to a particular field of inquiry that had not yet been pursued. In Chapter three, the methods used in developing the model of the problem are described in detail. The results of the work are presented in Chapter four, along with a discussion of them. Conclusions are drawn in Chapter five, in addition to

consideration of limitations and areas of future work.

2 Literature Review

2.1 Introduction

The literature review is primarily directed to three areas; a summary of the history and properties of fuel formulations (both metallic and oxide), the physical mechanisms relevant to fabrication and initial operation (materials science and engineering), and the mathematical models that exist to model these systems via computer simulation (phase field methods, finite element techniques).

2.2 History and Motivation

The first power reactors, and almost all since, have employed oxide formulations of uranium for their fuel source [1]. The ceramic oxides have shown excellent chemical compatibility and good radiological performance in Generation II and III reactors [2]. A fuel intended to replace UO_2 must be shown to be superior to the type currently used, making an understanding of UO_2's properties and performance important. [2] and [3] proved valuable in extracting figures of merit to consider in evaluating a fuel design. In particular, these resources were used to estimate UO_2's thermal conductivity, k, and how this parameter varied over a range of temperatures. It is understood that a significant drawback of these formulations is the low value of k and its tendency to be depressed at high temperatures. This typically results in operating restrictions to avoid unsafe thermally-induced feedback loops. Another figure of importance is in comparing the initial heavy metal density, also provided by these references. We will see that most oxide formulations

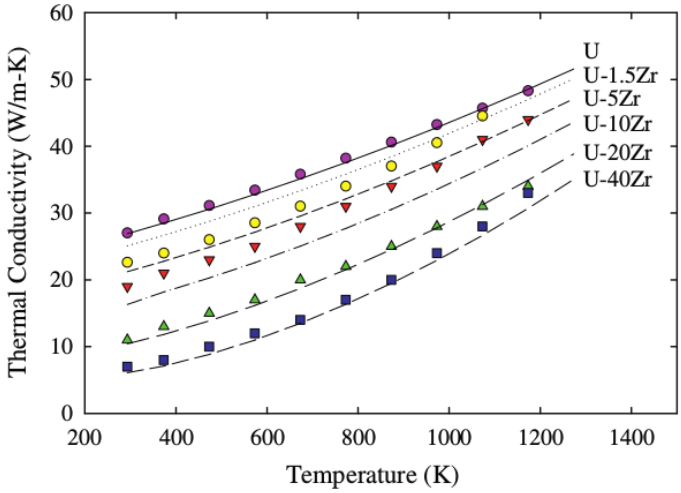

(a) U-rich U-Zr alloys

Figure 2.1: Thermal conductivity as a function of temperature for some experimental U-Zr alloys. Improvement with increasing temperature is observed, which is desirable in some accident scenarios. Image source [4]

suffer from a relatively low quantity of usable uranium in comparison to the metallic forms. While this study is motivated primary in developing an understanding of metallic fuels, the discoveries made may also be of use to the oxide fuel community. In the selection of simulation parameters, as will be discussed later, it is thus important to consider values that provide some benefit to both fuel types, making an understanding of UO_2's behavior important. In particular, [2] provides a typical value of $k = 5$ $[W/(m \cdot K)]$ for UO_2 fuels utilized in power reactors.

U-Zr alloys possess properties that make them thermally superior to UO_2 forms, and are chemically compatible with the less common, but much researched, sodium-based coolant systems [5, 6]. With superior conductivity, metal-fueled cores can be made more compact for a given thermal power output, possibly with a corresponding decrease in upfront capital investment. We see also from [4] that an improvement in conductivity occurs with an

increase in temperature. This is beneficial in unplanned overpower or transient conditions, where fuel melting becomes a concern. We also observe the metallic fuels as having superior initial heavy metal density to oxide fuels.

In [6], we find a detailed description of the testing of metallic fuel rods in research reactors, including various failure modes in response to irradiation. Specifically, we find evidence of the early swelling problem often resulting in cladding failure, even at moderate burnup. A solution to this problem is also described: an oversized fuel pin plenum and increased cladding and gap thickness survive fuel swelling until continuous vapor networks formed within the fuel pellet emerge, providing a venting path for the gases into the plenum and limiting further expansion. While this technique of over-strengthening the cladding and then venting all products into the plenum enabled high burnup operation (as high as 20 at.% for some cases), neutronic and thermal hydraulic considerations make this configuration undesirable, although an explanation of such is beyond the scope of this work. For this study, this undesirable swelling is the core issue we seek to remedy, or to at least understand the mechanism by which swelling occurs. Other considerations, such as phase segregation of the zirconium [6], fuel-cladding chemical interaction [5], and other fuel failure modes [6] are not considered herein.

2.3 Experiments

Techniques for the fabrication of metallic fuel forms are described in [7]. We see that traditional casting techniques have been used in the past, but that modern powder metallurgy methods are being developed that enable new capability in terms of microstructure specification for a given application. The formation of depleted uranium powders is

described, with the ability to specify average grain size. Of most interest is the use of an electric arc-discharge through a thin uranium-alloy wire. The wire is melted at the tip, and the molten component ejected into an inert chamber and subsequently solidified into small spheres, forming a fine powder. The distribution of sphere sizes can be controlled. The sintering process of metals is described in [8, 9, 10, 11]. A loose collection of particles is contact-pressure formed into a "green" part that is close to the final geometry of the desired part, but is not monolithic. Under various conditions (a combination of high temperature and pressure), this green part is then "sintered" to induce grain growth and coalescence, until the individual particles within the green part fuse. This process is analogous to the way pottery is produced in clay-based materials. Typically, the sintering temperature is much lower than the melting point of the material. While this is not of vital importance for uranium alloys, it proves essential for elements with high melting points, such as tungsten.

The initial porosity used for beginning our simulations was found to be an important parameter in the study. From the literature [2], this value (TD) is typically in the range 92-96% for UO_2 fuels. There were no articles found that specified this value for the various cast U-Zr alloys (although cast parts typically have minimal porosity), but the hot isostatic pressing process is know to produce metal parts with TD values >90% [12]. Indeed, with the various powder metallurgy techniques, a wide range of starting porosity values are feasible, with iron foams at TD = 20% having been produced [13]. Additionally, work in [14] demonstrated the ability to control both the average grain and pore size, as well as producing a distribution of pore sizes. The desire to produce systems with pores of various sizes adds additional requirements to the fabrication routes. As spheres are closely packed in the pressing step, the arrangement is such that small pores exist between these hard spheres

naturally. "Pore forming" compounds are added to the metal powder before cold pressing into the green form. The pore formers are typically much larger than the bulk material spheres. In a pre-sintering step, these pore formers are eliminated and escape the structure, leaving large voids within the green compact. Later sintering preserves these voids, and by careful pore former selection, it is possible to control the porosity distribution of the sintered piece. This technique provides an avenue for the production of the microstructures investigated in this work.

The work done in [15] proved to be of central importance in providing direction and parameters for the present work. In the article, the author details experiments in which UO_2 samples with different pore structures were sintered and their porosity loss compared. The discovery that the configuration with pores of two primary size ranges produced superior resistance to the otherwise rapid loss of pore space motivated a desire to understand the mechanism via simulation in the present work. As duplication is a valued goal in new domains, it was important to select parameters in this work that produced configurations comparable to those tested in [15]. We adopt similar ratios of large and small void size, as well as relative quantity of each. Additional, similar initial TD values are tested. The reference is similar to the present work in that only thermal effects were considered, not the effects of radiation. Here, we simulate extended annealing as a result of high temperatures exclusively.

2.4 Simulation

As computer systems have grown to enable the solution of larger and more complex problems, methods have been simultaneously developed to model physical systems in ways

that are amicable to solution via simulation. This trend has been maintained in the materials science domain, where computer simulation has been used to predict materials before they have been physically realized in the laboratory [16]. In order to represent the behavior of materials at the level of their microstructure, that is, at sizes that are not observable without magnification, new mathematical models have been developed that attempt to capture the different physical forces and interactions that are dominant.

The Monte-Carlo based Potts model of [17] was applied to the sintering process to address multi-grain growth in solid bodies in [18, 19]. The treatment of vacancies is probabilistic; a vacancy site is randomly relocated and the energy change measured. If the change is favorable, the relocation is kept and the system updated. Unfavorable changes are rejected, but a penalty term is applied that makes a future event more likely. The Potts model with such modification has been shown to provide a good overall description of the sintering process, but its handling of vacancy motion is problematic for this work, in which we expect vacancy motion to play a key role in controlling the densification kinetics. We are therefore motivated to select a model that effects vacancy motion in a more realistic manner.

The phase field family of models is used in this work, and will be described in detail in the methods section (Chapter 3). Their development in the literature is described here. The work of [20, 21, 22] introduced the method as applicable to a range of materials problems, and was applied to a binary alloy system to demonstrate the phase separation of two solids. Further development in [23, 24, 25] showed their application to the sintering process, and later continued annealing. The treatment of conserved and non-conserved entities, such as phase, grain orientation, and others through the use of a continuously varying "order parameter" lends well to discretization on a computational grid, a process which is commonly

performed in many fields. The ability to resolve phase interfaces is particularly advantageous, and various physical effects can be introduced by modification of the free energy functional with new parameters.

Several numerical techniques exist to solve coupled nonlinear systems such as those developed from the phase field methods. The finite difference approach is the most simplistic, but limited in handling complex geometries, and will not be considered further. Spectral techniques in development provide very fast and adaptable solution paths, but at this point in time require in-house development of the software, adding additional difficulties to the validation of simulation results. The finite element method, while relatively slow in the solution when compared to the other techniques, is extremely flexible with regards to problem geometry and physical time and length scales. The MOOSE project [26, 27] adds to this flexibility by providing a large codebase that combines high speed meshing and solution tools with physics packages, including modules for the phase field methods. This codebase allows the present work to sidestep most of the numerical implementation details, and instead focus on model development and interpretation of results. In addition, much of the MOOSE codebase is well validated, with many examples of its successful application to large, complex systems [28, 29]. MOOSE's organization and workings will be further described in the methods section (Chapter 3).

2.5 Densification Studies

Sintering as a process has been well studied, and accurate models exist that predict microstructure and material properties after a given sintering schedule has been performed. The further evolution of these pre-sintered systems in reactor environments is less well un-

derstood, owing to the difficulties created by handling highly radioactive fuel material. Experience with this process is necessary, however, in discerning the morphological changes that result in rapid densification and subsequent swelling of U-Zr alloy fuel formulations. In this regard, we refer to the process as re-sintering or densification studies (also, extended annealing), owing to the fact that the microstructure's starting point is not that of many spheres which are tightly packed, but a continuous structure with pores of various sizes existing within. Of primary interest becomes the phenomena of vacancy transport within the system.

3 Methods

3.1 Introduction

The main focus of this thesis is the execution of simulations that seek to represent the physics of extended annealing of solid bodies after they have been formed via powder metallurgy. We examine the changes that occur to the material's microstructure as thermal effects drive the evolution of the system. To this end, we need two tools to accomplish this goal; a computational model that represents the physics of interest, and a solution technique to transform our problem to one that is solvable on the computer. The phase field family of equations provides us this model, while the finite element method via MOOSE enables the solution via computer. Both of these tools will be described in this section, as well as the additional post-processing procedures utilized.

3.2 Modelling and Solution of Materials Problems

3.2.1 Phase Field Methods

We employ the phase field methods [22], originally introduced by Cahn and Hilliard. These methods are popular in solidification and phase separation problems. A common task of this model is in representing the phases present at a given point in the system (solid, gas, crystal structure, etc). Representation of this phase is done via a so called "order parameter" that is continuous throughout the system. The value of this order parameter (or order parameters) determines what phase is present at a given location. The order parameter takes on a singular value at a given point. Its variation in a two phase system, such as shown

in Figure 3.1, enables easy representation of both bulk and interface regions. Within the bulk of a phase, the order parameter, ϕ, takes on a value of +1 or -1, but interfaces assume intermediate values. The interface width is related to the rate of change of ϕ as it transitions to indicate a phase change.

In Figure 3.2, we show a description of the underlying lattice assumed by the Cahn-Hilliard model, which according to the evolution equation that we will shortly introduce, treats the discrete nature of vacancy transport as a continuous phenomena. Through this image, we can imagine a pore as a very large collection of uninterrupted vacancies in the lattice, contained within the larger fuel system. Individual vacancies can be released from this large collection and freely propagate throughout the lattice. Rather that discretely tracking the location of these vacancies at each lattice point, we use a single parameter whose value varies in accordance with the species (vacancy or atom) present at each point.

$$\frac{\partial \phi}{\partial t} = \nabla \cdot M \nabla \frac{\delta F}{\delta \phi} \tag{3.1}$$

$$F(\phi) = \int_V [f_{loc}(\phi) + \kappa |\nabla \phi|^2 + E_d] dV \tag{3.2}$$

Equation 3.1 is the Cahn-Hilliard equation, cast in the form suitable for the solution of binary systems. ϕ is the order parameter representing the phase at a given location. In our case, we take $\phi = 1$ as representing fuel material and $\phi = -1$ as denoting the existence of a void space. F is the free energy functional, determining the driving forces in the system's evolution. M is a proportionality value known as the mobility with units $[L^2][T^{-1}]$. To constrain ϕ to the range [-1,1], we cast F to include a double well potential, taking the

form given in Equation 3.2. This ensures that the two phases (fuel and void) will evolve to separate when given the proper driving force, which in this case is the thermal energy provided by the annealing process. Here, f_{loc} is the local energy of mixing (the double well), the κ term the interface energy, and E_d a lumped parameter including any additional energy sources (electric, stress, etc). We require that the annealing process be isothermal, enabling a uniform value for the mobility. Furthermore, we do not consider thermal profiles that would enable a non-homogeneous mixture of the two constituents (uranium and zirconium), nor do we allow precipitation of alternative phases. With this, we can declare the mobility as constant.

In its fully general form, F is given by Equation 3.2. Of interest to us are the energies of mixing and the interfacial energy. Thus, we only consider the f_{loc} term along with the interface term to arrive at the specialized form used here, Equation 3.3. We also specialize the local free energy term to encourage phase separation, as provided by the double well.

$$F(\phi) = \int_V (\phi^4/4 - \phi^2/2 + \kappa|\nabla\phi|^2)dV, \tag{3.3}$$

Often, we model mixing and subsequent phase separation of two immiscible phases. The problem of vacancy transport in a one phase system can be imagined as such. Figure 3.4 shows the evolution of a two phase system that is immiscible, such as oil and water. The two phases are initially intimately mixed, indicated by the uniform light blue color in the first image. As the system evolves in time, small pure regions of each phase are formed (or nucleate), and then grow as nearby material is transported into these nucleated regions. These separate regions can interact with other pure regions to form larger pure bubbles, or

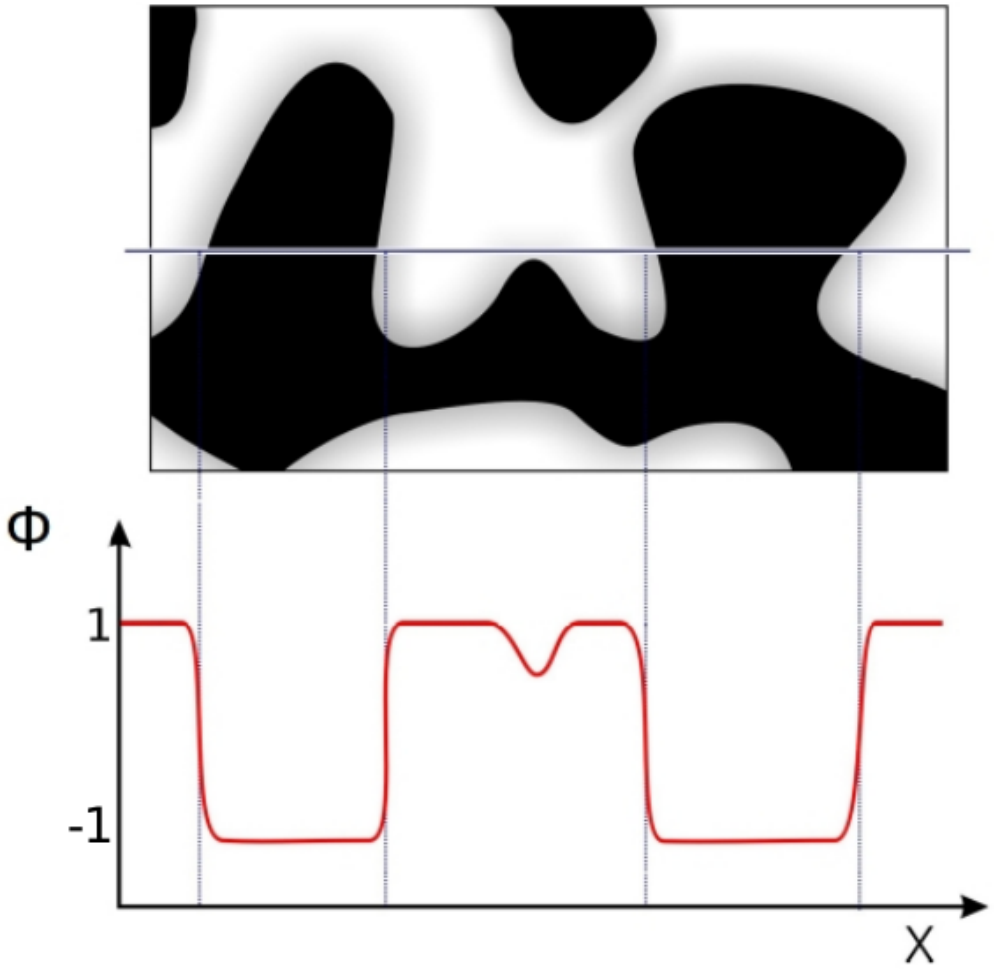

Figure 3.1: Example system containing two phases undergoing separation. The order parameter (ϕ) takes a value of 1 in the white region and -1 in the black region. The value changes rapidly, but smoothly, at the interfaces. Image source [30].

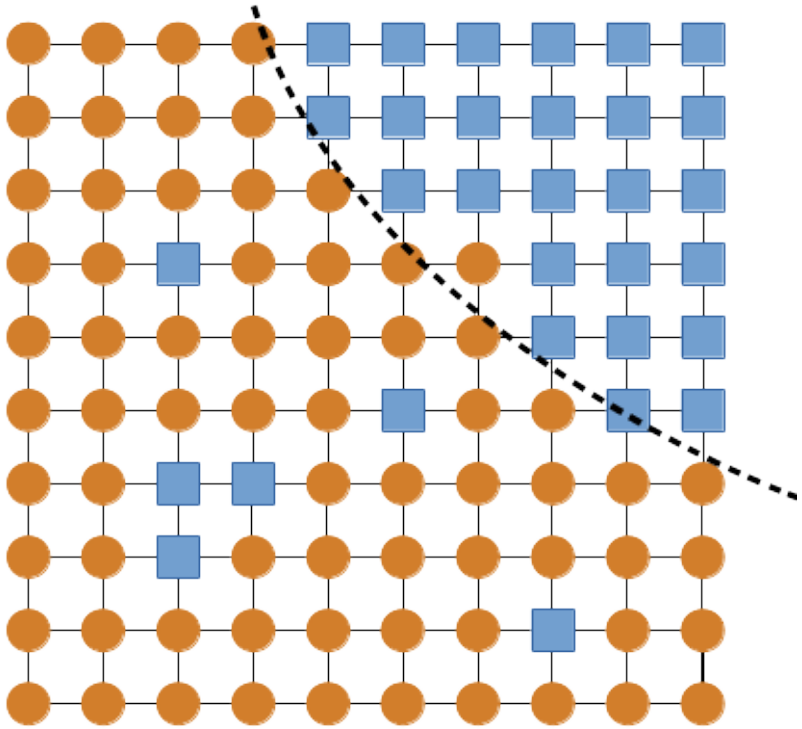

Figure 3.2: An illustration of the assumed atomic lattice, with orange circles denoting sites occupied by atoms, and blue squares representing a vacancy in the lattice. The dashed line represents the boundary of a pore contained within the system. Vacancies can travel through the lattice, exchanging places with nearby occupied sites. Through this mechanism, vacancies are able to exit the system during the densification process. We note here that the pore size is not to scale, but merely representative.

merely grow by diffusion of like components from nearby mixed regions to the surface of the nucleated bubble. This process will continue until only two regions of pure material exist, as such a configuration has the lowest interface area. As a large surface area is energetically unfavorable (the interface, not being pure, is energetically difficult to maintain), the configuration with minimum surface area is preferred.

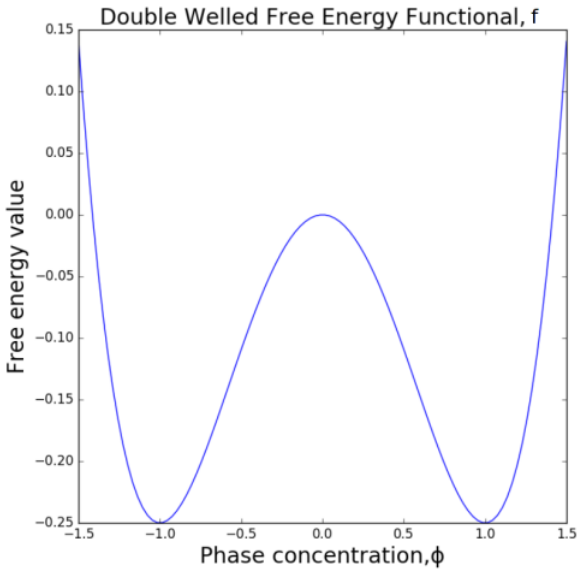

Figure 3.3: Double-welled free energy functional, f_{loc}. Minima at $\phi = \pm 1$ mean that full phase separation is most energetically favorable, and the system will attempt to evolve into such a configuration.

3.2.2 Polynomials by Example

Equations such as Equation 3.1 are difficult to solve on computers in general. It is usually necessary to cast such equations into an algebraic form that can be solved via basic operations, a task at which digital computers perform admirably. To this effect, we decompose the system we are trying to simulate (the fuel pin) into many smaller elements that are interconnected. By using sufficiently small elements, we can solve each individual location using simple algebraic relations, providing a good approximation of the true solution. This concept forms the basis of the finite element method, as we will now describe.

The concept of polynomial curve fitting provides a good starting point in describing the finite element method. In producing a polynomial that approximates a function, we take a set of points from the function, and seek a polynomial that passes through each point. There are many polynomial forms, with perhaps the simplest of the form:

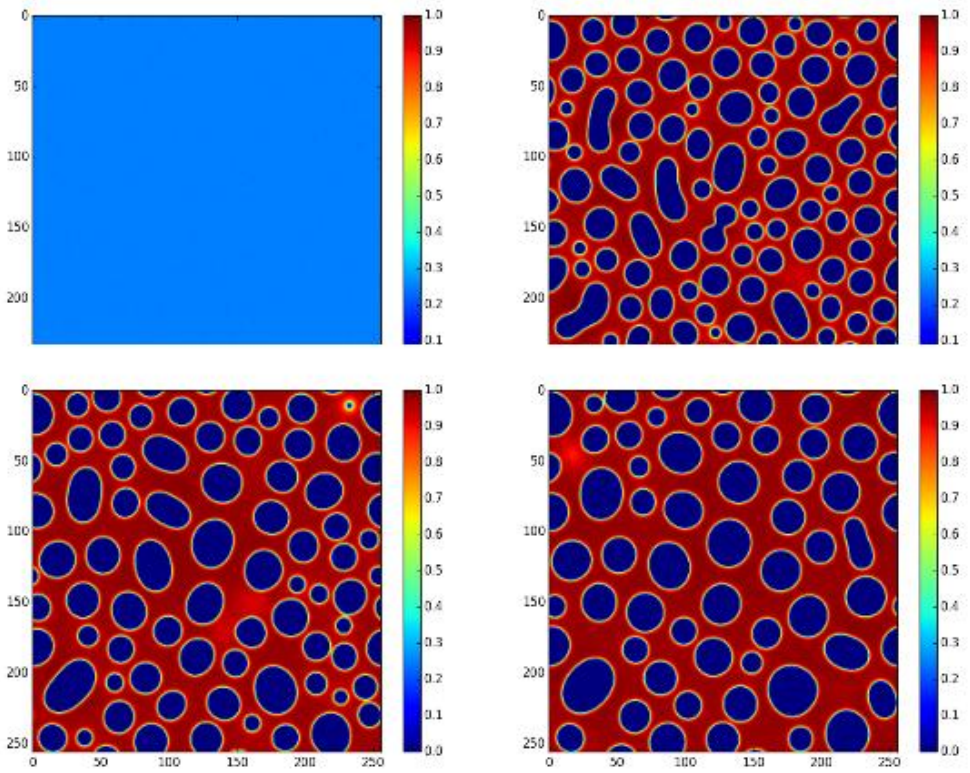

Figure 3.4: Evolution in time of a binary phase system. Initially well mixed, small bubbles of the blue phase nucleate and grow. After sufficient time, the two phases will entirely separate, leaving only a single interface. This configuration is energetically favored when using f_{loc} in a double welled form, as given in Equation 3.3.

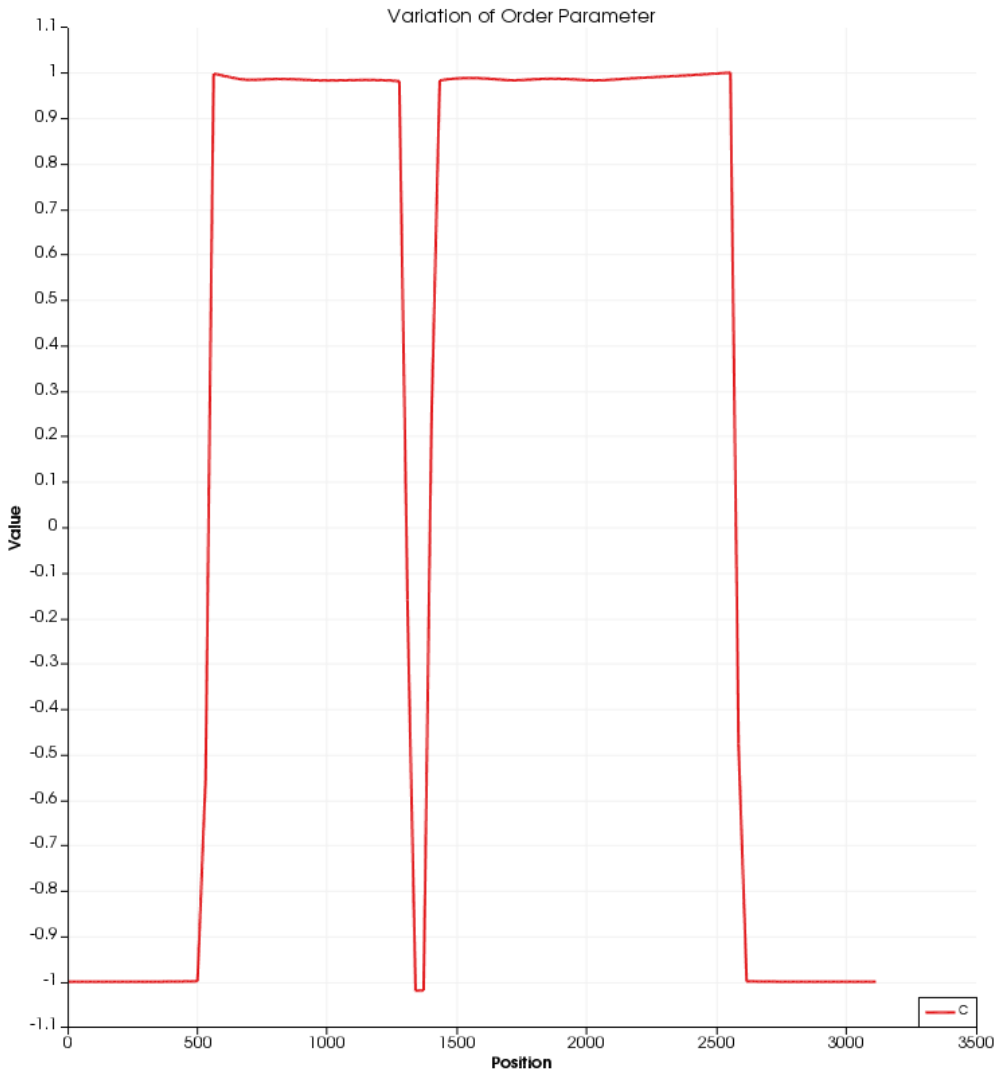

Figure 3.5: Variation of ϕ over a fuel/void system generated in this study. ϕ assumes a value of +1 within the fuel material, but rapidly changes to -1 over a few elements at the fuel/void interfaces. In this image, the region sampled starts at the boundary of the fuel, passes through a single void, and leaves the fuel region.

$$f(x) = \sum_{i=0}^{d} c_i \phi_i(x) \tag{3.4}$$

where the c_i form the coefficients and the ϕ_i are the basis functions. It is common for the basis functions to be of the form:

$$\phi_i(x) = x^i \tag{3.5}$$

We solve Equation 3.4 to obtain the values of the c_i's. If we are given n points, then selecting $d = n - 1$ (where d is the highest degree of the polynomial) will guarantee that $f(x)$ is unique and interpolary. It is best to illustrate the concept with a simple example. Consider the collection of points:

$$x_1, y_1 = (1, 2)$$

$$x_2, y_2 = (2, 1)$$

$$x_3, y_3 = (3, 2)$$

From Equation 3.4, our system of equations will be

$$y_i = c_1 + c_2 x_i + c_3 x_i^2, i = 1, 2, 3 \tag{3.6}$$

Keeping in mind our goal of solving for the coefficients c_i, we obtain the linear system:

$$\begin{bmatrix} 1 & 1 & 1 \\ 1 & 2 & 3 \\ 1 & 4 & 9 \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix} = \begin{bmatrix} 2 \\ 1 \\ 2 \end{bmatrix} \tag{3.7}$$

The solution to which is:

$$\begin{bmatrix} 5 \\ -4 \\ 1 \end{bmatrix} \tag{3.8}$$

This results in the solution function,

$$f(x) = 5 - 4x + x^2 \tag{3.9}$$

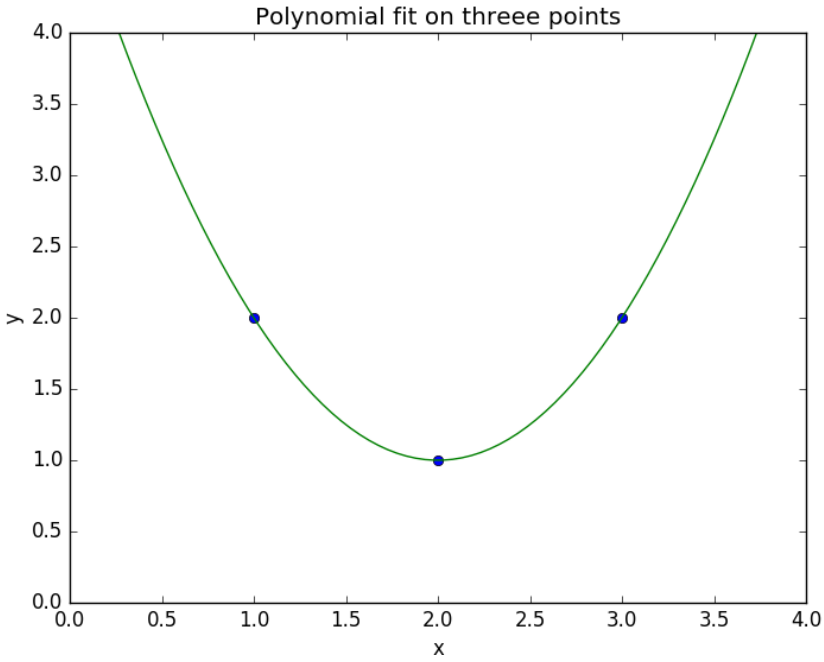

Figure 3.6: Polynomial fit of 3.4 to the points. Interpolation of the points is guaranteed, but accuracy outside of the point range is not.

Polynomial fitting as in the given example is expanded to be applied over the solution domain, with each polynomial being a globally valid component of the solution, but only providing a nonzero result at the element for which it represents. The domain is subdivided into many discrete parts, which are themselves represented by the original physical equation.

In order to allow the solution of these approximated elements, we must cast the original equation in its weak form.

3.2.3 Weak Formulations

With polynomial fitting, we can approximate continuous solutions with piecewise systems of equations. This enables us to solve difficult partial differential equations via solution of an algebraic system of equations. As an example, consider the 1-D steady-state conduction equation with a known source term,

$$-\nabla \cdot k\nabla u = f \tag{3.10}$$

We begin by recasting this equation in its weak form, that is, by reducing the degree of derivatives present and multiplying by a test function, to allow numerical flexibility. Start by rearranging all terms to one side:

$$-\nabla \cdot k\nabla u - f = 0 \tag{3.11}$$

Then, multiply each term by a test function, ψ:

$$-\psi(\nabla \cdot k\nabla u) - \psi f = 0 \tag{3.12}$$

Now integrate 3.12 over the problem domain, Ω:

$$-\int_{\Omega} \psi(\nabla \cdot k\nabla u) - \int_{\Omega} \psi f = 0 \tag{3.13}$$

It is useful here to recall Gauss's Divergence Theorem from vector calculus,

$$\int_{\Omega} \nabla \cdot \vec{g} dx = \int_{\partial\Omega} \vec{g} \cdot \hat{n} ds \tag{3.14}$$

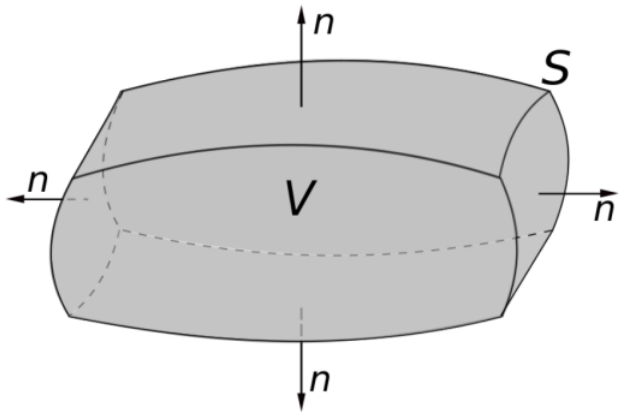

Figure 3.7: Divergence Theorem volume and surface. The time rate of change taken over the entire volume is equivalent to the total net flux across the surface, for a conserved vector field quantity. Image source [31]

Finally, we apply Equation 3.14 and the chain rule to the first term in Equation 3.13:

$$\int_{\Omega} \nabla\psi \cdot k\nabla - \int_{\partial\Omega} \psi(k\nabla u \cdot \hat{n}) - \int_{\Omega} \psi f = 0 \tag{3.15}$$

This constitutes the weak form of Equation 3.10. It is common to rewrite this equation using inner product notation, as will become useful when applied to MOOSE,

$$(\nabla\psi, k\nabla u) - < \psi, k\nabla u \cdot \hat{n} > -(\psi, f) = 0 \tag{3.16}$$

The next step to be performed in making the model suitable for solution via computer is by discretizing the domain, representing the points with interpolating polynomials. That is, we expand our solution variable,

$$u \approx u_h = \sum_{j=1}^{N} u_j \phi_j \tag{3.17}$$

where h is a nodal point on the finite element grid, and N is the number of nodes used. This is analogous to the expansion used in Equation 3.4, and by selecting the same functional form for ψ and ϕ, we arrive at the Galerkin finite element method. The weak form permits easier solution of the constituent equations, but is not strictly accurate. Thus, our problem changes from one of an exact solution to one of error minimization. We thus rewrite Equation 3.16 to reflect this fact:

$$R_i(u_h) = (\nabla\psi_i, k\nabla u_h) - < \psi_i, k\nabla_h u \cdot \hat{n} > -(\psi_i, f) \text{ for i = 1,...,N} \tag{3.18}$$

where $R_i(u_h)$ is the residual error vector.

By Equation 3.17, we have discretized all components except for the integrals appearing in Equation 3.15. In order to solve these on the grid, we relate the integral of a function to the sum of its element integrals,

$$\int_{\Omega} f(\vec{x}) d\vec{x} = \sum_{e} \int_{\Omega_e} f(\vec{x}) d\vec{x} \tag{3.19}$$

where Ω_e is a physical element in the domain.

The technique given here for the conduction equation can be extended to a large class of problems commonly encountered in science and engineering, including the Cahn-Hilliard equation (with some modification). Still remaining is a description of how it is implemented in MOOSE under the phase field physics module. In addition, we must specify the type of basis function we select to represent the solution. Also, the numerical algorithm chosen to

solve the system of equations is of some importance, as the equation system is large. These details will be explained after MOOSE is introduced, next.

3.3 MOOSE

3.3.1 MOOSE Basics

As taken from the mooseframework.org homepage,

> The Multiphysics Object-Oriented Simulation Environment (MOOSE) is a finite-element, multiphysics framework primarily developed by Idaho National Laboratory. It provides a high-level interface to some of the most sophisticated nonlinear solver technology on the planet... [32]

The purpose of MOOSE is to improve scientists' and researchers' access to advanced solution tools that utilize the finite element method. In the past, users typically had two options. While finite element meshing and linear algebra solution libraries were available, their use was complicated, being encumbered by the complexity of application development and the difficulties of parallel programming (if large problems were of interest). As an alternative, users could purchase licenses for expensive professional software, but these packages were often not sufficiently domain-specific to support the wide range of cutting edge research being performed.

MOOSE attempts to provide an alternative that is simultaneously open, capable, and easy for the user to modify. MOOSE implements the finite element method in a way that enables modular addition of physical phenomena at varying length and time scales. It provides a framework upon which users can specify the geometry and initial conditions of problems, invoke the appropriate physics packages already included, and specify solution techniques that lend well to their available hardware and particular problem type. The

addition of new physics kernels is also relatively straightforward, and typically does not require modification of any other part of MOOSE. MOOSE harnesses many modern software practices to improve collaboration with other researchers, and has an active community of users.

In this work, two main points of interaction with the MOOSE system were required to develop the capability of modeling annealed microstructures. The first involved the specification of the values of the field parameter, ϕ, to represent fuel and voids, respectively. We implement this with an initial condition class, denoted PorosityIC, to place voids at the desired locations for each microstructure configuration of interest. We will also need to invoke the appropriate phase field equations to solve ϕ over our domain. In MOOSE, physics "kernels" enable turn-key activation of various physical partial differential equations, including the Cahn-Hilliard equation.

3.3.2 Phase Field Implementations

The MOOSE phase field kernels (as we describe shortly) expect instruction on how to initialize the order parameter at each nodal point. The so called "CompactDataReader" class was written for this purpose. It takes a list of pore regions in the format

```
x-position , y-position , radius
```

and builds a list in MOOSE memory, which is then handed to PorosityIC. In the data file, we reserve the first entry to denote the location of the entire system, and assign a value of 1 to ϕ in this region to denote fuel. All other entries represent pores (or the exterior of the fuel) and receive an initial value of -1. Details of the algorithm used to select the pore

sites are discussed in Section 3.4.2.

In SmoothCircleBaseIC, the base class for initial conditions containing circles or spheres in MOOSE, the method "value()" is used by the mesh to assign ϕ's starting value at every nodal point. PorosityIC overrides value()'s operation with the following pseudo-code:

```
def value(gridPoint p):
    val = outval                        # given as -1, void
    dist = distance(p,circles[0])
    if (dist < circles[0].radii)        # within fuel, check for void
        val = inval                     # given as +1 by input file
        for (circle in circles[1:]):    # loop over all but first
            dist = distance(p,circle)
            if (dist < circle.radii):
                return outval           # in a void
    return val
```

This method is eventually called by MOOSE for every nodal point, p, in setting the initial value of ϕ. We store the boundary of the fuel in the first entry of "circles", with all subsequent entries being used to denote void locations. The cost of the void check is $O(N)$ (for each point p) in the worst case. For this work, the cost of the finite element calculations (eg. after initialization) was found to completely dominate this step.

The most general form of the Cahn-Hilliard equation used by MOOSE is given in Equation 3.20. We will specialize it to ignore additional external driving forces (E_d), and to consider only a single concentration variable ϕ.

$$\frac{\partial \phi_i}{\partial t} = \nabla \cdot M_i \nabla \left(\frac{\partial f_{loc}}{\partial \phi_i} + \frac{\partial E_d}{\partial \phi_i} - \nabla \cdot (\kappa_i \nabla \phi_i) \right) \tag{3.20}$$

In solution of Equation 3.20, ϕ is fourth-order in terms of its derivative. Thus, the usual technique of preparing PDE's for use with MOOSE is often modified. Introducing the chemical potential μ as the energy absorbed or released during a phase transition, we split the Cahn-Hillard equation into two parts, with the first seeking to solve ϕ w.r.t μ, and the second solving μ w.r.t ϕ. By flipping back and forth, we can update ϕ in time without explicitly needing the higher order derivatives. In MOOSE, this results in the following residual equations (where we have made the additional simplifications mentioned above),

$$\begin{aligned} R_\mu &= \left(\frac{\partial \phi}{\partial t}, \psi_m \right) + (M \nabla \mu, \nabla \psi_m) - \langle M \nabla \mu \cdot \vec{n}, \psi_m \rangle \\ R_\phi &= (\nabla \phi, \nabla (\kappa \psi_m)) - \langle \nabla \phi \cdot \vec{n}, \kappa \psi_m \rangle + \left(\left(\frac{\partial f_{loc}}{\partial \phi} - \mu \right), \psi_m \right) \end{aligned}$$

where f_{loc} is the double well potential. The parameters M and κ represent the vacancy mobility and interface coefficient, respectively, and are specified by the user in the MOOSE input file as material properties. All systems simulated were given a value of $M = \kappa = 16.0$.

With this splitting, the residual equation is formed by employing the process as described earlier to generate a weak form. Each term of the residual equation corresponds to a separate class, or "kernel", in MOOSE. The components are combined to generate the system of equations to be solved. In Table 3.1, the names of the kernels used are given, where we have simplified F to model our binary system. Also important is the boundary condition utilized in the solution of 3.21. The fuel bodies produced are surrounded by a slightly larger bounding box (which represents the environmaent) in which no fuel material

Name	Residual Term	Variable	Required Params
CoupledTimeDerivative	$(\frac{\partial \phi}{\partial t}, \psi_m)$	μ	-
SplitCHWRes	$(M\nabla\mu, \nabla\psi_m)$	μ	M
SplitCHParsed	$(-\kappa\nabla^2\phi) + \frac{\partial f_{loc}}{\partial \phi})$	ϕ	κ

Table 3.1: MOOSE Kernels invoked during solution of the binary system.

is present. No-flux boundary conditions are applied at the edges of this box.

3.3.3 Finite Element Implementations

Solution of the residual equation (Equation 3.18) requires the specification of two additional details. First, the interpolation functions must be specified. MOOSE supports a wide variety of shape functions. In many cases, the default "QUAD9" 2-D element is accurate and efficient. Its behavior is shown in Figure 3.8.

The second detail of necessity is the method by which the system of equations generated will be solved. While a deep understanding of the numerical techniques is not required to understand the results of this work, a short summary of the technique employed in MOOSE will be included here for completeness.

Newton's method is an iterative technique to finding an approximate root of a real valued function. In a single variable, it uses the derivative of the function and an evaluation at the current root guess to generate a better approximation, according to

$$x_{n+1} = x_n - \frac{f(x_n)}{f'(x_n)} \tag{3.21}$$

Newton's method has very good convergence properties for many problems, and is simple to implement. When the function's value and derivative are easily obtainable, it is an excellent choice for root finding. It can also be extended to a system of nonlinear equations and used to solve the residual equation given earlier, where the function roots minimize the residual error (in this case, the 2-norm of the error vector formed from each u_h component in the solution). In comparison to the other solution tools available in MOOSE, Newton's method has the best consistent performance on our particular problem, although the memory requirements were somewhat increased due to the need to store the Jacobian matrix. It was selected over the other primary nonlinear solver available in MOOSE (the Jacobian-Free Newton-Krylov method).

The extension of Newton's method to a nonlinear system is again best demonstrated by example. Consider the following nonlinear system in two dimensions,

$$f_1(x_1, x_2) = e^{x_1+x_2}$$

$$f_2(x_1, x_2) = sin(x_1)$$

We use a Taylor expansion around a point x_0,

$$\vec{F}(\vec{x}) = \vec{F}(\vec{x_0}) + \vec{J}(\vec{x}) * (\vec{x} - \vec{x_0}) + \Theta(x^2) \tag{3.22}$$

where $\vec{J}$ is the Jacobian,

$$\vec{J}(\vec{x}) \equiv DF(\vec{x}) = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \frac{\partial f_1}{\partial x_2} \\ \frac{\partial f_2}{\partial x_1} & \frac{\partial f_2}{\partial x_2} \end{bmatrix} \tag{3.23}$$

Dropping the truncation term (which makes the technique second-order accurate), we rearrange Equation 3.22 to solve for $\vec{x}$, the updated root approximation,

$$\vec{x} = \vec{x}_0 - (\vec{J}(\vec{x}_0))^{-1}\vec{F}(\vec{x}_0) \tag{3.24}$$

We can avoid taking the inverse of $\vec{J}$ by observing that there exists an s such that

$$\vec{J}(x_k)s = -\vec{F}(x_k) \tag{3.25}$$

Then, we use Gaussian elimination to obtain an updated x,

$$x_{k+1} = x_k + s \tag{3.26}$$

There are a wide range of techniques in solving the system from Equations 3.25 and 3.26 , which is large, sparse, non-symmetric, and linear. The Generalized Method of Minimum Residuals (GMRES), itself an iterative method, is fast and flexible, but requires preconditioning to be effective, as it belongs to the family of Krylov subspace methods. This is not a major issue, as we already posses the Jacobian for the nonlinear Newton solve, and can use this as a preconditioner. GMRES is the default choice in MOOSE, although others are supported. In summary, the overall solution process is as follows,

1. Produce the weak form of the nonlinear PDE to be solved.

2. Discretize the system and solution into nodal elements.

3. Rearrange the discretized system to a residual problem, in which we solve a system of nonlinear equations such that an error metric is minimized.

4. Use multivariate Newton's method as a root finder to find an approximate solution to the residual equations.

5. At each nonlinear Newton iteration, solve a linear system of equations, taking advantage of the Jacobian matrix as a preconditioner, via GMRES.

6. One time step has been solved, repeat for each time step.

3.4 Analysis

3.4.1 Parameters and Figures of Merit

Two key parameters were used to define the microstructure configuration. The first was the total initial porosity in relation to fuel material, expressed as a percentage,

$$TD = (1 - \frac{A_{pores}}{A_{system}}) * 100 \tag{3.27}$$

where

$$A_{pores} = A_{small} + A_{large} \tag{3.28}$$

This parameter is often referred to as the Initial Theoretical Density in the nuclear industry, and for oxide fuel pellets, a typical value is 92%. We expect configurations with a lower TD to require longer annealing times in order to remove all void space. We are not

Parameter	Range
$TD(\%)$	92,94,96
$\chi(\%)$	0,20,40,60,80,100

Table 3.2: Initial porosity characteristics.

sure, however, how bimodal systems will respond to varying TD values, so we simulate a range of configurations.

The second parameter that specifies microstructure configuration determines the ratio of small pores to all other pores in the compact, defined as

$$\chi = \frac{A_{small}}{A_{small} + A_{large}} * 100 \tag{3.29}$$

With this parameter, we control the degree to which the microstructure is bimodal. Configurations with $\chi = 100$ contain only small pores. These configurations are what are deemed to be comparable to fuel structures fabricated via traditional sintering methods, without any attempt at porosity control. As we decrease χ, we add more and more large pore regions, presumably via addition of pore formers that are subsequently eliminated during sintering. At $\chi = 0$, the only pores present are those with the larger diameter. This configuration is not necessarily realizable in practice, but is useful for comparison purposes (as will be seen in the results section). The combinations of TD and χ form our configuration space, and is enumerated in Table 3.2.

3.4.2 Initial Porosity Generation

At this point, it is important to mention that the MOOSE simulations are executed in two dimensions rather than three. Thus, spherical pores are reduced to circular pore regions. This approach is common in problems where the mechanisms are independent of the direction of propagation. In our case, a uniform and constant Mobility allow this simplification to hold. We will not consider 3-D geometries further, in the interest of avoiding undue computational expense.

We seek to model circular fuel pin geometries with two phases (fuel and void space). To generate systems with differing void placement, we adopt a stochastic approach in selecting the location of the voids. We randomly select a location to place a void within a defined bounding area. Keeping track of the quantity of small and large voids already placed, we select a size that keeps us close to our objective mixture, while also ensuring that the new void does not overlap with an existing void or go beyond the surface of the system. Once these criteria are met, we restart the procedure, adding pores until our desired initial porosity is reached. A simplified description of this algorithm is given in pseudo code below.

```
voids = []
small = large = 0
while ((small + large) < targetArea):
    chiCurrent = small / (small + large)
    if (chiCurrent < chiTarget):
        addVoid(small_r)
        small += pi*small_r**2
```

```
        else
            addVoid(large_r)
            large += pi*large_r**2


    addVoid(radius)
        x_pos = rand()
        y_pos = rand()
        if (!overlap(voids,x_pos,y_pos,radius)):
            voids.add(x_pos,y_pos,radius)
        else:
            addVoid(radius)
```

Occasionally, the system will be just short of the target porosity, but require a large void to meet the χ requirement. This results in a large pore being placed and causing the system to have a slightly lower porosity than expected (this can occur for small pores as well, but the effect is not as pronounced). We quantify this behavior by tracking the percent error between the expected and actual porosity. The error found here was never more than 3%, and was typically much smaller. A similar analysis is performed to ensure that the actual χ value is close to what is expected. We also note that in the implementation of the void placement algorithm, a limit is set on the number of times to try placing a void of either size, as in the extreme, highly porous systems may not have any space for another void. This resulted in the occasional failure to generate a configuration, which was always recoverable by simply running the algorithm again, with a different random number seed.

We note that the algorithm is $\Omega(N^2)$ in the best case, and can be worse, depending on the density of voids. While these simulations keep N at a tractable level (given the impressive speed of modern computers), larger simulations would require an improved algorithm for void placement. Such algorithms were briefly investigated and found to be promising, with expected complexities of $\Theta(NlogN)$.

3.4.3 Post-Processing

As the fuel compacts are evolved, we expect void space to move from the interior of the fuel structure to the surface. This has the effect of reducing the radius of the fuel cross section, a behavior we can directly measure as a means of comparing the shrinkage rates between different configurations. To this affect, we developed an analysis tool that tracks this change, using Paraview's [34] built-in LineSampler to record the extent of the fuel compact at each simulation time step. The general algorithm employed is as follows;

1. Sample a diameter of the fuel compact.
2. Record the outermost points at which ϕ transitions (based on the point at which ϕ goes below 0).
3. Measure the distance between the points as the fuel diameter.
4. Repeat a few times for each sample to average out small fluctuations, and produce an area.

The system sizes for each configuration are measured at each time step and then plotted for comparison. We also keep track of the error at two levels of granularity. By

averaging over a few line samples per test, we reduce our measurement error, which is important when fuel surface deformations occur. At a higher level, multiple runs for each configuration are averaged to account for "configuration" error, which arises due to the random placement of voids and the systems' relatively small size. Our error metric utilized here is the standard deviation.

In addition to a line sample, which produces shrinkage rates, we will also be interested in the average concentration of vacancies within the fuel region. By recording the value of ϕ at every point within the fuel, and ignoring void regions, we can average these measurements to arrive at the average vacancy concentration in the entire system. This average is not useful for understanding which particular regions of the system are vacancy rich, as we expect significant increases near pores and depression in pore-sparse areas. It is insightful, however, in measuring the overall vacancy transport gradient that is generated by the dissolution of the small pores into the fuel matrix. This will prove useful in comparing the behavior of the monomodal and bimodal configurations. A postprocessor was developed that measures this metric at each time step, again averaging over each run.

A brief description of the Paraview [34] software is warranted. This open source visualization package enables researchers to access powerful visualization and analysis tools. Its open and modular nature encourages extension; this fact proved valuable when developing the particular analysis techniques used for this work, as all the actual data extraction was abstracted by Paraview. We were solely responsible, therefore, for interpretation of the data returned. Paraview is developed by Kitware and Sandia National Laboratory.

3.4.4 Metaparameters

In development of the analysis tooling, a few parameters emerged as being important to the analysis itself, but not the results in particular. Some of those parameters are described here.

The system size is limited in practice by the number of nodes that could be reasonably simulated on the available computer hardware, and is selected to be as large as possible while still giving good grid resolution. Keeping the node count constant, a larger system would be resolved in less detail in comparison to a small system, but smaller systems were expected to be more sensitive to surface effects. A practical node count limit was found to be 4096 X 4096. A system size of 1100 x 1100 units is used, resulting in around 4 nodes per simulation length unit ($\Delta x = 0.268$).

Small pore regions were given a size of 5 units, and large pores 50. While the ratio of the two pore sizes is similar to that used in [15], the absolute size of the pores in comparison to the system diameter is impossible to replicate with the small finite system used. In this study, the ratio of large pore size to system size was 10^{-1}, whereas for physical fuel elements it is much smaller. This contributed to surface effects (as we will expound upon in the results section), but is not expected to detract from the validity of the results. Additionally, a parameterization of the system size is carried out to assess the impact this relatively low ratio has on the system evolution.

Each configuration enumerated in Table 3.2 was repeated 5 times for statistical reliability, with the intent of measuring surface effects. This count, referred to as N, was deemed to be sufficient, given the resource constraints.

3.4.5 Automation and Reusability

The entire process of fuel compact generation, MOOSE simulation, and the various data analysis steps is scripted to require minimal user management. Indeed, entirely new datasets, complete with analysis and graphical results, can be generated by adjustment of the desired parameters and only a few commands. The automated tasks include (in order of execution):

1. Generate system compacts for the combination space of TD and χ desired, with each configuration repeated N times for statistical soundness. This was automated via a bash shell script loop used to call the CompactGeneration routine. In this step, occasional compact generation failures required manual re-running of the offending member. In a typical batch of 90 samples, a few failures are expected.

2. Submit the samples to the HPC queue for simulation in MOOSE. A base MOOSE input file is read by the submission script and modified to point to the relevant compact file. All jobs are then submitted to the HPC job queue and executed by MOOSE.

3. After all the jobs have completed, the log files are checked for errors, rerun if necessary, and the results copied back to local storage for analysis. This step is manual.

4. Each configuration is line sampled twice at each time step. The results of the line samples are averaged and saved for analysis in the next step.

5. The averaged system diameters are converted to areas and the standard error for each configuration group computed. These values are then plotted, with user options

available to specify which configurations should be shown. This option is provided to make it easy to generate figures for analysis or publication.

6. A Paraview script is used to enumerate all points in each system, and compute the average concentration of vacancies in the non-pore regions at each time step. Again averaging for each configuration group, the results are plotted to figures.

7. Optional routines generate and display various error metrics used during analysis of the data. These error metrics are displayed over the time steps simulated.

A flowchart visualizing the process is given in Figure 3.10

3.4.6 Execution Costs

For a full configuration space run with N=5 samples per configuration, the 90 MOOSE simulations require around 1 week of real time on the Millett Research Group queue (32 Intel Xeon-equipped nodes). Each job is typically run on 4 nodes. The compact generation step is completed for all configurations in minutes, and the post-processing step in a few hours. Once postprocessing has been completed, plots can be reconfigured in a few seconds. Each sample requires around 2GB of storage space, or around 180GB for the entire experiment.

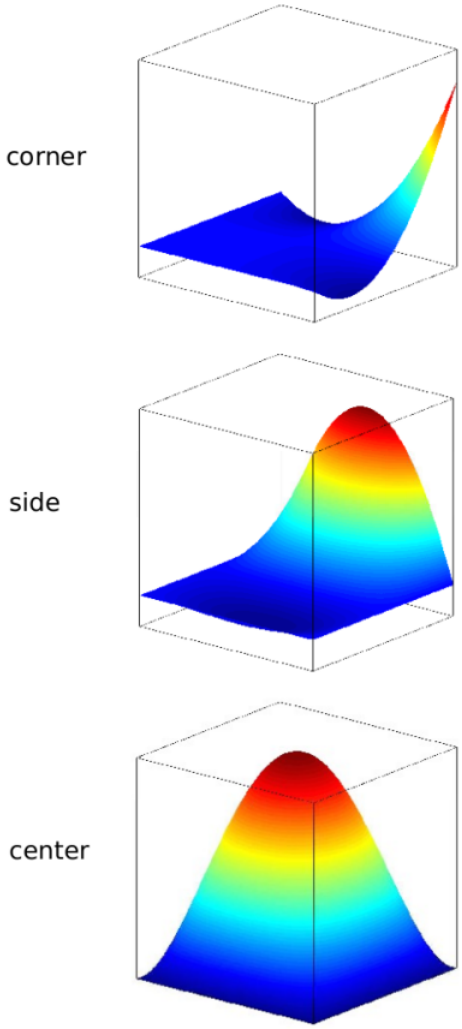

Figure 3.8: Bi-quadratic Lagrange Shape functions in 2-D, showing 3 of the 9 nodal types used for a finite element. Image source [33]

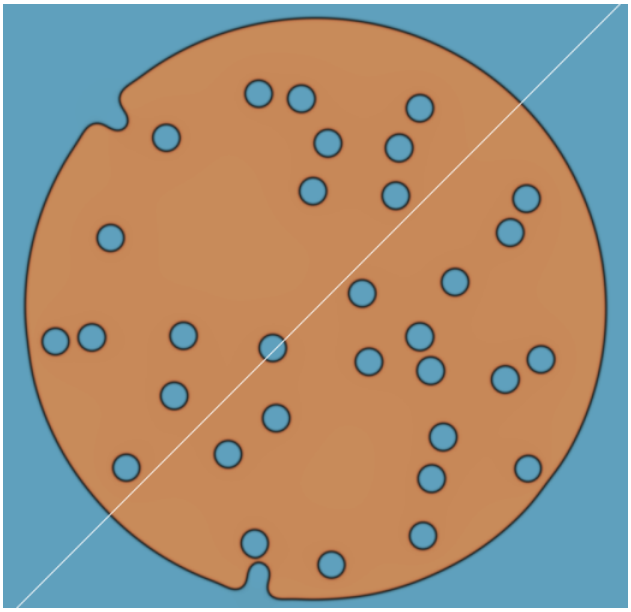

Figure 3.9: Example line sample object in Paraview. The value of ϕ is sampled along the white line at each time step. The outermost points at which ϕ goes to a depressed value are used to compute the size of the system.

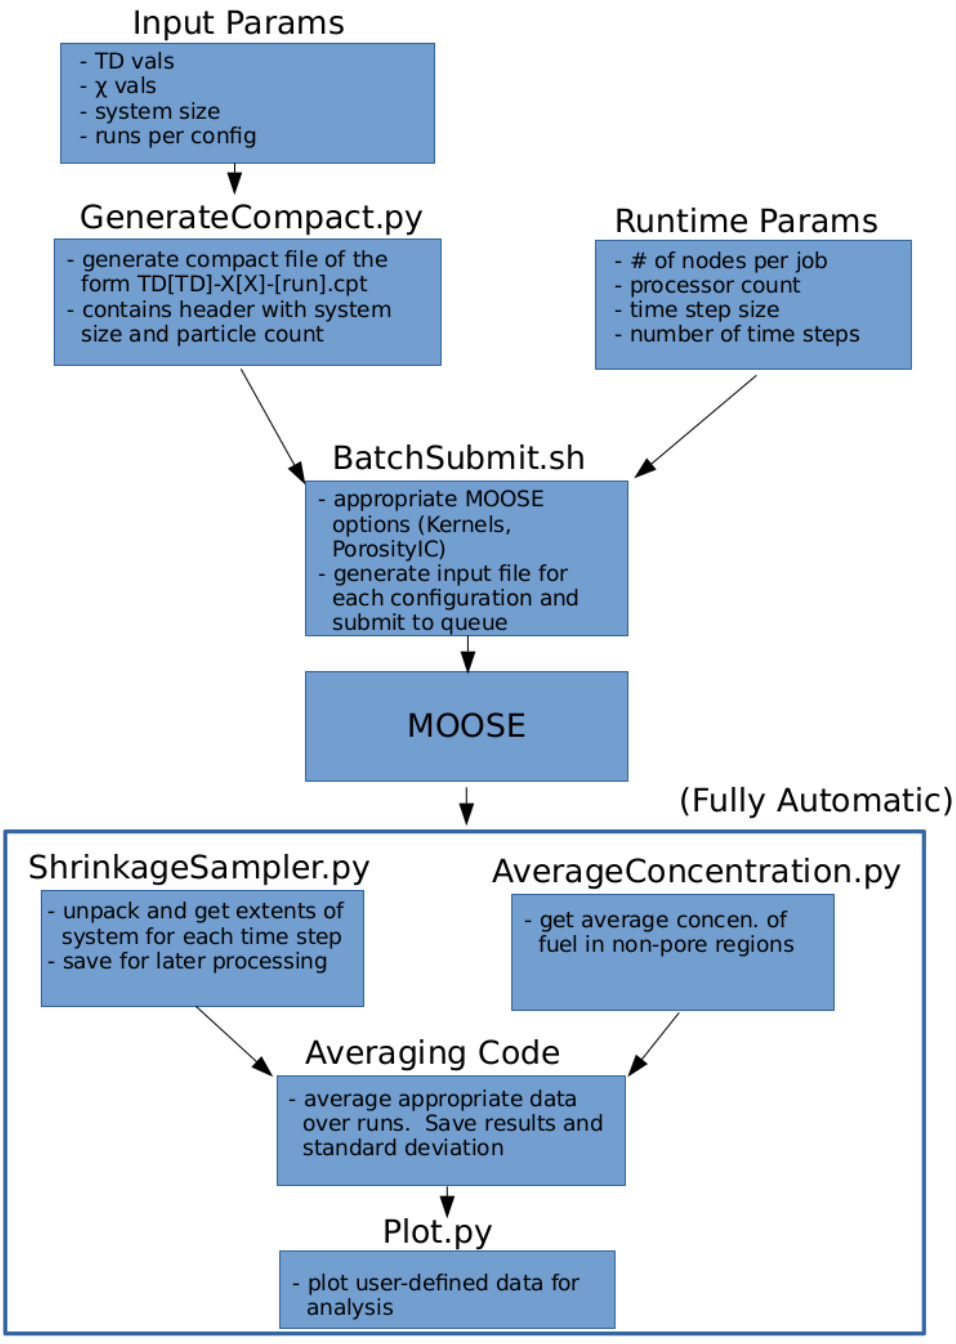

Figure 3.10: Steps involved in the generation, execution, and processing of the simulations.

4 Results

4.1 Introduction

This section summarizes the results obtained from the simulations, and provides an interpretation of them. On the whole, it is shown that the introduction of the bimodal pores greatly improves the ability of the microstructure to resist porosity loss. The time required to reach an equivalent fully dense state in comparison to microstructures without bimodal pores is extended. The greatest relative improvement is seen late in the annealing process. We present here selected results, primarily focused on the $\chi = 100, 60, 40$ and 0 cases. We consider $\chi = 100$ the baseline monomodal configuration fabricated via traditional techniques. The $\chi = 40$ and 60 cases are the most intriguing bimodal configurations. At $\chi = 0$, the configuration is comparable to the $\chi = 100$ case, but with larger pores present. It is useful in benchmarking the behavior of the other configurations. As a preliminary note, most images follow the color scheme of orange corresponding to fuel material, and blue to pore regions or the fuel exterior.

4.2 Qualitiative Results

4.2.1 Compact Generation

In Figure 4.1, a sampling of the microstructures generated for MOOSE by the compact generation algorithm are shown at the $TD = 92\%$ value. Here, we see both the size of the pores in comparison to that of the overall structure and the ratio of the two pore sizes. On close inspection of Figure 4.2, we observe that a few large pores have been placed near

the outermost surface of the fuel by random chance. This behavior is observed in many configurations, and during evolution contributes to deformation of the fuel surface. The pores produce sharp surface features on the fuel body, which are then smoothed out with continued evolution. This deformation behavior is not easily observable for structures with a majority of small pore regions, owing to their small size and subsequent reduced ability to modify the surface.

From the images, we become convinced that the systems are sufficiently random for our needs. The images shown indicate the ability to generate microstructures with the two parameters of interest (TD and χ) for the ranges given in the parameter space.

We are also concerned with the choice of grid resolution, as it is important that interfaces be resolved over a few elements. In Figure 4.4, we see that the interfaces are only resolved over one node; this is to be expected, as we assign the order parameter discretely at each point, and do not initially interpolate its value. By T=1, however, the interfaces have extended to the desired width (see Figure 4.5), and continue to maintain this size throughout the rest of the simulation. Note that when we parameterize the system size (near the end of the results section), the element count is adjusted to maintain the desired resolution, although a graphic of such will not be shown.

4.2.2 Shrinkage Behavior

Several example evolutions of the different microstructures are presented in Figure 4.6. The systems are first shown in the initial configuration, and then at equivalent time steps later in the simulation. In this image, all configurations displayed are at the same total initial porosity ($TD = 92\%$), but by T=80 it is clear that $\chi = 100$ has experienced almost complete

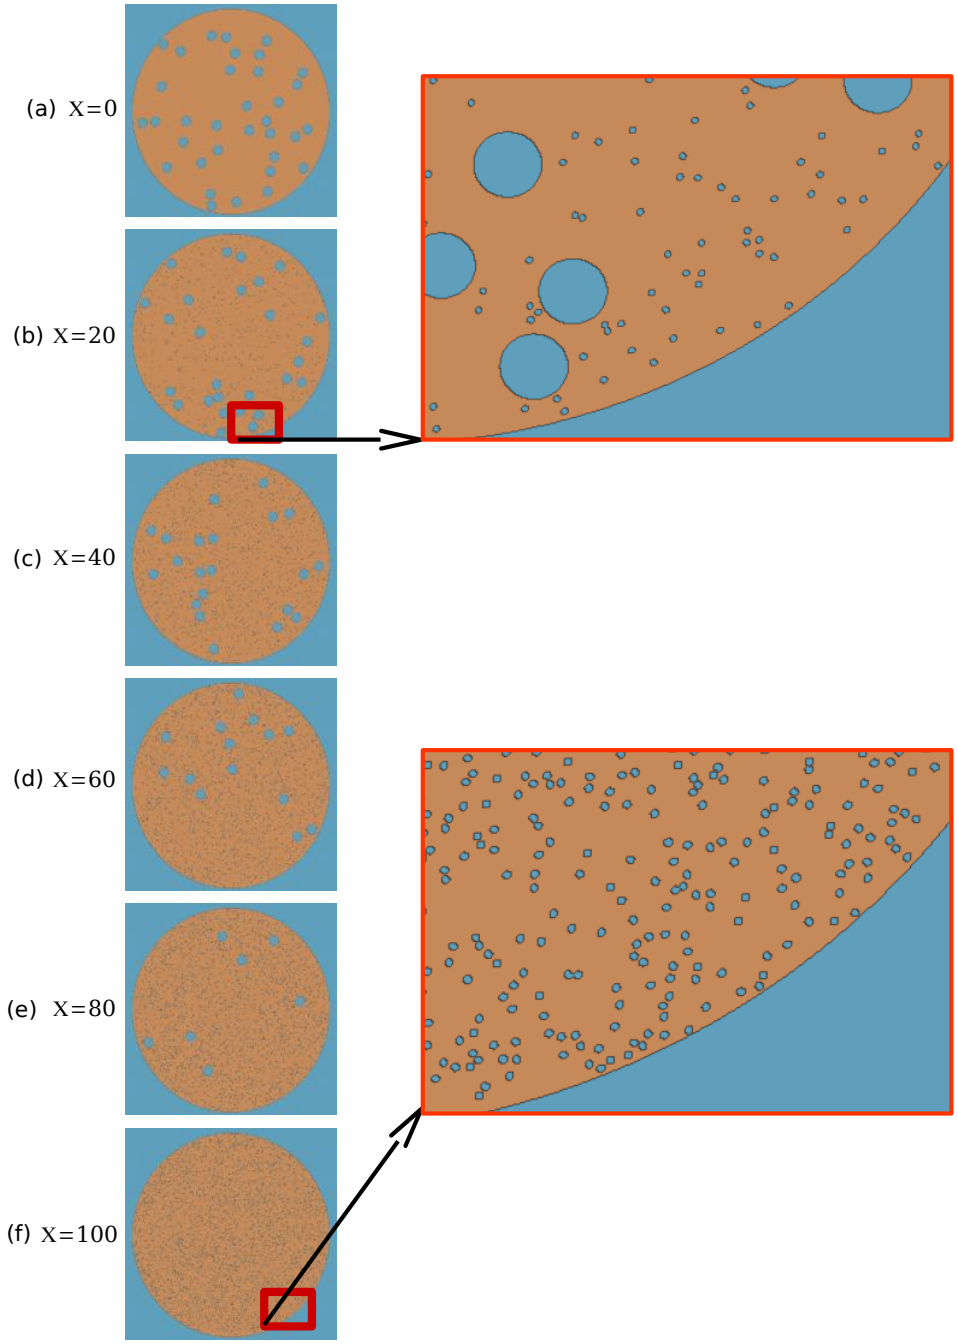

Figure 4.1: Example configurations generated by random void placement. Each configuration has convincingly random placement, based on visual inspection. Zoomed views highlight the relative sizes of the pores.

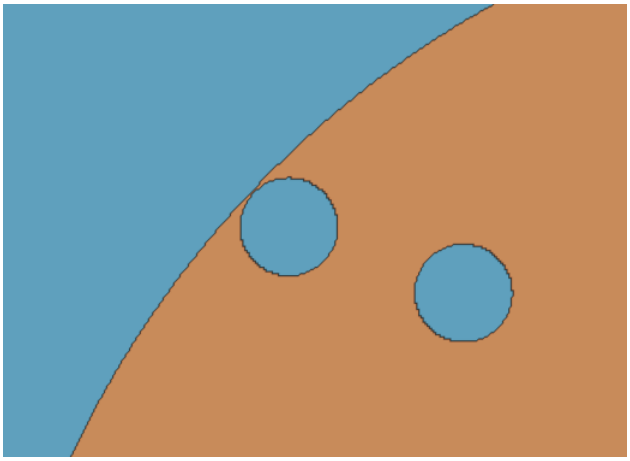

Figure 4.2: Zoomed view of a large pore in close proximity to the outer surface. The edge of the pore is only a few grid spacings from the fuel material termination. This often resulted in surface deformation, requiring several diameter samples to produce a reliable area estimate.

porosity loss, shown by the elimination of the yellow regions. The other configurations have persisting pores. This is evident at the remaining TD values as well, with the evolution to fully dense in general occurring in a shorter period of time, owing to the decreased total pore space that must be eliminated. The small vacancy sites, initially concentrated in the small pores, are rapidly dissolved into the fuel matrix, and by T=1, they are mostly hidden from visual inspection, as shown in Figure 4.7. This is in part due to the time step value used, as a reduction shows greater persistence of these pores. As the generation algorithm employed resulted in sharp interfaces between the pores and fuel (the interface is initially resolved over only a single element), the corresponding interface free energy is excessively large, and simulation in time will immediately eliminate these pores except in the smallest of time steps. This fact was not found to be problematic, however, as simulation at a reduced time step (see Figure 4.8) did not result in significantly different behavior. The coarse time step was therefore adopted for computational expediency.

In some cases, especially for structures with many small pores, it is observed that

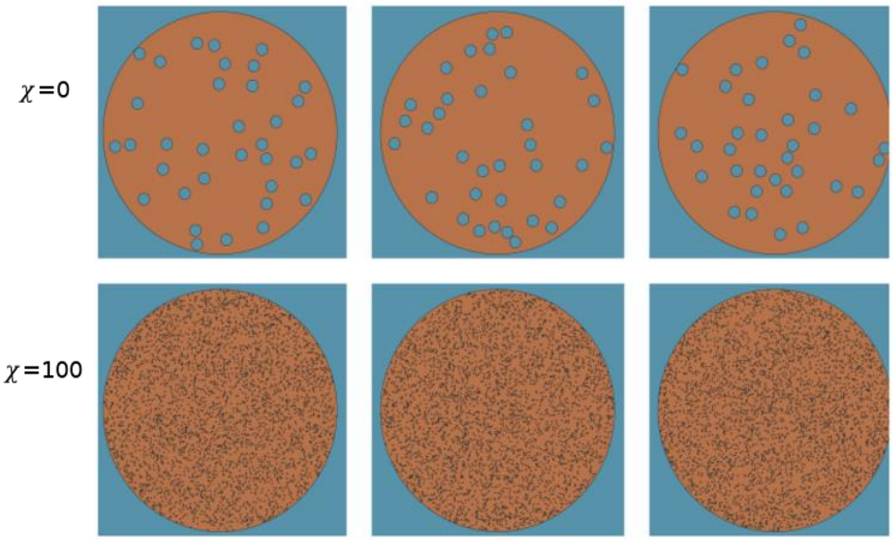

Figure 4.3: Several configurations at T=0 highlighting the random placement of pores. With sufficient samples for each configuration, a reasonably averaged spatial behavior is obtained, especially for the configurations with large χ.

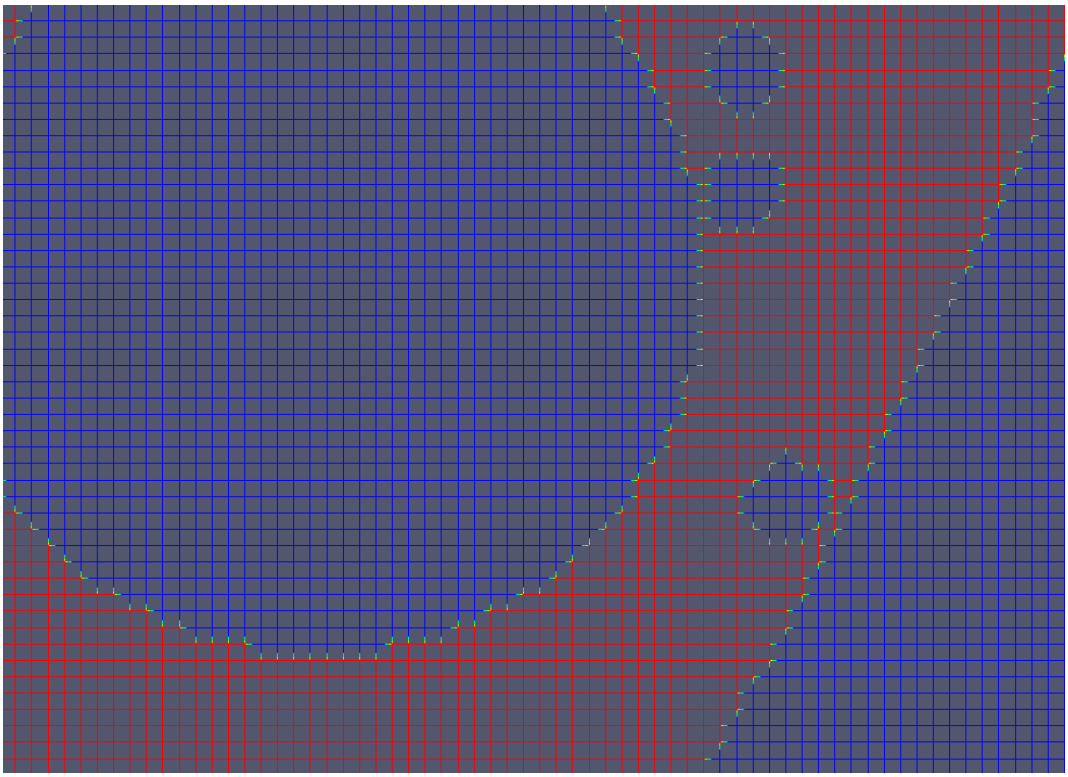

Figure 4.4: In the image, the finite element grid is shown overlaid on the concentration parameter. The bottom right of the image is the outer boundary of the fuel compact, with pores of both sizes visible within. Initially, the interfaces are sharp, owing to the method by which ϕ is assigned in PorosityIC. Here, we have adjusted the color scheme to make the grid easier to see, with blue denoting a pore and red fuel.

intermediate sized pores would occasionally nucleate and grow for some time before eventually being eliminated (Figure 4.9). The high density of mobile vacancies sourced from the small pores makes such nucleation feasible, and once it occurs, temporary growth can be energetically favorable, as this reduces the curvature and hence the interface energy of these nucleated regions. While growth can occur, no nucleated region will posses lower energy than the system's outer surface, resulting in the eventual elimination of the temporary void. We can observe a similar effect for structures with large pores, some of which are also capable of temporary growth Figure 4.10. The vacancies sourced from small pore regions, having been

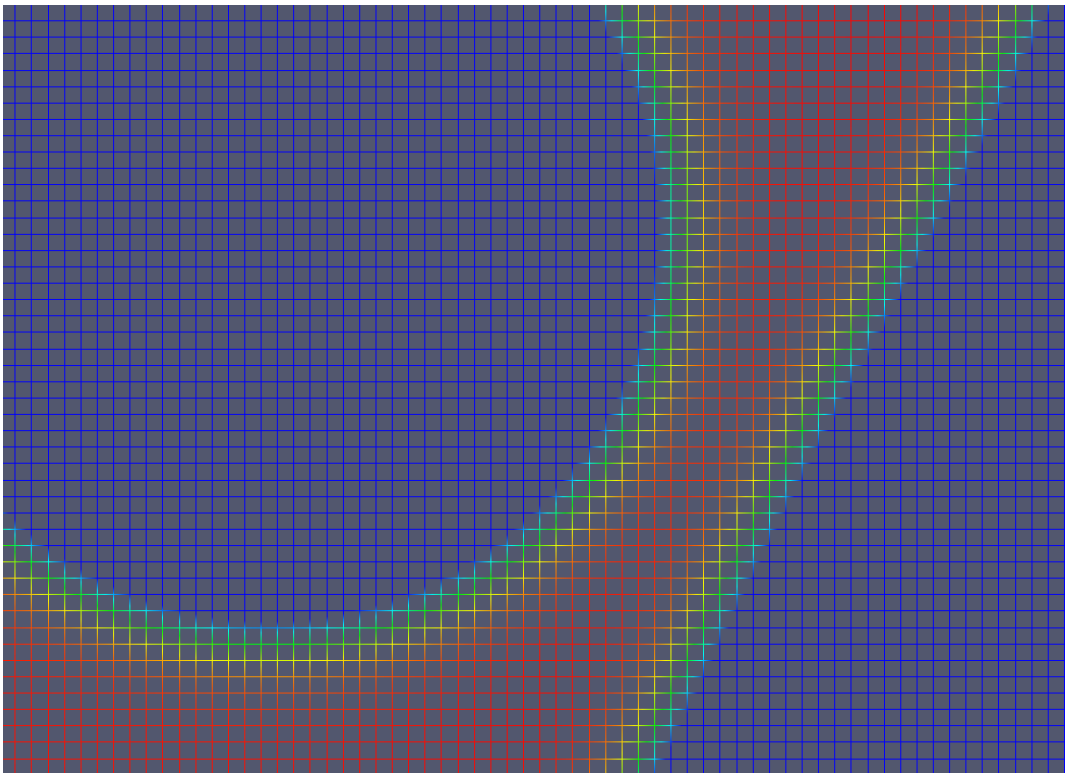

Figure 4.5: By T = 1, the interfaces have widened to the point at which they are resolved by around four grid elements. We know this to be sufficient in capturing the details of the Cahn-Hilliard model. Again, blue denotes pore space and red fuel, with intermediate values observed at the interfaces.

rapidly dissolved, migrate to the surface of the fuel but are interrupted by these large pore regions. Some thus contribute to large pore growth. These pores are again temporary, but on average persists for longer in the simulation than the intermediate nucleated sites, owing to their larger size.

4.3 Quantitative Results

4.3.1 Shrinkage Behavior

Plots are given comparing the shrinkage rates of selected configurations, along with the standard error bars, in Figures 4.11, 4.12, and 4.13. It is immediately evident that the

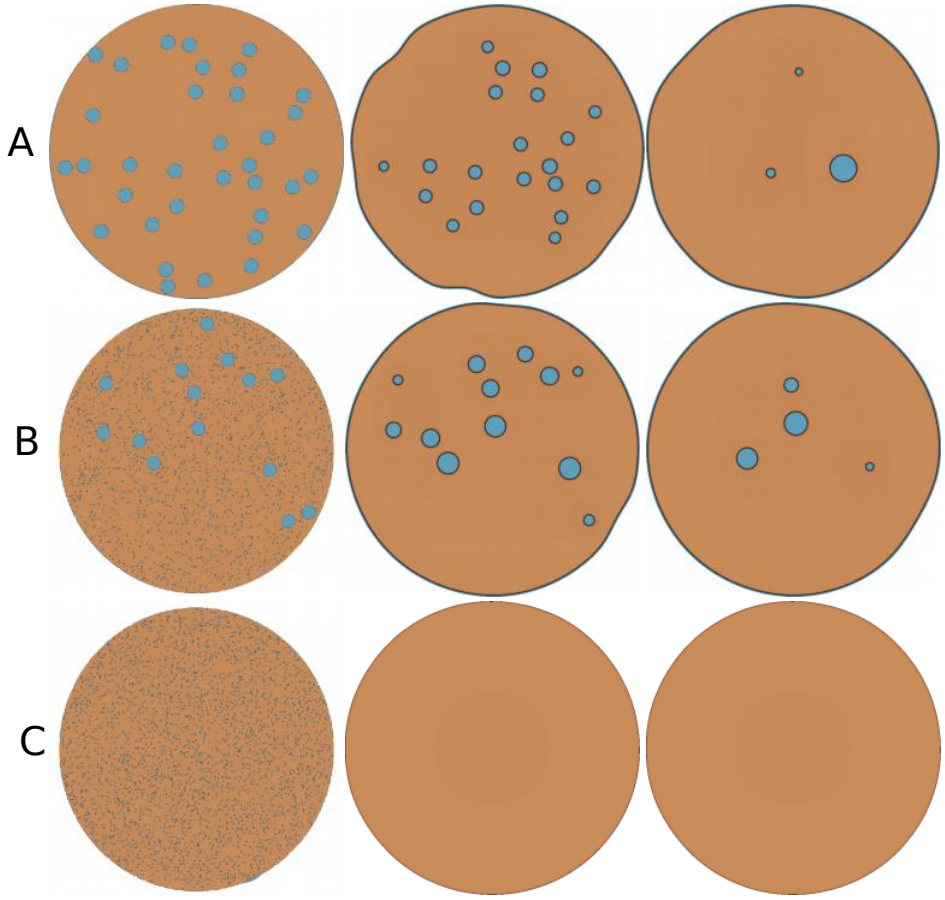

Figure 4.6: Evolution of three configurations ($\chi = 0, 60, 100$) at $TD = 92\%$. At T=0 (left column), the systems have not begun evolution. T=20 (middle column) shows the small pores have completed dissolution within the fuel matrix, with some surface deformation apparent for $\chi = 60$ and 0. At T=80 (right column), significant void space is lost for all configurations.

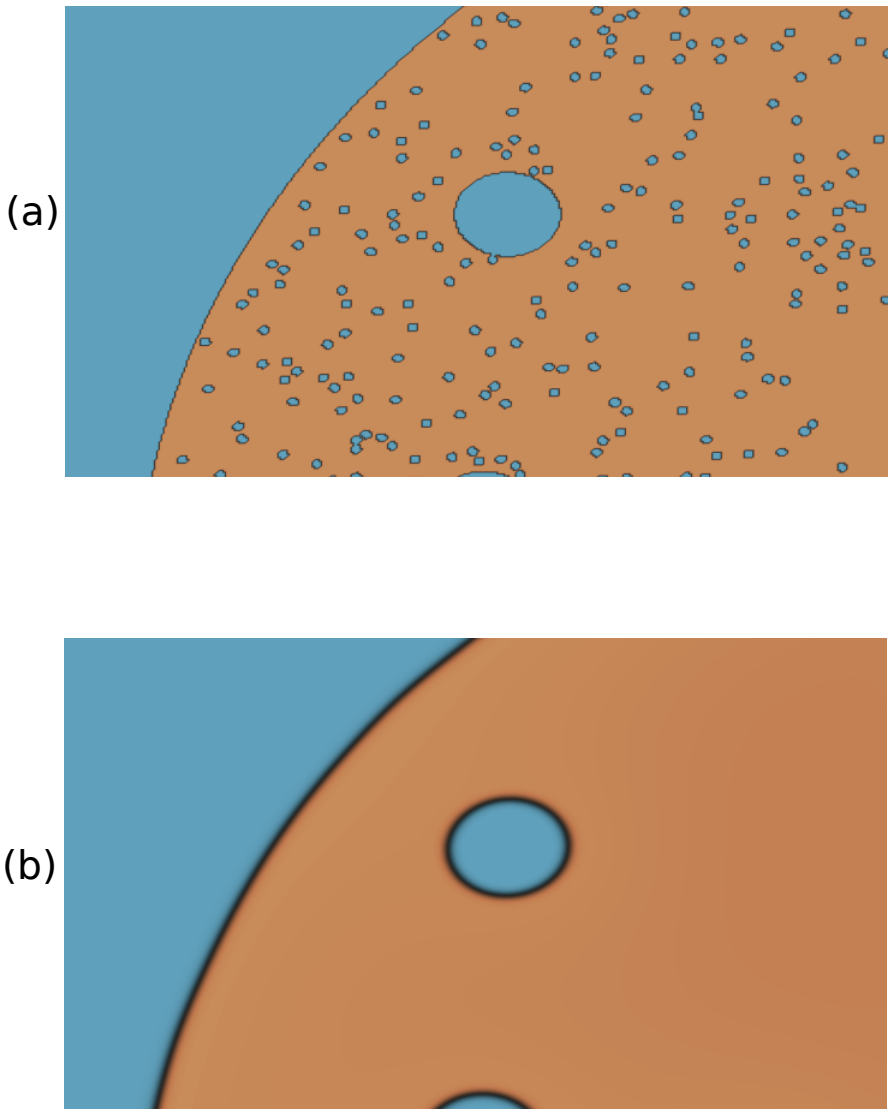

Figure 4.7: At T=0 (a), the theoretical as-fabricated condition is shown. One time step later, all visual evidence of the existence of the small pores has disappeared. Despite this fact, the bimodal configurations exhibited different behavior versus the monomodal structures.

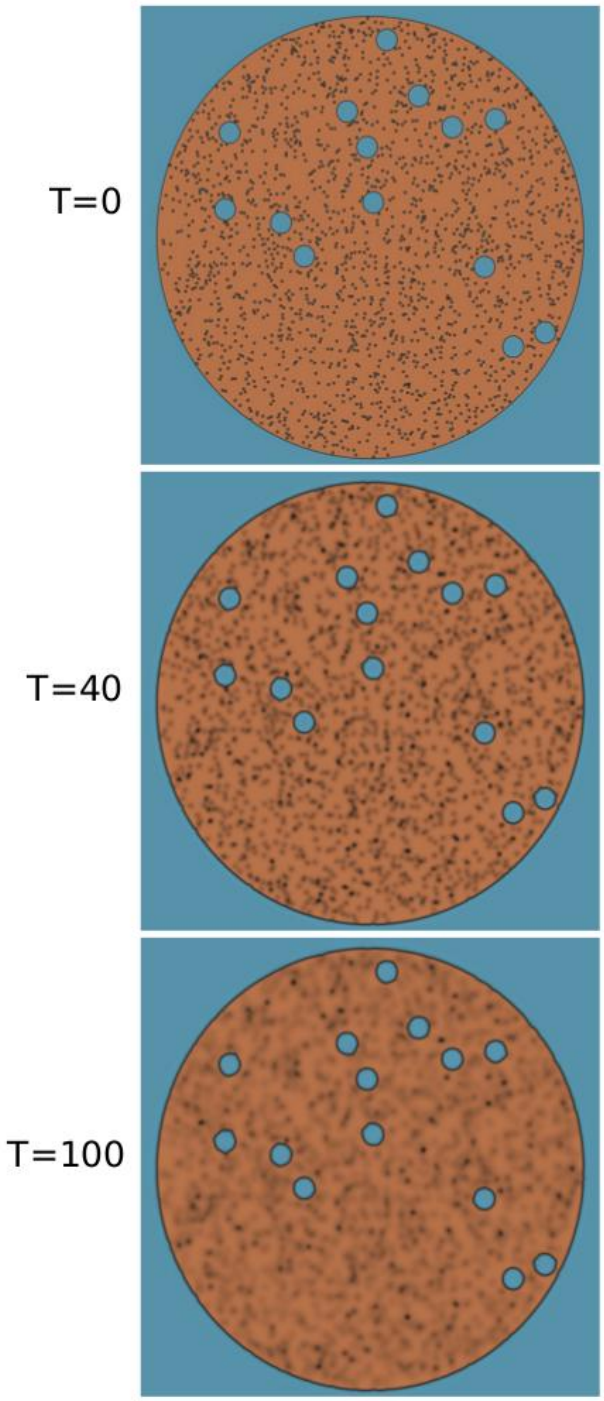

Figure 4.8: At a reduced time step of dt=0.001, vacancy dissolution from the small pores is captured in greater detail. The blurred features are evidence of early diffusion, which is required for absorption in the extremely vacancy rich region directly surrounding the small pores. Figures from $TD = 92\%$ for $\chi = 60$.

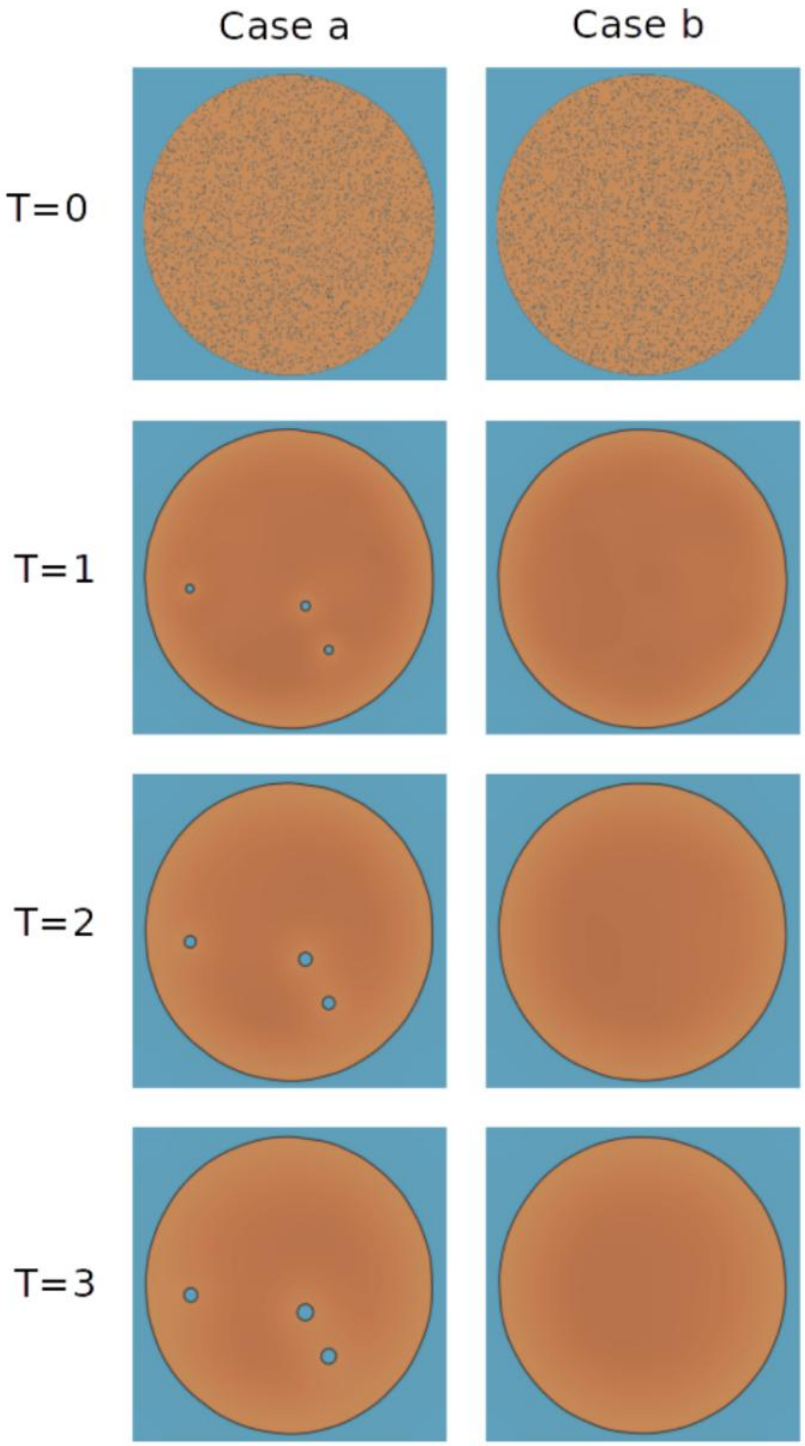

Figure 4.9: Nucleation of intermediate pores from vacancy diffusion. Two cases are exhibited here: in (a), pores are nucleated from vacancies soured from the small pore regions. In (b), no pores are nucleated.

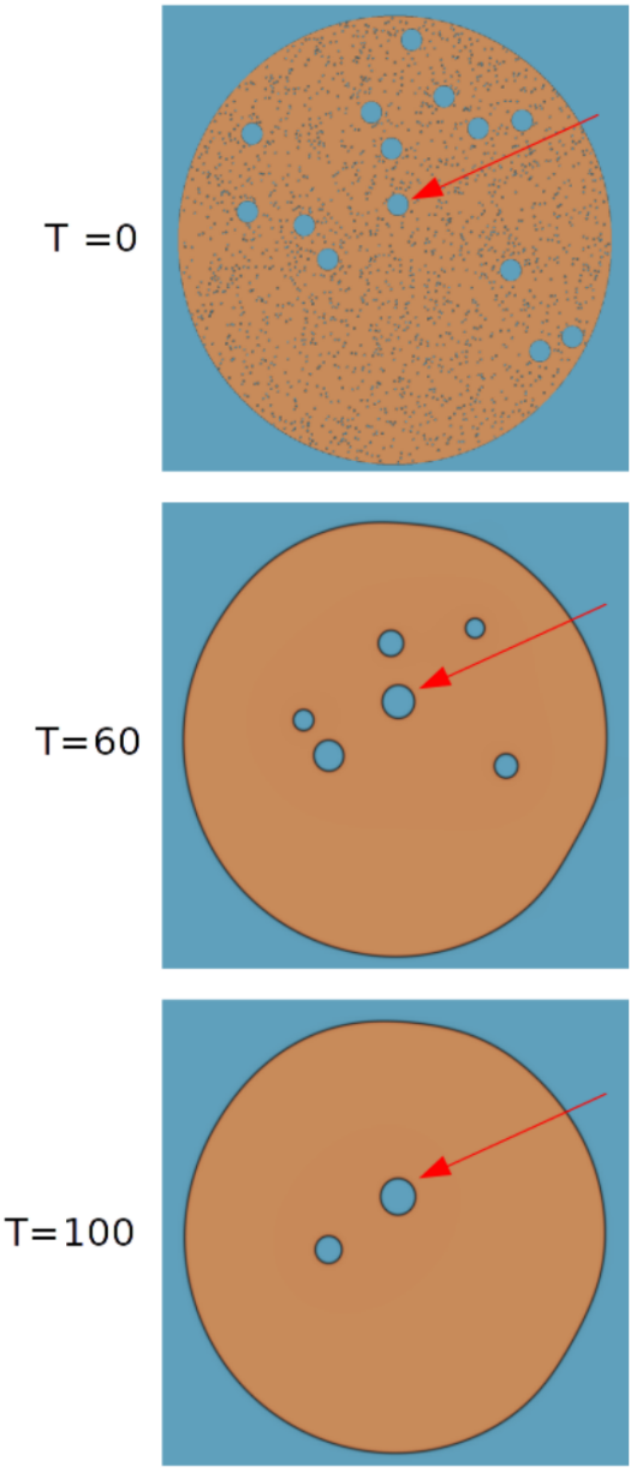

Figure 4.10: A large pore grows, partially as a result of free vacancies sourced from small pores. The evolution here is common in the bimodal configurations and is believed to contribute to their superior resistance to densification. Images from $TD = 92\%$ and $\chi = 60$.

$\chi = 100$ configuration experience rapid porosity loss across the configuration space. This is consistent with real-world behavior at annealing conditions; the elevated temperatures drive vacancy diffusion until eventually all pore space is lost. In the $\chi = 0$ case the behavior is the same, but the larger initial pore sizes mean that the interfaces are more energetically favorable than the smaller ones, resulting in more annealing time required to eliminate them. In fact, we observe that many systems at this χ value do not reach the fully dense state by the end of the simulation time. The bimodal configurations, best exemplified by the configurations with $\chi = 40$ and $\chi = 60$, in all cases outperformed the baseline monomodal structures ($\chi = 100$) in regards to their porosity loss rates.

At $TD = 94\%$, significantly larger error bars for $\chi = 0$ made comparison between it and the bimodal configurations difficult. The error observed is attributed to two primary effects. The first is concerned with the random placement of the voids, captured in two error parameters. The first of these parameters is the average spacing between pores (Figure 4.14). For the $TD = 94\%$ case, the variance in this parameter is greatest, and is believed to help contribute to the variation in observed shrinkage rate. Another parameter of interest is the average distance of pores to the surface of the fuel (Figure 4.15). Again, the $TD = 94\%$ cases exhibits the greatest variability.

While these parameters explain some of the noise exhibited in the data, the primary cause is believed to be the simulation size and corresponding surface effect. Because the diffusion path is relatively short, any deviation among configurations will have a strong influence on the overall kinetics of the system. To remedy this problem, more samples of the current size could be run, in the hopes of establishing a more consistent dataset. Also investigated is an increased system size, but hardware limitations require a reduced

parameter space. Given the current analysis, the trends are sufficiently resolvable among the configurations of greatest interest that additional runs were not strictly required.

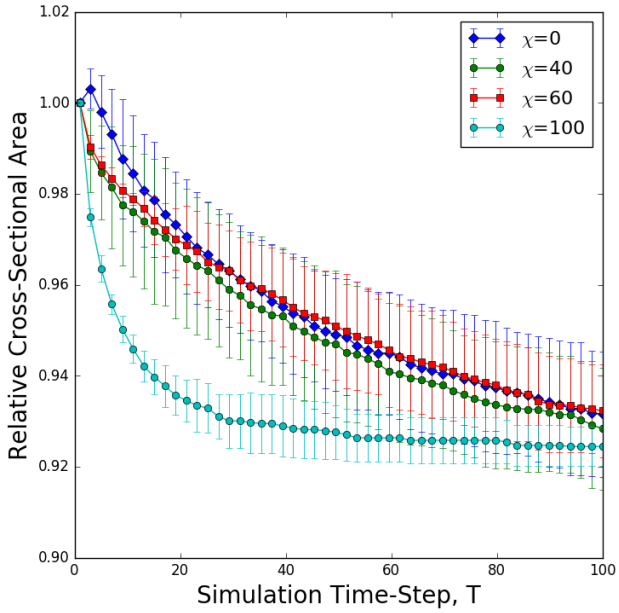

Figure 4.11: Area change rate at $TD = 92\%$. The $\chi = 100$ case experiences rapid loss of porosity in comparison with $\chi = 0$, owing to the smaller pore size. The intermediate cases, $\chi = 40$ and 60, follow the $\chi = 0$ behavior, although significant pore area is taken by the small pores.

4.3.2 Vacancy Concentration

In Figures 4.16, 4.17, and 4.18, the vacancy concentration throughout the entire system is averaged, excluding regions that are classified as pores. With this, we have a measure of the solution concentration of vacancy sites within the fuel body, and can make observations about the diffusion of these vacancies and their subsequent transport. Starting again with the $\chi = 100$ case, we see a rapid initial spike in the dissolution concentration, followed by a reduction to a small constant value as nearly all the vacancies migrate to the fuel's outer surface. At this point (around T=40 for most cases), the systems have reached

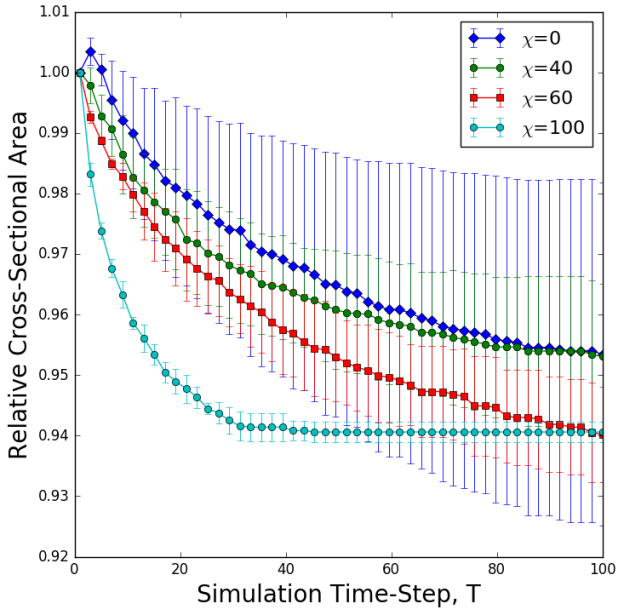

Figure 4.12: At $TD = 94\%$, there is significant noise in the $\chi = 0$ case. Trends for the other χ values at this TD are consistent with their analogues, however.

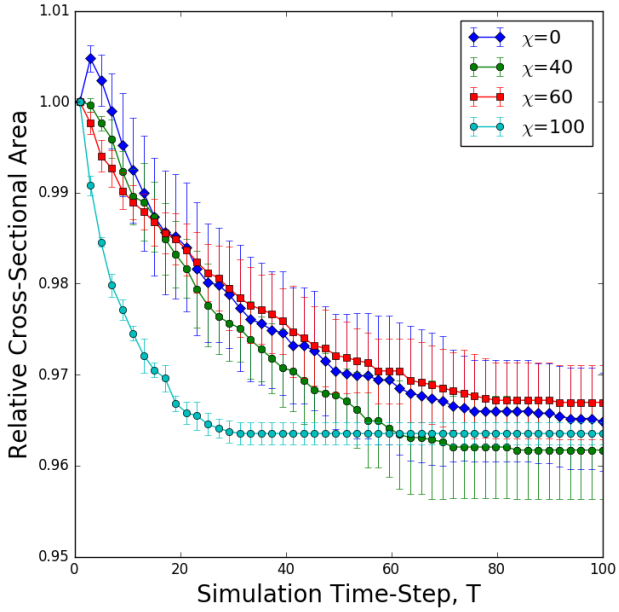

Figure 4.13: $TD = 96\%$ shows bimodal configurations that outperform the $\chi = 0$ case in long term behavior. Here, the $\chi = 100$ experiences the most rapid loss of porosity.

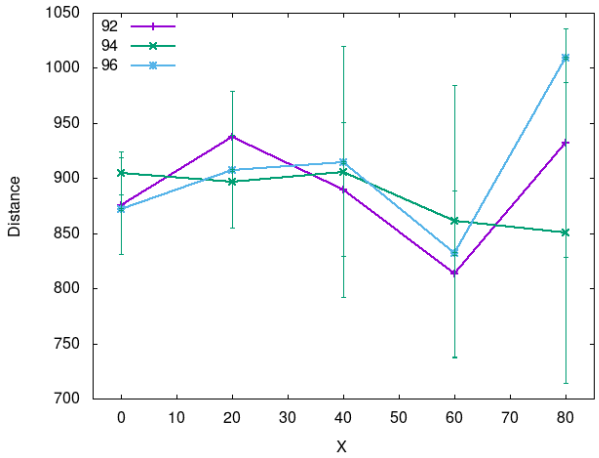

Figure 4.14: The interpore error metric, showing the average variation on the distance between large pores at initialization for each configuration. The variation shows little applicability to shrinkage kinetics, but highlights the $TD = 94\%$ case in that the greatest variability is seen here. This helps elucidate the noise seen in the shrinkage behavior for this case.

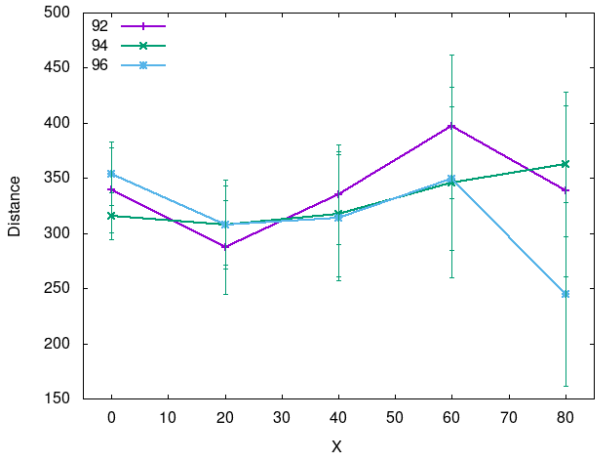

Figure 4.15: Here, the average distance of a large pore to the surface of the system is shown at initialization. Again, the greatest variance is seen for the $TD = 94\%$ case.

the fully dense state, with very few vacancies remaining.

This behavior is contrasted with the $\chi = 0$ case, where the sharp increase in concentration is subdued, and a longer "tail" exists where the vacancies remain dissolved within the fuel for some time. A spike is expected, of course, owing to the sharp interface created by the initialization algorithm. An interpretation of this behavior considers the surface of the pores at the initial state. High gradients drive the vacancy sites from the pore surface into the bulk fuel region. As these vacancies migrate away, the pore surface area decreases, resulting in progressively lower available reaction surface area, slowing pore shrinkage. As the system is evolved, more and more vacancies are removed from the pores, until with sufficient annealing time all are lost and the pores are destroyed. The fact that these large pores are capable of persisting for extended periods of time is noted.

At intermediate χ values, the advantage provided by the large pores is remarkable. Specifically, the bimodal configurations show spiking that is very similar to the magnitude experienced by $\chi = 0$, even though half of the pore area is occupied by small pore regions. As the small pores dissolve (a fact that is visually evident), the vacancies sourced from these regions are not immediately lost to the environment. The fact that the average vacancy concentration is higher suggests one of two behaviors: the vacancies congregate at the large pores, or quickly escape the fuel system. The compacts' exhibiting lower shrinkage rates seems to negate the second idea, and the frequent growth of the large pores supports the first theory. Furthermore, the long term behavior of the vacancy concentration is also similar to the $\chi = 0$ case. This indicates that considerable void space still exists within the fuel, and is being removed at considerably lower rates.

In all but one data series, the error observed here is minimal. The exception is with

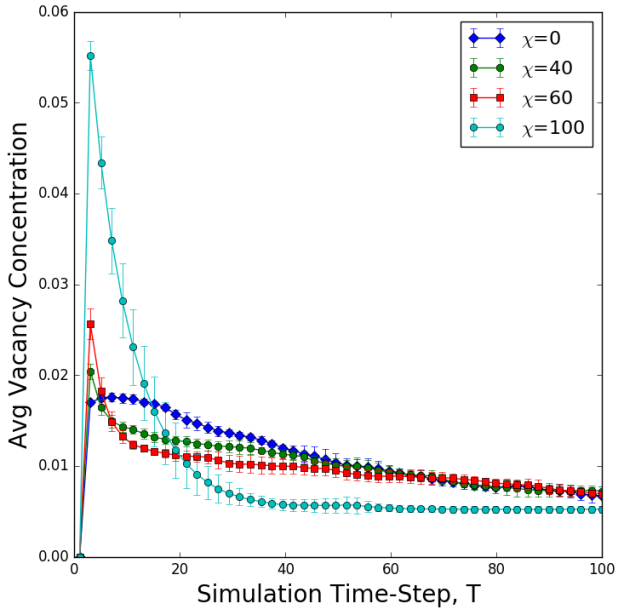

Figure 4.16: Measured in non pore regions, the average vacancy concentration spikes at early time steps in the $\chi = 100$ case, leading to large vacancy diffusion gradients and subsequently rapid transport. The effect is not as pronounced for the other configurations. The persistence of porosity is evidenced by the higher sustained vacancy concentration at extended time steps. Data shown for the $TD = 92\%$ case.

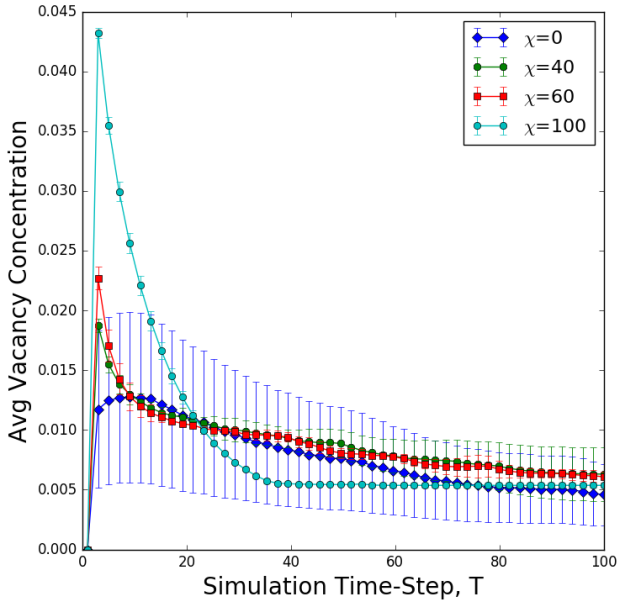

Figure 4.17: At $TD = 94\%$, high variance is again observed for the $\chi = 0$ case. The other configurations show good precision, however. The behavior otherwise follows the $TD = 92\%$ case.

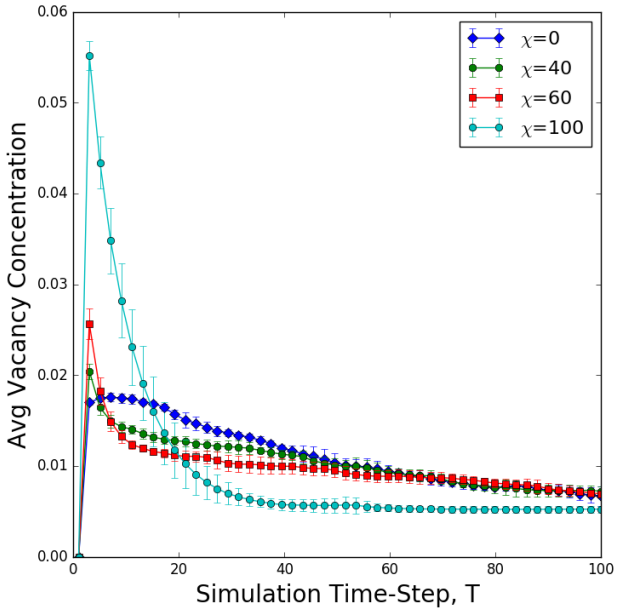

Figure 4.18: The bimodal configurations ($\chi = 40$ and 60) perform as good or better than the $\chi = 0$ case at $TD = 96\%$. Full densification occurs earlier than the other TD cases (when the plots reach an asymptote), owing to the low total initial porosity.

the $TD = 94$, $\chi = 0\%$ case. This error is again attributed to the random placement of the pores and small relative system size.

4.3.3 System Size Parameterization

All results shown to this point are for systems with the same Diameter, namely D = 1024 (dimensionless length units). We have stated previously that the size ratio between the large pores and the total system size is low (10^{-1}), and wish to explore other system sizes to understand if surface effects are contributing to our observation of the bimodal effect. Table 4.1 describes the configurations simulated; specifically, systems with sizes both halved and doubled in comparison to the original size are simulated.

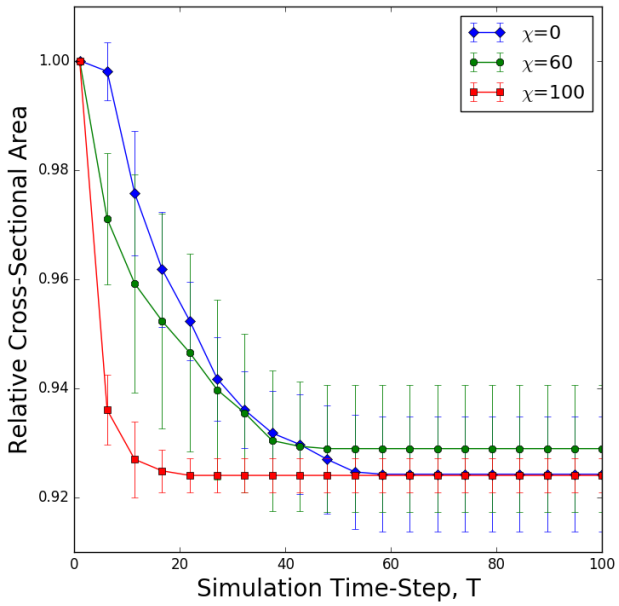

Figure 4.19: At a reduced system size of D = 512, rapid densification for all systems is observed. This is expected, as the average distance a vacancy must travel to escape the system is reduced. The bimodal configuration ($\chi = 60$), continues to follow the behavior of $\chi = 0$.

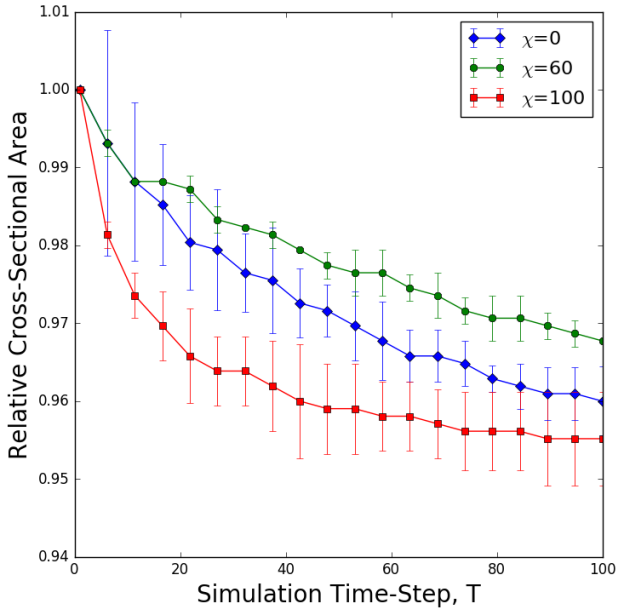

Figure 4.20: Doubling the system size (D = 2048) shows the bimodal configuration continuing to outperform the monomodal case. All systems exhibit slower kinetics than the base system size, owing to the much large diffusion path required.

System Size	Parameters
D = 512	TD=92, χ = 0,60,100
D = 1024	All configs.
D = 2048	TD=92, χ = 0,60,100

Table 4.1: Size parameterization space.

4.3.4 Discussion

Configurations with only small pores represent a microstructure that can be fabricated using standard sintering techniques. Densification in these configurations is rapid, both from experimental experience, and from the results of this simulation. The initial high total surface area of the pore-solid interface results in considerable excess interfacial energy. This imbalance drives the production of large diffusion gradients that enable rapid vacancy dissolution and subsequent transport to the surface of the body. Because regions of pore re-nucleation are relatively sparse (in terms of the area they subsume, rather than their frequency of occurrence), the diffusion pathway for vacancies is rather direct, and combined with large diffusion gradients results in rapid transport and the resulting densification.

With purely large pores, we again consider the location of vacancies, but must be reminded that most of them are located in the interior of the large pores. Thus, the quantity of vacancies available for diffusion is limited by surface area of the evolving pores. Furthermore, relative stability of the lower curvature interface results in pores with greater longevity. The corresponding total driving force given to the transport of the vacancies is considerably lower for the duration of the simulation, resulting in lower densification rates.

This result is also intuitive; larger pores are expected to require longer annealing times to eliminate. Occasional re-nucleation of intermediate pores or temporary growth of existing pores is observed, and further slows the densification process.

The bimodal cases of greatest interest are at $\chi = 40$ and 60. Here, we see that the small pores present in these systems are quickly eliminated, just as for the $\chi = 100$ case. The existence of large pores in these structures inhibits transport of their vacancies in two ways. In the first, vacancies that are bound for the surface of the system are often interrupted in their straight-line path out of the system. This has the effect of increasing the diffusion length required for their removal. Secondly, some vacancies are absorbed into large pores after being dissolved in the fuel matrix. Clearly, these vacancies will require much longer to complete their diffusion out of the system, as large pores are often stable for considerable time. We can see from Figures 4.16, 4.17, and 4.18 that at $\chi = 100$, a considerable spike in vacancy concentration is associated with the large diffusion gradients and rapid transport of the vacancies. These spikes are considerably subdued in the bimodal cases, even though significant total pore volume is occupied by the small pores. This fact, that the existence of large pores impedes the ability of vacancies sourced from small pores to freely diffuse within the system, is the key benefit of the bimodal pore structures simulated. The bimodal configurations also show increased values of vacancy concentration at extended simulation time. The result here is somewhat counter-intuitive: the existence of gradients, while detrimental at early time steps, is evidence of continued porosity transport at longer annealing time. The bimodal configurations are slower to eliminate these gradients, and thus densify over a much longer period of time than the $\chi = 100$ case.

The results of the size parameterization study are encouraging. The systems originally

simulated have on the order of 1000 total pores (in the case of $\chi = 100$), whereas realistic fuel pellets have millions of pores to reach the equivalent porosity. This discrepancy is a cause of concern in that the systems simulated may be exhibiting the improved densification behavior purely as a consequence of surface effects. From Figure 4.19 and 4.20, we continue to observe the bimodal configurations outperforming those with no porosity control. Particularly, we see no trend in the improvement as the system size is varied. Interestingly, the bimodal configuration outperforms the $\chi = 0$ case by a significant margin. The cause of this is not clear, although we caution that in the parameterization, fewer total runs are performed.

Fabrication routes for production of the bimodal configurations are under current research, as mentioned in the literature review. The elimination of small pores from the fuel compacts is difficult (barring samples that are already fully dense), but the addition of large pores through the pore forming technique can be deliberate. With this, we envision configurations analogous to the $\chi = 40$ or 60 case, and with the corresponding densification performance comparable to the values found in the results. The ability of the bimodal configurations to outperform the monomodal small configurations motivates further study of other configurations that may provide even greater benefit.

4.4 Validation

An important step in any study is validation of results. For physical experiment, quantification of error is necessary but not sufficient in testing a theory of some phenomena. Often, simulation is used to verify a model and ensure that the theory is sound, rather than just correct on the given data by chance. In simulation, the burden is perhaps greater, as the complex interactions that make up the physical world are greatly simplified into a model

that is computationally tractable. Thus, results on the computer, while they may be correct from the standpoint of a given model, may not extend into the real world if sufficient care is not taken in understanding that model's assumptions and applicability. In addition to the model being correct, we must also verify that a given program used in implementing a model does so honestly, that is, no programming errors must exist that give erroneous results. These two sources of error must be considered in any computational study.

The code correctness problem is well considered by the choice of MOOSE as the primary analysis tool. The finite element method is well developed, both mathematically and in software. The particular physics packages utilized from MOOSE are also validated by several external studies. In addition, the large community surrounding MOOSE and continual development on it provides considerable confidence that results generated by MOOSE are a faithful and accurate reproduction of the model described in the methods section, which is itself mathematically well supported.

Consideration of the model's applicability to the problem at hand is more difficult. Certainly, the Cahn-Hilliard equations have been applied with considerable success before. Some assumptions made in simplifying the annealing process are cause for concern, these assumptions are addressed in the discussion on future work. We are encouraged to consider the results as insightful, although perhaps not directly applicable to a specific fuel design. We are fortunate in that in tandem with this computational study, a fabrication effort is underway that seeks to reproduce our results in depleted-uranium fuel pins. Cooperation with this work has enabled both studies to advance in similar directions, which will considerably improve the ability to make direct comparisons and validate each group's results.

5 Conclusion

5.1 Summary

There is considerable interest in both academia and industry directed towards next generation nuclear power systems. These systems are expected to operate at lower cost, and to do so with a greater degree of reliability and safety. Additionally, they are expected to be free of some of the restraints imposed on current systems. Closing the nuclear fuel cycle, either in part or in whole, is a major goal, and key to this goal is the design of systems that enable significantly higher burnup. Metallic fuels exhibit an attractive combination of utility and efficiency, owing to their advantageous combination of high initial heavy metal density and excellent thermal conductivity in comparison to oxide fuel forms. This combination makes them a viable candidate to consider in a next-generation design.

Zirconium-alloyed uranium fuels have been used in research reactors in the past, and significant operating experience with them is possessed. Although good chemical compatibility with some reactor designs is enjoyed, their tendency to swell excessively at high temperature has made their adoption in power reactors problematic, with the primary mitigation strategy (increase in the fuel pin plenum size) leaving much to be desired. Further development is necessary in order to realize the greatest benefit of these fuels.

Modern material design is often driven by control of material morphology and behavior at the μm length scale, known as the microstructure. Control of microstructure features through novel fabrication routes, such as powder metallurgy and sintering, allows for precise specification of various material properties. In this work, we specifically target the size and

distribution of porous regions within the fuel. Previous work by [15] has experimentally demonstrated oxide fuel formulations with bimodal pore size distributions as exhibiting superior resistance to densification during sintering and subsequent extended annealing, partially reproducing the environment expected in high power reactor operation. This fact, if extensible to metallic fuel formulations, could provide a pathway to fuel formulations with specific microstructure specifications that alleviate the aforementioned limitation, and enable their use in production reactors.

We have developed a computational model capable of simulating the extended annealing of solid porous bodies that can be easily extended to specific metallic fuel formulations. The Cahn-Hilliard equation of the phase field family of methods provides a mathematical description of the annealing problem, and its implementation in the MOOSE finite element framework provides flexibility and scalability. Analysis tools utilizing Paraview's visualization capability were developed that are capable of quantifying the relative densification rates of a range of microstructure configurations, including monomodal configurations typical of standing casting and sintering techniques, and bimodal microstructures with intentional porosity added. A large parameter space is explored, with consideration of statistical variation included.

It is shown that bimodal configurations provide superior resistance to late stage densification in comparison to structures with monomodal small pores. The addition of large pores is thought to inhibit the transport of vacancies to the outer surface of the system, which follow diffusion pathways with the purpose of minimizing the surface free energy created by phase boundaries (fuel-pore interfaces). Without large pores, the vacancies are free to propagate to the fuel surface, and densification to near theoretical maximum levels occurs

quickly. Large pores obstruct these pathways, both extending the distance over which diffusion must occur and providing collection sites that are energetically favorable for temporary pore vacancy accumulation. While these enhanced pore regions are eventually eliminated, their persistence at extended annealing times is valuable in preventing rapid densification.

5.2 Future Work

Future work will address the influence of crystallographic orientation and size of grains within a metallic fuel configuration. The grain boundaries of the material are expected to have some effect on the diffusion pathways presented to vacancies, as it is common for the mobility of vacancies along grain boundaries to be considerably enhanced. Thermal gradients as a result of nuclear heat generation and convective cooling of the fuel pin surface will alter the kinetics of vacancy transport, primarily by adjustment of the effective mobility of the vacancies, but also considering compositional changes of the U-Zr alloy across the body cross section. The inclusion of a temperature gradient is a straightforward addition within the given MOOSE framework, but consideration of compositional changes and crystal orientation will require some development. Additionally, the inclusion of radiation-induced dislocation production will also act to modify vacancy mobility. Simulation of such behavior is possible within MOOSE, but its development is still limited. Fission gas production should be investigated through irradiation experiments. A follow on study for this, as well as physical irradiation of example fuel pellets is planned.

The work that has been described herein is part of a broader Multiscale study involving multiple members (Deo, Georgia Institute of Technology as P.I. 2, Mcdeavitt, Texas A&M as P.I. 3, Mariani, Idaho National Laboratory as P.I. 4). Microstructure parameters

that are realistic for a depleted-uranium alloy fuel will be provided by P.I. 2 in a continuing study to produce more realistic predictions. The fuel configurations of greatest interest will be fabricated and annealed by P.I. 3 for verification of results, and eventual irradiation with P.I. 4.

Bibliography

[1] H. Bailly, D. Mnessier, C. Prunier, and S. Simpson, *The nuclear fuel of pressurized water reactors and fast neutron reactors design and behaviour*. Lavoisier, 1999.

[2] N. E. Todreas and M. S. Kazimi, *Nuclear systems*. Taylor & Francis, 2011.

[3] G. K. J. Lanning D. D., Beyer C. E., *FRAPCON-3 updates, including mixed-oxide fuel properties*. The Commission, 2005.

[4] Y. e. a. Kim, "Thermal conductivities of actinides (u,pu,np,cm,am) and uranium-alloys (U-Zr, U-Pu-Zr, and U-Pu-TRU-Zr)." *Journal of Nuclear Materials*, vol. 445, p. 272, 2014.

[5] T. Sofu, "A review of inherent safety characteristics of metal alloy sodium-cooled fast reactor fuel against postulated accidents." *Nuclear Engineering and Technology*, vol. 47, no. 3, pp. 227–239, 2015.

[6] S. B. Walters L.C., Kittel J.H., "Performance of metallic fuels and blankets in liquid-metal fast breeder reactors," *Nucl. Technol.; (United States)*, vol. 65:2, May 1984.

[7] A. Solomon and S. Mcdeavitt, "Preparation of uranium-zirconium alloy nuclear fuel elements by pm," *Metal Powder Report*, vol. 47, no. 11, p. 55, 1992.

[8] W. D. Kingery and M. Berg, "Study of the initial stages of sintering by viscous flow, evaporation–condensation, and self-diffusion," *Sintering Key Papers*, pp. 367–382, 1990.

[9] D. L. Johnson and I. B. Cutler, "Diffusion sintering: I, initial stage sintering models and their application to shrinkage of powder compacts," *Journal of the American Ceramic Society*, vol. 46, no. 11, pp. 541–545, 1963.

[10] R. Raj and R. Bordia, "Sintering behavior of bi-modal powder compacts," *Acta Metallurgica*, vol. 32, no. 7, pp. 1003–1019, 1984.

[11] E. Arzt, "The influence of an increasing particle coordination on the densification of spherical powders," *Acta Metallurgica*, vol. 30, no. 10, pp. 1883–1890, 1982.

[12] H. Atkinson and S. Davies, "Fundamental aspects of hot isostatic pressing: An overview," *Metallurgical and Materials Transactions A*, vol. 31A, 2000.

[13] H. . Glsoy and R. M. German, "Sintered foams from precipitation hardened stainless steel powder," *Powder Metallurgy*, vol. 51, no. 4, pp. 350–353, 2008.

[14] S. Mcdeavitt and A. Solomon, "Hot-isostatic pressing of U-10Zr by grain boundary diffusion and creep cavitation. part 2: Theory and data analysis," Jan 1997.

[15] G. Maier, H. Assmann, and W. Drr, "Resinter testing in relation to in-pile densification," *Journal of Nuclear Materials*, vol. 153, pp. 213–220, 1988.

[16] W. J. A., D. R. Mumm, and A. Mohraz., "Microstructural tunability of co-continuous bijel-derived electrodes to provide high energy and power densities." *J. Mater. Chem. A*, vol. 4, no. 3, pp. 1000–007, 2016.

[17] R. B. Potts, "Some generalized order-disorder transformations," *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 48, no. 1, pp. 106–109, Jan 1952.

[18] V. Tikare, M. Braginsky, and E. A. Olevsky, "Numerical simulation of solid-state sintering: I, sintering of three particles," *Journal of the American Ceramic Society*, vol. 86, no. 1, pp. 49–53, 2003.

[19] M. Braginsky, V. Tikare, and E. Olevsky, "Numerical simulation of solid state sintering," *International Journal of Solids and Structures*, vol. 42, no. 2, pp. 621–636, 2005.

[20] M. P. Fix G.J., Ed. A. Fasano, *Free Boundary Problems: Thoey and Applications*. Pitman, 1983.

[21] L. J.S., "Models of pattern formation in first-order phase transitions," pp. 165–186.

[22] J. W. Cahn and J. E. Hilliard, "Free energy of a nonuniform system. i. interfacial free energy," *The Journal of Chemical Physics*, vol. 28, no. 2, pp. 258–267, 1958.

[23] L.-Q. Chen, "Phase-field models for microstructure evolution," *Annual Review of Materials Research*, vol. 32, no. 1, pp. 113–140, 2002.

[24] Y. U. Wang, "Computer modeling and simulation of solid-state sintering: A phase field approach," *Acta Materialia*, vol. 54, no. 4, pp. 953–961, 2006.

[25] Y. Wang, Y. Liu, C. Ciobanu, and B. R. Patton, "Simulating microstructural evolution and electrical transport in ceramic gas sensors," *Journal of the American Ceramic Society*, vol. 83, no. 9, pp. 2219–2226, 2004.

[26] M. R. Tonks, D. Gaston, P. C. Millett, D. Andrs, and P. Talbot, "An object-oriented finite element framework for multiphysics phase field simulations," *Computational Materials Science*, vol. 51, no. 1, pp. 20–29, 2012.

[27] D. Gaston, C. Newman, G. Hansen, and D. Lebrun-Grandi, "Moose: A parallel computational framework for coupled systems of nonlinear equations," *Nuclear Engineering and Design*, vol. 239, no. 10, pp. 1768–1778, 2009.

[28] A. Jokisaari, P. Voorhees, J. Guyer, J. Warren, , and O. Heinonen, "Benchmark problems for numerical implementations of phase field models," *Computational Materials Science*, vol. 126, pp. 139–151, 2017.

[29] X.-M. Bai, M. R. Tonks, Y. Zhang, , and J. D. Hales, "Multiscale modeling of thermal conductivity of high burnup structures in UO2 fuels," *Journal of Nuclear Materials*, vol. 470, pp. 208–215, 2016.

[30] Cenna, "Phase field order parameter."

[31] Chronholm144, 2006.

[32] "Moose framework," http://mooseframework.org/.

[33] D. Schwen, "Moose workshop: Moose-pf/marmot training," Nashville,TN, 2015.

[34] Ahrens, James, Geveci, Berk, Law, and Charles, "Paraview: An end-user tool for large data visualization, visualization handbook," pp. 1000–007, 2005.

[35] K. Asp and J. gren, "Phase-field simulation of sintering and related phenomena a vacancy diffusion approach," *Acta Materialia*, vol. 54, no. 5, pp. 1241–1248, 2006.

[36] R. L. Coble, "Sintering crystalline solids. i. intermediate and final state diffusion models," *Sintering Key Papers*, pp. 55–67, 1990.

[37] R. M. German, "Prediction of sintered density for bimodal powder mixtures," *Metallurgical Transactions A*, vol. 23, no. 5, pp. 1455–1465, 1992.

[38] G. L. Messing and G. Y. Onoda, "Sintering of inhomogeneous binary powder mixtures," *Journal of the American Ceramic Society*, vol. 64, no. 8, pp. 468–472, 1981.

[39] E. A. Olevsky, "Theory of sintering: from discrete to continuum," *Materials Science and Engineering: R: Reports*, vol. 23, no. 2, pp. 41–100, 1998.

[40] J. Pan, H. Le, S. Kucherenko, and J. Yeomans, "A model for the sintering of spherical particles of different sizes by solid state diffusion," *Acta Materialia*, vol. 46, no. 13, pp. 4671–4690, 1998.

[41] F. Wakai, M. Yoshida, Y. Shinoda, and T. Akatsu, "Coarsening and grain growth in sintering of two particles of different sizes," *Acta Materialia*, vol. 53, no. 5, pp. 1361–1371, 2005.

[42] O. Tange, "Gnu parallel - the command-line power tool," *;login: The USENIX Magazine*, vol. 36, no. 1, pp. 42–47, Feb 2011. [Online]. Available: http://www.gnu.org/s/parallel

[43] D. C. Crawford, D. L. Porter, and S. L. Hayes, "Fuels for sodium-cooled fast reactors: Us perspective," *Journal of Nuclear Materials*, vol. 371, no. 1-3, pp. 202–231, 2007.

[44] A. Moore, C. Deo, M. Baskes, and M. Okuniewski, "Atomistic mechanisms of morphological evolution and segregation in U-Zr alloys," *Acta Materialia*, vol. 115, pp. 178–188, 2016.

[45] G.-Y. Huang and B. D. Wirth, "First-principles study of bubble nucleation and growth behaviors in α U-Zr," *Journal of Physics: Condensed Matter*, vol. 24, no. 41, pp. 415–404, 2012.

[46] J. Rest, "An analytical study of gas-bubble nucleation mechanisms in uranium-alloy nuclear fuel at high temperature," *Journal of Nuclear Materials*, vol. 402, no. 2-3, pp. 179–185, 2010.

[47] S. Ahn, S. Irukuvarghula, and S. M. Mcdeavitt, "Thermophysical investigations of the uranium–zirconium alloy system," *Journal of Alloys and Compounds*, vol. 611, pp. 355–362, 2014.

[48] A. Sangjoon, I. Sandeep, and M. S. M., "Microstructure of α-u and δ-uzr2 phase uraniumzirconium alloys irradiated with 140-kev He ion-beam," *Journal of Alloys and Compounds*, vol. 681, pp. 6–11, 2016.

[49] D. B. Lancaster, R. L. Marsh, D. B. Bullen, P. Holger, S. C. Erwin, and A. E. Levin, "A technique for the selection of the fuel pin diameter for a uranium/zirconium alloy-fueled pressurized water reactor," *Nuclear technology*, vol. 97, 1992.

[50] E. S. Fisher, "The first 30 years of research on the physical properties of α-U," *Journal of Alloys and Compounds*, vol. 213-214, pp. 254–261, 1994.

[51] W. J. V. Sames, "Uranium metal powder production , particle distribution analysis , and reaction rate studies of a hydride - dehydride process." Texas A&M University.