

Fast kinetic Monte Carlo simulations: implementation, application, and analysis

by

Kristofer G. Reyes

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Applied and Interdisciplinary Mathematics)
in The University of Michigan
2013

Doctoral Committee:

Professor Peter Smereka, Co-Chair
Professor Joanna Mirecki-Millunchick, Co-Chair
Assistant Professor Henry Boateng
Professor Selim Esedoğlu
Professor Divakar Viswanath

© Kristofer G. Reyes 2013
All Rights Reserved

To my parents, my wife and daughter, for their constant love and support.

ACKNOWLEDGEMENTS

To my advisors, Peter and Joanna: thanks for your amazing guidance, support, encouragement and of course, the incredible amount of patience you have shown me. I have learned so much from you both – not only in the what appears in this dissertation, but in how to approach mathematics and science in general. As I progress to this new stage in life, I will constantly strive to emulate you both. Thanks for making these past few years rewarding and exciting.

I would like to thank the members of my committee: Professor Selim Esedoğlu and Divakar Viswanath. I would like to also acknowledge Tim Schulze, Arvind Basakaran, and Henry Boateng for their helpful conversations and insights. Thanks to Bennett Fauber for all the computer help and for looking the other way when I ran scripts on the screms login node.

I have learned a lot from my collaborators, and so they have my gratitude. Thanks especially to Denis Nothorn and Matt Dejarld.

I acknowledge the NSF for their funding support. This work was supported by NSF support grants DMS-0854870 and DMS-1115252.

Lastly, I would like to thank my family. My parents instilled in me a great curiosity of the world, and let me explore even at the detriment of various things around the house. I'm sorry for all the baking soda and vinegar you had to buy because of me. I cannot give enough thanks to my incredible wife , who has shown me limitless support, especially when I wavered and doubted myself.

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGEMENTS	iii
LIST OF FIGURES	vi
LIST OF TABLES	ix
LIST OF ALGORITHMS	x
CHAPTER	
I. Introduction	1
1.1 A brief overview of atomistic modeling	6
II. Kinetic Monte Carlo	13
2.1 The Kinetic Monte Carlo Model	13
2.1.1 Path connectedness	16
2.1.2 Liquid phase	18
2.1.3 Detailed balance	21
2.2 The KMC algorithm	22
2.3 Hash table based caching	25
2.3.1 Neighborhood locality	27
2.3.2 Hash table based rate caching	28
2.3.3 Finding optimal hash functions by simulated annealing	34
III. Application: Liquid Ga droplet epitaxy and crystallization	38
3.1 Liquid droplet epitaxy and crystallization	38
3.2 KMC model parameters	39
3.2.1 Calibration	41
3.3 Simulation results	44
3.3.1 Droplet epitaxy	47
3.3.2 Droplet crystallization	53
3.3.3 Core-shell structures	55
3.4 Morphological dependence on temperature and flux	61
3.4.1 Three fundamental processes in droplet crystallization	61
3.4.2 Quantum dots and nanorings	63
3.4.3 Nucleation of polycrystalline shells	66
3.4.4 Monocrystalline shells and a Mullins-Sekerka instability	67

IV. Further applications: nanowire growth by the Vapor-Liquid-Solid method and sintering of granular material	72
4.1 The Vapor-Liquid-Solid method for nanowire growth	72
4.1.1 Catalytic reaction rate	76
4.1.2 Liquid-solid interface mobility	77
4.1.3 Computational statistics	81
4.2 Sintering of granular material	81
4.2.1 Simulation Results	83
4.2.2 Computational statistics	90
V. Conclusion	92
5.1 Future work	94
BIBLIOGRAPHY	97

LIST OF FIGURES

Figure

2.1	An exchange event, along with the corresponding intermediate state. Atoms of species σ, τ (colored red and blue, respectively) exchange positions. The intermediate state is defined by replacing the exchanging atoms with an intermediary species $\iota = \sigma \wedge \tau$, colored purple.	15
2.2	An atom hop that results in a disconnected configuration.	17
2.3	A globally path-connected configuration on a hexagonal lattice that violates the local connectedness property. The gray line indicates a path between atoms x and y . The local neighborhood of atom B is outlined in red.	18
2.4	The ensemble average of N_{pattern} , the number of distinct neighborhood patterns within the atom configuration at time t during a nanowire simulation.	29
2.5	The ensemble average of C_{pattern} , the cumulative number of distinct neighborhood patterns seen up to step t during the nanowire simulation. The dashed red line indicates the linear asymptotic behavior of C_{pattern}	29
2.6	CPU time versus the number of MC steps performed during the nanowire simulations. Simulations without caching implemented are plotted in red squares. Green triangles show simulations where caching is used with an unoptimized hash function. Blue circles show simulation times when an optimal hash function is used.	33
2.7	The indexing of atoms in a neighborhood Σ of radius 2 and two-bit identifiers assigned to each species.	35
2.8	The number of initial conditions $C_{\mathcal{T}}(A)$ versus the number of iterations during the simulated annealing. Three independent trials are shown.	37
3.1	(Left) The Zincblende crystal structure. Obtained from [1]. (Right) The analogous 1+1 dimensional lattice used in the simulations. Nearest neighbor bonds are indicated by solid lines. Next-nearest neighbor bonds are shown as dashed lines. Ga and As atoms are colored red and green, respectively.	40
3.2	Substrate termination phase diagram as a function of deposition ratio and temperature, obtained from simulations and experiments. Those growth conditions resulting in a mostly Ga-terminated substrate are indicated by red squares, while green circles label As-terminated ones. The blue points above indicate the conditions where the transition from Ga to As termination occurred experimentally. The blue curve indicates the boundary between Ga and As-terminated given by equation 3.3.	45

3.3	Droplet epitaxy experimental results showing typical nanostructures observed over a range of substrate temperatures and As ₄ BEP.	46
3.4	Liquid Ga droplet diameter and linear density as a function of temperature.	48
3.5	The power law dependence of Ga droplet radius vs. $\frac{D_{\text{Ga}}}{\ell^2 F_{\text{Ga}}}$	49
3.6	Attachment and detachment events at the liquid-solid interface. This figure illustrates the attachment and detachment of an As atom in liquid Ga onto and from a perfectly flat liquid-solid interface, along with the intermediate state for the transitions. The intermediate species $\text{H} = \text{Ga} \wedge \text{As}$ is colored blue. The black lines indicate the relevant bonds that contribute to the change in energy between the initial and intermediate state.	51
3.7	Liquid droplets grown at $T = 200, 300, 350^\circ\text{C}$ and $F_{\text{Ga}} = 0.1$ ML/s. Here and throughout the paper, Ga and As atoms initially belonging to the substrate are colored red and green, respectively. Ga and As atoms deposited throughout the simulation are colored purple and blue, respectively.	52
3.8	Left panel: model snapshots of liquid droplet crystallization at various times for $T = 275^\circ\text{C}$ and $F_{\text{As}} = 0.06$ ML/s resulting in a compact quantum dot. Ga and As atoms from the original substrate are colored red and green, respectively. Ga atoms deposited to form liquid droplets are colored purple, while As atoms deposited during crystallization are blue. Right panel: AFM images of the GaAs growth fronts in partially crystallized droplets after 10 seconds (top), 40 seconds (middle) and 90 seconds (bottom). Crystallization was obtained at $T = 150^\circ\text{C}$ and 5×10^{-7} Torr As pressure.	54
3.9	Nanorings formed at $T = 375^\circ\text{C}$ and $F_{\text{As}} = 0.10, 0.20$ and 0.40 ML/s.	55
3.10	Simulation snapshots at times $t = 0, 1.8, 2.4$ seconds after crystallization, $T = 150^\circ\text{C}$ and $F_{\text{As}} = 0.8$ ML/s.	57
3.11	Simulation snapshots of a quantum dot annealed at high temperature at time t after temperature was increased. (A) The dot after exposure to As deposition at $F_{\text{As}} = 0.80$ ML/s and temperature $T = 150^\circ\text{C}$. This results in a polycrystalline GaAs shell trapping a liquid Ga core. (B)-(D) Temperature is increased to $T = 350^\circ\text{C}$ and the atoms rearrange in order to fully crystallize the liquid core.	59
3.12	Snapshots of liquid core formation at times $t = 1.4, 1.8$ and 2.0 seconds after crystallization at temperature $T = 350^\circ\text{C}$ and deposition rate $F_{\text{As}} = 1.0$ ML/s.	60
3.13	Morphological Dependence on Growth Conditions. This nanostructural phase map summarizes simulation results of droplet epitaxy and crystallization at various As deposition rates and temperatures. The three boundary curves indicate theoretically derived critical conditions delineating the simulation results and obtained in this section. The left-most, black line is given by Equation (3.18). The middle, red curve corresponds to Equation (3.35). The right-most, blue curve is given by Equation (3.23).	62
3.14	Schematic of kinetic processes that determine nanostructural development.	63

4.1	Simulation results of VLS nanowire growth, varying ρ^L between 0.75 and 1.25 eV. Liquid material is colored green, original solid substrate materials is purple. Material that was deposited from the vapor and subsequently reacted to solid phase is blue.	76
4.2	Simulation results of nanowire growth, for $\epsilon_{\text{ex}}^{\text{NL}} = 0.9, 1.10$ and 1.40 eV.	78
4.3	Snapshots of trajectories for solid atoms inside the liquid droplet for $\epsilon_{\text{ex}}^{\text{NL}} = 0.9, 1.1$ and 1.4 eV, over an interval of 0.1 seconds. Each straight line segment corresponds to 9.5×10^{-8} seconds.	79
4.4	An example of nanowire kinking.	80
4.5	Snapshots of nanowire kinking and facet nucleation. From left to right: (1) Start of nanowire kinking, with flat liquid-solid interface. (2) The nucleation of a new facet. (3) Growth of the facet and droplet motion. (4) Nucleation of another facet.	80
4.6	The initial configuration of the sintering simulations presented in this paper, obtained by a random sphere packing algorithm. The material density is 80%. Different colors correspond to different orientations.	84
4.7	Final configurations of sintering simulations after 100 seconds at $T = 327^\circ\text{C}$, for $\rho = 1.25, 1.50, 1.75$ and 2.00 eV. Different colors correspond to different orientations.	85
4.8	Grain coarsening statistics vs ρ . (Left) The number of grains vs. time. (Right) The mean disc-equivalent radius $\langle R \rangle$ vs. time. Best-fit power law models $\langle R \rangle = at^b + \langle R_0 \rangle$ are shown in dashed lines, and the fitted value b is indicated for each plot. Each plot corresponds to a different value of ρ	87
4.9	Long-term grain coarsening statistics after 10,000 simulated seconds for $\rho = 1.25$ eV and $T = 327^\circ\text{C}$. (Left) Final configuration. (Right) Number of grains vs. time, in the interval between 100 and 10,000 seconds.	87
4.10	Power law b vs ρ	89
4.11	Final configurations of sintering simulations after 100 seconds with $\rho = 1.25$ eV, for $T = 227, 327, 377$ and 427°C . Different colors correspond to different orientations.	91
4.12	Grain coarsening statistics vs temperature. (Left) The number of grains vs. time (initial configuration when $t = 0$ is not shown). (Right) Mean disc equivalent radius $\langle R \rangle$ vs. time. Best fit power-law models $\langle R \rangle = at^b + \langle R_0 \rangle$ are shown in dashed lines, and the fitted value b is indicated for each plot. Each plot corresponds to a different value of T	91

LIST OF TABLES

Table

3.1	Pairwise nearest and next-nearest neighbor bonding energies used in the GaAs system.	41
4.1	Pairwise bonding energies between the vapor, liquid and solid species used in the nanowire simulations.	74

LIST OF ALGORITHMS

Algorithm

1	The basic KMC main loop	23
2	The KMC algorithm for sampling the rate table by binary search.	26
3	The KMC algorithm for updating the rate table by propagating changes in leaf nodes upwards through their branches	26
4	Querying the rate cache by a linear, open-addressing hash table.	32

CHAPTER I

Introduction

Nanofabrication is the growth of nanometer to micrometer scale structures and has important applications in electronics and opto-electronic devices. Because of the short length scale, such structures are grown using atomistic processes that rely on thermodynamic properties of the materials used. Such processes include surface diffusion, adsorption and desorption. Depending on the relative rates of these processes, the resulting nanostructures can exhibit a broad range of phenomenology, leading to a corresponding range in the macroscopically observed electronic and optical properties of the prepared sample. Understanding the dynamic processes inherent in the fabrication of nanostructures at an atomistic resolution is therefore important in the development of materials with desirable properties. As such, the use of atomistic modeling and simulation compliments experimental, high resolution microscopy techniques such as transmission electron and atomic force microscopy. Moreover, atomistic modeling provides users a method of fine-tuning energetic parameters such as atomic bond strengths without reference to a specific physical material, allowing the study of material properties with respect to such parameters.

The difference in time and length scales between the atomistic processes outlined above and the mesoscales observed in structures of interest poses an inherent problem

in simulations. For example, atomistic surface diffusion of a single Ga atom has an associated diffusion length scale on the order of nanometers and a time scale on the order of microseconds [2]. As we shall see, such a process is important in the formation of GaAs nanostructures that are of order 100 nm in size and are grown over periods of minutes. Nanowires often achieve heights on the order of tens of micrometers over a period of hundreds to thousands of seconds [3–6]. The separation of scales imply that atomistic techniques must perform a large number of local events in order to simulate the self-assembly of these mesoscale structures. As such, it is important that the algorithms and implementations underlying any atomistic simulation be efficient.

The approach taken in this dissertation is Kinetic Monte Carlo (KMC), a technique generally regarded as well suited for mesoscale simulations, examples of which were described above. In the realm of simulation and modeling of physical systems, it sits between molecular dynamics, which handles small scale systems with a high degree of accuracy, and continuum methods, which are often used for larger scale systems at the expense of obscuring atomistic details. One of the first applications of KMC was in simulating the Ising model [7]. It has since been used to model a wide range of phenomena including the growth mechanisms of thin films [8–10], island nucleation and growth [2, 11], chemical vapor deposition [12], and catalytic oxidation [13]. Much work has been done incorporating elastic strain effects to simulate strained heteroepitaxial systems [14–18]. Off-lattice simulations include that of Al(100) crystal growth [19] and the modeling of 2D cluster diffusion [20]. Apart from applications in materials science, a similar method known as Dynamic Monte Carlo has been used to study protein folding [21, 22].

Because of its ability to simulate the formation of mesoscale structure efficiently,

KMC is being applied to increasingly complex systems. Examples of such systems are described in this work, where not only do we attempt to simulate even larger scale structures, but we introduce multiple species and multiple phases. We also depart from the traditional Solid-on-Solid (SOS) model, in favor of one that allows for the formation of more complicated morphologies. The central thesis of this dissertation is that KMC can be used to model a broad range of large scale structures in an efficient and flexible manner and as such, is a useful tool in the analysis of nanofabrication and materials science in general.

This dissertation presents an efficient KMC implementation and its application to three example systems. It illustrates how the KMC model can be used to analyze the resulting morphological structure in the context of KMC model parameters. As our primary example, we consider liquid Ga droplet epitaxy and crystallization by an As flux. There, Ga atoms are deposited on a GaAs substrate, and nucleate liquid droplets. The droplets are then exposed to As and crystallize into GaAs nanostructures. We show how parameters are calibrated in order to capture a range of phenomena, including correct surface termination, droplet size and density, droplet etching and the parametric dependence on the resulting nanostructure. As an extension to the KMC model, a unifying analytic model is developed to describe the critical conditions for the existence of nanostructure.

As further illustration of our implementation's ability to handle a broad range of phenomena, we present two more example applications. The second example describes simulations of nanowire growth by the Vapor-Liquid-Solid (VLS) method. In VLS, a liquid droplet on a solid substrate is exposed to a vapor. Atoms in the vapor phase are incorporated into the droplet and undergo a catalyzed vapor to solid reaction. The resulting solid atom then precipitates onto the liquid-solid interface,

resulting in nanowire growth. We study the effect of energy parameters that affect the rates of vapor to solid reactions and atom mobility at the liquid-solid interface. Simulations show a variety of nanowire morphologies, including broad, tapered bases and nanowire kinking. Atomic trajectories are also studied to gain insight on the growth modes leading to the observed morphologies. As a third application, we present simulations of sintering, by which we mean the evolution of porous granular material. In this system, grains evolve by atomistic changes in orientation at grain interfaces and by surface diffusion along grain-pore interfaces. We show how the KMC treatment of this system is similar to that of other atomistic grain growth models such as the Monte Carlo Potts model, but incorporates both orientation changes and surface diffusion within a proper kinetic framework. The amount of grain coarsening is then studied as a function of temperature and the energy barrier for orientation changes.

The outline of the dissertation is as follows. Chapter II introduces the underlying KMC model and algorithms. It details the basic notions of activation energies and rates underpinning KMC as a continuous time Markov chain (CTMC). Section 2.1.1 motivates and details the alternative to the SOS constraint used throughout the simulations. Section 2.1.2 details how the liquid phase is identified and treated in order to capture reasonable behavior with our on-lattice model. Section 2.1.3 describes the principle of detailed balance in the context of CTMC, and describes how activation energies are selected in order to satisfy this requirement. In section 2.2, we explore the underlying algorithms of KMC and describe efficient methods for sampling and updating the Markov chain. Section 2.3 details an important aspect to our specific KMC implementation: using a rate cache to eliminate redundant rate calculations. We motivate its use, and show how its implementation leads to

a significant performance gain. Key to the efficiency of such a rate cache is the selection of a system-specific hash function, In Section 2.3.3, we describe a simulated annealing search for optimal hash functions.

In Chapter III, we turn our attention to the first application of our KMC implementation – Ga droplet epitaxy and crystallization. We describe the choice of model parameters in Section 3.2, including simulations of GaAs homoepitaxial thin film growth in Section 3.2.1. Droplet epitaxy simulation results are presented in Section 3.3. There, we study the nanostructural dependence on growth conditions. We present experimental results exhibiting this dependence, along with corresponding simulations (Section 3.3.1). The model can capture the broad range of nanostructures observed in the experiments with the correct qualitative dependence on growth parameters. Moreover, simulations predict the presence of Ga/GaAs core-shell structures, which are difficult to observe unambiguously in experiments. This is described in Section 3.3.3. The existence of both monocrystalline and polycrystalline shells suggest two independent mechanisms for their formation: a morphological instability of the crystallization front or nucleation at the vapor-liquid interface. We show by simulation that nucleation-induced shell structures may be recrystallized into fully crystallized GaAs islands by annealing at high temperature. In the case of the instability-driven shell formation, we provide simulation evidence to suggest the presence of a Mullins-Sekerka instability. Lastly, we develop a fully analytical model (Section 3.4) that describes the existence of the structures observed and their dependence on growth conditions. The theoretical model agrees well with simulation and experimental results.

In Chapter IV, we present simulations of the two other example systems we consider in this work: nanowire growth and sintering. Nanowire simulation results are

presented in Section 4.1, which details a study of the effect of energy parameters on nanowire growth modes. Lastly, Section 4.2 details the sintering simulations. There, we measure the effect of energy parameters and temperature on coarsening statistics. Critical coarsening phenomena is observed in Section 4.2.1.

Sections of this dissertation appear verbatim in [23] and [24].

1.1 A brief overview of atomistic modeling

Before we describe kinetic Monte Carlo, it is instructive to provide a overview of atomistic modeling in general. Doing so motivates the specific components of KMC and illustrates why KMC is appropriate for the systems we wish to study. Among the primary techniques in atomistic modeling is Molecular Dynamics (MD). In MD, atoms occupy positions in \mathbb{R}^3 with no preference to specific points in space (as opposed to lattice-based techniques, which are described below). Atoms interact with each other via an atomic potential $\Phi(\mathbf{x}_1, \dots, \mathbf{x}_N)$ that measures the total potential energy of atoms occupying positions $\mathbf{x}_1, \dots, \mathbf{x}_N \in \mathbb{R}^3$ in space. Often the potential Φ is expressed as the sum of n -body interaction potentials $\phi(\mathbf{x}_1, \dots, \mathbf{x}_n)$, where $\mathbf{x}_1, \dots, \mathbf{x}_n$ are a collection of the n atom positions in the system. Typically, n is a small number. When $n = 2$, the potential $\phi(\mathbf{x}_1, \mathbf{x}_2)$ is called a pair potential, and the total potential energy is obtained by summing over every pair of atoms in the system:

$$(1.1) \quad \Phi(\mathbf{x}_1, \dots, \mathbf{x}_N) = \frac{1}{2} \sum_{\mathbf{x}_1} \sum_{\mathbf{x}_2} \phi(\mathbf{x}_1, \mathbf{x}_2).$$

An example of a typical pair potential is the Leonard-Jones potential:

$$(1.2) \quad \phi(\mathbf{x}_1, \mathbf{x}_2) = \epsilon \left[\left(\frac{\sigma}{r} \right)^{12} - \left(\frac{\sigma}{r} \right)^6 \right],$$

where

$$r = \|\mathbf{x}_1 - \mathbf{x}_2\|_2,$$

and ϵ, σ are parameters. This potential depends only on the distance r between the two interacting atoms and consists of a repulsive term $(\frac{\sigma}{r})^{12}$ that is dominant for small r and an attractive term $(\frac{\sigma}{r})^6$ dominant when r is large. The values of ϵ and σ determine the strength of the repulsion and attraction, and the value of the equilibrium length that minimizes the potential energy. Using pair or higher order n -body potentials is computationally expensive *a priori*, as it costs $\mathcal{O}(N^n)$ to sum over every $\binom{N}{n}$ subsets $\{x_1, \dots, x_n\}$ of size n . However, many potentials depend solely on the distances between atoms, and such dependence may be negligible over longer length scales. As such, in practice computational speed-ups may be gained by employing cut-off or nearest-neighbor techniques in which interactions between atoms sufficiently far apart are neglected. For example, for the Leonard-Jones potential above, auxiliary bookkeeping may be made so that every atom position \mathbf{x}_i is associated with a list \mathcal{N}_i containing the location of every other atom within a fixed radius. Then the summation in Equation (1.1), which takes quadratic work, may be replaced with one in which the inner summation is no longer over every atom in the system, but over the lists \mathcal{N}_i . Assuming the size of such lists are constant with respect to N , the computation of the potential energy (1.1) would be linear in the system size.

Given an initial configuration of atom positions $\mathbf{x}_1, \dots, \mathbf{x}_N$ and atom velocities $\mathbf{v}_1, \dots, \mathbf{v}_N$, an MD simulation proceeds by integrating

$$m\dot{\mathbf{v}}_i = \mathbf{F}_i$$

$$\mathbf{v}_i = \dot{\mathbf{x}}_i$$

for force term \mathbf{F}_i . The force term is often derived from Newtonian dynamics, in

which forces arise from gradients of the potential energy:

$$\mathbf{F}_i = -\nabla_{\mathbf{x}_i} \Phi,$$

or Langevin dynamics:

$$\mathbf{F}_i = -\nabla_{\mathbf{x}_i} \Phi - \gamma M v_i + \sqrt{2\gamma k_B T m} R(t),$$

where k_B is Boltzmann's constant, T is temperature, γ is a damping factor, and $R(t)$ is a Gaussian process. The advantage of the latter is its ability explicitly control temperature.

The graph of $\Phi(\mathbf{x}_1, \dots, \mathbf{x}_N)$ over the high-dimensional configuration space \mathbb{R}^{3N} is called the potential energy surface. For systems with a large number of particles, the topography of this landscape is quite complicated, and consists of several local minima separated by saddle points. The trajectories of atoms in \mathbb{R}^3 lead to a corresponding trajectory on the potential energy surface

$$\Phi(t) = \Phi(\mathbf{x}_1(t), \dots, \mathbf{x}_N(t)).$$

An examination of this trajectory during a typical MD simulation illustrates a main disadvantage of the MD technique. As a simulation progresses the trajectory $\Phi(t)$ vibrates within a basin of a local minimum of the energy surface. Macroscopically, such vibrations do not qualitatively differ from the configuration described by the local minimum. Often, we are interested in events in which the trajectory crosses over a saddle point of the energy surface, indicating a transition from one basin to another. Such events are rare events in MD, so an MD simulation must expend a large number of time steps in order to witness a qualitatively significant transition. Typical MD simulations use a time step of one femtosecond (10^{-15} s), while the time scale for a process such as e.g. Ga diffusion on GaAs is on the order of microseconds

(10^{-6} s) so that MD simulations would require billions of time steps in order to simulate the diffusion of an adatom over its characteristic length scales. Because of the computational burden associated to MD, this technique is usually restricted to simulating systems with a small number of atoms on the time scale of nanoseconds.

One way to address this disadvantage is to explicitly model the transitions between basins, viewed now as states in which the corresponding local minimizers are seen as representative configurations for the basins that they occupy. The correct rate constants at which these transitions occur is dictated by transition state theory. The rate constant $r(X \rightarrow Y)$ describes the transition from state X to Y and is given by the equilibrium flux (under the Boltzmann distribution) from state X through the saddle point connecting X and Y . Such constants are of Arrhenius form:

$$(1.3) \quad r(X \rightarrow Y) = \nu_{X,Y} \exp \left[\frac{E(X \wedge Y) - E(X)}{k_B T} \right],$$

where $\nu_{X,Y}$ is a rate pre-exponential factor and $X \wedge Y$ is the intermediate state described by the saddle point. The energy $E(X)$ is the energy of the local minimizer corresponding to state X , and can be obtained from any configuration within that minimizer's basin by gradient descent. The energy $E(X \wedge Y)$ represents the energy of the saddle point.

The state space \mathcal{X} identified with the set of local minimizers along with the transition rates $r(X \rightarrow Y)$, for $X, Y \in \mathcal{X}$ and the assumption that the evolution process is Markovian results in a continuous time Markov chain (CTMC). A realization of this CTMC leads to a straight forward simulation of the rare event transitions between qualitatively different atomic configurations, provided that a list of states accessible from any particular state or a list of saddle points and their respective energies can be obtained. This realization of the CTMC is Kinetic Monte Carlo (KMC) in its most general form. Given a state X , an accessible state Y is sampled according to the

rates $r(X \rightarrow Y)$ and the simulation time is incremented by the expected minimum waiting time from state X

$$dt_X = \frac{1}{\sum_Y r_{X,Y}},$$

assuming exponentially distributed waiting times, which is implied by the Markov property. This process is iterated until the desired simulation time is achieved.

Enumerating saddle points, and hence allowable transitions, and their respective energies poses a difficulty that was not present in MD techniques. Saddle point finding methods must be therefore employed. Common techniques to do so include the nudged elastic band method [25] and the dimer method [26]. Saddle point methods are computationally expensive. As an alternative to searching for saddle points, explicit transitions can be instead *defined a priori* via local, atomistic events. These local events alter a configuration in a small volume of space, keeping the rest of the configuration static. For example, in KMC a typical process is the diffusion of an adatom on the surface of a film. The local, atomistic event that corresponds to this process is an atom hop, the movement of an atom a small distance in a particular direction. The intermediate state for this event can be defined by removing the diffusing adatom. While this approach gives up generality by requiring the *a priori* specification of events, the computational speed-up gained by avoiding saddle point searches makes it attractive in the context of simulating large-scale, complex problems.

One outstanding problem in this presentation of KMC is that the energies of the local minimizer and of a saddle point are obtained indirectly by gradient descent. This too is computationally expensive when performed in the high-dimensional configuration space. One variant that attempts to address this is bond-counting, on-lattice kinetic Monte Carlo (BC-LKMC). In BC-LKMC, it is assumed that local

minimizers correspond to configurations in which atoms reside on a crystallographic lattice. Each atom is considered bonded to its neighbors on the lattice, and each such bond is assigned a bonding energy. The total energy of a configuration is the sum of the bonding energies within the system. Transitions between configurations are restricted to local events in which atoms are allowed to move only between neighboring lattice positions or the species of a single position is altered. Intermediate states are also placed on lattice, and the energy of these states is calculated in an identical manner. Due to the local nature of events and the direct representation of a configuration's energy in terms of atom bonds, the energy difference in the transition rate (1.3) can be obtained in constant time by identifying the small number of bonds that are different between an initial and intermediate configuration. This leads to an reduced computational burden by avoiding the gradient descent computations necessary if the energy is not so directly calculated. Therefore, the main advantages of BC-LKMC are:

1. a discretization of state space corresponding to local minimizers of the potential energy landscape;
2. a modeling of only rare event transitions between these local minimizers;
3. an explicit enumeration of the allowable transitions as local changes of the lattice occupancy pattern;
4. a local calculation of energy differences;

We have thus presented the development of BC-LKMC as a sequence of steps taken to provide a computationally efficient simulation that captures the relevant underlying physics behind its motivation. More can be said about further computational considerations in BC-LKMC, including the choice of appropriate data

structure used to sample events, as is described in Section 2.2. Indeed, much of the work in developing an efficient BC-LKMC algorithm is in this choice [27, 28]. Often overlooked, however, is the calculation of event rates in Equation (1.3). Albeit constant-time, this calculation must be performed several times each Monte Carlo step, and contributes a significant portion of floating point operations. As such, rate calculation is often a considerable bottle neck in a BC-LKMC implementation. A major focus of this work is in addressing this problem.

CHAPTER II

Kinetic Monte Carlo

2.1 The Kinetic Monte Carlo Model

The underlying model used throughout this work is an example of BC-LKMC described above and hereafter simply referred to simply as KMC. Most generally, atoms occupy positions on a graph $G = (\Lambda, E)$, typically embeddable in $S^1 \times \mathbb{R}$ for the 1+1 dimensional simulations with periodic boundary conditions presented in this work. Positions are either unoccupied or occupied with one of S atomic species, $\sigma_1, \dots, \sigma_S$. For example, in the GaAs system the species represent vacuum (unoccupied), Ga or As, while in the nanowire simulations, the different species correspond to vacuum, vapor, liquid and solid phases. In the sintering simulations, the different species correspond to different orientations of the sintered material. Edges in the graph represent nearest-neighbor bonds between atoms. A secondary graph $G_{nn} = (\Lambda, E_{nn})$ may also be specified in which the edges E_{nn} represent next-nearest neighbor bonds. This naturally extends to tertiary or higher-order bonding by the further specification of more edges. Throughout this work, shall only consider those graphs that are crystallographic lattices with periodic boundary conditions.

One of five atomistic events can occur at each step of the simulation. Atoms may desorb from the surface, adsorb onto vacant positions on the surface, diffuse

along the surface (modeled by the exchange of an atom with a vacant position), exchange species with a neighboring occupied position (an atom-atom exchange) or undergo a reaction that modifies an atom's species. All events alter a configuration by changing the species of at most two lattice positions, affecting a small number of local neighborhoods in which the two positions belong.

An energy $E(X)$ is obtained from an atom configuration X by bond counting. Bonds between neighboring atoms x and y are assigned an energy $\gamma(\sigma(x), \sigma(y))$ that depends solely on the species $\sigma(x), \sigma(y)$ of the atoms. The pairwise energies $\gamma(\cdot, \cdot)$ are specific to the system to be simulated, and are treated as parameters. The energy $E(X)$ of an entire configuration X is the sum of all bonds within that configuration

$$E(X) = \frac{1}{2} \sum_{x \in X} \sum_{\substack{y \in X \\ y \text{ a neighbor of } x}} \gamma(\sigma(x), \sigma(y)).$$

All event transitions $X \rightarrow Y$ between states X, Y are assigned an activation energy of the form

$$(2.1) \quad E_a(X, Y) = E(X \wedge Y) - E(X) + \epsilon(X, Y),$$

where $X \wedge Y$ is a well-defined intermediate state and $\epsilon(X, Y)$ is an *additional energy barrier*. As will be discussed momentarily, the additional energy barrier is event and neighborhood specific. Rates for the transition $X \rightarrow Y$ are given by

$$(2.2) \quad r(X \rightarrow Y) = R_0 \exp \left[-\frac{E_a(X, Y)}{k_B T} \right],$$

which are of the form in Equation (1.3) in which

$$\nu_{X,Y} = R_0 \exp \left[-\frac{\epsilon(X, Y)}{k_B T} \right],$$

where $R_0 = 10^{13} \text{ s}^{-1}$ is a constant prefactor. Each of the five events are assigned a method for determining an intermediate state as well as an additional energy barrier.

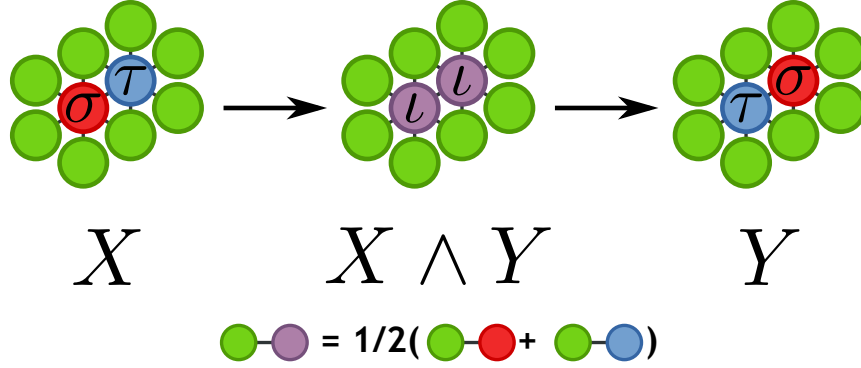


Figure 2.1: An exchange event, along with the corresponding intermediate state. Atoms of species σ, τ (colored red and blue, respectively) exchange positions. The intermediate state is defined by replacing the exchanging atoms with an intermediary species $\iota = \sigma \wedge \tau$, colored purple.

For example, when the transition $X \rightarrow Y$ represents an atom hopping to an adjacent, vacant position (representing the atomistic analog of surface diffusion), the additional energy barrier $\epsilon(X, Y)$ is set to zero, while the intermediate state is the one obtained by removing the diffusing atom from the system. As such, the difference in bonding between X and $X \wedge Y$ is precisely the bonds attached to the diffusing atom. Therefore, the activation energy $E_a(X, Y) = E(X \wedge Y) - E(X)$ represents the bonds broken in order to perform the hop. The case of two exchanging atoms is depicted in Figure 2.1. There, the two atoms have species σ, τ , respectively. The intermediate state is defined by replacing the two exchanging atoms with atoms of an intermediate species, denoted $\sigma \wedge \tau$. To calculate the energy $E(X \wedge Y)$ of the intermediate state, bonding energies between the intermediate species and any other atom are obtained by averaging:

$$(2.3) \qquad \gamma(\cdot, \iota(\sigma, \tau)) = \frac{1}{2}(\gamma(\cdot, \sigma) + \gamma(\cdot, \tau))$$

For this exchange event, the additional energy barrier in the form (2.1) describes the extra barrier necessary to exchange two atoms and will be denoted ϵ_{ex} throughout this paper.

In the case of a reaction event in which an atom of species σ changes to one of species τ , a similar intermediate state can be defined by replacing the reacting atom with one of intermediate species $\sigma \wedge \tau$. Bonding energies between this intermediate species and any other atom is computed as in Equation (2.3). The additional energy barrier for a reaction is denoted as ρ throughout this paper.

Remark II.1. Stronger bonds are energetically advantageous, and should result in a lowering of the energy. Therefore, bonding energies $\gamma(\cdot, \cdot)$ are regarded as negative values. However, the energies reported in this dissertation are presented as positive numbers as a matter of convenience. In order to get the correct positive energy difference $E(X \wedge Y) - E(X)$, this quantity is correspondingly negated in subsequent calculations. That is, we maintain the convention that $\gamma(\cdot, \cdot) \geq 0$ and compute an activation energy

$$E_a = E(X) - E(X \wedge Y) + \epsilon(X, Y)$$

whenever using the bonding positive energies.

2.1.1 Path connectedness

In the systems we consider, many complex morphologies and growth mechanisms are observed in experiments. For example in nanowire growth, experimentally observed phenomena such as diffusion along the length of the wire and nanowire kinking (a change in growth direction) must be captured in our model. As such, the SOS constraint, which dictates that atoms must be positioned directly above other atoms, is insufficient. An alternative constraint must be made, however, in order to prevent the unphysical detachment of “molecular satellites” from the main system. This is illustrated in Figure 2.2, in which an atom hop results in a disconnected configuration. Atom configurations are hence required to be *path-connected* in the graph




Figure 2.2: An atom hop that results in a disconnected configuration.

theoretic sense. That is, there must exist a chain of atoms through neighbor bonds between any pair of atoms in a configuration. Configurations containing overhangs and vacancies in the bulk are therefore possible, unlike in the SOS model, allowing our model to capture a broader range of morphologies. Any event resulting in a disconnected neighborhood is disallowed by assigning the rates of such events to be zero.

Checking the connectedness property is a non-trivial task, though it is well known how to do so. For example, it can be verified by a simple breadth-first traversal of the subgraph $G_{\text{occupied}} = (\Lambda_{\text{occupied}}, E_{\text{occupied}})$ consisting of those vertices $\Lambda_{\text{occupied}} \subseteq \Lambda$ that are occupied by an atom and the corresponding edges $E_{\text{occupied}} \subseteq E$ in which they belong. Alternatively, it can be determined from the zeroth Betti number $b_0(G_{\text{occupied}})$ of the graph G_{occupied} , which may be calculated by an eigenvalue computation of the adjacency matrix $A(G_{\text{occupied}})$.

The above methods for determining connectedness illustrate that the property is inherently global and is hence infeasible to check efficiently. We approximate it by instead requiring connectedness on small, local neighborhoods (or subgraphs) about every atom. This is an approximation of global path-connectedness, though local connectedness is not a necessary condition for the corresponding global property. This is illustrated in Figure 2.3. Atoms x and y are globally connected, but are not




Figure 2.3: A globally path-connected configuration on a hexagonal lattice that violates the local connectedness property. The gray line indicates a path between atoms x and y . The local neighborhood of atom x is outlined in red.

locally connected within the neighborhood of x . The size of the local neighborhood used to check this constraint affects the structures seen in the simulations. This size should be selected large enough to capture the relevant structures for the specific application being simulated. In the extreme case, the size of a local neighborhood approaches the entire domain, capturing global connectedness as a special case. Typically, the local neighborhood $B_L(x)$ about an atom x is taken as the set of atoms with graph distance less than some small number, L

$$B_L(x) = \{y \in \Lambda : d(x, y) \leq L\},$$

where $d(x, y)$ is the graph distance on G .

2.1.2 Liquid phase

As described in the introduction, we wish to simulate systems in which there is a liquid phase. In both the GaAs and VLS nanowire systems, nanostructural growth is mediated by the liquid phase, and so modeling reasonable physics in the liquid is important to obtaining accurate simulations. Therefore, care must be taken in order to capture the correct behavior in the liquid phase, and a first step in doing so is in identifying those regions that constitute the liquid. This is especially im-

portant in the GaAs system. There, Ga plays two roles: first being a constituent in the GaAs crystal, treated as a solid, and second, Ga atoms make up liquid Ga droplets. To crystallize the droplet, an As atom must diffuse through it and nucleate at the liquid-solid interface. Atomistically, As diffusion through the droplet is a sequence of atom-atom exchanges. Such exchanges may also appear during the attachment/detachment of atoms at the liquid-solid interface, but these events are physically different from As diffusion through the droplet. At the event level (e.g. atom-atom exchange), the model cannot distinguish between those events that occur in the liquid phase and those that do not. Additional, neighborhood-specific information must be incorporated into the model in order to address this.

The model identifies liquid neighborhoods, and applies different additional energy barriers for events occurring within such neighborhoods. While it is not difficult to identify neighborhoods in the bulk of their respective phases – e.g. any neighborhood of sufficient size consisting of only Ga atoms constitutes the “liquid” phase, while Ga and As arranged in a crystalline pattern constitute the “solid” phase – it is not so clear cut when neighborhoods contain a mixture of material, as in the case of regions near liquid-solid interfaces. At an atomistic level, the geometry of such interfaces is complex and is often far removed from the idealized sharp interface that delineates the bulk of solid and liquid phases. The technique for identifying liquid neighborhoods must be robust with respect to this complication. In the model, we use a simple population count. Within a neighborhood, the number of atoms that constitute a liquid phase (e.g. Ga in the GaAs system) are tallied. If the number of such atoms is more than some threshold of the neighborhood volume, then the neighborhood is deemed liquid. This method is impartial to the geometry of the interface and it ensures the preservation of detailed balance in a straight forward

manner, as described in the next section.

The different physics occurring within a liquid neighborhood as opposed to away from one is captured by assigning special additional energy barriers for events taking place within such a neighborhood. That is, each event $X \rightarrow Y$ can be further sub-categorized according to neighborhood, and the additional energy barrier $\epsilon(X, Y)$ may be written as

$$\epsilon(X, Y) = \begin{cases} \epsilon^L(X, Y) & \text{if event occurs in a liquid neighborhood;} \\ \epsilon^{\text{NL}}(X, Y) & \text{otherwise.} \end{cases}$$

For example, in the nanowire simulation, the vapor to solid reaction has an additional energy barrier ρ . Depending on whether the reaction is to take place in the liquid or elsewhere, ρ is assigned a value ρ^L or ρ^{NL} , respectively, such that $\rho^L < \rho^{\text{NL}}$. The lower energy barrier for a reaction in the liquid phase models the catalytic effect of the droplet on the vapor to solid reaction, an important feature in VLS growth.

As a matter of simplicity, liquid is modeled on-lattice. Despite this geometric constraint, the model captures a reasonable set of physics pertinent to a liquid phase. The steps outlined above imply that liquid phases can be assigned alternative energy barriers to differentiate those events occurring within that phase. This allows us to model different diffusivities as well as the role of a liquid catalyst. Moreover, because our model is a bond-counting one, the simple difference in coordination within a liquid neighborhood and outside of one induces a chemical potential between phases. Most importantly, because the model satisfies detailed balance (as described in the next section), our simulations proceed by the same thermodynamic driving force present in physical systems. That is, simulations naturally move toward configurations with lower energy. From this, thermodynamic properties of the liquid phase such as surface tension forces at the vapor-liquid interface and solubility within liquid droplets are

emergent properties arising within the simulation.

2.1.3 Detailed balance

As stated above, KMC is the simulation of a CTMC with transition rates $r(X \rightarrow Y)$ between states $X, Y \in \mathcal{X}$, a state space corresponding to atom configurations. For those events that maintain the current canonical ensemble, that is those events that preserve species counts, we ensure that their rates satisfy detailed balance:

$$(2.4) \quad \pi(X)r(X \rightarrow Y) = \pi(Y)r(Y \rightarrow X),$$

where π is the canonical Boltzmann distribution

$$\pi(X) = \frac{1}{Z} \exp \left[-\frac{E(X)}{k_B T} \right].$$

The detailed balance constraint is a sufficient condition to ensure that the equilibrium distribution of the Markov chain is π . In this distribution, low energy states have maximum likelihood, and so the Markov chain is driven to minimize energy. This the thermodynamic driving force behind the simulations. That the rates of the form (2.2) with activation energies as in (2.1) satisfy detailed balance is a matter of straight forward calculation, provided we impose a symmetry between intermediate states and additional energy barriers:

$$X \wedge Y = Y \wedge X$$

$$\epsilon(X, Y) = \epsilon(Y, X).$$

The constraint that $\epsilon(X, Y) = \epsilon(Y, X)$ means the additional energy barrier used in an event must be the same as the one used for its reverse event. This leads to a subtle complication when specifying alternative additional energy barriers in the liquid phase. In determining whether to apply $\epsilon^L(X, Y)$ or $\epsilon^{\text{NL}}(X, Y)$ to an event

$X \rightarrow Y$, we must do so using a property of the affected atoms' neighborhoods that is invariant under the event. This ensures that the same categorization will be applied for the reverse event $Y \rightarrow X$. The proposed population count outlined above is an example of such a property.

2.2 The KMC algorithm

As stated in the introduction, in its most general form KMC is a simulation of a CTMC and may be implemented as such. Given a discrete state space \mathcal{X} , an initial state X_0 , and rates $r(X \rightarrow Y)$ between states, the CTMC evolves according to three steps. First, a transition $X_i \rightarrow Y$ is sampled according to the probability distribution

$$P_{X_i}(Y) = \frac{1}{Z_i} r(X_i \rightarrow Y),$$

where the partition function is the summation

$$Z_i = \sum_Y r(X_i \rightarrow Y),$$

over all states Y accessible from X_i . Second, the current configuration is altered $X_{i+1} = Y$ according to the sampled transition. Because the transition times $X_i \rightarrow Y$ are each exponentially distributed with mean $\frac{1}{r(X_i \rightarrow Y)}$, the expected minimum waiting time is given by

$$dt_{X_i} = \frac{1}{Z_i},$$

and the simulation time is incremented by this amount. Lastly the transition rates $r(X_{i+1} \rightarrow Y)$ and partition function Z_{i+1} are retrieved or calculated in some fashion. This main loop is summarized in Algorithm 1. With the on-lattice assumption, transitions $X \rightarrow Y$ between atom configurations X and Y are the atomistic events outlined above, and are parameterized by an atom/event pair (x, e) , where e describes a local event of atom x (e.g. x desorbs, exchanges with a specific neighbor, etc). By

this local property and bond-counting assumption, differences in energies, and hence rates by Equation (2.2) are local quantities. We shall write $r(X; x, e)$ whenever we wish to emphasize the local nature of the transition rates and $r(X \rightarrow Y)$ when the general CTMC interpretation of KMC is used.

Algorithm 1: The basic KMC main loop

```

1 Initialize configuration  $X_0$  and rateTable;
2 for  $i = 0, 1, 2, \dots$  do
3    $(e, x) = \text{Sample}(\text{rateTable})$  ;
4    $(X_{i+1}, \text{affectedAtomList}) = \text{PerformEvent}(X_i, e, x)$  ;
5    $\text{Update}(\text{rateTable}, \text{affectedAtomList}, X_{i+1})$  ;

```

If L is sufficiently large, then the rates of any event involving atom x can be written as a function the species of atoms in $B_L(x)$, collected in an array

$$\Sigma(x) = (\sigma(y_1), \dots, \sigma(y_n)),$$

where y_1, \dots, y_n is some linear ordering of elements in $B_L(x)$. The value of L is selected so that $B_L(x)$ contains all atoms used in the calculation of any rates for events involving x , in addition to the neighborhood used in determining the local connectedness property. For the simulations in this paper, $L = 2$. The array $\Sigma(x)$ is called the *local neighborhood pattern of x* .

Because the number of configurations is large, it is infeasible to store the entire set of rates $r(X \rightarrow Y)$ for every pair $X, Y \in \mathcal{X}$. Instead, for the current configuration X_i a subset of rates $\{r(X_i, \rightarrow Y)\} = \{r(X_i; x, e)\}$ called the rate table is maintained. Because a transition $X_i \rightarrow X_{i+1}$ occurs in every iteration of the KMC main loop, this rate table must be updated at each step. Since state transitions correspond to local atomistic events, changing a configuration affects the species of at most two lattice positions, and subsequently a small number of local neighborhood patterns $\Sigma(x)$ for affected atoms x whose neighborhoods contain the changed lattice positions. Because

the rates $r(X_i; x, e)$ are a solely a function of $\Sigma(x)$, the rate tables $\{r(X_i; x, e)\}$ and $\{r(X_{i+1}; x, e)\}$ are identical for all but a few entries corresponding to those atoms x whose neighborhood pattern have changed after the event.

There are different methods for sampling from and updating the rate table, depending on how it is maintained. For example, the table may be implemented as an array, indexed by atom/event pairs (x, e) . Updating the table consists solely on updating the individual entries in the array, and is an $\mathcal{O}(1)$ operation. Similarly, updating the partition function can be done in $\mathcal{O}(1)$ time:

$$Z_{i+1} = Z_i + \sum_{\text{affected atoms } x} \sum_e r(X_{i+1}; x, e) - r(X_i; x, e).$$

Sampling from the array can be done by the rejection algorithm. An atom/event (x, e) pair is first sampled uniformly. Then a random number p is sampled uniformly between 0 and 1, and the event is selected if $p \leq \frac{r(X_i; x, e)}{Z_i}$. Otherwise it is rejected and another event pair is sampled. If $q(x, e) = \frac{r(X_i; x, e)}{Z_i}$, then the probability of accepting some event is given by

$$\begin{aligned} \mathbb{P}(\text{Accept}) &= \sum_{(x, e)} \frac{q(x, e)}{N_e N}, \\ &= \frac{1}{N_e N}, \end{aligned}$$

where $N_e = 5$ is the number of events and N is the number of atoms. As such, the expected number of samples that must be made before an atom/event pair is accepted is $\mathcal{O}(N)$.

A more balanced approach in sampling and updating may be taken by using a binary tree, in which leaves correspond to the rates $r(X_i; x, e)$ for every atom/event pair (x, e) and each internal node stores the sum of its two children [27]. This method is an efficient form of inverse transform sampling [29], which in generality

states that if Y is a uniform random variable between 0 and 1, and $F_X(x)$ is the cumulative distribution function (CDF) of a random variable X , then the random variable $F^{-1}(Y)$ and X are identically distributed, giving a way of sampling X by means of a uniform random variable. The binary tree structure, which encodes the unnormalized CDF for atom/event pairs, is an efficient way to perform the inversion by means of binary search. A uniform random variable Y representing a sample cumulative mass is selected between 0 and Z_i . Starting from the root node of the tree, the algorithm compares the cumulative masses C_L and C_R stored in the node's left and right children, respectively. If $Y < C_L$, then the algorithm recurses on the left child with sample cumulative mass Y . Otherwise, it recurses on the right child with sample cumulative mass $Y - C_L$. The recursion is terminated once the algorithm reaches a leaf node, corresponding to an atom/event pair (x, e) which is returned as the sampled pair. The running time for sampling is determined by the height of the tree, which is $\mathcal{O}(\log N)$. After an event is performed, the leaves for the affected atoms are updated, invalidating the quantities stored in the internal nodes for all branches terminating at such leaves. Therefore, the partial sums along these branches must be updated, which is done by propagating changes upward from the leaves to the root node. As such, updating the tree is also performed in time $\mathcal{O}(\log N)$. Algorithms 2 and 3 summarize the sampling and updating methods using a binary tree.

2.3 Hash table based caching

Previous work in addressing the efficiency of KMC has been in the proper data structure choice of the rate table, an example of which was presented above and in [27] using a binary tree. Another example of this is the use of inverted linked-list

Algorithm 2: The KMC algorithm for sampling the rate table by binary search.

input : Binary tree `rateTable`
output: A sampled atom/event pair (x, e)

- 1 `currentNode` \leftarrow `rateTable.headNode`;
- 2 Sample $Y \sim \text{Unif}[0, \text{currentNode.value}]$;
- 3 **while** `currentNode` *is not a leaf node* **do**
- 4 **if** $Y < \text{currentNode.leftChild.value}$ **then**
- 5 `currentNode` \leftarrow `currentNode.leftChild`;
- 6 **else**
- 7 $Y \leftarrow Y - \text{currentNode.leftChild.value}$;
- 8 `currentNode` = `currentNode.rightChild`;
- 9 **return** $(\text{currentNode.x}, \text{currentNode.e})$

Algorithm 3: The KMC algorithm for updating the rate table by propagating changes in leaf nodes upwards through their branches

Input: Binary tree `rateTable`, list of affected atoms `affectedAtomList`, current state X

- 1 **forall** atoms $x_i \in \text{affectedAtomList}$ **do**
- 2 **forall** events e **do**
- 3 Let `currentNode` be the leaf node corresponding to the atom/event pair (x_i, e) ;
- 4 `currentNode.value` $\leftarrow r(X; x_i, e)$;
- 5 **while** `currentNode` \neq `rateTable.headNode` **do**
- 6 `currentNode` \leftarrow `currentNode.rightChild` ;
- 7 `currentNode.value` \leftarrow `currentNode.leftChild.value` + `currentNode.rightChild.value` ;

structures to achieve constant time sampling and updating [28]. However, profiling the execution of a KMC implementation shows that the computation of rates (Algorithm 3, Line 4) contributes a significant amount to the execution time of the implementation. This is because the calculation consists of several floating point operations and calls to external libraries to compute rates of the form in Equation (2.2). Moreover, such rates must be computed several times for every step in simulations that require up to hundreds of billions of such steps. While much work has been done in reducing the computational complexity of sampling and updating, little attention has been made in the expensive, albeit constant-time calculation of rates.

One way to address this is by eliminating redundant rate calculations. As described in Section 2.1, rates for an atom x are a function of the species pattern $\Sigma(x)$ describing the atom’s local neighborhood. Atoms with identical species patterns are assigned identical rates. We denote the collection of these rates as $\mathbf{r}(\Sigma(x))$ to emphasize its dependence solely on the species pattern, and not expressly on the atom x itself. If the rates $\mathbf{r}(\Sigma)$ of a pattern Σ have been previously computed during a simulation run, it is unnecessary to recompute the rates whenever Σ is encountered again. Instead, the rates of frequently occurring patterns should be cached for reuse. Such a caching technique would only prove effective if the KMC algorithm revisits the same neighborhood patterns frequently during a simulation. This is similar to the principle of data locality studied by computer scientists. In this context, data locality is exploited by caching; storing recently accessed blocks of memory in a low-latency cache with the expectation that such blocks will be accessed frequently.

2.3.1 Neighborhood locality

As evidence to show that this repetition does indeed occur, we consider the nanowire simulations described in Section 4.1. At regular steps throughout the simu-

lation, the program determined the number of distinct neighborhood patterns N_{pattern} within the current configuration. Figure 2.4 plots N_{pattern} as a function of the Monte Carlo steps t , ensemble averaged over 128 identical trials. The graph shows that N_{pattern} is slowly growing, and remains less than 400 neighborhoods after 400 million steps. This suggests that, within a fixed configuration, there is much redundancy between neighborhood patterns, implying spatial neighborhood locality.

While N_{pattern} measures the number of neighborhood patterns active at some time, we may also determine the cumulative number C_{pattern} of distinct neighborhoods seen up to a particular time. This is shown in Figure 2.5 as the solid blue line. We observe that the growth is sub-linear, but asymptotically approaches linear growth as time increases. The dashed red line indicates the linear best fit of this asymptotic growth. Its slope indicates that in equilibrium, the rate of change $\frac{dC_{\text{pattern}}}{dt}$ is approximately 8.4×10^{-5} new neighborhoods per MC step. That is, a new neighborhood is encountered once every 12,000 MC steps. Therefore, there is much repetition in neighborhood pattern between configurations, implying temporal neighborhood locality. The temporal and spatial locality of neighborhood patterns therefore implies that the use of caching is indeed advantageous.

2.3.2 Hash table based rate caching

Because the rates of events depend on the occupancy pattern of local neighborhoods, it is desirable to have a precomputed look-up table of all the rates, indexed by pattern. Naively indexing by the species of atoms in a neighborhood results in a look-up table that is prohibitively large. In the nanowire simulations, the patterns $\Sigma(x)$ describe the occupancy of lattice positions in $B_2(x)$, a total of 19 atoms. Given there are 4 species and 10 events that may occur at each position, this results in a look-up table with $10 \times 4^{19} \approx 2^{41}$ floating point entries, though we have seen that

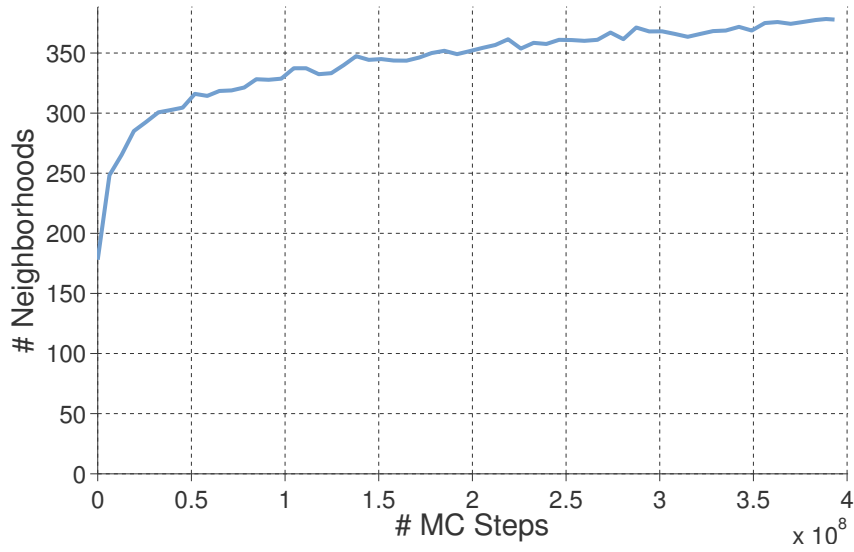


Figure 2.4: The ensemble average of N_{pattern} , the number of distinct neighborhood patterns within the atom configuration at time t during a nanowire simulation.




Figure 2.5: The ensemble average of C_{pattern} , the cumulative number of distinct neighborhood patterns seen up to step t during the nanowire simulation. The dashed red line indicates the linear asymptotic behavior of C_{pattern} .

only a small set of such entries are ever used during the simulations. One reason why the naive indexing results in a large table is that it includes neighborhoods that do not satisfy the local connectedness property, and hence will never occur in the simulations. If the number of neighborhood patterns that adhere to this constraint is sufficiently small, a clever encoding of patterns could be used as an index to the precomputed rate table. However, if there are a large number of species or large local neighborhoods used to check connectedness, then such an encoding method may also be infeasible. Another source of redundancy in the naive implementation is that symmetries of the lattice structure are not accounted for by a simple indexing. Such symmetries result in patterns associated with identical sets of rates, up to relabeling. Developing techniques characterizing such symmetries is system specific, and depends on the number of species and the lattice used. Indeed, much work has been done in pattern recognition of local neighborhoods to address this problem [30].

We use an alternative method of characterizing commonly occurring patterns apart from the *a priori* large set Ω of possible patterns by storing rates of encountered patterns on-the-fly in a hash table. Previous work along these lines used a Zobrist hash to encode an atom configuration as a whole, which recovers the entire rate table associated to the configuration [31]. This method ignores the spatial and temporal neighborhood locality inherent in simulating systems of a crystalline nature in addition to the observation that due to the local nature of a KMC transition $X \rightarrow Y$, the rate table for Y is identical to that of X except for a small number of entries. By encoding neighborhood patterns within a configuration rather than the configuration itself, we take advantage of the locality inherent in KMC. Stored within the hash table are portions of the rate table associated to particular neighborhood patterns, instead of the entire rate table for a given configuration. Storing

by local neighborhood implies a smaller number of possible states to encode, leading to a smaller hash table, faster hash key construction and comparison. This speed is important when placed in the context of the alternative, the explicit calculation of rates. The caching technique leads to a performance gain only if it can retrieve entries from the table faster than computing rates outright. This method is similar to that used in a KMC variant known as Self-Learning KMC, in which atomistic events (either specified explicitly or obtained from molecular dynamics simulations) are placed in a database for use in the KMC algorithm [20, 32].

When a configuration is altered, the local neighborhoods of a small number of atoms are affected, and the rates for the corresponding neighborhood patterns must be recalculated before the update algorithm is applied to the rate table. When the rates $\mathbf{r}(\Sigma)$ of a pattern Σ are needed, the hash table is queried to see if $\mathbf{r}(\Sigma)$ has been previously computed. If not, it must be computed and stored in the hash table, otherwise the rates are copied to the appropriate entries in the rate table. The hash table is queried and populated by linear open addressing, using a hash function

$$h : \Omega \rightarrow \{0, \dots, M - 1\},$$

where $M \ll |\Omega|$ is the size of the table. When the hash table is queried for rates $\mathbf{r}(\Sigma)$, a hash index $i = h(\Sigma)$ is generated and the i -th entry in the hash table is checked. If that entry is empty, then the pair $[\Sigma, \mathbf{r}(\Sigma)]$ is computed and stored at entry i . If it is not empty, it contains a pair $[\Sigma', \mathbf{r}(\Sigma')]$. If $\Sigma = \Sigma'$, then the rate $\mathbf{r}(\Sigma') = \mathbf{r}(\Sigma)$ is returned, otherwise, the hash index is incremented $i = i + 1$, and the procedure is iterated until either $\mathbf{r}(\Sigma)$ is found or an empty entry is encountered, in which case the rates $\mathbf{r}(\Sigma)$ must be calculated and stored in that position. The hash algorithm is summarized in Algorithm 4. The number of times the index i is incremented directly impacts the efficiency of the hashing procedure and depends on runs of non-empty

entries in the table, called chains. The efficiency of Algorithm 4 depends linearly on the average chain length, so it is important to select a hash function that minimizes this quantity. In our implementation, a parameterized hash function $h = h_A$ was used, where A is a set of parameters to the hash function. The parameters A were selected to minimize chain length. The selection of optimal parameters A is discussed in Section 2.3.3. In the hash table implementation, we employ a simple hash table eviction policy, clearing the entire table when a user-specified maximum capacity was reached. As such, the number C_{pattern} is the relevant quantity in deciding the size of the table.

Algorithm 4: Querying the rate cache by a linear, open-addressing hash table.

Input : A local neighborhood Σ
Output: The index i into `hashtable` where the rates $\mathbf{r}(\Sigma)$ are stored

```

1 Compute initial hash index,  $i = h(\Sigma)$ ;
2 while index  $i$  not returned do
3   if hashtable [ $i$ ] == null then
4     Compute  $\mathbf{r}(\Sigma)$ ;
5     hashtable [ $i$ ] = [ $\Sigma$ ,  $\mathbf{r}(\Sigma)$ ];
6     return  $i$ ;
7   else
8     [ $\Sigma'$ ,  $\mathbf{r}(\Sigma')$ ] = hashtable [ $i$ ];
9     if  $\Sigma' == \Sigma$  then
10      return  $i$ ;
11     else
12       $i = i + 1 \pmod{M}$ ;

```

When implemented with an optimal hash function, the caching procedure leads to performance gains in practice. Figure 2.6 plots the CPU time against the number of MC steps for several independent trails of nanowire growth simulations, varying the number of MC steps performed and the method of caching used. Red squares indicate those simulations in which no caching was used. There, when rates $\mathbf{r}(\Sigma)$ were required by the simulation, they were computed explicitly. The figure indicates a linear trend, and the best-fit slope suggests that without caching, simulations take

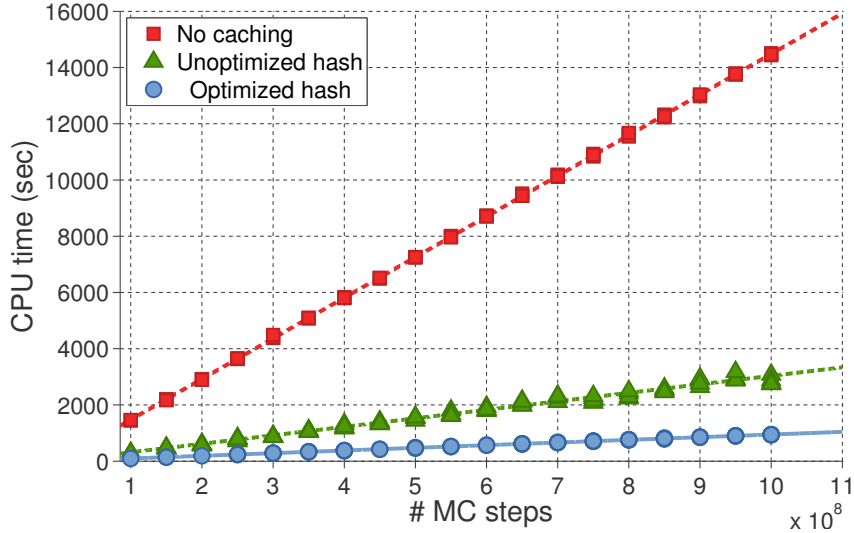


Figure 2.6: CPU time versus the number of MC steps performed during the nanowire simulations. Simulations without caching implemented are plotted in red squares. Green triangles show simulations where caching is used with an unoptimized hash function. Blue circles show simulation times when an optimal hash function is used.

about 14.5×10^{-6} seconds per MC step. Depicted in blue triangles are times for simulations that implement the caching procedure, but with an unoptimized choice of parameters A . In this case, each simulation step requires 3.0×10^{-6} seconds per MC step, a performance gain by a factor of about 4.8 compared to the simulations without caching. If an optimal hash function is used, the simulations take 0.9×10^{-6} seconds per MC step, a gain by a factor of 3 compared to the unoptimized case and a factor of 16 compared to simulations without caching. This is depicted by the green circles in the figure.

In the 700 million MC steps of a particular simulation run, there were 11.4584 billion queries into the hash table, indicating that about 16 neighborhoods were updated per MC step, on average. Of those queries, 11.4583 billion neighborhoods (99.9991%) were previously computed and stored in the hash table, requiring no extra computation. The average chain length in the hash table was approximately 1.7 nodes for the choice of optimal hash function.

2.3.3 Finding optimal hash functions by simulated annealing

As described above, neighborhood patterns are hashed by a hash function to produce an index into the hash table. In this context, neighborhoods $\Sigma \in \Omega$ are treated as bit strings packed into $W = \frac{\log_2 S |B_L|}{64}$ 64-bit words. Figure 2.7 illustrates how this is done. Each position in a neighborhood is assigned an index that linearly orders the positions of the neighborhood. The same ordering is used for every neighborhood. For example, in the figure, the center atom is assigned the index 0, indicating it is in the first position in the linear ordering. The atom coordinated to the northeast of this center atom is assigned the next index 1, indicating it is in the second position of the linear ordering. Each atom in a neighborhood is assigned a $\log_2 S$ bit identifier according to that atoms species. These bits are then concatenated together in the linear ordering, resulting in a bit strings of length $|B_L| \log_2 S$ bits. For example, the neighborhood in the Figure 2.7 can be represented by the bit string

$$\Sigma = 10101010010110011110001010010101010110,$$

assuming the 0-th atom position corresponds to the leftmost two bits, etc.

The bit string is hashed by a function $h_A : \Omega \rightarrow \{0, \dots, N - 1\}$ of the form

$$h_A(x) = \bigoplus_{i=1}^W (x_i \gg a_{i,1}) \oplus (x_i \gg a_{i,2}) \oplus (x_i \gg a_{i,3}) \pmod M,$$

where \oplus and \gg are bit-wise XOR and (non-cyclic) right shift, respectively. The parameters $A = (a_{ij})$ are shift values such that

$$-64 < a_{ij} \leq 64,$$

with the convention that $a_{ij} < 0$ is a left-shift.

Optimal shift parameters are selected with respect to a training set $\mathcal{T} \subset \Omega$ of neighborhoods. The training set \mathcal{T} is generated empirically by running the simula-

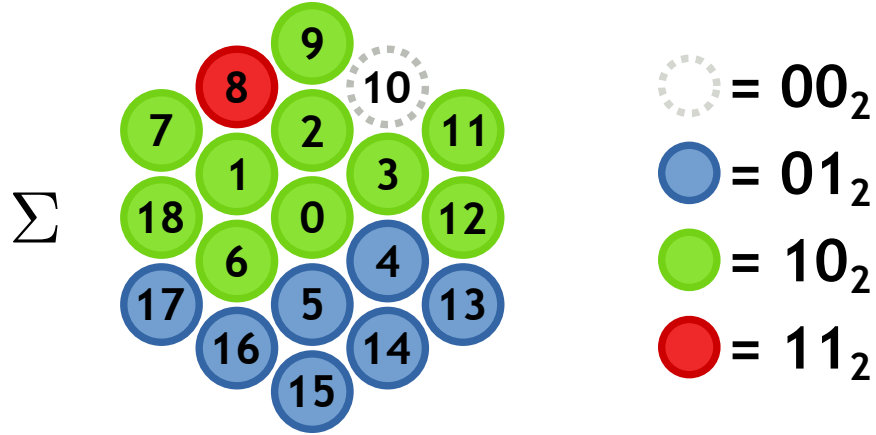


Figure 2.7: The indexing of atoms in a neighborhood Σ of radius 2 and two-bit identifiers assigned to each species.

tions using an unoptimized hash function. At regular steps, the set of neighborhoods stored in the hash table is sampled uniformly, and a small proportion of this set is added to \mathcal{T} as the simulation progresses. Repeated keys are removed from \mathcal{T} . The size of \mathcal{T} is chosen large enough to provide a representative sample of neighborhood patterns observed during a simulation. For example, in determining the optimal parameters for the nanowire simulations, a training set of size $|\mathcal{T}| = 96000$ distinct neighborhoods was used. Each choice of parameters A partitions \mathcal{T} into cosets

$$\mathcal{T}_y = \{X \in \mathcal{T} : h(X) = y\},$$

of configurations with identical hash value. We define a quantity $C_{\mathcal{T}}$ to be the number of such cosets of size larger than 1. That is, $C_{\mathcal{T}}$ is the number of configurations X for which there is at least one other configuration Y that hashes to the same value, i.e. a collision in the hash table. Minimizing collisions is an effective heuristic in reducing chain length, and hence we must choose parameters that minimize $C_{\mathcal{T}}$.

The time to perform a brute force search for optimal parameters over the entire space is of order $\mathcal{O}\left(S^{\frac{21}{64}|B_L|} |\mathcal{T}|\right)$. For the nanowire simulation, this is approximately 550 million calculations. While this brute-force calculation is feasible, we observe

that the above running time scales poorly with S and B_L . Indeed, increasing r from 2 to 3 results in a brute-force calculation of approximately $2^{41} \approx 2$ trillion calculations. For sintering and grain growth simulations, in which the number of species corresponds to the number of orientations a grain may possess, S can be on the order of 100,000 [33]. In this case, the brute-force calculation requires $2^{103} |\mathcal{T}|$ calculations, which is infeasible to perform.

Rather than using the brute-force calculation, good shift values are found by a simulated annealing algorithm with Metropolis selection. The simulated annealing algorithm runs as follows. Starting at a random $A = (a_{ij})$ where each a_{ij} is selected uniformly in $\{-63, \dots, 64\}$, a mutation is performed

$$A' = A + \Delta,$$

where $\Delta = (\delta_{ij})$, δ_{ij} being independently and uniformly selected integers between -4 and 4. The mutation is kept according to a Metropolis rule based on the energy $C_{\mathcal{T}}(A)$. That is, the probability $p(A \rightarrow A')$ of keeping the mutation is given by

$$p(A \rightarrow A') = \begin{cases} 1 & \text{if } C_{\mathcal{T}}(A') < C_{\mathcal{T}}(A), \\ \exp[-\beta(C_{\mathcal{T}}(A') - C_{\mathcal{T}}(A))] & \text{if } C_{\mathcal{T}}(A') \geq C_{\mathcal{T}}(A), \end{cases}$$

where $\beta = \frac{1}{1000}$ represents inverse temperature (units arbitrary). The above transition probabilities satisfies detailed balance with respect to a Boltzmann distribution

$$\pi(A) \propto \exp[-\beta C_{\mathcal{T}}(A)].$$

Therefore, in steady state the simulated annealing algorithm will select parameters that approximately minimizes $C_{\mathcal{T}}(A)$. The algorithm is run for several choices of initial values.

Figure 2.8 shows three independent simulated annealing trials using the same training set \mathcal{T} of over distinct 96,000 neighborhoods sampled during the nanowire

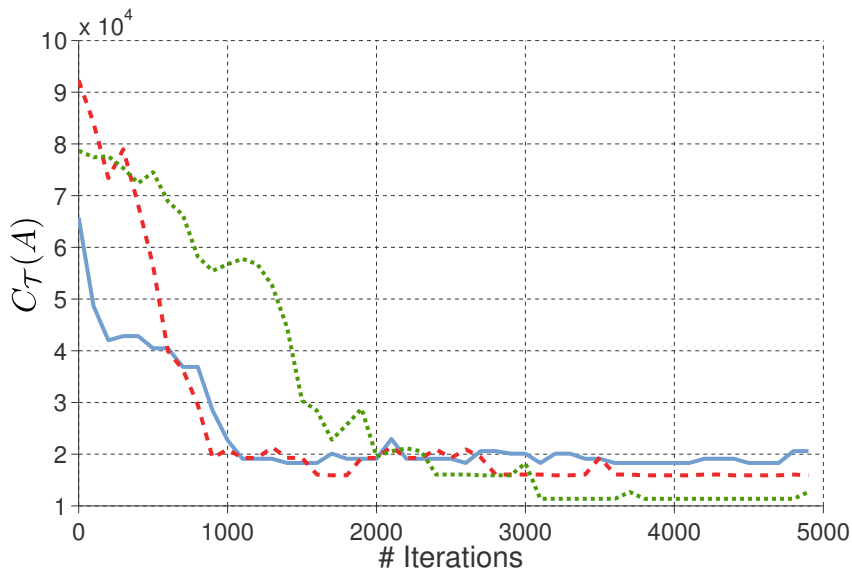


Figure 2.8: The number of initial conditions $C_{\mathcal{T}}(A)$ versus the number of iterations during the simulated annealing. Three independent trials are shown.

simulations. Different, independent initial shift widths A were selected for each trial. We see that the simulated annealing algorithm converges after 5000 iterations.

The specific choice of optimal hash parameters depends on the system to be simulated. For each of the three systems studied in this work, the simulated annealing search was performed off-line, apart from the actual KMC simulation. This results in three system-specific hash functions. As an alternative, the simulated annealing search could be made in-line, performed along side the simulation. A random choice of parameters are selected during the initial stages of the simulation. As the simulation progresses, the training set \mathcal{T} is built. A simulated annealing search is then performed to construct optimal hash functions based on the actual neighborhoods encountered for the specific simulation run. When the search is performed periodically throughout the simulation, we obtain an adaptive hash function responsive to the changing sets of neighborhood typical in simulations involving several experimental stages (e.g. droplet epitaxy and crystallization).

CHAPTER III

Application: Liquid Ga droplet epitaxy and crystallization

3.1 Liquid droplet epitaxy and crystallization

In this chapter we detail simulations and analysis of Ga droplet epitaxy (DE) and crystallization. Specifically, we characterize the nanostructural dependence on growth conditions. As stated in the introduction, Ga droplet epitaxy and crystallization proceeds in two steps. First, Ga atoms are deposited on a GaAs substrate. The Ga atoms diffuse along this substrate, and eventually coalesce to form liquid Ga droplets. The diffusion length scale depends on the Ga deposition rate and temperature, and determines the size and density of the droplets. During this stage, regions of GaAs underlying the droplets become dissolved into the droplets, resulting in an etching of the GaAs substrate. The amount of etching is dictated by the solubility of GaAs inside liquid Ga, which depends on temperature. After the Ga atoms are deposited and the system is allowed to re-equilibrate, an As flux is introduced to induce crystallization of the droplets. Those As atoms that incorporate into the droplet nucleate along the liquid-solid interface, resulting in a GaAs growth front advancing into the droplet. Arsenic atoms are also deposited away from the droplet. As we shall see, this induces a wicking of Ga atoms away from the droplet, and the competition between these two processes determine the resulting nanostructure in

large part.

As described above, several important processes play a role during DE and crystallization, making its simulation more complicated than previous work in GaAs systems. Prior work focused on homoepitaxial film growth and studied associated phenomena such as step density [34] and growth modes of GaAs films [9]. In these simulations, surface diffusion of adatoms played the central role, and because of the stoichiometric nature of epitaxial growth, a simple cubic lattice and the SOS constraint sufficed in modeling key aspects of this process. Droplet epitaxy however, poses several issues that cannot be captured by earlier models. During DE, the system is inherently non-stoichiometric in that the relative concentrations of Ga and As atoms on the surface are different from one another. Processes other than surface diffusion, such as events within liquid and at the liquid-solid interface, play a key role and cannot be captured by a simple SOS model. Lastly, DE results in nanostructures on the order of 10s to 100s of nm so that the simulations must be performed efficiently within large domains.

3.2 KMC model parameters

In reality, GaAs has a zincblende crystal structure [35], which is depicted in the left panel of Figure 3.1. The zincblende structure is composed of two interwoven face centered cubic (FCC) lattices, each occupied by either Ga or As. The two lattices are connected by diagonal Ga-As nearest neighbor bonds. Every atom is bonded to four nearest neighbors. In the simulations, an analogous 1+1 dimensional lattice is used, and is illustrated in the right panel of Figure 3.1, with red and green atoms representing Ga and As respectively. Such a configuration depicted in the figure could represent the interface between the GaAs substrate and a liquid Ga droplet.




Figure 3.1: (Left) The Zincblende crystal structure. Obtained from [1]. (Right) The analogous 1+1 dimensional lattice used in the simulations. Nearest neighbor bonds are indicated by solid lines. Next-nearest neighbor bonds are shown as dashed lines. Ga and As atoms are colored red and green, respectively.

Each lattice position has four diagonal nearest neighbor bonds (indicated by solid lines in the figure) and up to four lateral next-nearest neighbor bonds. The presence of the lateral next-nearest neighbor bonds mimics the FCC sublattices present in Zincblende. To encourage the segregation of species to their respective sublattice, next-nearest neighbor bonds are only assigned between atoms of the same species. In this way, it is energetically favorable for a Ga atom to be coordinated diagonally by As atoms and laterally by other Ga atoms, and similarly for As atoms. This results in a thermodynamically emergent crystalline structure. For example, the As atom *A* in Figure 3.1 has four Ga nearest neighbors but only three next-nearest bonds due to the presence of a Ga atom above it. The Ga atom *B* has three As nearest neighbors and one Ga nearest neighbor, illustrating the fact that the model allows for nearest neighbor bonds between atoms of the same species. This atom is also fully coordinated laterally with four next-nearest neighbor Ga-Ga bonds.

Due to the multi-species nature of the model, the several bonding energies must be specified. Nearest neighbor bonding energies are given by three parameters $\gamma(\text{Ga}, \text{Ga})$, $\gamma(\text{Ga}, \text{As})$ and $\gamma(\text{As}, \text{As})$ representing Ga-Ga, Ga-As and As-As bonds,

$\gamma(\text{eV})$	Ga	As	$\gamma_{\text{nn}}(\text{eV})$	Ga	As
Ga	0.3	0.5	Ga	0.3	-
As	0.5	0.1	As	-	0.1

Table 3.1: Pairwise nearest and next-nearest neighbor bonding energies used in the GaAs system.

respectively. Next-nearest neighbor bonds are only assigned between atoms of the same species and are denoted $\gamma_{\text{nn}}(\text{Ga}, \text{Ga})$ and $\gamma_{\text{nn}}(\text{As}, \text{As})$. The parameter values used in the simulations are summarized in Table 3.1. These energy values imply that the Ga-As bonds are energetically preferred, while weak As-As bonds effectively eliminate excess As by desorption. The specific values in Table 3.1 were calibrated using the homoepitaxy (Section 3.2.1) and liquid droplet simulations (Section 3.3) to match qualitative experimental assumptions and observations. Within the homoepitaxy simulations, the Ga-As bond strengths were tuned in order to observe a transition from rough island formation in the low temperature regime to a smooth step-flow growth mechanism in high temperatures [36]. In the case of liquid droplet simulations, we tune the Ga-Ga bond strength to match liquid droplet statistics such as droplet width and number density. As-As bonds were selected to be nominally weak.

3.2.1 Calibration

GaAs homoepitaxy simulations

As an initial application and calibration of the model, we simulated GaAs film growth at various growth conditions. It is known that the surface reconstruction of GaAs (i.e. a structure on the surface, which is different from that of the bulk) depends on both the temperature and the relative deposition rates of Ga and As [37]. In 1+1 dimensions, simulations cannot reproduce surface reconstructions. Instead, we measured surface termination and its dependence on temperature and incoming deposition rates. The surface termination describes the species (either Ga or As) of

the surface, and can be viewed as a coarsening of surface reconstruction information. That is, each surface reconstruction can be classified as Ga or As terminated. The simulation results were then compared to experimental data.

To simulate GaAs substrate growth, both Ga and As were deposited simultaneously on initially flat, As-terminated substrate. The Ga deposition rate F_{Ga} was fixed at 0.37 ML/s, while the As deposition rate F_{As} was varied so that the deposition ratio $F_{\text{As}}/F_{\text{Ga}}$ ranged between .5 and 10. Rates (reported in ML/s) describe the rate at which atoms are added to the system. The observed stoichiometric growth of the film is an emergent property of the model rather than explicitly enforced through deposition rates. This is manifest in the stoichiometry observed over a broad range of deposition ratios.

The temperature was varied between 427°C and 727°C. Five monolayers of total material were deposited, and surface Ga concentration was measured at regular intervals during the deposition of the last two monolayers. Figure 3.2 is a surface termination phase diagram as a function of deposition rate and temperature for both experimental and simulation results. Red squares indicate the conditions where simulations show a predominantly Ga-terminated surface, while green circles are conditions yielding a predominantly As-terminated surface. Experimentally determined transition from the As-terminated 2×4 to the Ga-terminated 4×2 reconstructions are shown as blue diamonds and indicate a good agreement between simulation and experimental data.

Along with the Ga-As bonding energy, another relevant parameter that controls surface termination is the additional energy barrier for an As desorption event, denoted μ_{As} . Varying this parameter effectively shifts the above phase diagram horizontally. The value $\mu_{\text{As}} = 1.1$ eV is fitted to experiments. This value, combined with

the energy values in Table 3.1 yield a total activation energy of

$$(3.1) \quad E_a = 2\gamma(Ga, As) + \gamma_{nn}(As, As) + \mu_{As}$$

$$(3.2) \quad = 2.2 \text{ eV},$$

for the desorption of an As adatom from a Ga terminated substrate, comparing favorably to experimental results [38, 39]. Moreover, the specific parameter value does not significantly impact the qualitative shape of the phase diagram. That is, independent of μ_{As} , the simulations capture a constant critical deposition ratio in the low temperature regime and its transition to an increasing critical ratio as temperature increases.

Analytic expressions for critical boundaries

Here we illustrate how the KMC model can be analyzed to derive expressions for critical phenomena, in this case the boundary between Ga and As terminated surfaces. This “model of a model” approach occurs frequently throughout this dissertation. During homoepitaxial growth, several processes occur, including surface diffusion and As desorption, but to first order we may approximate the system in a quasi-static deposition/ desorption-limited regime. In this regime, the transition between the As and Ga-terminated surface occurs when the amount of Ga on the impinging upon the surface (given by F_{Ga}) is equal to the net rate of As growth (given by $F_{As} - R_{\text{desorb}}$), where R_{desorb} is the desorption rate of As on a Ga terminated surface. Because this rate depends on the atomistic neighborhood of the desorbing As atom, we define R_{desorb} as the harmonic average of desorption rates of an As adatom and an As atom coordinated with an additional nearest or next-nearest As atom (representing the desorption at a step-edge site). The desorption of an As atom that is more fully coordinated than these two cases is considered unlikely, and is hence

not included in the harmonic average. The activation energy for the desorption of an As adatom is given by Equation (3.1), while the activation energy for the second case is similarly obtained as

$$2\gamma(Ga, As) + 2\gamma_{nn}(As, As) + \mu_{As} = 2.3 \text{ eV}.$$

Then, R_{desorb} is given by the harmonic average

$$R_{\text{desorb}} = 2R_0 \exp \left[-\frac{2\gamma(Ga, As) + \gamma_{nn}(As, As) + \mu_{As}}{k_B T} \right] \left(\exp \left[\frac{\gamma_{nn}(As, As)}{k_B T} \right] + 1 \right)^{-1}$$

The critical criterion given by a balance between the rates of incoming Ga and As growth, is given by

$$(3.3) \quad \frac{F_{As}}{F_{Ga}} = 1 + \frac{R_{\text{desorb}}}{F_{Ga}},$$

and is depicted in Figure 3.2 by the dashed blue line. We see that this boundary agrees well with both simulation and experiments.

3.3 Simulation results

The main thrust of the work in the GaAs system was to accurately capture the processes relevant to liquid droplet formation and crystallization. Experimentally, a range of nanostructures are observed depending on As flux and temperature [40–45]. The goal of the corresponding simulation work was therefore twofold: first was to simulate the range of structures observed experimentally and second, use the simulations to describe the macroscopically observed dependences at an atomistic level. The experiment and simulations proceed in two main steps, which are at the basis of the droplet fabrication of GaAs quantum nanostructures [40]. First, Ga is deposited onto the substrate by MBE, forming liquid Ga droplets. Next the droplets are crystallized by the introduction of an As flux. The resulting nanostructures may




Figure 3.2: Substrate termination phase diagram as a function of deposition ratio and temperature, obtained from simulations and experiments. Those growth conditions resulting in a mostly Ga-terminated substrate are indicated by red squares, while green circles label As-terminated ones. The blue points above indicate the conditions where the transition from Ga to As termination occurred experimentally. The blue curve indicates the boundary between Ga and As-terminated given by equation 3.3.




Figure 3.3: Droplet epitaxy experimental results showing typical nanostructures observed over a range of substrate temperatures and As_4 BEP.

range from fully crystallized compact quantum dots, to nanorings and even etched holes. Figure 3.3 shows experimentally obtained GaAs structures over a range of temperature and As flux (reported here as As_4 beam equivalent pressure (BEP)).

The simulations were performed in a similar manner to the experiments outlined above. However, in order to simplify the analysis, the temperature was maintained constant throughout each simulation. This is justified by the observation that the Ga deposition temperature affects droplet density but does not play a relevant role in determining droplet crystallization dynamics. Simulations proceed as follows:

1. Ga atoms are deposited on a flat, As-terminated GaAs substrate at a rate of 0.1 ML/s at temperature T until 4.0 monolayers are deposited;
2. The system is then annealed for 60 seconds in the absence of deposition;
3. After annealing, an As flux is introduced by the deposition of As at a rate of

F_{As} ML/s until the system attains equilibrium.

The growth parameters T and F_{As} were varied in order to study their effect on the resulting morphology. The temperature T ranged between 150°C and 350°C, while the As deposition rate F_{As} ranged between 0.1 ML/s and 4 ML/s. Within this range of growth parameters, we are able to simulate the formation of a variety of nanostructures similar to those observed experimentally.

3.3.1 Droplet epitaxy

During the first phase, Ga atoms are deposited, which in turn form liquid droplets. As Ga atoms wet the surface, Ga-on-Ga diffusion is the relevant process. The rate of this diffusion is determined by the atomistic hopping of a Ga atom on a Ga-terminated surface, and is given by

$$R_{\text{Ga}} = R_0 \exp \left[-\frac{E_{\text{Ga}}}{k_B T} \right],$$

where $E_{\text{Ga}} \approx 0.90$ eV is the implied average energy barrier taken by considering the harmonic average R_{Ga} of the hopping rates of a Ga adatom and a Ga atom that is singly coordinated by another Ga atom. As is evident in the above form for the hopping rate, statistics such as droplet size and density depend on temperature. As temperature is increased, average droplet diameters increase, while the density of droplets correspondingly decrease, phenomena captured in simulations. This is also shown in experiments [46]. Figure 3.4 is an Arrhenius plot of simulated average droplet diameters and linear density as a function inverse temperature. Statistics are ensemble averages over 64 trials for each choice of temperature after 2.9 ML of Ga deposition. Droplet diameters are obtained by considering the autocorrelation function

$$A(t) = \int_0^W h(x+t)h(t)dx,$$

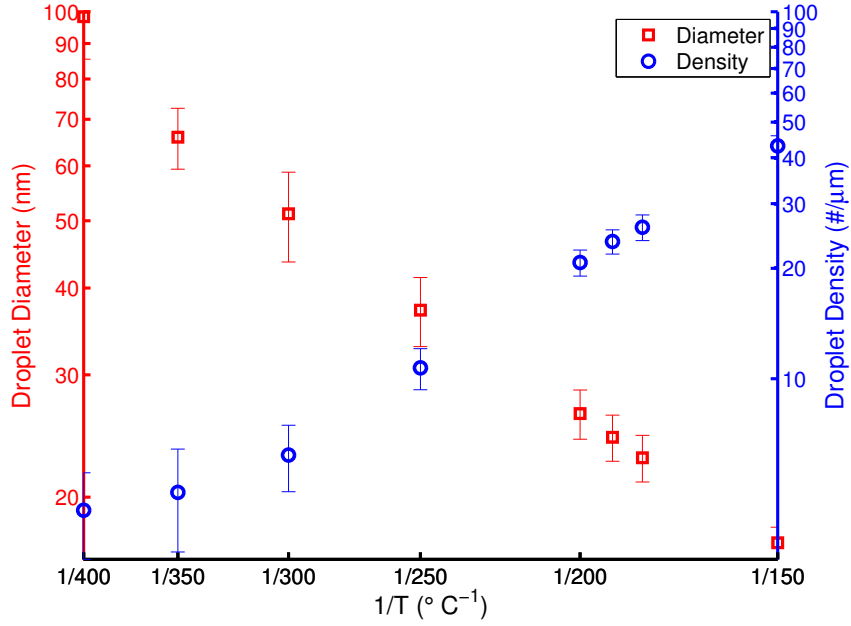


Figure 3.4: Liquid Ga droplet diameter and linear density as a function of temperature.

where $h(x)$ is the centered height profile of the final configuration of a trial and W is the width of the domain. The droplet radius $r \in [0, W]$ was obtained as the smallest value such that $A(r) < e^{-1}$. Droplet densities were determined by the number of local maxima in the $h(x)$ divided by W . We observe the correct trends in these statistics with respect to temperature – increasing temperature leads to an increase in droplet diameter and a decrease in droplet densities. The radius of a Ga droplet obtained from the deposition of 4.0 ML of Ga atoms at temperature T and flux F_{Ga} was found empirically from simulations to follow the power law involving the temperature and flux dependent dimensionless parameter $\frac{D_{\text{Ga}}}{\ell^2 F_{\text{Ga}}}$, where D_{Ga} is the diffusion coefficient for Ga-on-Ga diffusion,

$$D_{\text{Ga}} = \ell^2 R_{\text{Ga}},$$

and ℓ is the lattice spacing. That power law is given by

$$(3.4) \quad r = \ell r_0 \left(\frac{D_{\text{Ga}}}{\ell^2 F_{\text{Ga}}} \right)^\alpha,$$




Figure 3.5: The power law dependence of Ga droplet radius vs. $\frac{D_{\text{Ga}}}{l^2 F_{\text{Ga}}}$.

where $r_0 = 11.34$ (# atoms) and $\alpha = 0.182$ are empirically obtained values. Figure 3.5 shows the radius statistics obtained by simulation (blue circles) along with Equation (3.4) (dashed red line).

During droplet formation, simulations show that liquid Ga etches into GaAs, and the amount of etching is regulated by temperature, as illustrated in Figure 3.7. Higher temperatures result in more significant etching. This is in agreement with experimental observations [41, 47]. As the droplet etches into the substrate liquid Ga atoms displace substrate As atoms, which subsequently attach near the triple junction. In addition, some of the displaced substrate material is wicked out of the droplet in a step-flow growth mode. The relevant model parameter controlling the effect of etching is the additional barrier for atom-atom exchanges at the liquid-solid interface $\epsilon_{\text{ex}}^{\text{NL}}$. The value $\epsilon_{\text{ex}}^{\text{NL}} = 0.7$ eV was selected to fit qualitative experimental observations on the amount of etching occurring at various temperatures.

The amount of etching present during this phase is determined by the equilibrium concentration of As in liquid Ga. The As concentration c_{As} satisfies the reaction law

$$\frac{\partial c_{\text{As}}}{\partial t} = \ell^{-2} R_{\text{detach}} - c_{\text{As}} R_{\text{attach}},$$

where R_{detach} and R_{attach} are the atomistic rates for the detachment and attachment processes, respectively. The equilibrium concentration c_0 is then given by

$$(3.5) \quad c_0 = \ell^{-2} \frac{R_{\text{detach}}}{R_{\text{attach}}}$$

$$(3.6) \quad = \ell^{-2} \exp \left[\frac{E_{\text{attach}} - E_{\text{detach}}}{k_B T} \right],$$

where E_{attach} and E_{detach} are the energy barriers for the attachment and detachment events, respectively. To determine the values of these barriers, we consider the attachment and detachment of an As atom onto a flat liquid-solid interface, as illustrated in Figure 3.6. We denote the states before and after the detachment as X, Y , respectively. The change in energy $E(X) - E(X \wedge Y)$ between states X and intermediate state $X \wedge Y$, describing the energy barrier for attachment can be computed from a small number of relevant bonds, indicated by black lines in the figure. The total activation energy for attachment $E_{\text{attach}} = \epsilon_{\text{ex}}^{\text{NL}} + E(X) - E(X \wedge Y)$ is then given by

$$(3.7) \quad E_{\text{attach}} = \epsilon_{\text{ex}}^{\text{NL}} + 4\gamma(\text{Ga}, \text{As}) + 4\gamma(\text{Ga}, \text{Ga}) + 2\gamma_{\text{nn}}(\text{Ga}, \text{Ga})$$

$$- 8\gamma(\text{Ga}, \text{H}) - 5\gamma_{\text{nn}}(\text{Ga}, \text{H}) - \gamma_{\text{nn}}(\text{As}, \text{H}),$$

$$(3.8) \quad = \epsilon_{\text{ex}}^{\text{NL}} - \frac{1}{2}\gamma_{\text{nn}}(\text{Ga}, \text{Ga}) - \frac{1}{2}\gamma_{\text{nn}}(\text{As}, \text{As}),$$

$$(3.9) \quad = 0.5 \text{ eV},$$

where $\text{H} = \text{Ga} \wedge \text{As}$ is the intermediate species used in a Ga/As exchange. We may calculate the activation energy for detachment $E_D = \epsilon_{\text{ex}}^{\text{NL}} + E(Y) - E(X \wedge Y)$ in a

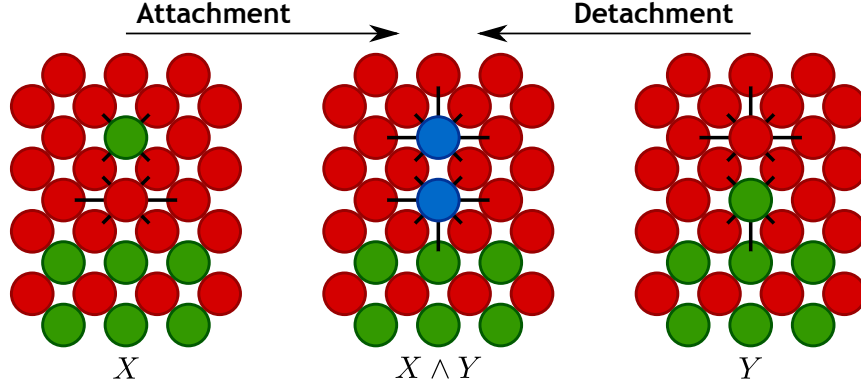


Figure 3.6: Attachment and detachment events at the liquid-solid interface. This figure illustrates the attachment and detachment of an As atom in liquid Ga onto and from a perfectly flat liquid-solid interface, along with the intermediate state for the transitions. The intermediate species $H = \text{Ga} \wedge \text{As}$ is colored blue. The black lines indicate the relevant bonds that contribute to the change in energy between the initial and intermediate state.

similar manner:

$$(3.10) \quad E_{\text{detach}} = \epsilon_{\text{ex}}^{\text{NL}} + 4\gamma(\text{Ga}, \text{As}) + 4\gamma(\text{Ga}, \text{Ga}) + 3\gamma_{\text{nn}}(\text{Ga}, \text{Ga}) + \gamma_{\text{nn}}(\text{As}, \text{As})$$

$$(3.11) \quad - 8\gamma(\text{Ga}, \text{H}) - 5\gamma_{\text{nn}}(\text{Ga}, \text{H}) - \gamma_{\text{nn}}(\text{As}, \text{H}),$$

$$= \epsilon_{\text{ex}}^{\text{NL}} + \frac{1}{2}\gamma_{\text{nn}}(\text{Ga}, \text{Ga}) + \frac{1}{2}\gamma_{\text{nn}}(\text{As}, \text{As}),$$

$$(3.12) \quad = 0.9 \text{ eV}.$$

The forms of the attachment and detachment barriers E_{attach} and E_{detach} indicate what physically occurs during the transitions, namely the formation/removal of one next-nearest neighbor Ga-Ga bond and one next-nearest neighbor As-As bond. The equilibrium concentration c_0 is therefore given by

$$c_0 = \ell^{-2} \exp \left[-\frac{\gamma_{\text{nn}}(\text{Ga}, \text{Ga}) + \gamma_{\text{nn}}(\text{As}, \text{As})}{k_B T} \right].$$

This concentration, and hence the amount of etching, increases with increasing temperature.




Figure 3.7: Liquid droplets grown at $T = 200, 300, 350^\circ\text{C}$ and $F_{\text{Ga}} = 0.1 \text{ ML/s}$. Here and throughout the paper, Ga and As atoms initially belonging to the substrate are colored red and green, respectively. Ga and As atoms deposited throughout the simulation are colored purple and blue, respectively.

3.3.2 Droplet crystallization

Once the liquid droplet has formed and come to equilibrium, an As flux is introduced to initiate crystallization. Arsenic atoms deposited near or on the droplet diffuse through the liquid quickly [48] and attach most typically near the triple-junction. Such crystallization results in a growing GaAs front and the droplet is crystallized inward. If no nucleation occurs at the vapor-liquid interface and the GaAs fronts coalesce, a fully crystallized quantum dot forms in place of the liquid droplet. Such is the case for moderate temperatures and deposition rate.

Figure 3.8 is a sequence of simulation snapshots illustrating the crystallization of a liquid droplet resulting in a quantum dot (left panel), along with analogous atomic force microscopy images of the GaAs fronts obtained experimentally (right panel). The simulation images in the figure illustrate a typical quantum dot grown at $T = 275^\circ\text{C}$ and $F_{\text{As}} = 0.06$ ML/s. The general trend is that a Ga drop forms once enough Ga has been deposited on the surface (Figure 3.8A), followed by crystallization near the vapor-liquid-solid triple junction upon exposure to As flux (B,C). As crystallization progresses, the liquid Ga is consumed, resulting in a fully crystallized quantum dot (D). Experimental images (Figure 3.8, right panel) were obtained from individual samples prepared according to the procedure outlined above, varying As exposure time. The unreacted liquid Ga was removed from the samples before imaging by selective wet etching [42], thus showing GaAs fronts at various stages during crystallization. The experimental images confirm the growth mechanism observed in the simulations during crystallization.

Besides compact quantum dots, other nanostructures are predicted by the simulation by considering a broad range of temperatures and As fluxes. For example, if F_{As} is sufficiently low or T is sufficiently large, the simulations show the formation of




Figure 3.8: Left panel: model snapshots of liquid droplet crystallization at various times for $T = 275^\circ\text{C}$ and $F_{\text{As}} = 0.06$ ML/s resulting in a compact quantum dot. Ga and As atoms from the original substrate are colored red and green, respectively. Ga atoms deposited to form liquid droplets are colored purple, while As atoms deposited during crystallization are blue. Right panel: AFM images of the GaAs growth fronts in partially crystallized droplets after 10 seconds (top), 40 seconds (middle) and 90 seconds (bottom). Crystallization was obtained at $T = 150^\circ\text{C}$ and 5×10^{-7} Torr As pressure.




Figure 3.9: Nanorings formed at $T = 375^\circ\text{C}$ and $F_{\text{As}} = 0.10, 0.20$ and 0.40 ML/s.

nanorings upon crystallization. The exact morphology of the nanorings is sensitive to the growth conditions. Figure 3.9 shows the morphological dependence of the rings on F_{As} , fixing $T = 375^\circ\text{C}$. At low As deposition rate ($F_{\text{As}} = 0.1$ ML/s), broad and short nanorings form. As F_{As} is increased the nanorings become more compact and taller so that at $F_{\text{As}} = 0.4$ ML/s the resulting structure resembles a “pitted” quantum dot. This compares well to the experimental results in Figure 3.3, e.g. the transition in structure between Figure 3.3c and Figure 3.3f as As_4 BEP is lowered. Figure 3.3c resembles the pitted quantum dot structure of Figure 3.9c. If the BEP is lowered, the resulting nanostructure in Figure 3.3f is a broad and shallow disk surrounding a pit, resembling Figure 3.9a. Similar structures and their dependence on both As deposition rate and temperature have been reported in the literature [49, 50].

3.3.3 Core-shell structures

In addition to quantum dots and nanorings, simulations show the existence of Ga/GaAs core-shell structures. These structures consist of liquid Ga being completely surrounded by GaAs. In the low temperature/high deposition rate regime, the GaAs shells are polycrystalline. However, in a higher temperature/lower flux regime, the simulations show the formation of a shell in registry with substrate. It

will be shown below that the first case is the result of nucleation of GaAs at the vapor-liquid interface whereas the second case results from a Mullins-Sekerka instability of the crystallization growth front.

Nucleation

In the low temperature and high As deposition rate regime nucleation of GaAs clusters near the vapor-liquid is significant. This results in the formation of a polycrystalline GaAs shell surrounding a liquid Ga core, as illustrated in Figure 3.10. In the figure, a liquid droplet grown at $T = 150^\circ\text{C}$ is crystallized by an As flux, deposited at a rate of 0.8 ML/s. Nucleation at the vapor-liquid interface occurs within seconds upon crystallization (Figure 3.10B). The liquid core in the final configuration (C) is completely surrounded by a GaAs shell after 2.4 seconds, preventing any further crystallization of the liquid.

The presence of polycrystalline GaAs, with a high number of grain boundaries and stacking faults that this implies, reduces the optical qualities of the grown sample and is hence is undesirable. Annealing at higher temperature may remove such defects because grain boundaries and stacking faults provide fast diffusion paths for the liquid Ga trapped within the shell. Such paths are accessible at higher temperatures. Moreover, thermal fluctuations of the nuclei can effectively serve to dissolve the shell, which are characteristically thin in this regime. Therefore such configurations may be annealed at a high temperature to remove defects. Indeed, the simulations bear this out. Figure 3.11 shows a quantum dot with polycrystalline shell resulting from the crystallization of a liquid droplet at an As deposition rate of $F_{\text{As}} = 0.80$ ML/s and temperature $T = 150^\circ\text{C}$. The dot is then annealed at a higher temperature $T = 350^\circ\text{C}$. The initial configuration (3.11A) shows the droplet prior to recrystallization. Temperature is then increased, maintaining the same As flux.




Figure 3.10: Simulation snapshots at times $t = 0, 1.8, 2.4$ seconds after crystallization, $T = 150^\circ\text{C}$ and $F_{\text{As}} = 0.8$ ML/s.

Ga atoms move along grain boundaries toward the surface, resulting in a broadening of the dot. The thin shell dissolves, resulting in liquid Ga exposed to As (B, C). Within two seconds, the droplet becomes fully crystallized (D) into a shallow GaAs island, absent of any defects.

The simulations suggest that low-quality polycrystalline structures can be annealed at high temperatures to improve their crystalline quality. Indeed, this has been observed experimentally [51, 52], where high temperature annealing was shown to increase the optical properties of the grown samples. As a consequence, experiments and simulations suggest that high quality nanostructures can be grown and crystallized at lower temperatures, provided an *in-situ* annealing phase is also performed after growth.

Mullins-Sekerka instability

As described above, simulations show that crystallization performed at low temperature and high As flux result in nucleation of polycrystalline GaAs shells surrounding a liquid Ga core, resulting from nucleation of GaAs at the vapor-liquid interface. In the high temperature regime, such surface nucleation does not occur. However, for crystallization at sufficiently high fluxes and high temperature, simulations show the existence of liquid Ga core structures surrounded by GaAs shells in registry with the substrate. As a consequence, such shells are monocrystalline and result from a mechanism separate from surface nucleation. Such structures have been verified experimentally.

By examining the formation of such structures in the simulations, we propose that these shells are driven by an instability at the liquid-solid growth front. Simulation snapshots in Figure 3.12 illustrate the growth mechanism behind this. When this phenomenon occurs no surface nucleation is observed; instead, the growth of

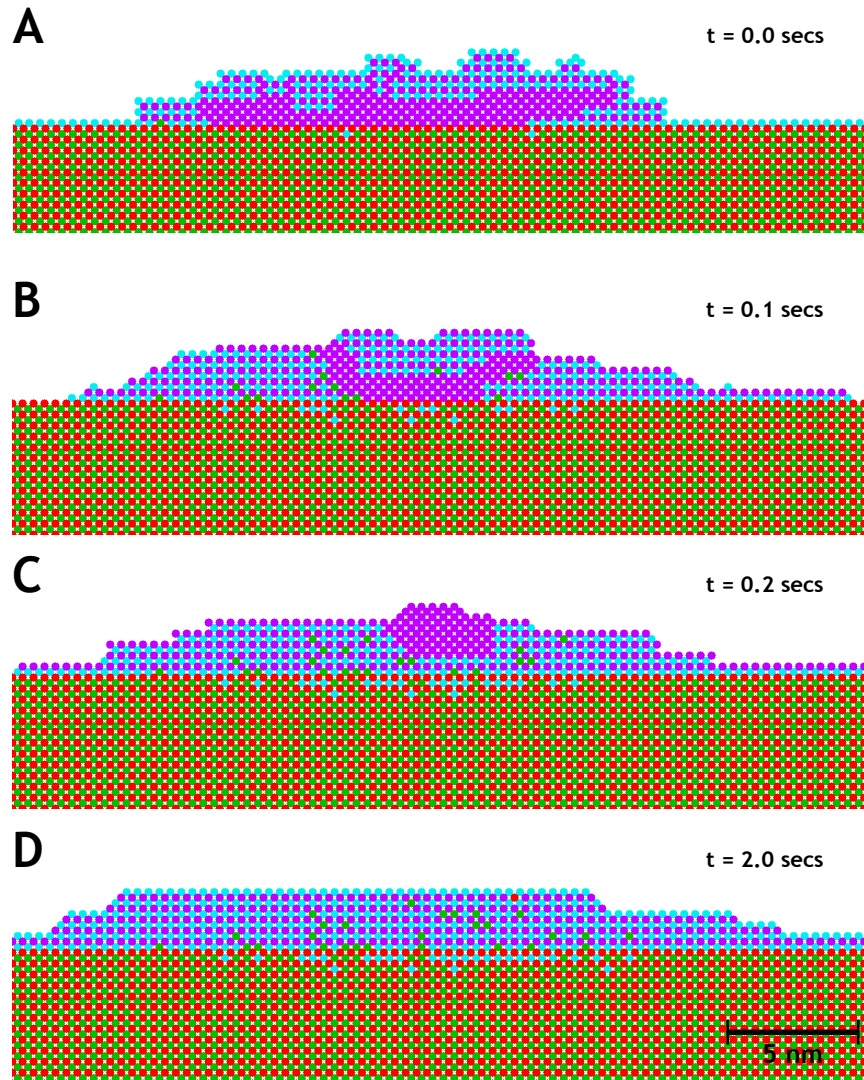


Figure 3.11: Simulation snapshots of a quantum dot annealed at high temperature at time t after temperature was increased. (A) The dot after exposure to As deposition at $F_{\text{As}} = 0.80$ ML/s and temperature $T = 150^\circ\text{C}$. This results in a polycrystalline GaAs shell trapping a liquid Ga core. (B)-(D) Temperature is increased to $T = 350^\circ\text{C}$ and the atoms rearrange in order to fully crystallize the liquid core.




Figure 3.12: Snapshots of liquid core formation at times $t = 1.4, 1.8$ and 2.0 seconds after crystallization at temperature $T = 350^\circ\text{C}$ and deposition rate $F_{\text{As}} = 1.0$ ML/s.

the GaAs front undergoes an instability at the liquid-solid interface characterized by unstable undulations of the solid growth front. Such instabilities grow along the droplet-vapor surface until they have completely surrounded the liquid Ga subsequently preventing further crystallization.

This behavior suggests the presence of a Mullins-Sekerka (MS) instability during crystallization [53, 54], which implies that only perturbations of sufficiently large wavelength experience unstable growth. Therefore, droplets must be sufficiently large to accommodate perturbations of appropriate wavelength in order for the instability to manifest. The effect of temperature, As flux and droplet radius have on the presence of the MS instability is given in Section 3.4.

It is important to note that simulations suggest that, unlike those formed by surface nucleation, the core-shell structure formed due to the MS instability cannot

readily be annealed and recrystallized into an epitaxial nanodot. In this case, shell is in registry with respect to the substrate. Therefore, there are no defects along which liquid Ga can move so that they may crystallize upon exposure to the As vapor. Crystallization of such liquid cores may still occur, however, if the shell is thin enough to dissolve upon annealing thus exposing the liquid Ga to the As flux, however instability-induced shells are characteristically thicker than those resulting from surface nucleation according to simulations.

3.4 Morphological dependence on temperature and flux

Both simulations and experiments have shown that droplet epitaxy can result in a wide range of morphologies depending on the growth conditions. By varying F_{As} and T , we have established their effect on the resulting nanostructures observed in simulations. From this data, it is clear that compact islands form at low temperature and, as the As overpressure is reduced or the temperature is increased the nanostructures become rings. In cases of large F_{As} , core-shell structures are observed. The simulation results are summarized in the structural map given in Figure 3.13. In this section, we will appeal to physical and mathematical arguments to further explain the simulation results. In particular, the solid lines in Figure 3.13 that delineate the morphological structure will be derived in this section.

3.4.1 Three fundamental processes in droplet crystallization

We argue that various morphologies observed both in experiments and simulations can be explained in the context of three key processes active during nanostructure formation. In the first Ga atoms in the liquid drop will be “wicked” out of the droplet onto the substrate by capillary-type forces when exposed to an As overpressure. These forces arise as it is energetically favorable for As atoms on the surface to




Figure 3.13: Morphological Dependence on Growth Conditions. This nanostructural phase map summarizes simulation results of droplet epitaxy and crystallization at various As deposition rates and temperatures. The three boundary curves indicate theoretically derived critical conditions delineating the simulation results and obtained in this section. The left-most, black line is given by Equation (3.18). The middle, red curve corresponds to Equation (3.35). The right-most, blue curve is given by Equation (3.23).




Figure 3.14: Schematic of kinetic processes that determine nanostructural development.

become fully coordinated with Ga atoms. Next, As atoms deposited near or on the droplet diffuse rapidly [48] through the liquid and attach on a growing GaAs front at the liquid-solid interface, crystallizing the droplet epitaxially. Finally, As atoms may also nucleate near the vapor-liquid interface. These three processes: wicking, crystallization and nucleation are illustrated in Figure 3.14. The relative rates at which these processes occur depend on growth conditions and will determine the resulting morphology.

3.4.2 Quantum dots and nanorings

The formation of nanorings is a competition between the wicking and crystallization processes. The crystallization process results in GaAs fronts that grow until the liquid Ga is consumed. If the fronts coalesce before this occurs, the resulting nanostructure is a compact quantum dot. If however, the liquid Ga is consumed before the fronts come together, nanorings result.

To compare the rates of the two processes, we establish expressions for the velocity of the GaAs front as well as the rate at which Ga atoms are wicked out of the droplet. Assuming quasi-static deposition, the velocity v_n of the GaAs front under an As deposition rate F_{As} scales according to that rate:

$$(3.13) \quad v_n = g\ell F_{As},$$

where ℓ is the atomic spacing of the lattice and g is factor that depends shape of the

liquid region and possible difference in the As adsorption probability between the droplet itself and its surroundings. For simplicity, $g = 1$. In the wicking process, Ga atoms are driven from the droplet by the deposition of As on the surface away from the droplet. The time scale of the wicking process is determined by the deposition rate of As

$$(3.14) \quad \tau = \frac{1}{F_{\text{As}}},$$

while the diffusion length scale is given by

$$(3.15) \quad \lambda = \sqrt{\frac{D'_{\text{Ga}}}{F_{\text{As}}}},$$

where D'_{Ga} is the diffusion coefficient of Ga diffusing on a mostly Ga terminated surface, in the presence of As deposition:

$$D'_{\text{Ga}} = \ell^2 R_0 \exp\left[-\frac{E'_{\text{Ga}}}{k_B T}\right].$$

The energy barrier $E'_{\text{Ga}} = 1.26$ eV describing the diffusion of Ga on surface not purely Ga terminated is obtained from the harmonic average of the diffusion rates of Ga-on-Ga diffusion and Ga-on-As diffusion. By prescribing a constant value E'_{Ga} that describes the effective energetic barrier for diffusion on a mostly Ga-terminated surface, we have assumed that the diffusivity throughout the domain is constant. In reality, diffusivity is spatially dependent due to differences in surface reconstruction near and away from the liquid [55]. This dependence is exaggerated in extreme conditions for temperatures less than 150°C and high As overpressure, but away from such growth conditions, a constant effective diffusivity is a reasonable approximation

Equations (3.14) and (3.15) yield a velocity v_w for the wicking of Ga atom from the liquid droplet

$$(3.16) \quad v_w = \frac{\lambda}{\tau} = \ell \sqrt{R_0 F_{\text{As}}} \exp\left[-\frac{E'_{\text{Ga}}}{2k_B T}\right].$$

A fully crystallized quantum dot forms under conditions where the crystallization process is dominant ($v_n \gg v_w$). However, if the wicking process is sufficiently fast ($v_n \sim v_w$), then as described above, the fronts may fail to coalesce, resulting in a nanoring. The critical configuration separating these two scenarios occurs when the crystallization fronts are tangent to each other. To express this critical condition quantitatively, consider the volume of unconsumed liquid Ga at time t :

$$(3.17) \quad V(t) = V_0 - 2\ell v_w t - 2\rho_1(v_n t)^2,$$

where $\rho_1 = \frac{\pi}{4}$ is a shape constant describing the geometry of the crystallization front (which we model as a quarter-circle) and

$$V_0 = \rho_2 r^2,$$

is the initial volume of the liquid droplet of radius r , $\rho_2 = \frac{3\pi}{4}$ being the shape constant that describes the geometry of the droplet. As a first-order approximation, the value of ρ_2 is selected as the average of the shape constant describing a circular and half-circular droplet.

The positive root t_f of equation (3.17) describes the time when all Ga has been consumed. The length of the crystallized front is then given by $v_n t_f$, and the critical condition may be expressed as

$$v_n t_f = r$$

Using equations (3.4), (3.13), and (3.16), this critical configuration can be written in terms of the growth conditions as

$$(3.18) \quad F_{\text{As}}^r = \left(\frac{2F_{\text{Ga}}^\alpha}{(\rho_2 - 2\rho_1)gr_0} \right)^2 R_0^{1-2\alpha} \exp \left[\frac{2\alpha E_{\text{Ga}} - E'_{\text{Ga}}}{k_B T} \right].$$

If $F_{\text{As}} > F_{\text{As}}^r$ then compact structures will result on other hand if $F_{\text{As}} < F_{\text{As}}^r$ nanorings will result. Figure 3.9 illustrates the transition from nanorings to compact quantum

dots with increasing F_{As} . The black line in Figure 3.13 shows a plot of F_{As}^r vs. $1/T$. The agreement with the simulation data is quite good – it separates compact structures (black dots in the figure) from nanorings (black circles).

A crucial assumption has been made in the argument presented here, namely that no nucleation takes place within the Ga droplet and the crystallizing front moves in a stable fashion (i.e. no Mullins-Sekerka instability occurs). In what follows we shall examine in what regimes in parameter space these factors play an important role.

3.4.3 Nucleation of polycrystalline shells

Here we will demonstrate that when the As deposition rate is sufficiently large and temperature is low enough, the dominant process will be nucleation at the vapor-liquid interface. As a consequence, the wicking process makes a negligible contribution to the morphology. The rate of the nucleation process at the vapor-liquid interface may be estimated by considering the As concentration $c_{\text{As}}(x, y, t)$ within a domain of liquid Ga in contact with a GaAs substrate and As flux. For simplicity we will consider the domain to be rectangular of height H with periodic boundary conditions in the horizontal direction. Assuming quasi-static deposition at a rate F_{As} ML/s and temperature T , the concentration satisfies

$$(3.19) \quad \begin{aligned} \nabla^2 c_{\text{As}} &= 0, \\ \ell D_{\text{As}} \left. \frac{\partial c_{\text{As}}}{\partial y} \right|_{y=H} &= F_{\text{As}}, \\ c_{\text{As}}|_{y=0} &= c_0 \end{aligned}$$

where c_0 is the equilibrium As concentration above a flat liquid-solid interface given in Equation (3.6) and the diffusion coefficient

$$D_{\text{As}} = \ell^2 R_0 \exp \left[-\frac{\epsilon_{\text{ex}}^{\text{L}}}{k_B T} \right],$$

describes the diffusion of As through liquid Ga, $\epsilon_{\text{ex}}^{\text{L}} = 0.7$ eV being the energy barrier for diffusion throughout the liquid droplet. The system has a linear solution

$$(3.20) \quad c_{\text{st}} = \nu y + c_0,$$

where $\nu = \frac{F_{\text{As}}}{\ell D_{\text{As}}}$. Substituting the expressions for D_{As} and c_0 above we obtain the equilibrium concentration c_{surf} at the surface of the domain

$$(3.21) \quad c_{\text{surf}} = \frac{F_{\text{As}} H}{\ell^3 R_0} \exp\left[\frac{\epsilon_{\text{ex}}^{\text{L}}}{k_B T}\right] + \ell^{-2} \exp\left[\frac{E_{\text{attach}} - E_{\text{detach}}}{k_B T}\right].$$

Nucleation is most likely to occur where the As concentration is the largest, near the vapor-liquid interface. This means that nucleation will occur when c_{surf} is larger than some critical concentration:

$$(3.22) \quad c_{\text{surf}} \geq c_0 \exp\left[\frac{E_{\text{nuc}}}{k_B T}\right],$$

where $E_{\text{nuc}} = 0.01$ eV is the nucleation barrier of GaAs in liquid Ga, treated as a fitting parameter. Replacing H in Equation (3.21) with the expression for the droplet radius in Equation (3.4), the critical condition (3.22) for surface nucleation can be expressed in terms of T and F_{As} as:

$$(3.23) \quad F_{\text{As}}^n = \frac{R_0^{1-\alpha} F_{\text{Ga}}^\alpha}{r_0} \exp\left[\frac{E_{\text{attach}} - E_{\text{detach}} - \epsilon_{\text{ex}}^{\text{L}} + \alpha E_{\text{Ga}}}{k_B T}\right] \left(\exp\left[\frac{E_{\text{nuc}}}{k_B T}\right] - 1\right).$$

If the As deposition rate exceeds F_{As}^n then GaAs crystallites will form at the vapor-liquid interface, as illustrated in Figure 3.10. The blue line in Figure 3.13 is a plot F_{As}^n as function of $1/T$. It accurately predicts the presence of nucleation in simulation results (blue triangles).

3.4.4 Monocrystalline shells and a Mullins-Sekerka instability

In the case where the rates of both the wicking and nucleation processes are negligible, the crystallization process dominates, resulting in the growth of crystallization

fronts at the triple point and in registry with the substrate. As observed in Section 3.3.3, this growth can be unstable due to a Mullins-Sekerka instability leading to GaAs shells epitaxial to the substrate surrounding a liquid Ga core. Perturbations to the planar growth front of sufficiently long wavelength experience this instability.

We show the existence of this instability by a linear perturbation analysis of the system (3.19). The normal velocity of the planar growth front is given by

$$(3.24) \quad v_n = \ell^2 D_{\text{As}} (\nabla c \cdot \mathbf{n})|_{y=0}.$$

Using the expression (3.20), along with the above expression for the normal velocity, we may write $v_n = \eta\nu$, where $\eta = \ell^2 D_{\text{As}}$.

We consider a perturbation of the interface at $y = 0$ of the form

$$h(x, t) = h_0 \sin(kx) e^{\omega t}.$$

For convenience, we describe the concentration as $c(x, y, t) = c_{st}(y) + u(x, y, t)$, where

$$u(x, y, t) = u_k(y) \sin(kx) e^{\omega t}.$$

Then, from the system (3.19), we conclude that

$$(3.25) \quad \nabla^2 u = 0,$$

with boundary condition

$$(3.26) \quad u_y|_{y=H} = 0.$$

We may linearize the boundary conditions at the perturbed interface near $y = 0$. By the Gibbs-Thomson relation [53], we may write

$$(3.27) \quad c|_{y=h(x)} = c_0 \exp \left[\frac{\ell^2 \gamma \kappa(x, t)}{k_B T} \right]$$

$$(3.28) \quad \approx c_0 + \sigma \kappa(x, t)$$

$$(3.29) \quad \approx c_0 + \sigma |h_{xx}|$$

where γ is the interface energy per unit length, $\kappa(x, t)$ is the curvature of the height profile and

$$\sigma = c_0 \frac{\ell^2 \gamma}{k_B T}.$$

The value $\gamma = 0.1$ can be obtained directly from the model by considering the energy difference per unit length between an infinite GaAs crystal and a configuration obtained by introducing a Ga liquid interface through the crystal. Specifically, it is the value of an As-As next-nearest neighbor bond that is lost when the liquid interface is introduced. The left hand side of the boundary condition at $y = 0$ in (3.19) may be linearized as

$$(3.30) \quad c|_{y=h(x)} = c_{st}(h(x)) + u(x, h(x), t),$$

$$(3.31) \quad \approx c_{st}(0) + \left. \frac{dc_{st}}{dy} \right|_{y=0} h(x) + u(x, 0, t) + \left. \frac{\partial u}{\partial y} \right|_{y=0} h(x)$$

$$(3.32) \quad = c_0 + \nu h(x) + u(x, 0, t),$$

where we have ignored the higher order term $\left. \frac{\partial u}{\partial y} \right|_{y=0} h(x)$. The linearized version of the boundary condition is obtained by equating (3.29) and (3.32):

$$(3.33) \quad u(x, 0, t) = \sigma |h_{xx}| - \nu h.$$

The normal velocity of the perturbed profile consists of the planar velocity term $\eta\nu$ and the normal velocity h_t of the perturbation, yielding the boundary condition

$$\eta (\nabla c \cdot \mathbf{n})|_{y=h(x)} = \eta\nu + h_t,$$

from which we may remove the planar growth terms to write the boundary condition solely in terms of u :

$$\eta (\nabla u \cdot \mathbf{n})|_{y=h(x)} = h_t.$$

A similar linearization may be performed for this boundary condition near $y = 0$. To first order, $\mathbf{n} = (-h_x, 1)$, so that $\nabla u \cdot \mathbf{n} \approx u_y$, ignoring the $u_x h_x$ term. Then

$$\begin{aligned}\eta (\nabla u \cdot \mathbf{n})|_{y=h(x)} &= \eta u_y|_{y=h(x)} \\ &= \eta u_y|_{y=0}\end{aligned}$$

where the last terms comes from expansion about $y = 0$, ignoring higher order terms. Hence the boundary condition describing the normal velocity may be linearized as

$$(3.34) \quad \eta u_y|_{y=0} = h_t.$$

Together, equations (3.25), (3.26), (3.33), and (3.34) may be solved, resulting in the following relationship between ω and k :

$$\begin{aligned}\omega &= \ell^4 R_0 \exp\left[-\frac{\epsilon_{\text{ex}}^{\text{L}}}{k_B T}\right] k \tanh(kH) \\ &\quad \times \left(\frac{F_{\text{As}}}{\ell^3 R_0} \exp\left[\frac{\epsilon_{\text{ex}}^{\text{L}}}{k_B T}\right] - \frac{\gamma}{k_B T} \exp\left[\frac{E_{\text{attach}} - E_{\text{detach}}}{k_B T}\right] k^2 \right).\end{aligned}$$

Solving for the critical wave number k_c , occuring when $\omega = 0$, and setting $\Lambda_c = \frac{2\pi}{k_c}$, the critical wavelength for the MS instability is given by

$$\Lambda_c = 2\pi \ell^{3/2} \sqrt{\frac{R_0 \gamma}{F_{\text{As}} k_B T}} \exp\left[\frac{E_{\text{attach}} - E_{\text{detach}} - \epsilon_{\text{ex}}^{\text{L}}}{2k_B T}\right],$$

In order to accommodate perturbations that experience the MS instability, droplets necessarily must have radius on the order of Λ_c , i.e. $r \geq C\Lambda_c$, for some constant C . This critical condition may be expressed in terms of F_{As} and T using the above equation along with the model for droplet radius (3.4):

$$(3.35) \quad F_{\text{As}}^{\text{MS}} = C^2 \ell \left(\frac{2\pi F_{\text{Ga}}^\alpha \ell}{r_0} \right)^2 \frac{\gamma}{k_B T} R_0^{1-2\alpha} \exp\left[\frac{E_{\text{attach}} - E_{\text{detach}} - \epsilon_{\text{ex}}^{\text{L}} + 2\alpha E_{\text{Ga}}}{k_B T}\right].$$

If $F_{\text{As}}^{\text{MS}} < F_{\text{As}} < F_{\text{As}}^n$ then liquid cores will form via a Mullins-Sekerka instability.

Figure 3.13 shows a plot of $F_{\text{As}}^{\text{MS}}$ vs $1/T$. The scaling constant $C = \frac{1}{8}$ was selected to

best match simulation results, though its specific value do not affect the qualitative shape of the boundary curve, and in particular its slope. The simulation results when a liquid core was observed are plotted as red triangles. We observe that the theoretical curve slightly underestimates the instability within the simulations. This underestimation is inherent to our model and can be attributed to discrete effects. Such effects on nucleation and hence instabilities of the type outlined above are indeed well-studied in the context of diffusion limited aggregation [56, 57].

CHAPTER IV

Further applications: nanowire growth by the Vapor-Liquid-Solid method and sintering of granular material

4.1 The Vapor-Liquid-Solid method for nanowire growth

As another application of the model and implementation, we have simulated nanowire growth by the Vapor-Liquid-Solid (VLS) method. In VLS growth, a liquid catalyst droplet is placed on a solid substrate. Vapor material is then introduced and is subsequently adsorbed onto the vapor-liquid interface. As the vapor is incorporated into the droplet, the liquid serves to catalyze the reaction in which vapor is converted to solid material. This solid material then nucleates at the liquid-solid interface, resulting in one-dimensional growth underneath the droplet. Common materials used in nanowire growth are silicon and liquid gold droplet as a catalyst [58]. A combination of Silane (SiCl_4) and hydrogen gas composes the vapor material. Within the gold material, the catalyzed reaction



results in the precipitation of Si, leading to its nucleation at the liquid-solid interface.

Experimental work in nanowires grown by the VLS method was introduced in [59]. Since then, the characterization of nanowire formation and morphology has

been studied in experiments, and a wide range of phenomena have been observed. Examples of such phenomena include nanowire tapering and a change in growth direction, called kinking. For example, the phenomenon of nanowire tapering in Si nanowires grown by the VLS method was studied in [60]. There, it was observed that suppressing the diffusion of the liquid catalyst by the introduction of oxygen gas resulted in a negligible amount of tapering. In another study, the appearance and magnitude of tapering was shown to be controlled by vapor flux and temperature [61] in the case of self-catalyzed nanowires. Kinking in GaAs nanowires can also be controlled by a consideration of the deposition ratio between Ga and As [62]. Experiments have also exhibited an explicit control of nanowire kinking by vapor pressure to produce multiple kink sites on Si nanowires [63].

Some previous work regarding the simulation and modeling of nanowire growth used continuum models [64,65]. In [64], the phenomena of nanowire kinking, wherein the nanowire is observed to change growth directions, was studied in the context of faceting at the liquid-solid interface and liquid droplet statics. There, external perturbations are added to the growth of the nanowire to drive such kinking. In the KMC model, such perturbations are built-in, and we show that vulnerability to unstable growth as a result of these perturbations is governed by certain energetic parameters. Such unstable growth results in nanowire kinking, consistent with the theory in the continuum model. There has also been previous work in the Monte Carlo simulation of nanowires [66], which focused on the effect of experimental parameters such as temperature and deposition rate. Here we take a complimentary approach, studying the effect of energetic parameters on growth kinetics and mechanisms.

In our simulations, we consider a system on a hexagonal lattice initially consisting of a liquid droplet of material L on a solid substrate of material S exposed to

γ (eV)	V	L	S
V	0.10	0.10	0.10
L	0.10	0.35	0.40
S	0.10	0.40	0.50

Table 4.1: Pairwise bonding energies between the vapor, liquid and solid species used in the nanowire simulations.

a vapor of material V. While not tuned to a specific physical system, the pairwise bonding energies in Table 4.1 and used in the simulations exhibit reasonable qualitative properties. Namely the model specifies strong bonds (0.5 eV) within the solid material, relatively weaker bonds (0.40 eV) between liquid material and even weaker bonds (0.35 eV) bonds between the solid and liquid materials to encourage phase segregation. The vapor phase is considered weakly interacting with other phases, and all bonds involving a vapor atom are assigned a bonding energy of 0.10 eV. Only vapor atoms are allowed to desorb, and the additional energy barrier for desorption of vapor atoms was set to $\mu_V = 0.50$ eV.

The transition $X \rightarrow Y$ between configurations X, Y representing a vapor to solid reaction of a single atom is assigned an activation energy barrier

$$(4.2) \quad E_a(X, Y) = E(X) - E(X \wedge Y) + \rho(V, S),$$

where $X \wedge Y$ is defined by replacing the reacting vapor atom with one of intermediate species $V \wedge S$ and $\rho(V, S)$ is the additional barrier for a reaction event. As with the exchange barrier $\epsilon_{\text{ex}}^{\text{NL}}$, we may assign different values ρ^{L} and ρ^{NL} to ρ depending on whether the reaction occurs within either a liquid neighborhood or otherwise, respectively. In the former case, the total activation energy barrier is

$$\begin{aligned} E_a &= E(X) - E(X \wedge Y) + \rho^{\text{L}} \\ &= 6\gamma(V, L) - 6\gamma(V \wedge S, L) + \rho^{\text{L}} \\ &= 3(\gamma(V, L) - \gamma(L, S)) + \rho^{\text{L}}. \end{aligned}$$

The energy barriers ρ^L and ρ^{NL} describe the additional energy barrier for reactions contributing to vapor-liquid-solid and vapor-solid growth, respectively. The case when $\rho^{NL} > \rho^L$ therefore models the catalytic effect of the liquid droplet on the vapor-solid reaction, which is essential in VLS growth. Detailed modeling of the different reaction pre-cursors and resultants (e.g. Cl and H in the reaction (4.1)) are not included in the model.

As an illustrative study of nanowire formation mechanisms, we examine the role of ρ^L and $\epsilon_{\text{ex}}^{NL}$, the additional energy barrier for atom exchanges at the liquid-solid interface. Two sets of simulations were performed. First, growth simulations were run to study the long-term growth and macroscopic properties of nanowires over several values of ρ^L and $\epsilon_{\text{ex}}^{NL}$. The initial profile consisted of a hemispherical liquid droplet of diameter 64 atoms placed on a flat solid substrate. The domain itself was 512 atoms wide, with periodic boundary conditions. Vapor material was deposited uniformly throughout the domain until 512 ML of atoms were deposited, though much of the material desorbs before the vapor-solid reaction occurs. The second set of simulations were performed to study the trajectory of individual solid atoms resulting from a vapor-solid reaction and yield microscopic insight on the mechanisms leading to the macroscopic properties observed in the first set of simulations. Here, the initial configuration is a nanowire grown as described above. Vapor atoms are then deposited as in the growth simulations, and the trajectory of the first solid atom resulting from a vapor-solid reaction is maintained until that atom becomes static, signifying its incorporation onto the wire. These trajectory simulations were run over several values of ρ^L and $\epsilon_{\text{ex}}^{NL}$, and for each choice of parameters, 512 independent and identical trials were run in order to get meaningful statistics by ensemble averaging. All simulations occurred at fixed temperature $T = 350^\circ\text{C}$ and vapor deposition rate




Figure 4.1: Simulation results of VLS nanowire growth, varying ρ^L between 0.75 and 1.25 eV. Liquid material is colored green, original solid substrate materials is purple. Material that was deposited from the vapor and subsequently reacted to solid phase is blue.

$$F_V = 0.25.$$

4.1.1 Catalytic reaction rate

The parameter ρ^L determines the energy barrier for, and hence the rate of a catalyzed vapor to solid reaction. Figure 4.1 shows growth simulation results for various values of ρ^L . In the case of low ρ^L (0.75 eV), a large, broad nanowire results in contrast with the high ρ^L (1.25 eV) case, in which no growth is observed. Hence the parameter ρ^L affects the growth rate and width of the nanowire. Additionally, significant tapering is observed when ρ^L is small. Such tapering has been reported in experimental literature [67,68]. Axial and radial growth is present in the case where $\rho^L = 0.75$ eV. When $\rho^L = 1.00$ eV, axial growth is present, but not as significant as the previous case, and radial growth is negligible. The presence and control of tapering is in qualitative agreement with what is seen in experiments, as described above.

4.1.2 Liquid-solid interface mobility

The parameter $\epsilon_{\text{ex}}^{\text{NL}}$ affects events that occur at the liquid-solid interface, and hence describes the mobility of that interface. Among the processes that depend on $\epsilon_{\text{ex}}^{\text{NL}}$ are attachment onto, detachment from and diffusion along the liquid-solid interface. Assuming a perfectly flat interface, the associated barriers for these events are

$$\begin{aligned} E_{\text{attachment}} &= \epsilon_{\text{ex}}^{\text{NL}} + 2\gamma(L, S) - \gamma(L, L) - \gamma(S, S) \\ &= \epsilon_{\text{ex}}^{\text{NL}} - 0.2\text{eV}; \\ E_{\text{detachment}} &= \epsilon_{\text{ex}}^{\text{NL}} - 2\gamma(L, S) + \gamma(L, L) + \gamma(S, S) \\ &= \epsilon_{\text{ex}}^{\text{NL}} + 0.2\text{eV}; \\ E_{\text{diffusion}} &= \epsilon_{\text{ex}}^{\text{NL}}. \end{aligned}$$

Changing $\epsilon_{\text{ex}}^{\text{NL}}$ therefore scales the rates for these three processes by the same amount. In the context of other processes, however, the effect of $\epsilon_{\text{ex}}^{\text{NL}}$ can be seen to alter the resulting nanowire growth modes. In particular the presence of nanowire kinking, which is observed experimentally as described above, is governed in part by this parameter. Figure 4.2 shows simulation results for various values of $\epsilon_{\text{ex}}^{\text{NL}}$, the parameter that controls atom-atom exchanges at the liquid-solid interface. For a small value $\epsilon_{\text{ex}}^{\text{NL}} = 0.90$ eV, the wire grows perpendicular to the substrate and the liquid-solid interface is flat. There is more significant nanowire kinking in the case when $\epsilon_{\text{ex}}^{\text{NL}} = 1.10$ eV. In this case, the liquid-solid interface is not flat asymmetrically faceted. When $\epsilon_{\text{ex}}^{\text{NL}} = 1.40$ eV there is once again little nanowire kinking, but the interface is faceted. In this case, the faceting is symmetric.




Figure 4.2: Simulation results of nanowire growth, for $\epsilon_{\text{ex}}^{\text{NL}} = 0.9, 1.10$ and 1.40 eV.

Diffusion-length scale and interface geometry

The growth simulations suggest that the geometry of the liquid-solid interface affects the growth direction of a nanowire. By examining the trajectory statistics, we may gain some insight on this role. Figure 4.3 show typical trajectories for $\epsilon_{\text{ex}}^{\text{NL}} = 0.90, 1.10$ and 1.40 eV, respectively. In the case of low $\epsilon_{\text{ex}}^{\text{NL}}$, we observe that a solid atom diffuses through the liquid droplet and upon initial attachment to the liquid-solid interface, may diffuse along that interface. When $\epsilon_{\text{ex}}^{\text{NL}}$ is increased however, the diffusion along the interface is limited. When a solid atom attaches onto the liquid-solid interface, it is static, indicating its incorporation into the droplet. As such, the growth of the interface tends toward diffusion-limited behavior as $\epsilon_{\text{ex}}^{\text{NL}}$ is increased. For large values of $\epsilon_{\text{ex}}^{\text{NL}}$ (e.g. $\epsilon_{\text{ex}}^{\text{NL}} = 1.40$ eV) the attachment rate is small, so that a solid atom spends most of its time diffusing and such trajectories are characteristically long. Indeed, the figure suggests that the solid atom visits the entirety of the liquid droplet before becoming static, avoiding the interface due to a large attachment barrier.




Figure 4.3: Snapshots of trajectories for solid atoms inside the liquid droplet for $\epsilon_{\text{ex}}^{\text{NL}} = 0.9, 1.1$ and 1.4 eV, over an interval of 0.1 seconds. Each straight line segment corresponds to 9.5×10^{-8} seconds.

Nanowire kinking

The simulation snapshots in Figure 4.2 suggests that the geometry of the liquid-solid interface plays a role in nanowire growth direction. In the case of a flat interface or growth along the vapor-liquid-solid triple junction, the resulting nanowire grows perpendicular to the substrate, while nanowire kinking is observed in the presence of asymmetric faceting. Figure 4.4 shows another example of nanowire kinking. By examining this nanowire simulation, we can observe the evolution of its configuration during the change in growth direction. Simulation snapshots of this process are illustrated in Figure 4.5, along with enlarged images of the liquid-solid interface. There, we see the nucleation of a horizontal facet (left most panel) and its growth. In the third panel, the horizontal facet has become the largest along the interface, and at that stage the nanowire has changed growth directions.

We can view the kinking phenomenon as originating from the inability of the diffusion and detachment/attachment processes at the liquid-solid interface to effectively relax and flatten nucleation along that interface within a reasonable time-scale. This results in the nucleation and growth of new facets. The roughness of this interface is characterized by the diffusion length scale along the interface, which we have seen to be affected by $\epsilon_{\text{ex}}^{\text{NL}}$. When $\epsilon_{\text{ex}}^{\text{NL}}$ is small (e.g 0.9 eV), the interface remains flat, and no kinking is observed. When this is increased however, the the growth of new facets




Figure 4.4: An example of nanowire kinking.




Figure 4.5: Snapshots of nanowire kinking and facet nucleation. From left to right: (1) Start of nanowire kinking, with flat liquid-solid interface. (2) The nucleation of a new facet. (3) Growth of the facet and droplet motion. (4) Nucleation of another facet.

becomes accessible, and kinking is observed.

4.1.3 Computational statistics

Over the entire set of growth simulations presented in this section, the average number of Monte Carlo steps required for a full simulation was 28 billion steps, while the maximum number was 117 billion steps. The corresponding average and maximum CPU times were 10 and 45 hours respectively.

4.2 Sintering of granular material

As another application of and illustration of the flexibility of the KMC model, we simulate sintering, a process in which a powder of material is heated below its melting point, allowing the powder to coalesce. This process may be modeled as the evolution of porous granular material; the grains signifying regions of identical crystallographic orientation. As such, the model is not unlike atomistic grain growth ones such as the Monte-Carlo Potts model [69]. The fundamental difference is in the presence of vacancies within the material, resulting in the contribution of surface diffusion along the interface of such vacancies toward the evolution of the system.

Previous work in simulating sintering and growth models span a broad range of length and time scales. Molecular dynamics simulations of sintering often involve small number of nanoclusters of specific materials such as tungsten and TiO_2 [70,71]. In [72], a kinetic Monte Carlo simulation of sintering involving only the surface diffusion of atoms between nanoclusters was studied. Larger scale simulations that include surface diffusion and grain growth were studied in [73]. There the simulations were done using Metropolis dynamics. Atomistic grain growth models that do not incorporate surface diffusion has been simulated using the Monte-Carlo Potts model and can be performed in a domain larger than the previous examples. Two

dimensional simulations in a domain consisting of 200×200 lattice sites have been performed [74, 75], over a time scale of up to 7000 Monte Carlo steps. More recent work simulated the model in three dimensions, within a domain of size $250 \times 250 \times 250$ lattice sites, over a time scale of 1000 steps. In comparison, continuum models of grain growth reach even larger time scales. There, systems with up to 670,000 initial grains can be simulated over a time scale in which only 4000 grains remain [33, 76, 77].

In this dissertation, we simulate porous granular material in a 512×512 atom domain, with over 8000 initial grains over a period of one hundred to ten thousand simulated seconds, corresponding to billions of Monte Carlo steps. In the simulations, each atom is assigned one of S species, indicating a particular orientation of the atom. Pairwise bonding energies between atoms of species σ_i and σ_j , are given by

$$\gamma(\sigma_i, \sigma_j) = \begin{cases} \gamma_0 & i \neq j, \\ \gamma_1 & i = j. \end{cases}$$

Throughout this section, we fix $\gamma_0 = 0.25$ eV and $\gamma_1 = 0.50$ eV, encouraging atoms to phase segregate. Only surface diffusion and reaction events (here seen as changes in orientation) are allowed. Unlike the other two example systems we have previously considered, no special neighborhood is assigned to signify a liquid phase. As such, only one reaction energy barrier ρ is specified. As before, the intermediate state for a reaction event in which an atom changes from orientation σ_i to σ_j is obtained by replacing that atom with one of intermediate species $\sigma_i \wedge \sigma_j$ and bonding energies involving this intermediate species are calculated according to Equation (2.3).

As an example, we consider the event $X \rightarrow Y$ where an atom of species σ changes orientation to species τ . Suppose that atom has n neighbors of species σ and the remaining (assuming a hexagonal lattice) $6 - n$ neighbors are of species τ . Such a reaction would occur on the boundary between two grains of orientation σ and τ , and

signifies an attempted advance of the τ grain into the σ grain. Then the activation energy of this reaction event is

$$\begin{aligned} E_a &= E(X) - E(X \wedge Y) + \rho \\ &= n\gamma(\sigma, \sigma) + (6 - n)\gamma(\sigma, \tau) - n\gamma(\sigma \wedge \tau, \sigma) - (6 - n)\gamma(\sigma \wedge \tau, \tau) + \rho \\ &= (n - 3)(\gamma_1 - \gamma_0) + \rho, \end{aligned}$$

and the corresponding rate may be written as

$$R(X, Y) = \omega \exp \left[-\frac{n\tilde{\gamma}}{k_B T} \right],$$

where $\tilde{\gamma} = \gamma_1 - \gamma_0$ and

$$\omega = R_0 \exp \left[\frac{3\gamma - \rho}{k_B T} \right].$$

This rate is similar to that given in the Monte Carlo Potts model for grain growth, with interfacial energy $\tilde{\gamma}$ and pre-exponential factor ω . One key difference is in the probability of selecting such an event. Typical Monte Carlo Potts models use the Metropolis selection criterion based on the energy difference between states X and Y . In that model, an atom is selected uniformly to change orientation, and this change is accepted with probability 1 if it results in a decrease of energy. Otherwise, it is kept with probability proportional to that of $R(X, Y)$ above. In doing so, the Monte Carlo Potts model ignores kinetic time scales, which is acceptable in the case where only one type of event may occur, namely orientation changes. With the KMC model, the rate of such orientation changes can be put in the context of the rates of other kinetic processes, such as surface diffusion.

4.2.1 Simulation Results

The simulations were performed on a 512×512 hexagonal lattice with periodic boundary conditions in both directions. The initial conditions were obtained by




Figure 4.6: The initial configuration of the sintering simulations presented in this paper, obtained by a random sphere packing algorithm. The material density is 80%. Different colors correspond to different orientations.

a random sphere packing algorithm with material density specified at 80%, and is depicted in Figure 4.6. Each color in the figure represents one of $S = 7$ orientations an atom may possess. Simulations were performed to study the effect of ρ and temperature T on the coarsening time scale of the system.

The effect of ρ

In the first series of simulations, the temperature T was fixed at 327°C, while the parameter ρ was allowed to vary between 1.25 eV and 2.00 eV. The system was then annealed for a total of 100 simulated seconds. In the high ρ regime, surface diffusion was the dominant process, consisting of 99.9% of the events when $\rho = 2.00$ eV. This percentage drops to 70.8% when $\rho = 1.25$ eV. The final configuration for each choice of ρ is depicted in Figure 4.7. The figure suggests that the time scale for coarsening




Figure 4.7: Final configurations of sintering simulations after 100 seconds at $T = 327^\circ\text{C}$, for $\rho = 1.25, 1.50, 1.75$ and 2.00 eV. Different colors correspond to different orientations.

increases along with increasing ρ . Qualitatively, the amount of coarsening is more significant in the cases when $\rho = 1.25$ and 1.50 eV than for the remaining two cases.

The utility of sintering is in recrystallizing material from a powder in order to grow material with a small amount of grains, minimizing the effects of grain boundary defects. As such, much experimental and analytical work on sintering and grain growth have focused on coarsening statistics, and in particular mean grain radius $\langle R \rangle$. Such works have demonstrated the fact that the growth of this radius with respect to time scales according to a power law $\langle R \rangle \sim t^b$. For grain growth models that do not model the presence of pores in the material, and hence do not incorporate surface diffusion, analysis predicts an exponent value $b = \frac{1}{2}$ [74]. However, several

experimental results yields a power law in which b is between $\frac{1}{4}$ and $\frac{1}{2}$. For example, in [78], the recrystallization of lead yielded an exponent of $b = 0.4$. In [79], the sintering of cadmium oxide powder results in an exponent of $b = \frac{1}{3}$.

The plots in Figures 4.8 show two statistics of interest with respect to coarsening. The left panel presents the number of grains within a configuration as a function of simulation time, over the values of ρ considered above. The plot shows a decay in this number for all values of ρ , though the rate of decay is more significant for $\rho \leq 1.50$. For such values of ρ , the number of grains quickly decay from 8600 to under 2000 within the first ten seconds of the simulations, and for these values of ρ , we observed qualitatively from the snapshots in Figure 4.7 that the amount of coarsening is more substantial than the case when $\rho > 1.50$. The right panel plots the mean equivalent disc radius $\langle R \rangle$ of the grains within a configuration. That is, $\langle R \rangle$ is the empirical mean value of the quantity $\sqrt{\frac{A}{\pi}}$, where A is the area of a grain, averaged over the entire set of grains within a configuration. The radius behaves according to a power law with respect to time. Along with this empirical average radius are the best fit power law functions of the form

$$\langle R \rangle = at^b + \langle R_0 \rangle,$$

obtained by the non-linear least squares method. The value $\langle R_0 \rangle = 2.25$ unit lengths is the mean radius for the initial condition. For $\rho = 1.25$ and 1.50 , the mean radius scales as $\langle R \rangle \sim t^{0.25 \pm 0.02}$. When $\rho = 1.75$ eV, the best-fit power law yields a value of $b = 0.52$, while $b = 0.21$ for $\rho = 2.00$.

As the number of grains decreases, the system approaches a meta-stable configuration. For example, consider the case where $\rho = 1.25$ eV (blue circles in Figure 4.8), left panel. Within the first 10 seconds of simulation time, the number of grains decreases from over 8000 to under 1000. After 100 seconds, the number of grains

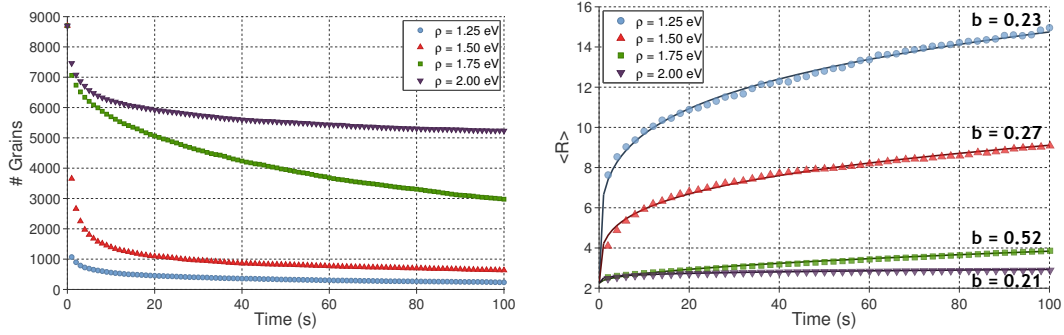


Figure 4.8: Grain coarsening statistics vs. ρ . (Left) The number of grains vs. time. (Right) The mean disc-equivalent radius $\langle R \rangle$ vs. time. Best-fit power law models $\langle R \rangle = at^b + \langle R_0 \rangle$ are shown in dashed lines, and the fitted value b is indicated for each plot. Each plot corresponds to a different value of ρ

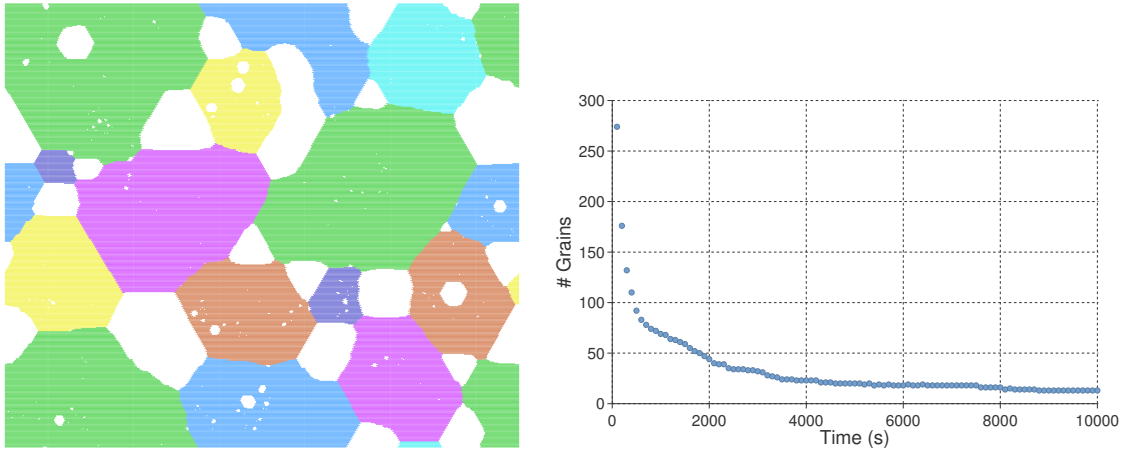


Figure 4.9: Long-term grain coarsening statistics after 10,000 simulated seconds for $\rho = 1.25$ eV and $T = 327^\circ\text{C}$. (Left) Final configuration. (Right) Number of grains vs. time, in the interval between 100 and 10,000 seconds.

decreases to under 500. We may examine the long-term behavior in this case to understand the system's approach to meta-stability. Figure 4.9 depicts the behavior of the system over a duration of 10,000 simulated seconds. The left panel shows the final configuration, characterized by larger grains than the system possessed after 100 seconds. The number of grains within the system over the interval between 100 and 10,000 seconds is shown in the right panel. This graph shows that the number of grains decreases to under 13 grains during this time.

Grain radius and critical phenomena

The utility of sintering is in recrystallizing a powder in order to grow material with a small amount of grains, minimizing the effects of grain boundary defects. As such, much experimental and analytical work on sintering and grain growth have focused on coarsening statistics which describe the rate at which this recrystallization occurs. Of particular interest is in the mean grain radius $\langle R \rangle$. Such work have demonstrated the fact that the growth of this radius with respect to time scales according to a power law $\langle R \rangle \sim t^b$. For grain growth models that do not consider the presence of pores in the material, and hence do not incorporate surface diffusion, analysis predicts an exponent value $b = \frac{1}{2}$ [74]. However, several experimental results yields a power law in which b is between $\frac{1}{4}$ and $\frac{1}{2}$. For example, in [78], the recrystallization of lead yielded an exponent of $b = 0.4$. In [79], the sintering of cadmium oxide powder results in an exponent of $b = \frac{1}{3}$. We observe that the above simulations yield an exponent consistent with the range outlined above and suggest that the presence of pores and surface diffusion along their interfaces provides some explanation regarding the discrepancy between analytic results and experiments.

The plots in Figures 4.8 suggest that the coarsening behavior for $\rho = 1.75$ is different from the other cases. For that case, the average grain radius scales like $t^{1/2}$ while the other cases scale like $t^{1/4}$. We performed further simulations over a larger set of ρ , each time calculating the mean power law exponent b over an ensemble of 16 trials for each choice of ρ . The dependence of b on ρ is plotted in Figure 4.10. We see that for all but a small interval of ρ values, the power law scales with exponent $b \approx \frac{1}{4}$. However, within the interval $1.6 \leq \rho \leq 1.8$ eV, the growth rate increases to near or above $b = \frac{1}{2}$. That is, for this critical interval, grain growth proceeds much faster than for other values of ρ .




Figure 4.10: Power law b vs ρ .

Effect of temperature

In the second series of simulations, the effect of temperature on coarsening was studied in an analogous way to the study above. The energy parameter ρ was fixed at 1.25 eV, and the temperature was varied between 277°C and 427°C. Figure 4.11 shows the configuration after 100 seconds for each choice of temperature. As the temperature is increased, the amount of coarsening is increased, resulting in larger and fewer grains, though a significant amount of coarsening is observed at all temperatures. In the high temperature case $T = 427^\circ\text{C}$, only a few dominant grains remain after 100 seconds. Coarsening statistics can be similarly obtained for these simulations, as is shown in Figure 4.12. We observe that the decay in the number of grains (left panel) is qualitatively similar over all choices of T . In the high temperature case $T = 427^\circ\text{C}$, the number of grains decreases from an initial value of 8600 (not shown in the plot) to under 500 within the first few seconds of the simulation. In the right panel, the mean grain radius is plotted, along with the best-fit power law for each choice of temperature. The power law varies from $b = 0.18$

when $T = 277^\circ\text{C}$ to $b = 0.34$ when $T = 427^\circ\text{C}$. The apparent periodicity observed in high temperatures is a result of the advance of a large grain into a smaller one. For two conjoined grains, the rate at which either grain's radius changes behaves according to the law

$$\frac{dR}{dt} = \frac{1}{2\pi R} v_n L,$$

where L is the interface length between the two grains and v_n is the normal velocity of that interface. As a result, as the smaller grain is consumed by the larger one, the radius of the smaller grain decreases much faster than the increase in the larger grain's radius. This implies that the expectation $\langle R \rangle$ is lowered as the advance occurs. When the smaller grain disappears, its contribution to the lower mean radius is removed, resulting in a jump in its value.

4.2.2 Computational statistics

The simulation depicted in 4.9, in which 10,000 seconds were simulated, required 42 billion Monte Carlo steps and 19 hours CPU time. For the other sintering simulations, the average number of Monte Carlo steps was 4.4 billion steps while the maximum was 30 billion steps. The corresponding average and maximum CPU time was 4.6 and 15.4 hours




Figure 4.11: Final configurations of sintering simulations after 100 seconds with $\rho = 1.25$ eV, for $T = 227, 327, 377$ and 427°C . Different colors correspond to different orientations.

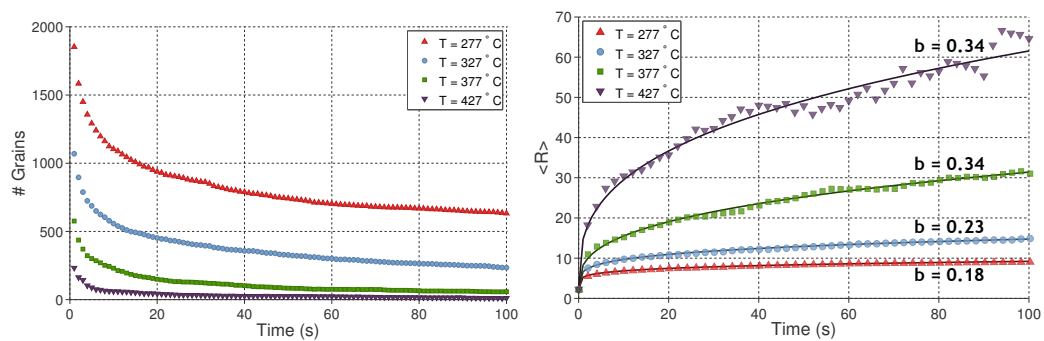


Figure 4.12: Grain coarsening statistics vs temperature. (Left) The number of grains vs. time (initial configuration when $t = 0$ is not shown). (Right) Mean disc equivalent radius $\langle R \rangle$ vs. time. Best fit power-law models $\langle R \rangle = at^b + \langle R_0 \rangle$ are shown in dashed lines, and the fitted value b is indicated for each plot. Each plot corresponds to a different value of T .

CHAPTER V

Conclusion

Kinetic Monte Carlo offers an attractive alternative to other atomistic modeling and simulation techniques such as molecular dynamics. It is capable of simulating larger scale processes and is therefore applicable to a broad range of systems of interest to materials scientists. As KMC is being applied to more complex systems, a careful study of efficient implementations must be made. While much work has been done in the certain algorithmic components of KMC, little attention has been focused on the constant time calculation of rates. This calculation however, is a significant portion of the floating point operations in KMC and is repeated hundreds of billions of times during simulation. As such, it is an often over-looked bottleneck. We addressed this by presenting a hash table based caching procedure designed to eliminate the redundant calculation of transition rates for neighborhood patterns frequently seen in a simulation. In order for this procedure to be effective, a “good” hash function must be specified. Instead of explicitly constructing one, we used a parameterized form of a hash function and designed a simulated annealing algorithm to search for optimal parameters. When implemented, the hash table based caching procedure leads a significant performance speed-up, allowing us to to perform large scale simulations in a reasonable amount of time.

With an efficient and flexible implementation in hand, we are able to simulate several example systems. As first example, we used KMC to model a GaAs system and are able to simulate the relevant processes in GaAs homoepitaxial film growth, droplet epitaxy and crystallization. As calibration and an initial validation of the model, we presented simulation results in GaAs homoepitaxy, exhibiting a good match with experimental data for surface termination experiments. As our main focus, we also modeled the formation and crystallization of Ga droplets resulting in GaAs nanostructures. Experimentally, such structures depend on growth conditions such as As flux temperature. Our simulation results are in good qualitative agreement with these experimental observations. We are able to reproduce the different nanostructures observed in experiments with the correct trends in flux and temperature – exhibiting a continuum of structures ranging between compact nanodots to broad nanorings. Simulations also suggest the existence of Ga/GaAs core/shell structures and yield insight in their formation. We observed two mechanisms behind the formation of core/shell structures. First is a simple nucleation of GaAs clusters at the liquid solid interface. Such nucleation results in a polycrystalline shell, which upon further high temperature annealing can be recrystallized into an epitaxial dot. This is confirmed by both simulations and experiments. The second method for core/shell structures is by a Mullins-Sekerka instability of the GaAs crystallization front. Shells formed via this instability are necessarily monocrystalline and epitaxial to the substrate. We developed a unifying model on top of the KMC model that encapsulates all the phenomena outlined above. The model identifies three processes active during crystallization. By specifying the critical conditions for the phenomena above, the model is able to delineate the existence of particular morphological structure in good agreement with simulations. With this model, we are able to precisely

state the dependence of structure on our KMC model parameters.

As further illustration of the KMC model, we simulated two more example systems. We presented simulation results of nanowire growth that capture a wide range of observed phenomena, including tapered bases, faceting of the liquid-solid interface and nanowire kinking. The amount and type of faceting was shown to be controlled by the parameter $\epsilon_{\text{ex}}^{\text{NL}}$, and a simple model for the diffusion length scale was developed to explain this influence. The simulations suggest that facet nucleation at the liquid-solid interface provides a mechanism for such growth modes, in agreement with past results from continuum models. We lastly presented simulation results of sintering, over several values of temperature and ρ , the additional energy barrier for changing orientations. The coarsening statistics of the grains show that the disc equivalent mean radius behaved as a power law with respect to time. That is $\langle R \rangle = at^b + \langle R_0 \rangle$ and that depending on the choice of parameters ρ and T . For most values of ρ^L , the mean radius scaled as $t^{\frac{1}{4}}$, however for values in a critical interval, we observed accelerated coarsening that scaled like $t^{\frac{1}{2}}$.

5.1 Future work

Given the ability of our KMC implementation to model a broad range of systems, we have been applying our work to even more systems, not described in this dissertation. Indeed, simulations have or are currently being run to model catalyst-free GaAs nanowires on a Si substrate, Ga droplet epitaxy on grooved Si substrates, and GaAs quantum dots capped overgrowth layers of GaAs. Future work in this regard are to produce physical scale simulations in a reasonable amount of time and perform an analogous study on parameters that was done on the basic GaAs system. One problem in the systems with Si is the low energy barrier for diffusion on a Si

substrate, leading to a dominance in surface diffusion in the events performed. The small barrier problem is ubiquitous in all atomistic simulations and our simulations are no different in this regard. The rate caching technique developed here somewhat mitigates the dominance of the diffusion process because of the repetition in neighborhoods implied by atom hopping. However, much work needs to be done regarding this in order to provide simulations that run quickly. Continuum/atomic coupling methods could prove advantageous here. Indeed work has already been done in modeling surface diffusion probabilistically in the context of KMC [80]. Future work would involve coupling such quasi-continuum models within our specific KMC framework through a proper utilization of the rate cache.

Even with the rate caching technique, there is much work that may be done to improve the performance and memory requirements of our KMC implementation. Specifically, the rate cache can be used as the main data structure for sampling and updating in a method similar to that proposed in [81]. There it is proposed that through that by organization rates by similar value, sampling and updating can be done in constant time. The cache, however sorts rates by neighborhood, but an analogous technique can be implemented for constant-time execution of a Monte Carlo step. In this way, the cache is no longer treated as an auxiliary data structure, but as the main one used to sample and update in KMC. As future work, this technique could be incorporated into the implementation, leading to even further performance gains.

Apart from the GaAs system, the two other examples presented here were presented as an illustration of the flexibility and efficiency of our model. More analysis should be done on the observed phenomena presented in the above parameter studies done for the nanowire and sintering examples. This is especially important in the

nanowire case, as experimentalists are starting to grow a variety of nanowire structures such as self-catalyzed GaAs nanowires. Understanding the precise relationship between parameters and wire growth modes is therefore important in complimenting experimental work, and providing a theoretical justification for the structures and growth mechanisms seen and proposed by experimentalists

BIBLIOGRAPHY

BIBLIOGRAPHY

- [1] <http://commons.wikimedia.org/wiki/File:Sphalerite-unit-cell-depth-fade-3D-balls.png>. Accessed: 2013-04-19.
- [2] Kristen A Fichthorn and Matthias Scheffler. Island nucleation in thin-film epitaxy: a first-principles investigation. *Physical review letters*, 84(23):5371–5374, 2000.
- [3] J. Kikkawa, Y. Ohno, and S. Takeda. Growth rate of silicon nanowires. *Applied Physics Letters*, 86(12):123109, 2005.
- [4] James R. Maiolo, Brendan M. Kayes, Michael A. Filler, Morgan C. Putnam, Michael D. Kelzenberg, Harry A. Atwater, and Nathan S. Lewis. High aspect ratio silicon wire array photoelectrochemical cells. *Journal of the American Chemical Society*, 129(41):12346–12347, 2007.
- [5] Y.Q. Chang, D.P. Yu, Y. Long, J. Xu, X.H. Luo, and R.C. Ye. Large-scale fabrication of single-crystalline mn_3o_4 nanowires via vapor phase growth. *Journal of Crystal Growth*, 279(12):88 – 92, 2005.
- [6] Yiying Wu and Peidong Yang. Direct observation of vaporliquidsolid nanowire growth. *Journal of the American Chemical Society*, 123(13):3165–3166, 2001.
- [7] A.B. Bortz, M.H. Kalos, and J.L. Lebowitz. A new algorithm for monte carlo simulation of ising spin systems. *Journal of Computational Physics*, 17(1):10 – 18, 1975.
- [8] A. Ishii and T. Kawamura. Monte Carlo simulation of homoepitaxial growth on two-component compound semiconductor surfaces. *Surface Science*, 436(13):38 – 50, 1999.
- [9] T. Kawamura and A. Ishii. Role of As during homoepitaxial growth on GaAs(001) studied using Monte Carlo simulation. *Surface Science*, 438(13):155 – 161, 1999.
- [10] T. Kawamura and A. Ishii. Monte Carlo simulation of recovery process after MBE growth on GaAs(100). *Surface Science*, 493(13):438 – 446, 2001.
- [11] GS Bales and DC Chrzan. Dynamics of irreversible island growth during submonolayer epitaxy. *Physical Review B*, 50(9):6057, 1994.
- [12] CC Battaile, DJ Srolovitz, and JE Butler. A kinetic monte carlo method for the atomic-scale simulation of chemical vapor deposition: Application to diamond. *Journal of applied physics*, 82(12):6293–6300, 1997.
- [13] Karsten Reuter and Matthias Scheffler. First-principles kinetic monte carlo simulations for heterogeneous catalysis: Application to the co oxidation at RuO₂(110). *Phys. Rev. B*, 73:045433, Jan 2006.
- [14] F. Much, M. Ahr, M. Biehl, and W. Kinzel. A kinetic monte carlo method for the simulation of heteroepitaxial growth. *Computer Physics Communications*, 147(12):226 – 229, 2002. Proceedings of the Europhysics Conference on Computational Physics Computational Modeling and Simulation of Complex Systems.

- [15] Tim P. Schulze and Peter Smereka. Simulation of Three-Dimensional Strained Heteroepitaxial Growth Using Kinetic Monte Carlo. *Communications in Computational Physics*, 10(5):1089–1112, 2011.
- [16] Arvind Baskaran, Jason Devita, and Peter Smereka. Kinetic monte carlo simulation of strained heteroepitaxial growth with intermixing. *Continuum Mechanics and Thermodynamics*, 22:1–26, 2010.
- [17] Arvind Baskaran and Peter Smereka. Mechanisms of stranski-krastanov growth. *Journal of Applied Physics*, 111(4):044321, 2012.
- [18] Tim P. Schulze and Peter Smereka. An energy localization principle and its application to fast kinetic monte carlo simulation of heteroepitaxial growth. *Journal of the Mechanics and Physics of Solids*, 57(3):521 – 538, 2009.
- [19] Graeme Henkelman and Hannes Jónsson. Long time scale kinetic monte carlo simulations without lattice approximation and predefined event table. *The Journal of Chemical Physics*, 115:9657, 2001.
- [20] Abdelkader Kara, Oleg Trushin, Handan Yildirim, and Talat S Rahman. Off-lattice self-learning kinetic monte carlo: application to 2d cluster diffusion on the fcc(111) surface. *Journal of Physics: Condensed Matter*, 21(8):084213, 2009.
- [21] Jeffrey Skolnick and Andrzej Kolinski. Dynamic monte carlo simulations of a new lattice model of globular protein folding, structure and dynamics. *Journal of molecular biology*, 221(2):499–531, 1991.
- [22] Eugene Shakhnovich, G Farztdinov, AM Gutin, and Martin Karplus. Protein folding bottlenecks: A lattice monte carlo simulation. *Physical review letters*, 67(12):1665–1668, 1991.
- [23] Kristofer Reyes, Peter Smereka, Denis Nothorn, Joanna Mirecki Millunchick, Claudio Somaschini, Sergio Bietti, Stefano Sanguinetti, and Cesare Frigeri. A unified model of droplet eptiaxy for compound semiconductor nanostructures: experiments and theory. *Phys. Rev. B*, 87, 2013.
- [24] Kristofer Reyes. Fast kinetic Monte Carlo simulations using hash table based caching with applications to nanowire growth and sintering. 2013. In preparation.
- [25] Graeme Henkelman, Blas P. Uberuaga, and Hannes Jonsson. A climbing image nudged elastic band method for finding saddle points and minimum energy paths. *The Journal of Chemical Physics*, 113(22):9901–9904, 2000.
- [26] Graeme Henkelman and Hannes Jonsson. A dimer method for finding saddle points on high dimensional potential surfaces using only first derivatives. *The Journal of Chemical Physics*, 111(15):7010–7022, 1999.
- [27] James L. Blue, Isabel Beichl, and Francis Sullivan. Faster Monte Carlo simulations. *Phys. Rev. E*, 51:R867–R868, Feb 1995.
- [28] Tim P. Schulze. Efficient kinetic monte carlo simulation. *Journal of Computational Physics*, 227(4):2455 – 2462, 2008.
- [29] Luc Devroye. *Non-uniform random variate generation*. Springer-Verlag, 1986.
- [30] Alexander Stukowski. Structure identification methods for atomistic simulations of crystalline materials. *Modelling and Simulation in Materials Science and Engineering*, 20(4):045021, 2012.
- [31] D.R. Mason, T.S. Hudson, and A.P. Sutton. Fast recall of state-history in kinetic monte carlo simulations utilizing the zobrist key. *Computer Physics Communications*, 165(1):37 – 48, 2005.

- [32] Oleg Trushin, Altaf Karim, Abdelkader Kara, and Talat S. Rahman. Self-learning kinetic monte carlo method: Application to cu(111). *Phys. Rev. B*, 72:115401, Sep 2005.
- [33] Matt Elsey, Selim Esedoglu, and Peter Smereka. Diffusion generated motion for grain growth in two and three dimensions. *Journal of Computational Physics*, 228(21):8015 – 8033, 2009.
- [34] T. Shitara, D. D. Vvedensky, M. R. Wilby, J. Zhang, J. H. Neave, and B. A. Joyce. Step-density variations and reflection high-energy electron-diffraction intensity oscillations during epitaxial growth on vicinal GaAs(001). *Phys. Rev. B*, 46:6815–6824, Sep 1992.
- [35] Marvin L. Cohen and T. K. Bergstresser. Band structures and pseudopotential form factors for fourteen semiconductors of the diamond and zinc-blende structures. *Phys. Rev.*, 141:789–796, 1966.
- [36] C. Orme, M. D. Johnson, J. L. Sudijono, K. T. Leung, and B. G. Orr. Large scale surface structure formed during GaAs(001) homoepitaxy. *Applied Physics Letters*, 64(7):860–862, 1994.
- [37] Akihiro Ohtake. Surface reconstructions on GaAs(001). *Surface Science Reports*, 63(7):295 – 327, 2008.
- [38] B.W. Liang and C.W. Tu. A study of Group-V element desorption from InAs, InP, GaAs and GaP by reflection high-energy electron diffraction. *Journal of Crystal Growth*, 128(14):538 – 542, 1993.
- [39] Hiroshi Yamaguchi and Yoshiji Horikoshi. As desorption from GaAs and AlAs surfaces studied by improved high-energy electron reflectivity measurements. *Journal of Applied Physics*, 71(4):1753–1759, 1992.
- [40] K. Watanabe, N. Koguchi, and Y. Gotoh. Fabrications of GaAs quantum dots by modified droplet epitaxy. *Japanese Journal of Applied Physics*, 39(2A):L79–L81, 2000.
- [41] Zh. M. Wang, B. L. Liang, K. A. Sablon, and G. J. Salamo. Nanoholes fabricated by self-assembled gallium nanodril on GaAs(100). *Applied Physics Letters*, 90(11):113120, 2007.
- [42] C. Somaschini, S. Bietti, N. Koguchi, and S. Sanguinetti. Fabrication of multiple concentric nanoring structures. *Nano Letters*, 9(10):3419–3424, 2009.
- [43] C. Somaschini, S. Bietti, N. Koguchi, and S. Sanguinetti. Shape control via surface reconstruction kinetics of droplet epitaxy nanostructures. *Applied Physics Letters*, 97(20):203109, 2010.
- [44] Takaaki Mano, Takashi Kuroda, Stefano Sanguinetti, Tetsuyuki Ochiai, Takahiro Tateno, Jongsu Kim, Takeshi Noda, Mitsuo Kawabe, Kazuaki Sakoda, Giyuu Kido, and Nobuyuki Koguchi. Self-assembly of concentric quantum double rings. *Nano Letters*, 5(3):425–428, 2005.
- [45] T. Kuroda, T. Mano, T. Ochiai, S. Sanguinetti, K. Sakoda, G. Kido, and N. Koguchi. Optical transitions in quantum ring complexes. *Phys. Rev. B*, 72:205301, 2005.
- [46] Ch. Heyn, A. Stemann, A. Schramm, H. Welsch, W. Hansen, and Á. Nemesics. Regimes of gaas quantum dot self-assembly by droplet epitaxy. *Phys. Rev. B*, 76:075317, 2007.
- [47] Ch. Heyn, A. Stemann, M. Klingbeil, Ch. Strelow, T. Kppen, S. Mendach, and W. Hansen. Mechanism and applications of local droplet etching. *Journal of Crystal Growth*, 323(1):263 – 266, 2011.
- [48] V. Gorokhov, T. Dedegkaev, Yu. Ilyin, V. Moshnikov, A. Petrov, Yu. Sosov, and D. Yaskov. The investigation of P- and As diffusion in liquid gallium. *Crystal Res. and Technol.*, 19:1465–1468, 1984.

- [49] C. Somaschini, S. Bietti, S. Sanguinetti, N. Koguchi, and A. Fedorov. Self-assembled GaAs/AlGaAs coupled quantum ring-disk structures by droplet epitaxy. *Nanotechnology*, 21(12):125601, 2010.
- [50] C. Somaschini, S. Bietti, A. Fedorov, N. Koguchi, and S. Sanguinetti. Outer zone morphology in GaAs ring/disk nanostructures by droplet epitaxy. *Journal of Crystal Growth*, 323(1):279 – 281, 2011.
- [51] T Mano, M Abbarchi, T Kuroda, C A Mastrandrea, A Vinattieri, S Sanguinetti, K Sakoda, and M Gurioli. Ultra-narrow emission from single GaAs self-assembled quantum dots grown by droplet epitaxy. *Nanotechnology*, 20(39):395601, 2009.
- [52] L. Cavigli, S. Bietti, N. Accanto, S. Minari, M. Abbarchi, G. Isella, C. Frigeri, A. Vinattieri, M. Gurioli, and S. Sanguinetti. High temperature single photon emitter monolithically integrated on silicon. *Applied Physics Letters*, 100(23):231112, 2012.
- [53] G. S. Bales and A. Zangwill. Morphological instability of a terrace edge during step-flow growth. *Phys. Rev. B*, 41:5500–5508, Mar 1990.
- [54] B. Caroli, C. Caroli, and B. Roulet. Instabilities of planar solidification fronts. In C. Godr che, editor, *Solids far from equilibrium*, pages 155–296. Cambridge University Press, 1992.
- [55] Toshiro Isu, Masayuki Hata, and Akiyoshi Watanabe. Real-time μ -RHEED observations of GaAs surfaces during growth with alternating source supply. *Journal of Crystal Growth*, 111(14):210 – 215, 1991.
- [56] G. Russo, L. M. Sander, and P. Smereka. Quasicontinuum Monte Carlo: A method for surface growth simulations. *Phys. Rev. B*, 69:121406, 2004.
- [57] T. A. Witten and L. M. Sander. Diffusion-limited aggregation, a kinetic critical phenomenon. *Phys. Rev. Lett.*, 47:1400–1403, 1981.
- [58] J. Westwater, D. P. Gosain, S. Tomiya, S. Usui, and H. Ruda. Growth of silicon nanowires via gold/silane vapor–liquid–solid reaction. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 15(3):554–557, 1997.
- [59] R. S. Wagner and W. C. Ellis. Vapor-liquid-solid mechanism of single crystal growth. *Applied Physics Letters*, 4(5):89–90, 1964.
- [60] Suneel Kodambaka, James B. Hannon, Rudolf M. Tromp, and Frances M. Ross. Control of si nanowire growth by oxygen. *Nano Letters*, 6(6):1292–1296, 2006.
- [61] Shadi A. Dayeh, Edward T. Yu, and Deli Wang. Iiiv nanowire growth mechanism: v/iii ratio and temperature effects. *Nano Letters*, 7(8):2486–2490, 2007. PMID: 17608541.
- [62] Hannah J. Joyce, Qiang Gao, H. Hoe Tan, Chennupati Jagadish, Yong Kim, Melodie A. Fickenscher, Saranga Perera, Thang Ba Hoang, Leigh M. Smith, Howard E. Jackson, Jan M. Yarrison-Rice, Xin Zhang, and Jin Zou. High purity gaas nanowires free of planar defects: Growth and characterization. *Advanced Functional Materials*, 18(23):3794–3800, 2008.
- [63] B. Tian, P. Xie, T.J. Kempa, D. C. Bell, and C. M. Lieber. Single-crystalline kinked semiconductor nanowire superstructures. *Nature Nanotechnology*, 4, 2009.
- [64] K. W. Schwarz and J. Tersoff. Elementary processes in nanowire growth. *Nano Letters*, 11(2):316–320, 2011.
- [65] K. W. Schwarz and J. Tersoff. From droplets to nanowires: Dynamics of vapor-liquid-solid growth. *Phys. Rev. Lett.*, 102:206101, May 2009.

- [66] AG Nastovjak, IG Neizvestny, IL Shwartz, and Z.S. Yanovitskaya. Monte carlo simulation of growth of nanowiskers. *Semiconductors*, 44(1):127–132, 2010.
- [67] Sergiy Krylyuk, Albert V. Davydov, and Igor Levin. Tapering control of si nanowires grown from sicl₄ at reduced pressure. *ACS Nano*, 5(1):656–664, 2011.
- [68] J. W. Dailey, J. Taraci, T. Clement, David J. Smith, Jeff Drucker, and S. T. Picraux. Vapor-liquid-solid growth of germanium nanostructures on silicon. *Journal of Applied Physics*, 96(12):7556–7567, 2004.
- [69] F. Y. Wu. The potts model. *Rev. Mod. Phys.*, 54:235–268, 1982.
- [70] Vishal N. Koparde and Peter T. Cummings. Molecular dynamics simulation of titanium dioxide nanoparticle sintering. *The Journal of Physical Chemistry B*, 109(51):24280–24287, 2005.
- [71] Amitava Moitra, Sungho Kim, Seong-Gon Kim, Seong Jin Park, Randall M. German, and Mark F. Horstemeyer. Investigation on sintering mechanism of nanoscale tungsten powder based on atomistic simulation. *Acta Materialia*, 58(11):3939 – 3951, 2010.
- [72] R.A Sutton and G.B Schaffer. An atomistic simulation of solid state sintering using monte carlo methods. *Materials Science and Engineering: A*, 335(12):253 – 259, 2002.
- [73] Michael Braginsky, Veena Tikare, and Eugene Olevsky. Numerical simulation of solid state sintering. *International Journal of Solids and Structures*, 42(2):621 – 636, 2005.
- [74] M.P Anderson, D.J Srolovitz, G.S Grest, and P.S Sahni. Computer simulation of grain growth kinetics. *Acta Metallurgica*, 32(5):783 – 791, 1984.
- [75] Elizabeth A. Holm, James A. Glazier, David J. Srolovitz, and Gary S. Grest. Effects of lattice anisotropy and temperature on domain growth in the two-dimensional potts model. *Phys. Rev. A*, 43:2662–2668, 1991.
- [76] Matt Elsey, Selim Esedoglu, and Peter Smereka. Simulations of anisotropic grain growth: Efficient algorithms and misorientation distributions. *Acta Materialia*, 61(6):2033 – 2043, 2013.
- [77] Matt Elsey, Selim Esedolu, and Peter Smereka. Large-scale simulation of normal grain growth via diffusion-generated motion. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, 467(2126):381–401, 2011.
- [78] G.F. Bolling and W.C. Winegard. Grain growth in zone-refined lead. *Acta Metallurgica*, 6(4):283 – 287, 1958.
- [79] V. Petrovic and M.M. Ristic. Isochronal and isothermal grain growth during sintering of cadmium oxide. *Metallography*, 13(4):319 – 327, 1980.
- [80] Jason P. DeVita, Leonard M. Sander, and Peter Smereka. Quasicontinuum monte carlo simulation of multilayer surface growth. In Axel Voigt, editor, *Multiscale Modeling in Epitaxial Growth*, volume 149 of *ISNM International Series of Numerical Mathematics*. Birkhuser Basel, 2005.
- [81] T. P. Schulze. Kinetic monte carlo simulations with minimal searching. *Phys. Rev. E*, 65:036704, Feb 2002.