Growth Simulations of Single Crystal Perovskite Alloys

A thesis submitted in partial fulfillment of the requirement
for the degree of Bachelor of Science with Honors in
Physics from the College of William and Mary in Virginia,

by Charles Tahan

Accepted for ______________________________
(Honors, High Honors, or Highest Honors)

____
Advisor: *Shiwei Zhang*[1,2]

____
Henry Krakauer[1]

____
Virginia Torczon[3]

____
Gina Hoatson[1]

Williamsburg, Virginia
May 2000

[1] Physics Department, The College of William and Mary

[2] Applied Science Department, The College of William and Mary

[3] Computer Science Department, The College of William and Mary

Abstract

Electromechanical actuators directly transform input electrical energy into mechanical energy. A recent research breakthrough with single crystals of relaxor ferroelectrics (perovskite alloys), which exhibit these piezoelectric properties, has produced samples of these materials that exhibit electromechanical strains reaching two percent—an order of magnitude higher than conventional polycrystalline piezoelectrics. Because of wide-ranging applications, questions as to the fundamental structure and growth mechanisms of these novel materials (with the goal of growing bigger crystals, faster) have drawn much attention.

Computer simulations, especially via Monte Carlo (MC) means, have proved very valuable in deducing the physical processes involved in crystal growth. Up till now, though, MC has been mostly limited to highly simplified (Ising-like) models, which are not adequate for the materials in question. For example, long-range electrostatic interactions, which are believed to have a strong effect on the structure of these types of crystals, are neglected in all MC growth calculations of which we are aware. Based on previous work, we add these long-range Coulomb interactions and develop an algorithm, based on the so-called kinetic Monte Carlo (KMC) algorithm, to model the growth processes of relaxor single crystals. Our new simulation has proven successful in exhibiting the growth and ordering behavior expected for the crystals "grown." An analysis of the new results, both qualitative and quantitative, concludes.

Acknowledgments

To my learned master, Shiwei Zhang, and his partner in arms, Henry Krakauer, I owe great thanks. I chose my honors thesis project based on the people I'd work with, and I think I chose very wisely. Thanks to Dr. Kossler for the office space and desk: how else could I have spent all waking hours working in Small? Perhaps out of necessity, I've grown rather fond of the place…I don't think things have changed much since the time of Jefferson; Small is still the best thing on campus. And finally, thanks to all of my committee members; I hope you enjoy the show…

Contents

List of Figures

List of Tables

Part I

Background and Methodology

Chapter 1

Introduction and Motivation

Santa Claus has a lot of names. It may be that the more important something is—the more widely it is used or known among different communities—the greater the number of descriptive words appear to represent it. Actuators offer a striking parallel. You may have heard the word actuator before, or perhaps transducer, *piezoelectric*[1], motivator, or even some other name. Although these varying terms may provide unique information as to the particular material or application involved, they all refer to a greater class of materials that have some amazing properties.

Electromechanical actuators directly transform input electrical energy into mechanical energy and vice-versa. To help understand the odd properties we are discussing here, imagine a match-box sized plate, with electrodes attached, made of one of these so-called actuating materials. Squeeze it between your fingers. Doing so

[1]Italicized words are further defined or explained in Appendix B. For example, piezoelectricity, via its Greek roots, literally means "pressure electricity."

causes an electrical impulse through the electrodes. Conversely, place the plate in an electric field or run a voltage through the electrodes; the material will actually expand! This electric field induced strain, known as the piezoelectric effect and quantified by the *piezoelectric coefficient*, is the most important parameter for actuator performance [2]. It is not a great leap to assume that this type of material, which usually comes in a *crystallic* or *ceramic* form, might have some use somewhere.

It took almost 50 years to find one. The discovery of the piezoelectric effect by Pierre and Jaques Curie in 1880 remained little more than a curiosity until the end of the first World War [1]. The product of steel plates, quartz slabs, and quite a bit of tinkering by a French Professor named Langevin, the first primitive, sonar-based, submarine detector was born [1]. And so it began. The piezoelectric properties of quartz are well used today; its high frequency oscillations probably drive the watch on your wrist. These are only two examples of many applications that these almost magical materials have found a place in, from the military to cell phones to speakers. Not surprisingly, much research has gone into the discovery, categorization, and analysis of these materials. On the theoretical side, analytical complexity and a lack of adequate computational ability have hampered progress. Today, the vast majority of piezoelectrics are still not well understood.

Very recently S.E. Park and T.E. Shrout at the Pennsylvania State University discovered a new class of these materials in the form of complex insulating single crystal *perovskite alloys* with E-field to strain ratios an order of magnitude higher than conventional polycrystalline piezoelectrics [2]. The catch is that they are difficult and expensive to grow. The applicational allure of these novel materials, coupled with the realization of our theoretical shortcomings, has prompted a renewed theoretical interest in the field.

Computational techniques offer an avenue of investigation not possible in the pre-"microprocessor revolution" era. Computer simulations provide a highly controllable, though simplistic, model-based world in which composition and environment parameters can be changed easily and cheaply. Predictions from simulations can be fed into experiments, whose output can be used to improve theory, and in turn the simulation model itself. But in order to choose a preferential computational technique, we must narrow the investigation process and decide what characteristic of the system we want to study.

The growth processes of these crystals are what interest us here. From a practical perspective towards application, the real goal is to grow bigger crystals, faster. The question of why these materials act as they do, what governs their piezoelectric properties, is an important though secondary question in our context. The structural properties, composition, and constituent arrangement of these crystals, however, are of vital interest; our growth simulation must result in a physically relevant and possible, if highly idealized, crystal.

The objective of crystal growth simulation is to create a model that describes the dynamics of crystal growth as *stochastic processes* such as *adsorption*, *evaporation*, and *surface migration* [4]. The kinetic Monte Carlo (KMC) method [34], one of several *simulation techniques*, is used to simulate the relaxation processes of systems away from equilibrium (e.g. growth processes) and has been applied successfully to crystal growth and surface/interface phenomena [5]. Although promising, traditional models of KMC growth simulations are too simplistic to adequately simulate the physics of these new crystals. Here we introduce a model which includes long-range ion-ion interactions and incorporate it into an enhanced KMC method. We then develop, implement, and test the enhanced method in a crystal growth simulation.

In 1998 Vanderbilt and Bellaiche [3] proposed that the *crystal ordering* of the crystals discovered at Penn State, indeed of a whole class of perovskite alloys, is governed predominately by the electrostatic interactions between ions in the *crystal lattice*. The very simplistic model they introduced proved surprisingly successful in predicting the ordering of various perovskite alloys across a spectrum of compositional constituents in *equilibrium Monte Carlo simulations*. Their equilibrium method provides no growth information however; our step is to fuse their simple ionic model with the KMC method.

Our enhanced KMC method introduces long-range *Coulomb forces* between various crystal *species* inhabiting the growing crystal lattice. A novel algorithm is presented which, in addition to the ability to handle multiple species, is much more efficient than traditional KMC algorithms—overcoming some of the difficulties inherent to the long-range interaction problem.

In the succeeding pages we describe and then further develop the kinetic Monte Carlo method, based on the aforementioned ionic model and generalized to incorporate either long-range or short-range (nearest-neighbor) forces depending on the potential chosen. We first outline the traditional approach, its limitations and deficiencies. We then introduce a new KMC algorithm, implement it, and test it against the traditional short-range method for qualitative and quantitative comparison. Long-range interactions are then turned on—signalling the incorporation of a novel *Ewald summation* for finite height technique. Qualitative and quantitative results of the completed program and a guide for future work conclude this thesis.

Chapter 2

Background

Crystallography, as a distinct field of study, represents over three centuries of investigation and discovery. In an effort to keep the reader motivated towards the exciting results ahead, we will limit our discussion here, as much as possible, to the relevent, necessary information. For the expert, this chapter can be skipped.

2.1 Crystals, Ceramics, and Nomenclature

The crystalline state of a substance is characterized by a time-invariant, regular, three-dimensional arrangement of atoms in space [6]. A necessarily broad definition, crystals are usually identified with their most prominent physical properties, such as translucency to light and their tendency to break, smoothly, along definite angles. Indeed, most solids can be considered crystalline, the biggest exceptions being *overcooled liquids* such as glass and *amorphous solids*. Also relavent are polycrystalline materials or ceramics which are aggregates of randomly oriented small crystals, usually of different sizes and irregular shapes. For many applications, including ultrasonics or high frequency oscillators, single crystals are much preferred over ceramics: their properties, not surprisingly, have much better frequency stability and lower acoustic losses [12].

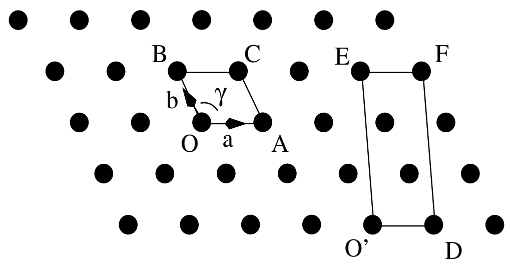

Figure 2.1: Two unit cells and the lattice they create. $OACB$ is a primitive unit cell; $O'DFE$ is not.

The principal theoretical constructs for describing a single crystal are the *lattice* and the *unit cell*. To simply illustrate their meaning, refer to the two-dimensional example in Figure 2.1. It should be easy to convince yourself that any point in the lattice is the same as any other point (assuming the lattice is infinite). The positions of all the lattice points of the two-dimensional lattice are reached by drawing all possible vectors of the form

$$\mathbf{r} = i\mathbf{a} + j\mathbf{b} \tag{2.1}$$

from the origin, where i and j take on all possible integer values, positive, negative, and zero [7]. A unit cell is defined by these lattice vectors; stacked together they generate the entire crystal lattice. In Figure 2.1, two unit cells have been drawn;

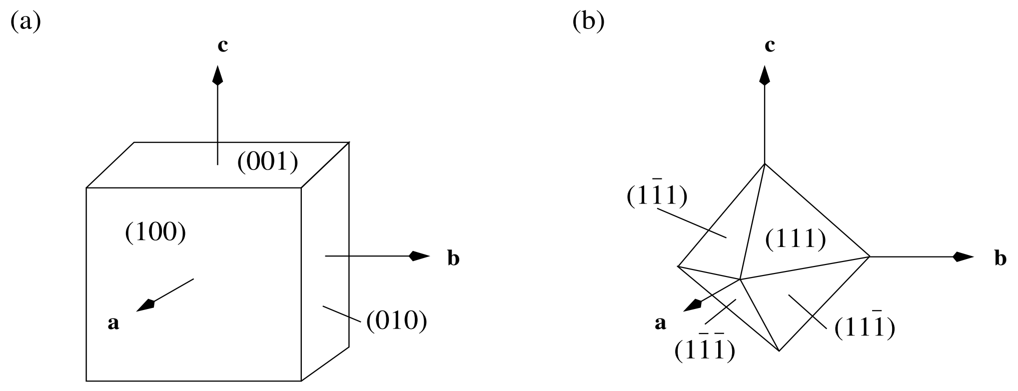

Figure 2.2: Two examples of different crystals in the cubic system: (a), a cube; (b), an octahedron.

many more can be imagined. The smallest possible unit cell is prefixed primitive. A crystal lattice is simply a lattice where each point is labeled with an atom or ion. Real single crystals are described in this way; although finite, the sheer number of atoms involved makes assuming an infinite lattice acceptable. Single crystals generally do not have any macroscopic region misoriented (crystallography) with respect to any other region by more than a few degrees. Despite the extreme regularity by which these materials are classified, they certainly are not perfect. In fact, some of their most interesting and useful properties come from their imperfections. The study and categorization of defects as well as many other properties will not be dealt with here. These are best left to a good Solid State reference [7, 8].

Rather than a broad (and lengthy) introduction to crystal lingo, we will precede by introducing the crystal types of interest in this thesis, defining concepts progressively. Of the seven *crystal systems*, sorted by shape and axial restrictions, the cubic system is perhaps the easiest to visualize. Fortunately, it is the one we are concerned with.

The cubic system is defined as having four threefold axes at equal angles to **a**, **b**, and **c** (along the body diagonals) where **a**, **b**, and **c** are perpendicular to each other. This means that from any one of these axes, there are three symmetry operations (rotations) that will return the system to the original configuration. If we take the simple example of Figure 2.2(a), we notice that the cube also has three fourfold axes along **a**, **b**, and **c**, and six twofold axes halfway between **a** and **b**, **b** and **c**, and **c** and **a** (parallel to the face diagonals). The trio of numbers enclosed in parenthesis labeling each face of the crystal are known as *Miller indices*. Defined by the inverse of the axes coordinate, (001) refers to the crystal face in the **a**-**b** plane at $c = 1$, (100) the **c**-**b** plane at $a = 1$, and (010) the **a**-**c** plane at $b = 1$. (Note: $1/\infty = 0$). Barred numbers are simply the negative. Miller indices can be quite powerful: replace the parenthesis with brackets and you're referring to the direction; for example, [001] would be the **c** direction.

2.1.1 Piezoelectricity and Piezoelectric Crystals

Piezoelectricity arises only in insulating materials which lack a center of symmetry. In crystalline materials, by far the largest group of materials showing piezoelectric behavior, the polarization is changed by applying a stress or by changing the electric field. The reverse is also true: an electric field will cause a strain in the material. Figure 2.3, courtesy of [7], shows how a uniaxial stress can cause a spontaneous dipole moment, thus an induced electric field. A centrosymmetric crystal cannot be piezoelectric because no combination of unified stresses will produce a separation of the center of gravity of the positive and negative charges in the crystal lattice [1]. The piezoelectrics that are a concern of this thesis are known as relaxor *ferroelectrics*. Ferroelectrics are the electrical analogues of ferromagnets. These particular materials are a subset of a larger group of materials called perovskite alloys. All ferroelectric materials are piezoelectric but the converse is not true (e.g. quartz).

2.2 Perovskite Alloys

Perovskite alloys are common compounds usually containing a mixture of alkaline and transition metals. The "new" single crystal relaxor ferroelectrics we are concerned with are described most generally by the chemical formula ABO_3. Here A is usually a fixed Group II metal, and B has a fractional compositional freedom containing metals from Groups II-VII. Perhaps a better formula is $A(BB')O_3$; some analysis of

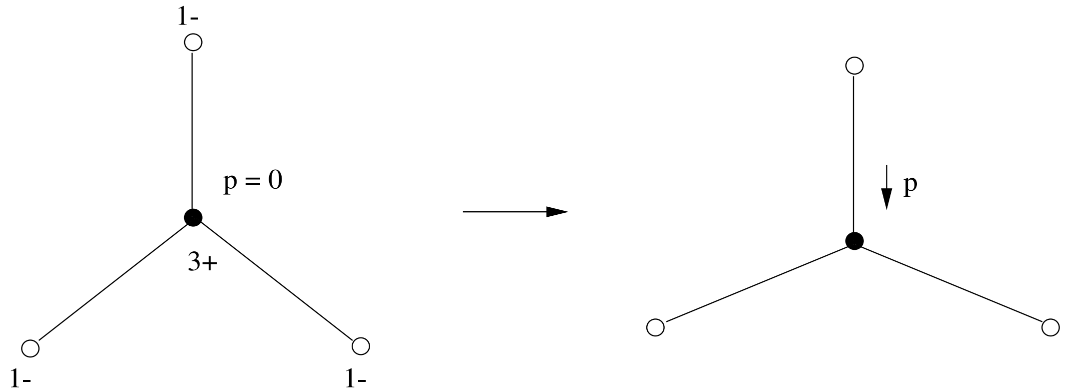

Figure 2.3: Illustration of a way in which a molecule can obtain a dipole moment as a result of uniaxial strain where the plus, minus signs reflect charge.

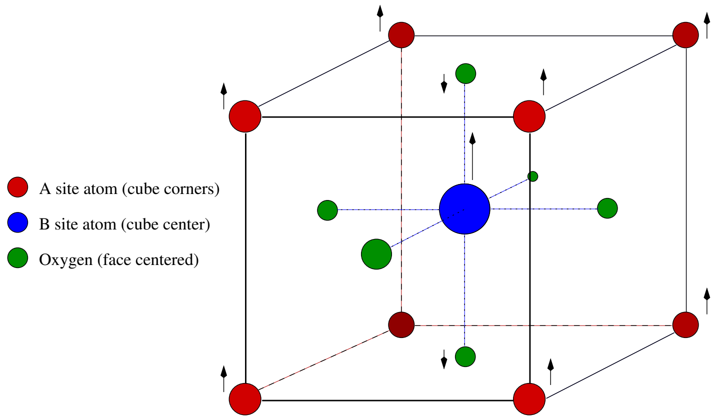

Figure 2.4: Unit cell of cubic perovskite alloy. The arrows indicate the relative sizes and directions of the displacement of the ions along the z direction.[7]

the actual materials will elucidate this. As foreshadowed above, these crystals are cubic in nature. Figure 2.4 shows a unit cell. A B atom occupies the center, A atoms the corners, and an Oxygen atom is centered on all six of the cube's faces.

The two remarkable perovskite crystals introduced by Park and Shrout [2], the impetus of our research, seem a perfectly good place to start an example. $Pb(Zn_{1/3}Nb_{2/3})O_3$-$PbTiO_3$ (PZN-PT) and $Pb(Mg_{1/3}Nb_{2/3})O_3$-$PbTiO_3$ (PMN-PT) are both alloys, a combination of two ABO_3 formulas. We can thus generalize our chemical equation even further; $(1-x)Pb(BB')O_3 - xPb(B'')O_3$ expresses the new relation, where x is the percentage composition of $Pb(B'')O_3$ in the crystal. (BB′) and B′′ must each have charge +4. This is where the compositional freedom of B comes into play. In these two examples alone many stoichiometries can be built from our limited array of choices: A=$\{Pb^{2+}\}$, B=$\{Zn^{2+}, Mg^{2+}\}$, B′=$\{Nb^{5+}\}$, and $B'' = \{Ti^{4+}\}$. Simple algebra shows how the various fractional amounts of components satisfy the chemical charge neutrality rules. For example, the fractions of B (+2) and B′ (+5) must be 1/3 and 2/3, respectively. The key is that the A and O_3 atoms are fixed in the cubic lattice, while the B sites are occupied by varying elements (up to three in these two cases) whose distribution defines the uniqueness of the crystal. It turns out that the ordering of these B-site elements has a strong effect on the relaxor properties of the crystal [9] (which should not astonish you now that the nature of these materials is becoming clearer).

2.2.1 An Electrostatic Model of Atomic Ordering

PMN-PT and PZN-PT described above are known as divalent-A compounds because they have a constant A atom of +2 charge on all A sites in the crystal lattice. With the A and O sites unvarying, the compositional freedom of the B atoms allows for diverse lattice-configuration arrangements. Experiments have confirmed the presence of long range order (LRO) in the placement of the B atoms. Compositional fluctuations among the B sites are believed to play a central role in the "relaxor" behavior in Pb-based perovskite alloys [3].

LRO is defined just how it sounds; Figure 2.5 shows several examples of crystal ordering dependent on the B site atoms. Basically, the result is periodic layers parallel to the (111) crystal plane. LRO types are denoted by colon separated numbers where each number represents the repetition of identical layers before a switch. Thus 1:1 order looks like a 3-d checker-board, an alternation of $\pm n$ charged planes along the

[111] direction. That is, unless one of the layers is disordered, as in Figure 2.5(b). Keep in mind that all talk of disorder or order is relative to the lattice structure—not spatially but compositionally; spatial disorder would indicate a change in phase to a liquid or gas. Our essential simplification is that all atoms remain absolutely fixed relative to the idealized crystal lattice.

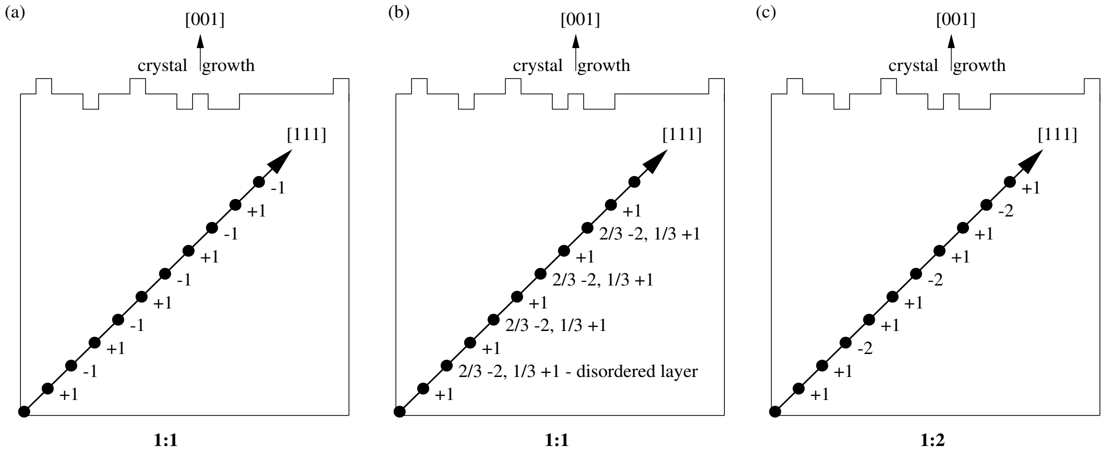

Figure 2.5: Examples of B-site ordering within the crystal lattice: (a), $[111]_{1:1}$ order, e.g. Rock Salt; (b), $[111]_{1:1}$ order with disordered odd layers; (c) $[111]_{1:2}$ order. Notice that the total system must be charge neutral (a consequence of our redefinition of the B-site charges, refer to Equation 2.3).

Bellaiche and Vanderbilt (BV) proposed a simple, though successful, model whereby the Coulomb interaction between ions is the dominant factor in determining the complicated ordering behavior of these complex perovskite alloys [3]. BV showed that up to a constant, the configurationally averaged total energy from Coulomb interactions in the crystal system depends only on the B-site charges, effectively reducing the system to the interaction between charge-normalized B atoms. The interaction energy in ABO_3 perovskite compounds is then expressed as

$$E = \frac{e^2}{2} \sum_{(l\tau) \neq (l'\tau')} \frac{Q_{l\tau} Q_{l'\tau'}}{\epsilon |\mathbf{R}_{l\tau} - \mathbf{R}_{l'\tau'}|}, \tag{2.2}$$

where $\mathbf{R}_{l\tau}$ is the position in the ideal cubic structure of the atom on site τ of cell l ($\tau =\{$A, B, O_1, O_2, $O_3\}$), and ϵ is a dielectric constant providing some screening effects. For alloys with Pb or Ba occupied A-sites, $Q_{l,A} = q_A = +2$ and $Q_{l,O} = q_O = -2$. $Q_{l,B}$ can be decomposed as

$$Q_{l,B} = q_B + \Delta q_l, \tag{2.3}$$

where the average valence q_B is equal to $+4$ while Δq_l depends on site l and the atom occupying it. The average charge of Δq_l over all B-sites, then, must vanish. We can further reduce the electrostatic energy to that depending on only the B sublattice (up to a constant),

$$E_B = \frac{e^2}{2\epsilon a} \sum_{l \neq l'} \frac{\Delta q_l \Delta q_l'}{|\mathbf{l} - \mathbf{l}'|}, \tag{2.4}$$

where $\mathbf{R}_l$ has been replaced by $\mathbf{l}a$, a being the lattice spacing constant and $\mathbf{l}$ the vector to a B lattice site.

With this B sublattice reduction in mind, BV introduced a B-site-only notation for divalent-A compounds with B-site alloying of which we will take full advantage. The notation uses periodic table element categories subscripted by percentage composition. Thus, PMN would be classified $\mathrm{II}_{1/3}\mathrm{V}_{2/3}$, while PMN-PT, and PZN-PT would be $\mathrm{II}_{\frac{1}{3}(1-x)}\mathrm{VI}_x\mathrm{V}_{\frac{2}{3}(1-x)}$. The next section will clarify.

Table 2.1: Common Group II-VII elements.

Group	Typical Elements	Charge
II	Be, Mg, Ca, Sr, Ba, Ra	2+
III	Sc, Y	3+
IV	Ti, Zr, Hf	4+
V	V, Nb, Ta	5+
VI	Cr, Mo, W	6+
VII	Mn, Re	7+
others	Zn^{2+}, Co^{2+}	

2.2.2 More Perovskites...

$Pb(Zr_{1-x},Ti_x)_3$ ceramics have been the mainstay for high performance actuator applications [2]. Ceramics, however, have large hysteresis in their strain vs. E-field behavior, and single crystals are much preferred in many types of applications. Despite the existence (and use) of readily available piezoelectrics such as quartz and lithium niobate, their inferior piezoelectric properties have lead to a search for better performing single crystals. Unsuccessful attempts at growing single crystal PZTs were the first step to Park and Shrout's investigation of PMN-PT and PZN-PT. Relaxor-PTs are readily growable compared to PZTs and their piezoelectric properties, as mentioned before, are quite impressive.

Any model we come up with, including BV's, should be able to predict the crystallographic ordering of Relaxor-PTs and similar materials. Obviously, for us to know whether our model adequately describes the physical system in question, we must be able to compare our simulation results with experimental data. Tables 2.2 and 2.3 outline some of the behavior observed in the equilibrium Monte Carlo simulations of

Table 2.2: A useful reference: different classes of alloys and their ordering behavior [3].

Alloy	**Type**	**Predicted Ordering**	**Examples & Evidence**
$IV_xIV'_{1-x}$	homovalent binary alloy having tetravalent B atoms	no order (all $\Delta q_l = 0$)	$P(ZrTi)O_3$ (PZT) [15]
$II_{1/3}V_{2/3}$	heterovalent binary	$[111]_{1:2}$ order in ground state	$Ba(Mg_{1/3}Nb_{2/3})O_3$ (BMN) [18], $Ba(Zn_{1/3}Ta_{2/3})O_3$ (BZT) [19, 20], $Ba(Mg_{1/3}Ta_{2/3})O_3$ (BMT) [21], $Ba(Sr_{1/3}Ta_{2/3})O_3$ (BST) [21], $Ba(Zn_{1/3}Nb_{2/3})O_3$ (BZN) [22], $Ba(Ni_{1/3}Nb_{2/3})O_3$ (BNN) [23], $Ca(Ca_{1/3}Ta_{2/3})O_3$ (CCT) [3, 24]
$II_{1/2}VI_{1/2}$, $III_{1/2}V_{1/2}$	heterovalent binary	rock salt type $[111]_{1:1}$ order, $\pm n$ where $n = 1, 2$	$Pb(Sc_{1/2}Ta_{1/2})O_3$ [9, 16], $Pb(In_{1/2}Nb_{1/2})O_3$ [15], $Pb(Yb_{1/2}Ta_{1/2})O_3$ [17], $Pb(Mg_{1/2}W_{1/2})O_3$ [15], $Pb(Co_{1/2}W_{1/2})O_3$ [15]
$II_{(1-x)/3}IV_xV_{2(1-x)/3}$	heterovalent ternary (heterovalent binary with tetravalent B-bulk)	$[111]_{1:2}$ or $[111]_{1:1}$, see Table 2.3	See Table 2.3

Table 2.3: The ordering behavior of heterovalent ternaries vs. percentage tetravalent B-bulk (x).

Alloy	**Ordering vs. x**	**Examples & Evidence**
$\mathrm{III}_{(1-x)/2}\mathrm{IV}_x\mathrm{V}_{(1-x)/2}$, $\mathrm{II}_{(1-x)/2}\mathrm{IV}_x\mathrm{VI}_{(1-x)/2}$	$[111]_{1:1}$ order becomes disordered with increasing x.	$(1-x)\mathrm{Pb}(\mathrm{Sc}_{1/2}\mathrm{Ta}_{1/2})\mathrm{O}_3 + x\mathrm{PbTiO}_3$ [24, 25]
$\mathrm{II}_{(1-x)/3}\mathrm{IV}_x\mathrm{V}_{2(1-x)/3}$	$[111]_{1:2}$ order for $x <$ 5%, $[111]_{1:1}$ order for $x >$ 5%, and disorder for large x.	$(1-x)\mathrm{Ba}(\mathrm{Mg}_{1/3}\mathrm{Nb}_{2/3})\mathrm{O}_3 - x\mathrm{BaZrO}_3$ [18], $(1-x)\mathrm{Ba}(\mathrm{Mg}_{1/3}\mathrm{Ta}_{2/3})\mathrm{O}_3 - x\mathrm{BaZrO}_3$ [26], $(1-x)\mathrm{Ba}(\mathrm{Zn}_{1/3}\mathrm{Ta}_{2/3})\mathrm{O}_3 - x\mathrm{BaZrO}_3$ [27], PMN-PT [*], PZN-PT [*]

* Experimental evidence does not confirm the 1:2 ordering phase.

BV and corresponding references to experimental confirmation [3]. They will provide a very good reference when the results are explored in the analysis of the simulation output; we will continue our description there (Chapter 5). It should be noted that the tables list results already predicted and confirmed or proved wrong by the equilibrium MC simulations run by BV. Our KMC growth simulation, which is based on the same long-range ionic model, should produce similar final order configurations, affected only by the growth processes of the crystal. It adds no new physics to the equilibrium ordering mechanism (more on this in Chapter 3). In addition, the example materials given in the tables are not necessarily single crystal perovskites (BMN-BZ, for example, is a microwave ceramic), but they are all ABO_3/B-site-reducible compounds veiled by the BV model.

2.3 Crystal Growth

The theory of crystal growth brings crystallography under the umbrella of thermodynamics and physical chemistry. An attempt at categorizing and describing very complex processes, growth theory over the decades has resorted to creating a set of "notions" explaining growth dynamics in certain situations. The fairly recent advent of computational techniques and reasonable computing power are helping to turn these empirically-gleaned rules of thumb into more quantitative simulations.

Crystals can be grown in many ways: from solution, vapor, or a "melt" (which is what we are concerned with (Section 2.3.1)). The common theoretical construct

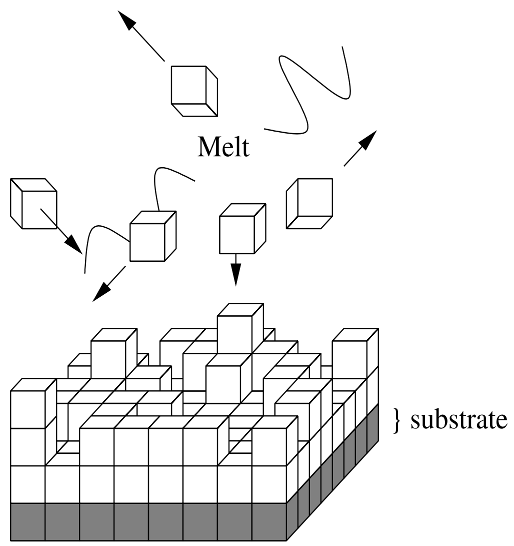

Figure 2.6: Growth from the melt.

involves planar or quasi-planar surface growth along a specified direction. Depending on the chemical potential difference, $\Delta\mu$, between the solid-liquid (or gas) interface, the crystal surface grows with some rate, R. Two basic and related questions are: what is the growth mode (Section 2.3.2) and what is the growth kinetics, i.e. how does the growth rate R depend on the driving force $\Delta\mu$ [10].

In our simulations of perovskite crystal growth, we are concerned with—as accurately as possible—modeling the growth processes of these specific materials. Not surprisingly, the dynamics depend on the way these crystals are grown in the laboratory, namely, crystal growth from the melt (Figure 2.6).

2.3.1 Laboratory Growth

Growing crystals is an art. Largely empirical in nature, growth techniques have evolved over many thousands of man-hours. Growing single crystals of Relaxor-PT alloys, which have the properties you want, is a significant undertaking (though not nearly as difficult as PZTs). A main goal of our research, though distant, is to contribute useful suggestions to this engineering effort, whose rewards may be vast. The principle growing method for these crystals is the high temperature flux technique [2]. Knowing how it works is important; our simulation should incorporate its rules as much as possible.

Depending on the crystal to be grown, relative quantities of highly purified powders of the elements and compounds involved are gathered and dry mixed using a tumbling mill. The mix is then loaded into a Platinum crucible (a bucket, in vulgar terminology) which is placed into an aluminum crucible, sealed, and melted in a tube furnace (1100-1200°C for PZN, PMN). The whole system is then slow cooled (1-5°C/h) and finally furnace cooled to room temperature. As the liquid is cooled, atoms within it start to stick to each other and to the *substrate*, an initial surface on the bottom of the bucket, whereby the crystal grows upward (gravity being what it is). Typical crystal sizes are 3-20 mm, their properties are a different matter. See [11, 13] for more information.

Kinetic Monte Carlo techniques have been shown to simulate well the growth processes of certain crystals under this growth technique [10, 29]. No prior simulations, however, have been done using models with long-range ionic interactions. Our goal is to incorporate such interactions in a multi-species model so as to simulate the growth processes of the new relaxor crystals, and ultimately to guide in the growth of better quality (better ordered) crystals.

2.3.2 Simulations

Kinetic Monte Carlo simulations of growth are based on simplified growth models which concentrate on a few important aspects of the processes involved, while disregarding the others. Although spatially continuous models are possible via MC, discrete models simplify things greatly and are the norm (see Figure 2.8(a)). Discrete models confine all events—processes within the simulation—to a fixed lattice which, in our case, is a simple 3-d array.

In KMC, an environment is set up where the elementary processes of growth can develop over time. In the case of crystal growth, this may be a three-dimensional fixed array of positions where an atom or molecule (dubbed an *adatom* in simulation terms) can occupy a given site. The three elementary processes that characterize most crystal growth simulations are adsorption, where an adatom sticks to the surface, evaporation, where an adatom on the surface desorbs or leaves the surface, and diffusion (or surface migration), where an adatom can migrate across the surface of the growing crystal. Our model deals only with the first two. (Diffusion is considered only a minor effect in the high temperature flux method of crystal growth [29]).

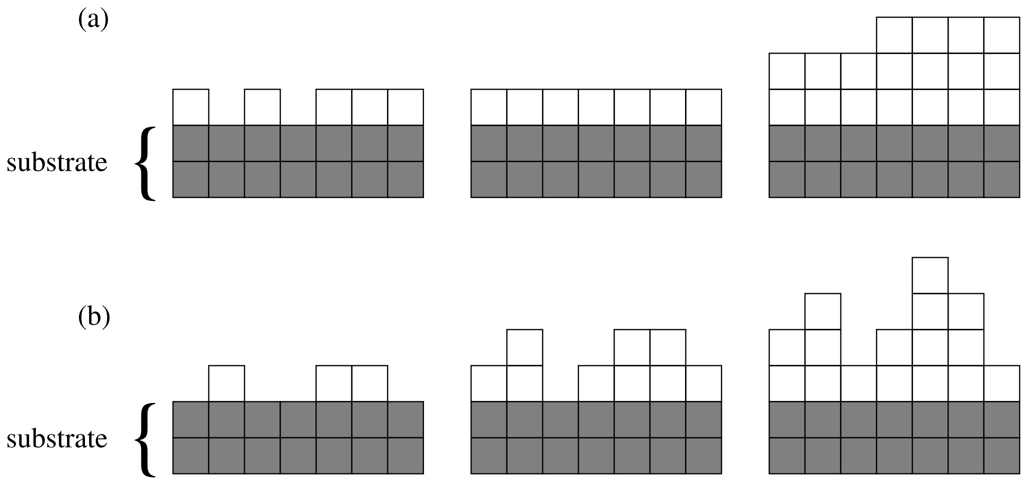

Figure 2.7: Modes of growth in 2-d (time increases from left to right): (a), nucleated or layer-by-layer; (b) rough.

As adatoms stick to the surface, the crystal will "grow." A primary feature of real crystal growth, which KMC has reproduced well, is that growth appears in different phases depending on the environmental parameters at the time. The *modes of growth* have been broadly categorized and are important in qualitatively checking any proposed growth simulation.

The interface between a crystal and the surrounding fluid changes from a thin transition region at low temperatures to a disordered mixture of solid and fluid at high temperatures [28]. This is because an adatom can adhere easily to a rough surface at almost any position, but it will have difficulty adhering to a smooth flat surface [29]. Figure 2.7 illustrates the primary growth modes. In *layer-by-layer* growth, a crystal layer tends to be completed before a new layer is started above it. *Three-dimensional* or rough growth, on the other hand, occurs when many crystal layers grow at the same time, forming hillocks and cavities on the crystal surface. The successful reproduction of the transition from one type of growth to another, a first order phase transition, is one of the great achievements of the traditional KMC method.

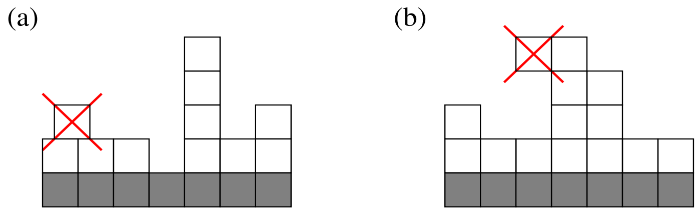

Figure 2.8: Growth simulation restrictions: (a), only discrete placement; (b), the Solid-on-Solid (SOS) restriction.

So, already we have restricted our growth to a discrete lattice without diffusion. Another restriction the model imposes is called the solid-on-solid (SOS) restriction. It imposes the rule that a potential adatom can only stick to the surface where another adatom is directly beneath it. In typical KMC crystal growth simulations, first an initial surface (substrate) is created, then one of two things can happen. An adatom on the surface can evaporate, or an adatom can appear and stick to a site on the surface. The KMC simulation only maintains the state of the completed crystal at any given time. In the KMC simulation, the presence of the melt is only manifested by adatoms appearing at a constant rate dependent on $\Delta\mu$. There is an unlimited supply of them.

There are a great deal of other physical growth effects which the simulation lacks, and the initial tendency is to disregard it as too trivial. Sometimes there is power in simplicity however; our focus on a select few dominant processes may pay off in a fundamental understanding of the system (it has in the past) [10, 29]. Of course, the ultimate validation of our model is by comparison with experiment.

The formulation of KMC then is to define the elementary processes of the system and how they relate to each other. Even with all these simplifications, the result is still too complex for a completely analytic solution. We must resort to simulations.

Chapter 3

The Physical and Mathematical Foundations of KMC

The kinetic Monte Carlo method simulates the stochastic processes of crystal growth by means of a set of probabilities characterizing these processes. Known as direct Monte Carlo (as opposed to indirect Monte Carlo which simulates *deterministic* processes) KMC has achieved much success in the simulation of relaxation processes, i.e., those away from equilibrium. The traditional KMC approach, derived from the *Ising model*, created an environment where all adatoms are the same. The short-range forces they felt were analogous to covalent bonds and there was no notion of charge.

In order to investigate the B-site ordering of single crystal perovskite alloys, our model requires the introduction of the concept of charge in the KMC method. Our simulation defines a number of species, differentiated by charge, which an adatom can be. In addition, we incorporate long-range interactions between species in the crystal lattice because of their apparent importance in determining the LRO prevalent to these systems [3].

Below we review the mathematical and physical foundations for the KMC method. We include in the model a general interaction term which can be either short-range or long-range. We allow for the possibility of short-range (or nearest-neighbor) interactions in order to help benchmark the algorithm that we will develop in Chapter 4.

3.1 Conceptual Overview

Before we exactly characterize the system we will be simulating, it is useful to take a step back and visually assess its structure, elements, and processes. Going in reverse order, its processes are two: adsorption and evaporation to or from the crystal surface. Its elements are adatoms, characterized entirely by charge, which interact with each other through short or long-range potentials and in cooperation with the definition of the processes. Finally, the structure of the simulation is a three-dimensional lattice where adatoms, easily visualized as cubes, can individually and singularly occupy lattice sites. The crystal formed is contiguous, and no holes are possible within its confines. This process will become embedded in your being by the end of Chapter 4, we promise.

3.2 General Derivation

The first thing to realize is that our simulation is only as good as our model. The model characterizes completely the system we are trying to simulate. If the physics is not in the model, it will not be in the simulation. Monte Carlo simulations allow us to investigate average properties of the physical system described completely by the model (this point cannot be emphasized enough). Because natural systems are usually complex and our abilities usually are not, models tend to be very simple. On the other hand, models are very powerful theoretical devices. Focused on a particular physical question, such as the ordering of B atoms in a perovskite alloy, they strip away unnecessary complexity in search of a simpler answer.

As introduced and predicted in Chapter 2, we use a simple ionic model to characterize our physical system. The Hamiltonian,

$$\mathcal{H}(\mathcal{C}) = \sum_{(l,l')} \mathcal{V}_{ll'} + \Delta\mu \sum_{(i,j)} h_{ij} = E(\mathcal{C}) + \Delta\mu \sum_{(i,j)} h_{ij}, \tag{3.1}$$

completely defines the crystal system we will be simulating. A description of the total energy of a physical system, the Hamiltonian in our case is a sum of two parts. The first part, $E(\mathcal{C})$, represents the total potential energy of a crystal configuration, $\mathcal{C}$. Explicitly, it is the sum over $\mathcal{V}_{ll'}$, where

$$\mathcal{V}_{ll'} = \begin{cases} C \sum\limits_{<l \neq l'>} \Delta q_l \Delta q_{l'}, & \text{nearest neighbor} \\ C \sum\limits_{(l \neq l')} \Delta q_l \Delta q_{l'} v_{lr}(\mathbf{l} - \mathbf{l}'), & \text{long-range} \end{cases} \tag{3.2}$$

(which will be explored more fully later) is the Coulomb potential any adatom, at $l = (i, j, k)$, in the three-dimensional crystal lattice feels from another adatom occupying a crystal lattice site $l' = (i', j', k')$. For long-range interactions, periodic boundary conditions in the x-y plane make this an infinite sum (to be dealt with later, where v_{lr} will be defined); for short range interactions (the traditional approach), the sum in $\mathcal{H}(\mathcal{C})$ is only over the the six nearest neighbors, denoted $< l \neq l' >$ (above, below, left, right, forward, behind). A crystal configuration, $\mathcal{C}$, represents any possible arrangement of species and empty space within the crystal lattice. For example, if $S = \{\mathcal{C}\}$ defines the set of all configurations, for a ridiculously small lattice size of $1 \times 1 \times 1$ with the number of species, $n_s = 2$ (say +1 and -1), then $|S| = 3$ and $S = \{\{+1\}, \{-1\}, \{\lambda\}\}$, where λ is defined as an empty site. Thus, for a crystal lattice of any appreciable size, there are many, many possible configurations. (Though you might guess that $|S| = (n_s + 1)^{W \times L \times H}$, it is actually somewhat less than this because under the SOS restriction there can be no hole within the crystal.)

The second part of the Hamiltonian, $\Delta\mu \sum_{(i,j)} h_{ij}$, requires a little deeper explanation. When an adatom moves from the melt to the crystal, aka adsorption, the energy of the system decreases. This energy difference, dubbed $\Delta\mu$, is defined as the difference in chemical potential between a growth unit in the fluid adjacent to the surface and in the solid phase ($\Delta\mu = \mu_s - \mu_l$); since the solid phase represents a lower energy existence, $\Delta\mu$ is *negative.* The magnitude of $\Delta\mu$, then, controls the "sticking" rate in the melt as will become readily apparent below. So, each time an adatom sticks, the energy of the total system decreases by $\Delta\mu$. Since h_{ij} is simply the height of the crystal as a function of (i, j) coordinate, $\Delta\mu \sum_{(i,j)} h_{ij}$ is simply the number of stuck or occupied sites in the lattice times the change in chemical potential, or more physically, the total loss in chemical potential energy caused by all adsorptions. With the Hamiltonian model now understood, we can dive into the formalism of KMC.

The goal of KMC is to represent faithfully the time evolution of the system. As in equilibrium MC simulations, the result must be a thermodynamically possible arrangement. If we look at crystal growth as a time evolution of crystal configurations, it gives us a powerful descriptive tool for deriving the KMC method.

Let us first define $P(\mathcal{C}, t)$ as a time-dependent distribution of configurations. Further, if we define $\mathcal{C}'$ as a crystal configuration with one additional adatom (i.e., after one adsorption), then the transition rate $w(\mathcal{C} \to \mathcal{C}')$ gives the rate of adsorption and $w(\mathcal{C}' \to \mathcal{C})$, the rate of evaporation. The KMC technique can be viewed as a method of solving the master equation [10]

$$\frac{\partial P(\mathcal{C}, t)}{\partial t} = -\sum_{\mathcal{C}} w(\mathcal{C} \to \mathcal{C}')P(\mathcal{C}, t) + \sum_{\mathcal{C}'} w(\mathcal{C}' \to \mathcal{C})P(\mathcal{C}', t). \tag{3.3}$$

Notice that this is just a mathematical description of the competing processes of adsorption and evaporation, relating to the crystal configuration, over time. In the equilibrium limit (as $t \to \infty$), the Boltzmann distribution,

$$P_{eq} = Z^{-1} \exp\left[\frac{-\mathcal{H}(\mathcal{C})}{kT}\right], \tag{3.4}$$

is necessarily a solution to the master equation. Taking further advantage of the equilibrium state, we find that the rate of adsorption must equal the rate of evaporation satisfying the detailed balance condition:

$$w(\mathcal{C} \to \mathcal{C}')P(\mathcal{C}) = w(\mathcal{C}' \to \mathcal{C})P(\mathcal{C}'), \tag{3.5}$$

and consequently from Eq. 3.4

$$\frac{w(\mathcal{C} \to \mathcal{C}')}{w(\mathcal{C}' \to \mathcal{C})} = \frac{P_{eq}(\mathcal{C}')}{P_{eq}(\mathcal{C})} = \exp\left[-\frac{\mathcal{H}(\mathcal{C}') - \mathcal{H}(\mathcal{C})}{kT}\right]. \tag{3.6}$$

The KMC method somewhat artificially separates the processes of adsorption and evaporation. In reality, both events depend on both the chemical and Coulomb potential energy changes involved. Here, we separate them. We choose the adsorption rate to be constant,

$$w(\mathcal{C} \to \mathcal{C}') = e^{\Delta\mu/kT}, \tag{3.7}$$

and derive the evaporation rate subsequently:

$$\begin{aligned} w(\mathcal{C}' \to \mathcal{C}) &= \frac{P_{eq}(\mathcal{C})}{P_{eq}(\mathcal{C}')} w(\mathcal{C} \to \mathcal{C}') \qquad (3.8)\\ &= \exp\left[\frac{\mathcal{H}(\mathcal{C}) - \mathcal{H}(\mathcal{C}')}{kT}\right] e^{\Delta\mu/kT} \\ &= \exp\left[\frac{\left\{E(\mathcal{C}) + \Delta\mu \sum\limits_{(i,j)} h_{ij}\right\} - \left\{E(\mathcal{C}') + \Delta\mu \sum\limits_{(i,j)} h_{ij} + \Delta\mu\right\}}{kT} + \frac{\Delta\mu}{kT}\right] \\ &= \exp\left[\frac{E(\mathcal{C}) - E(\mathcal{C}')}{kT}\right] \\ w(\mathcal{C}' \to \mathcal{C}) &= e^{-\Delta E/kT}. \qquad (3.9) \end{aligned}$$

This is not wholly unrealistic in that, for the type of growth we are simulating here—by the flux method—there is reason to believe that the adsorption rate is fairly constant [29] and dependent on the chemical potential difference. The quality of the simulation results will give further justification.

In Equation 3.9, ΔE represents the change in total potential energy in the crystal when an adatom evaporates from the surface. When adatoms evaporate, the total energy of the crystal tends to increase. Therefore, $E(\mathcal{C}) - E(\mathcal{C}')$ is, on the average, positive. Equation 3.9 reinforces our physical intuition that the greater the Coulomb attractive force a stuck adatom feels from within the crystal, the less likely it is to evaporate. It is not difficult to show that ΔE is given by

$$\Delta E = \mathcal{V}_o = \begin{cases} C \sum\limits_{<l' \neq o>} \Delta q_o \Delta q_{l'}, & \text{nearest neighbor} \\ C \sum\limits_{(l' \neq o)} \Delta q_o \Delta q_{l'} v_{lr}, & \text{long-range} \end{cases} . \tag{3.10}$$

where the vector $o \equiv (i, j, k)$ and Δq_o define the position and charge of the particle in question.

Look familiar? The change in E is simply the Coulomb potential inherent to a particular surface adatom (with charge Δq_o). Independently of whether we are considering only nearest neighbor interactions or long-range interactions, the potential is simply a sum over the familiar Coulomb equation, kq_1q_2/r, replacing the true atom charges with our relativized B-site charges. C contains all necessary constants (normalizing the distance between lattice cells to match r). And now, finally, our opportunity comes to explain v_{lr}.

3.3 The Ewald Summation for Finite Height

For long-range electrostatic interactions,

$$v_{lr} = \sum_{l=1}^{\infty} \frac{1}{|\vec{r_{iO}} - \vec{R_l}|}. \tag{3.11}$$

This $1/r$ summation over all lattice vectors (infinite because of periodic boundary conditions) is very slowly and conditionally convergent. Brute force methods in constant time are impossible. Another method must be found. The Ewald summation is a technique for summing the interaction between an ion and all its periodic images in a crystal lattice in constant time. It breaks the slowly converging sum into two rapidly converging complimentary sums: one in real space, the other in reciprocal space. The method also allows us to incorporate the charge neutrality of the system (remember our simulation is a B-site reduction of the complete system). The Ewald summation effectively subtracts a negative background term which can be thought of as a free electron contribution or the interaction of a site in the crystal matrix with its own images in the PBC-produced infinite lattice.

Derivations and implementations of the Ewald summation method in the integer-dimensional case are readily available [32, 33]. These are inapplicable to our situation. We need a general method by which the adatom's height above the plane can be accounted for in the Ewald method: an Ewald method for finite height. This problem has been solved by T.J. Walls, and we refer you to [31] for a detailed treatment of the solution.

The result is an equation for the potential, v_{lr}, as a function of $r \equiv l' - l$ defined for a +1—+1 charge pair. For our discretized model, all that is needed is a table of these potentials for all lattice vectors, $l' - l$, in the finite crystal matrix. Refer to Appendix D.

Armed with our evaporation and adsorption rates (Eqs. 3.9 and 3.7), a formulation of ΔE (Eq. 3.10), and a method for calculating long-range potentials (Section 3.3), we are ready to convert our simulation into an algorithm and implement it.

Part II

Algorithm and Implementation

Chapter 4

KMC Algorithms

Our starting point is the kinetic Monte Carlo algorithm introduced by Bortz, Kalos, and Lebowitz [5]. Although the algorithm is very effective in simulating the processes discussed in Section 3.2 for Ising-like systems, we will see that it becomes increasingly inefficient for the long-range system with increasing lattice size. We propose a new algorithm which overcomes these difficulties for long-range systems while sharing the simplicity of the BKL algorithm.

4.1 Initialization

The first step in a simulation program is to create and initialize the environment. We create the crystal matrix, a simple three-dimensional array of size $L \times L \times H$ with periodic boundary conditions (PBC) imposed, where L is the width and length of the matrix and H is the height. The "crystal" is initialized empty except for an initial surface, or substrate, constructed from the types of species defined in the simulation. The species are summarily defined by charge, no limit is placed on their number. The initial values of the parameters are set: the chemical potential difference, $\Delta\mu$, and the temperature, kT. Although not currently implemented, the simulation program can be updated so that these two environmental parameters can change over time during a simulation run. And that brings up a good point. How is time defined in a MC simulation? We define a *MC time step as* L^2 *events*—adsorptions or evaporations—and fractional in nature. Thus, after one event, simulation time would increase by $1/L^2$. For simplicity in describing the algorithms, $1/L^2$ MC step = 1 time step.

The simulation can be set to run for a given number of MC time steps. In addition, the number of layers to be grown can be set; the simulation will quit after this value has been reached. The crystal cannot outgrow its matrix. If the ceiling ($z = H$) or floor ($z = 0$) of the crystal matrix is reached, the simulation will stop. Only these four constraints control the time length of the simulation run.

The above initializations and constraints cap the beginning and end of the simulation; its body consists of a loop over the KMC algorithm. The algorithm will be given for one $1/L^2$ of an MC time step, a time step. Below we first review the BKL algorithm and explain why the algorithm as it stands is not effective for our model with long-range interactions. We then describe a new algorithm which allows us to effectively simulate our model. Both algorithms accurately simulate the evolution of the system according to the KMC method. Their differences lie in implementation and efficiency.

4.2 The BKL Algorithm

The basic idea behind any KMC algorithm is to convert the process rate equations (Eqs. 3.7, 3.9) into probabilities, and through continual realization of events (evaporations, adsorptions), explore the evolution of the system. This process is based on Monte Carlo methods, whereby a *Markov chain* is generated to sample the *state space* of all possible crystal configurations. The MC process ensures that a viable path is taken conforming to the Boltzmann distribution of the system.

The KMC algorithm proposed by Bortz et al. was a considerable improvement over the primitive KMC algorithm [10], and is how most KMC calculations are now done. After a surface site is randomly chosen, an event list (the BKL addition), is created wherein each event (evaporation, adsorption, etc.) is given a probability of occurrence. The probability is determined by dividing the calculated rate of the event process by the sum of the maximum rates for all processes. A random number, r, on the interval $[0, 1)$ is then generated and an event is chosen from the event list whose probability spread encompasses r. The sequence of steps is as follows:

(i) Randomly select an (i, j) pair from the L^2 possible surface sites.

(ii) Generate a list, E, of possible events at (i, j).
In our case, $E = \{adsorption, evaporation, no\ event\}$.

(iii) Define the probability each event will occur by normalizing the event rates:
$P_{evap} = \frac{w_{evap}}{w_{max}}$, $P_{adsorb} = \frac{w_{adsorb}}{w_{max}}$, $P_{no\ event} = 1 - P_{evap} - P_{adsorb}$.

(iv) Generate a random number $r \in [0, 1)$ and choose the first event E_i such that $\sum_{i}^{n} P_i \geq r$.

(v) Generate new configuration $\mathcal{C}$ based on chosen event E_i.

We set $w_{adsorb} = e^{\Delta\mu/kT}$, $w_{evap} = e^{-\Delta E/kT}$, and $w_{max} = e^{\Delta\mu/kT} + e^{-\Delta E_{min}/kT}$. ΔE_{min} is the largest potential a surface adatom can experience (for nearest-neighbor interactions, this is 5; for long-range interactions, it depends on $L^2 \times H$).

Since $\Delta\mu$ is constant, its contribution to w_{max} is constant no matter the lattice size or situation. For nearest-neighbor interactions $\Delta E_{min} = 5$, and its contribution is constant as well. Because the average potential of a potentially evaporated surface adatom will be less than ΔE_{min}, the probabilities P_{evap} and P_{adsorb} will not add to one: an evaporation or adsorption will not always occur.

In the case of long-range interactions, ΔE_{min}, being the sum contribution of all interactions with all particles in the lattice (with PBC), depends on the lattice size of the crystal. The interaction for evaporation at a given site (i, j) can vary a great deal dependent on the configuration $\mathcal{C}$, and can therefore be much smaller than ΔE_{min}. In fact, for small lattice sizes, ΔE_{min} even varies with system (matrix) size (although eventually for very large lattice sizes ($\gg 100$), ΔE_{min} will converge to a constant value). So the evaporation rates for many chosen (i, j) would be far less than that defined by ΔE_{min}. The probability of evaporation (or adsorption for that matter) will be driven to zero. (Choosing a ΔE_{min} for a specific matrix size is an added difficulty.) With the necessity now well defined, we developed a new algorithm which makes KMC with long-range interactions feasible.

4.3 A New Algorithm

The new algorithm goes a step beyond the BKL algorithm by creating an event list including every possible event for every possible surface site. The added complexity comes from the need to store and update an array of surface potentials, the calculation of an event list from these surface potentials (of length $2N$ now), and finding the potential event in this list (a non-trivial matter to do quickly for large N). The benefit comes in that an event is guaranteed to take place with each iteration of the algorithm and that the need for ΔE_{min}, whose large magnitude would drive most evaporation probabilities to zero and render the algorithm ineffective, is negated completely. We can guarantee this because evaporation/adsorption rates for all possible sites are normalized with respect to each other. The sum of the probabilities that an adsorption or evaporation can occur at any site are therefore one. The algorithm proceeds as below:

(i) Generate a list, E, of all possible events per time step. There are $2N$ possible events: an evaporation or an adsorption could happen on each of the $N = L^2$ surface sites.

(ii) Calculate the rates (w) of adsorption and evaporation for each site on the surface $(2N$ rates).

(iii) Normalize these $2N$ rates, giving probabilities, P_i, for evaporation/adsorption on site 1, on site 2, . . . on site N, which all add to unity.

(iv) Generate a random number $r \in [0, 1)$ and choose the first event E_i such that $\sum_i^{2N} P_i \geq r$. An event will always be chosen.

(v) Generate new configuration $\mathcal{C}$ based on chosen event E_i.

The rates, w, are defined as in Section 4.2.

Of course, the 5-step procedure as defined above gives little allusion to the complexity in implementing the simulation. The full source code can be found in Appendix E. To further clarify the process, we will trade an abstract description for a real life example: a simple walk through of the simulation program for an absurdly small system.

4.4 A 2×2 Walk Through

An excellent example system (with which you will become very well acquainted) is that of rock salt. Consisting of two species of equal quantity, equal and opposite charge, and exhibiting simple $[111]_{1:1}$ ordering for a wide range of environmental parameters, it is a common structure of many types of crystals. We begin by initializing the simulation,

$$\begin{aligned}
\Delta\mu &= -2.0, \\
kT &= 1.0, \\
L &= 2,\ H = 5, \\
V &= \{(0,0,0) : \infty,\ (0,1,0) : 1.378952,\ (1,0,0) : 1.378952,\ (1,1,0) : 1.142361; \\
&\quad (0,0,1) : 1.000000,\ (0,1,1) : 0.878519,\ (1,1,1) : 0.878519, \\
&\quad (0,0,2) : 0.823822;\ (0,1,2) : 0.500000,\ (1,0,2) : 0.495872,\ (1,1,2) : 0.495872, \\
&\quad (0,0,3) : 0.333333, \ldots\},
\end{aligned}$$

where $\Delta\mu$ and kT are given reasonable and relative values (more on their choice in the next chapter), L and H define the crystal matrix dimensions, and V is the Ewald table of $+1$—$+1$ potentials for different lattice vectors, (i, j, k), in this matrix. For

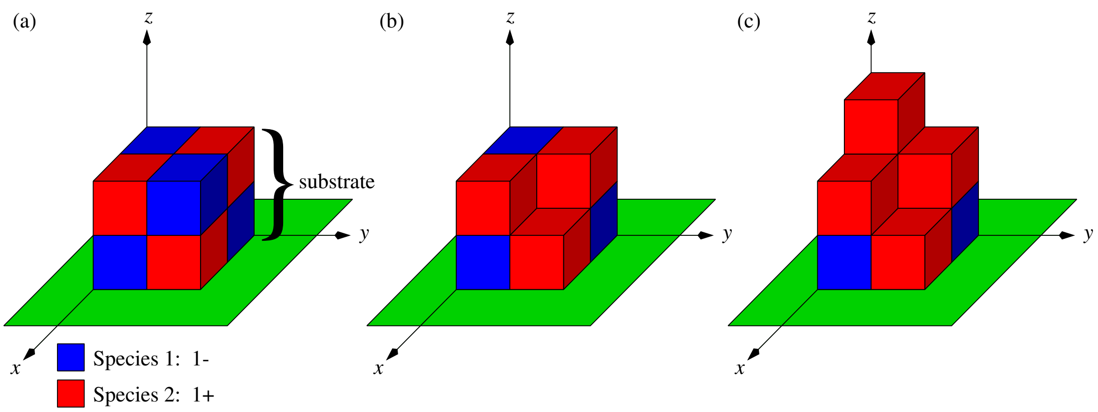

Figure 4.1: 2×2 Crystal Evolution: (a), the initial surface; (b), after one evaporation at site 3; (c), after one adsorption at site 0.

nearest-neighbor interactions, the potential table would consist simply of zeros with ones at (1,0,0), (0,1,0), and (0,0,1). That is how easy switching from long-range to short-range interactions is in our program. The initial surface is created of height two with perfect ordering, as in Figure 4.1(a).

Besides the crystal matrix itself, the simulation program maintains two other data types. The surface potential array, a not quite accurate title, holds the potential or binding energy (Equation 3.10) for each of the N possible next adatoms. Defined as

$$U_{ij} \equiv U_o = \sum_{(o \neq l')} \Delta q_o \Delta q_{l'} V(|\mathbf{o} - \mathbf{l}'|), \tag{4.1}$$

where $o = (i, j, k)$ is the position being considered for the surface adatom (note that k is omitted in the subscript of U because, due to the SOS restriction, (i, j) uniquely determines a surface position). The sum over $\mathbf{l}'$ is over all vectors in the present, *finite* crystal matrix, its initial state for configuration (a) will be

$$U(0) = \begin{pmatrix} -1.682327 & -1.682327 \\ -1.682327 & -1.682327 \end{pmatrix}.$$

Although preferentially described by a 2-d array, in our simple example, we label the surface sites 0 through 3.

The second data type, an event list P_m of length $2N$, maintains the normalized event probabilities for each surface site. It is defined for each algorithm step as

$$P_{m=1\ldots 2N} = \frac{w_m}{\sum\limits_m w_m}, \tag{4.2}$$

where

$$w_m = \begin{cases} e^{\Delta\mu/kT}, & \text{if } m \text{ is even,} \\ e^{U_{ij}/kT}, & \text{if } m \text{ is odd,} \end{cases} \tag{4.3}$$

and (ij) gives the (x, y) position of the $(m+1)/2$-th surface site. From this definition we see that even numbers of i correspond to adsorption event probabilities and odd numbers to those for evaporation. With the information contained and updated in these three data constructs—the crystal matrix, the surface potential array, and the event list, we can realize the new algorithm and simulate crystal growth.

Beginning with configuration (a) of Figure 4.1 and the surface potential array

$U(0)$ we can calculate the first event probability list,

$$P(0) = \left\{ \begin{array}{ccccc} m & (x,y,z) & Type & Rate & Probability \\ \hline 1 & (1,1,3) & adsorb & 0.135335 & 0.105311 \\ 2 & (1,1,2) & evap & 0.185941 & 0.144689 \\ 3 & (2,1,3) & adsorb & 0.135335 & 0.105311 \\ 4 & (2,1,2) & evap & 0.185941 & 0.144689 \\ 5 & (1,2,3) & adsorb & 0.135335 & 0.105311 \\ 6 & (1,2,2) & evap & 0.185941 & 0.144689 \\ 7 & (2,2,3) & adsorb & 0.135335 & 0.105311 \\ 8 & (2,2,2) & evap & 0.185941 & 0.144689 \end{array} \right\},$$

for all possible events. Visually confirmed, the adsorption rate is constant as expected. The evaporation rate happens to be constant across all possible event sites as well, because the potential across all surface adatoms is the same. The perfect crystal—the two layer initial surface—seems pretty happy where it is (at least for these choices of $\Delta\mu$ and kT). Nevertheless, according to our algorithm, an event must occur.

To choose from the list, we generate a random number $r = 0.885954$. Following step iv's formula, r selects event $m = 8$, an evaporation, to occur at site $(2, 2, 2)$. This corresponds to the crystal configuration of Figure 4.1(b). Thus, one iteration of the new algorithm is complete. Now we must recompute the surface potential array and begin again.

After one time step,

$$U(1) = \begin{pmatrix} -0.867998 & -1.181894 \\ -1.181894 & -2.000866 \end{pmatrix}.$$

A real energy landscape is starting to unfold. Notice that surface site $(1,1,2)$ has the least binding energy ($|\Delta E| = 0.86$), our event list should reflect this in a high probability of evaporation. Site $(2,1,1)$ conversely should have the least probability of evaporation. The event list after 1 time step then is

$$P(1) = \left\{ \begin{array}{lllll} m & (x,y,z) & Type & Rate & Probability \\ \hline 1 & (1,1,3) & adsorb & 0.135335 & 0.079156 \\ 2 & (1,1,2) & evap & 0.419791 & 0.245531 \\ 3 & (2,1,3) & adsorb & 0.135335 & 0.079156 \\ 4 & (2,1,2) & evap & 0.306679 & 0.179373 \\ 5 & (1,2,3) & adsorb & 0.135335 & 0.079156 \\ 6 & (1,2,2) & evap & 0.306679 & 0.179383 \\ 7 & (2,2,2) & adsorb & 0.135335 & 0.079156 \\ 8 & (2,2,1) & evap & 0.135218 & 0.079087 \end{array} \right\},$$

confirming our predictions. But this time an evaporation does not occur, for the newly generated random number, $r = 0.057315$, selects an adsorption at site $(1,1,3)$, and the crystal gains an adatom. The species type of the new adatom is determined from the relative quantities of the two species using another random number.

That concludes a two-event realization of our new algorithm, hopefully the process has become clearer. It should be said that the new algorithm is not limited to a fixed number of species or events. Either algorithm can be generalized easily to include other processes, such as diffusion.

4.4.1 Time Complexity

Since in our new algorithm, an event happens at every step, it becomes possible to fairly accurately calculate the time complexity of the simulation. If we define $N_T \approx L^2 \times H$ as the total number of adatoms in the crystal matrix, we will see that the algorithmic complexity of the simulation scales as $\mathcal{O}(N_T)$. For each time step, the bulk of the algorithm's time is spent doing three things: updating the surface potential array, creating the event probability list, and finding the selected event in this list.

Naively, one would assume that calculating the surface array would be of order N_T^2 because each of the N adatoms binding energies from all the adatoms in the lattice must be calculated. However, this is not necessary. Each time an event occurs on a particular surface site, whether its an adsorption or evaporation, the new energy contribution need only be added or subtracted to the $N-1$ other surface sites depending on the event type and charge of the new or vanquished particle. To update the changed site though, we still need to sweep through the entire lattice, adding N_T. The total time for updating the surface array is $\mathcal{O}(N_T + N)$. Creating the event list takes $\mathcal{O}(2(N+1))$ time steps, not $4N$ because all the adsorption rates are constant. This must be done twice, once for the rates, and once for the probabilities. Finally, to search the event list, linearly, would take $\mathcal{O}(2N)$, though significantly less than this most times. This time can be improved with a more advanced search technique. Although implemented linearly now, other solutions will become necessary with very large lattice sizes. (Currently, the linear search does not appear to contribute much to the running time of the simulation).

The total complexity then for one time step is given by $\mathcal{O}(N_T + N + 2(N+1) + 2N) = \mathcal{O}(N_T)$. Remember, though, that N_T, unlike N, is not constant throughout the life of the simulation. It will grow or shrink corresponding to how the crystal evolves. For a crystal with a positive growth rate, the simulation will progressively slow down.

Simulation Time versus Real Time

Our simulation definition of time, the MC time step, does not correlate with the real time evolution of a growing crystal system. The new algorithm ensures an event for each time step, all time steps being equal. In reality, the speed at which events occur will depend on the configuration of the surface (rough, etc.). It is possible, although not done here, to go back to the master equation (Equation 3.3) and derive a configuration dependent time relation [38, 39]. This allows for the translation of simulation time to system time (useful for experimental comparison) and is an avenue of future work (Chapter 7).

It can be shown that our short-range derivation from Chapter 3 is identical to the traditional non-ionic KMC method. This will prove valuable to us when we are analyzing and assessing the new algorithm, implemented via our ionic, long-range KMC method, in the next chapter.

Part III

Results and Analysis

Chapter 5

Analysis

Our crystal growth simulation, at its most basic level, is a realization of a two parameter model. The variation of the temperature of the system, kT (given in units of energy), and the adsorption driving rate, $\Delta\mu$, will control the changing characteristics of the growing crystal. Our analysis here will also focus on two things: first, proving that our simulations for both the nearest-neighbor and long-range models are correct (they display physically meaningful, verifiable, and realistic behavior); and second, contingent on the first, the exploration of the new and exciting phenomena the addition of long-range interactions brings to crystal growth simulations.

We will begin our investigations through one of the most typical of crystal configurations, that of rock salt. Although the bulk of our time will be devoted there, other system types—including those that encompass the Relaxor-PT materials—will be touched upon. The abilities and necessity of the long-range model become very evident here. Our conclusion will consist of a brief summary of the work, observations from our simulations, and an overview of needed tests that time and computing power would not permit.

5.1 Rock Salt Type Systems

Sodium Chloride (NaCl), in its single crystal form, is called rock salt. Composed of alternating layers of Na^{1+} and Cl^{1-} in both the [111] and $[\bar{1}11]$ directions, it typifies the crystal ordering of a wide variety of materials. It also represents the simplest crystal our ionic model can grow, providing an excellent testbed for analysis.

Heterovalent binaries such as described by $\mathrm{II}_{1/2}\mathrm{VI}_{1/2}$ ($\Delta q_B = \pm 2$) or $\mathrm{III}_{1/2}\mathrm{V}_{1/2}$ ($\Delta q_B = \pm 1$, see Table 2.2) are good perovskite examples of rock salt type ordering. Although the absolute value of the charge between the two systems is different, this only effects the magnitude of the effective energy scale factor ($n^2e^2/2\epsilon a$) [3], not the ordering or any crystal growth effects we are interested in. Thus we can safely choose $\Delta q_B = \pm 1$, labeling our simulations rock salt, and have our results be more universally applicable.

5.1.1 The Nearest-Neighbor Model

The nearest-neighbor ionic simulation itself represents a step forward over the traditional KMC method. Although it introduces the addition of electrostatic interactions between differently charged adatoms, the rock salt type system gives us the closest comparison to traditional crystal growth simulations that we have. We will therefore use it as a first test in verifying the legitimacy of the long-range simulation to come. Since the only difference between the nearest-neighbor and long-range simulations—as discussed in Chapter 4—is the electrostatic potential table used, correct behavior in the nearest-neighbor model should forecast the same for the long-range (assuming the Ewald table is accurate). The structure of the analysis here will match very closely the long-range discussion that succeeds it.

Choice of Parameters

The choice of parameters represents a tuning of the model towards regions of physically meaningful and interesting behavior, which can help us gain insights into the physical evolution of the system. In our case, the choice of $\Delta\mu$ reflects the energy scale of the crystal we are simulating. The solid-liquid interface energy should be similar to that of the electrostatic energy felt by a surface adatom from the rest of the crystal (for realistic growth simulations). Changes in kT shift the overall energy scale of the system. Thus the energy and parameter values can be renormalized with respect to each other, the preliminary goal being to discover relative values for the two parameters such that the system displays the behavior we are interested in.

As detailed in Chapter 3, we chose our adsorption rate (dependent on $\Delta\mu$ and kT) to be constant and derived the evaporation rate relative to it and dependent on the electrostatic potential of the adatom in question (and kT as well). It is necessary then to choose $\Delta\mu$ such that the adsorption and evaporation rates are consistent, relative to each other, for the type of growth behavior sought. Because the interaction energy scale of the nearest-neighbor model has no correlation with matrix size (the maximum potential an adatom can feel is always 5), our choices for $\Delta\mu$ and kT, and the behavior they represent, should remain consistent across crystal sizes. As we will see, this is not necessarily so for the long-range model.

Through a little reasoning and a little more trial-and-error, a region (in parameter-space) of interesting growth was found to lie between $\Delta\mu = 0$ and -2.5 for $kT \approx 1$. If we consider the value of ΔE_{min} for the nearest-neighbor model (i.e. 5) and that the average of ΔE in our evaporation rate equation for $kT = 1$ will typically be less than this, the choice of the $\Delta\mu$ region starts to gain some theoretical backing. Figure 5.1, which will be discussed more fully in the next section, shows KMC calculations of the growth rate behavior in this region.

We have created a simulation environment—confined to the restrictions of our model Hamiltonian (Equation 3.1)—whose physical nuances and oddities we can explore. And because the simulation is an exact (probabilistic) realization of the model, it corresponds directly to the real system (in as much as the model captures the real physics of the system). The first set of behaviors to investigate and verify against our physical and historical expectations is that of crystal growth.

Growth Rates and Qualitative Behavior

We define the growth rate of the crystal as the number of adatoms adsorbed divided by the number of MC time steps. It does not factor in the roughness of the crystal surface, though this is a possible future enhancement (Chapter 7). The simulation

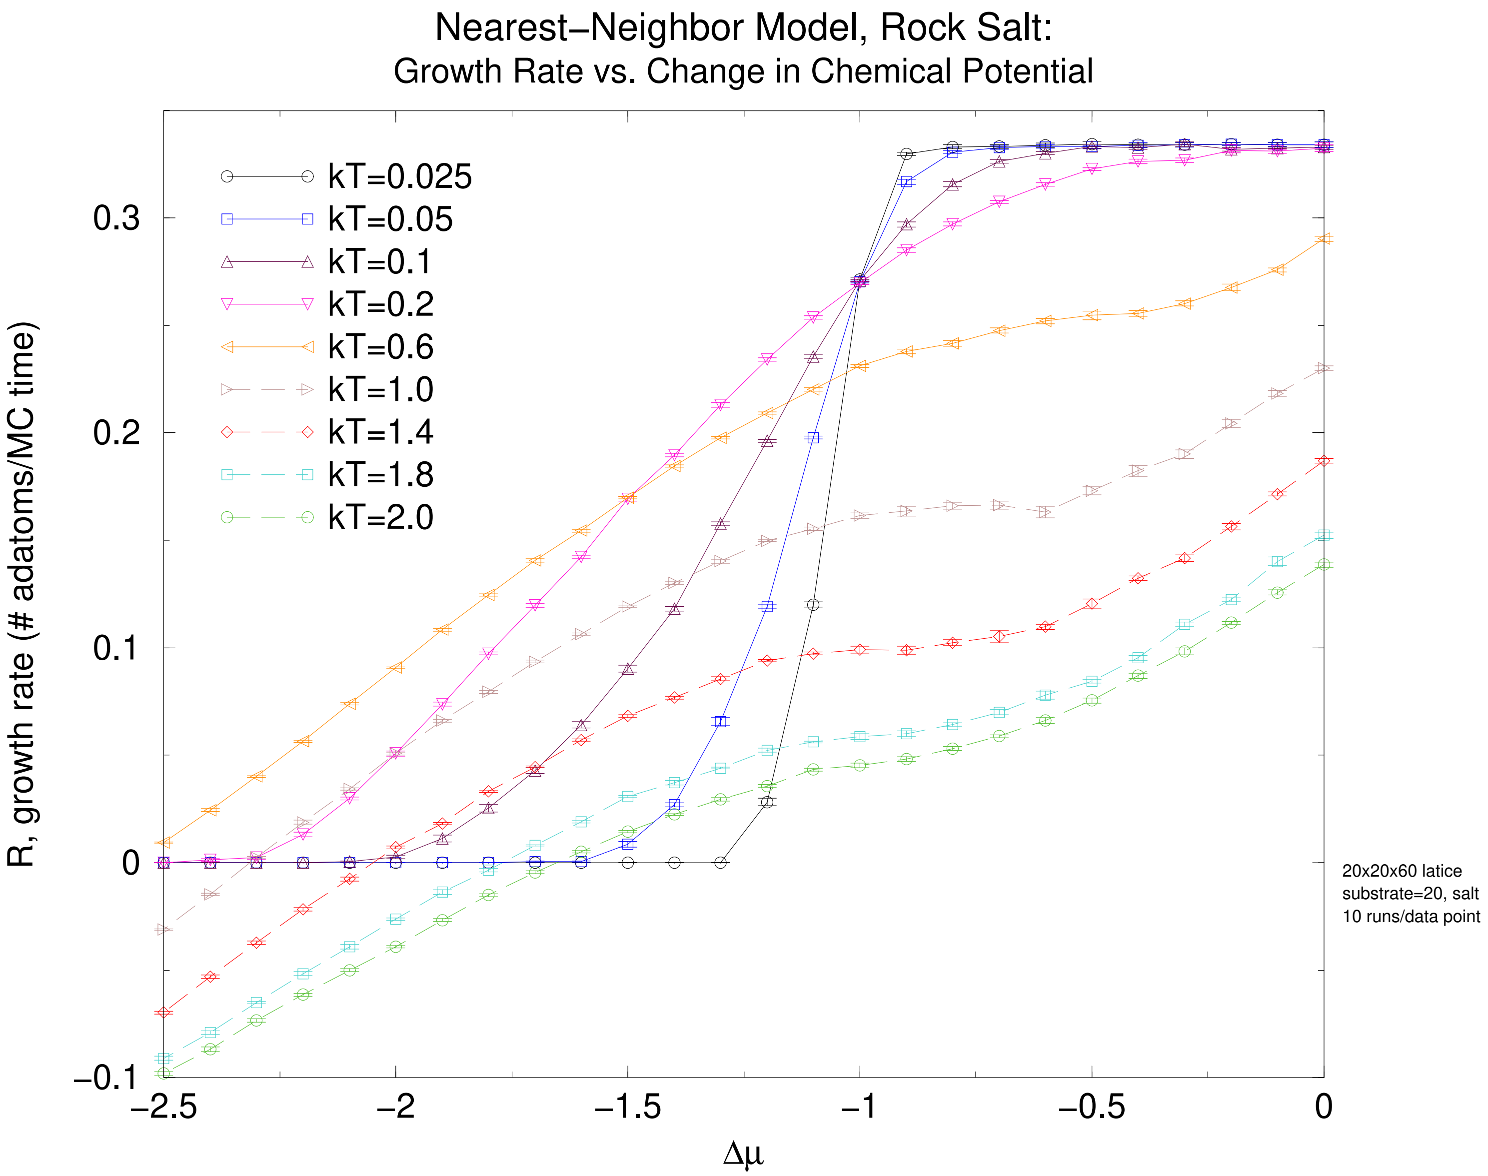

Figure 5.1: A range of nearest-neighbor growth rates for rock salt. The error bars are discussed in the text.

growth rate curves, over a range of temperatures, are given in Figure 5.1. The selected temperatures range from very, very cold ($kT = 0.025$) to hot ($kT = 2.0$). There's much to notice, discover, and learn from this figure and thankfully so; it took almost two days of computational time, spread across nine different computers, to construct it. A $20 \times 20 \times 60$ crystal matrix was used, a balance between statistical breath and time issues.

Each simulation ran for 200 MC steps or 50 layers (whichever came first) although the actual time varied, depending on the growth restrictions outlined in 4.1. Appendix C.2 goes through an analysis of the period of the random number generator (RNG) used, to ensure that it was adequate for the lengths of the run, i.e., it did not cause any under-estimation of the statistical errors. Every data point in Figure 5.1 represents the average growth rate value of ten separate crystal simulation runs. The error bars address the normal statistical error ($\sigma/\sqrt{n}$) relavent to probabilistic (MC) simulations. These growth rate curves, the product of many simulation runs over the parameter region, should display behavior consistent with the definition of the model and that of crystal growth in general.

Let us review. The rate of adsorption is given by $e^{\Delta\mu/kT}$ and the rate of evaporation by $e^{-\Delta E/kT}$. As $\Delta\mu$ increases, an adatom is more likely to stick to the crystal surface. As kT decreases, the rate of sticking will increase but more importantly, the "selectiveness" of evaporation will increase. A lower kT will in effect increase the energy differences ($\Delta E/kT$) between competing configurations. The direct result, as growth is concerned, will be that adatoms will increasingly favor having more neighbors instead of less (layer-by-layer growth vs. rough) and charge similar adatoms will seem even more repulsive. So we are left with the competing forces of $\Delta\mu$ and the temperature. If the rate of adsorption becomes overwhelming (very high $\Delta\mu$), adatoms will stick anywhere, no matter the location or ionic adversity, and the growth rate will be high. Alternatively, if the temperature becomes too high, the crystal will melt or get eaten away—the preferred phase becomes the liquid phase and negative growth is the consequence (notice this dynamical phase transition at $\Delta\mu \approx -1.1$). The dance between these two forces gives us the behavior in Figure 5.1.

Though we have saved most of the revelry of success—the exploration of the different surface growth modes—for the long-range section, we show here some representative results for short-range interactions. Figure 5.2 introduces you to the visualization environment, *tview*, written in-house to explore the crystal growth of our simulation (see Appendix E). Figures 5.3 and 5.4 pictorially illustrate the behavioral evolution of the crystal surface as a function of $\Delta\mu$ and kT. For the cases displayed here, we see a wide spectrum of growth modes: no growth, layer-by-layer growth, rough growth, and very rough growth. The ordering is almost perfect in all cases (a testament to the stability of the rock salt system and to the correctness of the implementation). Only for very high temperatures and $\Delta\mu$'s do we start to see imperfections.

Another interesting phenomena of these intensive cases is the presence of what appear to be "towers." Although in the extreme they justify our definition of excessive roughness, they have a very sound physical explanation. At very high sticking rates, crystal formations nucleate on the surface independent of one another and likely with several defects. As these initial formations grow, they will eventually meet and will not necessarily match up. That is, the alternation of charges will not fit the rock salt description. Since the crystal is growing so fast, there is not time to “fix” these imperfections (which may be very deep by now) with evaporations and the most energy efficient alternative is the creation of towers of alternating charge. A long-range example, Figure 5.14, showcases this. The presence of these towers brings us a step closer to real crystals, which commonly have macro-imperfections such as *dislocations*. In a non-finite, non-fixed world, two mismatched regions of the crystal would in fact still be connected, a *strain* in the crystal would be the result.

Comparison to Traditional KMC

The creation of towers is just one prediction that a traditional, non-ionic KMC simulation would miss. Because all adatoms are the same, the idea of an ordering mismatch is completely foreign. We might conclude, then, that a region of “reasonable roughness” for a traditional KMC simulation would be much wider than for the ionic model. This might explain the maximum growth rate of ≈ 0.3 in our region. Expanding the region of $\Delta\mu$ into the positive, although resulting in wholly unrealistic growth, may match more closely the growth rates of traditional KMC models. Do not mistake the flattening out of the very low temperature curves as an upper limit to our model, this may just be a temporary equilibrium situation for these values of kT. In any case, more runs are in order.

5.1.2 The Long-Range Model

The moment of truth has come. The incorporation of long-range electrostatic interactions into a KMC simulation was the first and primary research goal of this project and that on which all future goals depend. The growth of simple salt will serve as our first test of the implemented method. The complete simulation code, including an overview, can be found in Appendix E.

We begin with the same size crystal matrix and range of parameters as in our nearest-neighbor investigations, the resulting growth rate graph is shown in Figure 5.5. Immediately we see differences from the nearest-neighbor results, but we also see similar trends. Let us think about how the addition of long-range interactions reflects itself in the processes of adsorption and evaporation. Because the Coulomb contribution of all adatoms in the lattice are considered and not just the nearest-neighbors, the energy differences between potential crystal configurations ($|\Delta E|$) will be less. In other words, the energy difference between a surface adatom with five neighboring adatoms and one having only one will be much greater in the nearest-neighbor model than in the long-range model. The system's charge neutrality ensures that this difference is narrowed even further; all potentials will become closer to zero. All this contributes to a narrowing of the primary region of interest: where growth behavior spans, with a little leeway, the transition region from negative growth to

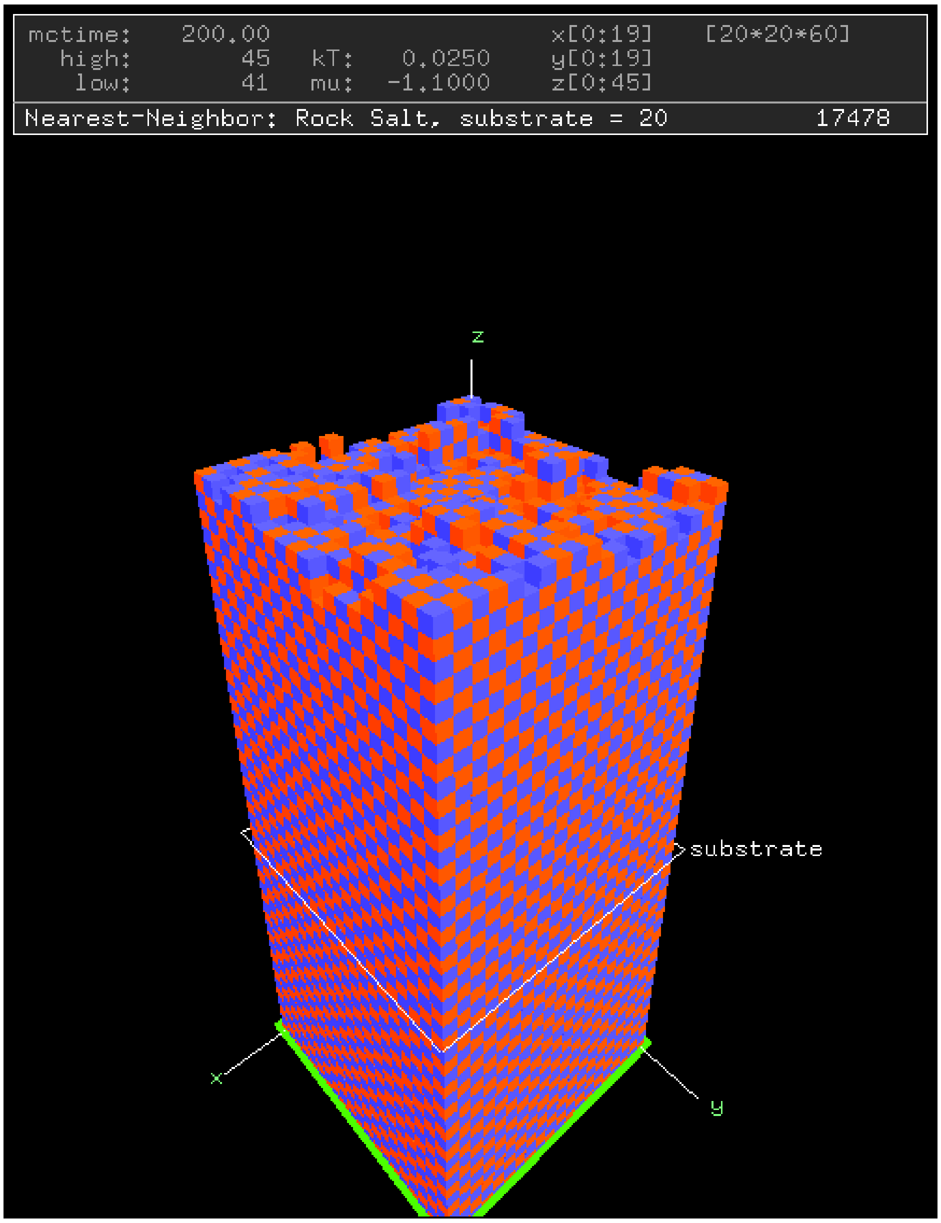

Figure 5.2: The Simulation Visualization Environment, *tview*: the crystal, in this case of dimensions 20×20 is displayed, as it grows, in three dimensions. The zoom and rotation of the matrix can be manipulated in real time. The white lines midway up the base of structure labeled "substrate" indicate where the initial surface ends and the growth processes began. The current value of MC time is displayed in the upper left corner along with the parameter values and the level of the highest and lowest layer of the crystal. Although the crystal matrix is finite, periodic boundary conditions (especially in the long-range model) allow us to imagine the x and y dimensions as infinite. In this case, for rock salt, blue cubes are of charge $1-$ and red cubes are of charge $1+$.

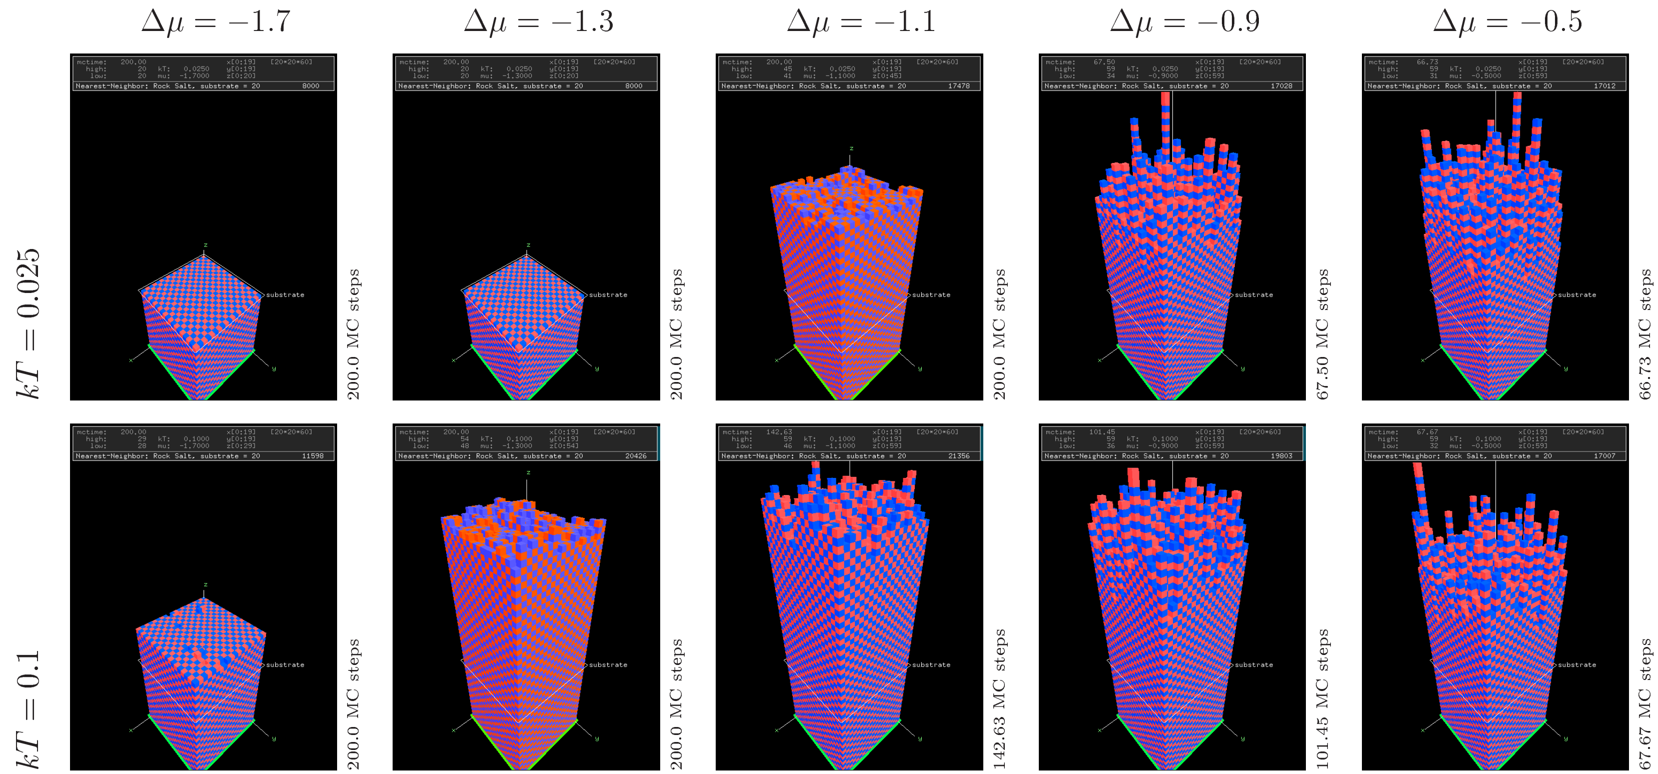

Figure 5.3: **Nearest-Neighbor** Simulations for Rock Salt (Page 1), $20 \times 20 \times 60$ crystal matrix: from left to right $\Delta\mu$ increases from low to high; temperature increases going downward, from very, very low to very low. For both temperatures we see a nice transition from no growth to layer-by-layer growth to rough growth—and very ordered, as the low temperatures would have us predict. Do not let the size of the crystal fool you, look at the number of MC steps and Figure 5.1 to realistically judge the growth rate. (Remember the simulations stops when the artificial bounds are reached (see Section 4.1).)

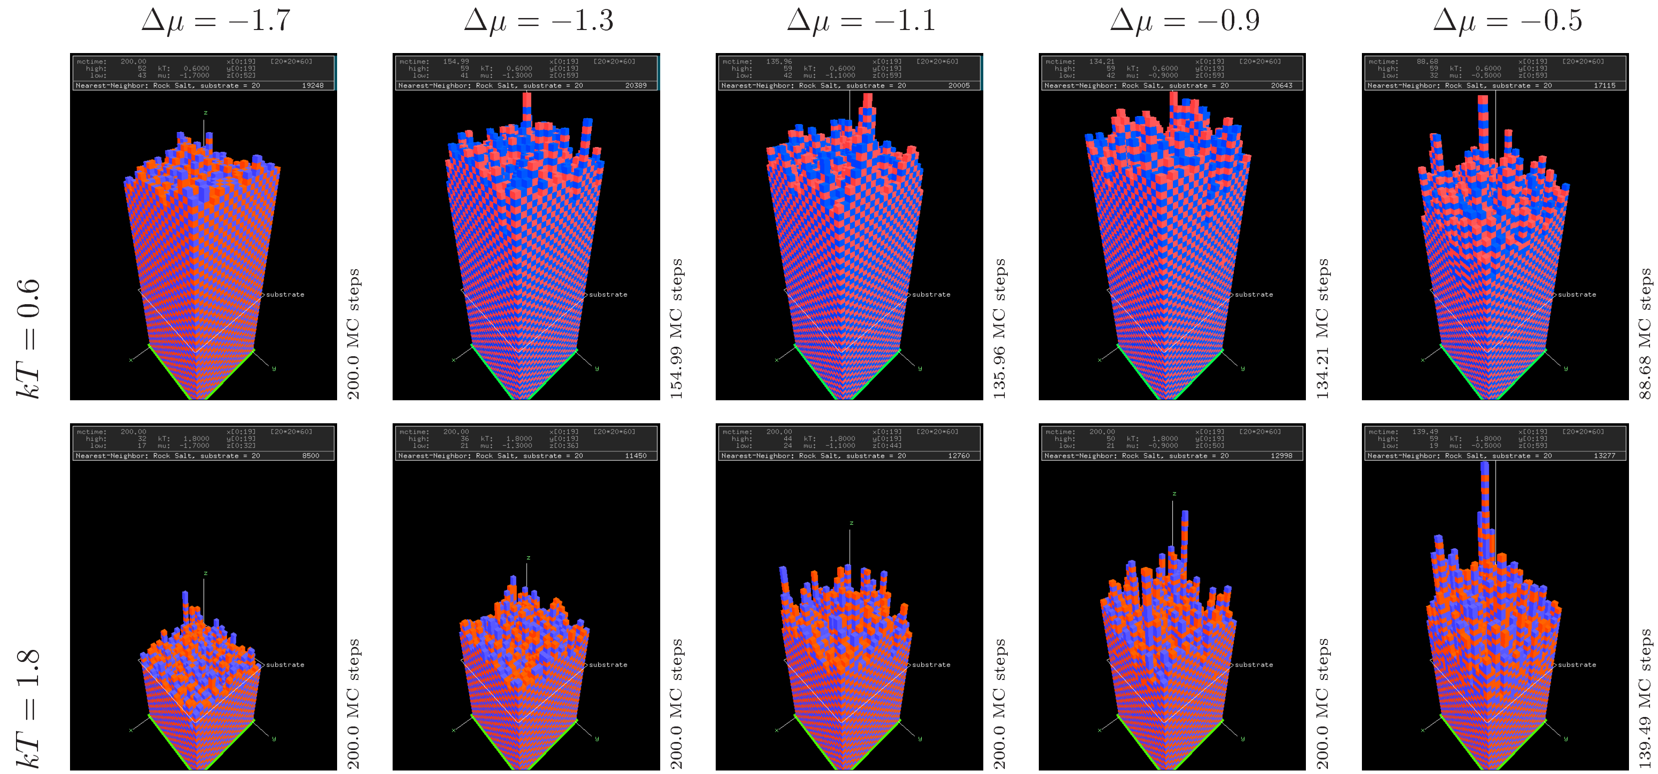

Figure 5.4: **Nearest-Neighbor** Simulations for Rock Salt (Page 2), $20 \times 20 \times 60$ crystal matrix: from left to right $\Delta\mu$ increases from low to high; temperature increases going downward, from moderately high to high. Notice that for these values of $\Delta\mu$ and kT we see the end of the layer-by-layer phase and then increasingly roughening growth. Particularly for $\Delta\mu = -0.5$, imperfections in the crystal structure are becoming apparent, along with gross roughness and "towers."

layer-by-layer growth to rough growth. As one would expect, this region for the long-range model, confirmed by visual investigations, is centered at the area of zero growth rate. Figure 5.6 shows this zoomed-in region. We see a very similar range of growth rates and not quite as distinctive behavior as in the nearest-neighbor case, but there are more issues to consider. Figures 5.7 and 5.8 give a pictoral spread like that given for the nearest-neighbor model, though with a squeezed set of $\Delta\mu$'s. The qualitative behavior is quite promising.

Matrix Size Effects

Because the evaporation rate in the long-range model depends on the energy contribution of all the adatoms in the lattice, it is easy to assume that the lattice size of the crystal matrix will affect it. This assumption would be correct, sort of. The electrostatic interaction range is really not infinite. Screening effects and mere distance will eventually cause the Ewald energy contribution of an adatom (with all its images) some finite distance away in the crystal matrix to be negligible. The Ewald sum accomplishes this by subtracting the uniform background (see Section 3.3). Thus, the contribution of all adatoms will not change with changes in very big matrix size (a result we would expect from the charge neutrality of the system). Although the Ewald method attempts to scale this behavior down for reasonably sized matrices, a finite size effect still becomes apparent. So, the assumption would be correct for the matrix sizes we have dealt with but incorrect for very big matrix sizes.

A direct consequence of this realization is that the growth rates calculated for the long-range system may not be completely trustworthy. As the crystal grows, the number of adatoms increases, and finite size effects come into play. Although for our range of growth this is only slight, it begged mentioning. This finite size scaling effect, however, can be measured versus increasing lattice sizes and eventually extrapolated to matrices of infinite size.

Long-Range Explorations

Here we will detail some representative long-range simulations which span Figures 5.9 through 5.14. Figure 5.9 shows an example of negative growth rate, the crystal is actually being evaporated away before our eyes. The resulting “boxy” formations are quite intriguing in that they have not been observed in the nearest-neighbor model. It is unclear, however, whether they represent finite size effects or physical phenomena. Ponder this though, the formation of miniature crystals in a deteriorating crystal is a good way to lower the energy of the system, at least temporarily. Only investigations for larger lattice sizes will prove it.

Figures 5.10 and 5.12 show layer-by-layer and rough growth respectively. Intertwined with them are simulations with the same parameters but run for longer times and more growth. The purpose is to demonstrate that the growth is robust against the time, or z-direction size for these growth lengths.

We conclude our pretty picture foray with an example of really rough growth. Notice the single stack gaps in the crystal lattice. Look closer. The crystal structures on either side of the hole are mismatched. An adatom in one of those crevices would be in a far less energetically favorable position than one on a tower. The excellent ordering of the towers confirms our reasoning.

5.1.3 Comparison

As hinted above, comparing the growth rates of the two models is difficult without increasing the matrix size and extending the values of $\Delta\mu$. Since time would not permit this (resorting to a supercomputer is likely), we will conjecture a bit. Notice how the growth rate curves for high temperature in the long-range model flip at about $\Delta\mu = -0.6$. It is reasonable to guess that this will happen for the nearest-neighbor model as well, at some value of $\Delta\mu$ lower than -2.5. It is also possible that the zero growth rates characteristic of the cold temperature curves in the nearest-neighbor model around $\Delta\mu = -2$ may turn negative as in the long-range model around $\Delta\mu = -1.25$. These are interesting further question that will have to be explored. The behavior of the long-range model, though, appears consistent with both physical reason and the nearest-neighbor model; we claim preliminary success.

5.2 Other Systems

Other crystal systems, whose ordering is far more interesting than that of rock salt (and also more relevant to the study of piezoelectric materials), are just as easy to investigate in our crystal growth simulation. In addition, they provide a more dramatic contrast between the long-range and nearest-neighbor models, illustrating

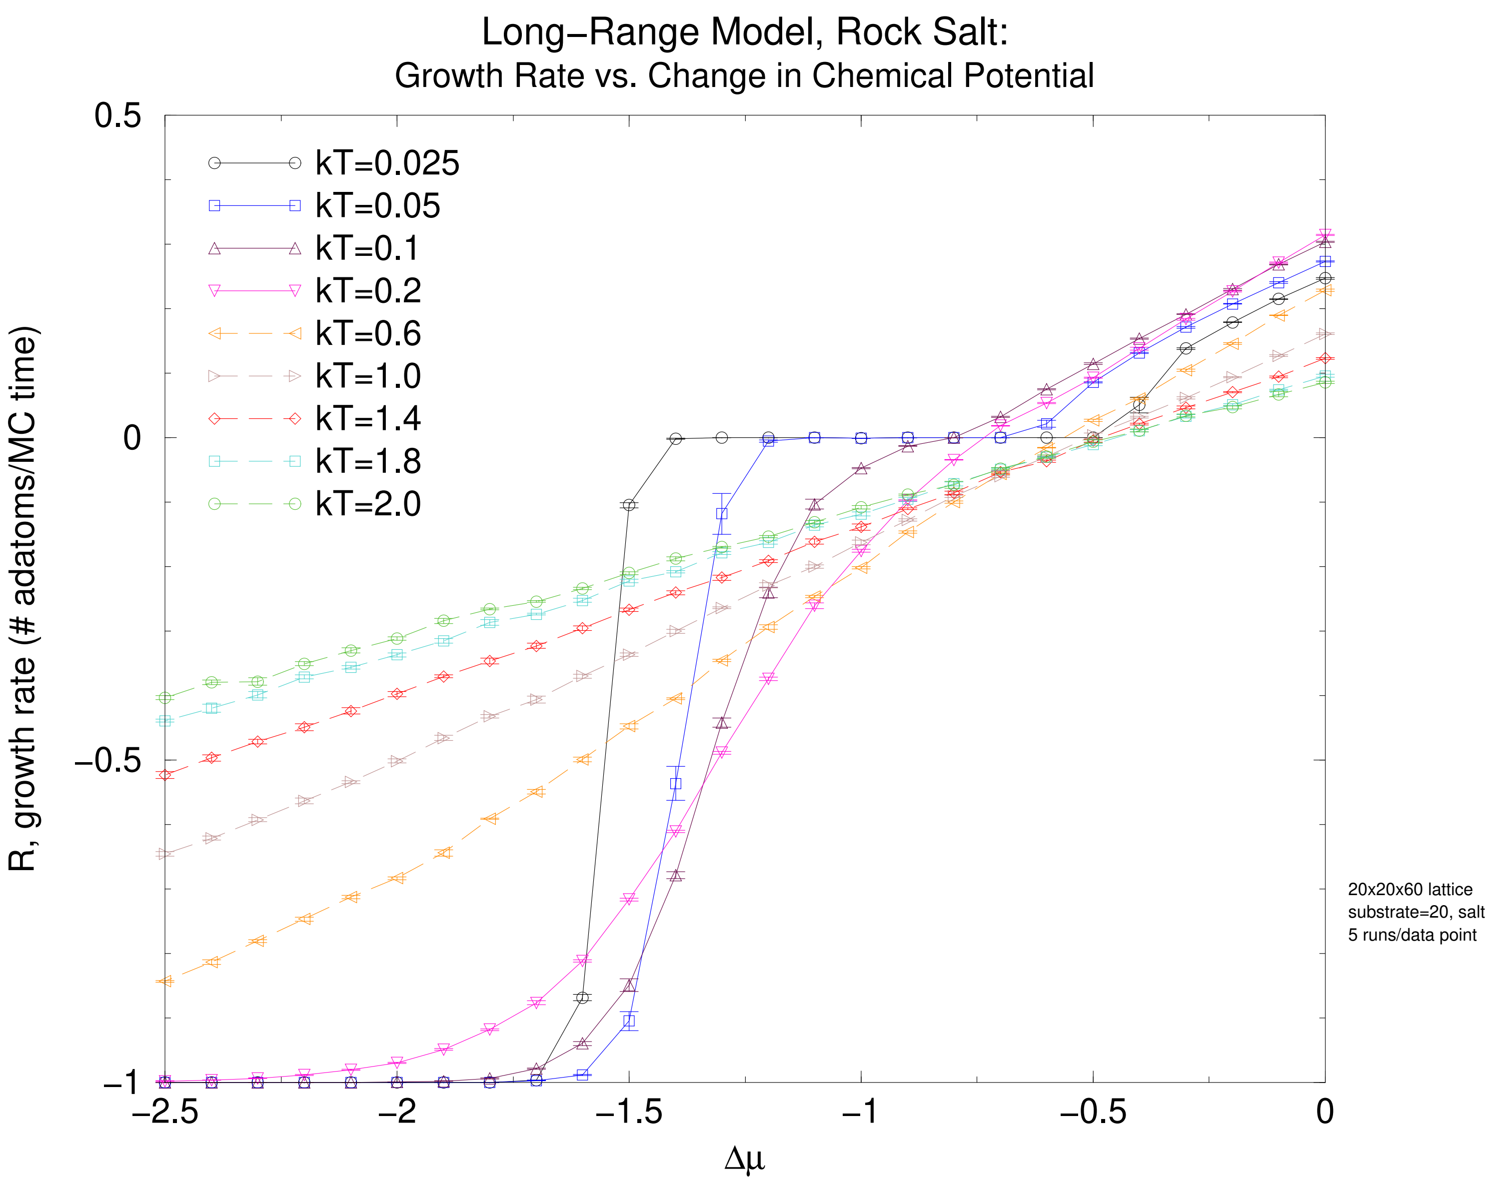

Figure 5.5: A range of long-range model growth rates for rock salt.

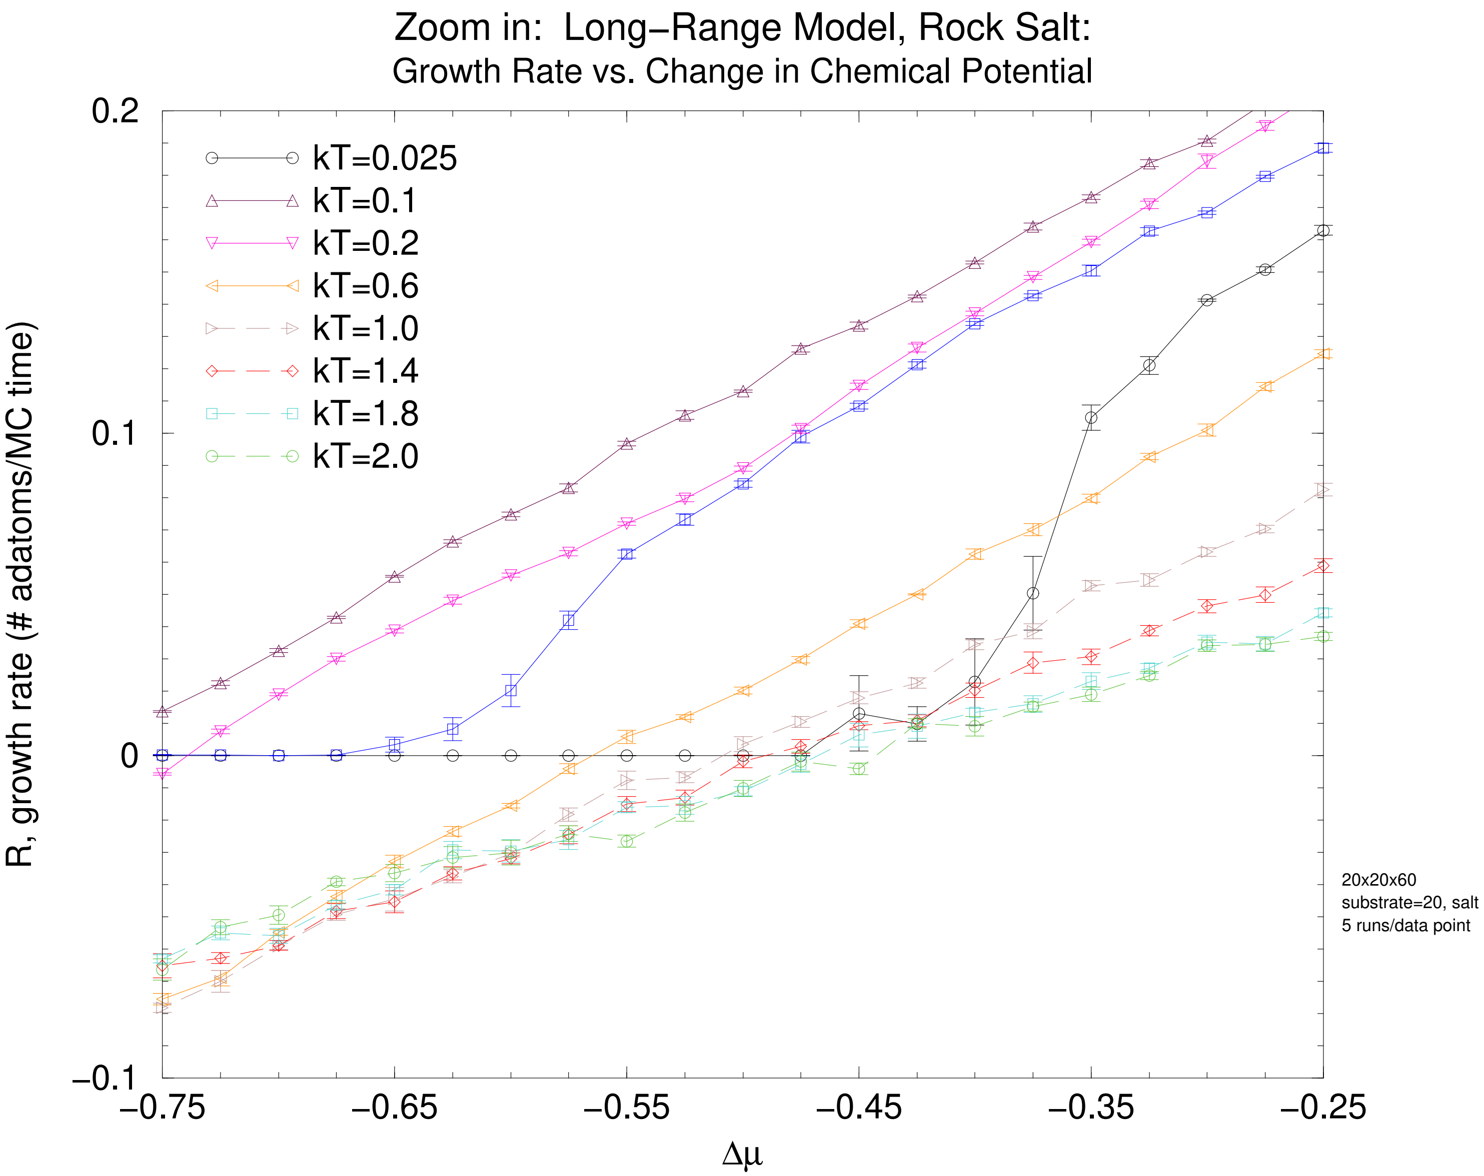

Figure 5.6: Zoom in to the narrow growth mode transition region of the long-range system.

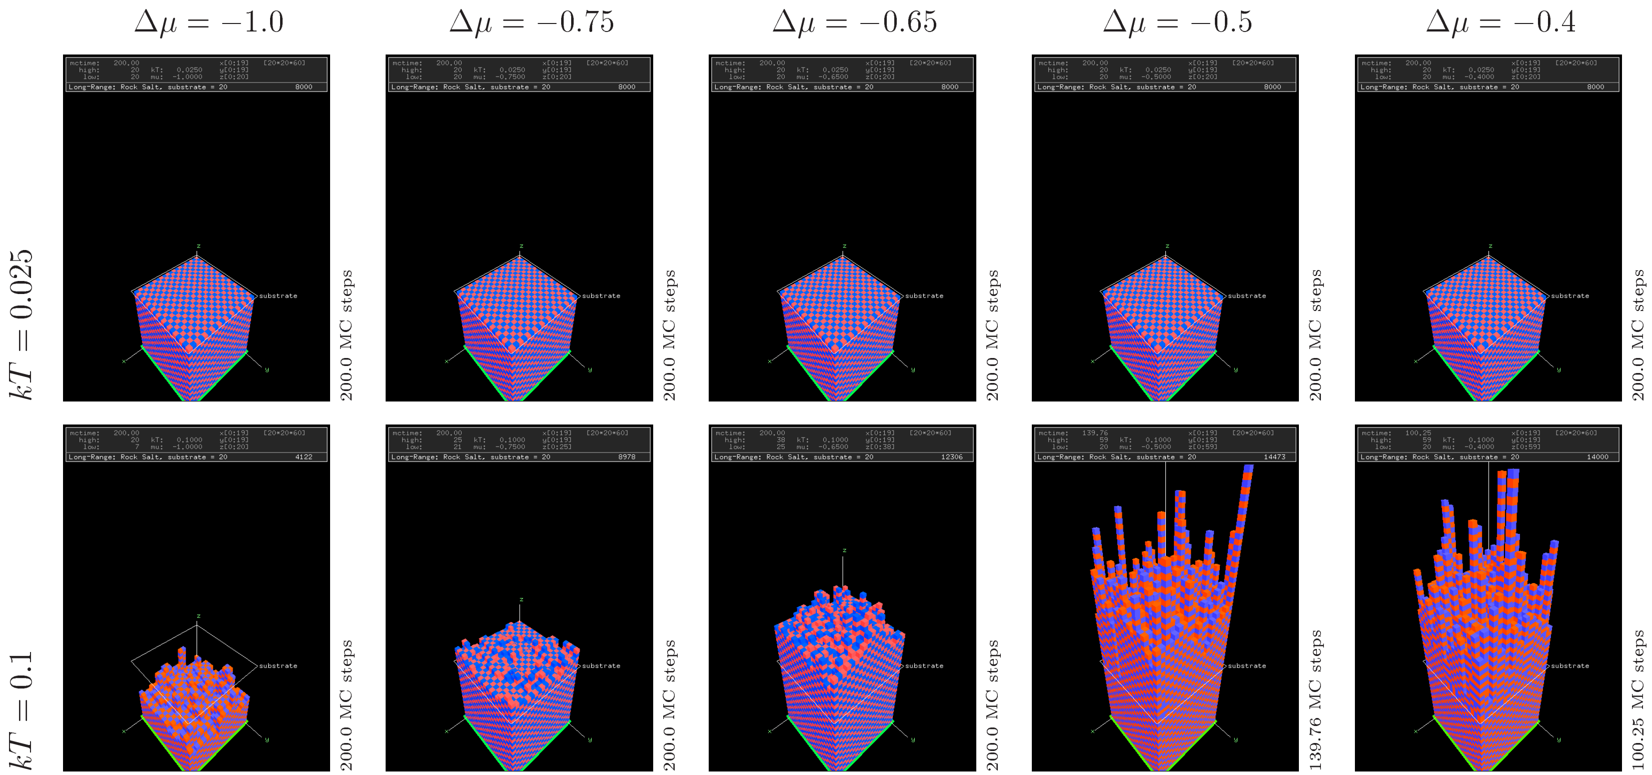

Figure 5.7: **Long-Range** Simulations for Rock Salt, $20 \times 20 \times 60$ crystal matrix: $\Delta\mu$ increases from left to right; temperature increases going downward. We are looking at the small transition region in Figure 5.6 where growth goes from negative to rough. The transition region, as explained in the text, is much narrower for the long-range model. Confirming our growth graphs, for $kT = 0.025$, there is absolutely no growth in this region. Again, do not let the size of the crystal fool you, look at the number of MC steps and Figure 5.5 to realistically judge the growth rate. (Remember the simulations stops when the bounds are reached.)

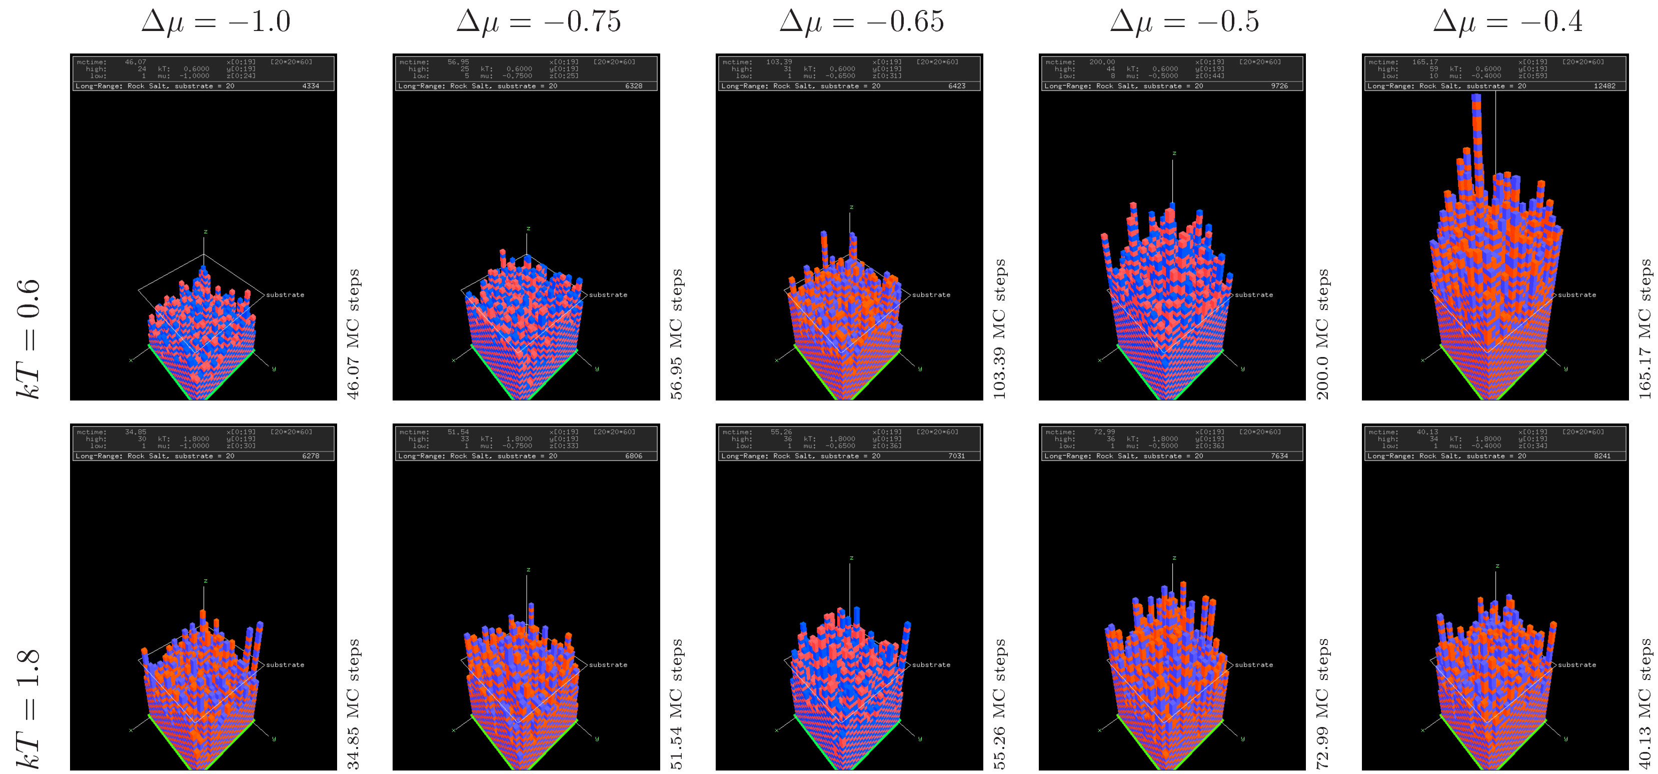

Figure 5.8: **Long-Range** Simulations for Rock Salt, $20 \times 20 \times 60$ crystal matrix: from top to bottom temperature increases from medium-high to very high; from left to right $\Delta\mu$ increases as on the previous page. These temperatures, however, don't permit layer-by-layer growth. Growth is continually very rough through the spectrum of μ.

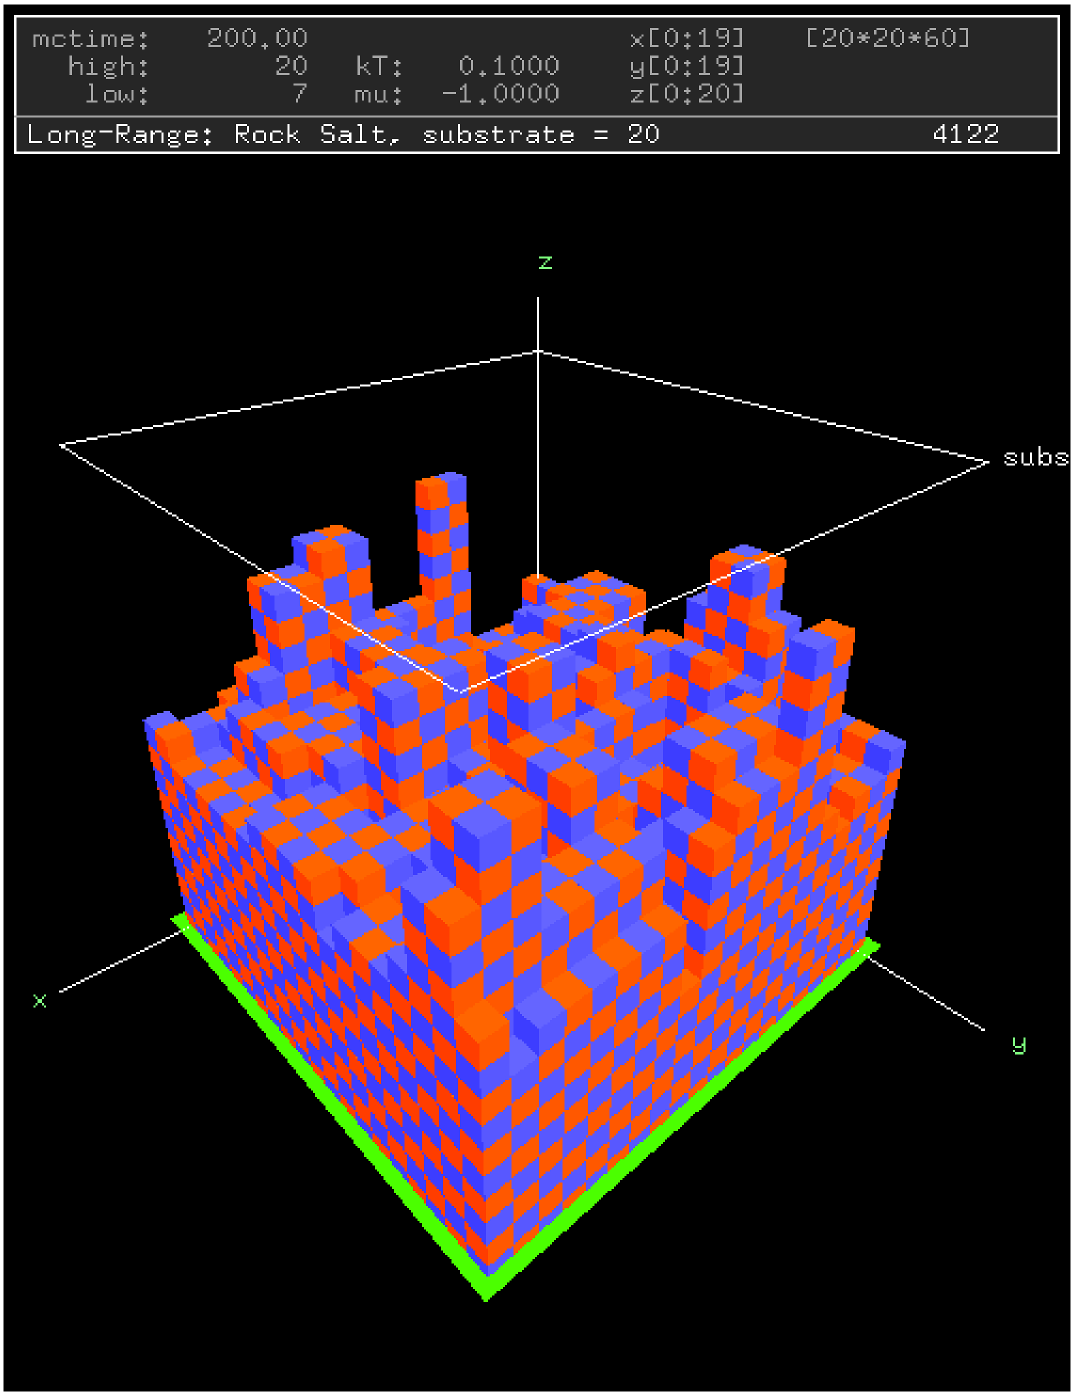

Figure 5.9: Long-Range Simulation for Rock Salt, $20 \times 20 \times 60$: negative growth at medium-low temperature ($kT = 0.1$) and low $\Delta\mu$ (-1.0), notice how the crystal gets eaten away at this temperature and chemical potential. Instead of a rough surface, it becomes "boxy"; adatoms form miniature solid crystals. *This behavior is not observed in the nearest-neighbor model.*

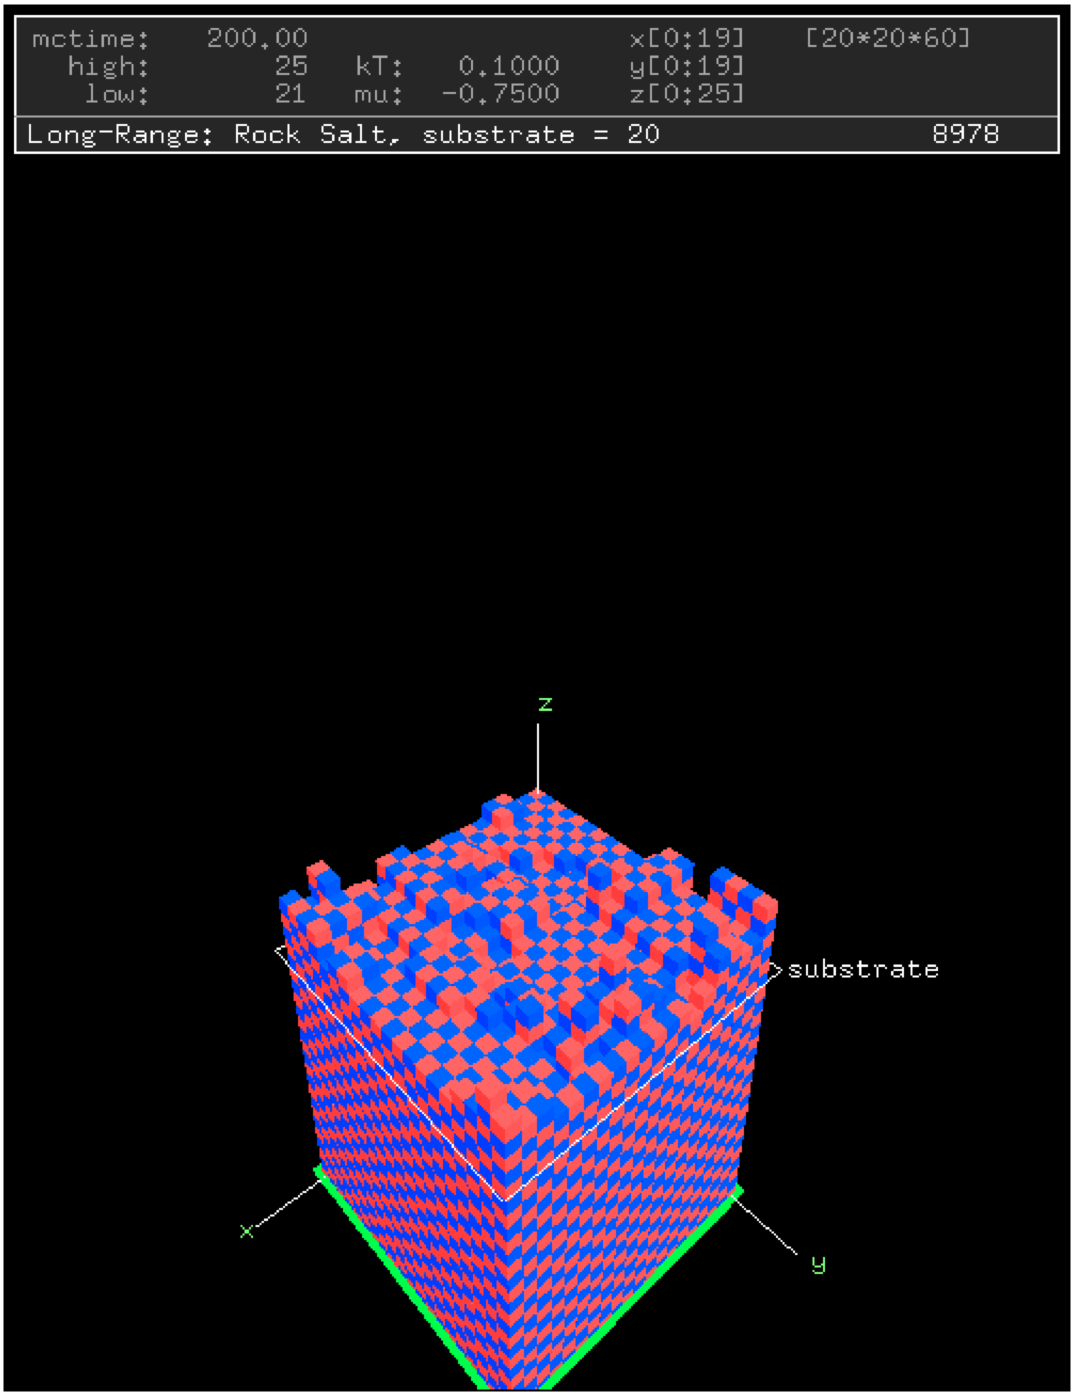

Figure 5.10: Long-Range Simulation for Rock Salt: layer-by-layer growth at a slightly higher μ ($kT = 0.1$ and $\Delta\mu = -0.75$).

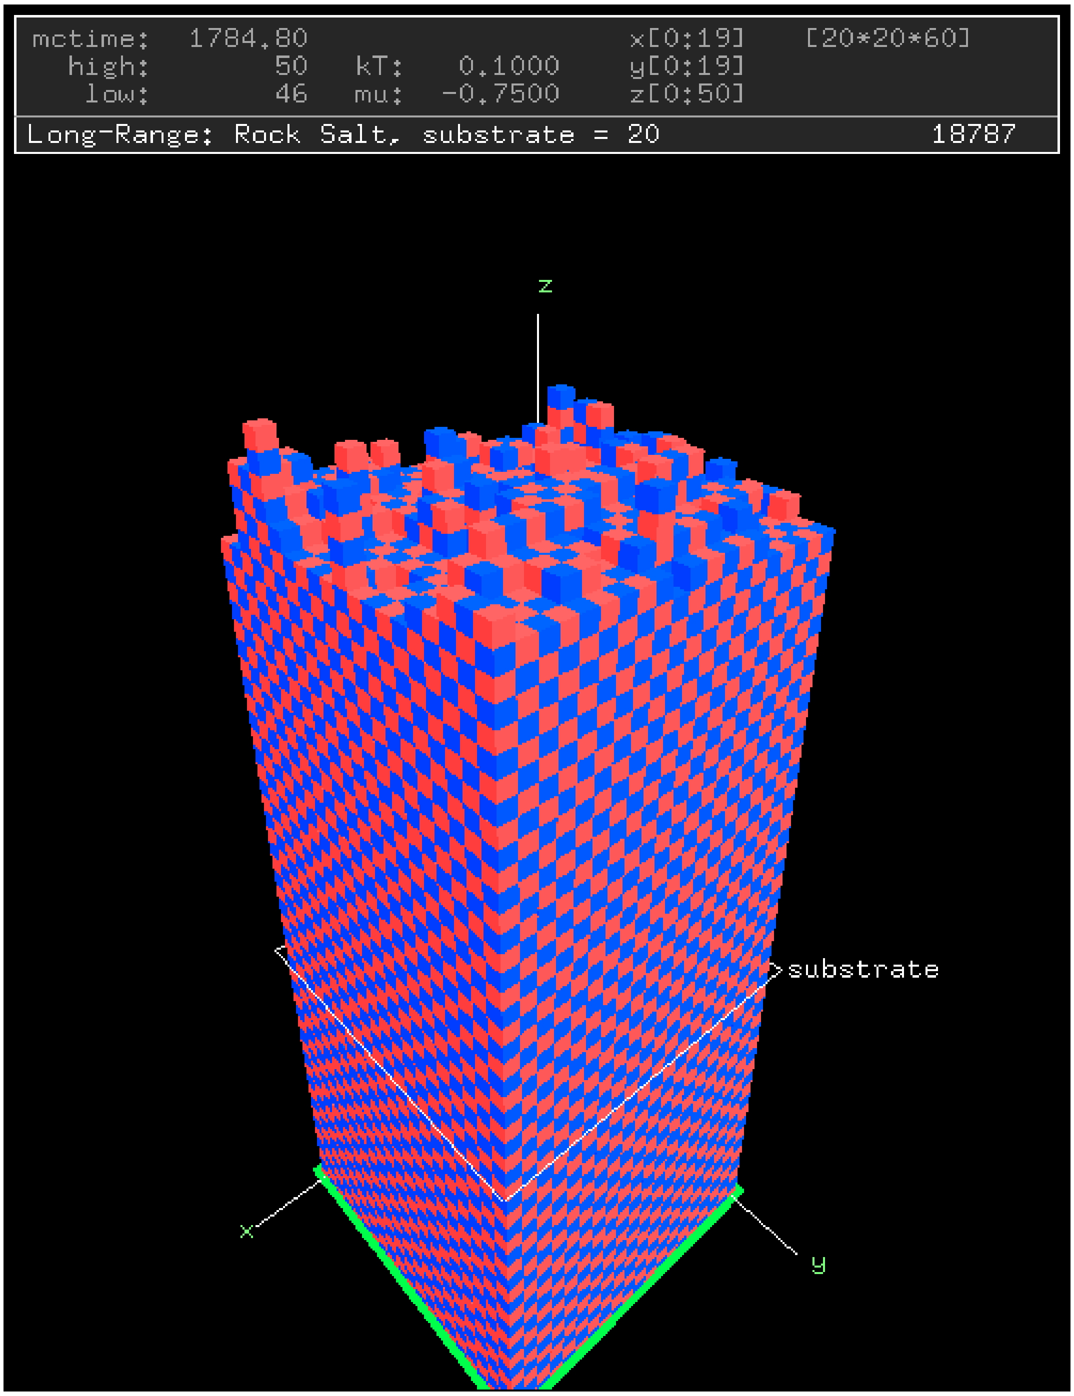

Figure 5.11: Long-Range Simulation for Rock Salt: layer-by-layer growth at a slightly higher μ ($kT = 0.1$ and $\Delta\mu = -0.75$). Simply a longer run version of the previous simulation.

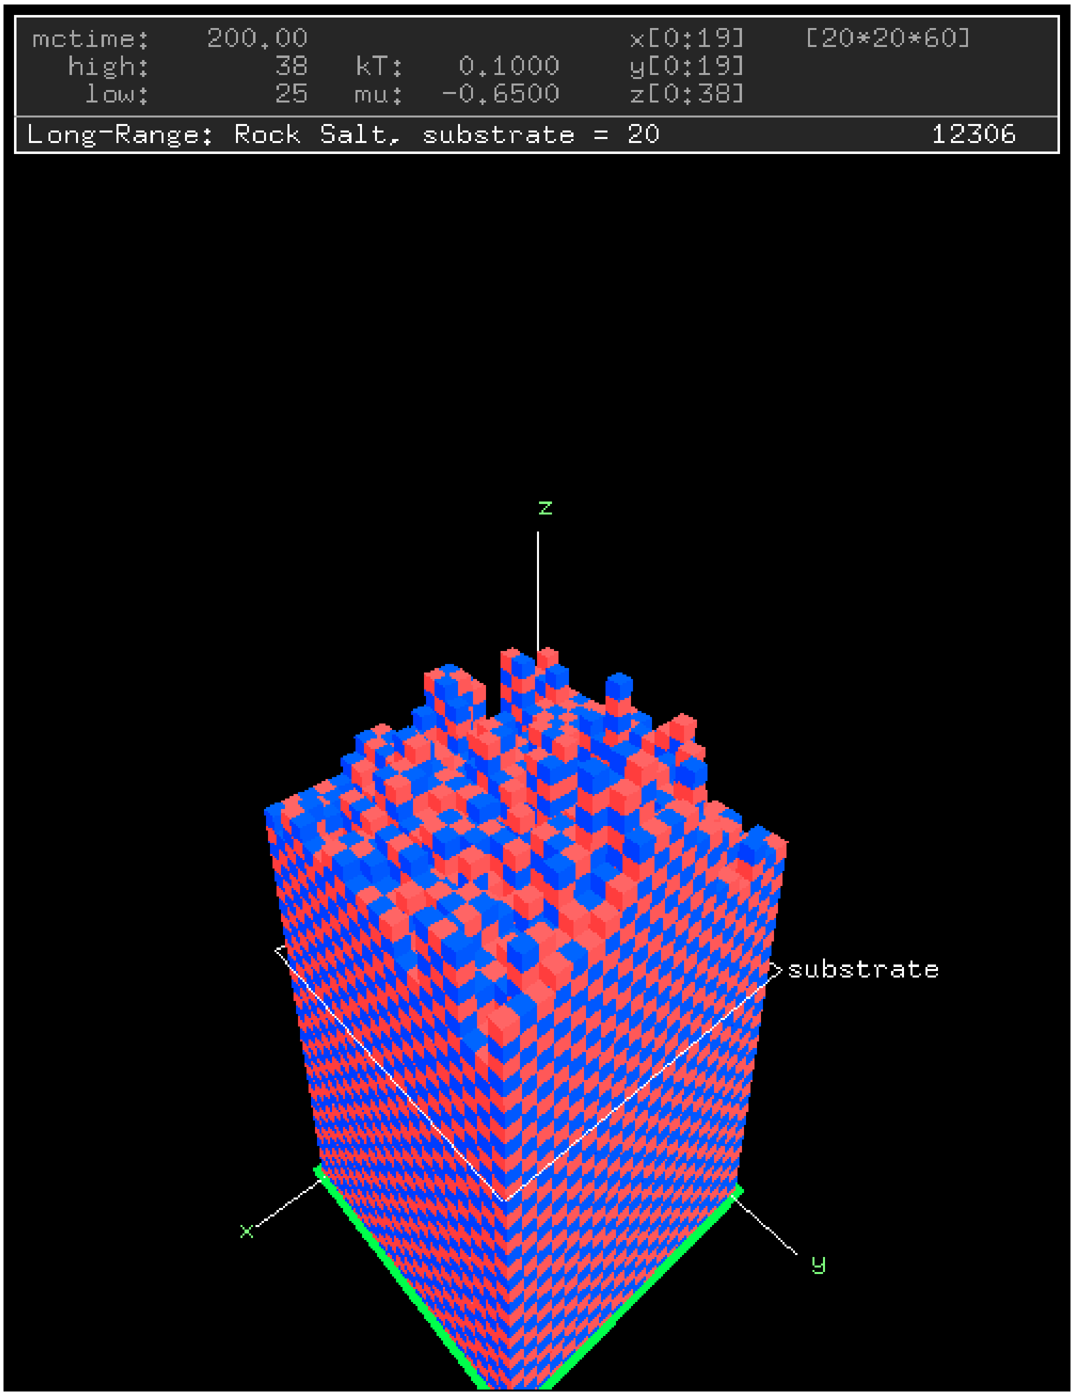

Figure 5.12: Long-Range Simulation for Rock Salt: the transition between layer-by-layer and rough growth ($kT = 0.1$ and $\Delta\mu = -0.65$).

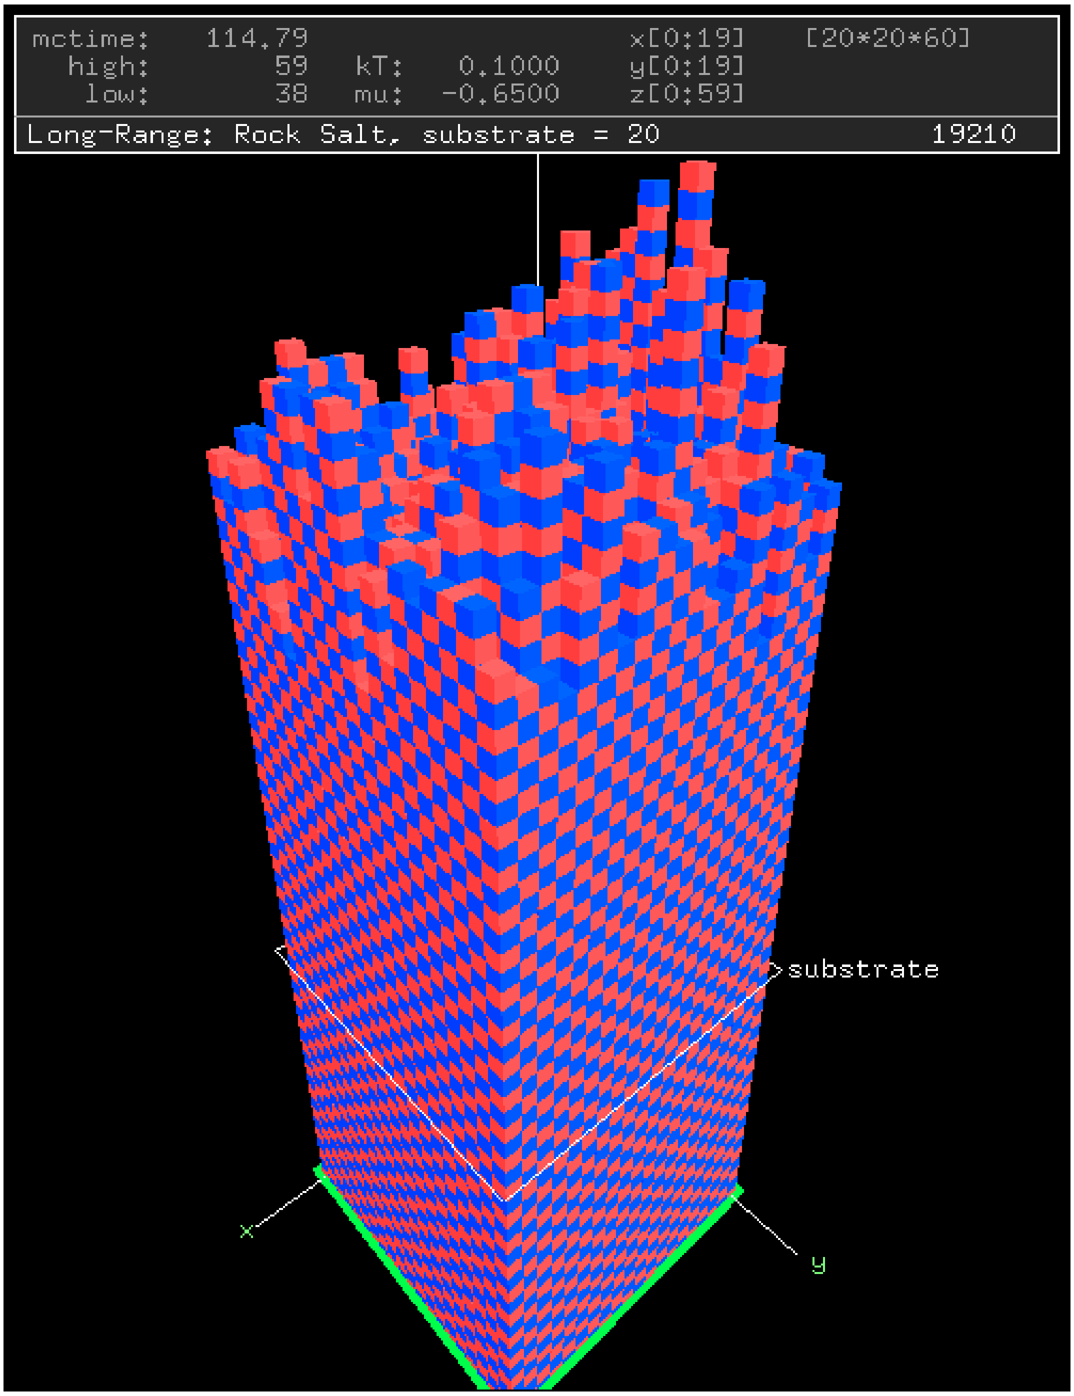

Figure 5.13: Long-Range Simulation for Rock Salt: rough growth ($kT = 0.1$ and $\Delta\mu = -0.65$). Simply a longer run version of the previous simulation.

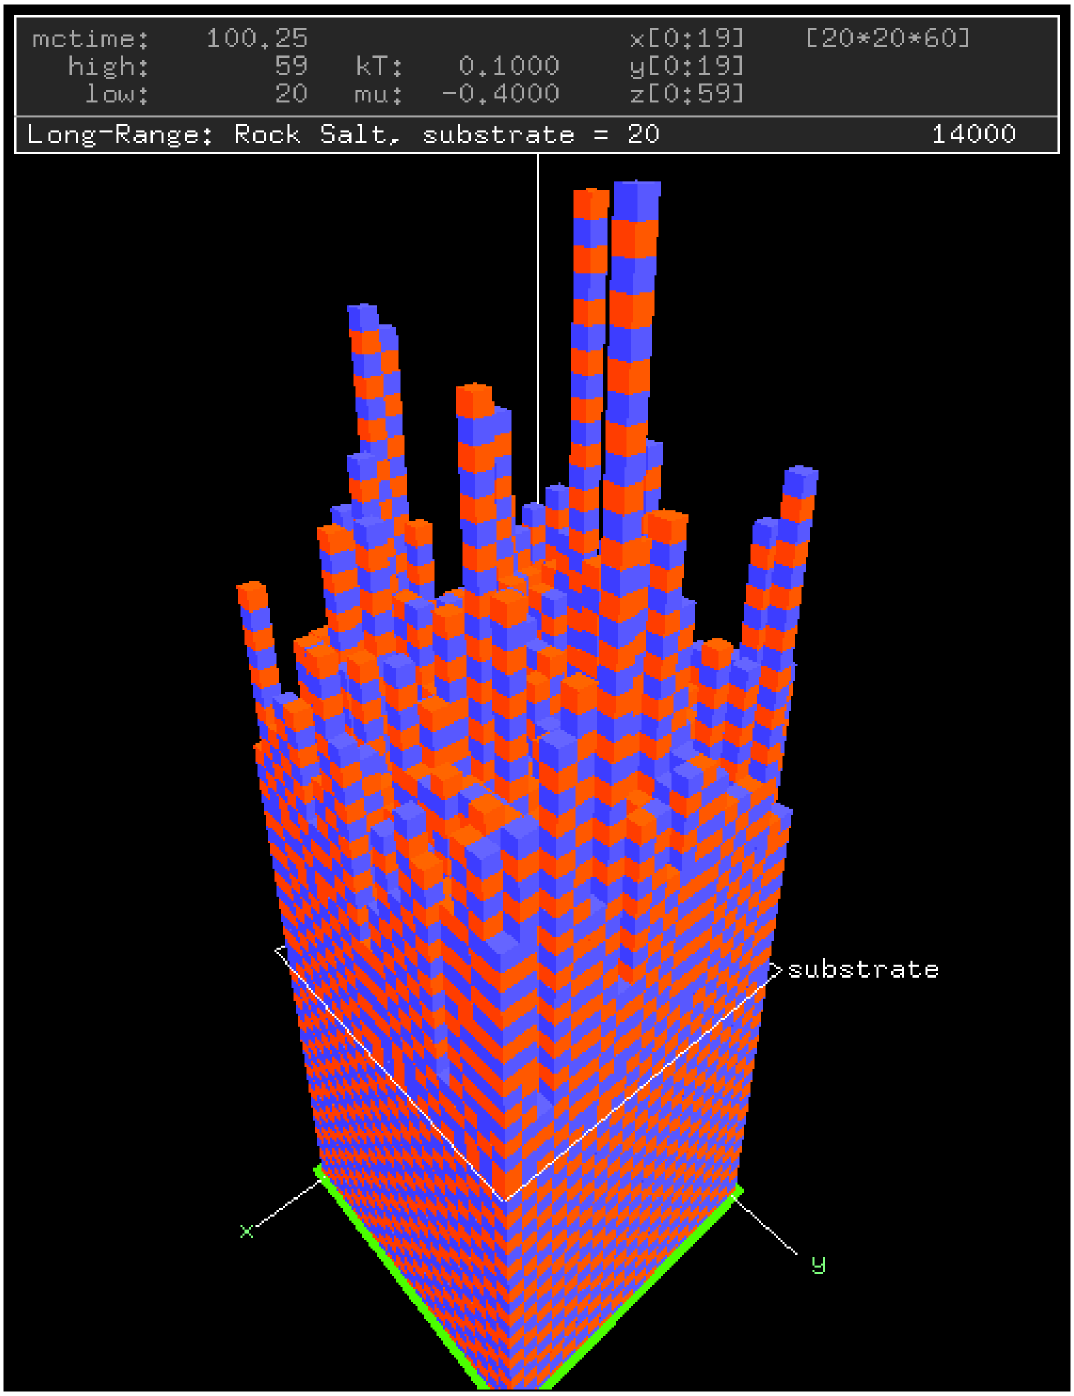

Figure 5.14: Long-Range Simulation for Rock Salt: very rough growth at low temperature and high μ ($kT = 0.1$ and $\Delta\mu = -0.4$). Notice the imperfections and the towers.

very clearly the inadequacies of the latter. Unfortunately, subsumed with verifying both the nearest-neighbor and the long-range models in the rock salt case, we have very little time to explore these exciting new applications of the enhanced algorithm and simulation. Instead, we give a teasing example of the power of the new simulation and a brief outline of the methods we will need for further scientific inquiry.

$II_{1/3}V_{2/3}$ Type Systems

$II_{1/3}V_{2/3}$ materials (described in Chapter 2) are named heterovalent binaries (which include BMN, BZT, ...) and display very interesting ordering behavior. The ground

(a) (b)

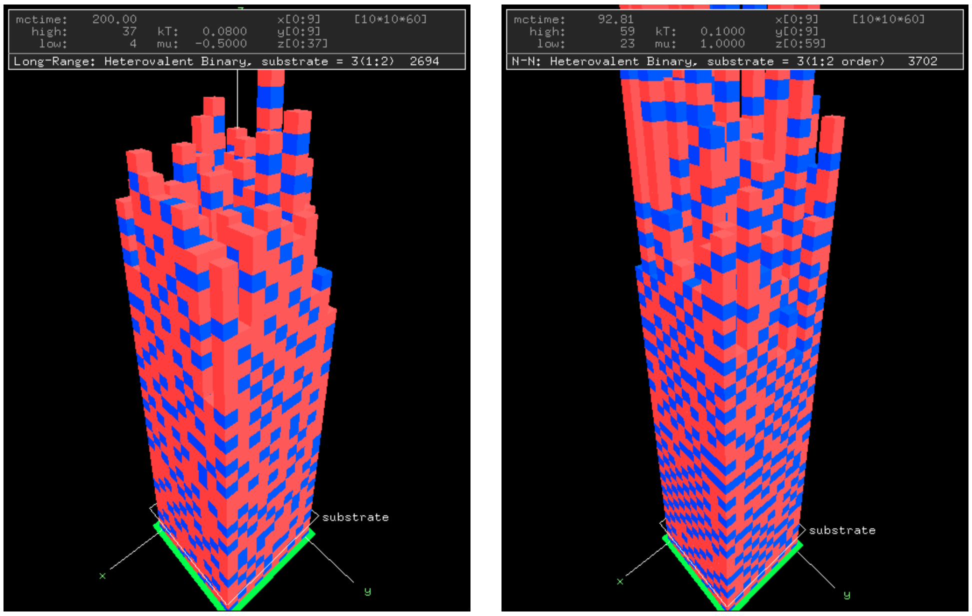

Figure 5.15: Long-Range and N-N Simulation for the Heterovalent Binary $II_{1/3}V_{2/3}$, $10 \times 10 \times 60$: notice that, although both systems show disorder, the long-range model has produced a charge neutral crystal whereas the nearest-neighbor model has not.

state crystal ordering of theses materials is predicted (and confirmed by equilibrium MC [3]) to be $[111]_{1:2}$. In our B-site reduced model, $\Delta q_1 = -2$ and $\Delta q_2 = +1$, where the subscript on q is the species identifier of this two species system. The quantity of species one is therefore 1/3 and that of species 2 is 2/3, maintaining the charge neutrality of the B sublattice. Because of time constraints, intensive investigation was not possible. It is interesting, however, to compare the simulation of this system using the nearest-neighbor model to the long-range model. Figure 5.15 illustrates.

Although full blown 1:2 ordering has not been discovered yet, the long-range model shows promise in how it simulates this more complex material type. The relative quantities of the two species seem adequately realized (in other words, the system is charge neutral) and pockets of 1:2 order hint at future success. The quantification of the level of this ordering is the subject of the next section.

Identifying Order

A key behavioral element of these more complicated materials is how the LRO changes as the relative quantities of species change (see Tables 2.2 and 2.3). To measure the degree of order we use the Fourier transform of the charge-charge correlation function,

$$\eta(\mathbf{k}) = \alpha \sum_{ll'} \Delta q_l \Delta q_{l+l'} \exp(-i\mathbf{k} \cdot l'). \tag{5.1}$$

Here α is a normalization factor, the sum runs over the charges in the B sublattice, and $\mathbf{k}$ is the wave vector of the unit cubic cell [3]. The larger the value of the order parameter, $\eta(\mathbf{k})$, the greater the LRO for that wave vector. For $[111]_{1:1}$ order,

$\mathbf{k}= 2\pi(\frac{1}{2},\frac{1}{2},\frac{1}{2})$; for $[111]_{1:2}$ order, $\mathbf{k}= 2\pi(\frac{1}{3},\frac{1}{3},\frac{1}{3})$. The quantitative exploration of ordering behavior, versus both temperature and species variations, will be an important aspect of any further research into the growth processes of these materials using this simulation.

5.3 Bigger, Better, and More of It

The immediate next step is to run more simulations, of bigger size, and across a greater range of parameters. In doing so, hopefully we can answer some of the questions posed in the previous paragraphs. The exploration of the new model has proved exciting and as we map the new physics it brings to KMC simulations, we can start linking our newfound knowledge to the growth processes themselves.

It is important to mention that although the strictly ionic model of BV has shown great success for describing the crystal ordering in equilibrium MC simulations, it may leave out physics relevant to growth simulations. This is a caution we must face as we push further the use of our simulation as a research tool. Early results are very good, however, and we look forward to discovering new phenomena the addition of long-range interactions brings to crystal growth simulations.

Chapter 6

Conclusions

The development, implementation, and realization of a kinetic Monte Carlo simulation with the addition of long-range ion-ion interactions appears to be a success. Detailed investigations (Chapter 5) into the growth rates of the long-range model for rock salt, as compared to the short-range or nearest-neighbor case, confirm the existence of growth behavior and phase transitions consistent with physical reason and prior research elsewhere. Although more simulations need to be run to both confirm and explore the results, the preliminary success of our new model/simulation represents a significant advance over the traditional non-ionic, nearest-neighbor KMC model. It brings us a step forward in simulating the growth processes of the relaxor single crystals that sparked our interest in this research.

Simulations of rock salt in the long-range model confirmed the presence of phase transitions predicted by crystal growth theory and witnessed in the traditional or nearest-neighbor model. We saw, among other things, that the addition of long-range ionic interactions squeezed this transition region to a narrower slice of $\Delta\mu$ than in the nearest-neighbor case. The explanation arose from the charge neutrality of the crystal system which causes, in the long-range case, potentially evaporated adatoms to feel similar electrostatic potentials across the surface, more so than in the highly discretized nearest-neighbor case. We also witnessed the results of finite size effects across differing lattice sizes and growth distances. This indicates the need for bigger crystal matrix sizes, which will eventually drive its influence to zero.

With positive results in hand from our limited investigations into the crystal growth and ordering behavior of rock salt (and thus heterovalent binary perovskite materials of similar construction), the next step is to expand our investigations for larger matrix sizes, more parameter values, and other types of perovskite crystals. Progressively larger crystal matrices will allow us to overcome the finite size effects pointed out in the last chapter for the long-range model. More simulations over a wider distribution of temperature and adsorption driving-rate values ($\Delta\mu$) will give us more information on the growth behavior of our simulated, ionic crystals. Finally, the move to other crystal types, which include the Relaxor-PT ferroelectric crystals of special interest to us, will represent a new phase of our research.

The investigation of the growth and ordering processes of these PMN-type systems initiates the research objective, outlined previously, to help guide in the growth of these crystals. Questions arise. Is the growth realistic? Does it match experiment? How do we relate—renormalize—our parameters to match the real ones? How do we relate simulation time to real time? These are all important questions (some of which will be addressed in Chapter 7) and follow from the more general question of whether the behavior exhibited in our simulation matches that of experiment. If it doesn't, then the model is incomplete and more physics needs to be added. If it does, than what can we learn?

In conclusion, we have developed, implemented, and realized a simple, ionic, long-range model and simulation built upon an enhancement of the standard KMC formalism. In the process we developed a new algorithm to overcome the difficulties the addition of long-range interactions introduces to the traditional KMC algorithm. Focusing on the primary processes of crystal growth (evaporation, adsorption, growth from the melt, etc.) and those governing the ionic crystal ordering (the BV model), the simulation has succeeded in displaying the expected growth and ordering behavior for the crystal types we explored. Only more testing will completely verify its validity.

Chapter 7

Deficiencies and Future Work

Any model, when compared to the complex real world system it attempts to mimic, will always come up short. Although there is a tendency to add more features to a simulation newly completed, any model—no matter the added complexity—is still a gross simplification. Indeed, the addition of more features, and the additional parameters and complexity that come with them, can serve as more of a negative influence rather than a positive one. It will become increasingly impossible to determine what the physically relavent and dominating processes are as related to the actual physical system. However, when confronted with a model that doesn't capture the physics of the system to an adequate level (self-chosen), the addition of new features/more physics becomes necessary. These are the issues we must weigh. What follows is a list of possible simulation augmentations, paths of future investigation, and some of the implementation issues that correspond.

7.0.1 Possible Shortcomings and Solutions

Diffusion

Numerous times throughout this thesis we have stated that diffusion is not an important process in crystal growth from th melt, our reason for excluding it from the simulation. Well, it might very well be. Although there is support for our claim [29], the only way to really find out is to implement diffusion in our simulation. A next-nearest-neighbor implementation of surface diffusion, a not too difficult to implement first step, is a likely addition we will make to our simulation.

The Solid-on-Solid Restriction

The solid-on-solid (SOS) restriction restricts the adsorption of adatoms to locations on the crystal surface where an adatom is located directly below it. This restriction implies that there can be no vacancies or defects within the bulk of the crystal lattice. The presence of these defects can be very important to the quality of the crystal grown and a study of defects versus growth rate would be very interesting. Removing the SOS restriction, however, is a very difficult proposition algorithmically.

Currently, with SOS, there are a constant number of surface sites (L^2). By turning SOS off, any part of the crystal "open" to the melt could be a possible site for adsorption or evaporation (still maintaining the discrete three-dimensional restriction of course). Additionally, the number of surface sites would not be constant across differing crystal configurations. For even moderately rough growth then, the number of surface sites can be enormous. Maintaining an accurate list of these surface sites and searching through it efficiently with each iteration of the algorithm is a very difficult programming challenge. Because the method of growth relevant to the new relaxor crystals is such that there is little chance for bulk defects (see Section 2.3.1), we will most likely not attempt this challenge.

7.0.2 Enhancements and New Investigations

Controlling Parameters Over Simulation Time

In laboratory crystal growth from the melt—as discussed in Section 2.3.1—temperature is decreased continually throughout the growth process. Although temperature is constant through our crystal growth simulation, changing it over simulation time is certainly possible and easily implemented. Once we have passed the verification phase and proceeded to the investigation of the growth of the new relaxor crystals, this functionality will be a certain addition.

[111] Direction Growth

Currently the crystal in our simulation grows in the [001] direction. This need not be the case however. A simple transformation of the basis vectors that define our crystal lattice (Equation 2.1) would allow for growth in other directions, [111] for example. A comparison of the growth and ordering behaviors in the [111] direction versus those in the [001] direction would be quite interesting. Although a non-trivial programming challenge, it could certainly be achieved. These investigations are a likely path of future work.

Appendix A

Symbols and Variables Listed

(In order of appearance.)

$\mathbf{a}, \mathbf{b}, \mathbf{c} \equiv$ crystallographic axes vectors

$l \equiv$ perovskite unit cell in the crystal lattice, sometimes reduced to just the B site

$\tau = \{A, B, O_1, O_2, O_3\} \equiv$ specific atom in cell l

$\mathbf{R}_{l\tau} \equiv$ the position in the ideal cubic structure of the atom on site τ of cell l

$\mathbf{l} \equiv$ the vector to site l in the lattice

$a \equiv$ the lattice spacing constant: the distance between lattice sites

$Q_{l\tau} \equiv$ the charge of the atom on site τ of cell l

$\Delta q_l = Q_{l,B} - q_B \equiv$ the relativized charge of B-site atom

$E_B \equiv$ the electrostatic energy of the B sublattice

$\mathcal{H} \equiv$ the Hamiltonian defining the total energy of the system

$(i, j, k) \equiv$ the position of an adatom in the 3-d crystal matrix

$\mathcal{V}_{ll'} \equiv$ the Coulomb potential felt by one adatom in the crystal lattice

$\Delta\mu = \mu_s - \mu_l \equiv$ the chemical potential difference between the solid and liquid phase

$h_{ij} \equiv$ the height of a (i, j) stack in the crystal

$\mathcal{C} \equiv$ a particular crystal configuration

$E(\mathcal{C}) \equiv$ the total electrostatic energy of a particular crystal configuration

$v_{lr} \equiv$ the long-range contribution to the Coulomb sum

$P(\mathcal{C}, t) \equiv$ the time-dependent distribution of configurations

$w(\mathcal{C} \rightarrow \mathcal{C}') \equiv$ the rate of adsorption

$w(\mathcal{C}' \rightarrow \mathcal{C}) \equiv$ the rate of evaporation

$Z \equiv$ the partition function

$kT \equiv$ the temperature of the system given in units of energy

$L \equiv$ the length and width of the crystal matrix

$H \equiv$ the height of the highest layer of the crystal

$V \equiv$ the electrostatic potential table

$U_{ij} \equiv$ the surface potential array

$P_i \equiv$ the event probability list

Appendix B

Definitions and Slight Explanations

(Although it may not seem that way, there is a method to the following order.)

condensed matter physics: Condensed matter physics is the study of the macroscopic properties of materials. It seeks to use the well-established laws of microscopic physics to predict the collective behavior of very large numbers of electrons, atoms, or molecules.

simulation techniques: *Molecular Dynamics* (MD) and *Monte Carlo* (MC) are by far the two biggest simulation techniques relavent to growth simulations. The term Monte Carlo refers to a class of computational methods whose essence is the invention of games of chance (random numbers) and whose behavior and outcome can be used to study some interesting phenomena. Developed to study the diffusion of neutrons in fissionable material, the Monte Carlo method is based on finding a probabilistic analogue to a deterministic (indirect MC) or stochastic (direct MC) mathematical problem and then solving it by a stochastic sampling experiment, ie simulation. Simulations provide information on model systems which is arbitrarily detailed, and whatever quantity the researcher may consider useful he may attempt to "sample" from the simulation.[35] Molecular Dynamics simulations, in contrast to Monte Carlo, deal with techniques to solve the classical equations of motion of particles in a system. Unlike MC which gathers average information about a system, MD—by integrating Newton's laws numerically—realistically simulates physical systems (at a big performance price). Kinetic Monte Carlo bridges the gap between the two by sampling relaxation or disequilibrium stochastic processes.

adatom: the label for the smallest considered object in a crystal growth simulation be it an atom, molecule, or some larger aggregate.

piezoelectric: Derived from the Greek as "pressure electricity", a piezoelectric is an actuating crystal which lacks a center of symmetry. See Section 2.1.1.

piezoelectric coefficient: The piezoelectric coefficient, usually given the symbol d_{ij}, represents the level of strain at a given electric field. It is the most widely used parameter describing actuator performance.

lattice: A lattice is an infinite, regular array of points in space.

unit cell: A unit cell is a region of space which when repeated by primitive translation vectors fills all space.

crystal: A crystal is a regular and finite array of objects.

crystal lattice: A *lattice* is an infinite, regular array of points in space. When an adatom is placed on these points in a regular way, the result is an infinite crystal lattice.

crystal systems: The seven crystal systems are: triclinic, monoclinic, orthohombic, tetragonal, cubic, trigonal, and hexagonal. All crystalline materials, except the cubic system, are optically anisotropic, meaning they look different depending on the direction of sight.

Miller indices: A crystallographic labeling mechanism, Miller indices are constructed from the reciprocals of the indices of the plane with the crystallographic axes. See Section 2.1.

ceramic: 1) Polycrystalline materials of inorganic oxides and salts. 2) A compound of metallic and nonmetallic elements, for which the interatomic bonding is predominantly ionic. See *polycrystalline materials*.

polycrystalline materials: aggregates of randomly oriented small crystals.

overcooled liquid and amorphous solids: Non-fluid materials with a very high degree of disorder.

perovskite alloys: Common compounds usually containing a mixture of alkaline and transition metals. They are described by the chemical formula ABO_3.

Here A is usually a fixed Group II metal, and B has a fractional compositional freedom containing metals from Groups II-VII. See Section 2.2.

ferroelectric: Ferroelectrics have a permanent electric dipole moment which may be reoriented by the application of an electric field. Ferroelectric behavior is only possible for crystal structures belonging to one of the 10 non-centrosymmetric point groups that possess a crystallographically unique direction.

adsorption: when an atom, molecule, or group of atoms (adatom) in the melt, vapor, or solution attaches itself to the substrate/crystal.

evaporation: when an adatom, due to internal energies within the crystal system like thermal fluctuations or electrostatic potential desorbs or is removed from the solid crystal.

surface migration: See *diffusion.*

diffusion: after an adatom has adsorbed, it tends to travel along the crystal surface until becoming bonded to two or more other adatoms; this is referred to as surface migration or diffusion.

crystal ordering: The positioning of ions in an ordered, repetitious pattern rather than in a random arrangement. Further related terms include long-range order (LRO), short-range order (SRO), and disorder. See Section 2.2.1

substrate: The substrate is the initial surface used to begin the nucleation process of crystal growth. Although usually arranged to match the type of growth, ie $[111]_1 : 1$ for rock salt, this need not be so. Nor does the substrate have to be made of the same constituents are the crystal to be grown.

modes of growth: The categorization of the roughness of the growing crystal surface falls to two basic types: *layer-by-layer growth*: where crystal layers grow one at a time (this includes nucleated growth) ; *three-dimensional growth:* where many layers grow at once forming a rough or hilly surface.

Metropolis Monte Carlo A means of importance sampling, the Metropolis or $M(RT)^2$ method generates a sequence of random states so that by the end of the simulation each state has occurred with the appropriate probability; this is called a *Markov chain*. For systems with an enormous number of possible configurations (such is in the case of crystal growth, see Section 2.1), it becomes impossible to adequately sample the huge state space—original MC breaks down. The Metropolis method offers a way out: the ability to importance sample as you go. In most situations today, when people refer to Monte Carlo, they are actually referring to Metropolis Monte Carlo.

Equilibrium Monte Carlo: Standard thermodynamical MC simulation is a method for the calculation of average values in a given equilibrium thermodynamical ensemble.[10] Equilibrium MC simulations of crystal systems seek only to find the crystal ordering resulting in the lowest energy state for a given temperature. They do not realistically simulate the physical processes that get the system to that order. . .hence the need for kinetic MC.

Coulomb force: A central mass force between two charged particles, the Coulomb or electrostatic force is defined by $F = kq_1q_2/r$ where k is a constant, q_1 and q_2 are the charges of particle 1 and 2 respectively, and r is the distance separating the two particles.

species: In our crystal growth terminology, species of adatoms are grouped by like charge. Thus, for a rock salt type system, there would be two species of charge $1+$ and $1-$.

Ewald summation: The Ewald summation is a technique for summing the interaction between an ion and all its periodic images in a crystal lattice. See Section 3.3.

deterministic processes: A process such as the outcome can be known from the given initial conditions (e.g. projectile motion).

stochastic processes: A process such that the outcome is random in nature (e.g. adsorption in crystal growth).

Ising model: The first computational model to successfully reproduce a system exhibiting a first order phase transition, the Ising model attempts to simulate the ferromagnetic properties of iron. Using a discrete lattice of spin magnetic moments with periodic boundary conditions, it defines an interaction energy between nearest-neighbor spins such that a more favorable energy state is reached by parallel alignment. The system evolution is realized, according to the temperature, through a Metropolis Monte Carlo simulation. The Ising model is the basis for many other computational models, including KMC.

Markov chain: A sequence of trials that satisfy: 1) the outcome of each trial belongs to a finite set of outcomes called the state space, 2) the outcome of each trial depends on the outcome of the trial that immediately precedes it.

state space: The set of all possible outcomes of an experiment.

crystal imperfections: *physical:* strain, dislocations, grain boundaries, twin planes, and stacking faults. *chemical:* the presence of a foreign atom or a vacancy in the crystal.

Appendix C

Hardware and Software

C.1 Equipment and Compilers

All simulations were done on either Pentium-class computers under Linux 2.2.5-22 or a PowerPC G3 running Linux 2.2.10. On the Pentiums, gcc version 2.95.2 19991024 was the compiler used, gcc version egcs-2.91.66 19990314 (egcs-1.1.2 release) on the PowerPC. Since standard data types (ie "double"s) and gnu compilers were used throughout, in addition to the huge number of runs, significant numerical error across machines is unlikely.

Visualization was done on the PowerPC, its 128-bit graphics card allowing for reasonable speed. The *OpenGL* copy-cat library *Mesa 3.1* was the basis for the graphics routines.

C.2 Random Number Generator

The random number generator used, the quality of which is very important for MC simulations, was designed and implemented in C++ by Park and Miller.[14] According to [14]'s documentation, the period of one of their random number generating "streams" is $2,147,482,647-1$. To see if this number is big enough, we must calculate the number of random numbers used in a given simulation to make sure that repetition is not possible. For a 20×20 crystal matrix simulation ran for 10 times for 200 MC steps each time, the number of random numbers used is approximately $n_r = 20^2 \times 10 \times 200 \times 2 = 1,600,000$ where the extra 2 comes from randomly choosing the species to adsorb at each time step (an overestimation). We appear to be OK for these matrix sizes.

Appendix D

An Ewald Table for Finite Height

Courtesy T.J. Walls, below is a sample Ewald table for a crystal matrix of size ten cubed. Because of periodic boundary conditions, we only need potential data for vectors half the size of the matrix in any direction. The table gives the +1–+1 Coulomb potential with images (as described in Section 3.3) between a site at $(0, 0, 0)$ and a site (i, j, k) away. Again, please refer to the work of [31] for more detailed information.

D.1 Ewald Table for a $10 \times 10 \times 10$ Matrix

(0,0,0)	∞	(0,1,0)	1.002284	(0,2,0)	0.509453	(0,3,0)	0.355849	(0,4,0)	0.293428	(0,5,0)	0.275790
(1,0,0)	1.002284	(1,1,0)	0.711578	(1,2,0)	0.458557	(1,3,0)	0.340103	(1,4,0)	0.286432	(1,5,0)	0.270734
(2,0,0)	0.509453	(2,1,0)	0.458557	(2,2,0)	0.370879	(2,3,0)	0.305674	(2,4,0)	0.269502	(2,5,0)	0.258086
(3,0,0)	0.355849	(3,1,0)	0.340103	(3,2,0)	0.305674	(3,3,0)	0.272563	(3,4,0)	0.250894	(3,5,0)	0.243522
(4,0,0)	0.293428	(4,1,0)	0.286432	(4,2,0)	0.269502	(4,3,0)	0.250894	(4,4,0)	0.237369	(4,5,0)	0.232511
(5,0,0)	0.275790	(5,1,0)	0.270734	(5,2,0)	0.258086	(5,3,0)	0.243522	(5,4,0)	0.232511	(5,5,0)	0.228472
(0,0,1)	1.000000	(0,1,1)	0.709333	(0,2,1)	0.456414	(0,3,1)	0.338085	(0,4,1)	0.284513	(0,5,1)	0.268852
(1,0,1)	0.709333	(1,1,1)	0.581710	(1,2,1)	0.419297	(1,3,1)	0.324711	(1,4,1)	0.278182	(1,5,1)	0.264177
(2,0,1)	0.456414	(2,1,1)	0.419297	(2,2,1)	0.350220	(2,3,1)	0.294835	(2,4,1)	0.262757	(2,5,1)	0.252448
(3,0,1)	0.338085	(3,1,1)	0.324711	(3,2,1)	0.294835	(3,3,1)	0.265331	(3,4,1)	0.245644	(3,5,1)	0.238882
(4,0,1)	0.284513	(4,1,1)	0.278182	(4,2,1)	0.262757	(4,3,1)	0.245644	(4,4,1)	0.233107	(4,5,1)	0.228585
(5,0,1)	0.268852	(5,1,1)	0.264177	(5,2,1)	0.252448	(5,3,1)	0.238882	(5,4,1)	0.228585	(5,5,1)	0.224799
(0,0,2)	0.500000	(0,1,2)	0.449278	(0,2,2)	0.362052	(0,3,2)	0.297397	(0,4,2)	0.261662	(0,5,2)	0.250412
(1,0,2)	0.449278	(1,1,2)	0.412295	(1,2,2)	0.343563	(1,3,2)	0.288598	(1,4,2)	0.256854	(1,5,2)	0.246672
(2,0,2)	0.362052	(2,1,2)	0.343563	(2,2,2)	0.304338	(2,3,2)	0.268030	(2,4,2)	0.244953	(2,5,2)	0.237218
(3,0,2)	0.297397	(3,1,2)	0.288598	(3,2,2)	0.268030	(3,3,2)	0.246487	(3,4,2)	0.231447	(3,5,2)	0.226161
(4,0,2)	0.261662	(4,1,2)	0.256854	(4,2,2)	0.244953	(4,3,2)	0.231447	(4,4,2)	0.221359	(4,5,2)	0.217682
(5,0,2)	0.250412	(5,1,2)	0.246672	(5,2,2)	0.237218	(5,3,2)	0.226161	(5,4,2)	0.217682	(5,5,2)	0.214545
(0,0,3)	0.333333	(0,1,3)	0.318059	(0,2,3)	0.284843	(0,3,3)	0.253198	(0,4,3)	0.232688	(0,5,3)	0.225752
(1,0,3)	0.318059	(1,1,3)	0.305105	(1,2,3)	0.276311	(1,3,3)	0.248114	(1,4,3)	0.229460	(1,5,3)	0.223088
(2,0,3)	0.284843	(2,1,3)	0.276311	(2,2,3)	0.256430	(2,3,3)	0.235717	(2,4,3)	0.221332	(2,5,3)	0.216293
(3,0,3)	0.253198	(3,1,3)	0.248114	(3,2,3)	0.235717	(3,3,3)	0.221954	(3,4,3)	0.211872	(3,5,3)	0.208236
(4,0,3)	0.232688	(4,1,3)	0.229460	(4,2,3)	0.221332	(4,3,3)	0.211872	(4,4,3)	0.204645	(4,5,3)	0.201977
(5,0,3)	0.225752	(5,1,3)	0.223088	(5,2,3)	0.216293	(5,3,3)	0.208236	(5,4,3)	0.201977	(5,5,3)	0.199644
(0,0,4)	0.250000	(0,1,4)	0.244099	(0,2,4)	0.229962	(0,3,4)	0.214673	(0,4,4)	0.203729	(0,5,4)	0.199835
(1,0,4)	0.244099	(1,1,4)	0.238775	(1,2,4)	0.225925	(1,3,4)	0.211874	(1,4,4)	0.201722	(1,5,4)	0.198090
(2,0,4)	0.229962	(2,1,4)	0.225925	(2,2,4)	0.215996	(2,3,4)	0.204845	(2,4,4)	0.196596	(2,5,4)	0.193606
(3,0,4)	0.214673	(3,1,4)	0.211874	(3,2,4)	0.204845	(3,3,4)	0.196700	(3,4,4)	0.190503	(3,5,4)	0.188221
(4,0,4)	0.203729	(4,1,4)	0.201722	(4,2,4)	0.196596	(4,3,4)	0.190503	(4,4,4)	0.185758	(4,5,4)	0.183986
(5,0,4)	0.199835	(5,1,4)	0.198090	(5,2,4)	0.193606	(5,3,4)	0.188221	(5,4,4)	0.183986	(5,5,4)	0.182396
(0,0,5)	0.200000	(0,1,5)	0.197411	(0,2,5)	0.190927	(0,3,5)	0.183448	(0,4,5)	0.177784	(0,5,5)	0.175704
(1,0,5)	0.197411	(1,1,5)	0.195001	(1,2,5)	0.188946	(1,3,5)	0.181930	(1,4,5)	0.176592	(1,5,5)	0.174627
(2,0,5)	0.190927	(2,1,5)	0.188946	(2,2,5)	0.183930	(2,3,5)	0.178045	(2,4,5)	0.173519	(2,5,5)	0.171841
(3,0,5)	0.183448	(3,1,5)	0.181930	(3,2,5)	0.178045	(3,3,5)	0.173420	(3,4,5)	0.169810	(3,5,5)	0.168461
(4,0,5)	0.177784	(4,1,5)	0.176592	(4,2,5)	0.173519	(4,3,5)	0.169810	(4,4,5)	0.166881	(4,5,5)	0.165778
(5,0,5)	0.175704	(5,1,5)	0.174627	(5,2,5)	0.171841	(5,3,5)	0.168461	(5,4,5)	0.165778	(5,5,5)	0.164764
(0,0,6)	0.166667	(0,1,6)	0.165449	(0,2,6)	0.162331	(0,3,6)	0.158615	(0,4,6)	0.155711	(0,5,6)	0.154625
(1,0,6)	0.165449	(1,1,6)	0.164293	(1,2,6)	0.161332	(1,3,6)	0.157793	(1,4,6)	0.155024	(1,5,6)	0.153987
(2,0,6)	0.162331	(2,1,6)	0.161332	(2,2,6)	0.158761	(2,3,6)	0.155671	(2,4,6)	0.153240	(2,5,6)	0.152327
(3,0,6)	0.158615	(3,1,6)	0.157793	(3,2,6)	0.155671	(3,3,6)	0.153103	(3,4,6)	0.151067	(3,5,6)	0.150299
(4,0,6)	0.155711	(4,1,6)	0.155024	(4,2,6)	0.153240	(4,3,6)	0.151067	(4,4,6)	0.149334	(4,5,6)	0.148678
(5,0,6)	0.154625	(5,1,6)	0.153987	(5,2,6)	0.152327	(5,3,6)	0.150299	(5,4,6)	0.148678	(5,5,6)	0.148063

(0,0,7)	0.142857	(0,1,7)	0.142260	(0,2,7)	0.140716	(0,3,7)	0.138842	(0,4,7)	0.137354	(0,5,7)	0.136792
(1,0,7)	0.142260	(1,1,7)	0.141687	(1,2,7)	0.140202	(1,3,7)	0.138399	(1,4,7)	0.136966	(1,5,7)	0.136424
(2,0,7)	0.140716	(2,1,7)	0.140202	(2,2,7)	0.138869	(2,3,7)	0.137247	(2,4,7)	0.135953	(2,5,7)	0.135464
(3,0,7)	0.138842	(3,1,7)	0.138399	(3,2,7)	0.137247	(3,3,7)	0.135839	(3,4,7)	0.134713	(3,5,7)	0.134285
(4,0,7)	0.137354	(4,1,7)	0.136966	(4,2,7)	0.135953	(4,3,7)	0.134713	(4,4,7)	0.133717	(4,5,7)	0.133339
(5,0,7)	0.136792	(5,1,7)	0.136424	(5,2,7)	0.135464	(5,3,7)	0.134285	(5,4,7)	0.133339	(5,5,7)	0.132979
(0,0,8)	0.125000	(0,1,8)	0.124700	(0,2,8)	0.123919	(0,3,8)	0.122964	(0,4,8)	0.122199	(0,5,8)	0.121908
(1,0,8)	0.124700	(1,1,8)	0.124409	(1,2,8)	0.123652	(1,3,8)	0.122725	(1,4,8)	0.121982	(1,5,8)	0.121700
(2,0,8)	0.123919	(2,1,8)	0.123652	(2,2,8)	0.122956	(2,3,8)	0.122102	(2,4,8)	0.121416	(2,5,8)	0.121156
(3,0,8)	0.122964	(3,1,8)	0.122725	(3,2,8)	0.122102	(3,3,8)	0.121336	(3,4,8)	0.120720	(3,5,8)	0.120486
(4,0,8)	0.122199	(4,1,8)	0.121982	(4,2,8)	0.121416	(4,3,8)	0.120720	(4,4,8)	0.120160	(4,5,8)	0.119947
(5,0,8)	0.121908	(5,1,8)	0.121700	(5,2,8)	0.121156	(5,3,8)	0.120486	(5,4,8)	0.119947	(5,5,8)	0.119741
(0,0,9)	0.111111	(0,1,9)	0.110958	(0,2,9)	0.110558	(0,3,9)	0.110066	(0,4,9)	0.109670	(0,5,9)	0.109519
(1,0,9)	0.110958	(1,1,9)	0.110808	(1,2,9)	0.110418	(1,3,9)	0.109937	(1,4,9)	0.109551	(1,5,9)	0.109403
(2,0,9)	0.110558	(2,1,9)	0.110418	(2,2,9)	0.110052	(2,3,9)	0.109601	(2,4,9)	0.109238	(2,5,9)	0.109100
(3,0,9)	0.110066	(3,1,9)	0.109937	(3,2,9)	0.109601	(3,3,9)	0.109187	(3,4,9)	0.108853	(3,5,9)	0.108726
(4,0,9)	0.109670	(4,1,9)	0.109551	(4,2,9)	0.109238	(4,3,9)	0.108853	(4,4,9)	0.108542	(4,5,9)	0.108424
(5,0,9)	0.109519	(5,1,9)	0.109403	(5,2,9)	0.109100	(5,3,9)	0.108726	(5,4,9)	0.108424	(5,5,9)	0.108308

Appendix E

Crystal Growth and Visualization Code: *tsim* and *tview*

Due to the quantity of time put into these two programs, we feel they've earned their names: *tsim*, the KMC simulation program; and *tview*, the *OpenGL*-based 3-d visualization program. A brief description and the source code for each are given below.

E.1 *tsim*

tsim realizes the new algorithm of Section 4.3, building an environment where many crystals can be "grown" and statistics can be gathered on any of a number of system characteristics (growth, order, ...). Based on a built in scripting mechanism, growth simulations of wildly differing characteristics can be run sequentially, many times, whereby each simulation's output—including 2-d and 3-d final crystal states, step-by-step evolution of the system data for *tview*, growth rates and averaged growth rates for multiple runs including error, etc.—can be controlled for the level of profuseness or brevity one needs (or has space for). A sample input script is given below.

```
"tsim" script file...2 runs: long-range rock salt; n-n 3 species
bmn-bt type system
#
2

ewald table file: ../ewald/20_20_60.ew
dimensions: 20 20 60
species: 2 +1 .5 -1 .5
grow num layers: 25
num MC steps: 2000
or num of runs: 1
kT: 0.1
chem: -0.65
SOS: true
full output?: true
init surface: surfaces/salt.20by20.20.surf
output: ../simulations/lr_salt

nn table file: ../ewald/50_50_150.fak
dimensions: 50 50 150
species: 3 -2 0.3 +1 0.6 0 0.1
grow num layers: 10
num MC steps: 200
or num of runs: 10
kT: 1.0
chem: -2.5
SOS: true
full output?: false
init surface: surfaces/1:2order.50by50.10.surf
output: ../simulations/nn_bmn-bt
```

The potential table input file, whether it be Ewald or nearest-neighbor, along with the initial surface configuration are stored elsewhere for maximum modularity (and speed concerns). The number of runs for the particular script entry can be set resulting in average values of system characteristics, such as growth rate, under one random-number-stream and seed (see Appendix C.2). The number of species and of what charge and quantity can be trivially changed, and the length of the simulation run, either by number of layers grown or by number of MC steps, equally so. Finally, the simulation parameter's themselves—assumed constant through the simulation for now—are inputted. The SOS toggle currently has no effect.

The source code, in C++, follows.