Macroscopic order from reversible and stochastic lattice growth models

by

Raissa Michelle D'Souza

B.S., University of Illinois (1991)

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 1999

Author .....
Department of Physics
August 27th, 1999

Certified by .....
Norman H. Margolus
Research Associate Professor, Boston University CCS
Research Affiliate, MIT Artificial Intelligence Laboratory
Thesis Supervisor

Certified by .....
Mehran Kardar
Professor of Physics
Thesis Supervisor

Accepted by .....
Thomas J. Greytak
Professor, Associate Department Head for Education

Macroscopic order from reversible and stochastic lattice growth models

by

Raissa Michelle D'Souza

Submitted to the Department of Physics
on August 27th, 1999, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

This thesis advances the understanding of how autonomous microscopic physical processes give rise to macroscopic structure. A unifying theme is the use of physically motivated microscopic models of discrete systems which incorporate the constraints of locality, uniformity, and exact conservation laws. The features studied include: stochastic nonequilibrium fluctuations; use of pseudorandomness in dynamical simulations; the thermodynamics of pattern formation; recurrence times of finite discrete systems; and computation in physical models. I focus primarily on pattern formation: transitions from a disordered to an ordered macroscopic state.

Using an irreversible stochastic model of pattern formation in an open system driven by an external source of noise, I study thin film growth. I focus on the regimes of growth and the average properties of the resulting rough surfaces. I also show that this model couples sensitively to the imperfections of various pseudorandom number generators, resulting in nonstochastic exploration of the accessible state space.

Using microscopically reversible models, I explicitly model how macroscopic dissipation can arise. In discrete systems with invertible dynamics entropy cannot decrease, and most such systems approach fully ergodic. Therefore these systems are natural candidates for models of thermodynamic behavior. I construct reversible models of pattern formation by dividing the system in two: the part of primary interest, and a "heat bath". We can observe the exchange of heat, energy, and entropy between the two subsystems, and gain insight into the thermodynamics of self-assembly.

I introduce a local, deterministic, microscopically reversible model of cluster growth via aggregation in a closed two-dimensional system. The model has a realistic thermodynamics. When started from a state with low coarse grained entropy the model exhibits an initial regime of rapid nonequilibrium growth followed by a quasistatic regime with a well defined temperature. The growth clusters generated display a rich variety of morphologies. I also show how sequences of conditional aggregation events can be used to implement reusable logic gates and how to simulate any digital logic circuit with this model.

Thesis Supervisor: Norman H. Margolus
Title: Research Associate Professor, Boston University CCS
Research Affiliate, MIT Artificial Intelligence Laboratory

Thesis Supervisor: Mehran Kardar
Title: Professor of Physics

Acknowledgments

I would first like to thank my advisors Norm Margolus and Mehran Kardar. Without their guidance, attention, and perseverance this work would not exist. I want to thank Norm resoundingly for sharing with me his visionary ideas, and for providing me with hours of stimulating conversation, freedom to explore all my ideas, and an environment unparalleled with regards to inspiration and fun. I want to particularly thank Mehran for always being willing to provide keen insights and for his patience with and belief in me. I especially appreciate the honest assessments of my personal thinking style and my ideas. These assessments will continually guide me on the path to becoming a more lucid, clear thinker.

I want to next thank the people who opened the doors of NE43 to me: Tom Knight, Gerry Sussman, and Hal Abelson. I want to sincerely thank them, not only for providing me with a stimulating research environment, but for sharing their wisdom, and showing me that it can be extremely fun to be a serious person.

I would like to thank the other members of the Information Mechanics group, especially Mark Smith and Harris Gilliam. Mark was an unofficial advisor to me at many times, passing on the oral lore of the IM group and the oral lore of doing computational physics experiments. Harris I thank for priceless hours of programming help as well as being an engaging office mate.

I would like to thank the people who helped me get here, namely my family. I want to thank my father, Frank, for showing me what it means to have personal conviction and courage, and how to demand the best of others by always demanding it in yourself. I want to thank my mom, Cecilia, and sister, Geri, for their unfettered support and encouragement. Their cheerful voices and awe and respect of what I was doing provided me with much inspiration in dark hours. I want to thank my partner in adventure, Jeff Hunter. His unconditional love for and support of me and his love of life are the very reasons why I have reached this point in my life. My fundamental respect and admiration will never falter, and I am forever grateful to him for teaching me the art of living and the beauty of loving.

I would like to thank the friends who have stood by me and stretched my mind in wonderful ways, especially: David Gloss, in summary, for having the soul of the Prophet; Janet Cahn for comic relief, moral support, and for joining me in countless whimsical and momentous diversions; Evelyn Fox Keller for her encouragement, incisive analysis, and for being an incontrovertible role model of how to lead an interesting independent life; Nancy Horn, Will Meyers, Krys D'Antonio, Nathan Cherson, and Cristina Lanzl all for their infectious vitality and for being amazing, diverse people who have nothing to do with MIT, and everything to do with keeping me sane over the years!

I would like extend my thanks to all of the scientists who have given me guidance and input, especially Bruce Boghosian, Yaneer Bar-Yam, and Yakov Kantor. Bruce's approval has served as my litmus test of good science. Yaneer's ideas contributed extensively to the chapters in this thesis dealing with stochastic growth. And, of course I thank Yaneer for his indelible enthusiasm which helped get this all started. Yakov I thank for teaching me many things about the subtleties of analyzing data

from numerical experiments.

From the early years I want to thank the "homework klatch", especially Joe Habib, David Abusch-Magdar, Charlie Katz, and Josh Erlich. I also extend my gratitude to all my office mates over the years—Jim McBride, Rava da Silveira, Dan Risacher, Ruben Agin, and Mike Biafore—for countless hours of stimulating discussion and critical input.

I would like to thank two of the most cheerful and pleasant people at MIT, Peggy Berkovitz and Anthony Zolnik. Their personal involvement has helped me navigate the waves of bureaucracy, and has made me always feel appreciated around here.

Finally I would like to thank George Homsy for showing me unwaivering emotional support over the past year, for being a constant catalyst to my growth—pushing me to define and sharpen my ideas—and for eagerly blazing down the trail of learning with me, wherever it leads. His input (from the spiritual counsel to the ergodic theory tutorials to the LaTeXhacking) has made this manuscript complete.

Reflecting on my years at MIT, I can earnestly say, the most rewarding aspect of being here is meeting the unique and profoundly interesting people.

This research was supported in part by the MIT Department of Physics, under the Departmental Minority Graduate Student Fellowship, and in part by the following government agencies: National Science Foundation (Grant Nos. DMR-93-03667 and DMS-95-96217), and Defense Advanced Research Projects Agency (Grant No. DABT63-95-C-0130).

To the memory of my father, Anthony Frank D'Souza. By his humble nature he would be more proud of this document than he was of any of his own illustrious accomplishments.

It is now night in Goa.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 The scope of this thesis

This thesis advances the understanding of how autonomous microscopic physical processes give rise to structure—the emergence of macroscopic order from microscopic growth rules. This is a large issue and a broad array of topics are herein addressed: local microscopic growth rules for pattern formation; stochastic fluctuations of nonequilibrium processes; use of pseudorandomness in dynamical simulations; emergence of order from deterministic, reversible microscopic rules; the thermodynamics of pattern formation; recurrence times of finite state discrete systems; and computation in discrete lattice systems. These topics overlap the fields of statistical physics, cellular automata, computational physics, nonlinear dynamics, and computer science.

A primary aim of this research is to gain insight into how microscopic mechanisms give rise to macroscopic correlations in both irreversible and reversible microscopic models. The irreversible stochastic models studied allow us primarily to understand average properties of growth interfaces, scaling behavior, and dynamic fluctuations. The reversible deterministic models studied allow us to gain insight into the thermodynamics of pattern formation and self-assembly. Instead of trying to understand these issues in the broadest context, we restrict our attention to a tractable subclass

of systems: discrete systems on regular spatial lattices.

We are concerned with physically motivated microscopic models. We want to put in constraints at the microscopic scale consistent with the microscopic properties of physics which persist at the macroscopic scale, primarily: locality, uniformity, and conservation laws. To this end we develop microscopic models with local dynamics, meaning that in a short time the evolution of a part of the system is influenced only by the state within a nearby region of space. The dynamics are also uniform, meaning that the entire system evolves in accord with the same rules. The use of discrete systems allows for exact conservation laws. We study both irreversible stochastic dynamics and invertible deterministic dynamics which adhere to these general principles. The motivation for studying invertible dynamics is two fold. First, the laws of physics are microscopically reversible. Second, entropy cannot decrease in reversible systems (as discussed below). If we coarse grain the system, entropy can increase, and we can observe the approach to thermodynamic equilibrium. In contrast, the irreversible models are open systems which allow us to avoid many of the subtleties attendant in making models conform to microscopic reversibility.

1.2 Types of reversibility

"Reversibility" can have different meanings, with subtle distinctions between them. The forms of reversibility of interest to us are: thermodynamic reversibility, statistical reversibility, and exact (or microscopic) reversibility. Thermodynamic reversibility is a property of a dynamics on the macroscopic states of a system. The other forms are properties of dynamics on distinct microscopic states of a system. I first discuss irreversibility and then briefly summarize the categories of reversibility listed above and their distinctions.

Irreversible dynamics are not backwards deterministic. In general, many distinct microscopic states may make a transition to the same target state under an irreversible

dynamics, so each state may not have a unique predecessor. At each point where two trajectories merge, we lose information about the past. Irreversible dynamics are characterized by the familiar concepts of fixed points, limit cycles, basins of attraction, and strange attractors.

A thermodynamically reversible process is one in which the value of the macroscopic entropy for a closed system does not change. The macroscopic entropy, proportional to the log of the number of microstates compatible with given macroscopic properties, is the definition of entropy typically used in statistical physics. The macroscopic properties are the thermodynamic variables, such as temperature, volume, and pressure. We can invert the time evolution of the macroscopic thermodynamic properties of a closed thermodynamically reversible system without addition of external work.

A statistically reversible dynamics is a probabilistic dynamics with semidetailed balance. Semidetailed balance means that the sum of the transition probabilities out of a particular microscopic state is equal to the sum of the transition probabilities into that same state. A consequence of semidetailed balance is that in equilibrium all accessible microscopic states are equally likely. Systems with this property can be used to model thermodynamic behavior.

By exact reversibility we mean the dynamics are exactly invertible. This is a stronger condition than thermodynamic reversibility or statistical reversibility. Each microscopic state has a unique predecessor and a unique successor: a complete history of the system can be generated from any state through which the system progresses. The probability to make a transition from one microscopic state to any other is unity for exactly one target state and zero for all others. When initialized with a particular probability distribution ($\{p_i\}$) of states, the shape of that distribution cannot change as the system evolves; the probability simply "moves around" between states. So the Gibbs entropy $S = k_B \sum_i p_i \ln p_i$ is constant. If we coarse grain the system, and thus lose information about the microscopic correlations, the entropy can

increase. Because of this property, we can use exactly reversible systems to model thermodynamic behavior while retaining microscopic reversibility as a fundamental principle of our dynamics. For a thorough discussion of entropy, thermodynamics, and coarse graining see, for example, Ref. [84].

1.3 Stochastic irreversible dynamics

Stochastic processes require a supply of randomness. In almost all standard computer implementations of stochastic processes the randomness is replaced by deterministic "pseudorandomness", traditionally using pseudorandom number generator algorithms. These algorithms are not physically motivated, nor are they local.

The irreversible models of crystal growth we study allow state transitions in only one direction: the probability of transitioning back to a state which has already been visited is zero. Thus there is a clear direction of time. Things grow, they never shrink. We study irreversible stochastic dynamics: a collection of random variables determines the evolution of the process. If we had a source of true randomness, and we wanted to make the process reversible, we would have to keep a complete history of the system: a running list of the outcomes of all the random events. The source of randomness is external to the simulated system as is the sink of information (as explained below). We study a general model of thin film growth, focusing on the regimes of growth and the average properties of the rough surfaces of the growth structures.

1.4 Physically motivated pseudorandomness

Physical processes at the macroscopic scale appear random and dissipative, but the laws of physics are microscopically reversible—there is no dissipation of information at the microscopic scale. Information cannot be destroyed. It can however be trans-

ferred to degrees of freedom which are not directly tracked and is usually ultimately transferred to thermal degrees of freedom.

To reconcile macroscopic "randomness" with microscopic determinism, we model systems which rely on the outcomes of random variables in a physically motivated way: the randomizing degrees of freedom are explicitly modeled as the outcome of a local, uniform, invertible process. We accomplish this by dividing the model into two subsystems: the system of primary interest, and a "heat bath" system which physically models the randomizing influences. The complete system is exactly reversible. We can observe the exchange of heat, energy, and entropy between these two subsystems.

1.5 Pattern formation

We wish to understand how local dynamics can give rise to large scale order. The types of questions which we might try to answer include: How does a growing bacterial colony organize into an "ordered" fractal? How do aggregating ice crystals grow into a symmetric branched snowflake? How does a growing diamond film form surface patterns and ordered facets?

Pattern formation is the transition from a disordered to an ordered macroscopic state (*i.e.*, a process with a decrease in coarse grained entropy). Thus it is conceptually difficult to reconcile exact reversibility, which cannot result in a change of the amount of information at the microscopic scale, with pattern formation, which requires a change at the macroscopic scale. To do this we make use of the heat bath technique introduced above to explicitly model how dissipation arises (*i.e.*, how information flows between the macroscopic and the microscopic degrees of freedom). We provide a clear example of how to reconcile the macroscopic irreversibility that gives rise to patterns with the microscopic reversibility adhered to by physical processes. We see pattern formation in part of the system accompanied by an increase of entropy in the

heat bath.

Using reversible dynamics gives us the potential for a realistic thermodynamics. In addition the dynamics is local—local heat flow and the local creation of coarse grained entropy. Thus we can do more than study simulated structures: We have a laboratory for studying nonequilibrium thermodynamic behavior. By modeling the entire process of pattern formation, including the dissipation, we hope to emulate the essential mechanisms of nature more realistically than do irreversible models.

1.6 Physics and computation

We are concerned with developing models of physics which obey essential constraints of physics. Physical processes, such as computation, must also operate within these constraints. The insights which we gain from developing computer models of physics can help us understand how to build physical models of computers. We show how some of the reversible models we develop can also be used to simulate digital logic and thus perform computations with no dissipation. This may help us understand how to make more efficient physical devices for information processing and improve our ability to harness physics for computation.

1.7 Summary of results

In Part I of the thesis I study a well known, stochastic lattice growth rule, Ballistic Deposition[125], which serves as a paradigm for nonequilibrium lattice growth and dynamic scaling behavior. An initial detailed numerical implementation of the BD model manifests anomalous macroscopic properties of the resulting growth interface that cannot be explained within the accepted theoretical framework: The simulated interfaces do not obey simple self affine scaling relations and those scaling exponents which can be measured do not agree with any theoretical prediction. In order to

identify the cause of the anomalies, I conduct an extensive, statistically rigorous numerical study of the model. This study demonstrates that the model couples to various pseudorandom number generators, resulting in nonstochastic exploration of the accessible phase space. Although the BD model has been studied for more than a decade, there is no consensus in the literature for the value of the scaling exponents for BD growth clusters. I suggest reasons for this and propose methods for better use of PRNG's in numerical simulations.

In Part II of this thesis, I focus on invertible dynamics on discrete systems. I propose a new model for lattice growth which is a thermodynamic extension of the well known Diffusion Limited Aggregation[129] model. This Reversible Aggregation model is local, deterministic, microscopically reversible, and has exact energy conservation. In addition the model has a realistic thermodynamics: We can observe a quasistatic increase of entropy in a regime with a well defined temperature.

The RA model is a model of reversible cluster growth via aggregation in a closed two-dimensional system. There are three species of particles in the RA model: gas, heat, and crystal. A diffusing gas particle may aggregate at the boundary of a crystal cluster. Latent heat is released during each aggregation event and is explicitly modeled by introducing a heat particle into a diffusing heat bath. Conversely if a heat particle contacts a singly connected cluster member it may be absorbed and that crystal particle will "evaporate" from the cluster, becoming a gas particle. When started with a dilute gas and a single crystal seed particle the model exhibits an initial regime of rapid nonequilibrium growth followed by a slow quasistatic regime with a well defined temperature. In the first regime the crystal rapidly grows, in the second the crystal slowly anneals. Growth clusters generated by this model display a rich variety of morphologies. I present a detailed study discussing the formulation of the model, the implementation, an analytic formulation (including the continuum limit of the dynamics), the regimes of growth, the scaling behavior of the clusters, and control of the cluster morphology.

The aggregation events in the RA model depend conditionally on the state of gas, heat, and crystal particles within the appropriate interaction range. The interaction is sufficiently complex that we expect *a priori* that the system is capable of universal computation: that universal logic gates can be built out of the aggregation interactions. Showing the RA model is capable of universal computation is not difficult, showing it is capable of *efficient* computation is more challenging. I show how to explicitly embed reusable computational primitives into the RA model and thus how to simulate any digital logic circuit using the RA model. The mobile gas and heat particles are the logic signals used in the computation. The paths these particles take are the wires. Sequences of conditional crystallization events are the basis of the logic gates. I also discuss the relationship of computation in the model to computation in real physical systems.

Lastly I present a study of general aspects of the dynamics of discrete, reversible systems, extending earlier work. For a dynamics on a finite size system, the initial state must eventually recur. If all states are on the same cycle, the dynamics is fully ergodic. If not, the dynamics partitions the space into distinct orbits. I study the random network limit, where any state in the system can map onto any other state, thus neglecting the constraint of locality. First the ensemble of all invertible dynamics on a discrete finite system is established, then the expectation value of the orbit lengths over the ensemble and the expected number of orbits per dynamics are established. The orbits for a typical dynamics can be labeled by a conserved quantity with the same range as an extensive variable, which suggests that dynamics with extensive conservation laws (such as CA's) are not excluded from the consideration of typical dynamics. Finally the ergodicity of invertible dynamics is discussed and contrasted with the ergodicity of unconstrained dynamics (the ensemble of dynamics which includes noninvertible ones). I show that invertible mappings are approximately ergodic, in a suitably defined statistical sense, whereas unconstrained mappings are not.

1.8 Placing the results in context

Of all the topics addressed and results obtained, I believe the most significant are those which deepen the current understanding of pattern formation in discrete, reversible systems. This is a field where little has been done in the past. By constructing a thermodynamic extension to a well studied model of crystal growth, we observe a rich variety of resulting growth morphologies and enable the study of thermodynamics of nonequilibrium systems.

The results showing that computation in physical models can be accomplished with exact control of the microscopic degrees of freedom, would be more significant if we could extend the result to show we can compute with only macroscopic control. With this type of understanding we might be able to compute with a growing crystal aggregate, or even an aggregate of growing cells.

This work has been great fun. Collectively it should be taken as a step toward understanding the nonequilibrium process of pattern formation. At the "meta-level" this is a small step toward relating microscopic information and computation with macroscopic physics.

These topics form a deep ocean. I have just begun to skim the surface.

Part I

Irreversible Models

Chapter 2

Review of stochastic growth models

2.1 Overview of Part I

2.1.1 Growth of rough surfaces

Spontaneous growth processes are responsible for the structure we observe in the natural world around us. Many of these processes involve growth of rough surfaces. Hence understanding rough growth surfaces of clusters and solidification fronts has become an increasingly relevant issue[42, 53, 91, 5, 56]. Two approaches have traditionally been taken. The first is developing discrete algorithmic models of growth processes, iterating rules to grow structures from the atomistic scale. The second is developing analytic continuum models of growth, using stochastic differential equations to represent the key processes responsible for the resulting surface morphologies. The average properties of the growth surfaces generated with the two approaches can be directly compared.

Rough growth surfaces can in general be described by concepts of fractal geometry and most rough surfaces exhibit self-affine scaling over a range of length scales (as

discussed in Sec. 2.2.2). Regimes of dynamic scaling have been found in most models of surface growth. Furthermore, many seemingly unrelated models of growth have similar scaling exponents, suggesting the existence of universality classes. These properties of surface growth models make them relevant to those who study modern statistical physics. These simple models also lend themselves to the study of non-equilibrium processes in general.

2.1.2 Summary of results presented in Part I

Ballistic Deposition (BD) is a prototypical discrete algorithmic model for interface growth and for exhibiting dynamic scaling behavior in nonequilibrium systems. It seems to be an accurate model of thin film growth. In BD, an aggregate is grown on a substrate of finite size L. Particles are sequentially added at randomly selected active sites on the surface of the aggregate. Active growth sites are defined in relation to the state of their nearest neighbors: at least one nearest neighbor must already be an aggregate member. Thus neighboring surface particles are spatially correlated. The aggregate typically begins as a flat interface, with no correlations. As particles are added the roughness of the growth interface, ξ, increases quickly, scaling with the time as $\xi \sim t^{\beta}$. The lateral correlation length, $\xi_{||}$, concurrently grows as a power of time up to the point where the entire surface is correlated, $\xi_{||} \sim L$. Once the entire surface is correlated, the roughness saturates, reaching its maximum value. In this steady state regime the roughness fluctuates but the average value no longer changes with time, instead it scales with the size of the underlying substrate: $\xi \sim L^{\chi}$. The resulting BD growth aggregates are compact clusters with a rough surface. The surface is expected to be self-affine and exhibit dynamic scaling behavior.

BD is typically investigated through computer simulations where randomness is replaced by use of deterministic Pseudorandom Number Generators (PRNGs). I implement a model of BD over a range of length and time scales to establish the scaling behavior of the growth interface. I investigate the initial roughening as a function of

time, the scaling of the surface roughness with length in the conjectured steady state regime, and the crossover time between these two behaviors. I am able to establish clear scaling behavior, yet the results are discrepant with the prevailing paradigm. First, the value of the roughness scaling exponent, χ, obtained is below the value for a random walk (i.e. $\chi < 1/2$). The value $\chi = 1/2$ is predicted by the prevailing theory (KPZ theory) and many models of growth obtain this exponent. Second, height fluctuations of the growing interface appear not to satisfy simple scaling. Third, a decrease in the surface roughness is observed in a conjectured steady state regime. A thorough search for corrections to scaling necessitates obtaining better statistics and the exploration of longer times into the growth. Although none of the corrections reveal any solutions, the additional statistics obtained make it clear that computer implementations of BD may be responsible for the discrepancies. Implementing the model with different PRNGs yields statistically inconsistent results.

The implicit assumption that PRNGs adequately represent true randomness is tested via a detailed statistical analysis of the width of the BD interface. We study the width of the interface over time scales orders of magnitude longer than the expected model relaxation time, yet much smaller than the period of the PRNG, and observe fluctuations which still appear to be correlated. Distinct dynamic behavior is observed for an implementation with a different PRNG, further indicating a strong coupling between the model and the PRNGs (even with PRNGs that pass extensive statistical tests). Thus we demonstrate a breakdown of basic sampling assumptions, and of the ergodic exploration of phase space. The details of the BD implementation and the statistical inconsistencies will be explicitly presented.

2.1.3 Organization of Part I

This section of the thesis begins with this review chapter. First I review the well known algorithmic models of discrete stochastic growth. Then I briefly review the concepts relevant to the dynamics of surface growth: roughening; dynamic scaling of

the growth interface; scaling of the mass of the clusters; correlations and saturation. I then discuss some continuum models of stochastic growth based on symmetry motivated differential equations. Chapters 3 and 4 contain the details of the numerical studies yielding the results summarized above.

2.2 Review of stochastic growth

I present here only a brief review of the concepts and models of stochastic growth. Many comprehensive reviews have recently appeared[124, 91, 74, 5, 56], as well as many collections of reprinted articles[117, 42, 53]. The growth processes I am studying are all $1+1$ dimensional (growth occurs along the columns of a discrete square lattice of finite size). Thus, for simplicity, the discussion below assumes $1+1$ dimensional growth unless otherwise explicitly stated.

2.2.1 Lattice models of rough surfaces

Complex disordered patterns resembling those seen in nature can easily be generated by simple probabilistic algorithms. Thus simple algorithmic models play a central role in our understanding of surface growth. In the interest of concreteness I will first describe various well known models of discrete stochastic growth.

2.2.1.1 The Eden model

The earliest model of cluster growth is the Eden model, which was introduced in 1958 as a model of bacterial colony or tissue culture growth[33, 34]. A two-dimensional lattice is initially empty except for one lattice site which is occupied. This site is a seed particle. The growth process is represented by choosing at random an unoccupied site which has an occupied nearest neighbor (*i.e.,* a perimeter site), and filling it in. This procedure results in a cluster with a compact internal structure but a rough surface. Here compact means uniform density (note there can be small holes inside

the bulk aggregate). Note that there are lattice effects. We are using a nearest neighbor growth rule on a square lattice. The growth velocity is slightly greater along the basis directions than along the diagonal directions, thus the cluster will asymptotically approach a diamond shape.

Many variants on this model have been introduced for a range of purposes such as modeling the progress of skin cancer[128], or simulating interface properties by starting the growth from a single straight line. There are several variants on the microscopic rules (for example chosing an empty nearest neighbor site of the cluster at random and filling it in, versus choosing a perimeter cluster member at random and then filling in a empty nearest neighbor site at random). These variants have different short time behaviors, yet in the long time and large size regime, they exhibit the same morphology and scaling behaviors[65].

2.2.1.2 Random Deposition

Perhaps simpler than the Eden model is the Random Deposition model (RD) of interface growth. It starts with a single straight line of particles occupying an otherwise empty two dimensional square lattice. At each update a column in the space is chosen at random and a particle added to the lowest unoccupied lattice site in that column. Unlike the Eden model, the columns are independent and uncorrelated. The RD model generates a completely occupied bulk structure with a rough surface, as shown in Fig. 2-1. There are no voids in the bulk.

Exact results from atomistic algorithm: Since there are no correlations between columns, each column grows independently. This makes the RD model easily amenable to analytic treatment—in fact, the RD model is exactly solvable. The state of the aggregate can be described completely by specifying the length of the substrate, L (*i.e.,* the number of columns), and the vertical coordinate of the highest occupied site in each column at time t into the growth, $h(x,t)$ (*i.e.,* the height of the interface

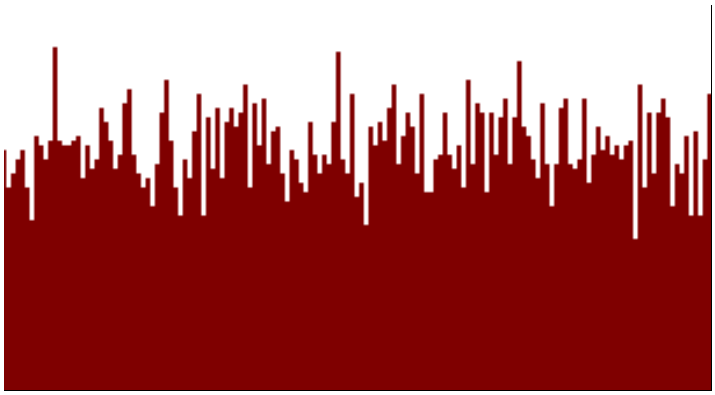

Figure 2-1: A typical RD model growth cluster. Note adjacent columns are uncorrelated.

in column x).

During each update of the space, every column grows by one unit of height with a probability $p = 1/L$. Thus the column heights follow a binomial distribution:

$$P(h,N) = \left(\begin{array}{c} N \\ h \end{array} \right) p^h (1-p)^{(N-h)}, \tag{2.1}$$

where $P(h,N)$ is the probability that a column has height h after the deposition of N particles. We define a length independent measure of time, $t = N/L$, and thus can easily show the expectation value of the height grows linearly in time:

$$Ex[h(t)] \equiv \sum_h hP(h,t) = Np = \frac{N}{L} = t. \tag{2.2}$$

The second moment also grows in time:

$$Ex[h^2(t)] \equiv \sum_h h^2 P(h,t) = Np(1-p) + N^2p^2 = \frac{N}{L}\left(1 - \frac{1}{L}\right) + \frac{N^2}{L^2}. \tag{2.3}$$

The standard deviation, $\xi(t)$, can serve as a measure of the width of the active growth

interface. The variance can be exactly determined to be

$$\xi^2(t) \equiv Var[h(t)] = Ex[h^2(t)] - Ex[h(t)]^2 = \frac{N}{L}\left(1 - \frac{1}{L}\right) \approx t. \tag{2.4}$$

Thus we have exactly solved for the average height, $Ex[h(t)]$, and average width of the growth interface, $\xi(t)$. The width scales with time, $\xi(t) \sim t^{\beta}$, with the scaling exponent $\beta = 1/2$. Note that in this model the width grows indefinitely with time, without saturation. Adjacent columns are independent hence values of adjacent heights are uncorrelated. Thus differences in height amongst adjacent columns are unconstrained and grow indefinitely. The surface is *not* self similar. The ideas of scaling and self similarity will be explained in Sec. 2.2.2.

Continuum treatment: The RD model is simple since there are no correlations between columns, and thus the values of the relevant growth parameters can be solved for exactly. For other growth models this is not the case, and alternative approaches must be applied. A useful approach is to begin with continuum stochastic growth equations. Despite having obtained an exact solution for the RD model, for pedagogical reasons I illustrate here the stochastic growth equation approach.

We introduce a differential equation to represent random deposition. The growth of the height can be described as

$$\frac{\partial h(x,t)}{\partial t} = \Phi(x,t), \tag{2.5}$$

where the flux $\Phi(x,t)$ is the number of particles per unit time arriving on the surface at position x and time t. Since the particle flux is random we can decompose the flux into two parts,

$$\frac{\partial h(x,t)}{\partial t} = F + \eta(x,t), \tag{2.6}$$

where F represents the average number and η the fluctuations of the average number. The noise η should be random with mean of zero, $\langle \eta(x,t) \rangle = 0$, and without any

correlations in space or time, meaning the second moment of the distribution should be of the form

$$Ex\left[\eta(x,t)\eta(x',t')\right] = 2D\delta(x-x')\delta(t-t'). \tag{2.7}$$

Two simple distributions that satisfy these conditions are the Gaussian distribution with probability density function:

$$\Pi\left[\eta(x,t)\right] = \frac{1}{\sqrt{4\pi D}}\exp\left[-\frac{\eta^2}{4D}\right], \tag{2.8}$$

(it is independent of x and t; η is independently Gaussian distributed at each site) and "bounded noise" in which $(+1)$ and (-1) are chosen with equal probability.

Integrating Eq. 2.5 over time we obtain

$$h(x,t) = Ft + \int_0^t d\tau\eta(x,\tau), \tag{2.9}$$

and for the expectation value of the height we obtain

$$Ex[h(t)] = Ft. \tag{2.10}$$

We can integrate for the expectation of $h^2(x,t)$ as well. Noting that $\langle\eta^2(x,t)\rangle = 2D$, we obtain

$$\xi^2(t) = Ex[h^2(t)] - Ex[h(t)]^2 = 2Dt. \tag{2.11}$$

Thus we obtain the same scaling of the growth from the continuum stochastic differential equations as from the microscopic algorithm.

2.2.1.3 Ballistic Deposition

Ballistic Deposition (BD) was introduced in 1959 by Vold as a model of growth of colloidal aggregates by sedimentation[125]. Typically the BD model is initialized with a single straight line of particles (the substrate) occupying an otherwise empty two

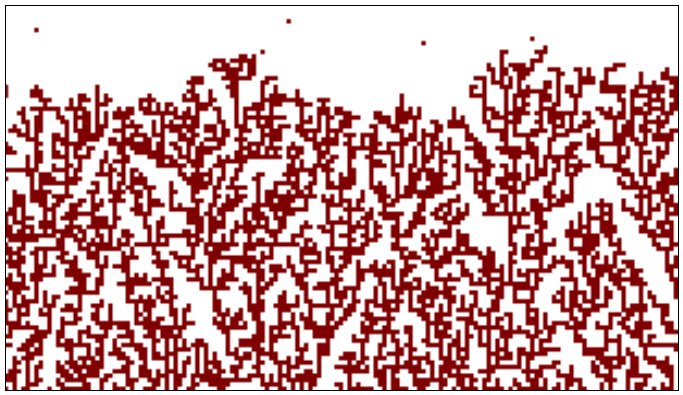

Figure 2-2: A typical BD model growth cluster. Note this implementation is for a parallel version of BD, were several particles are allowed to fall at once. The falling particles are shown as the small spots above the cluster.

dimensional square lattice. The substrate of length L, consists of discrete columns indexed by integer values x, with $1 \leq x \leq L$. A particle is placed above the substrate in a randomly selected column and descends along a straight vertical path until encountering a site on the surface of the existing cluster and sticking there. Once that particle aggregates, another particle is randomly placed and descends. Since particles stick upon first contacting the aggregate, growth can occur laterally (in the direction parallel to the substrate), resulting in voids and overhangs in the resulting aggregates. Initial interest in BD focused on the resulting porous bulk aggregate. But with the introduction of experimental techniques for growing thin crystal films, such as molecular beam epitaxy, interest has shifted from the bulk to the surface properties. A typical BD growth cluster is shown in Fig. 2-2.

The growth interface is defined by the maximum height occupied along each column, $h(x,t)$, where $h(x,t)$ takes on discrete integer values. Starting from a flat interface, $h(x, t=0) = 0$ for all x, the surface evolves by sequential addition of particles to randomly chosen columns. The index number of particles deposited is denoted by t', and the deposition time by $t = t'/L$. Each deposition event consists of choosing a column, $x(t')$, by a call to a PRNG, and updating the height in that column as

follows:

$$h(x(t'), t'+1) = \max[h(x(t')-1, t'), h(x(t'), t')+1, h(x(t')+1, t')]. \tag{2.12}$$

Thus the deposited particle occupies the highest empty site with one or more occupied nearest neighbor sites; this mimics the process of cluster aggregation. The stochastic process in this model is the random choosing of successive columns.

While the resulting aggregates are compact, their interface is rough, with fluctuations that are expected to be self similar at all scales. The width of the growth interface $\xi_L(t)$, on average increases following a power law behavior until reaching a steady asymptotic value, the magnitude of which depends on the underlying substrate size L. A good measure of $\xi_L(t)$ is the sample variance of the surface heights[1], $\{h(x,t)\}$,

$$\xi_L^2(t) = \frac{1}{L}\sum_{x=1}^{L}\left(h(x,t) - \overline{h(t)}\right)^2, \tag{2.13}$$

where $\overline{h(t)}$ is the sample mean of the surface heights at time t.

Results presented in this chapter all involve an implementation of the BD model. The full details of the model and the implementation will be given in Chaps. 3 and 4. Note that above we have added the subscript L to the width of the growth interface ξ. This is a reminder that the roughness is measured across the entire substrate from $x = 1$ to $x = L$. We will study ξ_L, but we will also study ξ_l—how the roughness scales along the substrate up to length $l \leq L$. Simple scaling suggests that $\xi_l \sim l^{\chi}$.

[1] We use the sample variance as an estimator of the true variance of the distribution. In order to be an unbiased estimator, the denominator in Eq. 2.13 should be $(L-1)$ instead of L. This subtlety only makes a difference for very short lengths, on the order of $L < 10$, which are not included in scaling considerations due to discreteness effects. As every reference on surface growth uses the definition given in Eq. 2.13, we stick with the convention.

2.2.1.4 Additional growth models

Countless variations on these models exist. One can alter the microscopic rules or the initial conditions. One additional growth rule which is relevant to our study is the Restricted Solid on Solid model[71]. The RSOS model is identical to the BD model, with the exception that growth events are only accepted if the difference in height between adjacent columns is within some acceptable range (usually the difference is constrained to be ± 1). Growth events are otherwise rejected. This yields a smoother growth interface (the rejection of events results in some local averaging). This also allows one to obtain statistics with tighter error bounds using less simulation time (*i.e.*, the time to converge upon the asymptotic behavior is reduced over the BD model).

There are countless other models of growth. One common model, Diffusion Limited Aggregation[129], will be discussed in Chap. 6.

2.2.2 Roughness, correlations, and scaling behavior

The growth models considered above all share certain fundamental characteristics. They model irreversible, far from equilibrium growth. They model stochastic roughening: the substrate or interface is initially flat; particles are added to the aggregate in random positions according to well specified microscopic rules; the width of the active growth interface increases as a function of time and the surface of the growth aggregate goes from being flat to rough. The surfaces can be described by a single valued function of position, $h(x,t)$, representing the highest occupied site in each column of the aggregate. Note that while there can be voids and overhangs in the bulk structure, all the bulk structures considered are compact (*i.e.,* have a uniform density). Two important quantities describing the surface are the sample average of

the surface heights at time t:

$$\overline{h(t)} = \sum_{x=1}^{L} h(x,t)/L, \tag{2.14}$$

and the sample variance of the surface heights at that time. Eq. 2.13. The variance gives a measure of the width of the active growth interface. Note if we begin with a theoretical continuum treatment (as we did for the RD model in Sec. 2.2.1.2), we use the expectation value and variance of a distribution function to describe the growth process. If we begin with a numerical treatment, we use the numerically determined sample mean and sample variance.

The surface initially starts out flat with zero width. As particles aggregate in random places, the surface roughens. It has been found that most of these surfaces are self affine—that the statistical properties of the surface are invariant under a scale transformation (looking at different length scales, the properties of the interface are equivalent within a rescaling). For example when we increase the length scale by a factor of b the response increases by a power of b:

$$F(bx) = b^{\gamma} F(x). \tag{2.15}$$

Note the distinction between self affine and self similar. Self affine fractals can be rescaled by transformations that require different changes in length scales in different directions (they are not isotropic). Self similar fractals can be rescaled via the same change of length scales in all directions.

Initially the width of the interface increases as a power of time,[2]

$$\xi(L,t) \sim t^{\beta} \quad \text{for} \quad t < t_x. \tag{2.16}$$

[2]Note here we are using the notation $\xi(L,t)$ rather than $\xi_L(t)$, since we are explicitly discussing the situation where L can vary.

However, due to the lateral correlations in surface heights and the finite size of the system, the width eventually saturates. The saturated value of the width scales with a power of the size of the system,

$$\xi(L,t) \sim L^{\chi} \quad \text{for} \quad t > t_x. \tag{2.17}$$

The crossover time, t_x, at which the interface crosses over from the behavior of Eq. 2.16 to the behavior of Eq. 2.17 depends on the system size:

$$t_x \sim L^z. \tag{2.18}$$

The exponent χ is called the roughness exponent; z is the dynamic exponent. It was pointed out by Family and Viscek[41] that the three exponents, β, χ, and z, are not independent and that the scaling forms for the growth and saturation of the width of the interface can be described by a scaling ansatz, similar to that applicable to critical systems;

$$\xi(L,t) = L^{\chi} f(t/L^z), \tag{2.19}$$

where $f(x) \sim x^{\beta}$ for $x \ll 1$ and $f(x) = const$ for $x \gg 1$. Thus $z = \chi/\beta$. Detailed use of the scaling function in Eq. 2.19 will be shown in the later sections of this chapter. We will use the function to collapse data from different length substrates onto one curve.

The RD model discussed in Sec. 2.2.1.2 can be solved exactly. It is a simple model where each column grows independently of the others. No information about other columns ever propagates in the system. In contrast the other models considered have nearest neighbor interactions. The result of a growth event in one column depends directly on the state of that column and also on the state of both neighboring columns. Thus columns are not independent, they are correlated—the surface heights contain information about neighboring heights. Initially the length over which spatially sep-

arated columns are correlated is small (this lateral correlation length is denoted by $\xi_{||}$). Yet, as the growth process continues, information about neighboring columns spreads laterally and $\xi_{||}$ increases as some power of time. For a finite length substrate the entire surface will eventually be correlated: $\xi_{||} \sim L$, where L is the length of the substrate. Once the correlation length reaches the maximum allowed length, the width of the growth interface saturates, being unable to increase anymore. Although the saturation of the width is merely postulated here, the continuum equations predict this saturation. Note that the perpendicular correlation length is related to the surface width ξ.

We have seen that the same general principles apply to many seemingly unrelated growth models. The growth interface initially roughens as a power of the time. The parallel correlation length grows initially as a power of time. Once the entire surface is correlated, both the roughness and the correlation length saturate, their values scale with the size L of the underlying substrate. We can classify the different stochastic growth models distinctly by the values obtained for their scaling exponents.

2.2.3 Continuum models of rough surfaces

The Random Deposition model is a simple model with no correlations, and as such is amenable to exact theoretical treatment both in its discrete algorithmic form and the continuum differential equation form. Other models are not so easy to solve for, other than by directly simulating them. This has its own complications, as will be seen throughout the remainder of this chapter. There have been attempts to theoretically predict growth properties using symmetry motivated differential equations. I will briefly review these analytic growth models.

2.2.3.1 General symmetry considerations

We expect to describe the evolution of the surface by a differential equation of the form

$$\frac{\partial h(x,t)}{\partial t} = G(h,x,t) + \eta(x,t). \tag{2.20}$$

Where $G(h,x,t)$ is a general function of height, position, and time. Symmetry considerations will allow us to put restrictions on this function. A basic symmetry we want the growth equation to posses is **(i)** translational invariance in the direction perpendicular to the growth direction. The growth should not depend in the value of x: In other words the equation should be invariant under the transformation $x \to x+\delta_x$. This means G cannot explicitly depend in x. **(ii)** The rule should be independent of where we define the value $h = 0$, and thus invariant under the transformation $h \to h + \delta_h$. This means G cannot explicitly depend on h, yet it can depend on a derivative of h, for example ∇h. **(iii)** The rule must be independent of where we choose the origin of time, thus G cannot explicitly depend on t. **(iv)** The rule must have reflection symmetry about the growth direction: it must be invariant to the transformation $x \to -x$. Thus G must only depend on terms that have an even number of derivatives in x. This together with constraint **(ii)** means G can have terms of the form $(\nabla h)^2$,$\nabla^2 h$, etc. **(v)** Finally, if growth is allowed only in the direction perpendicular to the interface (*i.e.,* no lateral growth is allowed) and the interface can be considered at or near equilibrium, we want the interface fluctuations to be similar about the mean interface height. Thus G should be invariant about the transformation $h \to -h$. Note this final constraint excludes the inclusion of a term of the form $(\nabla h)^2$ in the growth equation: Since the left hand side of Eq. 2.20 has an odd power of h all the terms on the right hand side must as well. Thus the simplest equation consistent with the symmetries to describe the fluctuations of the interface is

$$\frac{\partial h(x,t)}{\partial t} = \nu \nabla^2 h(x,t) + \eta(x,t). \tag{2.21}$$

2.2.3.2 Edwards-Wilkinson equation (*i.e.*, linear theory)

Equation 2.21 was introduced by Edwards and Wilkinson[35]. The Laplacian term $\nu\nabla^2 h$ acts to smooth out the interface. The prefactor ν corresponds to the surface tension: the greater the surface tension, the smoother the interface. The noise term acts to roughen the interface. Scaling arguments let us exactly determine the scaling exponents of this model. Note the surface is self affine, as defined by the scaling relations defined in Eqs. 2.17 and 2.18. If we rescale the horizontal direction $x \to xb$, then the vertical direction is rescaled as $h \to hb^{\chi}$, and time is rescaled as $t \to tb^z$. We should recover a growth interface which is statistically indistinguishable from the original interface under this rescaling. Likewise the growth equation should be indistinguishable under the rescaling. With the rescaling Eq. 2.21 becomes[3]

$$b^{\chi-z}\frac{\partial h}{\partial t} = \nu b^{\chi-2}\nabla^2 h + b^{-\frac{1}{2}-\frac{z}{2}}\eta. \tag{2.22}$$

To rescale the noise term, we have used general properties for the rescaling of delta functions, $\delta^d(a\vec{x}) = a^{-d}\delta(\vec{x})$. Multiplying both sides of Eq. 2.22 by $b^{z-\chi}$ we obtain

$$\frac{\partial h}{\partial t} = \nu b^{z-2}\nabla^2 h + b^{-\frac{1}{2}+\frac{z}{2}-\chi}\eta. \tag{2.23}$$

To ensure scale invariance, each side of Eq. 2.23 must be independent of b. Thus for the Edwards-Wilkinson equation in $1+1$ dimensions we obtain the exact values of the scaling exponents[4]:

$$z = 2, \tag{2.24}$$

and

$$\chi = \frac{1}{2}. \tag{2.25}$$

[3]Note, in d dimensions, the noise term scales as $\eta(\vec{x},t) \to b^{-(d+z)/2}\eta(\vec{x},t)$.

[4]Note that this scaling argument is almost identical for d dimensional growth, where d is the dimension of the substrate. The noise term scales as given in the footnote above. The value of the exponent $z = 2$ is unchanged. And the general expression for the roughness exponent is $\chi = (2-d)/2$.

These predicted values for the scaling exponents agree with the values obtained numerically for certain microscopic growth rules, for example Random Deposition with surface diffusion[40]. In the RD model each particle drops along a single column until reaching the lowest unoccupied height in that column. To add surface relaxation the newly fallen particle is allowed to diffuse along the surface up to a finite distance l, stopping when in the column with the lowest height. This introduces a diffusive smoothing mechanism. The agreement of the scaling exponents and similarities in the relaxation mechanisms (the Laplacian smoothing) lead us to believe that the Random Deposition model with surface diffusion is in the universality class described by the Edwards-Wilkinson equation, (Eq. 2.21).

2.2.3.3 KPZ equation (i.e., nonlinear theory)

The RD model, even with surface diffusion, is quite simple. Growth always occurs in the direction perpendicular to the initial flat substrate. In contrast, models such as BD allow for lateral growth (which leaves voids in the bulk aggregate) and so growth can occur in the direction parallel to the initial flat substrate. In these types of models, the invariance under the transformation $h \to -h$ is no longer preserved (symmetry property **(v)**). Kardar, Parisi, and Zhang were the first to notice this, and they added a term of the form $(\nabla h)^2$ to the Edwards-Wilkinson Equation. Note that a term of this form was not allowed for the simplest growth equation due to symmetry considerations. With the exclusion of symmetry property **(v)**, this is the simplest nonlinear term allowed. Aside from symmetry considerations, there are physical motivations for including a term of this form, mainly that growth occurs along the local normal to the interface (*i.e.*, lateral growth). Thus accounting for lateral growth we get the KPZ equation

$$\frac{\partial h(x,t)}{\partial t} = \nu \nabla^2 h + \frac{\lambda}{2} (\nabla h)^2 + \eta(x,t). \tag{2.26}$$

As in Eq. 2.21 the Laplacian term results in smoothing of the interface, and the prefactor ν corresponds to the surface tension. The nonlinear term accounts for lateral coarsening. Again, $\eta(x,t)$ is noise which accounts for the surface roughening.

The solution of the KPZ equation is more complex than that of the linear theory. Simple scaling arguments do not allow the determination of the exponents. The scaling exponents of the KPZ equation can be obtained exactly in one dimension by noting the existence of a Fluctuation-Dissipation theorem[5]. Insights into the behavior of this equation in other dimensions comes from renormalization group treatments.

Despite the fact that the growth model studied in this chapter (BD) is conjectured to belong to the KPZ universality class, the derivation of the solution to the KPZ equation is not relevant to the subsequent discussion. Ideas of scaling are much more relevant. Hence I gave a detailed derivation of the solution to the linear theory to get the reader familiar with concepts of scaling. For the solution to the KPZ equation (including the necessary background on the renormalization group), the reader is referred to any of the excellent reviews on this topic, for example [5].

The main consequence of the KPZ equation which is important to us is the predicted values of the scaling exponents for $1+1$ dimensional growth: $\chi = 1/2$, $\beta = 1/3$, and $z = 3/2$. We take a numerical approach, using computer simulations to generate BD clusters, hence we will be able to directly compare the exponents obtained empirically with those predicted analytically.

2.3 Implementation of stochastic growth models

Stochastic models of surface growth are typically investigated with computer simulations. Randomness is replaced by use of deterministic Pseudo Random Number Generators (PRNGs). The remainder of Part I of the thesis focuses on implementations of BD, and explores the implications of the use of PRNGs. The next two chapters are reprinted from journal articles. There is limited repetition in the intro-

duction section of each chapter with the review in this chapter.

Chap. 3 reports on a study of basic scaling properties of the BD model. Several anomalous results for the properties of the growth interface are obtained. During a thorough search for corrections to scaling, coupling between the BD model and the PRNG used in the implementation becomes apparent. Implementations with different PRNGs yield statistically discrepant results. This work also appears as Ref. [28].

Chap. 4 reports on an extensive, statistically rigorous study of BD. The purpose is to quantify the coupling between the BD model and PRNGs. We show a breakdown of basic statistical sampling assumptions, and of the ergodic exploration of phase space. This work also appears as Ref. [29].

Chapter 3

Anomalies in simulations of Ballistic Deposition[1]

Ballistic Deposition (BD) is a prototypical model for interface growth and for dynamic scaling behavior in non-equilibrium systems. BD is typically investigated with computer simulations where randomness is replaced by use of deterministic Pseudo Random Number Generators (PRNGs). In this study of BD, several results discrepant with the prevailing paradigm were observed. First, the value of the roughness exponent, χ, obtained is below the value for a random walk (i.e. $\chi < 1/2$). The value $\chi = 1/2$ is predicted by the KPZ equation, and many models of growth obtain this exponent. Second, height fluctuations of the growing interface appear not to satisfy simple scaling. Third, a decrease in the surface roughness is observed in a conjectured steady state regime. Computer implementations of BD may be responsible for the discrepancies. A coupling between the BD algorithm and a PRNG algorithm is identified, and statistically discrepant results are obtained for an implementation with a different PRNG.

[1]This work is reprinted from J. Mod. Phys. C **8** (4), 941–951 (1997), with references to other chapters added.

3.1 Overview

Understanding growth patterns, both of clusters and solidification fronts, has become an increasingly interesting problem, relevant to non-equilibrium processes in general[42, 53, 91, 5, 56]. Dynamic scaling characterizes many of these processes, and Ballistic Deposition[125] (BD) is a prototypical model for this class of system, and for interface growth in general. In the BD model free particles, following ballistic trajectories, encounter the active growth interface of the substrate at which point they aggregate. The resulting growth patterns are compact clusters with a rough surface, which may be an accurate model of thin film growth. The active growth interface exhibits dynamic scaling behavior.

BD has been investigated for more than a decade and the findings have defined pathways for countless successive studies of growth processes[91, 124, 56]. However it is disconcerting that a true consensus has never been reached on the value of the roughness exponent, χ. Several anomalies were observed in the implementation of BD reported in this study, namely a growth surface not described by simple scaling, a roughness exponent below the theoretically predicted value, and a decrease in the surface roughness in a conjectured steady state regime. Coupling to PRNGs is observed, which may resolve the anomalies and discrepancies in past work, and challenges the implicit assumption that randomness can be replaced by deterministic Pseudo Random Number Generators (PRNGs).

The organization of this chapter is as follows. A brief review of the values for the scaling exponents obtained in past studies will be presented. The implementation details and results are then discussed. Pathological situations which may cause long crossover times or a change in exponents are investigated. For example, long range spatial and/or temporal correlations in the pseudorandom sequence of numbers employed are sought for, but no evidence of such two-point correlations is found. Yet a dynamics in the value of the surface roughness is observed in a conjectured steady state regime, and a different dynamics is observed for an implementation with a dif-

ferent PRNG. We thus conclude that the BD model couples sensitively to, as yet undertermined, non-randomness of pseudorandom sequences, even sequences which pass all standard statistical tests.

3.2 Review of past results

As discussed in Sec. 2.2.1.3, the surface configuration for BD is completely described by the height, $h(x,t)$, along each position, x, of an underlying substrate, with t denoting the time duration of growth. Throughout this chapter growth on a one-dimensional substrate (of length L) is considered, for which exact theoretical predictions exist. The surface evolves as follows. A column along the substrate is chosen at random and a particle is added to the surface of that column at the height:

$$h(x,t'+1) = max[h(x-1,t'), h(x,t')+1, h(x+1,t')]. \tag{3.1}$$

(The deposited particle occupies the highest empty site with one or more occupied nearest neighbor sites.) Here t' is the number of individual deposition events and is proportional to t ($t = t'/L$). Beginning with an initially flat substrate, the width of the active growth interface, $\xi(L,t)$, increases from zero to an asymptotic value which depends on the underlying, finite size, length scale, L. A measure of $\xi(L,t)$ is the standard deviation of the surface heights, $\{h(x,t)\}$,

$$\xi^2(L,t) = \frac{1}{L}\sum_{x=1}^{L}(h(x,t) - \overline{h(t)})^2, \tag{3.2}$$

where $\overline{h(t)}$ is the mean height of the surface at time t.

It was pointed out by Family and Vicsek[41] that the scaling forms for the growth and saturation of the width of the interface can be described by a scaling ansatz,

similar to that applicable to critical systems;

$$\xi(L,t) = L^{\chi} f(t/L^{z}), \tag{3.3}$$

where $f(x) \sim x^{\beta}$ for $x \ll 1$ and $f(x) = const$ for $x \gg 1$. For short times, the width of the interface should increase as $\xi \sim t^{\beta}$. In the asymptotic regime the width of the interface should scale as $\xi \sim L^{\chi}$.

The analytical theory by Kardar, Parisi and Zhang (KPZ) describes the evolution of fluctuations on growing surfaces using a symmetry motivated differential equation[66]:

$$\frac{\partial h}{\partial t} = \nu \nabla^2 h + \frac{\lambda}{2}(\nabla h)^2 + \eta(\mathbf{x}, t). \tag{3.4}$$

Here, ν is related to surface relaxations, $(\lambda/2)(\nabla h)^2$ is introduced to account for lateral coarsening, and $\eta(\mathbf{x}, t)$ is white noise. In $1+1$ dimensions, the values of the scaling exponents can be obtained exactly from KPZ theory: $\chi = 1/2$ and $\beta = 1/3$. Note that $\xi(L,t) \sim L^{1/2}$ is equivalent to a random walk. While these exponents agree very well with simulations of several growth models (e.g. Restricted Solid on Solid (RSOS)[71], discussed further in the conclusions), it is puzzling that the reported values of the roughness exponent, χ, obtained for the BD model have all been less than the theoretically predicted value. Moreover the reported values of χ have a substantial range (greater than 15% of the lowest value), (for a summary see Table 3.1).

3.3 Implementation and results

In this study of BD, initially flat substrates of length varying from $L = 127$ to $L = 2047$ are considered. At each update a pseudo random number (PRN) is generated indicating which site along the substrate will have the next event. A particle is added to that column at a height described by Eq. 3.1. The source for PRNs is a C library

Table 3.1: A table of scaling exponents for two growth models, BD and RSOS, as determined by selected numerical investigations. Error bars are included when available.

Scaling Exponents from Selected Numerical Simulations			
Model	Reference	β	χ
BD	FV85[41]	0.30 ± 0.02	0.42 ± 0.03
	MRSB86[92]	0.331 ± 0.006	0.47
	M93[91]	0.33	0.45
	HHZ95[56]	0.31	not reported
	This study[28]	0.31 ± 0.02	χ_{loc} 0.42 ± 0.02 χ_{glob} 0.455 ± 0.015
RSOS	KK89[71]	0.332 ± 0.005	0.50
KPZ Theory	KPZ86[66]	1/3	1/2

subroutine "random()"[51], which is a non-linear additive feedback PRNG, initialized with a 256 bit seed, giving a repeat period of 2^{64} numbers. The lowest 20 to 24 bits of each returned number were shifted off, leaving only the highest (and least correlated) bits. Many PRNGs were investigated before this variant of random() was chosen, as it performed best in initial tests. All of the simulations were carried out on a desktop workstation, with the runs on the shortest length substrates requiring a few hours, and the runs on the longest, a few days.

The dynamic scaling exponent, β, is determined using a plot of the width of the interface, $\xi(L,t)$, versus time, for all of the substrate lengths L. Note one unit of time is L particle depositions. Consistent with previous studies (see Table 3.1) the results are $\beta = 0.29 \pm 0.01$ for times $3 \leq t < 100$, and $\beta = 0.31 \pm 0.02$ for times $100 \leq t < 2000$. We chose to report the value for the longer times, as the scaling interval is of greater absolute size.

Determination of the asymptotic roughness exponent χ is not as straightforward. The primary complication is that $\xi \sim L^{\chi}$ only in the regime where the correlation length has reached the full system size ($t \gg L^{z}$), which can be computationally prohibitive for large substrate lengths. The conservative estimate of time for "relaxation", $\tau = 10L^{z=\frac{5}{3}}$ is employed (which exceeds the expected scaling of $z = 3/2$).

This time corresponds to an average surface height $\bar{h} = 20L^{\frac{5}{3}}$.

Once in this regime, the local roughness of the surface is investigated by studying "windows" of length $l < L$ (i.e $\xi(l,t)$ for $l < L$). Finite size effects (imposed by periodic boundary conditions) cause a rounding of these curves, restricting the scaling regime to lengths $l < L/2$, where self-affine scaling suggests $\xi(l,t) \sim l^{\chi}$. The BD local surface roughness is compared to that of a random walk. We calculate the expectation value of the local surface roughness of a random walk over the ensemble of all possible walks of total length L:

$$\left\langle \xi(l,t)^2_{rwalk} \right\rangle = \left\langle \frac{1}{l} \sum_{x=1}^{l} \left(h(x) - \overline{h_l} \right)^2 \right\rangle . \tag{3.5}$$

Here $h(x+1) = h(x) + \delta_x$, $\delta_x = \pm 1$, $\sum_{x=1}^{(L-1)} \delta_x = 0$ (which implements periodic boundary conditions), and the correlations between $\delta'_x s$ are assumed uniform. This yields the scaling relation

$$\left\langle \xi(l,t)^2_{rwalk} \right\rangle \sim l \left(1 - \frac{l}{2L} \right), \tag{3.6}$$

for $l, L \gg 1$. The second term in Eq. 3.6 represents a finite size correction to scaling. Comparison with our results for BD is shown in Fig. 3-1, where the prefactor for the random walk in Eq. 3.6 was selected so as to agree with the BD simulations at the shortest lengths. The two curves are very distinct; Both appear to exhibit scaling behavior over more than one decade, however the scaling exponents differ. If the $\delta'_x s$ are chosen in agreement with the empirically obtained step height distribution of the BD simulations, no *ad hoc* prefactor is needed and similar results are obtained.

From curves of $\xi(l,t)$ in the asymptotic regime, three distinct scaling behaviors can be identified. For $3 \leq l < 20$ the relation $\xi(l,t) \sim l^{\chi_{loc_0}}$ is obtained, with $\chi_{loc_0} = 0.35 \pm 0.01$. For $30 < l < 400$ the relation $\xi(l,t) \sim l^{\chi_{loc}}$ is obtained, with $\chi_{loc} = 0.42 \pm 0.02$. Looking only at the longest length from each substrate the relation $\xi(L,t) \sim L^{\chi_{glob}}$ is obtained, with $\chi_{glob} = 0.455 \pm 0.015$. The data for all substrate

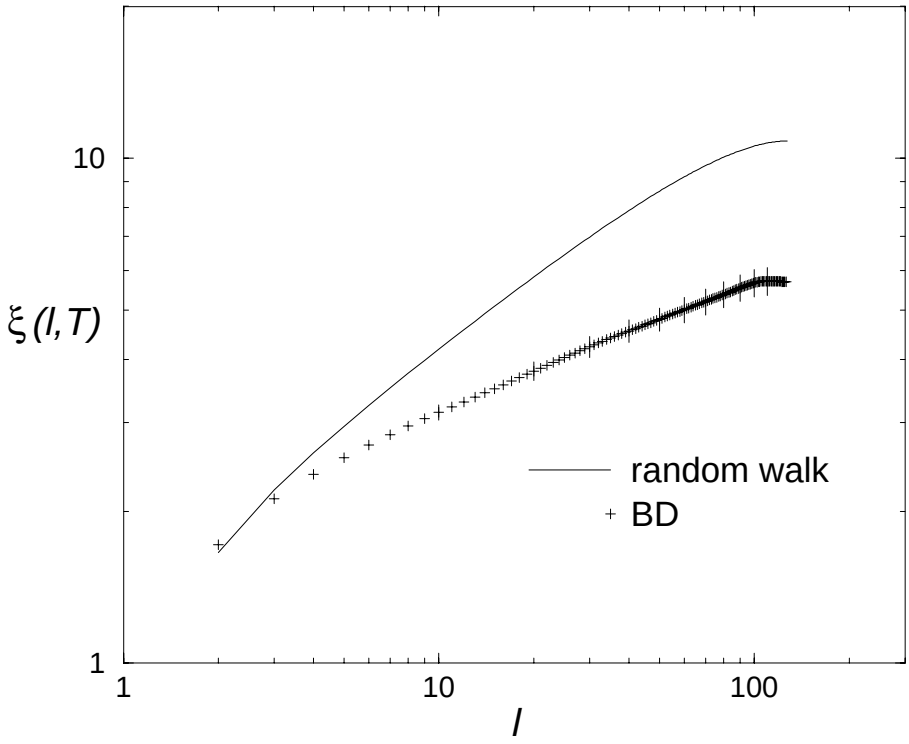

Figure 3-1: The width of the growth interface plotted as a function of increasing window size, for the BD simulations and a theoretical calculation of a random walk with periodic boundary conditions.

lengths is shown in Fig. 3-2, along with a comparison to the equivalent data for a random walk with periodic boundary conditions. For the random walk, only one scaling relation can be extracted $\xi(l,t) = l^{1/2}$, for both $l \leq L/2$ and $l = L$.

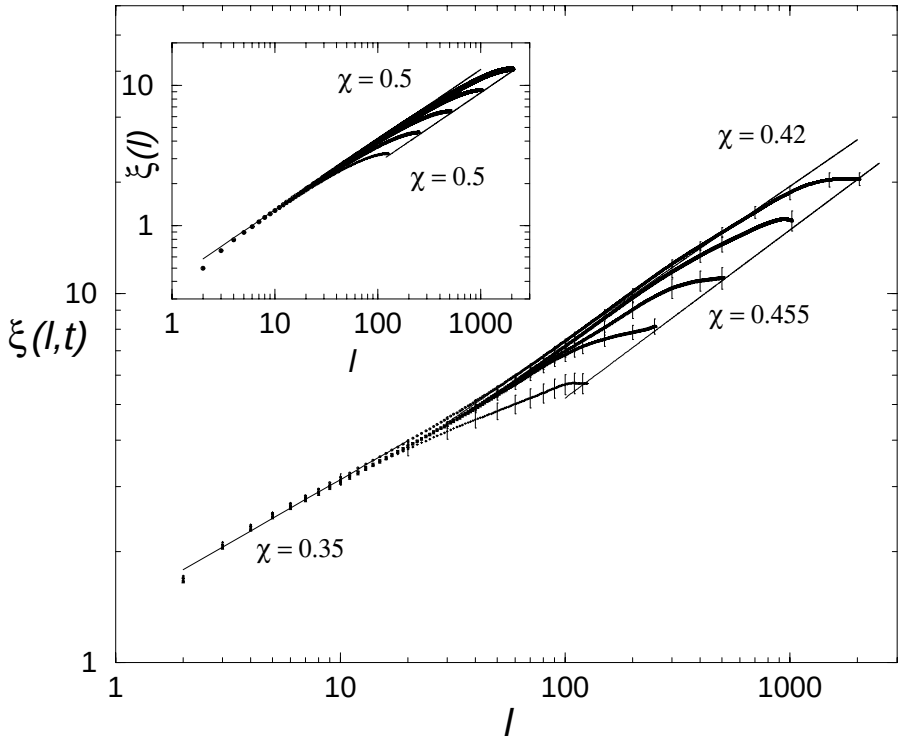

Figure 3-2: The width of the growth interface plotted as a function of increasing window size for the longest times simulated. The data from all five substrate lengths are included. Dependent upon the range of lengths we examine, three apparent scaling exponents may be obtained, as indicated by the solid lines. The insert is the corresponding data for a random walk with periodic boundary conditions.

In an attempt to formulate a consistent scaling picture, we collapse the data from the five different length substrates, using Eq. 3.3 (see Fig. 3-3). For $\xi(L,t)$ collapse is achieved by plotting $\xi(L,t)/L^{\chi_{glob}}$ *vs* t/L^z, with the exponents $\chi_{glob} = 0.45 \pm 0.02$ and $z = 1.45 \pm 0.03$. For $\xi(l,t)$, with $l = 0.1L$, collapse is achieved by plotting $\xi(l,t)/l^{\chi_{loc}}$ *vs* t/l^z, with the exponents $\chi_{loc} = 0.40 \pm 0.02$ and $z = 1.41 \pm 0.03$.

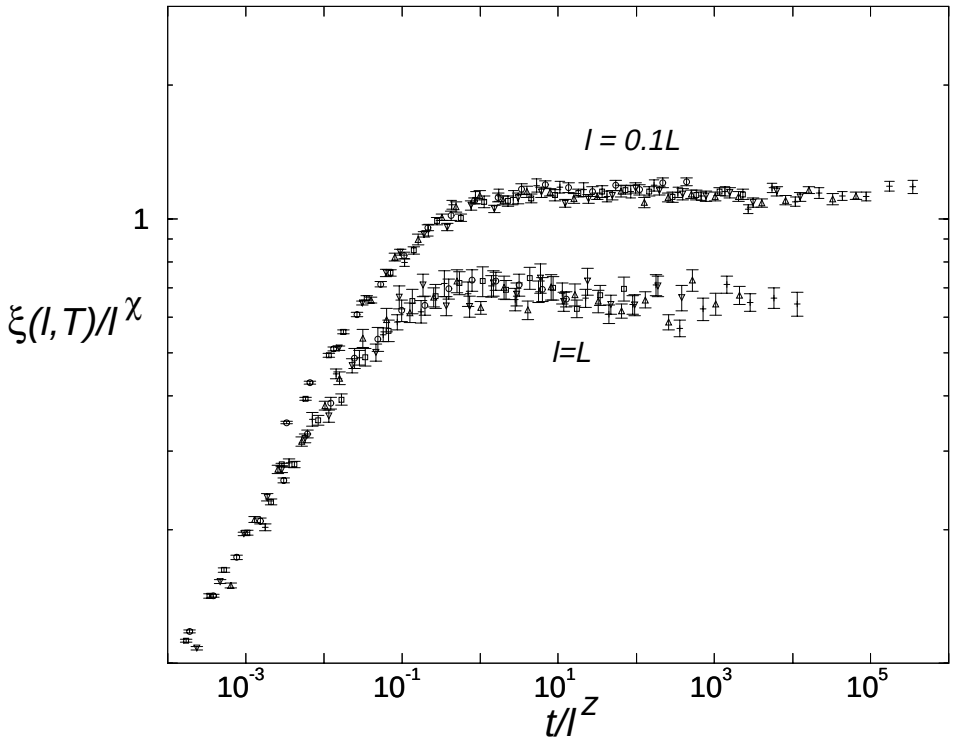

Figure 3-3: The scaling function, $f(\frac{t}{l^z})$, determined for data sampled at $l = 0.1L$ and at $l = L$. For the data sampled at $l = 0.1L$ collapse is achieved with the exponents $\chi_{loc} = 0.40 \pm 0.02$ and $z = 1.41 \pm 0.03$. For data sampled at $l = L$ collapse is achieved with $\chi_{glob} = 0.45 \pm 0.02$ and $z = 1.45 \pm 0.03$

Exponents obtained by the collapse are consistent with those obtained by a linear fit and with those reported in previous studies (see Table 3.1). The measurement of χ_{loc} agrees with the FV85 measurement of χ[41]. The measurement of χ_{glob} is consistent with the M93 and the MRSB86 measurements of χ[91, 92],. However, these results are puzzling on three counts. First $\chi_{loc} \neq \chi_{glob}$, *i.e.*, the surface is not self affine. Second, all of the exponents obtained are below the value predicted by the KPZ theory. Third, the characteristic identity $\chi + z = 2$ is *not* obeyed. These are not only violations of the KPZ scaling, but also most of its extensions (e.g. Medina, *et al.*[93]). Thus we are lead to ask, what kind of continuum model could explain a decrease in the values of the scaling exponents (i.e. a hypo-rough surface)?

3.4 Possible resolutions

One possibility is that the model suffers from a long crossover regime. For example the Wolf-Villain model of growth has a length dependence in the adjacent step height distribution, which persists for long times ($t \sim 10^4$), and long lengths ($L \sim 256$)[108]. The step height distribution for the BD model is a well behaved quantity, which reaches a steady state, length-independent value within the completion of a few monolayers. With reference to the work on crossover, it has been suggested that the structure factor is the only way of accurately determining exponents[110]. The structure factor is related to the Fourier transform of the height-height correlation function

$$S(\mathbf{k}, t) = \left\langle \hat{h}(\mathbf{k}, t)\hat{h}(-\mathbf{k}, t)) \right\rangle , \tag{3.7}$$

with $\hat{h}(\mathbf{k}, t) = L^{-d/2} \sum_{\mathbf{x}} [h(\mathbf{x}, t) - \overline{h(t)}] e^{i\mathbf{k}\cdot\mathbf{x}}$. Crossover behavior is manifested by a change in the scaling exponent of $S(\mathbf{k}, t)$ *vs* $\mathbf{k}$. There is no evidence of crossover in $S(\mathbf{k}, t)$ for our data, including the longest lengths and longest times simulated. The exponent obtained is consistent with $\chi_{glob} = 0.45$.

A second possibility is an intrinsic width correction[136], as introduced by Kertesz

and Wolf to the Eden Model. This correction accounts for voids in the bulk, and allows for a clear scaling regime to be obtained for the Eden model in 2+1 and 3+1 dimensions[69]. Introducing such a correction to our data destroys the scaling regime.

A third possibility is correlated noise. There has been extensive past theoretical and numerical work into growth models with positively correlated noise (Medina, *et al.*, is the first study)[93, 91, 56]. For noise with long-range correlations in space and/or time, the noise term in Eq. 3.4 is of the form $\langle \eta(\mathbf{x}, t) \rangle = 0$, with

$$\langle \eta(\mathbf{k}, \omega) \eta(\mathbf{k}', \omega) \rangle = 2D(k, \omega) \delta^{d-1}(\mathbf{k} + \mathbf{k}') \delta(\omega + \omega'), \tag{3.8}$$

where the noise spectrum $D(k, \omega)$ has power-law singularities of the form $D(k, \omega) \sim |\mathbf{k}|^{-2\rho} \omega^{-2\theta}$. For uncorrelated noise, the noise spectrum is a constant.

As shown in the original study, for spatially correlated noise (i.e $\rho > 0, \theta = 0$), χ increases and z decreases, preserving the identity of Galilean invariance, $\chi + z = 2$[93]. An extension to anti-correlated noise (i.e. $(\rho, \theta) < 0$) may show a decrease in the values of the scaling exponents and a breaking of that identity. It is clear that extreme forms of anti-correlated noise will result in a flat interface; Consider sequential updating of a BD growth algorithm, this generates a perfectly flat interface at each timestep. Hence it is reasonable that correlations in the noise may change the roughness of the surface.

We have conducted extensive tests to identify potential long-range spatial and/or temporal correlations in the PRNs generated by the function call `random()`. The PRNs produced pass every mathematical and physically motivated test, including those directly relevant to the growth algorithm. The Fourier transform of the noise in each column, and along the columns of the substrate, produce flat, white-noise spectra. The number of calls to each column are Poisson distributed (as discussed below). Waiting times between successive calls to each column are exponentially distributed. The simulated space was divided into checkerboard sublattices, but no

discrepancy between events on even or odd numbered sites was found. The autocorrelation function for every property of the generated surface tested is a simple decaying exponential. There is no bias for relevant successive events to be in the same column or left or right neighboring columns. Other researchers have recently pointed out some physical models which manifest the pathologies of certain PRNGs[122]. The PRNs produced by random() pass even those sensitive tests.

Nonetheless, a systematic decrease in the width of the BD growth interface is observed in the "asymptotic" regime (Fig. 3-4). For the longer lengths simulated, the decrease occurs even before the strictly defined asymptotic regime. In order to bound the decrease outside of statistical error, 200 independent runs on substrates of lengths $L = 127$ and $L = 511$ are studied. Note that all the data reported so far were obtained with 20 independent runs on each substrate length, and all error bars reported are statistical. A plot of $\xi(l,t)$ *vs* l, during the four greatest times, on the $L = 127$ substrate, is shown in Fig. 3-4. As the time increases, $\xi(l,t)$ decreases systematically. The error bars should be noted; The value of $\xi(l,t)$ for the longest time is over four standard errors away from the value for the shortest time. This effect may be due to anti-correlation in the system (where system refers to the BD algorithm coupled to the PRNG algorithm). The effect of anti-correlation in the PRNG on a growth algorithm is illustrated by a simple Poisson process, the Random Deposition (RD) model. Implementing RD also serves as a check of the BD computer code.

In the RD model the site filled by deposition is the next available height in the active column (not even nearest neighbor interactions exist). No boundary conditions are needed and hence there is no dependence on L. The surface height values should theoretically be Poisson distributed: ($\xi \sim t^{1/2}$, for all t). Fig. 3-5 is of $\xi(t)$ *vs* t for an RD algorithm implemented by altering only one line of our BD code (that line describing Eq. 3.1). Two different PRNGs were used: random() (described earlier), and rand(), a 16 bit version of the standard C-library subroutine[52]. At the close of

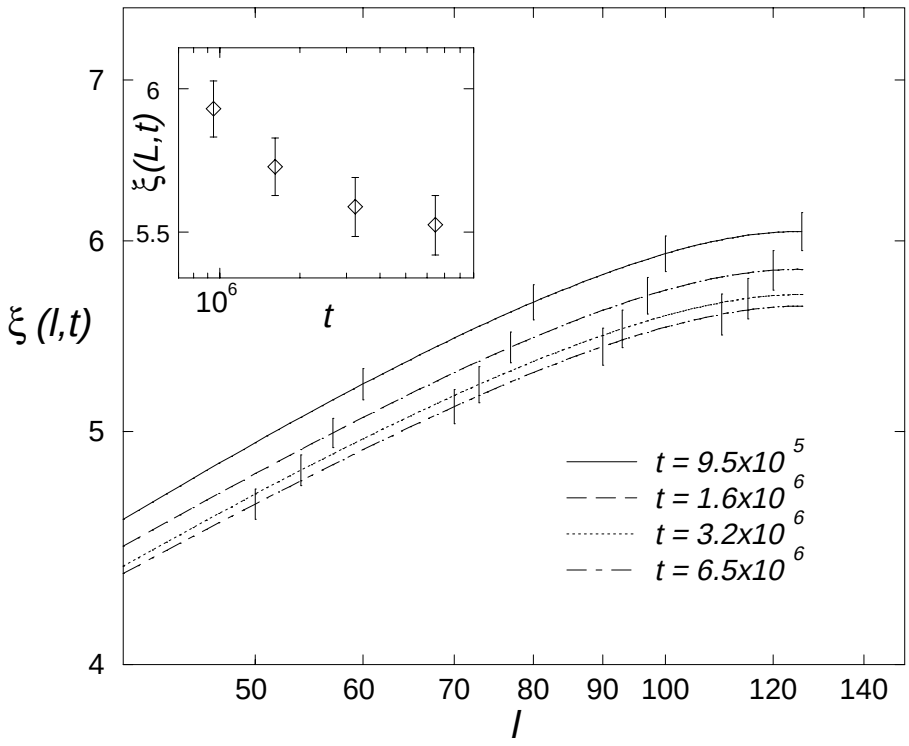

Figure 3-4: The width of the BD growth interface versus increasing window size, plotted for different times. It can be seen that as time increases, the width of the interface decreases systematically. The inset plot is the width of the interface for the largest window size, $l = L$, plotted with increasing time. Note the error bars included on all points.

one repeat cycle, rand() has sampled all numbers evenly, yielding a flat surface for RD. The results using random() agree with theory out to the longest times simulated, yet a systematic decrease is observed using rand(), as it is for the BD model using random(). As mentioned, the repeat period of random() is on the order of 10^{19}, almost ten orders of magnitude greater than the total number of calls made to it.

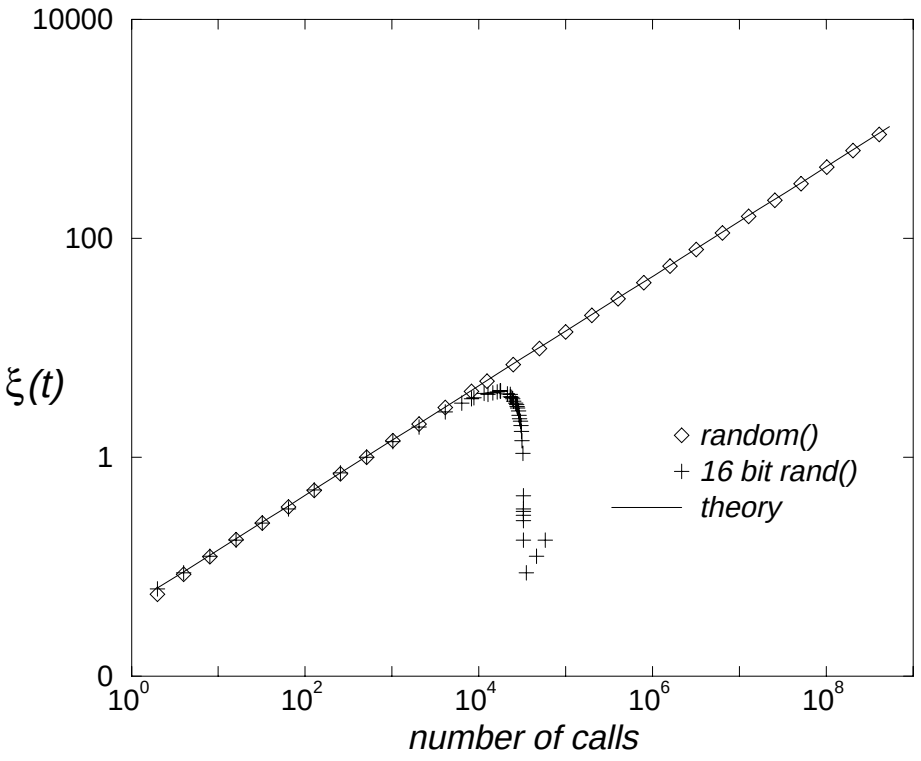

Figure 3-5: The width of the growth interface versus the number of function calls for simulations of RD using two different PRNGs. The solid line corresponds to the result for ideal random numbers, $\xi \propto t^{1/2}$. The cycle length of rand() is apparent since the interface width decreases to zero at the end of a cycle.

To quantify that the decrease of the surface roughness is a result of evolving the system with random() the full BD simulations were run, using the identical code, but a different PRNG, ran2() (which combines two distinct types of PRNGs in an attempt to eliminate correlations inherent to each one separately)[101]. The systematic decrease

was not observed, but several unsystematic fluctuations in the value of $\xi(l,t)$, outside of statistical error bounds, were observed. Selecting individual time samples for each length substrate, it is possible to construct a scaling curve, with the scaling exponents $\chi_{loc} = 0.45 \pm 0.01$, and $\chi_{glob} = 0.51 \pm 0.02$. Note that these are statistical error bounds, which do not address the issue of fluctuations. Constructing the curve with statistical outliers, changes the values of the exponents. We do not assert that these exponents are the "true exponents", but instead wish to focus on the fact that different PRNGs yield different results for an identical system, implemented with the identical code.

The more revealing comparison is between average values obtained by the two different PRNGs. For example, the average asymptotic values of $\xi(L,t)$ differ by more than three standard errors. A "t-test" comparing these average values fails at the 99% confidence level[43]. Statistically discrepant results for average values and a statistically distinct dynamics for the interface width fluctuations show a breakdown of basic sampling assumptions, and moreover, that the observed dynamics is not inherent to the BD model. Assuming the dynamics does not reside in the PRNGs, it must reside in a coupling between the BD and PRNG algorithms. This is in line with the observation that the short time scaling exponent (β) is not effected, but the asymptotic scaling exponent (χ) is; The effects of coupling may take some time to accumulate. A detailed statistical analysis quantifying the breakdown of sampling assumptions and ergodic exploration of phase space is reported in Chap. 4.

3.5 Conclusions and discussion

It is reasonable that the BD algorithm is more sensitive to correlations in PRNs than standard Monte Carlo (MC) algorithms. In standard MC, comparison to the Boltzmann probability causes rejection of some PRNs produced. In BD, all PRNs are used (in the sequence produced). It should be noted that in restricted models of growth (where physical constraints cause rejection of PRNs), such as RSOS, the theoretically

predicted scaling exponents are recovered with great precision in numerical simulations. It may be interesting to implement the BD algorithm with random rejection of PRNs. In addition, we have been considering the use of massive physical simulations (run on a special purpose cellular automata machine[87]) which have a vast amount of initial state, as a source of randomness.

BD is a sensitive physical test for correlations present in pseudorandom sequences, and it would be desirable to identify the exact nature of the correlations detected. Discrepancies in the reported values for the roughness exponent χ, and anomalies found in this study, may be attributed to distinct couplings between the BD and the PRNG algorithms. Disagreement in results generated by two different PRNGs is strong evidence for this and, more importantly, shows the dynamics observed in the asymptotic regime are not inherent to the BD model itself. Results from previous studies of BD have not indicated which PRNG was used, in addition many of these past simulations utilized power of two substrate lengths ($L = 2^n$), a system size for which PRNGs manifest their greatest pathologies[73]. As of yet we are not able to identify the universality class of the BD model. It appears that coupling to the PRNG algorithms becomes dominant before the steady state regime is achieved.

Chapter 4

Sensitivity of Ballistic Deposition to pseudorandom number generators[1]

Ballistic deposition (BD) serves as a prototype for studies of dynamic scaling phenomena in nonequilibrium growth processes. In BD, particles are sequentially added to a growing surface at randomly selected positions. The model is typically investigated by computer simulations where randomness is implemented by pseudorandom number generators (PRNGs). The implicit assumption that PRNGs adequately represent true randomness is tested in this study via a statistical analysis of the width of the BD interface. We study the width of the interface over time scales orders of magnitude longer than the expected model relaxation time, yet much smaller than the period of the PRNG, and observe fluctuations which still appear to be correlated. Distinct dynamic behavior is observed for an implementation with a different PRNG, further indicating a strong coupling between the model and the PRNGs (even with PRNGs that pass extensive statistical tests). Thus we demonstrate a breakdown of

[1]This work is reprinted from Phys. Rev. E **57** (5), 5044–5052 (1998), with references to other chapters added. My coauthors are Y. Bar-Yam and M. Kardar.

basic sampling assumptions, and of the ergodic exploration of phase space.

4.1 Introduction

Dynamic scaling phenomena in stochastic nonequilibrium systems have attracted increasing attention in recent years[90]. In theoretical models of open systems, the external influences are usually represented by random noise. In computer simulations of the models pseudorandom numbers (PRNs) implement the stochastic process. The use of deterministic PRN algorithms necessarily introduces some degree of correlation in the produced sequence of PRNs[83, 3]. While these correlations are probably irrelevant in most applications, they may in principle couple to the underlying dynamics of the simulated model, resulting in artificial behaviors.

A good example of open nonequilibrium behavior is provided by growth of aggregates through random addition of particles. The width of the resulting interface exhibits dynamic fluctuations, which have been studied extensively, analytically and numerically[42, 56, 5, 75]. A particularly simple model for this phenomena is Ballistic Deposition (BD)[125]: Particles are randomly placed above an aggregate growing on a substrate, they descend along a straight vertical path until they encounter a site on the existing cluster and stick there. The random placements of subsequent particles represents a stochastic process.

We present results of a numerical study of BD in which the potential coupling to PRNs is examined via statistical tests on the width of the growing interface. A dynamic change in a conjectured steady-state regime is observed, signaling a breakdown of ergodicity. The breakdown is quantified by demonstration of violations of the basic sampling assumptions. Statistically relevant inconsistencies occur repeatedly in the data. First, fluctuations statistically inconsistent with the steady-state distribution are observed (even when the time scale of observation is orders of magnitude greater than the expected relaxation time of the model). Second, values of

steady-state quantities averaged over different sets of initial seeds are in statistical disagreement. Third, average values of steady-state quantities, obtained by using two distinct pseudorandom number generators (PRNGs), are in statistical disagreement. These results lead to the conclusion that the observed dynamical fluctuations are not inherent to BD, but result from a coupling to the PRNG algorithms.

The manuscript is organized as follows. Section 4.2 provides the algorithmic details of the BD model and its implementation, with emphasis on the role played by the PRNs. In Sec. 4.3 probability distributions, sampling assumptions, and the construction of the statistical tests are discussed. The statistical inconsistencies of the numerical data with the basic sampling and ergodicity assumptions are presented in Sec. 4.4. Potential implications for the use of PRNs, particularly in the context of growth models, are discussed in the concluding Sec. 4.5.

4.2 Model and implementation

Given that the focus of this work is on the unwanted coupling of PRNGs with the underlying model, it is necessary to provide details of the numerical simulation in greater depth than usual. Here we shall review the BD algorithm, and its numerical implementation with specific PRNGs.

4.2.1 Ballistic Deposition

In the BD model of growth, free particles initiated at random positions above a one-dimensional substrate descend ballistically and stick upon first touching the surface of the growing cluster. The substrate of length L, consists of discrete columns indexed by integer values x, with $1 \leq x \leq L$. The growth interface is defined by the maximum occupied site along each column, $h(x,t)$, where $h(x,t)$ also takes on discrete integer values. Starting from a flat interface, $h(x,t=0) = 0$ for all x, the surface evolves by sequential addition of particles to randomly chosen columns. The index number

of particles deposited is denoted by t', and the deposition time by $t = t'/L$. Each deposition event consists of choosing a column, $x(t')$, by a call to a PRNG, and updating the height in that column as follows:

$$h(x(t'), t' + 1) = \max[h(x(t') - 1, t'), h(x(t'), t') + 1, h(x(t') + 1, t')]. \tag{4.1}$$

Thus the deposited particle occupies the highest empty site with one or more occupied nearest neighbor sites; this mimics the process of cluster aggregation. The stochastic process in this model is the random choosing of successive columns.

While the resulting aggregates are compact, their interface is rough, with fluctuations that are expected to be self-similar at all scales. The width of the growth interface $\xi_L(t)$, on average increases following a power law behavior until reaching a steady asymptotic value, the magnitude of which depends on the underlying substrate size L. A good measure of $\xi_L(t)$ is the variance of the surface heights, $\{h(x,t)\}$,

$$\xi_L^2(t) = \frac{1}{L} \sum_{x=1}^{L} \left(h(x,t) - \overline{h(t)} \right)^2, \tag{4.2}$$

where $\overline{h(t)}$ is the mean height of the surface at time t.

It was originally pointed out by Family and Vicsek[41] that the scaling forms for the growth and saturation of the width of the growing interface can be described by a dynamic scaling ansatz, similar to that applicable to critical systems. Kardar, Parisi and Zhang (KPZ)[66] introduced an analytic theory describing the evolution of fluctuations on growing surfaces, which has been successfully applied to several growth models. One consequence of KPZ theory is that the steady-state behavior for the interface fluctuations in one dimension should resemble a random walk; i.e. $\xi_L(t \to \infty) \propto L^{1/2}$.

4.2.2 Algorithmic details

Substrates of lengths L = 127, 255, 511, 1023, 2047, and 10007 are considered. At each update, a PRN is generated corresponding to a column along the substrate. A particle is added to that column at a height described by Eq.(4.1). Periodic boundary conditions are applied. The only subtlety is in mapping the PRN uniformly to a value between 0 and $L-1$, which is achieved as follows: The least significant bits of the PRN are shifted off, leaving a number between 0 and 2^n-1, with n chosen such that 2^n is the integer closest to, but greater than L. The PRN is rejected if it falls in the interval between L and 2^n-1. Variants on this scheme were tested (including use of all the bits of the PRN), yet similar results were obtained. The scheme described above was chosen because the algorithms that generate subsequent PRNs involve algebraic operations which cause carries from lower order bits to higher order bits; thus the higher order bits are influenced by two sources (the algebraic operation and the carries) and are expected to be less correlated.

As the system evolves, the variance of the surface heights, ξ_L^2, is calculated at selected times. Statistical errors are considered[45], and the associated standard error, σ_{ξ^2}, is also recorded. The values of the surface heights (contained in a 1-dimensional array of length L) are the only essential data for these calculations.

The onset of the asymptotic regime is estimated by the conservative criterion of a "relaxation" time $\tau \sim 10L^{z'}$, with $z' = 1.6$. This estimate is conservative since the exponent employed exceeds the expected value of the dynamic scaling exponent, $z = 3/2$, and this time is well beyond a qualitative judgment of the time required for saturation of ξ_L^2. Accounting for the measured growth rate, this relaxation time corresponds to an average surface height, $\overline{h(\tau)} = 20L^{z'}$. Exploring the asymptotic regime requires extensive computer power for the larger lengths investigated. All simulations were implemented on desktop workstations, with the shortest length systems requiring a few hours of run time, the longest requiring on the order of five days.

4.2.3 Randomness and PRNGs

PRNGs are algorithms for deterministically generating a string of bits, resembling a completely uncorrelated, and hence "random" string. Knowing the past and present values should give no information as to future outcomes of a truly random variable. Hence "deterministic randomness" is inherently unattainable. PRNGs are at best a practical substitute, and should be generally tested for the absence of undesired correlations. While two point correlations can be readily examined, there are a multitude of other subtle effects that are not in practice possible to measure. When considering which tests for correlations to conduct, it is advisable to include both the standard statistical tests, as well as physically motivated ones directly related to the particular model being implemented (see the example in the final paragraph of this section). It is also necessary to verify that any observed dynamic behavior is inherent in the simulated model, and not artificially introduced to the system by the PRNGs. Several physical models have been shown to couple to correlations in PRNGs[44, 122, 109]; we shall provide evidence that BD also belongs to this category.

Preliminary simulations with a simple PRNG resulted in various anomalies which will not be discussed in detail here, but one example is the occurrence of repeated patterns of surface configurations. Having identified the PRNG as the likely culprit, we decided to use more sophisticated PRNGs. Several were tested, and two were selected: random(), a Feedback Shift Register, employing a primitive trinomial of degree 64[51] (a "C-library" subroutine call); and ran2() which combines pseudorandomness produced by two distinct multiplicative congruential generators, and has been shown to reduce certain serial correlations inherent to each generator separately[101]. The developers of ran2() are confident enough to have offered $1000 to anyone who proves a deficiency in their PRNG!

Extensive tests for correlations in the PRN sequences were conducted, with emphasis on tests directly relevant to the BD growth algorithm. One such test is for any bias in the next relevant call. If the next growth site in a neighborhood is biased to the

left- or right-neighboring column of the last added particle, a thiner growth interface results. If it is biased towards the same column, a wider interface is generated. No discrepancies were found with any of the tests, a brief summary of which is included here: *(1)* Spatial and temporal Fourier transforms of the sequence of PRNs were consistent with white noise power spectrums, suggesting no two-point correlations. *(2)* Measurement of the number of calls to each column, and waiting times between subsequent calls to the same column, were consistent with Poisson statistics. *(3)* No spatial bias for the subsequent calls, as discussed above, was detected. *(4)* No bias was detected when the space was partitioned into sublattices. *(5)* Autocorrelation functions for natural surface height observables were decaying simple exponentials. The reader is referred to a previous manuscript for details.[28]

4.3 Probability distributions

We first obtain the unique steady-state probability distribution for $\xi_L^2(t > \tau)$. Once this distribution is known, each independent measurement of $\xi_L^2(t > \tau)$ can be considered an independent, identically distributed (iid) random variable, drawn with the associated probability. Using only this general assumption of unbiased sampling, we construct statistical tests which show that: implementations which are identical, except for use of different PRNGs, result in different values for average quantities of the growth interface; implementations with the PRNG studied most extensively in this work lack steady-state behavior; data obtained by averaging over several independent implementations are inconsistent with the underlying distribution.

The tests focus on the width of the growth interface in the steady-state regime $\xi_L^2(t > \tau)$. We shall discard the time argument in favor of a compact notation, and henceforth denote this variable by ξ_L^2. When it is necessary to deal with shorter times, the explicit time argument is included. The discussion also focuses on the $L = 127$ system size. We were able to explore the asymptotic regime for orders of

magnitude beyond the conjectured model relaxation time only for this shortest length investigated due to practical limits on computational resources. For the $L = 127$ system, adequate statistics could be obtained for times as large as $t = 1000\tau$.

4.3.1 The steady-state distribution

There is a unique, steady-state distribution, for the overall width of the growth interface, $P(\xi_L^2)$. It is shown in Fig. 4-1, as determined by sampling 350 realizations with different initial seeds for each PRNG; each realization was evolved to the asymptotic regime and ξ_L^2 was measured. Each realization was further evolved for ten autocorrelation times, and ξ_L^2 was measured again. This latter step was repeated 200 times. Hence the histogram shown in Fig. 4-1 was constructed with 7×10^4 data points for each PRNG. The data obtained by both PRNGs converge to the same empirical distribution, to well within statistical error. We will denote the average value of this distribution by μ, and its standard deviation by σ. For `random()` the values obtained are $\mu = 35.87$ and $\sigma = 20.85$. For `ran2()` the values obtained are $\mu = 35.74$ and $\sigma = 20.58$. Each independent observation of ξ^2, in the asymptotic regime, should be an iid random variable sampled from this distribution.

Before proceeding to the statistical tests, we briefly compare this distribution to previous ones obtained for growth models. The complete distribution function naturally contains much more information about the system than just the average value of ξ_L^2. The KPZ equation, as well as other exactly solvable models in 1-dimension,[92, 23] give rise to steady-state distributions which are identical to those of a random walk. The random walk distribution is expected to describe the BD model as well, assuming that it falls in the KPZ universality class. The theoretical distribution for the overall width of a random walk, with periodic boundary conditions, was calculated recently[46] and is shown in Fig. 4-1, overlaying the empirical BD distribution. There is a slight, but systematic disagreement between the BD histograms and the theoretical distribution for the random walk. The data for numerical implementations of the

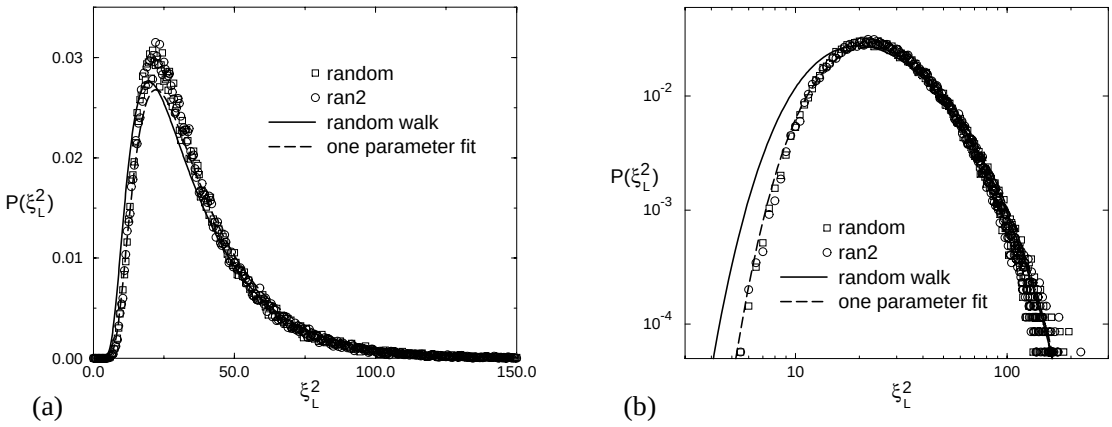

Figure 4-1: The empirical steady-state distribution function, $P(\xi_L^2)$, obtained for both PRNGs. The solid line corresponds to the theoretical distribution function for a random walk ($\alpha = 1/2$), the dashed line to a one parameter fit with $\alpha = 0.45$, where α is the roughness exponent, as plotted on (a) a linear scale, (b) a log-log scale.

KPZ equation (and for solid-on-solid growth models) have been successfully fitted to the random walk distribution[46, 105], confirming that this distribution does indeed describe systems in the KPZ universality class.

A *better* fit is obtained by following a phenomenological approach, introduced by Rácz and Plischke[105]. The width of the growth interface in the steady-state can be obtained from the structure factor $S(k) = \left\langle \hat{h}(k)\hat{h}(-k) \right\rangle$, as $\xi_L^2 = b^{-1} \sum_{k \neq 0} S(k)$, where $\hat{h}(k)$ is the Fourier transform of $h(x)$ (and hence $k = 2\pi m/L$, where m is an integer), and b is a constant with dimensions of inverse width squared. Assuming that $\{h(x)\}$ is continuous, and Gaussian distributed with a Kernel given by $S(k)$, the probability distribution for the width is calculated as[105],

$$P(\xi_L^2) = \int_{-i\infty}^{i\infty} \frac{d\lambda}{2\pi i} e^{\lambda \xi^2} \prod_{k \neq 0} \frac{bS^{-1}(k)}{\lambda + bS^{-1}(k)}. \tag{4.3}$$

For self-similar fluctuations, the power spectrum behaves as $S(k) \propto |k|^{-\gamma}$, where γ is related to the standard roughness exponent α by $\gamma = 2\alpha + d$, where d is the dimension of the substrate.

Figure 4-1 shows the phenomenological distribution function for the parameter value of $\alpha = 0.45$, alongside the random walk distribution function ($\alpha = 0.5$), and the numerical data. Figure 4-1(a) shows these functions plotted on a linear scale, Fig. 4-1(b) on a log-log scale. The phenomenological distribution shown is obtained by summing the residues of the first forty-five poles in the contour integral of Eq.(4.3). The sum of the residues for the first eighteen poles converges to the identical distribution throughout the regime considered, indicating that the result is essentially exact. The phenomenological distribution captures certain aspects of the numerical data with more fidelity than the random walk distribution. The value of $\alpha = 0.45$ is smaller than the KPZ prediction of $\alpha = 1/2$, but is consistent with previous values of the roughness exponent reported in numerical studies of BD[75, 41, 28, 92].

In regard to the statistical tests discussed in the remainder of this manuscript, the relevant result presented in this section is that both PRNGs converge to the same statistical distribution. The discrepancies with random walk behavior will not be further considered in this manuscript.

4.3.2 Distribution of averages

We are concerned with the consistency between independent measurements, and introduce the statistical tests that follow in order to test this (or more accurately to evaluate the lack of consistency). The tests assume approximately Gaussian distributed variables. Although a single measurement of the width of the interface (Fig. 4-1) is not a Gaussian variable, the probability distribution of the average width over many realizations is Gaussian (via the Central Limit Theorem[43, 14]). In simulations one typically considers average quantities, which also necessitates knowing the probability distribution of the average. Our simulations are of N independent realizations, so the relevant distribution is that of the width of the interface averaged over N independent samples, $\overline{\xi_L^2} = \sum_{i=1}^{N} \xi_L^2/N$.

We can construct the actual distribution function for the average over N iid sam-

ples, $P(\overline{\xi_L^2})$, from its Fourier transform denoted by $\tilde{P}_N(k)$ (and usually referred to as the characteristic function). The characteristic function of the average is related to the characteristic function of the individual samples ($\tilde{P}(k)$) by: $\tilde{P}_N(k) = \left[\tilde{P}(\frac{k}{N})\right]^N$. The characteristic function is the generator of the cumulants of the distribution, and the *nth* cumulant of the distribution $P(\overline{\xi_L^2})$ is related to the *nth* cumulant of the distribution $P(\xi_L^2)$ simply by $\left\langle (\overline{\xi_L^2})^n \right\rangle_c = \frac{1}{N^{n-1}} \left\langle (\xi_L^2)^n \right\rangle_c$.

Dealing explicitly with the first two cumulants of the probability distribution of the average,

$$\left\langle \overline{\xi_L^2} \right\rangle = \left\langle \sum_{i=1}^{N} \frac{1}{N} \xi_{L,i}^2 \right\rangle = \sum_{i=1}^{N} \frac{1}{N} \left\langle \xi_{L,i}^2 \right\rangle = \mu, \tag{4.4}$$

$$\begin{aligned} \left\langle (\overline{\xi_L^2})^2 \right\rangle_c &= \sum_{i=1}^{N} \left\langle (\frac{1}{N} \xi_{L,i}^2)^2 \right\rangle_c = \frac{1}{N^2} \sum_{i=1}^{N} \left\langle (\xi_{L,i}^2)^2 \right\rangle_c \\ &= \frac{\sigma^2}{N} \equiv \nu^2. \end{aligned} \tag{4.5}$$

Thus each independent observation of $\overline{\xi_L^2}$ is a random variable drawn from a distribution with mean μ and variance $\sigma^2/N \equiv \nu^2$.

As N increases in value, higher order cumulants go to zero, and the distribution approaches a Gaussian (as required by the Central Limit theorem [43, 14]). The function $P(\overline{\xi_L^2})$ is shown in Fig. 4-2, along with the numerical data for $P(\xi_L^2)$ (previously shown in Fig. 4-1), and a Gaussian distribution with mean μ and variance ν^2, for $N = 20$.

4.3.3 The χ^2 distribution

Each independent realization of $\overline{\xi_L^2}$ is approximately Gaussian distributed about the mean, μ, of the empirical distribution, with a variance ν^2 (see Fig. 4-2). Hence the normalized difference $\left(\overline{\xi_L^2} - \mu\right)/\nu$ should be a random variable sampled from a unit normal distribution (i.e. a Gaussian distribution with mean of zero, and unit variance). The sum of squares of M independently distributed unit normal random

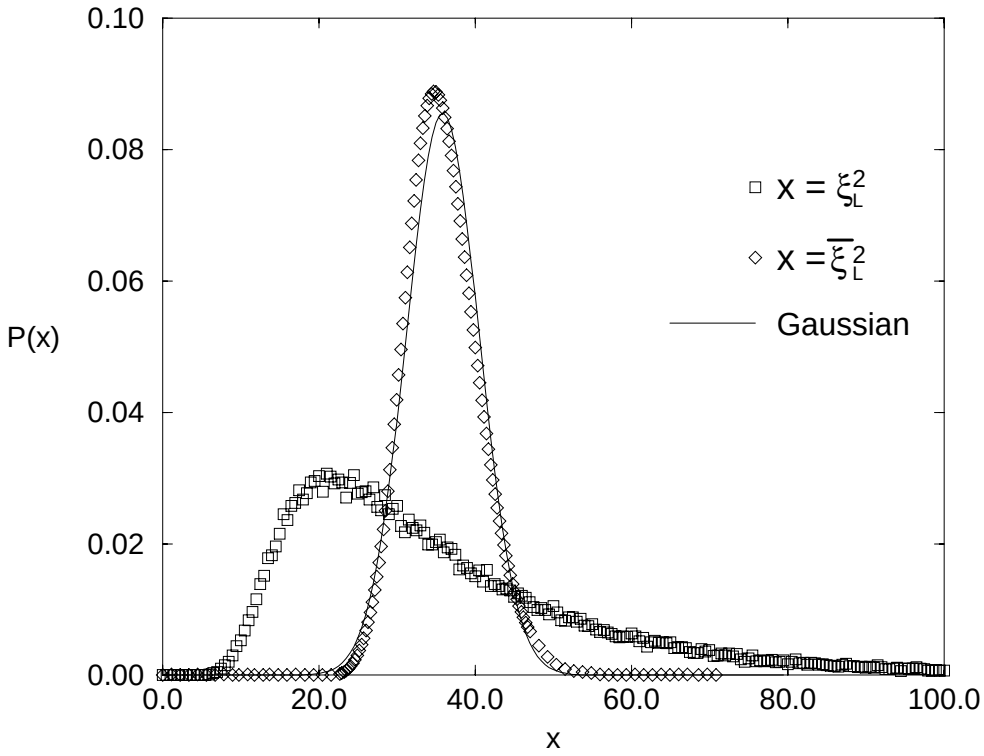

Figure 4-2: The distribution function for the individual samples, $P(\xi_L^2)$, the distribution for the average over $N = 20$ samples, $P(\overline{\xi_L^2})$, and a Gaussian approximation to the distribution for the average.

variables, denoted by χ^2_{ss},

$$\chi^2_{ss} = \sum_{i=1}^{M} \frac{\left(\overline{\xi^2_{L,i}} - \mu\right)^2}{\nu^2}, \tag{4.6}$$

follows a χ^2 distribution with M degrees of freedom[43, 14]. If the χ^2_{ss} statistic is sufficiently large, it is unlikely that all values in the sum are approximately unit normal distributed. The χ^2 test quantifies how unlikely; the test determines the probability that a number of value χ^2_{ss} or greater is drawn from a χ^2 distribution with M degrees of freedom. We denote this test as the χ^2_{ss} test and use it to determine the probability of the hypothesis that all the values of $\overline{\xi^2_L}$ in the conjectured steady-state regime were sampled independently from the same underlying distribution (shown in Fig. 4-2).

When performing a simulation one uses the values of the average and standard deviation obtained in the simulation as estimates of the average and standard deviation of the distribution function. In order to use the sample average $\left\langle \overline{\xi^2_L} \right\rangle_M = \left(\sum_{i=1}^{M} \overline{\xi^2_{L,i}}\right)/M$ in place of μ, and the weighted variance $s^2/(N-1) = \sum_{i=1}^{N}(\xi^2_{L,i} - \overline{\xi^2_L})^2/\left[N(N-1)\right]$ in place of $\sigma^2/N = \nu^2$, we refer to the t test. Note that these values obtained from our simulations should be unbiased estimators of μ and ν^2 respectively[43, 14], and represent more accurately the error bars obtained.

4.3.4 The t distribution

'Student'[104] first discussed the error introduced by estimating σ^2 with the sample standard deviation $s^2 = \sum_{i=1}^{N}\left(\xi^2_{L,i} - \overline{\xi^2_L}\right)^2/N$, and suggested the t test as an alternative statistical test[104]. Note that $s^2/\left(\sigma^2/N\right) = s^2/\nu^2$ follows a χ^2 distribution with $N-1$ degrees of freedom (the degrees of freedom are reduced by 1 as there is one constraint on the random variables: the mean value is equal to $\overline{\xi^2_L}$). Following 'Student' we can construct a statistic from the ratio of a unit normal distributed random variable $\mathcal{Z}$ to an independently χ^2 distributed random variable f^2_{N-1} with

$N-1$ degrees of freedom:

$$T_{N-1} = \frac{\mathcal{Z}}{\sqrt{f_{N-1}^2/(N-1)}}.$$

We label this variable T_{N-1}, as its probability density should follow a t distribution with $N-1$ degrees of freedom[43, 14]. The unit normal distributed random variable we are interested in is $\mathcal{Z} = \left(\overline{\xi_L^2} - \mu\right)/\nu$. The χ^2 distributed random variable is $f_{N-1}^2 = s^2/\nu^2$. Thus the corresponding t statistic is

$$T_{N-1} = \frac{\left(\overline{\xi_L^2} - \mu\right)/\nu}{\sqrt{s^2/\left[(\nu^2)\,(N-1)\right]}} = \frac{\left(\overline{\xi_L^2} - \mu\right)}{\sqrt{s^2/(N-1)}}. \tag{4.7}$$

We likewise define a second t statistic, useful for comparing two independent data sets. Consider two independent sets of M iid samples drawn from a probability distribution with mean μ and variance ν^2 (i.e. two independent sets of M iid realizations of $\overline{\xi_L^2}$). There are thus two independent measurements of the average value over the M iid samples, denoted by $\left\langle \overline{\xi_L^2} \right\rangle_{M,i}$, for $i = 1, 2$ (as defined earlier, but note the additional index i, used to designate the data set). Likewise there are two independent measurements of the variance over M samples, denoted by $\mathcal{S}_i^2 = \sum_{j=1}^{M} \left(\overline{\xi_L^2}_{,i,j} - \left\langle \overline{\xi_L^2} \right\rangle_{M,i}\right)^2 /M$, for $i = 1, 2$. The difference between two independent observations of average values is a random variable which converges to a Gaussian distribution with mean of zero and variance $2\nu^2/M$, in the limit of large M. Hence an approximately unit normal distributed random variable is $\mathcal{Z} = \left(\left\langle \overline{\xi_L^2} \right\rangle_{M,1} - \left\langle \overline{\xi_L^2} \right\rangle_{M,2}\right)/\sqrt{2\nu^2/M}$. An independent χ^2 distributed random variable (with $2(M-1)$ degrees of freedom) is $f_{2(M-1)}^2 = \mathcal{S}_1^2/(\nu^2/M) + \mathcal{S}_2^2/(\nu^2/M)$. The t statistic to compare two independent data sets is thus

$$T_{2(M-1)} = \frac{\left(\left\langle \overline{\xi_L^2} \right\rangle_{M,1} - \left\langle \overline{\xi_L^2} \right\rangle_{M,2}\right)}{\sqrt{(\mathcal{S}_1^2 + \mathcal{S}_2^2)\left(\frac{1}{M-1}\right)}}. \tag{4.8}$$

If the t statistic is sufficiently large, it is unlikely that the ratio of the sample mean

to the sample variance is an accurate estimator of the ratio of the theoretical mean to the theoretical variance. The "t test" measures the probability that a value the size of the t statistic or greater is drawn from the appropriate t distribution[43, 14]. In summary, the T_{N-1} test using the T_{N-1} statistic, measures the level of validity for the hypothesis that each sample in a set was drawn from the same underlying distribution. The $T_{2(M-1)}$ test using the $T_{2(M-1)}$ statistic, measures the level of validity for the hypothesis that samples in two distinct sets were drawn independently from the same underlying distribution. This latter test compares average values generated by the two distinct sets.

4.3.5 Distribution of extrema

The final class of statistical tests employed deals with extreme values found in sets of iid random variables. The extrema found in our data appear to present the largest deviation from the theoretical distribution. Of course, when dealing with extrema there is a selection effect, for which one has to correct. The probability of obtaining a particular value for the minimum in a set of M iid random variables can be readily calculated as[50]

$$p(x_{\min} = x) = M \times p(\overline{\xi_L^2} = x) \times \left(\int_x^\infty p(\overline{\xi_L^2}) d\overline{\xi_L^2} \right)^{M-1}. \tag{4.9}$$

Likewise the probability for obtaining a particular value for the maximum in a set of M iid random variables can be readily calculated as[50]

$$p(x_{\max} = x) = M \times p(\overline{\xi_L^2} = x) \times \left(\int_{-\infty}^x p(\overline{\xi_L^2}) d\overline{\xi_L^2} \right)^{M-1}. \tag{4.10}$$

The original probability distribution for the average width, $P(\overline{\xi_L^2})$, along with the distributions for the minimum and the the maximum in a set of $M = 10$ independent observations, are shown in Fig. 4-3.

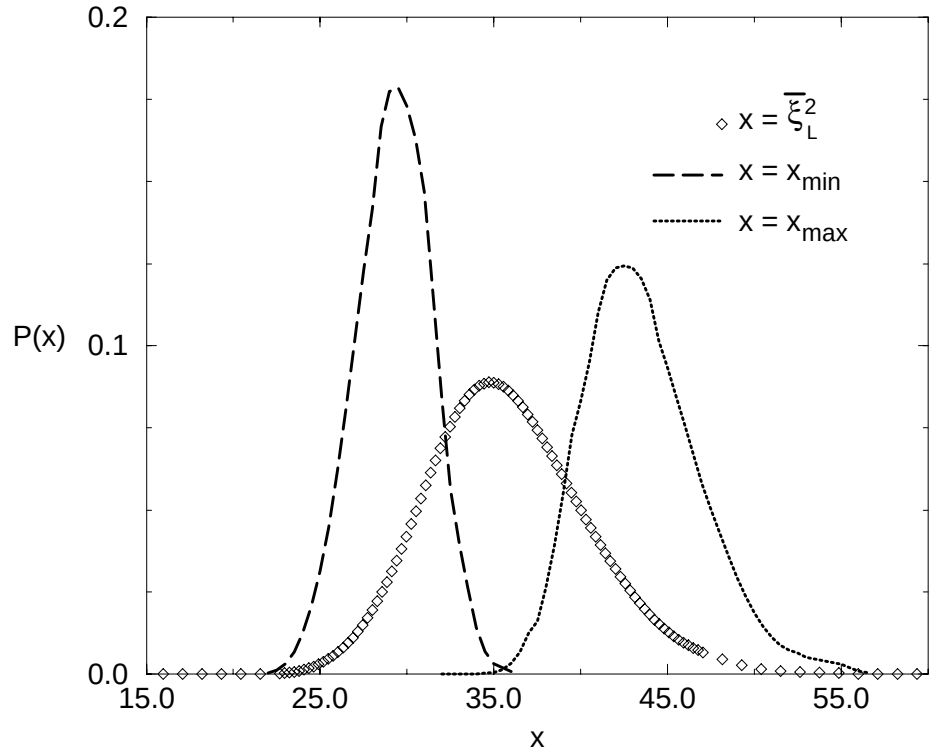

Figure 4-3: The probability distribution for the average width, $P(\overline{\xi_L^2})$, along with the distribution for the minimum and the maximum in a set of $M = 10$ independent observations of $\overline{\xi_L^2}$.

4.4 Results of statistical tests

We applied the statistical tests outlined above to the data obtained from our simulations. Results for the χ^2 test are reasonable, however the results for the t test and for the extremal values are highly unreasonable. Thus specific simulations may give highly anomalous results, inconsistent with the theoretical distribution.

4.4.1 Comparison of data at different times (in the asymptotic regime)

In the asymptotic regime, the width of the interface should saturate to a steady-state value. The average over 20 independent samples, $\overline{\xi_L^2}$, for an implementation with random(), is recorded at selected subsequent times, and a plot of this data is shown in Fig. 4-4. All of the values shown should be equal within statistical error, but there are large differences. In fact the greatest difference between two values is over four standard errors in magnitude. To quantify the significance of this difference, we apply the statistical tests discussed in Sec. 4.3, to this data set.

The χ^2_{ss} statistic, defined in Eq.(4.6), is calculated, but is not large enough to be significant: the probability that the null hypothesis of steady-state behavior is valid is 24%, as determined by the χ^2_{ss} test. We denote this as a pass at the 24% Confidence Level (CL), or likewise a fail at the 76% CL. In other words an event with this magnitude is expected to occur one in four times, therefore a pass at the 24% CL is not an unreasonable result.

The T_{N-1} test for steady-state behavior however does manifest discrepancies. 8 of the 10 points pass the T_{N-1} test above the 20% CL. But two points fail at the 94% CL, with one of these points continuing to fail at the 99.98% CL. An event that fails at the 99.98% CL is expected to occur only twice in $10,000$ times, we observe it once in ten times, indicating that the null hypothesis is highly suspect.

To ensure that the data sample shown in Fig. 4-4, is not a statistical fluke, and to

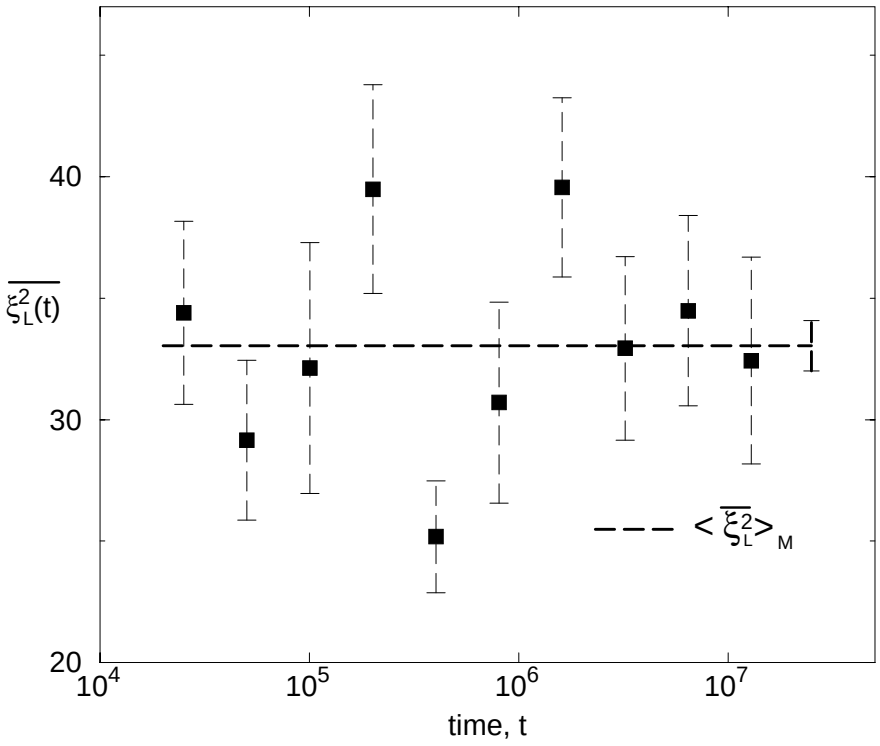

Figure 4-4: The average width of the growth interface, $\overline{\xi_L^2}$, as obtained for 20 independent samples using random(), shown at selected subsequent times in the asymptotic regime. The horizontal line corresponds to the average value $\left\langle \overline{\xi_L^2} \right\rangle_M$ of the $M = 10$ data points in this figure, with the associated error bar plotted at the extreme end of the line. Note that the expected relaxation time is less than $\tau = 2 \times 10^4$, and that the logarithmic scale spans roughly $10^3\tau$.

establish the repeatable and consistent absence of steady-state behavior, many more independent samples were generated. A total of 200 independent realizations were simulated, and the value of ξ_L^2 was measured for each realization at all of the times recorded in Fig. 4-4. The Gaussian distributed random variable $\overline{\xi_L^2}$ is constructed by splitting the 200 independent realizations into subsets of $N = 20$ and calculating the average of each subset. Hence there are ten subsets, each one of the random variable $\overline{\xi_L^2}$ sampled at the ten times shown in Fig. 4-4, for a total of 100 realizations of $\overline{\xi_L^2}$ in the conjectured steady-state regime. We refer to each subset as a data set throughout the remainder of the manuscript. The χ_{ss}^2 test is applied independently to each of the ten data sets. $T_{2(M-1)}$ test can be applied to any combination of two independent data sets. The T_{N-1} test is applied independently to each of the 100 realizations of $\overline{\xi_L^2}$. By conducting the statistical tests on the additional data sets we show that the results obtained for the original data set are systematic, as discussed below.

The χ_{ss}^2 statistic defined by Eq.(4.6) was calculated for each data set. The data sets all pass this test at greater than the 20% CL. Yet sufficiently many realizations of $\overline{\xi_L^2}$ fail the T_{N-1} test, to bring the null hypothesis of a steady-state into question. 14 of the 100 iid realizations of $\overline{\xi_L^2}$ fail the T_{N-1} test at the 90% CL, which in itself would not allow us to reach conclusions. However as the CL criteria is tightened beyond acceptable standards a surprising number of points still fail. At the 99% CL, 5 of the 100 points fail. These 5 points also fail at the 99.8% CL. At the 99.98% CL, 3 points fail. Finally, at the 99.995% CL, 1 of the 100 points fails. The meaning of the 99.98% CL is that probability theory predicts the occurrence of 2 such events out of $10,000$. Instead we observe 3 such events out of 100. Likewise the 99.995% CL corresponds to 5 events out of $100,000$. We observe 1 such event out of 100. Extreme "tail events" thus occur with a frequency which is more than two orders of magnitude greater than the laws of probability would indicate.

We now turn to a discussion of the average asymptotic value of each data set of $M = 10$ events (i.e. each set consists of the random variable $\overline{\xi_L^2}$ sampled at

10 consecutive times in the asymptotic regime). The value of $\left\langle \overline{\xi_L^2} \right\rangle_M$ generated by each set of data is not consistent with the other sets. The ten values range[2] from $\left\langle \overline{\xi_L^2} \right\rangle_M = 33.05 \pm 1.31$ to 37.39 ± 1.34. Randomly picking pairs of data sets to compare using a $T_{2(M-1)}$- test, we find several instances where the null hypothesis (that the two sets of data being compared were sampled independently from the same underlying distribution) fails at the 96% CL. Thus separate runs of a simulation frequently give statistically inconsistent results.

When all $N = 200$ samples are combined into one data set, the average quantities are consistent with the empirical distribution function shown in Fig. 4-1: $\left\langle \overline{\xi_{L,N=200}^2} \right\rangle_M = 35.65 \pm 0.24$. In addition, this data set of the average over 200 iid samples, passes the χ^2_{ss} test at the 97% CL. However, for the data sets of averages over $N = 20$ iid realizations, the disparity in average asymptotic values obtained shows that sampling a subset of 20 iid samples is not consistent with the empirical distribution. The values for averages and variances obtained from the subsets are not unbiased estimators of the empirical distribution function. Furthermore, the statistical tests discussed so far reject the hypothesis of steady-state behavior, even at time scales orders of magnitude greater than the conjectured model relaxation time of $\tau \sim 10L^{z'}$. Thus there is no steady-state behavior for the data generated with random(); instead there is an asymptotic dynamic behavior.

4.4.2 Comparison of extrema with the steady-state distribution

In the preceding section the occurrence of many extreme tail events was established. In this present section we show that the occurrence of tail events can skew the average

[2]Note that the data set plotted in Fig.4-4 is the set with the smallest asymptotic mean ($\left\langle \overline{\xi_L^2} \right\rangle_M = 33.05 \pm 1.31$). However two additional data sets obtained with random() have asymptotic means below 34.0. We choose to plot the set shown in Fig.4-4 as it was the first set obtained, as such we are consistent in comparing results with ran2(), where only one data set was generated.

values obtained.

Comparing the data shown in Fig. 4-4 to the expected probability distribution, shown in Fig. 4-2, it is observed that the data points are skewed to the left side of the expected distribution. Eight of the ten points are below the mean value, $\mu = 35.87$, and the lowest value, $\overline{\xi^2_{L\min}} = 25.18$, is in the left hand tail with less than 0.32% of the total area of the probability distribution function. The probability of obtaining a particular value for the minimum in a set of M iid random variables is described in Eq.(4.9). Using this formula we calculate the probability that the minimum of $M = 10$ iid random variables drawn from the distribution shown in Fig. 4-2 is less than or equal to $\overline{\xi^2_{L\min}} = 25.18$ is only $p(x \leq x_{\min} = 25.18) = 3.6\%$.

Again we wish to determine if the first data set is a fluke event, so a statistical analysis based on all ten data sets obtained with random() is warranted. We find that three of the ten data sets have minima which come from the extreme left hand tail of the distribution, and $p(x \leq x_{\min}) < 6\%$ for each. As the probability distribution for the minimum, $p(x_{min})$, is known (see Fig. 4-3), we can construct the probability for observing three such low probability events (from the left hand tail of the distribution) out of a total of ten events, and find this probability to be 1.7%. As such, we can state at the 98% CL that such results would not be obtained by random sampling.

A similar analysis can be carried out with respect to the maxima. In line with the observation that the data in Fig. 4-4 is skewed to the left side of the expected distribution, we find several values for maxima which are questionably low. The analogous probability for the maximum value in a set of M iid random variables is given by Eq.(4.10), and plotted in Fig. 4-3. Maximum values for four of the ten data sets are much smaller than expected, with each having probability $p(x \leq x_{\max}) < 12\%$. The probability of observing four such low probability events (from the left hand tail of the distribution) out of a total of ten events is 1.9%. Again we can state at the 98% CL that such results would not be obtained by random sampling.

4.4.3 Comparison of data from distinct PRNGs

To determine if the source of the observed asymptotic dynamic behavior resides in the PRNG random(), results using a second PRNG, ran2(), were also analyzed. ran2() is substantially slower than random(), hence the comparison is based on 20 independent samples for each PRNG (i.e. we use only the initial set of data[3] for random(), shown in Fig. 4-4).

For the shortest length scale implemented ($L = 127$), the time-series data for ran2() is self-consistent. The data passes the χ^2 test, the T_{N-1} test and the tests for extremal values. However at longer length scales the data for ran2() fails several statistical tests, making ran2() also suspect in simulations of BD. The tests performed at the longer length scales were adequate to show statistical inconsistencies[26], however not with the high level of rigor demonstrated by the tests on the data at the shortest length scales[4]

A direct comparison of data generated by the two different PRNGs for $L = 127$ is shown in Fig. 4-5. There are 18 sets of points that can be directly compared, including eight which were sampled at $t < \tau$. The $T_{2(M-1)}$ test for consistency between the two values at each time fails at the 90% level for three out of the eighteen sets of points. Most striking is the direct comparison of average asymptotic values obtained for each PRNG, $\left\langle \overline{\xi_L^2} \right\rangle_{M,\mathsf{random}} = 33.05 \pm 1.31$ and $\left\langle \overline{\xi_L^2} \right\rangle_{M,\mathsf{ran2}} = 37.06 \pm 0.80$. A $T_{2(M-1)}$ test for the equivalence of the asymptotic averages of the data sets for random() and ran2() fails at the 99% confidence level (the exact probability of failure is $p = 0.996$). Hence use of different PRNGs can yield statistically distinct values for averages.

[3]Note that the data set plotted in Fig.4-4 is the set with the smallest asymptotic mean ($\left\langle \overline{\xi_L^2} \right\rangle_M = 33.05 \pm 1.31$). However two additional data sets obtained with random() have asymptotic means below 34.0. We choose to plot the set shown in Fig.4-4 as it was the first set obtained, as such we are consistent in comparing results with ran2(), where only one data set was generated.

[4]The theoretical fits discussed in Sec. 4.3.1 are not precise enough to provide accurate measures of the cumulants of the empirical distribution function. Hence, the distribution for longer lengths cannot be assumed to be accurately described by a corresponding theoretical curve. Sampling sufficient data points to establish the empirical asymptotic distribution for the longer lengths is beyond the scope and intent of this manuscript.

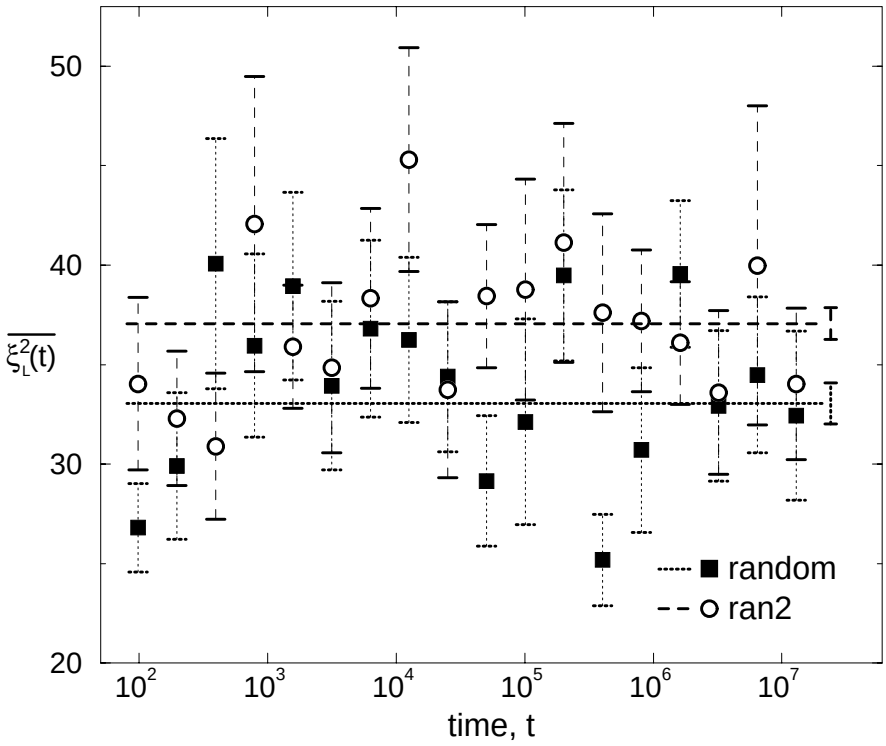

Figure 4-5: The average width of the growth interface as obtained by 20 independent samples for each PRNG. The horizontal lines correspond to the respective average asymptotic values, $\left\langle \overline{\xi_L^2} \right\rangle_{M,\text{random}}$ and $\left\langle \overline{\xi_L^2} \right\rangle_{M,\text{ran2}}$, with the corresponding error bars plotted at the extreme end of each line.

The statistical tests reject the hypothesis that the two PRNGs sample the same underlying distribution, despite the fact that the asymptotic distributions shown in Fig. 4-1 agree. There is a unique steady-state distribution which is obtained in the limit of large numbers of independent samples. However sampling of this distribution is nonstochastic, in that each sample average is not an unbiased estimator of the asymptotic distribution. Likewise the standard error for each sample does not lead to an unbiased estimator for the standard deviation of the asymptotic distribution.

4.5 Discussion & conclusions

The original impetus for this study was an indepth investigation of BD at long length and time scales. However, we encountered many features in the data that could not be easily explained; most notably, non-self-affine surface fluctuations. After searching for various corrections to scaling, which necessitated obtaining better statistics and exploration of longer times into the growth, the coupling to PRNGs became apparent, and motivated the detailed statistical analysis described in this manuscript.

It should be noted that there are discrepancies between values of the scaling exponents for BD reported in the literature[75, 41, 28, 92]. At this point, we can only speculate that these discrepancies are due to the differences in the implementations of BD. Conclusions about the scaling exponents can only be reliably reached once difficulties with PRNs are resolved.

In retrospect, it is not surprising that the BD algorithm is more sensitive to correlations in PRN sequences than standard Monte Carlo (MC) simulations. In standard MC, comparison to the Boltzmann probability causes rejection of PRNs at pseudorandom points in the sequence; hence a second independent source of pseudorandomness influences the dynamics. In BD, all PRNs are used in the sequence produced (with the exception of the very few cases discussed in Sec. 4.2.2). Similarly, in restricted solid-on-solid models of growth (where physical constraints cause rejection of PRNs),

the scaling exponents and the random walk distribution predicted by KPZ theory are recovered with great precision in numerical simulations[46, 71].

We have demonstrated that computer implementations of BD can couple to certain PRNG algorithms. Results statistically inconsistent with general sampling assumptions, and with the ergodic exploration of phase space were observed. Exploration of accessible phase space is not decoupled from the initialization of the PRNG. In addition, driving the dynamics of the system with different PRNGs results in sampling different areas of phase space. In conclusion, BD is a sensitive physical test of correlations in pseudorandom sequences. In general, PRNG algorithms can couple to models of stochastic, nonequilibrium phenomena. One must ensure that observed dynamical properties are inherent in the non-equilibrium model itself and not an artifact of coupling to PRNGs.

Part II

Reversible Models

Chapter 5

Discrete reversible models of physics

5.1 Overview

5.1.1 Why study discrete reversible systems?

Discrete systems played a key role in the development of statistical mechanics. The tools of combinatorics enabled the exact counting of states and measurement of entropy. From the analytic expression for the entropy, the expressions for the other thermodynamic quantities can be derived as can the conditions for thermodynamic equilibrium.

Discrete systems still play a key role today in understanding statistical dynamics, especially since they can be exactly implemented on digital computers. The stochastic growth models studied in Part I are classic examples of useful computer simulations. Also, one of the most widely theoretically and numerically studied systems in statistical physics remains the Ising model (a discrete spin system)[9, 10].

We are interested in discrete *invertible* dynamics, meaning that from any state in the dynamical evolution of the system we can recover the previous state exactly:

A complete history of the system can be generated from any state through which the system progresses. We are interested these invertible models since the laws of physics are microscopically reversible, and for the study of thermodynamic behavior. Entropy cannot decrease in reversible systems (as discussed in the Chap. 1). Yet if we coarse grain the system, entropy can increase, and we can observe the approach to thermodynamic equilibrium.

5.1.2 Reversibility and microscopic physics

All of the fundamental equations of physics are microscopically reversible. Nevertheless, there is ongoing controversy over the fundamental reversibility of the laws of physics[103, 79], which has persisted since the time of Boltzmann. Our daily experience of irreversibility is at the macroscopic scale which necessarily involves coarse graining and averaging over microscopic degrees of freedom, so we lack intuition about microscopic dynamics.

Current proponents of the fundamental irreversibility of physics initially based their arguments on the formation of "dissipative structures"[102]; The formation of stable structures necessitates the dissipation of information at the macroscopic scale. Yet this does not necessarily imply dissipation at the microscopic scale. These same proponents now base their arguments on phenomena not completely understood, such as nonequilibrium thermodynamics, and require extending the standard constructs of physics, such as needing to go "outside the Hilbert space" or to anomalous Fock spaces[103]. Yet the simple explanations for why we observe macroscopic dissipation in a microscopically reversible world, dating back to Boltzmann, have not been disproven. A guiding principle of scientific research remains Occam's razor: the simplest theory which explains the phenomena should be preferred. Joel Lebowitz is the main modern day proponent of Boltzmann's views. He has written several eloquent reviews of Boltzmann's writings (see *e.g.*, Ref. [78, 80, 79]).

5.1.3 Reversibility and pattern formation

Pattern formation is an intrinsically dissipative process, however it seems that the laws of physics are microscopically reversible: there is no dissipation of information at the microscopic scale. Information cannot be destroyed, but it can be transferred to other degrees of freedom and is usually ultimately transferred to thermal degrees of freedom. We are interested in simple systems which exhibit pattern formation through microscopically reversible dynamics. They model how dissipation arises (i.e., how information flows between the macroscopic and the microscopic degrees of freedom), and provide a clear example of how to reconcile the macroscopic irreversibility inherent in pattern formation with the microscopic reversibility adhered to by physical processes. By modeling the entire process of pattern formation, including the dissipation, we hope to emulate the mechanisms of nature more accurately than do irreversible models. Implementations of reversible microscopic models provide a laboratory in which to probe the thermodynamics of pattern formation.

5.1.4 Reversible dynamics and computation

We are concerned with developing models of physical phenomena which incorporate basic constraints of physics. Physical processes, such as computation, must also exist within these constraints. The same insights which help us construct computer models of physics can help us build physical models of computers.

Such models may have important practical consequences. For example understanding the physics of computation may enable computing at the atomic scale. The relevant fundamental constraint involves heat production and dissipation and necessitates the use of invertible information processing devices[76, 7]. For scaling arguments establishing the physical limits of dissipative computation see for example Ref. [47].

5.1.5 Cellular Automata

Cellular Automata (CA's) are fully discrete dynamical lattice systems: space is discrete, time is discrete, and the state at each discrete lattice site can take on only one of a fixed number of discrete values. Discrete dynamical lattice spin systems, such as dynamical Ising models[18, 19], are examples of CA's. The cellular automaton (CA) evolves through the synchronous application of a local deterministic rule; the state of each site is updated in parallel in accordance with a rule involving the state of that lattice site and the states of neighboring sites. The rules for the evolution are typically nonlinear. Since information can only propagate a finite distance with each update, we can study the effects of locality. CA's are local, nonlinear dynamical systems useful for investigating how correlations grow with time. They were originally introduced by von Neumann[126] to study self-reproducing systems.

In the standard language of physics, a cellular automaton is a discrete classical lattice field theory. CA's offer an alternative to partial differential equations for modeling spatial temporal dependencies and time evolution. In addition since they are fully discrete systems, the dynamics can have exact conservation laws, exact discrete symmetries, and exact invertibility.

CA's have been used to model general properties of dynamical systems[3], statistical physics[123, 131, 118, 112, 18, 21], self-organization[72, 58], chemical reactions[85], even traffic jams[96, 37, 130], and computer models of artificial "life"[2]. These are just a few applications, other relevant CA models will be discussed in the remainder of this chapter. Note that all the growth models discussed in Chap 2 can be implemented as CA's. For reviews and collections of the use of CA's in physics see Refs. [121, 3, 107, 16, 39, 25].

5.1.6 Overview of remainder of Part II

My motivation for studying discrete reversible systems is to study the thermodynamics of pattern formation. In Chap. 6, I present the formulation, implementation, and results of a thermodynamic model of the growth of clusters via the aggregation of particles. In the reminder of this chapter, I review the relevant reversible CA models of physics, focusing on lattice gas diffusion and hydrodynamics, and spin systems coupled to heat baths.

In Chap. 7, I show how to simulate digital logic with the RA model and discuss issues of computation in real physical systems. Finally in Chap. 8, I discuss issues of ergodicity and recurrence times for invertible discrete dynamics.

5.2 Reversible cellular automata models of physics

The laws of physics and CA's have fundamental similarities. They are local. They are deterministic. They are uniform. Physics has additional properties, such as microscopic reversibility and exact conservation laws. We want to put into our models the constraints at the microscopic scale consistent with the microscopic properties of physics that persist at the macroscopic scale. Some models of this sort have been shown to accurately reproduce physical phenomena in the macroscopic limit, independent of discreteness and lattice effects.

CA modeling offers some advantages over modeling with partial differential equations. A primary advantage is the ease of dealing with complex boundary conditions. The interaction at the boundary is included in the definition of the CA rule. As an example consider diffusive flow through porous media. The CA dynamics specifies the interaction of the diffusing particles and the boundary. In contrast using traditional partial differential equation approaches to diffusion requires iterating a finite difference equation, solving a boundary value problem, and repeating this process until we converge to a solution. The more complicated the boundary, the more difficult the

solution. CA models can handle complex boundaries with no increase in the required computing power.

5.2.1 Qualitative agreement with Physics

5.2.1.1 Conservation of mass: Diffusion

Consider a simple model of particles moving on a lattice. The particles move with unit velocity along the lattice directions, so that at integer times they occupy a lattice site. In the case of a two dimensional square lattice there are four possible transport directions (which we will call channels), corresponding to the four lattice directions. If we restrict ourselves to one particle per channel per site, then each site can accommodate up to four particles. Between one "streaming" phase (the portion of the dynamics when the particles are moving) and the next, the particles are occupying lattice sites and we allow an "interaction" phase. We want the interaction to conserve particle number. This will also conserve energy since all particles move with the same speed and all energy in this model is kinetic.

One of the simplest ways to implement such an interaction is to define particle "states" for each channel at each site, with the state "0" representing the absence of a particle and the state "1" the presence. To implement a reversible interaction, we can permute the states. For instance we can decide at random to rotate the channels at each site in a clockwise or counterclockwise fashion. Clockwise meaning the northbound channel becomes the eastbound; the eastbound the southbound; the southbound the westbound; and the westbound the northbound. After the permutation, the particles then stream to adjacent lattice sites and again they are permuted. Since we permute each site individually, the particles undergo simultaneous random walks along the lattice, while maintaining the exclusion principle (only one particle per channel per site)[1]

[1]Note the individual random walkers do not feel the presence of the other walkers. However, all

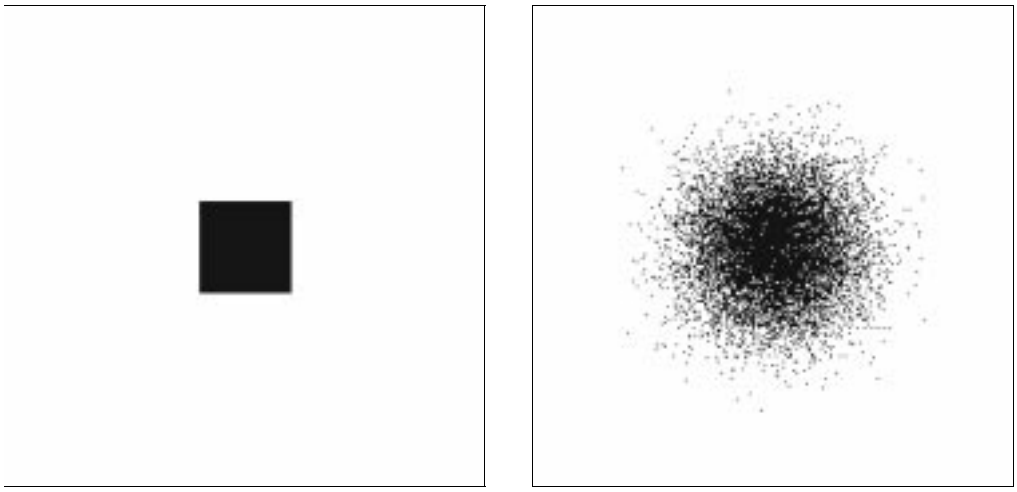

Figure 5-1: Diffusion simulated by a simple CA rule which conserves energy but not momentum. Particles move along the lattice directions of a two-dimensional square lattice. At each lattice site, at each update, the velocities of the particles at that site are permuted. Each site is permuted independently, yet the particles occupying the same site undergo the same permutation. Starting from an initial uniform block of particles, as shown in (a), we observe the diffusive behavior shown in (b), which corresponds to $t = 360$ complete updates of the lattice.

Such a dynamics conserves mass (no particles are ever destroyed or created), and thus energy (since particles all move with the same speed). It does not conserve momentum. Once we average over a few lattice sites, we can recover the exact statistical properties of diffusion from these particles undergoing random walks along the lattice. The quantitative behavior will be discussed in detail in Sec. 5.2.2.2. The qualitative behavior is shown in Fig. 5-1. We start from an initial uniform block of particles occupying an otherwise vacant lattice, as shown in Fig. 5-1(a). Qualitative diffusive behavior is apparent. After $t = 360$ updates of the full lattice we observe the configuration shown in Fig. 5-1(b).

Note the decision of whether to permute clockwise or counterclockwise, can be made by accessing the state of a binary random variable. We can superimpose a

the particles at one site leave that site correlated, as they have all undergone the same permutation. We initialize the system with up to one particle in each channel, and this constraint is preserved by the dynamics.

second lattice on the particle lattice, filled with random bits, to serve this function. This superimposed lattice is like a noninteracting heat bath. We use a deterministic, invertible dynamics on the heat bath bits, thus we can invert the dynamics, and recover the data used to make the permutations, and thus invert the random walks. CA models of diffusion will be discussed in more detail in Sec. 5.2.2.2, and reversible models of diffusion discussed in Chap. 6.

5.2.1.2 Conservation of mass and momentum: Wave phenomena

Let us add one more conservation law, conservation of momentum. For example consider the dynamics described above. Particles stream between the lattice sites. When they occupy the lattice sites, they interact. Above we considered a simple permutation which conserves mass. We now want to add in conservation of momentum. The only way to implement a nontrivial interaction of this sort on a square lattice is to allow two body collisions[2]: If two particles stream towards each other head on and occupy the same site, they each leave the interaction site with their velocities rotated by 90 degrees. The particles now "collide" in a way which conserves the overall momentum; There is zero net momentum before the collision and zero net momentum after the collision. We will still observe similar diffusive behavior[3]. In addition, we observe wave propagation, as illustrated in Fig 5-2. If we start with a pressure gradient such as a high density of particles everywhere in the space, but with a void in the center, we see an onrush of particles into the empty center and a round wavefront[4] of low pressure leaving this region. Thus adding one more relevant conservation results in a CA which supports a wave equation[121]. Note that this dynamics is exactly invertible: if we invert the particle velocities, while holding positions fixed,

[2]The only permutation of a single particle which conserves momentum is the identity permutation. Likewise for three particles on a square lattice only the identity permutation conserves momentum.

[3]In this case the diffusion coefficient depends in the density of particles, since it is the collisions which randomize the velocities. The mean free path is inversely proportional to the particle density.

[4]There are some artifacts of the square lattice which keep the wavefront from being isotropic. This will be discussed in Sec. 5.2.2.1.

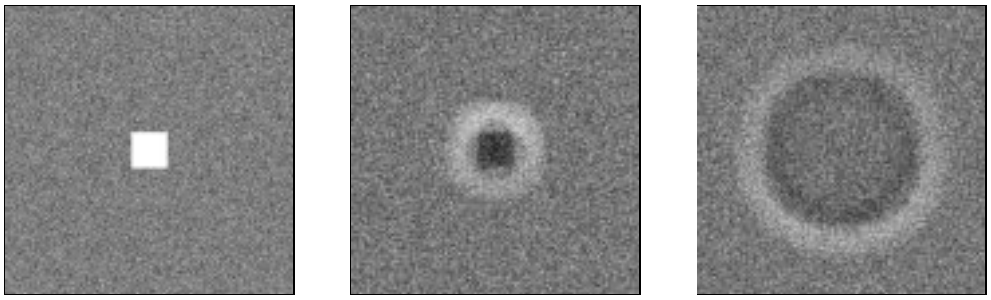

Figure 5-2: Wave propagation in a simple CA rule which conserves mass and momentum. (a) An initial density perturbation. The propagation of the wave front is shown in (b) and (c).

the collisions all happen in the reverse order. Note also that the collisions are their own inverse (applying the same collision twice in a row to the same particles results in the identity permutation).

We can also put in too many conservation laws and have "spurious invariants". Take the simple model which conserves mass and momentum described above (called the HPP lattice gas[59]). Since all collisions are head on, momentum is conserved along each row and along each column of the lattice.

5.2.2 Quantitative agreement with Physics

5.2.2.1 Lattice Gas hydrodynamics

Reversible cellular automata have been used to model various physical phenomena. The most well known are the lattice gas models of hydrodynamics, such as the HPP lattice gas described above. These models are important due to the ubiquity of hydrodynamic flow in nature, and because they provide intuition about the multiscale modeling approach. The macroscopic properties of the flow which are the defining characteristics of hydrodynamics are: conservation of mass, conservation of momentum, conservation of energy, and isotropy. The lattice gas approach answers the question: what is the simplest possible microscopic model which will reproduce these properties at the macroscopic scale?

In the lattice gas models discrete particles stream along the lattice and interact via collisions which conserve mass and momentum. The discrete conservation laws and the continuity equations for the system can be written exactly. By considering the average over an ensemble of similar systems the continuous conservation and continuity equations for the system can be established. With a few approximations, such as the assumption of molecular chaos and Taylor expansions, the partial differential equations which govern this system in the continuum limit can be ascertained. We will describe a lattice gas model in detail in Chap. 6. At that point the concepts of ensemble averaging and continuum descriptions of discrete dynamics will be explicitly discussed.

The HPP lattice gas has collisions on a square lattice which conserve mass and momentum. The macroscopic behavior of the lattice gas *approximately* obeys the Navier Stokes equation for fluid flow[59]. As mentioned on a square lattice such a dynamics has spurious invariants (*i.e.*, the momentum is conserved not only globally, but also along each row and each column of the lattice). Lattice effects also introduce anisotropy into the macroscopic behavior. If instead we use a hexagonal lattice with three body collisions, which mix momentum along the different lattice directions, are allowed, and the system can be shown to *exactly* obey the Navier Stokes equation for fluid flow[49].

The lattice gas approach has been applied widely to a variety of systems such as modeling phase separation of immiscible fluids[106, 11], modeling flow through porous media[107, 17], and even to the study of discrete forces[134, 135]. In order to ensure that the model will converge upon the correct thermodynamic equilibrium behavior, these lattice gas dynamics must obey semidetailed balance (although detailed balance is preferred). Since reversible dynamics provide a mechanism for transitioning backwards (and satisfy semi-detailed balance), there have been some attempts made to develop reversible algorithms for these extensions to the simple hydrodynamic models. To date the reversible implementations of complex fluids remain elusive[13].

As mentioned, the same constraints that we impose on the CA models to make them more like physics can help us understand how to make more efficient physical devices. A tangible example of the fruits of this approach is the development of a special purpose cellular automata machine, the CAM8. The CAM8 is based upon the lattice gas paradigm for modeling physics. The architecture and information flow in the machine directly maps onto that paradigm[87]. The lattice gas paradigm has been applied to an actual physical computer.

5.2.2.2 Diffusion

Diffusion implementations play an essential role in a variety of physical models. CA's can not only be used to model diffusion processes very accurately, they offer some advantages. The main advantage is the ease in dealing with complex boundary conditions; The interaction at boundary is included in the definition of the CA dynamics. These models can handle complex boundaries with no increase in the required computing power.

One illuminating result of using CA dynamics is that they are local and thus cannot reproduce exactly Fick's law for diffusion. Fick's law

$$\frac{\partial \rho}{\partial t} = D\nabla^2 \rho, \tag{5.1}$$

is non-local: gradients are sensed instantly everywhere in space. Chopard and Droz conducted a rigourous and exact study of a CA model for diffusion[15] and showed how locality alters Eq. 5.1. Instead of Fick's Law they obtain the telegraphist's equation, which describes a random walk with inertia:

$$\frac{\partial \rho}{\partial t} = D\nabla^2 \rho - \frac{D}{c^2}\frac{\partial^2 \rho}{\partial t^2}, \tag{5.2}$$

where $c \propto 1/\sqrt{2}$ is the speed of sound on the lattice. The telegraphist's equation emerges from the microscopic treatment of diffusion. It reflects the fact that infor-

mation cannot travel faster than the actual particles. In this way it is more realistic than Fick's law, however the difference only becomes important when the diffusion coefficient, D, is large (*i.e.*, when $D \gg c$).

Chopard and Droz model diffusion as synchronous random walks with exclusion on a two-dimensional lattice. Their model is essentially identical to the one presented in Sec. 5.2.1. They were the first to show how to derive the analytical solution for the diffusion algorithm. The dynamics is linear and is exactly described by a linear discrete Boltzmann equation. To extrapolate to the partial differential equations describing the system in the continuum limit requires averaging over an ensemble of similar systems. Typically the ensemble averaging requires the assumption of molecular chaos or the truncation of a hierarchy of correlations[12]. For the linear Boltzmann equation, the averaging can be done exactly. They solve the resulting equation analytically for certain boundary conditions, and observe true diffusive behavior, independent of the artifacts of discrete space and time. Note that details of ensemble averaging Boltzmann equations will be given in Sec. 6.3.

Along the lines of "spurious invariants" it is interesting to point out that the Chopard-Droz diffusion algorithm has one obvious one. Consider a checkerboard parity on the lattice. The four particle states which stream into a site at one time and interact have all come from one of the checkerboards (say the "red" one). There is no interaction between the "red" and the "black" checkerboards, so they are simulating diffusion on two independent lattices. In Fig. 5-1, we started both checkerboards in the identical state, so this is not apparent. If we had started with one checkerboard empty, at any point in the dynamics one of the two checkerboards would be empty.

I have developed a reversible diffusion algorithm, the details of which and the quantitative aggreement between the implementation and the predictions of the diffusion equation (Fick's law) will be discussed in Chap. 6. Our diffusion algorithm is also linear so it can be solved in the indentical manner. In addition our method is much more tractable from a computational standpoint and it preserves microscopic

reversibility.

5.2.2.3 Dynamical spin systems

Another area where reversible CA's have influenced our understanding of physics is the dynamics of spin systems. The Ising model[64] is perhaps the most studied model in statistical physics. It models a classical spin system on a lattice with nearest neighbor interactions. The Hamiltonian for the ferromagnetic Ising model is

$$\mathcal{H} = -\sum_{\langle i,j \rangle} J_{ij}\sigma_i\sigma_j - \sum_i H\sigma_i, \tag{5.3}$$

where the σ_i denote the spins, $\langle i,j \rangle$ the sum over nearest neighbors, $J_{ij} > 0$ the coupling between spins, and H an external magnetic field. For a two dimensional system, at high temperatures the entropic contribution to the free energy ($F = E - TS$) dominates and the system is disordered. The system undergoes a second order phase transition as the temperature is lowered. At the critical temperature the energetic contribution to the free energy becomes prominent and the spins begin to align.

There are various algorithms for generating equilibrium spin configurations, but the Ising model is inherently a static model. A physically motivated dynamics was introduced to the Ising model by Cruetz[18, 19], and independently discovered by other researchers[123, 100]. The dynamics is local, invertible, and conserves bond energy. In Cruetz's original model he considered "heat demons" which are able to absorb or supply a token of energy and thus allow spins to flip while maintaining the overall energy of the system constant. Since the total energy is constant, it is a microcanonical system. In the limit where the heat demons have infinite heat capacity (and if they move about the lattice) the system transitions continuously from microcanonical to canonical. The demons undergo a deterministic evolution hence the overall dynamics is invertible. Cruetz introduced the idea of storing information in a heat bath. Heat bath techniques have been used on occasion since then to construct

reversible computer models of physical phenomena[112].

More standard for CA simulations of the dynamical Ising model is to consider the microcanonical situation where there are no demons (or equivalently where the demons cannot accept or supply energy). Consider a two dimensional square lattice. If a spin is aligned with two of its neighbors and anti-aligned with the other two, the spin can flip without changing the energy. If we apply this dynamics in a checkerboard fashion, all the spins on the "red" squares can undergo the flip decision while those on the "black" squares stay fixed, leaving the total energy unchanged. Then the "black" squares update while the "red" are fixed. If we update the same checkerboard twice in a row, we exactly invert the dynamics. This model can be shown to support simple harmonic waves, which can be strikingly illustrated when the system is started from specific initial conditions[88].

I will discuss my own implementation of a heat bath system in the following chapter, Chap. 6, and an extension of microcanonical spin dynamics to the Potts model in Chap. 9.

5.2.2.4 Other physical models

Aside from the models reviewed above there are many other CA rules of interest for physical modeling. There is a an invertible one dimensional rule with the same Ising conservation law discussed above (invented by Bennett, and called the "scarves" rule). CA's have been used to model electromagnetic fields[111]. Recent efficient simulations of polymer dynamics have been accomplished with CA's[114, 113]. See Ref. [121] for a collection of other relevant models.

Chapter 6

Reversible Aggregation[1]

We introduce a lattice gas model of cluster growth via the diffusive aggregation of particles in a closed system obeying a local, deterministic, microscopically reversible dynamics. This model roughly corresponds to placing the irreversible Diffusion Limited Aggregation model (DLA) in contact with a heat bath. Particles release latent heat when aggregating, while singly connected cluster members can absorb heat and evaporate. The heat bath is initially empty, hence we observe the flow of entropy from the aggregating gas of particles into the heat bath, which is being populated by diffusing heat tokens. Before the population of the heat bath stabilizes, the cluster morphology (quantified by the fractal dimension) is similar to a standard DLA cluster. The cluster then gradually anneals, becoming more tenuous, until reaching configurational equilibrium when the cluster morphology resembles a quenched branched random polymer. As the microscopic dynamics is invertible, we can reverse the evolution, observe the inverse flow of heat and entropy, and recover the initial condition. This simple system provides an explicit example of how macroscopic dissipation and self-organization can result from an underlying microscopically reversible dynamics.

[1]This work is reprinted from Phys. Rev. E **60** (1) 264–274 (1999). My coauthor is N. H. Margolus.

We present a detailed description of the dynamics for the model, discuss the macroscopic limit, and give predictions for the equilibrium particle densities obtained in the mean field limit. Empirical results for the growth are then presented, including the observed equilibrium particle densities, the temperature of the system, the fractal dimension of the growth clusters, scaling behavior, finite size effects, and the approach to equilibrium. We pay particular attention to the temporal behavior of the growth process and show that the relaxation to the maximum entropy state is initially a rapid non-equilibrium process, then subsequently it is a quasistatic process with a well defined temperature.

6.1 Microscopic reversibility and pattern formation

Pattern formation is an intrinsically dissipative process[97], however the laws of physics are microscopically reversible: there is no dissipation at the microscopic scale. In this paper we describe a simple system which organizes into patterns through microscopically reversible dynamics, hence it also models how dissipation arises (i.e., how information flows between the macroscopic and the microscopic degrees of freedom). This system provides a clear example of how to reconcile the macroscopic irreversibility that gives rise to patterns with the microscopic reversibility adhered to by physical processes. Motivated by the desire to understand this general issue, we study specific details of the model, focusing on transitions in the resulting growth morphology and the approach to thermodynamic equilibrium.

We have previously observed several examples of reversible cellular automata dynamics which produce large scale order through microscopically reversible dynamics[121, 88, 27]. In contrast other research in the field of pattern formation has focused on irreversible microscopic mechanisms, with examples ranging from crystal growth[125, 71], to Turing patterns in chemical reactions[99], to patterns formed by growing bacterial

colonies[38], to kinetic growth problems[33, 34, 129].

Here we model cluster growth by reversible aggregation (RA), generalizing the irreversible Diffusion Limited Aggregation model (DLA)[129] to include contact with a heat bath. Particles, which are initially diffusing on a two-dimensional lattice, stick upon first contacting a cluster member and release heat which then diffuses about a superimposed lattice representing the heat bath. The two subsystems exchange only heat and together form a closed system. The release of heat transfers entropy from the aggregating system (which is becoming ordered) into the heat bath (which was initially empty). When the heat bath is nearly empty the model is essentially equivalent to the canonical DLA formulation (analogous to a supercooled gas crystalizing in a far from equilibrium situation). Hence the RA growth cluster initially resembles a typical DLA cluster. As the heat bath becomes populated, singly connected cluster members are able to absorb heat and evaporate. As the effect of evaporation becomes significant the RA and DLA models diverge. The RA dynamics is exactly invertible: at any point we can invert the dynamics and run backwards, observing the flow of heat from the heat bath back into the gas-crystal system until we recover the exact initial condition.

The population levels of the heat bath and of the aggregate initially grow linearly in time, quickly reaching stable values which remain very nearly constant for the remainder of the evolution. The energy of each subsystem is a function only of the population levels, independent of the physical configuration of the particles. Hence, once the population levels stabilize, the rate of energy exchange (which is entirely in the form of heat) between the heat bath and the gas-aggregate system becomes so slow that we can characterize the subsequent dynamics as a quasistatic process, with a well defined temperature at all times.

The aggregate mostly forms while the heat bath is at a lower temperature than in the quasistatic steady-state. Hence, after the population levels stabilize, the cluster slowly anneals. The cluster morphology, which initially resembles a DLA cluster,

gradually becomes more spread out and tortuous, until it ultimately resembles a branched polymer with quenched randomness. The two timescales that characterize the growth process are separated by two orders of magnitude. Initially, the population levels quickly reach a quasistatic steady-state. Subsequently, the aggregate slowly anneals until reaching the ensemble of configurations corresponding to the highest entropy macrostate (the branched polymer).

Aside from insight into microscopically reversible mechanisms that give rise to macroscopic patterns, the development of invertible dynamics and algorithms has technological significance in pushing down the barrier to atomic scale computing. Each bit of information erased at temperature T releases at least $T\Delta S = k_B T \ln 2$ units of heat into the environment[76]. Heat is created in proportion to the volume of the computer, yet heat leaves the computer only in proportion to the surface area. Hence, as logic gate density in computers increases, the use of an invertible dynamics (which does not erase information and hence does not need to produce heat) will be required to keep the mechanical parts from burning up[8, 47]. From a more pedagogical viewpoint, discrete computer models of reversible microscopic dynamics provide a laboratory for studying non-equilibrium statistical mechanics and the approach to equilibrium. These models let us explore physically plausible dynamics for non-equilibrium systems (i.e., discrete dynamics which are microscopically reversible and thus automatically obey Liouville's Theorem). A particularly instructive example of this approach is the formulation of a dynamical Ising model[19, 100, 123]. However, more widely used in physics are discrete, reversible models of fluid flow such as the HPP and FHP lattice gases[59, 49]. For a recent discussion of modeling physical phenomena with reversible computer models see Ref. [88]. For a recent discussion of macroscopic irreversibility and microscopic reversibility see Ref. [78]. For a recent discussion of techniques for the explicit construction of reversible models in statistical mechanics see Ref. [82]; but note that closely related techniques were discussed in the early 1980's by Fredkin (as discussed in Ref. [86]).

The initial sections of this manuscript describe our model; the middle contain a mathematical formulation of the model; the final, the empirical results. Specifically, Sec. 6.2 describes the detailed dynamics, including the subtleties of constructing an invertible dynamical model and implementation issues. In Sec. 6.3 we discuss the macroscopic limit of an analytic formulation of the model and establish the reaction-diffusion equations describing the system. In Sec. 6.4 we treat the reaction-diffusion equations in the mean field limit and compare predictions for equilibrium densities of particles to empirical measurements. Empirical measurements of temperature are presented in Sec. 6.5.1, with emphasis on the quasistatic nature of the annealing portion of the growth process. In Sec. 6.5.2 we study the evolution of the fractal dimension of the clusters and thus quantify the change in growth morphology as the clusters relax to the maximum entropy state. We conclude with a discussion of limitations and possible modifications of our model.

6.2 Modeling aggregation

6.2.1 Diffusion Limited Aggregation

Diffusion Limited Aggregation (DLA) [129] is a conceptually simple model which serves as a paradigm for some aspects of kinetic growth phenomena. Several comprehensive reviews of DLA have been written. In particular see Ref. [116] for a clear presentation of the basics, Ref. [77] for details on physical mechanisms, and Ref. [36] for a review of real-space renormalization group approaches to DLA.

The typical scenario for DLA begins with a vacant two-dimensional lattice initialized with a single stationary seed particle, which is the nucleation site for a growth cluster. Moving particles are introduced from the edges of the lattice, following random walks along the lattice sites. When a moving particle lands on a site adjacent to a stationary seed particle (an active site) it sticks (i.e., the moving particle aggregates and becomes a stationary seed particle). The frozen aggregate particles constitute

the solid (crystal) phase, and the moving particles constitute the gas phase. Aggregation hence consists of a particle undergoing an irreversible transition from gas to solid. Gas particles are usually introduced in a serial manner: only one gas particle is diffusing at a time. However, to take advantage of parallel computational resources, parallel models of DLA have been studied in which multiple particles are diffusing at once[127, 95]. In the dilute particle limit these models recover the serial DLA model exactly.

With the first aggregation event the DLA cluster grows from one single to two adjacent sites. The presence of the second cluster member eliminates certain paths along which random walkers could approach the first, with the effect that the probability of sticking at either end (tip) of the cluster is enhanced, whereas the probability of sticking along the edge of the cluster is reduced. As particles continue to aggregate creating new cluster tips and edges, the probability to stick at the tips continually outweighs the probability to stick along the edges. This leads to branching. A second influence on the growth morphology comes from shadowing: the probability for a particle to diffuse into the center of the growth cluster before encountering an active site becomes negligible as the cluster grows in size. Hence the outer tips grow most rapidly. As a result the growth aggregate rapidly assumes a bushy and branching, random fractal structure, resembling frost on a window pane, the branching of neurons, and many other branched structures found in nature.

6.2.2 Reversible Aggregation

Our goal is to introduce a reversible, deterministic model of growth by aggregation, where reversible means that from any state of our system we can recover the previous state exactly. We must address the subtleties of making each component step invertible, including steps which realize stochastic processes. As we discuss below, the same mechanisms that are employed in our model in order to ensure exact conservation of energy, particle number, and other constraints, also make it easy to incorporate in-

vertibility. The stochastic component of the model is diffusion, which is modeled as a sequence of invertible "random walks" based on a deterministic algorithm using an invertible pseudorandom number generator.

Information about exactly when and where a particle undergoes a phase transition is stored in the heat particles. The idea of storing information in a heat bath was introduced by Creutz[19] to explore the connection between the microcanonical and canonical ensembles for the dynamical Ising model. Heat bath techniques have been used on occasion since then to construct reversible computer models of physical phenomena[112].

6.2.2.1 Overview of the model

To construct the reversible aggregation (RA) model we begin with a parallel DLA system (similar to that described above) and add degrees of freedom at each site, corresponding to a distributed heat bath. The latent heat released during each aggregation event can then be explicitly represented. In the RA dynamics, whenever a random walking gas particle lands on a site with exactly one nearest neighbor crystal particle, it will stick only if there is room locally in the heat bath to accept the latent heat it will release as the particle transitions from the gas phase to the crystal phase. The heat is released in quantized units called heat particles, with one heat particle released for each aggregation event. These heat particles diffuse amongst themselves (i.e., they undergo random walks along the lattice sites, independently of the gas particles). Explicitly modeling the latent heat released upon aggregating provides a mechanism for modeling the inverse process: a diffusing heat particle which contacts a susceptible crystal particle (a crystal particle which has only one nearest neighbor crystal particle) is absorbed while the crystal particle evaporates to become a gas particle which then diffuses away.

The restriction on the dynamics that aggregation and evaporation events can occur only when exactly one nearest neighbor is a crystal particle means that only

one crystal bond is ever formed or broken when a single lattice site is updated. As each aggregate particle contributes one crystal bond to the aggregate, and there is no further potential energy contribution, the energy of the aggregate is a function only of the number of aggregate particles, independent of their configuration. Moreover this constraint has two direct implications for the growth morphology. The first is that evaporation can only occur for particles which are singly connected to the growth cluster, and so the aggregate cannot break off into disconnected clusters. The second is that it introduces an excluded volume (i.e., no closed loops can be formed), thus we might expect the equilibrium cluster configuration to be similar to that of a polymer. Note that introduction of an evaporation mechanism in the RA model mitigates the shadowing effect that was important in determining the DLA growth morphology: crystal particles within the cluster can evaporate, thus introducing gas particles into the interior of the aggregate.

6.2.2.2 The detailed dynamics

The RA model is constructed with 7-bits of state at each lattice site. One bit, $N_c(\vec{x},t)$, denotes the presence or absence of a crystal particle at that site (i.e., $N_c(\vec{x},t)=1$ indicates presence, $N_c(\vec{x},t)=0$ absence). Two bits, $N_g^\gamma(\vec{x},t)$ where $\gamma=\{1,2\}$, denote the presence or absence of each of two gas particles. Two bits, $N_h^\gamma(\vec{x},t)$, denote the presence or absence of each of two heat particles. The final 2 bits, $\xi_g(\vec{x},t),\xi_h(\vec{x},t)$, are independent binary pseudorandom variables. The dynamics of the model consists of two kinds of steps: diffusion steps alternating with interaction steps.

The same kind of diffusion process is applied to the gas and heat subsystems simultaneously and independently, while the crystal particles remain unchanged. A given diffusion step consists of two parts: mixing and transport. During the mixing portion of the step, a binary random variable is used to determine whether or not the

two particle bits of that species at the site $(\vec{x}, t)$ are interchanged:

$$\begin{aligned} N_i^1 &= (1-\xi_i)N_i^1 + \xi_i N_i^2 \\ N_i^2 &= (1-\xi_i)N_i^2 + \xi_i N_i^1, \end{aligned} \tag{6.1}$$

where $i = g$ or $i = h$. During the transport portion of the step, every site replaces its first particle bit ($\gamma = 1$) with the first particle of its neighbor a distance k away on one side, and its second particle ($\gamma = 2$) with the one from the same distance neighbor on the opposite side. At even time steps, we use horizontal neighbors (i.e., the diffusion moves particles horizontally):

$$\begin{aligned} N_i^1(\vec{x}, t+1) &= N_i^1(\vec{x} + k\hat{x}, t) \\ N_i^2(\vec{x}, t+1) &= N_i^2(\vec{x} - k\hat{x}, t). \end{aligned} \tag{6.2}$$

At odd time steps we use vertical neighbors (i.e., substitute $\hat{y}$ for $\hat{x}$ in Eq. (6.2)). The only differences between the gas and heat diffusion are (1) each uses a separate binary random variable to control its mixing, and (2) the distance of the neighbor particle to be copied, k, can be chosen separately for each subsystem—this allows us to independently control the diffusion constants for the heat bath and for the gas (c.f. [15]).

Diffusion steps alternate with steps in which the two subsystems—gas-crystal and heat bath—interact allowing aggregation and evaporation. The rule at a single lattice site during an interaction step is that exactly one particle can aggregate or evaporate provided that

(a) there is exactly one crystal particle at one of the four nearest neighbor sites

(b) there is room at the site for a crystal particle (for aggregation) or for another gas particle (for evaporation), and

(c) the heat bits at the site can absorb (for aggregation) or supply (for evaporation)

one unit of heat.

Since the gas and heat particles will undergo a mixing step before transport, it makes no difference which of the two available gas particle positions a crystal particle is moved into when it evaporates, or which of the two possible heat particle positions a unit of heat gets put into. Defining this precisely will, however, become important when we discuss invertibility.

The interaction rule described thus far would be sufficient if we updated just one lattice site at a time. If, however, all sites on the lattice are updated simultaneously, then the global dynamics no longer obeys the "single bond" constraint—that at any site where particles aggregate or evaporate exactly one crystal bond is formed or broken. For example, suppose that the tip of a crystal branch evaporates just as a gas particle condenses next to it. Each of these events would separately preserve the constraint, but the two simultaneous events result in the addition of a disconnected crystal particle which has no other crystal particle immediately adjacent to it. We can easily avoid this difficulty by holding the values at the adjacent sites fixed during a step in which we let the subsystems interact at a given lattice site, since the interaction step has a nearest neighbor range. In other words, we perform a *checkerboard updating*: all of the lattice sites in which the x and y lattice coordinates add up to an even number are updated using our single site interaction rule, while the odd sites are held fixed, and then vice versa. Since nearest neighbors are held fixed during an interaction, the constraint is obeyed locally and thus it is also obeyed globally. The overall dynamical rule is summarized in Table 6.1—the various phases of the rule are applied consecutively.

Every phase of the rule described in Table 6.1 can be inverted. The transport portion of the step can be run backwards by simply moving all particles back into the sites they came from (i.e., inverting the directions of the transport in Eq. 6.2). The mixing portion of the step is easy to invert, given the same "random" binary variables that were used to determine which pairs of bits were originally swapped.

1. interact gas/heat/crystal at even sites
2. mix gas and mix heat separately
3. transport gas and heat horizontally
4. interact gas/heat/crystal at odd sites
5. mix gas and mix heat separately
6. transport gas and heat vertically

Table 6.1: The various phases of one step of the RA dynamics. Each phase is applied over the entire lattice simultaneously.

We simply swap exactly those pairs again. The pseudorandom portion of the system (which supplies the random data) can simply follow some invertible dynamics that is independent of the rest of the system—the rest of the system looks at the state of this subsystem but does not affect it—so this pseudorandom subsystem can be run *backwards* independently of everything else.

Making the interaction steps invertible is also straightforward. When a single gas particle turns into a crystal particle, we put the heat token that is released into the corresponding heat bit (i.e., the heat particle with the same value of γ), and thus we remember which of the two gas particle bits was initially occupied. If the corresponding heat bit is already occupied, the particle is not allowed to aggregate (even if the other heat bit is unoccupied). Similarly, a crystal particle is allowed to evaporate only if it can move into the gas bit with the same value of γ as the heat token being absorbed. If there are two gas particles at a site we impose the constraint that the particle with $\gamma = 1$ attempts to aggregate first; likewise if there are two heat particles at a site, the crystal particle attempts to absorb the $\gamma = 1$ particle first. This does not introduce a bias to the growth since we are mixing the $\gamma = 1$ and $\gamma = 2$ variables in an unbiased manner at each time step. With these refinements,

our interaction rule *applied to a single site* is its own inverse: if we apply it twice (without a diffusion step in between) we get back the state we started with. Since the interaction rule is applied in a checkerboard fashion, sites are updated independently: if we apply the rule a second time to the same checkerboard, it will undo the first application at every site.

Thus an inverse step consists of applying the inverses of the rule-phases in the opposite of the order listed—once one phase is undone, the previous phase can be undone. Each inverse step undoes one step of the forward time evolution. As we watch the inverse evolution, we see each heat particle retrace its path, to be in exactly the right location at the right time to uncrystalize the crystal particle which originally released it. Particles un-aggregate and un-diffuse and un-evaporate in a manner that exactly retraces their behavior in the forward evolution.

6.2.2.3 Implementation

The RA model was implemented on a special purpose cellular automata machine, the CAM-8[87], which was designed to efficiently perform large-scale uniform, spatially arrayed computations. On this machine, all simulations must be embedded into a lattice gas framework[59, 49], in which uniform data movement (data-advection) alternates with processing each site independently (site-update). For a 2-dimensional model such as ours, sheets of bits move coherently during the advection phase: corresponding bits at each site all move in the same direction by the same amount. The boundaries are periodic—bits that shift past the edge of the lattice reappear at the opposite edge. After moving the bits, we perform the site update phase. During this phase, the bits that have landed at each lattice site are updated in a single operation by table lookup: the bits at each lattice site are used as an index into a table that contains a complete listing of which state should replace each possible original state. Both the data movement and the lookup table can be freely changed between one lattice updating step and the next.

Our model requires 7 bits of state to appear at each site in our $L \times L$ lattice. Using random data generated by a serial computer, the bits which correspond to the gas particles are initially randomly filled with a 4% density of particles and the bits which correspond to the binary random variables with a 50% density of particles. One crystal particle is placed at the center site of the lattice. The heat bath is initially empty.

The dynamics on the pseudorandom subsystem is very simple: each of the two random bit-planes (each consisting of all the η_g's or all of the η_h's) are simply shifted by some large amount at each time step. We could choose the amount and direction of each shift at random for each step of updating, using a reversible random number generator running on the workstation controlling the simulation. Instead, the simulations discussed here simply shift the bit planes by a large and fixed amount at each step, making sure that the the x and y components of the two shifts are all mutually co-prime, as well as being co-prime with the overall dimensions of the lattice. Thus to run the random subsystem backwards, we just reverse the direction of the shifts.

The checkerboard updating is accommodated by adding an eighth bit to each lattice site, and filling these bits with a checkerboard pattern of ones and zeros. In our rule the various subsystems are allowed to interact only at sites marked with a one. To change which checkerboard is marked for updating, we simply shift the checkerboard-marker bit-plane by one position in the $+x$ direction.

The rule described in Table 6.1 turns into two lattice-gas steps on CAM-8. The first three phases listed in Table 6.1 are done during one step, and the next three in the second step. The data movement is part of each step: note that each of the two steps uses the same lookup table applied to each lattice site, but slightly different data movement. To run backwards, we use the inverse lookup table, and the inverse data movement. Note that in the discussion of experimental results, everything in Table 6.1 is counted as a single step.

CAM-8's event counting hardware was used to monitor simulation parameters

while the simulations ran. Including event counting, the 8-processor CAM-8 performed about 10^8 site update operations per second for this model.

6.3 The macroscopic limit

The dynamics of the RA model, described in detail above, can be succinctly presented in an analytic framework. We develop this framework first in terms of the discrete space, time, and occupation number variables. We then ensemble average over the occupation numbers and take the continuum limit of the space and time variables to establish the reaction-diffusion equations for the system.

As discussed in Sec. 6.2.2.2 there are seven bits of state:

$$N_c(\vec{x},t), N_g^\gamma(\vec{x},t), N_h^\gamma(x,t), \xi_g(x,t), \quad \text{and} \quad \xi_h(x,t), \quad \text{where} \quad \gamma = \{1,2\}.$$

They correspond respectively to one bit of crystal, two bits of gas particles, two bits of heat particles, and two bits of random data. $N_i^\gamma(\vec{x},t) = 1$ indicates the presence of species i at location $\vec{x}$ and time t, in channel γ, and $N_i^\gamma(\vec{x},t) = 0$ indicates the absence. The absence or presence of a crystal particle is denoted by $N_c(\vec{x},t) = \{0,1\}$ respectively. The total number of particles of species i at time t present on the lattice is denoted by $\mathcal{N}_i(t) = \sum_{\vec{x},\gamma} N_i^\gamma(\vec{x},t)$, where the sum is over all of the lattice sites and the two particle channels.

There are no external sources or sinks for any of the three species represented (the gas, crystal, and heat species). Energy is only exchanged between the gas-crystal and the heat bath subsystems. Thus the complete system is isolated. Conservation of the total number of gas and crystal particles implies that

$$\mathcal{N}_g(t) + \mathcal{N}_c(t) = \mathcal{N}_g(0) + \mathcal{N}_c(0). \tag{6.3}$$

Conservation of the total energy of the system implies that

$$\mathcal{N}_g(t)\varepsilon_g + \mathcal{N}_c(t)\varepsilon_c + \mathcal{N}_h(t)\varepsilon_h = \mathcal{N}_g(0)\varepsilon_g + \mathcal{N}_c(0)\varepsilon_c + \mathcal{N}_h(0)\varepsilon_h, \tag{6.4}$$

where ε_i represents the energy (kinetic and potential) per particle of species i (notice that there is no configurational contribution to the energy of the crystal). As discussed in Sec. 6.2 each aggregation event releases one heat particle (likewise each evaporation event absorbs one heat particle), thus $\varepsilon_h = \varepsilon_g - \varepsilon_c$ and moreover

$$\mathcal{N}_h(t) = \mathcal{N}_c(t) - \mathcal{N}_c(0). \tag{6.5}$$

(Note that $\mathcal{N}_h(0) = 0$ and $\mathcal{N}_c(0) = 1$ in our experiments).

To facilitate the description of the dynamics, we introduce a functional equation which is $+1$ at any site where a particle is about to crystalize, -1 at a site where a particle is about to evaporate, and 0 otherwise. The functional,

$$\mathcal{F}^\gamma \left[N_g^\gamma(\vec{x},t), N_h^\gamma(\vec{x},t), N_c(\vec{x},t), \{N_c(\vec{x}+\hat{e}_k,t)\} \right],$$

is evaluated on a neighborhood of lattice sites surrounding some given position $\vec{x}$ at a given time t (the notation $\{N_c(\vec{x}+\hat{e}_k,t)\}$ refers to the set of values of N_c for the nearest neighbors of the point $\vec{x}$).

$$\begin{aligned}\mathcal{F}^\gamma(\vec{x},t) &= N_g^\gamma(\vec{x},t)[1-N_h^\gamma(\vec{x},t)][1-N_c(\vec{x},t)]\sum_{j=1}^{d} N_c(\vec{x}+\hat{e}_j,t)\prod_{k\neq j}[1-N_c(\vec{x}+\hat{e}_k,t)] \\ &- [1-N_g^\gamma(\vec{x},t)]N_h^\gamma(\vec{x},t)N_c(\vec{x},t)\sum_{j=1}^{d} N_c(\vec{x}+\hat{e}_j,t)\prod_{k\neq j}[1-N_c(\vec{x}+\hat{e}_k,t)]\,. \end{aligned} \tag{6.6}$$

Here $\hat{e}_j$ and $\hat{e}_k$ are the vector lattice directions of the nearest neighbor cells, and d the number of distinct lattice directions. For a two dimensional square lattice (i.e., the lattice used for the present implementation) $d = 4$ and the vector lattice directions

are $\{\hat{x}, -\hat{x}, \hat{y}, -\hat{y}\}$.

The first term in Eq. (6.6) equals 1 if a gas particle in channel γ is present at site $\vec{x}$ and time t, a heat particle in channel γ is absent at site $\vec{x}$ and time t, there is no crystal particle already at that site, and only one crystal particle is present at a nearest neighbor site. It is zero otherwise. The second term equals 1 if there is no gas particle in channel γ present at site $\vec{x}$ and time t, there is a heat particle in channel γ present at site $\vec{x}$ and time t, there is a crystal particle present at that site, and only one crystal particle is present at a nearest neighbor site. It is zero otherwise. The first and second terms are mutually exclusive (a heat particle in channel γ cannot be simultaneously present and absent, nor can a gas particle).

The dynamics consists of making the changes indicated by $\mathcal{F}^1$ and then $\mathcal{F}^2$, then applying a random permutation to mix $\gamma = 1$ and $\gamma = 2$, and then performing the streaming step to move the particles. The permutation mixes the $N_i^1(\vec{x}, t)$ and $N_i^2(\vec{x}, t)$ components in an unbiased way, so it is simpler to discuss the dynamics of a combined variable, $N_i(\vec{x}, t) = N_i^1(\vec{x}, t) + N_i^2(\vec{x}, t)$. Likewise, if we let

$$\mathcal{F}(\vec{x}, t) = [\mathcal{F}^1(\vec{x}, t) + \mathcal{F}^2(\vec{x}, t)][1 - \mathcal{F}^1(\vec{x}, t)\mathcal{F}^2(\vec{x}, t)/2],$$

the interaction portion of the dynamics at a single lattice site can be written

$$\begin{aligned} N_c(\vec{x}, t+1) &= N_c(\vec{x}, t) + \mathcal{F}(\vec{x}, t) \\ N_g(\vec{x}, t+1) &= N_g(\vec{x}, t) - \mathcal{F}(\vec{x}, t) \\ N_h(\vec{x}, t+1) &= N_h(\vec{x}, t) + \mathcal{F}(\vec{x}, t). \end{aligned} \tag{6.7}$$

The $[1 - \mathcal{F}^1(\vec{x}, t)\mathcal{F}^2(\vec{x}, t)/2]$ factor in the definition of $\mathcal{F}(\vec{x}, t)$ reflects the fact that only one transition occurs at a given site, even if two gas or two heat particles are present.

To construct continuous variables from the discrete ones, we consider the average short-term behavior of the system over an ensemble of independent realizations which

all have the same set of local particle densities. For each discrete variable, we let $n_i(\vec{x},t) \equiv \langle N_i(\vec{x},t)\rangle$, and denote the average of the functional as $\langle \mathcal{F}\rangle$. (This technique of averaging over many independent realizations, i.e., establishing the one particle density function, is commonly used to derive the lattice Boltzmann equation starting from discrete particle models of hydrodynamics[12, 62]).

With this notation, the average propagation of the gas and heat particles can be expressed as

$$n_i(\vec{x},t+1) = \frac{1}{4}\left[n_i(\vec{x}-\hat{x},t) + n_i(\vec{x}-\hat{y},t) + n_i(\vec{x}+\hat{x},t) + n_i(\vec{x}+\hat{y},t)\right]. \tag{6.8}$$

To establish the continuum limit, we Taylor expand. The terms involving the first derivatives cancel, leading to the result

$$\begin{aligned} n_i(\vec{x},t+1) &= n_i(\vec{x},t) + \sum_j \frac{|\Delta x|^2}{4}\frac{\partial^2}{\partial x_j^2}\left(n_i(\vec{x},t)\right) + O(\Delta x^3) \\ &= n_i(\vec{x},t) + \frac{|\Delta x|^2}{4}\nabla^2 n_i(\vec{x},t), \end{aligned} \tag{6.9}$$

where $i = g$ or $i = h$. Note that to order Δt the above equation is the standard diffusion equation[2]

$$\frac{\partial}{\partial t} n_i(\vec{x},t) = \frac{|\Delta x|^2}{4\Delta t}\nabla^2 n_i(\vec{x},t). \tag{6.10}$$

As discussed in Sec. 6.2.2.2, we can control the length of each diffusion step separately for the heat and for the gas particles. The heat particles execute random walks composed of individual steps of length k, whereas the gas particles execute walks of step length unity. Thus if the $|\Delta x|^2$ that appears in Eq. (6.9) and in Eq. (6.10) refers to the gas subsystem, then $|\Delta x|^2$ for the heat subsystem (and hence its diffusion constant) is a factor of k^2 larger.

[2]Note there is a more detailed derivation of the continuum limit for the diffusion process[15], discussed in detail in Sec. 5.2.2.2. For our purposes, keeping the equation to order Δt and $(\Delta x)^2$ is sufficient.

To proceed further, we will make the approximation that our average variables are independent. Then we are allowed to replace the average of a product of variables by the product of the average for each variable: $\langle ab \rangle = \langle a \rangle \langle b \rangle$, for a, b independent variables. This is the assumption of molecular chaos, which is also used to derive the lattice Boltzmann equation. With this approximation the average of the functional $\mathcal{F}^\gamma$ is simply

$$\begin{aligned} \langle \mathcal{F}^\gamma \rangle &= \left\{ n_g^\gamma(\vec{x},t) \left[1 - n_h^\gamma(\vec{x},t)\right] \left[1 - n_c(\vec{x},t)\right] - \left[1 - n_g^\gamma(\vec{x},t)\right] n_h^\gamma(\vec{x},t) n_c(\vec{x},t) \right\} \\ &\times \sum_j n_c(\vec{x} + \hat{e}_j, t) \prod_{k \neq j} \left[1 - n_c(\vec{x} + \hat{e}_k, t)\right]. \end{aligned} \tag{6.11}$$

Similarly, we can write down an expression for $\langle \mathcal{F} \rangle$.

To obtain the continuum limit of these averaged equations, we again use Taylor series approximations. In the diffusive regime, $\Delta t \sim (\Delta l)^2$, so we truncate the expansions at these appropriate orders. Let $\tilde{\mathcal{F}}$ be the continuum limit of $\langle \mathcal{F} \rangle$ (which we won't write out explicitly). Then from Eq. (6.7), we obtain

$$\frac{\partial}{\partial t} n_c(\vec{x},t) = \frac{1}{\Delta t} \tilde{\mathcal{F}}, \tag{6.12}$$

The other reaction-diffusion equations for our system can be obtained by proceeding as we did in Eq. (6.9). For example, under the full dynamics (which consists of both the interaction and diffusion phases)

$$\begin{aligned} & n_g(\vec{x}, t+1) + n_c(\vec{x}, t+1) \\ &= n_c(\vec{x},t) + \frac{1}{4} \left[n_g(\vec{x} - \hat{x}, t) + n_g(\vec{x} - \hat{y}, t) + n_g(\vec{x} + \hat{x}, t) + n_g(\vec{x} + \hat{y}, t) \right], \end{aligned} \tag{6.13}$$

since any particles present at a site at time $t+1$ were either already there at time t, or moved there. Expanding this exactly as in Eq. (6.9) and using Eq. (6.12), we get

$$\frac{\partial}{\partial t} n_g(\vec{x},t) = \frac{|\Delta x|^2}{4 \Delta t} \nabla^2 n_g(\vec{x},t) - \frac{\partial}{\partial t} n_c(\vec{x},t), \tag{6.14}$$

and similarly,

$$\frac{\partial}{\partial t} n_h(\vec{x},t) = \frac{k^2|\Delta x|^2}{4\Delta t}\nabla^2 n_h(\vec{x},t) + \frac{\partial}{\partial t} n_c(\vec{x},t). \quad (6.15)$$

Note that if $k = 1$, and we add the last two equations the variable $n_g(\vec{x},t)+n_h(\vec{x},t)$ obeys the diffusion equation, unaffected by the interaction between the subsystems (i.e. if we remove the distinctions between gas and heat, the combined variable simply diffuses without interacting).

To test the consistency between the microscopic diffusive dynamics implemented in our model and macroscopic descriptions given by Eq. (6.14) and Eq. (6.15), we empirically measured the diffusion coefficient for gas and for heat particles as they diffuse about the space. Each particle should execute a random walk. The variance of the distance from the origin in the $\hat{x}$ or $\hat{y}$ direction, σ_i^2, is proportional to the diffusion coefficient in that direction, D_i, where $i = g$ or $i = h$. The exact relation is $D_i = \sigma_i^2/4p$, where p is the number of steps taken. For an unbiased random walk, the variance of the net displacement from the origin is $\sigma_i^2 = k^2 p$, thus $D_i = k^2/4$. For the gas particles ($k = 1$) we find $D_g = (0.996 \pm 0.009)/4$. For the heat particles, with $k = 3$, we find $D_h = (9.00 \pm 0.08)/4$. Thus the ratio of the heat to the gas diffusion length $D_h/D_g = 9.0 \pm 0.1$, agreeing with the theoretically predicted value of k^2.

6.4 The mean field limit

The mean field limit corresponds to the "well-stirred reaction," meaning that the reacting species are uniformly spread throughout the space, and thus each particle feels the presence of the mean concentration of each species. For our system in equilibrium the gas particles and the heat particles *are* uniformly distributed throughout the space; it is only the crystal particles which do not obey this assumption. A uniform distribution means that there are no concentration gradients ($\nabla n_i(\vec{x},t) = 0$ for all $\vec{x}$ and i), and thus $\nabla^2 n_i(\vec{x},t) = 0$. Also we can drop the explicit $\vec{x}$ notation from the

argument of the variables: $n_i(\vec{x},t) = n_i(t)$.

Once the population level of the heat bath has reached the quasistatic steady-state, the concentrations of all three species will remain essentially constant and the systems will have a well defined temperature from then on, as discussed in Sec. 6.5.1. We denote the time to reach the steady-state (i.e. the time for the subsystems to reach the same temperature) as $\tau_{\rm T}$. We can now drop the explicit time notation from the arguments of the variables in steady-state: $n_i(t > \tau_{\rm T}) = n_i$. In this regime $dn_i/dt = 0$ and likewise $dn_i^\gamma/dt = 0$, thus $< \mathcal{F}^\gamma >= 0$:

$$\begin{aligned} 0 &= \langle \mathcal{F}^\gamma \rangle \\ &= 4\left[n_g\left(1-n_h\right)\left(1-n_c\right) - \left(1-n_g\right)n_h n_c\right] \cdot \left[n_c\left(1-n_c\right)^3\right]. \end{aligned} \tag{6.16}$$

There are three solutions to Eq. (6.16). Each solution corresponds to fixed point of the dynamics. Only one is in the regime of interest. The fixed point at $n_c = 0$ corresponds to the presence of only gas particles. The fixed point at $n_c = 1$ is not allowed by the aggregation conditions (the aggregate can not have any closed loops). The remaining fixed point predicts that the equilibrium condition is $n_g\left[1-n_h\right]\left[1-n_c\right] = \left[1-n_g\right]n_h n_c$. Noting that in the mean field limit ensemble averages equal spatial averages (i.e., $n_g(t) = \mathcal{N}_g(t)/L^2$), the constraints described in Eq. (6.3) and Eq. (6.5) can be written respectively as $n_g(t) = \mathcal{N}_g(0)/L^2 - n_c(t) + 1/L^2 \approx \mathcal{N}_g(0)/L^2 - n_c(t)$, and $n_c(t) = n_h(t) + 1/L^2 \approx n_h$. After incorporating these relations the equilibrium condition can be expressed as

$$\frac{n_c}{1-n_c} = \mathcal{N}_g(0)\frac{1}{L^2} + O(n_c^3). \tag{6.17}$$

Figure 6-1 is a plot of the equilibrium value $n_c/(1-n_c)$ versus the initial density of gas particles, $\mathcal{N}_g(0)/L^2$. The solid line is the mean field prediction, Eq. (6.17). The points were obtained empirically from our simulations of three different system sizes, $L = 128,\ 256$, and 512. The agreement between the three system sizes should be

noted.

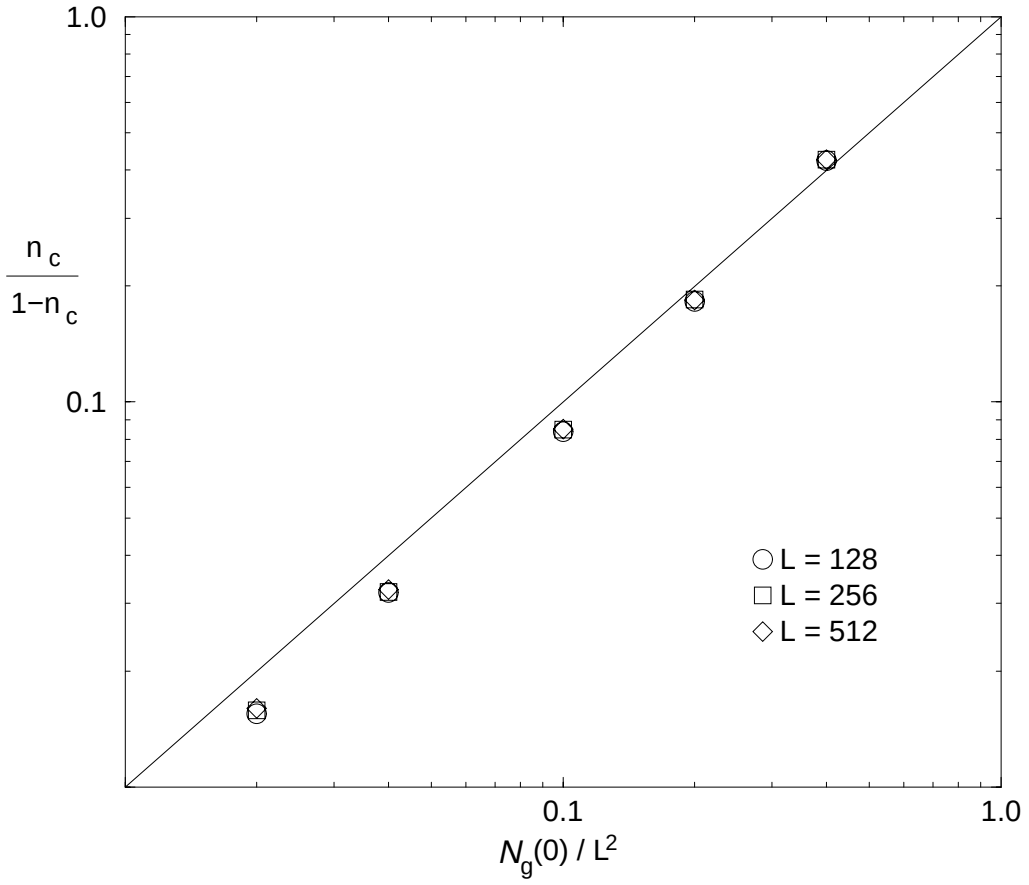

Figure 6-1: The empirically determined equilibrium value of $n_c/(1-n_c)$ as a function of the initial gas density, $\mathcal{N}_g(0)/L^2$, for systems of size $L = 128,\ 256,$ and 512. The solid line is the mean field prediction. Note that the error bars are the same size as the points.

The mean field approach makes predictions about the overall density of the system (hence the equilibrium temperature, as described below), but it does not make predictions about the growth morphology.

6.5 Empirical analysis

6.5.1 Temperature

The gas-aggregate subsystem and the heat bath subsystem together form a thermodynamically isolated system. These two subsystems are allowed to exchange energy only between themselves, and this energy is purely in the form of heat (ΔQ). As discussed in Sec. 6.2.2.1 and Sec. 6.3 the energy of the aggregate is a function only of the number of aggregate particles and is independent of their configuration. The total internal energy of the gas and heat particle species also is a function only of the number of particles of each species. Hence if invariant average population densities are achieved there is no further net exchange of heat between the subsystems, and they have by definition attained the same temperature.

The standard expression for the temperature of a two-level system[63], such as the heat bath in the RA model, follows directly from combining the definition of temperature ($1/T = \Delta S/\Delta E|_V$) with the microcanonical definition of entropy ($\Delta S = k_B \ln(\Omega_f/\Omega_i)$, where Ω denotes the number of microstates consistent with the macroscopic variables):

$$\frac{1}{T} = -\frac{k_B}{\varepsilon_h} \ln\left(\frac{n_h}{1-n_h}\right). \tag{6.18}$$

Directly computing the temperature of the other subsystem is more difficult, but we can infer its temperature from that of the heat bath (note the gas particles are free to diffuse over the crystal, hence there is no change in the accessible volume for the heat particles as the crystal changes size and conformation: the crystal does no work on the gas, $PdV = 0$).

The approach to temperature equilibrium and a closeup of the subsequent fluctuations in temperature are shown in Fig. 6-2. Figure 6-2(a) plots the mean occupancy of the heat bath versus the time step into the simulation, with the corresponding temperature (in units of $k_B T/\varepsilon_h$) displayed on the right vertical axis. The initial growth is linear, with a slope of about 1.8. It then levels off near the quasistatic

steady-state density of $n_h = 0.031$ (indicated in the figure as the dashed horizontal line). The results are averages over several independent realizations for three different system sizes, $L = 128,\ 256,$ and 512. Note that the three systems reach the same steady-state densities and hence the same temperature, but that the time to equilibrate depends on the system size. The data for the three systems was collapsed onto one curve by rescaling the time axis by the factor L^z, with $z = 1.8$. The time to reach the equilibrium temperature is $\tau_{\rm T} \sim 10L^z$. Note that this scaling behavior has an exponent which is slightly smaller than the diffusion exponent: the diffusion time is proportional to L^2. As discussed in Sec. 6.5.2 the fractal dimension at time $\tau_{\rm T} \approx 1.8$: the time seems to scale with the fractal dimensionality instead of the Euclidean dimensionality of the space.

To study the details of the subsequent fluctuations we focus on the largest system size, 512×512. As mentioned, during the initial period the growth of the heat bath population (and the size of the aggregate) is linear. It levels off at about 8270 particles on average, in a time which is less than 10^6 steps. The population continues to grow extremely slowly after this point, rising by an average of 1.8 ± 0.1 particles every 1×10^7 steps, as determined by a linear regression based on about 8000 data points taken at equally spaced intervals in the regime where $t > \tau_{\rm T}$. The probability that the slope is actually zero is 3×10^{-8}, as determined by a t-statistic comparing the ratio of the obtained slope to the sum of squares differences. Figure 6-2(b) shows a scatter plot of every third of the 8000 data points, overlayed by a straight line indicating the results of the linear regression on all 8000 points (only every third point is shown for visualization purposes: showing all the points results in a dense black cloud). The actual number of particles in the heat bath is indicated on the left axis, the corresponding temperature is given on the right. Although the temperature of the heat bath is not constant, it is very nearly so. Once the population levels stabilize the subsequent dynamics (i.e., the relaxation to the maximum entropy state for the crystal) is clearly a quasistatic process. The crystal does continue to exchange heat

with the heat bath when it anneals, but the net heat exchange is essentially zero (the net heat exchange rate is $\sim 2 \times 10^{-7}$ particles per update of the space).

We measured also the fluctuations in population levels of the heat bath for the $L = 256$ and 128 systems, in the corresponding regimes. Consistent with the $L = 512$ system, we find the population rises by an average of 2.1 ± 0.1 particles every 1×10^7 steps.

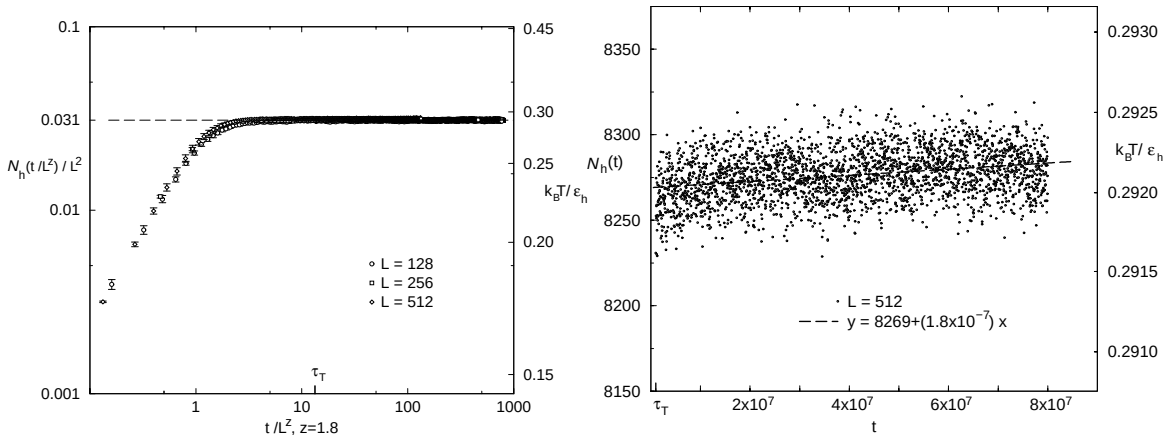

Figure 6-2: (a) The mean density of heat bath particles as a function of time into the simulation, plotted for $L = 128$, 256, and 512. The corresponding temperature is given on the right vertical axis. The initial growth of the heat bath density is linear, with a slope of about 1.8. Note that the steady-state density of the heat bath, and hence the temperature, is equal for all three system sizes, yet the time to equilibrate scales with the system size as $\tau_T \sim 10L^{1.8}$. (b) The actual average values of the total population of the heat bath as a function of time, for every third time measured beyond τ_T. The dashed line is the result of a linear least squares regression on all of the data. Note the slight drift upward with time, of about 2 particles per 10^7 steps.

6.5.2 Fractal dimension

The aggregate formed primarily while the heat bath contained less energy than its equilibrium level. Hence, if we continue running the dynamics, the cluster anneals; it evolves from a DLA-like cluster to a quenched branched random polymer structure. To quantify the cluster structure we calculate the fractal dimension of the aggregate

and specifically how the fractal dimension changes as a function of the time into the simulation.

We measure the fractal dimension using a box-counting procedure which requires that we first establish the center of mass of the growth aggregate (which is typically not the initial seed particle, as the center of mass diffuses about the space as the aggregate anneals). An imaginary window box of edge length l is defined and centered on the center of mass. The number of lattice sites within that window that contain a crystal particle, $\mathcal{N}_c(l)$, is tallied. The window size is increased and the count retallied. This procedure is iterated until the number of crystal particles contained no longer increases with window size. Before saturation, the number of particles contained should increase with some power of the window size

$$\mathcal{N}_c(l) \propto l^{d_f}. \tag{6.19}$$

The exponent d_f is the fractal dimension.

The RA cluster should initially resemble a parallel, irreversible DLA cluster of equivalent mass. Figure 6-3(a) shows a typical RA cluster for the $L = 512$ system at the time $t = \tau_{\mathrm{T}}$, which is the time for the mass of the RA cluster to stabilize at essentially its final mass ($\mathcal{N}_c \simeq 8270$). Figure 6-3(b) shows a typical DLA cluster of equivalent mass. Both systems were initialized with a 4% density of diffusing gas particles. The gas particles still present at this stage of the evolution are shown as the small dots in the figure. Note that for the RA system the gas particles are distributed throughout the space, yet for the DLA system very few gas particles penetrate the region defined by the edges of the cluster. The RA cluster has experienced enough annealing by the time $t = \tau_{\mathrm{T}}$ to have a fractal dimension different than that of the DLA cluster, yet the radii of both clusters are comparable and approximately equal to a quarter of the length of the system ($r \approx L/4$). The RA cluster morphology is still far from its final morphology.

Figure 6-4 shows the box-counting results obtained for both models in the regime described above and pictured in Fig. 6-3. The top curve is for DLA, the bottom for RA. Both models were implemented on a $L = 512$ size system. The vertical axis is the mass contained in the window, $\mathcal{N}_c(l)$, the horizontal is the window size l. The curve for the DLA system is the average of 10 independent DLA clusters of mass $\mathcal{N}_c \simeq 8270$. The curve for the RA system is the average of ten independent RA clusters sampled at time $t = \tau_{\rm T}$. The slope of the curve corresponds to the fractal dimension and was determined via a linear least squares fit. Consistent with past numerical studies of DLA[116, 57], we find that the fractal dimension for DLA clusters is $d_f^{\rm DLA} = 1.71 \pm 0.01$ (for difficulties associated with determining the fractal dimension of DLA see the detailed discussion in Ref. [36]). For the RA clusters the fractal dimension is $d_f^{\rm RA}(t = \tau_{\rm T}) = 1.81 \pm 0.03$. A line with this slope is shown overlaying each respective curve.[3]

The RA cluster is less dense than the DLA cluster in the area immediately surrounding the initial aggregation site, however the radii of both clusters are comparable. Several of the aggregate particles in the RA cluster have annealed away from the center to occupy the region between the center and the edge of the cluster. Hence, at the time depicted in Fig. 6-3, the RA cluster has a higher fractal dimension than the DLA cluster. The only constraints on the cluster are the number of particles and the connectivity. As there are more ways to have connected clusters of a specified particle number in an area of larger radius, the RA cluster evolves from the dense, bushy DLA-like structure shown in Fig. 6-3, to a tenuous structure which occupies more of the available lattice (with the initial increase in fractal dimension being a transient behavior). We ultimately expect to observe a diffuse structure with just a few meandering vines which can access more of the available configuration space.

[3]Note that in the limit where the cluster size approaches infinity, a cluster grown via a parallel implementation of DLA will exhibit a crossover from fractal to two-dimensional. In our simulations of parallel DLA the initial density of gas particles is dilute enough, and the cluster size small enough, that we do not see any crossover effects and we obtain smooth scaling curves as shown in Fig. 4.

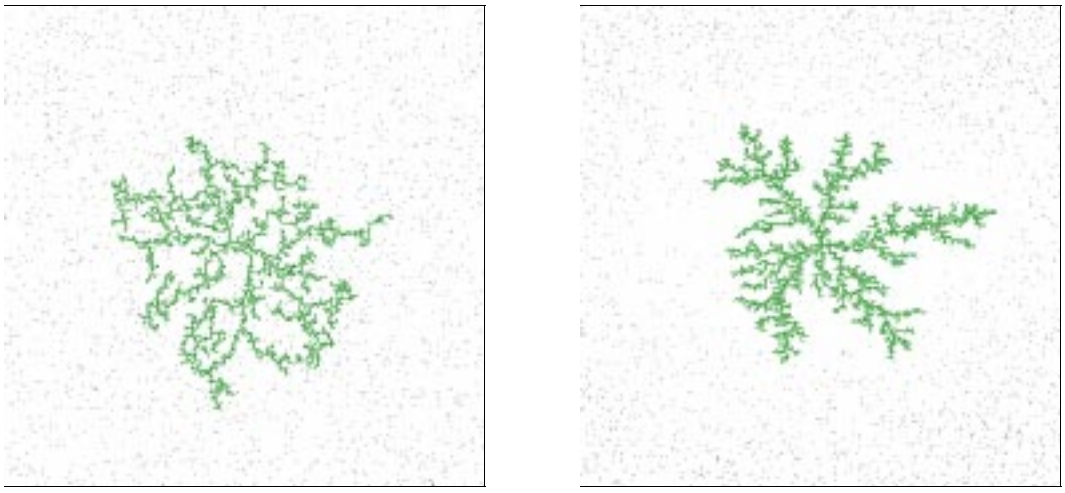

Figure 6-3: Two growth clusters of the same mass, $\mathcal{N}_c \sim 8270$. (a) A cluster grown via the RA model, pictured at time $t = \tau_{\rm T}$, where $\tau_{\rm T}$ is the time for the heat bath and gas-aggregate system to reach the same temperature. (b) A parallel DLA cluster. Note the gas particles, which are shown as the small dots. For the RA system the gas particles are distributed throughout the space, yet for the DLA system very few gas particles penetrate the region defined by the edges of the cluster.

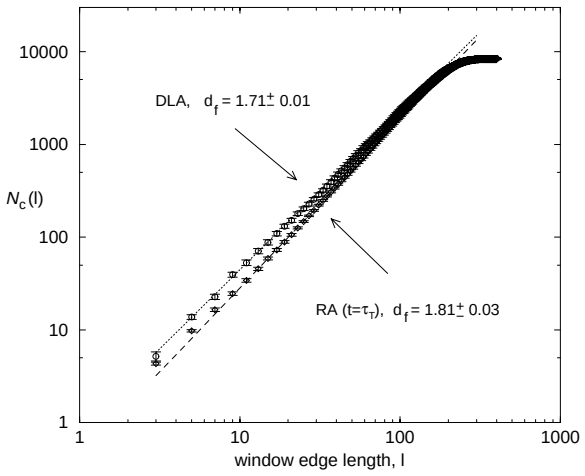

Figure 6-4: The number of aggregate particles contained in a box of length l, as a function of l. The slope of the line is the fractal dimension. The top curve is for parallel DLA clusters of mass $\mathcal{N}_c \simeq 8270$. The bottom curve is for the RA clusters sampled at $t = \tau_{\rm T}$. Examples of these clusters are pictured in Fig. 6-3.

As the time into the simulation advances, the density of the growth aggregate decreases, the radius of the aggregate increases, and hence the fractal dimension decreases. Figure 6-5 shows a typical RA growth cluster at the time $t = 80\tau_{\mathrm{T}}$ timesteps. Note that the structure does resemble meandering vines. Also the radius of the cluster is comparable to half of the lattice size ($r \approx L/2$).

Figure 6-6 is a plot of the average fractal dimension for RA clusters as a function of time into the simulation, for all three system sizes. The measurements reported below are averages over 5 independent realizations for the $L = 512$ system, 10 independent realizations for the $L = 256$ system, and 10 for the $L = 128$ system (i.e., averages over either 5 or 10 independently generated, large clusters). The data points and standard errors shown in the plots are the average and standard error over the set of independent realizations.

The fractal dimension is initially very close to the fractal dimension for DLA. We then observe a slight increase in the fractal dimension as the cluster center begins to anneal (an example is the RA cluster shown in Fig. 6-3), then a gradual decrease in the fractal dimension until it converges upon an equilibrium value. The solid line is drawn to denote the equilibrium value upon which results for the three system sizes are converging, $d_f = 1.6$. Using Flory-type scaling arguments it has previously been shown that a quenched branched polymer obeys the scaling relationship $N \sim R^{d_Q}$, with $d_Q = [2(D+2)]/5$[22]. Here R represents the characteristic end-to-end distance of a polymer, and D the dimension of the space. R can be taken in direct analogy to l in Eq. (6.19), which defines the end-to-end distance of the window of interest. For $D = 2$ the exponent $d_Q = 1.6$. We should note that an exact result was obtained for quenched polymers in $D = 2$, $d_Q = 1.64$[98], which is slightly larger. Flory-type scaling has also been studied for annealed branched polymers and the scaling exponent determined to be $d_A = (3D+4)/7$[54]. For $D = 2$, $d_A = 1.43$. One might expect to observe a crossover from quenched to annealed behavior for the equilibrium RA growth clusters as we go from the large to the small system sizes, but we did not

see this for the system sizes investigated.

Note that the time axis in Fig. 6-6 is rescaled by $L^{1.8}$ in order to match that of Fig. 6-2. Neither the fractal dimension nor the equilibrium temperature exhibit finite size effects as far as we can determine within the precision of our measurements.

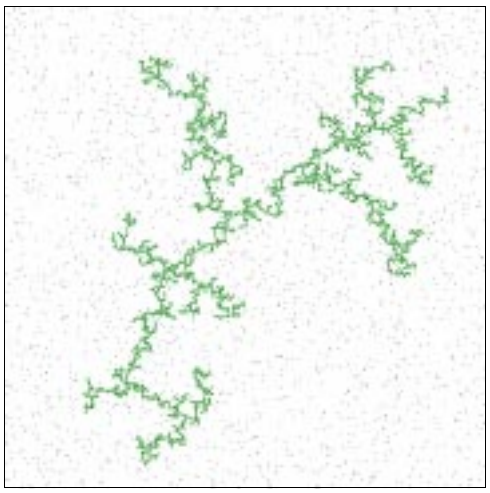

Figure 6-5: A growth cluster grown via the RA model, pictured at time $t = 80\tau_{\mathrm{T}}$. The fractal dimension for this cluster ($d_f = 1.63 \pm 0.02$) has seemingly reached the asymptotic value.

6.6 Discussion

We have presented a microscopically reversible model which exhibits macroscopic pattern formation. In this model, entropy initially grows rapidly with time, and then subsequently grows exceedingly slowly—the slow relaxation can be characterized as a quasistatic isothermal process. The morphology of the aggregate formed by this dynamics changes markedly with time, evolving from a pattern having a conformation and fractal dimension similar to that of an irreversible DLA system, to a pattern characteristic of a branched quenched random polymer.

The RA model is an extension of the standard DLA model: we model the latent heat released when a gas particle aggregates in addition to modeling the gas and

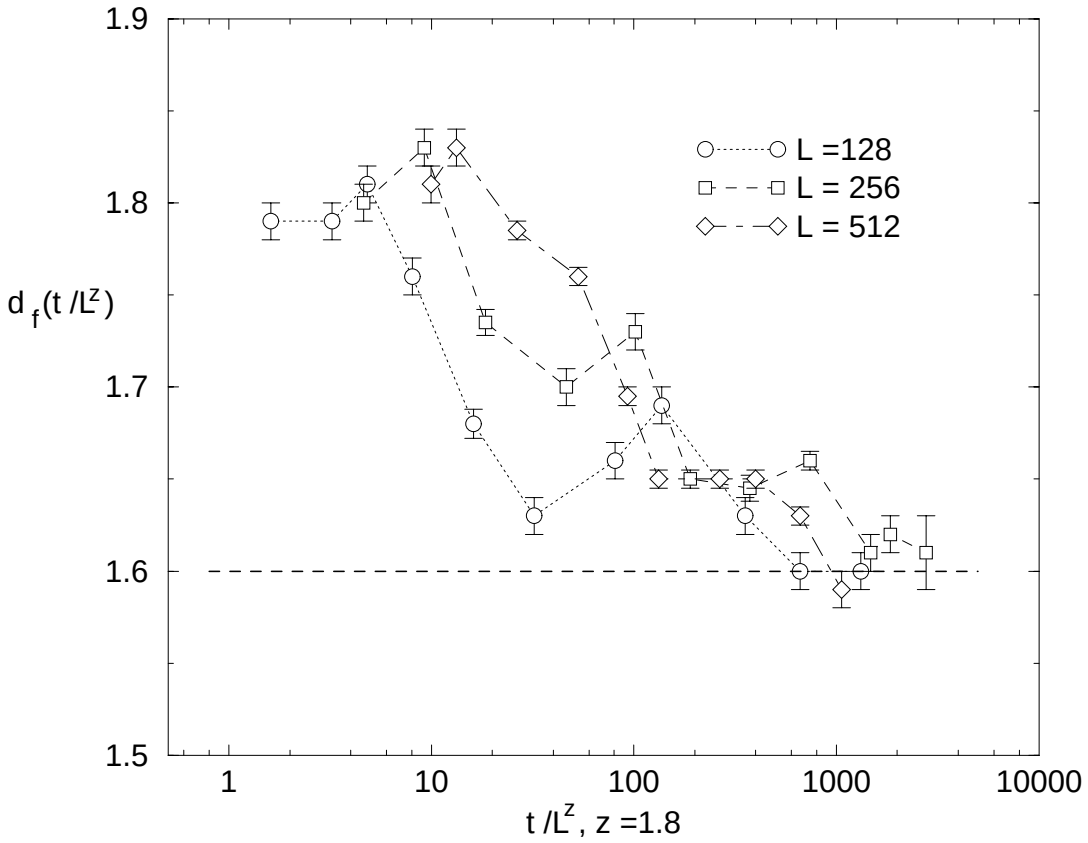

Figure 6-6: The average fractal dimension of the RA growth clusters as a function of time into the simulation, for all three system sizes.

crystal particles. Since the dynamics is local and microscopically reversible, we are realistically modeling the flow of heat and the creation of entropy in this system and thus we model the thermodynamic behavior of growing clusters.

The model presented is simple and amenable to theoretical analysis. Given the constraints that we set out, an even simpler model could be constructed with just a single heat bath particle at each site, and a single gas particle at each site. In this case, diffusion would be performed by block partitioning[88]. This simpler model would have two fewer bits of persistent state at each lattice site than the RA model, and would be slightly easier to analyze theoretically. It would, however, be less computationally efficient: for a given lattice size, the volume available to the gas and heat particles would be reduced, but the computational effort required by each step of the simulation would be unchanged.

There are some simple variants of the RA model which might merit study. For example, we have only investigated situations in which the temperature is set by the size of the final aggregate. It would be interesting to study morphology in situations where there is independent control of temperature and aggregate size. It would also be interesting from a thermodynamic perspective to modify the model by introducing a gas-crystal exclusion: gas particles would collide with the aggregate, but not diffuse over it. Thus there would be an excluded volume for the gas particles, and the crystal would do work on the gas as it grew.

6.7 Controlling growth morphology[4]

We are interested in investigating the relation between the diffusion length of the heat particles and the branching ratio of the growth aggregate. With control over the branching ratio one should be able to control many aspects of the growth morphology. Previous studies have focused on the relative effects of noise versus lattice anisotropy

[4]Note this section does *not* appear in the Phys. Rev. E **60** (1) 264–274 (1999).

to gain some control over the growth morphology[116, 55].

Until now we have been discussing a closed two dimensional system with periodic boundary conditions. If we modify the boundary conditions of the heat bath lattice from being periodic to open, we essentially place the heat bath in contact with a reservoir at zero temperature and release the heat particles into this reservoir once they reach the edges of the heat bath. Note by doing this we give up full reversibility (*i.e.*, the heat sink is irreversible). We can control the relative time it takes for the heat particles to reach the heat sink through independent control of the heat and gas diffusion lengths and thus control the effective temperature in which the aggregate grows. The relative tunability of the two diffusion fields allows us to observe a rich variety of growth structures with a continuous transition from structures resembling uniform growth, through invasion percolation, to classical DLA growth. Figure 6-7 shows the typical resulting growth morphologies. Another parameter which we can

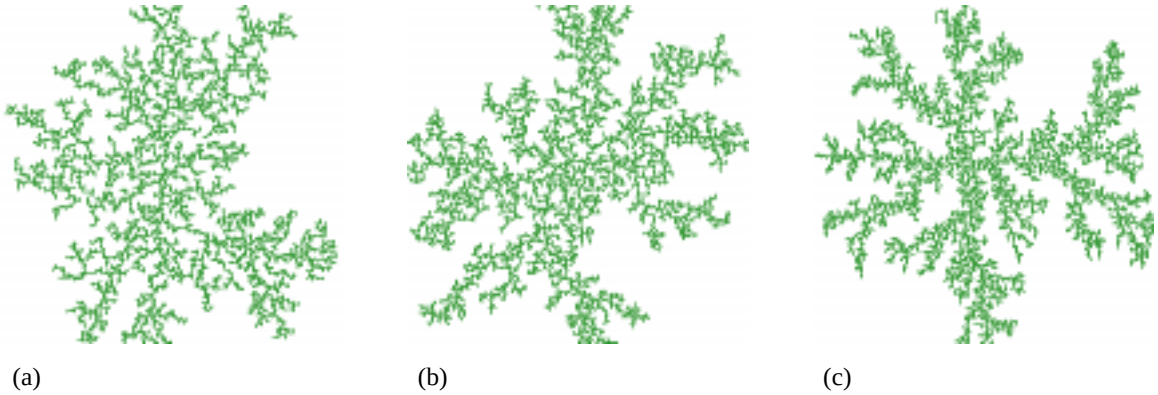

Figure 6-7: We modify the RA model by implementing a heat bath with open, rather than periodic, boundary conditions. The independent control of the heat and gas diffusion lengths allows us to tune the effective temperature in which the aggregate grows. (a) For $k = 1$ the interior of the aggregate is no longer fractal, but two dimensional. (b) For $k = 10$ we observe aggregate morphology resembling invasion percolation in diamond cells[81]. (c) For $k = 100$ we recover DLA growth.

control is the initial density of gas particles, $\mathcal{N}_g(0)/L^2$. In Fig. 6-8 we illustrate the range of growth morphologies possible as we vary both the initial density and the ratio of the heat to gas diffusion lengths. We show increasing initial gas density along

the horizontal direction and increasing heat diffusion length along the vertical axes.

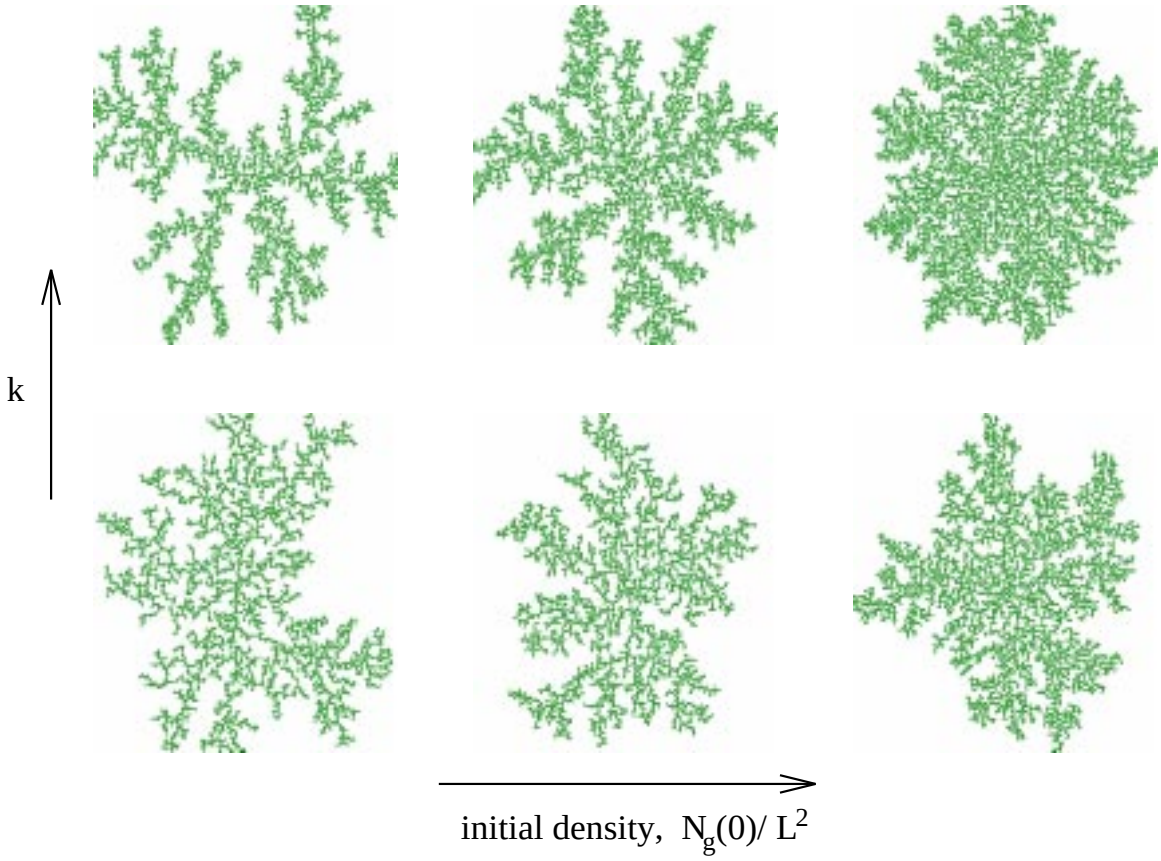

Figure 6-8: A range of RA growth morphologies as the initial gas density is increased (along the horizontal direction) and the heat diffusion length is increased (shown along the vertical direction).

Chapter 7

Simulating digital logic with the Reversible Aggregation model[1]

We are concerned with understanding the implicit computation occurring in a physical model of cluster growth, the Reversible Aggregation model. The RA model is a lattice gas model of reversible cluster growth in a closed two-dimensional system, which captures basic properties of physics such as determinism, locality, energy conservation, and exact microscopic reversibility. There are three species of particles in the RA model: gas, heat, and cluster. A diffusing gas particle may aggregate when contacting the boundary of a connected cluster. Latent heat is released during each aggregation event and is explicitly modeled by introducing a heat particle into a diffusing heat bath. Conversely a cluster particle at the boundary of the connected cluster may absorb a heat particle and evaporate, becoming a diffusing gas particle.

Allowing ourselves complete control over all the initial conditions of the model, we show that the RA model can simulate any logic circuit, and hence perform any computation. The mobile gas and heat particles are used as logic signals. The paths

[1]This work will appear in the *Proceedings of the Santa Fe Institute workshop on Constructive Cellular Automata* (Oxford University Press), in press. My coauthors are G. E. Homsy and N. H. Margolus.

these particles take are the wires. Sequences of conditional aggregation events form the basis of the logic gates. We show how to embed a universal single use gate into the dynamics of the model, then show how to construct a reusable universal gate, showing the system is capable of space-efficient computation. We show how to build arbitrary logic circuits by interconnecting gates. This requires steering and routing the signals, delaying them, and letting them cross. Finally we briefly discuss the relationship of computation in the RA model to computation in real physical systems.

7.1 Overview

We examine the computational capabilities of a physical model of cluster growth, the Reversible Aggregation (RA) model[31], which captures basic properties of physics such as determinism, locality, energy conservation, and exact microscopic reversibility. The RA model is a lattice gas model of reversible cluster growth in a closed two-dimensional system. It was introduced as a microscopically reversible physical model for studying the thermodynamics of cluster growth and pattern formation. By microscopically reversible we mean that from any state in the system we can recover the previous state exactly. There are three species of particles in the RA model: gas, heat, and cluster. A diffusing gas particle may aggregate at the boundary of a connected cluster; if it lands next to a single nearest neighbor cluster particle, it may aggregate, becoming a cluster particle and enlarging the connected growth cluster. Latent heat is released during each aggregation event and is explicitly modeled by introducing a heat particle into a diffusing heat bath. Conversely if a heat particle contacts a singly connected cluster particle it may be absorbed and that cluster particle will "evaporate" from the connected cluster, becoming a gas particle. When started with a dilute gas and a single cluster seed particle the model exhibits an initial regime of rapid nonequilibrium growth followed by a slow quasistatic regime with a well defined temperature. In the first regime a connected cluster rapidly grows, in

the second the connected cluster slowly anneals.

The present study, showing a construction for computational universality in the RA cellular automaton seems fitting for these proceedings of a workshop on "Constructive Cellular Automata". The mobile gas and heat particles are the logic signals used in the computation. The paths these particles take are the wires. We show how to steer and route the signals, how to delay them, and how to allow signals to cross each other. Aggregation events occur only at sites with one nearest neighbor which is already a cluster particle. By routing a control signal through a potential aggregation site we can conditionally create new potential aggregation sites. Such sequences of conditional aggregation events are the basis of our logic gates. We show how to embed into the dynamics of the RA model a universal single use gate, and then how to embed a reusable universal gate. We show how to interconnect gates and thus how to build arbitrary digital logic circuits, proving that the RA model is capable of space-efficient computation.

In this manuscript we take a microscopic perspective; that is we give ourselves control over all of the microscopic degrees of freedom, namely the initial positions of all the gas, cluster, and heat particles, and detailed control over the microscopic parameters controlling the pseudorandom motion of the gas and heat particles. With microscopic control and synchronous time evolution we can compute with this discrete lattice system. A more general issue is understanding computation in real physical systems where we have control only over macroscopic degrees of freedom and where we cannot depend on perfect synchronization. In the final section we will address issues of whether we have abstracted concepts from our microscopic dynamics which apply to computation in real physical systems.

7.2 The Reversible Aggregation model

The RA model is a reversible, deterministic model of cluster growth in a closed two dimensional lattice system. It extends the canonical Diffusion Limited Aggregation (DLA) model of cluster growth on a lattice[129]. The DLA model is an irreversible, deterministic model with two particle species: gas and cluster. The gas particles follow a pseudorandom walk along the lattice sites, resulting in diffusive behavior at the scale of several lattice sites. If a diffusing gas particle contacts a stationary cluster member, it aggregates, itself becoming a stationary cluster member. DLA is a serial model: Only one gas particle diffuses at a time. A few parallel versions of DLA, with multiple particles diffusing at once, have been considered[127, 95, 68, 94]. The parallel model we will extend begins with a uniform dilute configuration of gas particles. They aggregate, but no more are ever added to the system. Before a substantial fraction of the particles aggregate, this parallel version of DLA is equivalent to the serial version[127, 95]. When a large cluster has formed, the simulated structure can be compared to structures found in nature, typically by comparing the fractal dimensions. DLA is not a thermodynamic model: Particles stick irreversibly, so there is no notion of detailed or semi-detailed balance.

The RA model extends the DLA model by placing a parallel DLA model in contact with a heat bath, implemented as a field of diffusing heat particles, or "tokens", on a superimposed lattice. Both gas and heat particles diffuse throughout the system independently, each on their own lattice, under pseudorandom dynamics. They diffuse freely over the cluster and through empty space. The gas and heat particles interact through aggregation and evaporation events along the boundary of the connected cluster.

A potential aggregation site for a gas particle is a site unoccupied by cluster, with exactly one nearest neighbor occupied by cluster. Upon reaching such a site a gas particle will aggregate, becoming a cluster particle and releasing a heat particle, which represents the latent heat of crystallization. This interaction is contingent on there

being room locally in the heat bath to accept the heat particle. Explicitly modeling the latent heat released upon aggregation provides a mechanism for modeling the inverse process of evaporation. Similarly, a potential evaporation site is a site occupied only by a singly connected cluster particle. A heat particle arriving at such a site is absorbed by the cluster particle which evaporates, becoming a gas particle. This is contingent on there being no gas particle already at that site.

The model is implemented with a two phase rule. Diffusion steps, in which the particles move, alternate with interaction steps in which the gas, heat, and cluster interact, allowing aggregation and evaporation. The interaction rule is as given above. The details of the diffusion implementation are given in Sec. 7.4.1. In order to facilitate a parallel updating of the space, we divide the system into checkerboard sublattices. Since our interaction range is nearest neighbor, we can update all the even parity sites while holding the odd parity sites fixed and vice-versa.

In our cluster growth simulations, we begin with an empty heat bath, thus only aggregation can occur. We observe rapid, nonequilibrium growth of the cluster and concomitant increase in the population of the heat bath. The occupancy of the heat bath (hence also the mass of the cluster) quickly reaches steady state, meaning evaporation and aggregation events are equally likely. At this point the system has attained a single well defined temperature, despite the fact it has not yet reached thermodynamic equilibrium. During the subsequent slow approach to thermodynamic equilibrium we observe a quasistatic annealing of the cluster. The cluster morphology, as quantified by the fractal dimension, undergoes a transition from the typical ramified pattern observed for irreversible models of diffusive cluster growth (resembling frost on a window pane), to the highest entropy macrostate allowed for a connected cluster in a finite volume: a branched polymer. Figure 7-1 shows the typical cluster morphology in each of the two regimes. The small grey dots also shown are the gas particles.

The dynamics of the RA model captures a number of properties of realistic microscopic physical dynamics such as locality, conservation of energy, determinism, and

microscopic reversibility. Since aggregation and evaporation are both allowed and heat is explicitly modeled, any transition between two states may occur in either direction. This gives us a realistic thermodynamics: When started from a low entropy state (*e.g.,* with an empty heat bath), entropy increases and the system approaches a state of detailed balance, or thermodynamic equilibrium. Since we realistically model thermodynamic variables—local heat flow and the creation of entropy—we can do more than study simulated structures: We have a laboratory for studying nonequilibrium thermodynamic behavior of growing clusters of particles. For a detailed discussion of the thermodynamics, the temperature, and the evolution of the growth morphology see Ref. [31].

Until now we have been discussing a closed two dimensional system with periodic boundary conditions. If we modify the boundary conditions of the heat bath lattice from being periodic to open, we essentially place the heat bath in contact with a reservoir at zero temperature and release the heat particles into this reservoir once they reach the edges of the heat bath. We can control the relative speed with which this happens through independent control of the heat and gas diffusion lengths and thus control the effective temperature in which the aggregate grows. The relative tunability of the two diffusion fields allows us to observe a rich variety of growth structures with a continuous transition from structures resembling uniform growth, through invasion percolation, to classical DLA growth.

7.3 Computation in reversible systems

Computation was long considered to be an inherently dissipative process, requiring the "decision of a two-way alternative and elementary transmittal of one unit of information" (von Neumann as quoted in Ref. [7]). A quantitative understanding of the mechanism of dissipation came with Landauer's work on erasing a bit of information[76]. Erasure requires the transfer of information from computational to

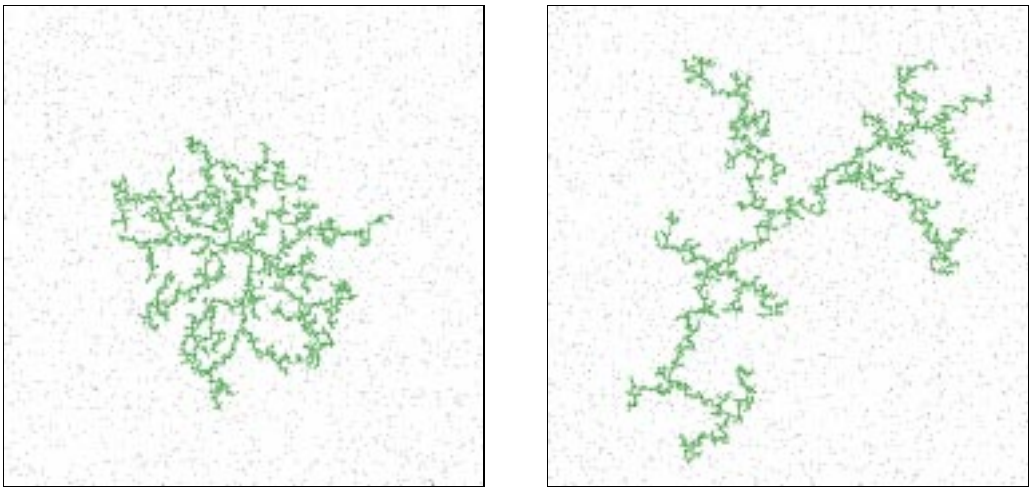

Figure 7-1: (a) An RA cluster of approximately 8270 particles, pictured at the time $t = \tau_{\mathrm{T}}$, which is when the heat bath and the gas-aggregate system first reach the same temperature. Note the grey dots represent the diffusing gas particles. (b) An RA cluster with the same number of particles pictured at time $t = 80\tau_{\mathrm{T}}$. The fractal dimension for this cluster ($d_f = 1.63 \pm 0.02$) has apparently reached the asymptotic value, and is equivalent to the fractal dimension for a quenched branched polymer.

other degrees of freedom and normally ultimately to thermal degrees of freedom. The lower bound on the heat produced by erasing one bit is $kT \ln 2$. More than a decade later Bennett showed that erasure is not necessary: Computation can in principle be performed with no dissipation (*i.e.*, no loss of information)[7]. Bennett's proof was based on an abstract Turing machine. He suggested RNA transcription (in the limit where the rate of transcription approaches zero), as a possible example of a dissipationless digital process. Interest in physical models of computation began around this same time, focusing on prototype Brownian motion computers such as Bennett's RNA model[70, 7].

After yet another decade, Fredkin introduced the concept of conservative logic and introduced a universal conservative logic gate[48]. His goal was to formulate laws of computation more like the laws of microscopic physics, with particular emphasis on microscopic reversibility and exact conservation laws. Conservative logic gates are reversible, and the total number of one's on the wires of a closed system is conserved.

The output of each gate is a permutation of its inputs. The Fredkin gate is a three input, three output conservative logic gate which implements a conditional swap. A schematic is shown in Fig. 7-2. The standard logical primitives ("and","not", and "fanout") can easily be built out of a Fredkin gate by supplying some constant inputs. The existence of a universal, reversible logic gate means it is possible to implement any digital computation out of these gates without ever erasing information. Moreover it makes reversible computation and circuit design look very similar to conventional computation and circuit design (but using different primitive logic elements).

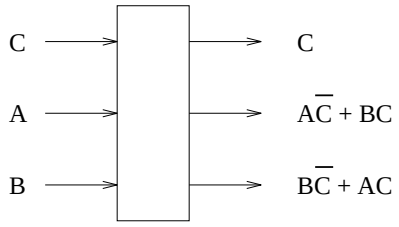

Figure 7-2: The Fredkin gate: a conservative logic gate which performs a conditional swap. If the value of the signal C is true, signals A and B are interchanged; otherwise A and B go straight through. This is a reversible operation which conserves ones: the number of ones entering and leaving the gate is the same.

Fredkin was also interested in embedding his conservative logic computations into realistic physical systems. In Fredkin's Billiard Ball Model (BBM) of computation[48], finite diameter moving balls are the signals. Collisions between balls implement logic gates. The trajectories of the balls are the wires. Strategically placed mirrors reflect the balls, implementing bends in the wires. The initial positions of the mirrors and the balls must be carefully chosen to allow for precise control and synchronization of signals. The BBM Cellular Automaton (BBMCA) is a universal, reversible CA modeled after the BBM[86]. It uses pairs of particles as the "balls", with the spacing between particles corresponding to the finite diameter of the balls.

7.4 Computation in the RA model

The scheme proposed for computation in the RA model is similar to that for the Billiard Ball Model. They both use conserved particles as signals and timing and synchronization are crucial. The conditional aggregation events in the RA model are sufficiently complicated that we expect *a priori* that the model is capable of computation. We can show this capability if we allow ourselves to explicitly specify the initial state of all the microscopic degrees of freedom of the system. We must specify the initial positions of all of the gas, cluster, and heat particles, the positions of the "mirrors" which control the diffusion (as discussed below), and the parameters controlling the motion of these mirrors (as discussed in Sec. 7.4.3.6).

We will first show how to implement wires and delays, and how to route signals. Using these elements and the RA model interaction we show the structure of a simple one-time-use logic gate, and thus prove the system is capable of computing combinatorial logic functions. However, we are interested in more general circuits: those with reusable gates, feedback, and memory elements. To this end we introduce a more complicated dyadic signaling implementation and describe a reusable universal logic gate. We then discuss issues of interconnection and signal crossing. With these in hand, we can build a simulator for any digital logic computer. We demonstrate the construction technique for a simple example circuit.

7.4.1 Signal routing and delay

In the RA model the gas and heat particles undergo pseudorandom walks, implemented using a lattice gas transport algorithm. For each particle species (gas and heat), there are two transport channels moving in opposite directions along one of the principle lattice directions. At consecutive updates of the system, we alternate between lattice directions (i.e., transport is along the $+\hat{x}$ and $-\hat{x}$ directions on odd time steps and along the $+\hat{y}$ and $-\hat{y}$ directions on even time steps). This scheme

can be extended to arbitrary dimensions by introducing additional substeps along each additional lattice direction. A particle remaining in one channel exclusively will follow a diagonal path through the lattice as shown in Fig. 7-3(a).

To simulate diffusion we cause the particle to switch between the two channels at random. We include a pseudorandom number field: a "random" binary variable at each site at each time $\eta(\vec{x}, t)$. If $\eta(\vec{x}, t) = 1$ the particle switches channels. At the end of each update we spatially permute the η values in a deterministic and invertible manner so as to have fresh random bits at each site, while maintaining constant the probability that $\eta = 1$. The permutation must be deterministic, so we can invert the dynamics when running the model in reverse. This allows us to recover the data used to make switching decisions so we can unswitch the particles and invert the "random" walks. If we use the identity permutation, the η values remain fixed and so does the particle motion (which depends on these values).

If the pseudorandom bit fields are initially filled randomly or pseudorandomly with ones and zeros, the gas and heat particles switch between channels in an unbiased manner and simulate large scale diffusion (c.f. Ref. [31] for a quantitative discussion of diffusion coefficients and a comparison of theoretical and empirical results). If instead, the $\eta(\vec{x}, t)$ bits and their motion are precisely controlled, the one bits act as deliberately placed mirrors, switching the gas and heat particles between transport channels at determined locations. An example with fixed mirrors is shown in Fig. 7-3(b). Between encounters with mirrors, the particles "stream" along a given channel uninterrupted. The gas and heat particles are controlled by separate η bits, so the gas and heat particles are reflected by separate mirrors.

Timing and synchronization are crucial to our logic scheme. To adjust timing, we can use delay loops. A delay loop can be constructed from a collection of mirrors, placed to implement a sequence of reflections. Using the transport algorithm described above, each particle takes a step in the horizontal and then the vertical directions. Since it takes at least four steps for a particle to return to its original

location, delays must be a multiple of four. Figure 7-3(c) shows a delay loop of eight time steps.

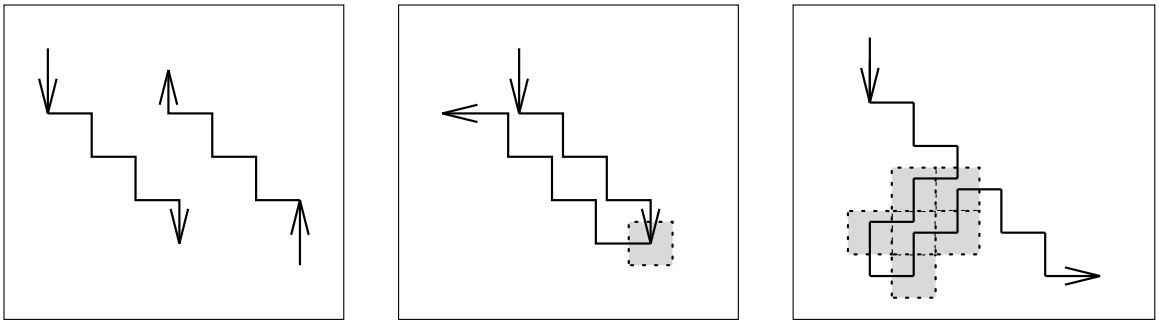

Figure 7-3: (a) Streaming in channel 1 and channel 2. (b) A mirror. (c) A delay loop of eight time steps.

7.4.2 Universality: a simple gate

Consider a simple gate, as shown in Fig. 7-4. If we count only gas and heat particles entering and exiting at A, B, C, and D (*i.e.* not counting cluster particles), this gate conserves particle number. The shaded squares represent cluster, and positions of the mirrors are implied by the particle paths. Signal A must precede B in time. This gate has the following truth table:

A	B	C	D
gas	gas	heat	heat
gas	heat	heat	gas
heat	gas	gas	heat
heat	heat	gas	heat

If we take the signal convention that a gas particle represents one logical truth value, and a heat particle the other, then clearly this gate is universal. Say, for instance, gas represents true and heat represents false. Then $C = \bar{A}$, and $D = A\bar{B}$.

So, if we can build arbitrary networks of gates such as this, then we can build arbitrary logic circuits.

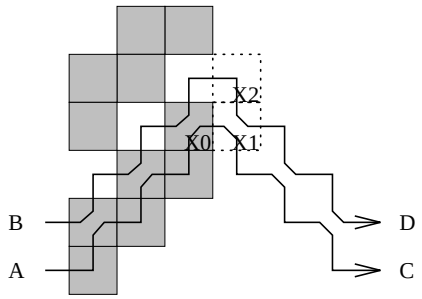

Figure 7-4: A simple, non-reusable, universal logic gate. Cluster is shown by shaded squares. Signal paths are shown as wiggly arrows. Mirror locations are implied by the paths. A and B enter as shown, with A preceding B in time. If A is gas it aggregates at X1, releasing heat. If A is heat it evaporates X0, becoming gas. If A has aggregated at X1, then if B is gas it aggregates at X2, yielding heat, and if B is heat it re-evaporates X1, yielding gas. If A has evaporated X0 and B is gas, B re-aggregates X0 yielding heat, and if B is heat, it does not interact and remains heat. Recall both gas and heat particles diffuse freely over the cluster.

There is a subtlety here, however. This gate is universal, yet it is not reusable. So in fact we may only build combinatorial logic circuits, and not those with feedback. If we wish to simulate the operation of a universal Turing machine, we must "unroll" the operation of the TM. That is, we simulate the time evolution of the machine's state by computing each state combinatorially from the previous state. This spreads the time progress of the computation out spatially, requiring more logic levels for each step we wish the machine to execute. So with a polynomial amount of space, we may simulate the machine's action for a polynomial number of steps.

We can go further, though. We wish to simulate normal reusable digital logic, so that we can build arbitrary logic circuits, with feedback, memory, etc. To this end, we propose a scheme for reusing gates which utilizes matched pairs of aggregation and evaporation events. It is a dyadic signaling scheme in which pairs of particles, appropriately delayed, are routed through the same gate so as to clean up after the computation (i.e., remove the state from the gate) leaving the gate ready for reuse.

7.4.3 A reusable gate, gate interconnection, and circuits

The gate we choose to implement is the "switch gate": a two input, three output conservative logic gate which is universal [48]. It has a control input, which we will call B, and a switched input, which we will call A. The B input passes through the gate unchanged. The A input is routed to one of two outputs, conditional on the state of B. A schematic is shown in Fig. 7-5.

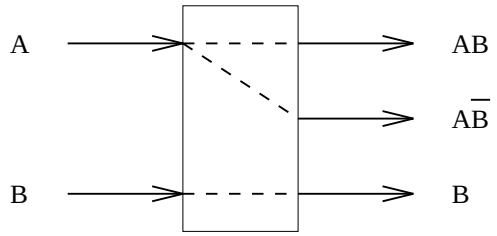

Figure 7-5: The schematic diagram of a switch gate. If the input B is true, the signal A exits the top output. If B is false, A exits the middle output.

7.4.3.1 Summary of the gate implementation

Our first consideration in implementing such a gate is that the gate must be reusable, or stateless, as mentioned above. To realize this, we adopt a dyadic signaling convention, in which the presence of a "one" is represented not by arrival of a single particle, but by a gas particle followed four time steps later by a heat particle. This allows us to use aggregation events to implement interactions without leaving permanent changes in the structure of the gate: If a gas particle enters an input and leaves a cluster bit behind, then the corresponding heat particle follows, cleaning up the aggregation event, and leaving the gate in its original state.

We implement a switch gate in the RA model as follows. We first set up an initial condition with a few cluster bits forming a simple aggregate, with only one potential aggregation site, X1. A second site, X2, will become a potential aggregation site if and only if X1 becomes occupied by cluster. There are input paths for signals A and B. The path for signal B is routed through X1, and that for A through X2.

The action of the gate is as follows: We route the gas particle of B through X1, and delay the heat particle of B for now. If B is true (that is, the gas and heat portions of B are indeed present), the gas aggregate at X1, and X2 now becomes a potential aggregation site. If B is false, nothing happens and X2 is not a potential aggregation site. Now we arrange, by appropriate choice of delays, for both particles of A to pass through X2 while X1 is (in the $B = 1$ case) occupied by cluster. If B is false, A passes through X2 unchanged (*i.e.* with its heat portion following its gas portion). If B is true, the gas portion of A aggregates at X2 yielding heat, and the heat portion of A evaporates the cluster particle at X2, yielding gas again. Thus, B being true essentially reverses the order of gas and heat particles exiting X2. When both the gas and heat portions of A have had a chance to interact at X2, we send the heat portion of the B signal through X1. If B is a one, the heat signal encounters the cluster particle at X1 and evaporates it, thus restoring the connected cluster to its original state and leaving the gate ready for reuse. Note if B is false nothing happens, which also leaves the gate ready for reuse.

The particles for A exit X2 in the same direction regardless of the state of B, but in different temporal orders. To obtain the two outputs $A\bar{B}$ and AB on different spatial paths, we use η fields which change with time to switch the gas and heat particles onto two different paths according to their timing. We then delay the heat on the AB output path, to conform with our dyadic signaling convention that heat follows gas. A schematic diagram of this interdependence of events is shown in Fig. 7-6.

7.4.3.2 Interconnect and parity

There are some subtleties involved in being able to route any output of any gate to any input of any other gate. In particular, there is a parity defined on particles: Since they are constrained to move alternately vertically and horizontally, we may draw a checkerboard on our lattice and separate the particles into those starting at time zero on a red square ("red" particles), and those starting on a black square

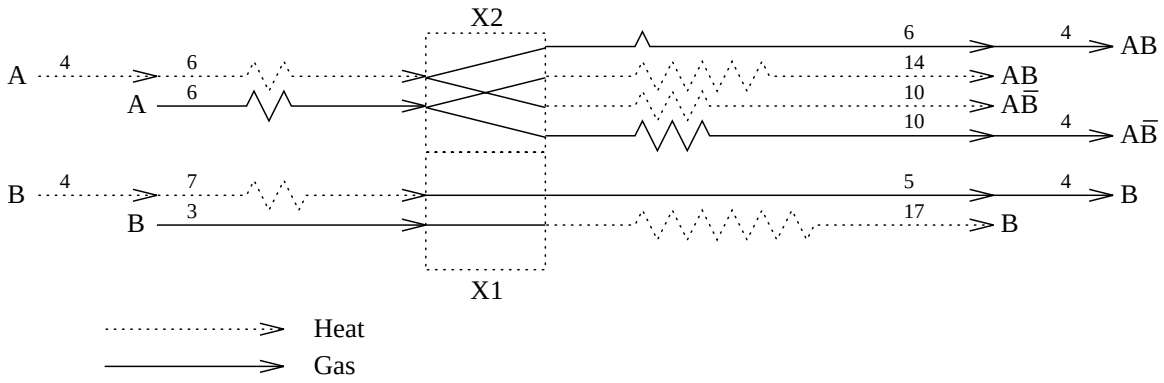

Figure 7-6: A schematic representation of the switch gate. Arrows indicate signal paths for heat and gas particles, with delays annotated. Signals A and B enter simultaneously from the left at time zero. The heat particle for each signal trails the gas particle by four time steps. If B is true, X1 is occupied from time three through time 11. In this case, A aggregates and evaporates from site X2 at times six and ten respectively. The signal paths are indicated by the top two lines in X2. If B is false, X1 is never occupied. In this case, both A and its cleanup signal pass through without interacting, and exit at the appropriate times on a different path. The signal paths in this case are indicated by the bottom two lines in X2. Output delays are chosen to place the signals at the gate outputs at a time independent of input values and signal path. Note that all paths have an identical length of 20.

("black" particles)[2]. Since the first move is horizontal for all particles, any red particle will move horizontally off its red square, and will always move horizontally off any red square. Similarly any black particle will always move horizontally off a black square. These types of particles cannot be interconverted, and a path followed by one type cannot be followed by the other. Each gate input follows a specific path, thus requiring a particular color particle. Likewise, each gate output produces a particular color particle. So if we had a gate output producing a black particle and another gate's input expecting a red particle, there would be no way of directly connecting that input to that output. To solve this, we constrain all our inputs and outputs to be on black squares of the checkerboard.

Even with all inputs and outputs on black squares, we have a synchronization problem: Our gates require synchronous signal arrivals, so sometimes we need to delay individual signals to satisfy this requirement. Note delay loops are only available in multiples of four time steps (see Fig. 7-3). Consider any two input signals to our circuit. If they are ever to interact in any gate, they must not only start on the same color square, but they must also be routable to each other in a multiple of four time steps; otherwise they will never be able reach the inputs of the gate simultaneously. So we introduce an additional constraint that all inputs to our circuit be routable to each other in multiples of four time steps. This constraint can be met simply by placing all inputs at time zero on a sublattice of a checkerboard: a square sublattice with separation distance two (see Fig. 7-7). All inputs are now separated by paths whose lengths are multiples of four, and in fact particles can travel from one site of the sublattice to an adjacent one in exactly four time steps. This also guarantees that all outputs of our circuit will be routable to each other in multiples of four time steps (since gate departure times as well as arrival times are synchronous). Note that by restricting our inputs to be on this sublattice, we have somewhat reduced the usable volume of space. This simplifies our discussion, however, and introduces only

[2]Note, this discussion is independent of the checkerboard updating scheme discussed in Sec. 7.2.

a constant factor slowdown and size increase.

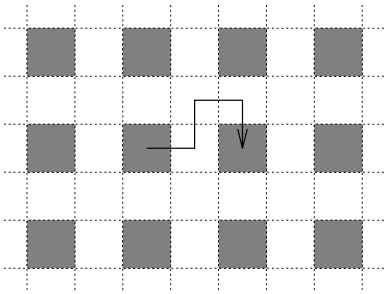

Figure 7-7: The RA lattice, with the circuit input/output sublattice shown in shaded squares. A particle is shown traveling from one sublattice site to an adjacent one in four time steps.

7.4.3.3 Details of the switch gate implementation

The details of the RA model switch gate are shown in Fig. 7-8, which depicts the initial state of the lattice and the paths taken by the signals. As mentioned above we set up an initial condition with a few cluster bits forming a simple aggregate. X1 and X2 are indicated by heavy dashed lines. One can see that X1 is initially a potential aggregation site, whereas X2 is a potential aggregation site if and only if X1 is occupied by cluster. The "extra" cluster bits prevent spurious aggregation at undesired locations.

The paths taken by the particles are shown as the wavy lines, and are determined by the placement of mirrors (indicated by the various shades of gray as shown in the legend). At the left are the two input signal paths; at the right are the three outputs. Note that all inputs and outputs lie on the sublattice described in Sec. 7.4.3.2, thus guaranteeing routability.

One can verify by "walking through" the paths by hand, that the actual delay times are as specified in Fig. 7-6. If B is true, X1 is occupied between times three and 11. In this case, if A is also true, a heat particle exits X2 at time six, and a gas at time ten. If B is false and A is true, a *gas* particle exits X2 at time six and a *heat* at time ten. The gas and the heat so produced travel to a toggling mirror, arriving

at times ten and 14 respectively. The toggling mirror begins by deflecting gas in the down direction and heat in the up direction at time zero, and toggling gas and heat directions every four steps thereafter. Hence, if B is true and a heat arrives at time ten and a gas at time 14, they are both deflected in the up direction, since the mirror toggles in between them at time 12. Conversely, if B is false, the gas arriving at time ten and the heat arriving at time 14 are both deflected *downward*, again since the mirror toggles at time 12. The extra eight steps of heat delay in the upper exit path are to delay the heat so that it follows the gas, in accordance with our signaling convention.

The longest path through the gate requires sixteen time steps, so we have imposed extra delay on some signal paths so that the propagation delay through the gate is exactly sixteen on all paths. This will be convenient when building more complex circuits, since we will not have to worry about delaying signals to compensate for differences in gate delay. The paths in the figure are longer than the gate propagation time so that both the gas and heat particles can be shown entering and exiting the gate.

7.4.3.4 The reverse switch gate

In order to build general logic circuits we will need to use switch gates in both the forward and reverse directions[48]. The switch gate used in reverse is a three input two output gate which performs the inverse logic operation. The three inputs must of course be appropriately correlated for this to be possible. Since the dynamics of the RA model is reversible, we can build the reverse gate by sending the signals through the outputs of the original gate, but with the opposite signaling convention: heat preceding gas. Converting between forward and reverse signaling conventions is simple: We need only impose an extra delay of eight steps on the gas so that it follows the heat by four time steps. We may convert back to the forward signaling convention by similarly delaying the heat.

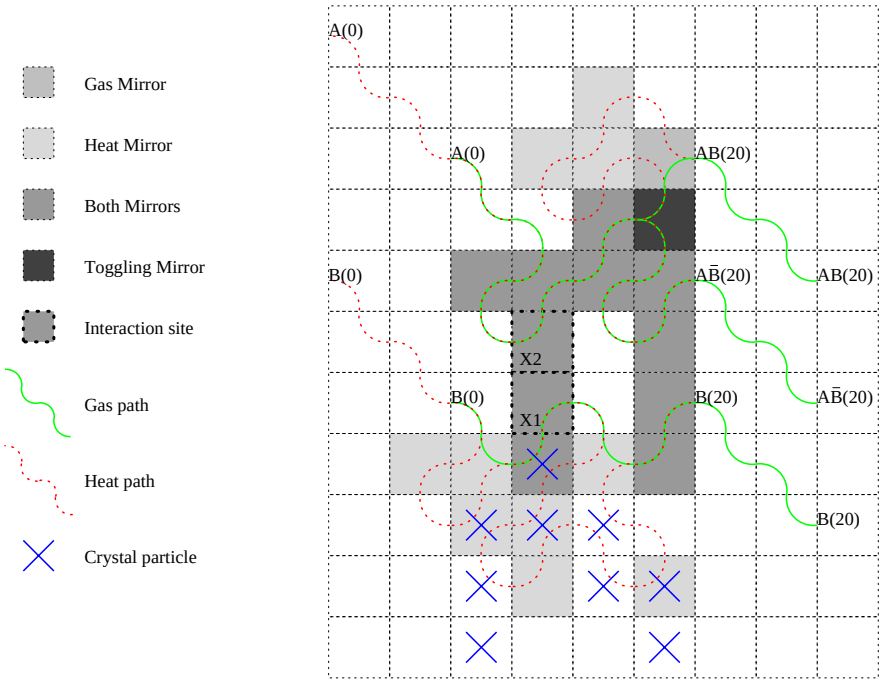

Figure 7-8: A detailed picture of an RA switch gate. Signal entry and exit times are given in parentheses after the signal name.

7.4.3.5 Crossing wires

We showed above that, observing parity constraints in gate placement, any output can be connected to any input. However we have yet to show that signals can cross—a necessary detail for connecting any output to any input. Consider two signals, one traveling in a diagonal downward and to the right, the other traveling in a diagonal upward and to the right. The first signal will travel purely in one channel (channel 1 in Fig. 7-3(a)), encountering no mirrors. The second signal must flip channels at each time step and thus encounter a mirror at every site. Thus for these signals to cross there must both be mirrors and be no mirrors at the two lattice sites they will both encounter. A way to implement this is to delay one signal by four time steps, allowing the first signal to propagate past the relevant two sites, then toggle the mirrors. The second signal now encounters the correct mirror configuration when it occupies the relevant two sites. A "cross gate" so constructed is shown in Fig. 7-9. The shaded

sites are the ones where the mirrors are toggled. Note that we have delayed the first signal by four time steps after its encounter with the toggling sites, so that both signals leave the cross gate synchronously.

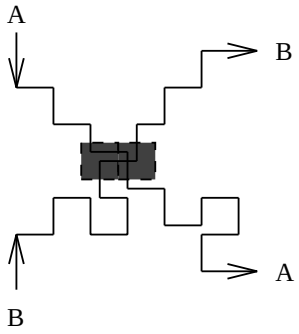

Figure 7-9: A "cross gate". The signal traveling upward is delayed by four time steps while the signal traveling downward passes through the two shaded lattice sites. The mirrors at the shaded sites are then toggled. The upward signal now encounters the correct mirror configuration to pass through as indicated. The first signal is delayed by four steps after passing through the shaded sites so that both signals leave this gate synchronously.

7.4.3.6 A sample circuit

The RA model has been implemented on a special purpose cellular automata machine, the CAM8[87]. A detailed description of this implementation can be found in Ref. [31]. For constructing the logic gates we have discussed in this paper, we use the CAM8 implementation of the RA model, but with the modification that the initial state and the dynamics of the pseudorandom bits are precisely specified. In the original model the permutation of the η bit fields was simply a displacement (shift) in $\hat{x}$ and $\hat{y}$ by a prespecified amount at each update step. Here, where we wish to exercise detailed control over the microscopic dynamics, we choose the displacements more restrictively. We displace the η bit fields for both gas and heat by half the size of the space in the horizontal direction only. Further, this displacement occurs only on every fourth time step. This allows us to implement the toggling mirrors easily. We place all the circuitry and one set of mirrors in the left half of the space and place

a second set of mirrors, with the toggling mirrors complemented, in the right half. Note that since the shifts of the pseudorandom bit planes are specified as part of the initial condition, we are still just manipulating our initial condition in order to effect computation in the RA model.

To see the detailed action in the CAM8 simulation, consider a single switch gate. The operation of this gate in three of the four input cases is shown in Fig. 7-10. Note that A takes the topmost output path if and only if $B = 1$.

Figure 7-10: Three cases in the operation of a switch gate. Left: $A = 0$, $B = 1$. Center: $A = 1$, $B = 0$. Right: $A = 1$, $B = 1$. Particles are heavily shaded, cluster is black, and paths are lightly shaded. The path shading is only for every alternate time step, to make it simpler to resolve distinct paths by eye. The null case, $A = 0$, $B = 0$ is omitted.

We have shown how to implement a reusable, universal logic gate, how to route and delay signals, and how to let signals cross so that output from any can serve as input to any other gate. Thus we can build any Boolean logic circuit we wish. As a simple example of connecting gates together we construct an identity circuit from two back to back switch gates, the second running in reverse, followed by a cross gate. The schematic of the circuit is shown in Fig. 7-11. The CAM8 implementation is pictured in Fig. 7-12. The black squares are the stationary cluster bits. The lightly shaded squares mark the trace of where the signals have traveled (the wires). The signals (heavily shaded squares) exit the gate at the right. Note that in the space between the two switch gates, the gas particles are delayed by eight so the signals enter the reverse gate with the opposite signal convention. In order not to clutter the

figure we did not reinvert the signals, so they leave with the inverse signal convention of heat first.

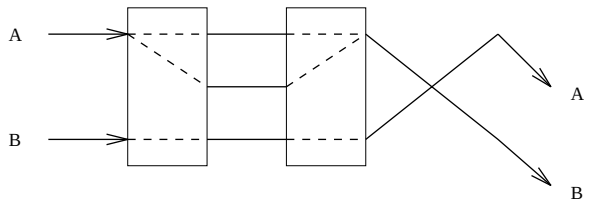

Figure 7-11: A circuit composed of two switch gates—one forward, one reverse—followed by a signal crossover.

Figure 7-12: A CAM8 implementation of an identity gate (composed of two switch gates back to back) and a crossover. The shaded squares are the wires, visualized every other time step. The signals have exited the circuit at the right.

7.5 Discussion and conclusions

7.5.1 Summary

We have investigated the computational capabilities of a specific physical model of cluster growth, where we have control over all of the microscopic degrees of freedom. The dynamics can compute any sequence of digital logic operations if the system is initialized to a precisely specified state. Thus the RA model can simulate any other digital dynamics.

The moving particles are the logic signals, the paths they take are the wires. Synchronization of signal arrival times at specific sites requires careful layout of mir-

rors which route and delay signals. By routing particles through specified interaction sites, we can build logic gates from conditional aggregation and evaporation events. These constructs are relatively straightforward to implement, allowing us to build a simple logic gate as shown in Sec. 7.4.2. This gate is changed by the interaction—it can be used only once. Showing that the RA model can support reusable gates adds complication. We have many degrees of freedom; choosing an appropriate signaling convention is crucial. We choose a dyadic signaling convention which allows us to construct reusable logic gates, as discussed in Sec. 7.4.3.3.

7.5.2 Computation in real physical systems

We have exhibited computations in a discrete lattice system that depend on exactly synchronous updating and complete control of all microscopic degrees of freedom. Have we abstracted anything useful for computation in real physical systems?

The RA model captures some essential aspects of real cluster growth. Perhaps the most relevant to computation is the conditional nature of aggregation (*i.e.*, the presence of nucleation sites only at the perimeter of a growing cluster along with the absence of heat in the local environment). This conditional interaction allows us to build logical primitives that do not depend on exact synchronization.

Computation in the RA model can be accomplished with exact control of the microscopic degrees of freedom. However, in real physical systems, including electronic computers, we typically have control only of the macroscopic degrees of freedom. For the macroscopic dynamics of a system to be universal, the microscopic dynamics must necessarily be universal: If we cannot compute with complete control over the system, we cannot hope to compute with less control. Thus we can consider the construction given in this paper as a "warmup" for addressing the "larger" issue of macroscopic universality. The analogy to thermodynamics is straightforward. A gas can do mechanical work moving a piston despite the fact we know nothing of the individual gas particles; We know only aggregate quantities. The question then is,

can we get *computational* work out of a system with control only of the macroscopic degrees of freedom? With this type of understanding we might be able to compute with a growing bacterial colony or a growing crystal aggregate, and exert detailed control over the structures produced[1].

A first step in understanding macroscopic computation in the RA model is to test the robustness of the system's ability to compute when it is subjected to an actual stochastic dynamics. For instance we can study the situation where the particles follow truly random walks or where the interactions are probabilistic, yet we maintain control over the other degrees of freedom. If we replace each single gas and heat particle with a dilute cloud of particles we may be able to conditionally aggregate and evaporate with high probability. Directing the macroscopic motion of clouds of particles may require us to add momentum conservation to the RA model, making it even more realistic. Even so, restoring the gates and signals to their starting state would require some dissipative cleanup process. This would entail a constant throughput of "power" (*i.e.*, the addition and removal of particles).

Chapter 8

Orbit structures of discrete invertible dynamics

We are interested in the general theory of discrete dynamical systems with exact symmetries and conservation laws (as will be discussed briefly in "Directions for future work", Sec. 9.2.4). In the previous sections we studied specific examples of such physically motivated microscopic models which emulate macroscopic physical behavior. We used these models to study aspects of dynamic scaling and growth phenomena, and thermodynamics. A property of a dynamics important for the study of thermodynamics is ergodicity (*i.e.*, the property that all states are actually visited during the evolution of the system).

As a first step toward a general theoretical treatment of ergodicity, we study the cycle structures of invertible dynamics on systems with a finite number of discrete states. Invertible dynamics are a subclass of all dynamics. They permit only one-to-one mappings of the set of states onto itself. We examine the random network limit where any state can map onto any other state, and thus ignore the constraint of locality. This general case is amenable to theoretical treatment and it includes CA's as a special case.

We present some general results for the structure of the ensemble of allowed dy-

namics, the expected cycle time for a point chosen uniformly at random from that ensemble, and the expected number of distinct cycles for a dynamics. We also present results for the expected ergodicity of the dynamics. The results are summarized in detail in Sec. 8.2.

8.1 General cycle structures

Consider a discrete state space consisting of M bits. This system has $N = 2^M$ distinct states. As a specific example, consider a lattice of M sites, where each site M_i takes on one of two values (for instance $M_i = \{0, 1\}$, or equivalently $M_i = \{-1, +1\}$). Now consider an arbitrary deterministic dynamics on this state space. We can describe the dynamics fully by specifying the successor of every state. In an unconstrained dynamics, a state may have more than one predecessor. But for a invertible dynamics there must exist a unique inverse, so each state must have one and only one predecessor (*i.e.*, the dynamics is a one-to-one and onto mapping). An invertible dynamics necessarily divides the state-space into distinct cycles, or orbits. To see this, start with an arbitrary state $\mathcal{S}_i$ and enumerate the temporal sequence of successive states. The state $\mathcal{S}_i$ must recur, since the state space is finite and the dynamics is one-to-one and onto. As soon as it does, we reach a cycle, forming a closed orbit for that portion of the state space. One extreme possibility is for all states to be on one orbit, corresponding to a fully ergodic system. At the other extreme, each state is on an orbit which includes only itself, corresponding to the identity mapping.

For an unconstrained deterministic dynamics the cycle structure is quite different and involves many familiar concepts from the field of nonlinear dynamics: fixed points, limit cycles, basins of attraction, strange attractors, etc. The generic diagram of trajectories through state space for one basin of attraction looks like a network of branches leading into a single closed loop, see for example Fig. 8-1. Unless a state is located on the closed loop (*i.e.*, the limit cycle), once it is visited it will never be

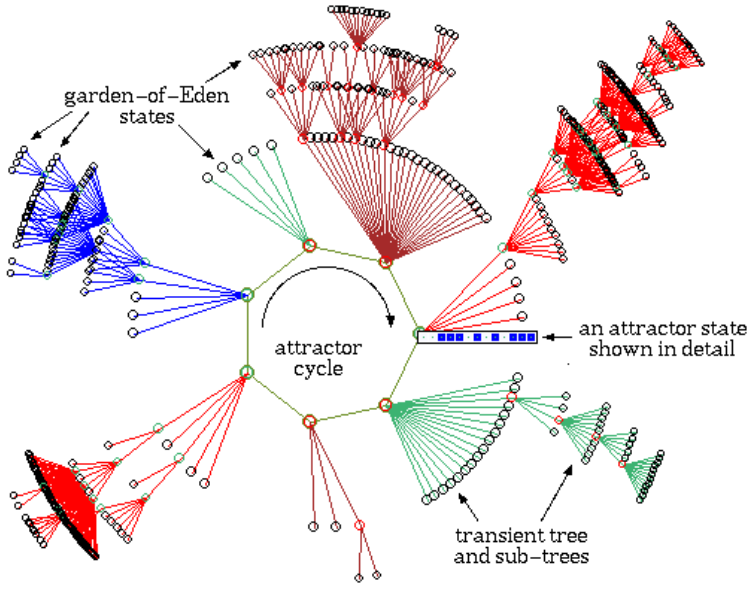

Figure 8-1: An example of trajectories through state space for one basin of attraction. This looks like a network of branches leading into a single limit cycle. Note this figure is reprinted from Ref. [133].

revisited. Hence, until a limit cycle is reached, there is a contraction in the number of accessible states with each step of the evolution.

The cycle structure for an invertible deterministic dynamics consists of a collection of simple closed loops (*i.e.*, a collection of simple orbits). The orbits are distinct. There is no merging of trajectories. Once started on a specific orbit, the system remains on that orbit. Using the language of nonlinear dynamics, we note that the orbits are not only the basins of attraction, they are the limit cycles, and that the only fixed points are orbits of length one. All the states on an orbit remain accessible throughout the evolution of the system so there is no contraction in the number of accessible states.

8.2 Summary of results

We present some exact results for the structure of the ensemble of allowed dynamics, showing first how to count the number of invertible dynamics for a discrete system of

finite size. If a particular state from this ensemble of dynamics is chosen at random, the expectation value of the length of the orbit on which that state is located can be exactly calculated. Also exactly countable is the number of orbits in the ensemble of dynamics and the number of orbits of each specified length. We will then consider the ensemble average over the set of all possible dynamics to derive the expected number of orbits for a dynamics chosen uniformly at random. This analysis will suggest that each orbit of a typical dynamics can be characterized by an extensive quantity which is conserved during the time evolution of the system. Finally we discuss issues of ergodicity and specific ergodicity. We compare all these results to the equivalent results for unconstrained random mappings. We show that invertible mappings are approximately ergodic, in a suitably defined statistical sense, whereas unconstrained mappings are not. The average trajectory for an invertible dynamics tends to visit half of all possible states before cycling. The average trajectory for an unconstrained dynamics tends to visit the square root of the number of possible states.

Most of the results presented in this chapter have been known for some time to distinct communities of scientists. Mathematicians have been interested in both unconstrained and invertible discrete mappings, primarily due to their interesting probability distributions and the variety of quantities which are not self averaging[1][43, 60, 61]. The random number generator community has been interested in invertible discrete mappings due to their long cycle times[73]. The physics community has been interested in unconstrained mappings. The random network limit of an unconstrained map is equivalent to a fully connected Boolean network[67], and the probability distributions of random networks, particularly the distribution of basins of attraction, relate directly to properties of spin glasses[24, 6]. To date physicists have not focused on invertible random maps, other than some interest in discrete versus continuous Hamiltonian maps[32]. Invertible random maps should be of interest as they are mi-

[1] One quantity which is not self averaging is the average cycle time which increases indefinitely with the size of the state space.

croscopically reversible, they have longer evolution times, and they approach fully ergodic in the asymptotic limit.

Here we bring together some of the known results in a novel context, and we present a new result for the expected specific ergodicity of a randomly chosen unconstrained dynamics. Our interpretation of results is novel. We note that even though CA's have extensive conserved quantities, this fact alone is not enough to distinguish them from typical invertible dynamics. We also show that for invertible dynamics the distribution of values for the specific ergodicity is peaked at the limiting value. Thus an overwhelming fraction of invertible dynamics have specific ergodicities close to this limiting value and approach fully ergodic. We conjecture that most reversible CA's are close to ergodic since they are members of the set of invertible dynamics, and our empirical observations are in agreement with this. In contrast the distribution of the specific ergodicity for unconstrained dynamics has an intermediate average value, so we get no information about the width and shape of the distribution.

8.3 Ensemble of dynamics and average cycle times

Let us assign arbitrary indices to the states $\mathcal{S}$ ranging from $\mathcal{S}_1$ to $\mathcal{S}_N$. We can specify a mapping by specifying the successor of each state. The collection of distinct mappings forms the ensemble space of dynamics. Figure 8-2 depicts the ensemble space for the set of invertible dynamics (as will be discussed below in Sec. 8.3.1).

Consider a state $\mathcal{S}_{j,k}$ chosen uniformly at random from the ensemble of dynamics (*i.e.*, state $\mathcal{S}_j$ subject to dynamics d_k). We define $\tau_{j,k}$ to be the number of distinct successor states of $\mathcal{S}_{j,k}$. If the dynamics is invertible, $\tau_{j,k}$ represents the number of points on the orbit containing state $\mathcal{S}_{j,k}$. If the dynamics is unconstrained, $\tau_{j,k}$ represents the number of distinct states on the trajectory starting at state $\mathcal{S}_{j,k}$. We are interested in the quantity τ_j which is the average over all dynamics of $\tau_{j,k}$. It can be regarded as a random variable over states, representing the number of successors

of state $\mathcal{S}_j$ averaged over all dynamics.

8.3.1 Invertible dynamics

Note that for an invertible mapping each successor must be unique. If we start from an arbitrary state selected at random, there are N possible successors. $N-1$ possible successors remain for the next state to be mapped, and so on until only one possible successor remains for the Nth state selected to be mapped. One-to-one and onto mappings on a finite state space are permutations. Thus there exist a total of $D = N!$ invertible discrete dynamics on a state space of size N.

Figure 8-2 depicts the ensemble space of all possible dynamics. A few of the dynamics are explicitly enumerated, including how the permutation breaks the state-space up into distinct orbits (note that we are using the fixed set of arbitrary state labels). The first dynamics, where each element is mapped onto itself, is the identity. Each element is on an orbit of length one. In contrast for the second dynamics illustrated (which we call "shift-right") all of the states are contained on one orbit. This dynamics is fully ergodic. The third dynamics ("shift-right by two") partitions the space into two distinct orbits, each of length $N/2$, where one orbit contains all even labeled states, the other contains all odd labeled states. Note that the orbit structures as drawn are meant to be abstractions indicating the length the orbits and topology of the state space. Any dynamics with an identical topology is represented by the same orbit structure (*i.e.*, the orbit structures define an equivalence class for the dynamics).

We can calculate average quantities for the ensemble of dynamics. Consider again a state $\mathcal{S}_{j,k}$ chosen uniformly at random from the ensemble of dynamics (*i.e.*, the state $\mathcal{S}_j$ subject to dynamics d_k). The probability that $\mathcal{S}_{j,k}$ is mapped onto itself under dynamics d_k is $1/N$. The probability it is mapped onto a state other than itself is $(N-1)/N$. The probability that $\mathcal{S}_{j,k}$ is mapped onto itself at the second iteration is the probability it is *not* mapped onto itself on the first iteration times the

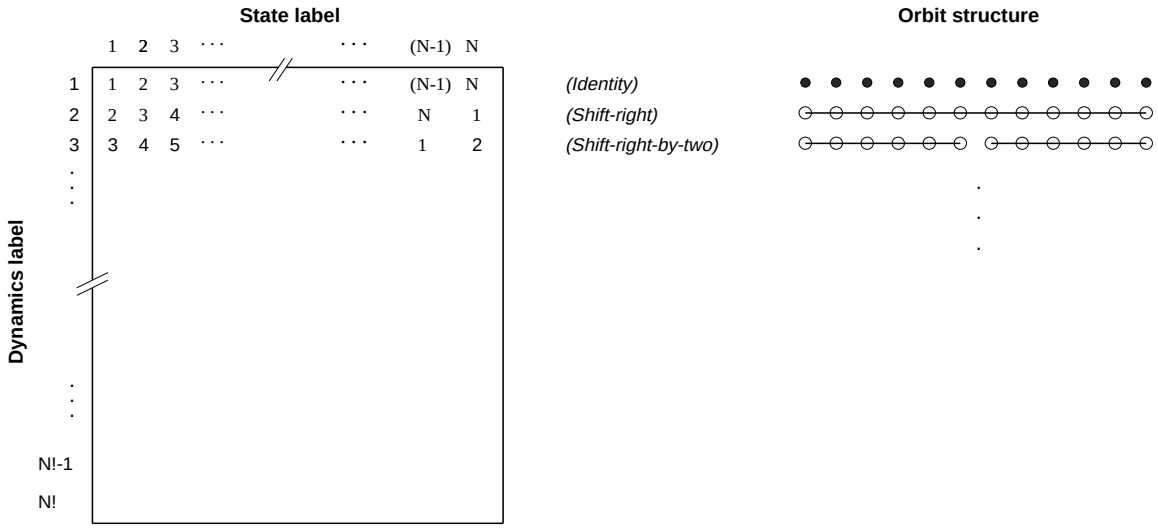

Figure 8-2: The ensemble space of invertible dynamics and allowed orbits. The state labels are listed across the top of the table shown on the left hand side. Note the state labels are arbitrary but consistent—they are assigned once and persist. The state onto which state $\mathcal{S}_j$ is mapped by dynamics d_k is shown in line k and column j of the table. Each dynamics partitions the space into orbits. The orbit lengths and topology of the state space are shown on the right hand side of the figure. Orbits of one point are drawn as solid circles. Longer orbits are drawn as a line through open circles. Even though not explicitly drawn, the two endpoints of such a line connected. Note the state labels along the orbits are not explicitly specified.

conditional probability it *is* mapped onto itself on the second iteration: $[(N-1)/N] \cdot [1/(N-1)] = 1/N$. Note that the first state we encounter on an orbit remains a valid successor state, subsequent states do not—the dynamics must first revisit the original point on the orbit before revisiting any subsequent points. The probability that $\mathcal{S}_{j,k}$ is not mapped onto itself at the second iteration is $[(N-1)/N] \cdot [(N-2)/(N-1)] = (N-2)/N$. Table 8.1 shows the probabilities of both outcomes as a function of number of iterations. The probability that $\mathcal{S}_{j,k}$ will map onto itself at iteration t, and thus close the orbit (meaning $\tau_j = t$), is uniformly distributed with respect to t:

$$p_{\tau_j}(\cdot) = 1/N. \tag{8.1}$$

Thus the randomly chosen state $\mathcal{S}_{j,k}$ is just as likely to be on an orbit of any length between $\tau_j = 1$ and $\tau_j = N$. The expected length of the orbit containing an arbitrary initial state (chosen uniformly at random from the ensemble of dynamics), averaged over all states in the ensemble is

$$\langle \tau_{j,k} \rangle = \frac{1}{ND} \sum_{k=1}^{D} \sum_{j=1}^{N} \tau_{j,k}. \tag{8.2}$$

Since the number of states with $\tau_{j,k}$ equal to some particular value t is given by $(ND)\, p_{\tau_j}(t)$, we can regroup the sum as follows:

$$\begin{aligned} \langle \tau_{j,k} \rangle &= \frac{1}{ND} \sum_{t=1}^{N} (ND)\, p_{\tau_j}(t) \cdot t \\ &= \sum_{t=1}^{N} p_{\tau_j}(t) \cdot t. \end{aligned} \tag{8.3}$$

Using the expression in Eq. 8.1,

$$\langle \tau_{j,k} \rangle = \frac{1}{N} \sum_{t=1}^{N} t = \frac{1}{N} \cdot \frac{N(N+1)}{2} = \frac{N+1}{2}. \tag{8.4}$$

Table 8.1: Probability that a cycle will be reached, or will not be reached, as a function of the number of iterations, for both invertible and unconstrained dynamics.

Iteration number, t	Invertible		Unconstrained	
	Probability orbit closes, $p_{\tau_j}(t)$	Probability orbit does not close, $q_{\tau_j}(t)$	Probability trajectory cycles, $p^u_{\tau_j}(t)$	Probability trajectory does not cycle, $q^u_{\tau_j}(t)$
1	$\frac{1}{N}$	$\frac{N-1}{N}$	$\frac{1}{N}$	$\frac{N-1}{N}$
2	$\frac{N-1}{N}\cdot\frac{1}{N-1}=\frac{1}{N}$	$\frac{N-1}{N}\cdot\frac{N-2}{N-1}=\frac{N-2}{N}$	$\frac{N-1}{N}\cdot\frac{2}{N}$	$\frac{N-1}{N}\cdot\frac{N-2}{N}$
3	$\frac{N-2}{N}\cdot\frac{1}{N-2}=\frac{1}{N}$	$\frac{N-2}{N}\cdot\frac{N-3}{N-2}=\frac{N-3}{N}$	$\frac{N-1}{N}\cdot\frac{N-2}{N}\cdot\frac{3}{N}$	$\frac{N-1}{N}\cdot\frac{N-2}{N}\cdot\frac{N-3}{N}$
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.
n	$\frac{1}{N}$	$\frac{N-n}{N}$	$\frac{(N-1)!}{(N-n)!}\cdot\frac{n}{N^n}$	$\frac{(N-1)!}{(N-n-1)!}\cdot\frac{1}{N^n}$
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.
N	$\frac{1}{N}$	$\frac{N-N}{N}=0$	$(N-1)!N^{-(N-1)}$	N^{-N}

8.3.2 Unconstrained dynamics

Invertible dynamics are a subclass of all possible dynamics. Arguments analogous to those made in Sec. 8.3.1 can be made for the ensemble of unconstrained dynamics. Since states do not need to be mapped uniquely, there are always N choices for the successor of any state, thus there are $D^u = N^N$ possible dynamics. If a state recurs, a limit cycle is reached. The probability over all dynamics that there are t elements in the set of distinct successor states for a point $\mathcal{S}_{j,k}$ chosen uniformly at random from the ensemble of unconstrained dynamics proceeds as did the analysis for the invertible dynamics. It is outlined in Table 8.1, and results in

$$p^u_{\tau_j}(t) = \frac{(N-1)!}{(N-t)!} \cdot \frac{t}{N^t}. \tag{8.5}$$

The asymptotics of this distribution can be established. Using Stirling's approximation ($N! \approx (2\pi N)^{\frac{1}{2}} \cdot N^N \cdot e^{-N}$) and some algebra it can be shown that in the limit where N is large,

$$p^u_{\tau_j}(t) \approx \frac{t}{N} \exp\left(-\frac{t^2}{2N}\right). \tag{8.6}$$

Thus the expectation value of the time to reach a closed cycle starting from a point chosen uniformly at random from the ensemble of unconstrained dynamics can be calculated. In the limit where N is large

$$\begin{aligned} \left\langle \tau^u_{j,k} \right\rangle &\approx \int_1^N t \cdot p^u_{\tau_j}(t) \cdot dt = \frac{1}{N} \int_1^N t^2 \cdot \exp\left(-\frac{t^2}{2N}\right) \cdot dt \\ &= \left(\frac{\pi N}{2}\right)^{\frac{1}{2}}. \end{aligned} \tag{8.7}$$

Comparing this with the expected time to reach a closed cycle for the invertible case (Eq. 8.4) we see a randomly chosen state subject to an invertible dynamics will have a much longer cycle time than a randomly chosen state subject to an unconstrained dynamics.

When many-to-one mappings are allowed the orbit will close whenever a state previously visited is revisited. The cycle structures consist primarily of "wispy tails" leading to short limit cycles (a network of branches leading into a single closed loop). The structures are visually striking to the point where an "atlas" of cycle structures for unconstrained dynamics has been chronicled [132].

8.4 Structure of orbits for invertible dynamics

Recall that there are $D = N!$ possible invertible dynamics on a discrete state space of N points. Each dynamics must account for all N points. Thus there are $N \cdot N!$ total points in the collection of dynamics, as shown in Fig. 8-2. From this we can deduce $\mathcal{O}(\tau)$, the total number of orbits of length τ, and the total number of orbits, $\mathcal{O} = \sum_{\tau'} \mathcal{O}(\tau')$. Consider the number of points which are on an orbit of length τ. The following are equivalent: (1) the number of orbits of length τ times the number of points on each orbit, $\mathcal{O}(\tau) \cdot \tau$; (2) the number of total points times the probability that a point is on an orbit of length τ: $N \cdot N! \cdot p_{\tau_j}(\tau)$. Recall that in the invertible case τ_j is a random variable representing the average over all dynamics of the length of the orbit containing state $\mathcal{S}_j$, and that the probability $p_{\tau_j}(\cdot)$ was determined in Eq. (8.1). Equating the two expressions,

$$\mathcal{O}(\tau) \cdot \tau = N \cdot N! \cdot p_{\tau_j}(\tau) = N \cdot N! \cdot \frac{1}{N}. \tag{8.8}$$

Thus the total number of orbits of length τ,

$$\mathcal{O}(\tau) = \frac{N!}{\tau}. \tag{8.9}$$

The total number of orbits, in all dynamics,

$$\mathcal{O} = \sum_{\tau'=1}^{N} \mathcal{O}(\tau') = N! \sum_{\tau'=1}^{N} \frac{1}{\tau'} \approx N! \left(\ln N + \gamma_E\right) \approx N! \ln N. \tag{8.10}$$

Note $\gamma_E = 0.5772...$ is Euler's constant, which by definition is the difference between the discrete and continuous "integration" of the natural logarithm in the limit where N approaches infinity[115]. We can now compute the probability density function for the length of an orbit chosen uniformly at random from the entire ensemble:

$$p_\tau(t) = \frac{\mathcal{O}(t)}{\mathcal{O}} \approx \frac{N!}{t} \cdot \frac{1}{N! \ln N} = \frac{1}{t \ln N}. \tag{8.11}$$

So, averaged over all orbits (which we denote as $\langle \cdot \rangle_{\mathcal{O}}$), the expected length of a randomly chosen orbit is

$$\langle \tau \rangle_{\mathcal{O}} = \sum_{t=1}^{N} t \cdot p_\tau(t) \approx \sum_{t=1}^{N} t \cdot \frac{1}{t \ln N} = \frac{N}{\ln N}. \tag{8.12}$$

Note the distinction between the expected length of a uniformly chosen orbit $\langle \tau \rangle$, Eq. (8.12), and the expected length for the orbit of a uniformly chosen point $\langle \tau_{j,k} \rangle$, Eq. (8.3).

We can calculate the expected number of orbits per dynamics, $\langle \omega \rangle$, for a dynamics chosen uniformly at random,

$$\langle \omega \rangle = \frac{\mathcal{O}}{D} \approx \frac{N! \ln N + N! \gamma_E}{N!} = \ln N + \gamma_E. \tag{8.13}$$

We can distinguish the orbits by labeling each one with a numerical index. That value of the index will be the same for all points on that orbit (it is conserved). The expected number of distinct indices is $\langle \omega \rangle \approx \ln N$. Note that a natural extensive quantity of such a system is the sum of all the bits, $\sum_i M_i$, a number which ranges between zero and $M = \log_2 N - 1$, and has $\log_2 N$ distinct values. Hence the expected number of orbits grows as the number of bits in the system: $\log_2 N \sim \ln N$.

CA's have extensive conserved quantities, thus we would expect this extra structure would distinguish CA's from typical invertible dynamics. Yet the size range of conserved quantities in CA's is of the order of the size range of the number of orbits

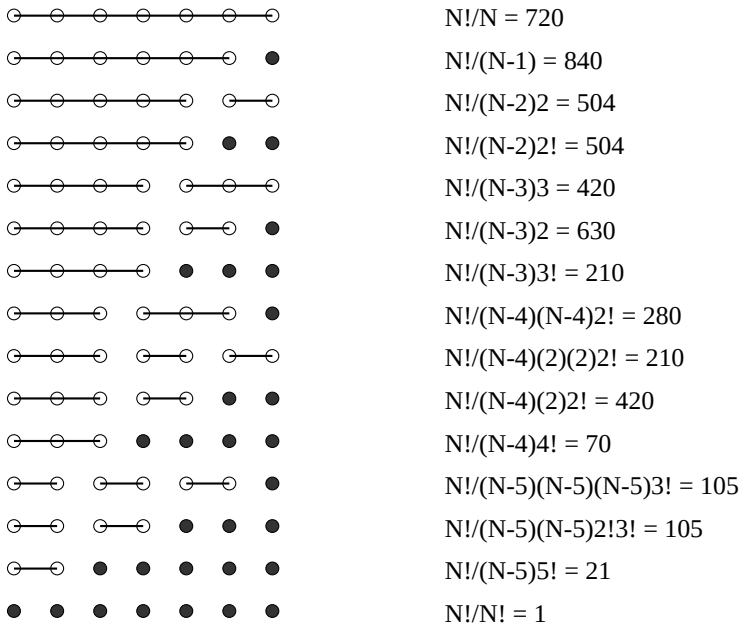

Figure 8-3: Left: The orbit structures for invertible dynamics on a state space of size $N = 7$. Note the structures define an equivalence relation on dynamics: two dynamics with the same number and length of the orbits (and hence the same topology of the state space) are in the same class. Right: The number of dynamics in the equivalence class with the specified orbit structure.

for arbitrary invertible dynamics. Thus, against our initial intuitions, the existence of extensive conserved quantities is not sufficient to distinguish CA's from typical invertible dynamics.

8.5 Ergodicity

A fully ergodic dynamics is one which visits every state in the state-space before cycling; it has only one orbit of length $\tau = N$. As illustrated in Fig. 8-3 the number of fully ergodic dynamics equals the number of orbits of length N: $\mathcal{O}(N) = (N-1)!$. The fraction of invertible dynamics which are fully ergodic is thus

$$(N-1)!/N! = 1/N. \tag{8.14}$$

In the limit where the state-space becomes very large ($N \to \infty$) this fraction approaches zero. Consider however a less restrictive measure, the specific ergodicity, which is an intensive measure of how efficiently a given dynamics explores the state space.

The entropy of the state space is proportional to the natural logarithm of the number of states. For convenience we define the constant of proportionality to be unity (instead of Boltzmann's constant, k_B). So the total entropy of the state space for each dynamics is $S_{\text{tot}} = \ln N$. This is the amount of information necessary to specify a particular state $\mathcal{S}_j$ from the collection of N states. The entropy associated with the set of successors of state $\mathcal{S}_j$ subject to dynamics d_k is $S_{\text{phase}} = \ln \tau_{j,k}$. Consider the expectation value of this quantity over all the states in the dynamics[2], $E_k(\ln \tau_{j,k})$, which gives the average value of S_{phase} for this dynamics. We define specific ergodicity

[2]Note, the expectation value over the ensemble of dynamics is denoted by $\langle\cdot\rangle$, as is the expectation value over all points in the dynamics (which one is intended can easily be determined from context). The expectation value over all states in a specific dynamics is denoted by $E_k(\cdot)$.

for a dynamics d_k as the ratio of the two entropies:

$$\eta_k = \frac{E_k\left(S_{\text{phase}}\right)}{S_{\text{tot}}} = \frac{E_k\left(\ln \tau_{j,k}\right)}{\ln N}. \tag{8.15}$$

Clearly, $0 \leq \eta_k \leq 1$. In the case where d_k is invertible this reduces to the definition of specific ergodicity introduced by Toffoli[119, 120] as a measure of how efficiently a system can "count". His motivation was to find an intensive measure of the amount of computation performed by a dynamics.

We are interested in the expectation value of the specific ergodicity over the ensemble of dynamics, $\langle\eta\rangle$:

$$\langle\eta\rangle = \frac{\langle E_k\left(\ln \tau_{j,k}\right)\rangle}{\ln N}, \tag{8.16}$$

where the brackets indicate an average over the ensemble of all dynamics. Note that the average over the ensemble of dynamics of the average over all points in each dynamics is the same as the average over all points, $\mathcal{S}_{j,k}$, in the ensemble space:

$$\langle E_k\left(\ln \tau_{j,k}\right)\rangle = \frac{1}{DN}\sum_{k=1}^{D}\sum_{j=1}^{N}\ln \tau_{j,k} = \frac{1}{DN}\sum_{t=1}^{N}\left(DN\right)p_{\tau_j}(t)\ln t = \sum_{t=1}^{N}p_{\tau_j}(t)\ln t, \tag{8.17}$$

where we have again written the double sum as in Sec. 8.3.1. The final expression is the ensemble average over all the points, just as in Eqs. 8.3 and 8.7. Thus

$$\langle\eta\rangle = \frac{\langle\ln \tau_{j,k}\rangle}{\ln N}. \tag{8.18}$$

8.5.1 Invertible dynamics

For an invertible dynamics τ_j is a random variable representing the average over all dynamics of the number of distinct states on the orbit containing $\mathcal{S}_j$. In Eq. 8.1 we found that the probability a point in the ensemble is on an orbit of any length is uniformly distributed with respect to the orbit length, which ranges from 1 to N.

Thus:

$$\langle \eta \rangle = \frac{\langle \ln \tau_{j,k} \rangle}{\ln N} = \frac{\sum_{t=1}^{N} p_{\tau_j}(t) \ln t}{\ln N} = \frac{\sum_{t=1}^{N} \ln t}{N \ln N} = \frac{\ln N!}{N \ln N}. \tag{8.19}$$

Taking the asymptotic limit and using Stirling's approximation for the logarithm,

$$\lim_{N\to\infty} \langle \eta \rangle = \lim_{N\to\infty} \frac{N \ln N - N}{N \ln N} = 1. \tag{8.20}$$

Thus in the limit where the size of the state space is large, knowing on which orbit a point is located gives almost no reduction in the uncertainty for locating a specific of the N total points. Since the average value of the specific ergodicity approaches the limiting value of unity, the distribution of values must become sharply peaked for large N. That is, an overwhelming fraction of invertible dynamics must have $\eta \approx 1$. Therefore, in the limiting case where N is large, any dynamics in the ensemble of possible dynamics will approach fully ergodic with high probability.

8.5.2 Unconstrained dynamics

We can also calculate the expected specific ergodicity of an unconstrained dynamics. Since we are interested in how effectively the system can "count" we are interested in knowing how many distinct states are visited by a trajectory. If we choose a point at random from the ensemble of possible dynamics, the probability to form a cycle at iteration t for large N is given by Eq. 8.6. Thus for an unconstrained dynamics

$$\begin{aligned}
\lim_{N\to\infty} \langle \eta^u \rangle &= \lim_{N\to\infty} \frac{\langle \ln \tau_{j,k} \rangle}{\ln N} \\
&= \lim_{N\to\infty} \frac{1}{\ln N} \int_1^N (\ln t)\, p^u_{\tau_j}(t) dt \\
&= \lim_{N\to\infty} \frac{1}{\ln N} \int_1^N \frac{t}{N} (\ln t) \exp\left(-t^2/2N\right) dt \\
&= \lim_{N\to\infty} \frac{1}{\ln N} \left[\frac{1}{2}\left(\ln N + \ln 2 - \gamma_E\right)\right] \\
&= \frac{1}{2}.
\end{aligned} \tag{8.21}$$

Thus in the limit where the state space is large, knowing the trajectory on which a state is located reduces the uncertainty in locating the specific state. Since the average trajectory for an unconstrained dynamics visits on the order of the square root of the total number of states (Eq. 8.7), if we use an unconstrained trajectory as a counter we make use of only about half of the available bits.

Recall that in the invertible case the distribution of values for the specific ergodicity is sharply peaked at the limiting value. In contrast, for unconstrained dynamics, the average value of the specific ergodicity is equal to an intermediate value (actually the middle value). This gives us no information about the width and shape of the distribution function for the η^u's.

We summarize the results for unconstrained and reversible dynamics in Table 8.2. Most of the results summarized have been known to various communities for several years. The definition and calculation of the specific ergodicity for an unconstrained dynamics is an original contribution. Note on the table we have an additional quantity, l_j. This is the length of the limit cycle on which point S_j subject to dynamics k is located. For invertible dynamics the limit cycles are exactly the orbits. For unconstrained dynamics the average length of the limit cycle is one half of the average trajectory length.

8.6 Extensions of these general considerations

The random network limit considered here neglects the property of locality, so we would expect this treatment not to apply specifically to CA's. In addition CA's have extensive conserved quantities. We would expect this extra structure to divide the state space into more orbits than typical. However we have shown that the expected number of orbits for a typical dynamics has the same range as the values of the extensive conserved quantities. Furthermore we can label each orbit and that label is conserved by the dynamics. Thus the existence of extensive conservation laws in

Table 8.2: Comparison of invertible and unconstrained random mappings.

Invertible	Unconstrained
$D = N!$	$D^u = N^N$
$p_{\tau_j}(t) = \frac{1}{N}$	$p^u_{\tau_j}(t) = \frac{(N-1)!}{(N-t)!}\frac{t}{N^t}$
$p_{l_j}(t) = \frac{1}{N}$	$p^u_{l_j}(t) = \sum_{k=t}^{N} \frac{(N-1)!}{(N-k)!N^k}$
$\langle \tau_{j,k} \rangle = \frac{N+1}{2}$	$\left\langle \tau^u_{j,k} \right\rangle = \left(\frac{\pi N}{2}\right)^{\frac{1}{2}}$
$\langle l_{j,k} \rangle = \langle \tau_{j,k} \rangle = \frac{N+1}{2}$	$\left\langle l^u_{j,k} \right\rangle = \frac{1}{2}\left\langle \tau^u_{j,k} \right\rangle = \left(\frac{\pi N}{8}\right)^{\frac{1}{2}}$
$\langle \omega \rangle = \ln N + \gamma_E$	$\langle \omega^u \rangle = \frac{1}{2}\ln N + O(1)$
$\lim_{N\to\infty} \langle \eta \rangle = 1$	$\lim_{N\to\infty} \langle \eta^u \rangle = \frac{1}{2}$

reversible CA's does not *a priori* preclude them from being typical. Moreover, since the distribution of values for the specific ergodicity is peaked at the limiting value, an overwhelming fraction of invertible mappings approach fully ergodic. We conjecture that most reversible CA's approach fully ergodic.

Our empirical observations for CA's with small neighborhoods of interaction also suggest that reversible CA's are typical members of the ensemble of all invertible discrete dynamics, while irreversible CA's are not typical members of the ensemble of all unconstrained discrete dynamics. Reversible CA's with small neighborhoods on average have cycle times longer than we could ever observe. In contrast, irreversible CA's with small neighborhoods of interaction tend to quickly reach a quiescent state, and increasing the size of the space does little to increase the amount of time before the system becomes quiescent.

Chapter 9

Conclusions

9.1 Discussion and conclusions

I have examined issues that lie at the foundations of statistical mechanics: how to reconcile macroscopic dissipation with microscopic reversibility; fair and unbiased sampling of state space for a system driven by an external pseudorandom number generator; the increase of course grained entropy and the approach to equilibrium; and ergodicity of deterministic dynamics on discrete systems. I have gained some insights into the mechanisms of pattern formation and computation in physical systems, and an understanding of how to develop computational models of physics.

These models may help make the statistical physics community aware of the interest of information preserving dynamics. Just as discrete models played a key role in the development and understanding of statistical mechanics, information preserving dynamics may play a key role in the development and understanding of the informational underpinnings of physics.

9.2 Directions for future work

The work presented in this thesis consists primarily of complete articles published in journals. Along the way I have explored several ideas which have yet to be fully developed, and I have implemented models which have yet to be fully understood analytically. I present some of these here, and also ideas for future research. Note that some ideas for the extension of the work described in this thesis can also be found in Secs. 6.6 and 7.5.

9.2.1 Stochastic growth models

We have begun to study an extension of the BD model[4], which incorporates the selective etching and the subsequent redeposition of particles. The "multiple regrowth" model was formulated as an attempt to simulate diamond film growth via the ballistic deposition of carbon in a hydrogen plasma. There are two species of particles: sticking and etching. The sticking particles are analogous to the carbon particles, the etching, analogous to the hydrogen. Both species fall ballistically and stick upon first contacting the growth surface (just as in the BD model). The etching particles however immediately evaporate taking with them singly bonded sticking particles. Based on the ratio of etching to sticking particles we observe a smooth transition in growth morphology from dendritic, to columnar, to crystalline. We have conducted a preliminary study on the scaling properties of the three regimes, but, as of yet, have no conclusions for the values of the scaling exponents. It would be good to have the issues with coupling of ballistic growth models to PRNGs straightened out before proceeding with a detailed numerical study. Sample morphologies for the three regimes are shown in Fig. 9-1.

The BD model appears to be especially sensitive to correlations in PRNGs. It couples to PRNGs that have not been shown to be problematic in the past, despite being used extensively. A reasonable conjecture is that the BD model coupling to the

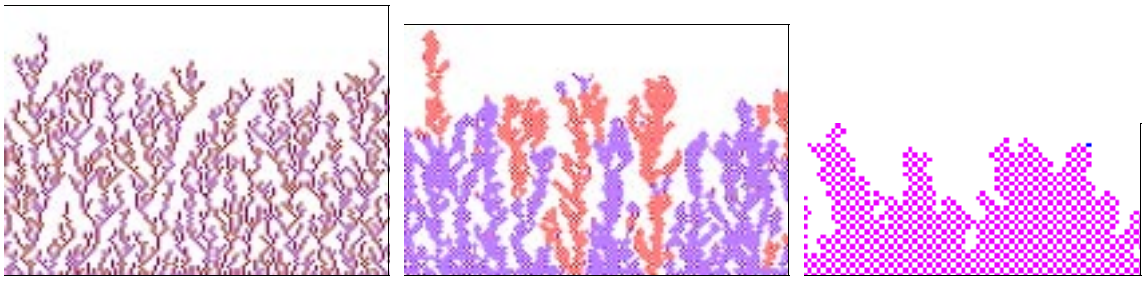

Figure 9-1: Examples of the three regimes of growth: (a) Dendritic growth, (b) columnar growth, (c) smooth growth.

PRNGs depends upon it using every number produced in the sequence it in which it is produced. It would be interesting to introduce a mechanism for the rejection of PRNs at pseudorandom points in the sequence. One approach would be to use a second PRNG as a filter for rejecting PRNs generated by the first, and thus a second independent source of pseudorandomness would influence the dynamics. But note, combining generators is a process with many caveats[73].

9.2.2 Reversible Aggregation

We discuss possible extensions to the RA model in Sec. 6.6, such as implementing a model with independent control of temperature and aggregate size, or a model with an excluded volume for the gas particles (requiring the crystal to do work on the gas as it grows). Introducing a variant which included surface tension of the aggregate may connect the RA model more directly to the standard models of growth, such as DLA and interface growth. An obvious avenue of study is to implement a three-dimensional version of the RA model. These are but a few of the wide range of modifications to the basic RA model which may prove interesting.

9.2.3 Exact versus statistical reversibility

In collaboration with Norman Margolus, we have begun to investigate the differences between exact microscopic reversibility and statistical reversibility, primarily noting

the effects on macroscopic behavior. One way to do this is to take a reversible cellular automata rule which has a long time evolution and interesting macroscopic behavior and modify the rule from being microscopically reversible to being statistically reversible. For example with the RA model, we could use a noninvertible diffusion algorithm. In this case, it would be very surprising to find differences. More illuminating would be to consider asynchronous updating of the space (*i.e.*, selecting lattice sites at random to be updated). An interesting system in which to test this approach is the spin system discussed below (the "SAME" ball) which forms a stable three dimensional structure.

We would also like to develop simple CA models of standard thermodynamic systems, such as a gas of particles pushing on a piston. Explicit microscopically reversible models of these systems could illuminate details of quasistatic processes.

9.2.4 The SAME CA

Studying dynamical spin systems may lead to an understanding of discrete dynamical equations of motion. We would like be able to derive these equations based on discrete symmetries, conservation laws, and constraints. An interesting system deserving of further study is the SAME cellular automaton[88]. This is an Ising-like model. There is one bit of state at each lattice site and each site has a bond with each of its nearest neighbors. If the state at the site and the state of all of the neighbors are the "same" (*i.e.*, they are all 0's or all 1's), then the state at that site is complemented. We apply the dynamics in a checkerboard fashion, so all the sites on the "red" squares are updated simultaneously while the "black" squares are held fixed, and vice versa. By using the checkerboard updating scheme, we never modify nearest neighbors while sites are updated and we preserve exact reversibility (applying the rule twice in a row to the same sublattice inverts the dynamics).

When started from a small block of randomness in a lattice otherwise uniformly filled with all 0's or 1's, the system evolves initially by sending off uniform wave

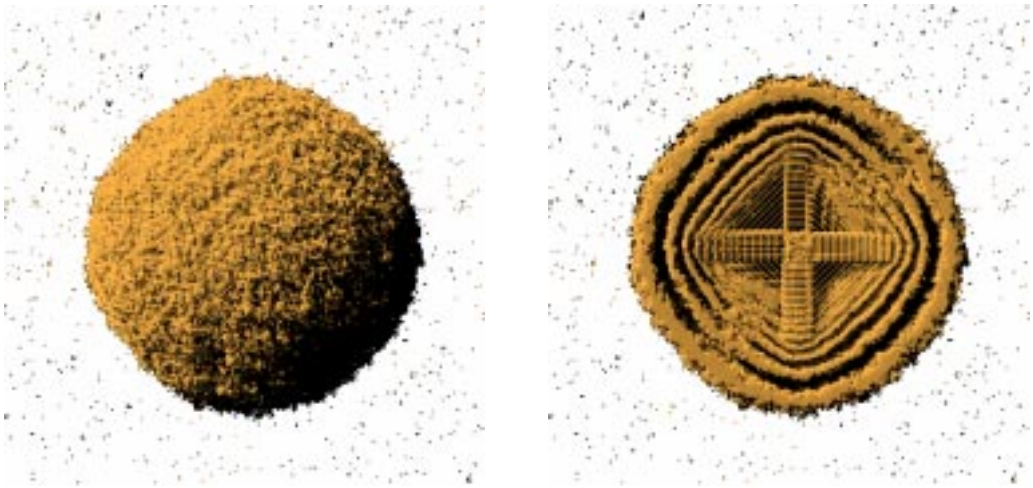

Figure 9-2: The "SAME" rule implemented on a three dimensional square lattice of size $256 \times 256 \times 256$. When started from a small block of random bits in an otherwise uniformly filled lattice the systems evolves to a stable structure resembling a pyramid embedded in fluctuating concentric shells. A rendering of the three dimensional structure is shown (a) from the outside, and (b) with the front half of the structure removed.

fronts originating from the random block. The fronts from the opposite sides of the wave collide at the edges of the space and destroy the order, allowing interior waves to fluctuate. In two dimensions this process continues until all the wave fronts have "peeled away", and the system is completely disordered. In three-dimensions the long time behavior is strikingly different. Initially the waves collide and annihilate, but there is a "surface tension" which keeps waves within a certain radius of the origin from fluctuating out to the edge of the lattice. These waves experience microscopic fluctuations, but stay macroscopically ordered. The resulting structure resembles a pyramid embedded in fluctuating concentric shells, as shown in Fig. 9-2. There are small microscopic fluctuations, but the stable macroscopic structure will persist indefinitely.

Note that a few distinct CA rules produce this identical stable structure. These rules must have similar constraints, conservation laws, and symmetries. A conservation of the dynamics discussed above is, for example, if we label half of the bonds at

each site ferromagnetic and half of the bonds antiferromagnetic, then the bond energy is conserved by the "flip if same" dynamics. We have observed that constraints and conservation laws can maintain the system in an ordered, metastable state, for a prolonged time. Although there is no current theory for the dynamics of CA's with conservation laws, some work was done more than a decade ago on identifying invariants in CA's[100, 89]. Ref. [88] also contains discussion of invariants in CA's. Other researchers have developed a "computational mechanics" approach[20, 58], attempting to use the concept of regular languages (borrowed from automata theory in computer science) to classify patterns produced by CA's.

Aside from being useful for studying conservation laws, this may be an ideal system to use for testing the difference between exact and statistical reversibility. It has also been suggested that this system could be an example of a statistical system with non-extensive entropy.

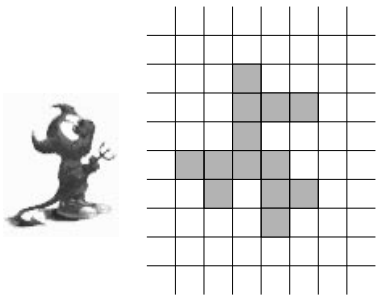

Very small, but lively.

—James Clerk Maxwell, describing his demon.

Bibliography

[1] H. Abelson, T. F. Knight, Jr., G. J. Sussman, and friends. Amorphous computing, MIT AI Laboratory. *http://swissnet.ai.mit.edu/ switz/amorphous/white-paper/amorph-new/amorph-new.html*, 1995-present.

[2] C. Adami. *Introduction to artificial life.* Springer, New York, 1998.

[3] Y. Bar-Yam. *Dynamics of Complex Systems.* Addison Wesley, Reading,MA, 1997.

[4] Y. Bar-Yam and C. T. Capraro. Thin film growth with selective etching: Multiple regrowth model for diamond films. *Computational Materials Science*, 1:169–176, 1993.

[5] A. L. Barabasi and U. E. Stanley. *Fractal Concepts in Surface Growth.* Press Syndicate of the Univ of Cambridge, New York, 1995.

[6] U. Bastolla and G. Parisi. Attraction basins in discretized maps. *J. Phys. A*, 11:3757–3769, 1997.

[7] C. H. Bennett. Logical reversibility of computation. *IBM Journal of Research and Development*, 17:525–532, 1973.

[8] C. H. Bennett. Thermodynamics of computation. *Int. J. of Theor. Phys.*, 21:905–940, 1982.

[9] S. M. Bhattacharjee and A. Khare. Fifty years of the exact solution of the two-dimensional Ising model by Onsager. *Current Science*, 69(10):816–821, 1995.

[10] K. Binder. Applications of Monte Carlo methods to statistical physics. *Rep. Prog. Phys.*, 60(5):487–559, 1997.

[11] B. M. Boghosian, P. V. Coveney, and A. N. Emerton. A lattice-gas model of microemulsions. *P. Roy. Soc. Lond. A*, 452:1221–1250, 1996.

[12] B. M. Boghosian and W. Taylor. Correlations and renormalization in lattice gases. *Phys. Rev. E*, 52(1):510–554, 1995.

[13] B. M. Boghosian, J. Yepez, and N. H. Margolus. Personal communincation.

[14] R. S. Burington and Jr. D. C. May. *Handbook of Probability and Statistics with Tables.* McGraw-Hill Book Co., New York, 1970.

[15] B. Chopard and M. Droz. Cellular automata model for the diffusion equation. *J. Stat. Phys.*, 64(3/4):859–892, 1991.

[16] B. Chopard and M. Droz. *Cellular automata modeling of physical systems.* Cambridge University Press, Cambridge, U.K., 1998.

[17] P. V. Coveney, J. B. Maillet, J. L. Wilson, P. W. Fowler, O. Al-Mushadani, and B. M. Boghosian. Lattice-gas simulations of ternary amphiphilic fluid flow in porous media. *Int. J. Mod. Phys. C*, 9(8):1479–1490, 1998.

[18] M. Creutz. Microcanonical monte carlo simulation. *Phys. Rev. Lett.*, 50(19):1411–1414, 1983.

[19] M. Creutz. Deterministic ising dynamics. *Annals of Physics*, 167:62–72, 1986.

[20] J. P. Crutchfield. The calculi of emergence—computation, dynamics and induction. *Physica D*, 75(1–3):11–54, 1994.

[21] J. P. Crutchfield and D. P. Feldman. Statistical complexity of simple one-dimensional spin systems. *Phys. Rev. E*, 55(2):R1239–R1242, 1997.

[22] M. Daoud, P. Pincus, W. H. Stockmayer, and T. Witten. Phase separation in branched polymer solutions. *Macromolecules*, 16(12):1833–1839, 1983.

[23] B. Derrida and M. R. Evans. Exact correlation functions in an asymmetric exclusion model with open boudaries. *J. Physique*, I3:311, 1993.

[24] B. Derrida and H. Flyvbjerg. The random map model: a disordered model with deterministic dynamics. *J. Physique*, 48:971–978, 1987.

[25] G. D. Doolen, editor. Special issue of *Physica D*, Cambridge, MA, 1989. MIT Press.

[26] R. M. D'Souza. unpublished.

[27] R. M. D'Souza. Reversible pattern formation in Ising-like cellular automata. *manuscript in preparation.*

[28] R. M. D'Souza. Anomalies in simulation of nearest neighbor ballistic deposition. *Int. Jour. of Modern Physics C*, 8(4):941–951, 1997.

[29] R. M. D'Souza, Y. Bar-Yam, and M. Kardar. Sensitivity of ballistic deposition to pseudorandom number generators. *Phys. Rev. E*, 57(5):5044–5052, 1998.

[30] R. M. D'Souza, G. E. Homsy, and N. H. Margolus. Simulating digital logic with the Reversible Aggregation model. In C. Moore and D. Griffeaths, editors, *Proceedings of the Santa Fe Institute workshop on constructive cellular automata.* Oxford University Press, 1999. In press.

[31] R. M. D'Souza and N. H. Margolus. Thermodynamically reversible generalization of diffusion limited aggregation. *Phys. Rev. E.*, 60(1):264–274, 1999. Available as cond-mat/9810258.

[32] D. J. D. Earn and S. Tremaine. Exact numerical studies of Hamiltonian maps—iterating without roundoff error. *Physica D*, 56:1–22, 1992.

[33] M. Eden. In H. P. Yockey, editor, *Symposium on information theory in Biology*, New York, 1958. Pergamon Press.

[34] M. Eden. A two-dimensional growth process. In J. Neyman, editor, *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, volume 4, pages 223–239, Berkeley, 1961. University of California Press.

[35] S.F. Edwards and D.R. Wilkinson. The surface statistics of a granular aggregate. *Proc. Roy. Soc A*, 381:17–25, 1982.

[36] A. Erzan, L. Pietronero, and A. Vespignani. The fixed-scale transformation approach to fractal growth. *Rev. Mod. Phys.*, 67(3):545–604, 1995.

[37] J. Esser and M. Schreckenberg. Microscopic simulation of urban traffic based on cellular automata. *Int. J. Mod. Phys. C*, 8(5):1025–1036, 1997.

[38] E. Ben-Jacob *et al.* Response of bacterial colonies to imposed anisotropy. *Phys. Rev. E*, 53(2):1835–1843, 1996.

[39] P. Manneville *et. al.*, editor. Springer proceedings in physics, vol. 46, New York, 1989. Springer-Verlag.

[40] F. Family. Scaling of rough surfaces: effects of surface diffusion. *J. Phys. A*, 19:L441, 1986.

[41] F. Family and T. Vicsek. Scaling of the active zone in the Eden process on percolation networks and the ballistic deposition model. *J. Phys. A*, 18:L75, 1985.

[42] F. Family and T. Vicsek. *Dynamics of Fractal Surfaces.* World Scientific, Singapore, 1991.

[43] W. Feller. *Introduction to Probability and Statistics*, volume 1. John Wiley & Sons, Inc., New York, 1957.

[44] A. M. Ferrenberg and D. P. Landau. Monte carlo simulations: Hidden errors from "good" random number generators. *Phys. Rev. Lett.*, 69(23):3382–3384, 1992.

[45] A. M. Ferrenberg, D. P. Landau, and K. Binder. Statistical and systematic errors in monte carlo sampling. *Jour. of Stat. Phys.*, 63(5/6):867–882, 1991.

[46] G. Foltin, K. Oerding, Z. Rácz, R. L. Workman, and R. K. P. Zia. Width distribution for random-walk interaces. *Phys. Rev. E*, 50(2):R639–R642, 1994.

[47] M. P. Frank and T. F. Knight, Jr. Ultimate theoretical models of nanocomputers. *Nanotechnology*, 9:162–176, 1998.

[48] E. Fredkin and T. Toffoli. Conservative logic. *Int. J. of Theor. Phys.*, 21(3/4):219–253, 1982.

[49] U. Frisch, B. Hasslacher, and Y. Pomeau. Lattice-gas automata for the Navier-Stokes equation. *Phys. Rev. Let.*, 56(14):1505–1508, 1986.

[50] J. Galambos. *The Asymptotic Theory of Extreme Order Statistics*. Robert E. Krieger Publishing Co., Malabar, FL, 1987.

[51] *Gnu C Library, Sun Release 4.1.* Original code is copyright Regents of the University of California, 1983.

[52] *Gnu C Library, Sun Release 4.1.*

[53] C. Godrèche, editor. *Solids Far From Equilibrium*, Cambridge, 1991. Cambridge University Press.

[54] A. M. Gutin, A. Yu. Grosberg, and E. I. Shakhnovich. Polymers with annealed and quenched branchings belong to different universality classes. *Macromolecules*, 26(6):1293–1295, 1993.

[55] H. Röder H. Brune, C. Romainczyk and K. Kern. Mechanism of transition from fractal to dendritic growth of surface aggregates. *Nature*, 369:469–471, 1994.

[56] T. Halpin-Healy and Y.-C. Zhang. Kinetic roughening phenomena, stochastic growth, directed polymers and all that. *Phys. Rep.*, 254(4-6):215–414, 1995.

[57] T. C. Halsey and M. Leibig. Theory of branched growth. *Phys. Rev. A*, 46(12):7793–7809, 1992.

[58] J. E. Hanson and J. P. Crutchfield. Computational mechanics of cellular automata: An example. *Physica D*, 103(1–4):169–189, 1997.

[59] J. Hardy, O. de Pazzis, and Y. Pomeau. Molecular dynamics of a classical lattice gas: transport properties and time correlation functions. *Phys. Rev. A*, 13(5):1949–1961, 1976.

[60] B. Harris. Probability distributions related to random maps. *Ann. Math. Stat.*, 31:1045–1062, 1960.

[61] B. Harris. The early history of the theory of random mappings. *Advances in applied probability*, 24:768, 1992.

[62] B. Hasslacher. Discrete fluids. *Los Alamos Science*, 1987.

[63] K. Huang. *Statistical Mechanics, 2nd ed.* Wiley, New York, 1987.

[64] E. Ising. *Z. der Physik*, 31:253, 1925.

[65] R. Jullien and R. Botet. Surface thickness in the Eden model. *Phys. Rev. Lett.*, 54:2055, 1985.

[66] M. Kardar, G. Parisi, and Y.-C. Zhang. Dynamic scaling of growing interfaces. *Phys. Rev. Lett.*, 56 (9):889–892, 1986.

[67] S. A. Kauffman. *The Origins of Order*. Oxford University Press, New York, 1993.

[68] H. Kaufman, A. Vespignani, B. B. Mandelbrot, and L. Woog. Parallel Diffusion Limited Aggregation. *Phys. Rev. E*, 52:5602–5609, 1995.

[69] J. Kertesz and D. Wolf. Noise reduction in Eden models: II. surface structure and intrinsic width. *J. Phys. A*, 21:747–761, 1988.

[70] R. W. Keyes and R. Landauer. Minimal energy dissipation in logic. *IBM Journal of Research and Development*, 14:152–157, 1970.

[71] J. M. Kim and J. M. Kosterlitz. Growth in a restricted solid-on-solid model. *Phys. Rev. Lett.*, 62(19):2289–2292, 1989.

[72] J. S. Kirkaldy. Spontaneous evolution of spatiotemporal patterns in materials. *Rep. Prog. Phys*, 55(6):723–795, 1992.

[73] D. E. Knuth. *The Art of Computer Programming*, volume 2/ Seminumerical Algorithms. Addison-Wesley, Reading, MA, 1969.

[74] J. Krug. *Origins of scale invariance in growth processes*. PhD thesis, Heinrich-Heine-Universität Düsseldorf, Nov 1994. (This is a Habilitationschrift thesis).

[75] J. Krug and H. Spohn. Kinetic roughening of growing surfaces. In C. Godrèche, editor, *Solids Far From Equilibrium*. Cambridge University Press, 1991.

[76] R. Landauer. Irreversibility and heat generation in the computing process. *IBM Jour. Res. Dev.*, 3:183–191, 1961.

[77] J. S. Langer. Instabilities and pattern formation in crystal growth. *Rev. Mod. Phys.*, 52(1):1–28, 1980.

[78] J. L. Lebowitz. Microscopic reversibility and macroscopic behavior: Physical explanations and mathematical derivations. In J. J. Brey, J. Marro, J. M. Rubi, and M. San Miguel, editors, *Lecture Notes in Physics*. Springer, 1995.

[79] J. L. Lebowitz. Microscopic origins of irreversible macroscopic behavior. In A. Gervois, D. Iagolnitzer, M. Moreau, and Y. Pomeau, editors, *Statistical Physics, Invited papers from STATPHYS 20*, pages 516–527, 1999. Special issue of *Physica A* **263** (1–4).

[80] J. L. Lebowitz. Statistical mechanics: A selective review of two central issues. *Rev. Mod. Phys.*, 71:S346–S357, 1999.

[81] R. Lenormand. Pattern growth and fluid displacements through porous media. *Physica A*, 140:114–123, 1986.

[82] D. Levesque and L. Verlet. Molecular-dynamics and time reversibility. *J. Stat. Phys.*, 72(3-4):519–537, 1993.

[83] M. Li and P. Vitányi. *An Introduction to Kolmogorov Complexity and It's Applications*. Springer-Verlag, New York, 1993.

[84] M. C. Mackey. *Time's arrow: The origins of thermodynamic behavior*. Springer-Verlag, New York, 1992.

[85] A. Malevanets and R. Kapral. Microscopic model for the FitzHugh-Nagumo dynamics. *Phys. Rev. E*, 55(5):5657–5670, 1997.

[86] N. H. Margolus. Physics-like models of computation. *Physica D*, 10:81–95, 1984.

[87] N. H. Margolus. CAM-8: a computer architecture based on cellular automata. In A. Lawniczak and R. Kapral, editors, *Pattern Formation and Lattice-Gas Automata*. American Mathematical Society, 1996.

[88] N. H. Margolus. Crystalline computation. In A. Hey, editor, *Feynman and Computation*. Addison-Wesley, 1998.

[89] Norman H. Margolus. *Physics and Computation*. PhD thesis, Massachusetts Institute of Technology, May 1987.

[90] A. McKane, M. Droz, J. Vannimenus, and D. Wolf, editors. *Scale Invariance, Interfaces, and Non-equilibrium Dynamics*. New York, 1995.

[91] P. Meakin. The growth of rough surfaces and interfaces. *Phys. Rep.*, 235:189–289, 1993.

[92] P. Meakin, P. Ramanlal, L. M. Sander, and R. C. Ball. Ballistic deposition on surfaces. *Phys. Rev. A*, 34 (4):5091–5103, 1986.

[93] E. Medina, T. Hwa, M. Kardar, and Y.-C. Zhang. Burgers equation with correlated noise: Renormalization-group analysis and applications to directed polymers and interface growth. *Phys. Rev. A*, 39(6):3053–3075, 1989.

[94] K. Moriarty, J. Machta, and R. Greenlaw. Parallel algorithm and dynamic exponent for diffusion-limited aggregation. *Phys. Rev. E*, 55(5):6211–6218, 1997.

[95] T. Nagatani. Unsteady diffusion-limited aggregation. *J. Phys. Soc. Jpn.*, 61(5):1437–1440, 1992.

[96] K. Nagel. Life times of simulated traffic jams. *Int. J. Mod. Phys. C*, 5(3):567–580, 1994.

[97] G. Nicolis and I. Progogine. *Self-Organization in Nonequilibrium Systems*. Wiley-Interscience, New York, 1977.

[98] G. Parisi and N. Sourlas. Critical behavior of branched polymers and the Lee-Yang singularity. *Phys. Rev. Lett.*, 46(14):871–874, 1981.

[99] J. E. Pearson. Complex patterns in a simple system. *Science*, 261:189–194, 1993.

[100] Y. Pomeau. Invariant in cellular automata. *J. Phys. A*, 17:L415–L418, 1984.

[101] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C: the art of scientific computation.* Cambridge University Press, Cambridge, 2nd edition, 1992.

[102] I. Prigogine. *From Being to Becoming.* W. H. Freeman and Company, San Francisco, 1980.

[103] I. Prigogine. Laws of nature, probability and time symmetry breaking. In A. Gervois, D. Iagolnitzer, M. Moreau, and Y. Pomeau, editors, *Statistical Physics, Invited papers from STATPHYS 20*, pages 528–539, 1999. Special issue of *Physica A* **263** (1–4).

[104] W. Gossett publishing under the pseudonym 'Student'. *Biometrika*, 6:1, 1908. For a modern discription of 'Student's work see e.g. M. G. Bulmer, *Principles of Statistics* Dover Publications, Inc., New York, 1979.

[105] Z. Rácz and M. Plischke. Width distribution for 2+1-dimensional growth and deposition processes. *Phys. Rev. E*, 50:3530, 1994.

[106] D. H. Rothman and J. M. Keller. Immiscible Cellular Automata. *J. Stat. Phys.*, 52(3–4):1119–1127, 1988.

[107] D. H. Rothman and S. Zaleski. *Lattice-gas cellular automata: simple models of complex hydrodynamics.* Cambridge University Press, Cambridge, U.K., 1997.

[108] M. Schroeder, J. Shore M. Siegert, D. Wolf, and M. Plischke. Scaling of growing surfaces with large local slopes. *Europhys. Lett.*, 24(7):563–568, 1993.

[109] L. N. Shchur, J. R. Heringa, and H. W. J. Blöte. Simulation of a directed random-walk model: the effect of psuedo-random-number correlations. *Physica A*, 241(3-4):579, 1997.

[110] M. Siegert. Determining exponents in models of kinetic surface roughening. *Phys. Rev. E*, 53(4):3209–3214, 1996.

[111] N. R. S. Simons, M. Cuhaci, N. Adnani, and G. E. Bridges. On the potential use of cellular-automata machines for electromagnetic-field solution. *Int. J. Numer. Model. Electron. N*, 8(3/4):301–312, 1995.

[112] M. A. Smith. *Cellular Automata Methods in Mathematical Physics*. PhD thesis, Massachusetts Institute of Technology, May 1994.

[113] M. A. Smith and Y. Bar-Yam. Cellular automata simulations of pulsed field gel electrophoresis. *Electrophoresis*, 14:337–343, 1993.

[114] M. A. Smith, Y. Bar-Yam, Y. Rabin, B. Ostrovsky, C. H. Bennett, N. Margolus, and T. Toffoli. Parallel-processing simulation of polymers. *Comp. Poly. Sci.*, 2(4):165–171, 1992.

[115] J. Spanier and K. B. Oldham. *An Atlas of Functions*. Harper and Row, New York, 1987.

[116] H. E. Stanley. Fractals and multifractals: The interplay of physics and geometry. In A. Bunde and S. Havlin, editors, *Fractals and Disordered Systems*, New York, 1996. Springer-Verlag.

[117] H. E. Stanley and N. Ostrowsky. *On Growth and Form*. Martinus Nijhoff Publishers, in cooperation with NATO ASI Series, Dordrecht, 1986.

[118] T. Toffoli. Cellular automata as an alternative to (rather than an approximation of) differential equations in modeling physics. *Physica D*, 10:117–127, 1984.

[119] T. Toffoli. Four topics in lattice gases: Ergodicity; Relativity; Information Flow; and Rule compression for parallel lattice-gas machines. In R. Monaco, editor, *Discrete Kinetic Theory, Lattice Gas Dynamics and Foundations of Hydrodynamics*, pages 343–354. World Scientific, 1989.

[120] T. Toffoli. Action, or the fungibility of computation. In A. Hey, editor, *Feynman and Computation*. Addison-Wesley, 1998.

[121] T. Toffoli and N. H. Margolus. *Cellular Automata Machines: A New Environment for Modeling*. MIT Press, Cambridge, MA, 1987.

[122] I. Vattulainen, T. Ala-Nissila, and K. Kankaala. Physical tests for random numbers in simulations. *Phys. Rev. Lett.*, 73(19):2513–2516, 1994.

[123] G. Vichniac. Simulating physics with cellular automata. *Physica D*, 10(1/2):96–116, 1984.

[124] J. Villain. Continuum models of crystal growth from atomic beams with and without desorption. *Journal de Physique I*, (1):19–42, 1991.

[125] M. J. Vold. A numerical approach to the problem of sediment volume. *Jour. Colloid Sci.*, 14:168–174, 1959.

[126] J. von Neumann. *Theory of Self-Reproducing Automata*. University of Illinois Press, Urbana, Illinois, 1966.

[127] R. F. Voss. Multiparticle fractal aggregation. *J. Stat. Phys.*, 36(5/6):861–872, 1984.

[128] T. Williams and Bjerknes. *Nature*, 236:19, 1972.

[129] T. A. Witten and L. M. Sander. Diffusion-Limited Aggregation, a kinetic critical phenomenon. *Phys. Rev. Lett.*, 47(19):1400–1403, 1981.

[130] D. E. Wolf. Cellular automata for traffic flow. *Physica A*, 263(1–4):438–451, 1999. *Appears in* Statistical Physics, Invited papers from STATPHYS 20.

[131] S. Wolfram. Statistical mechanics of Cellular Automata. *Rev. Mod. Phys.*, 55:601–644, 1983.

[132] A. Wuensche. *The global dynamics of cellular automata: An atlas of basin of attraction fields of one-dimensional cellular automata.* Addison-Wesley, Reading, MA, 1992.

[133] A. Wuensche. Discrete dynamical networks and their attractor basins, 1998.

[134] J. Yepez. Lattice-gas crystalization. In Y. H. Qian, editor, *Discrete Models for Fluid Mechanics*, pages 255–294, 1995. Special issue of *J. Stat. Phys.* **81** (1/2).

[135] J. Yepez. *Lattice-gas dynamics.* PhD thesis, Brandeis University, 1997.

[136] J. G. Zabolitzky and D. Stauffer. Simulation of large Eden clusters. *Phys. Rev. A*, 34(2):1523–1530, 1986.