Growth Phenomena in Cellular Automata

JANKO GRAVNER
Mathematics Department
University of California
Davis, CA 95616
e-mail: `gravner@math.ucdavis.edu`

Article outline:

Glossary

Asymptotic density:

The proportion of sites in a lattice occupied by a specified subset is called asymptotic density, or, in short, density.

Asymptotic shape:

The shape of a growing set, viewed from a sufficient distance so that the boundary fluctuations, holes, and other lower order details disappear, is called the asymptotic shape.

Cellular automaton:

A cellular automaton is a sequence of configurations on a lattice which proceeds by iterative applications of a homogeneous local update rule. A configuration attaches a state to every member (also termed a site or a cell) of the lattice. Only configurations with two states, coded 0 and 1, are considered here. Any such configuration is identified with its set of 1's.

Final set:

A site whose state changes only finitely many times is said to fixate, or attain a final state. If this happens for every site, then the sites whose final states are 1 comprise the final set.

Initial set:

A starting set for a cellular automaton evolution is called intial set, and may be deterministic or random.

Metastability:

Metastability refers to a long, but finite, time period in an evolution of a cellular automaton rule, during which the behavior of the iterates has identifiable characteristics.

Monotone cellular automaton:

A cellular automaton is monotone if addition of 1's to the initial configuration always results in more 1's in any subsequent configuration.

Nucleation:

Nucleation refers to (usually small) pockets of activity, often termed nuclei, with long range consequences.

Solidification:

A cellular automaton solidifies if any site which achieves state 1 remains forever in this state.

1 Definition of the Subject and Its Importance

In essence, analysis of growth models is an attempt to study properties of physical systems far from equilibrium (e.g., [53] and its more than 1300 references). Cellular automata (CA) growth models, by virtue of their simplicity and amenability to computer experimentation [60, 68], have become particularly popular in the last 30 years in many fields, such as physics [63, 60, 15], biology [18], chemistry [51, 15], social sciences [7], and artificial life [52]. In contrast to voluminous empirical literature on CA in general and their growth properties in particular, precise mathematical results are rather scarce. A general CA theory is out of the question, since a Turing machine can be embedded in a CA, so that examples as "simple" as elementary one-dimensional CA [17] and Conway's Game of Life [8] are capable of universal computation. Even

the most basic parametrized families of CA systems exhibit a bewildering variety of phenomena: self-organization, metastability, turbulence, self-similarity, and so forth [24, 44, 23, 1]. From a mathematical point of view, CA can be rightly viewed as discrete counterparts to partial differential equations, and so they are able to emulate many aspects of the physical world, while at the same time they are easy to experiment using widely available platforms (from many available simulation programs we mention just [68]).

Despite their resistance to traditional methods of deductive analysis, CA have been of interest to mathematicians from their inception and we will focus on the rigorous mathematical results about their growth properties. The scope will be limited to CA with deterministic update rule — random rules are more widely used in applications [18, 51, 7], but fit more properly within probability theory (however, see e.g. [37] for a connection between deterministic and random aspects of CA growth). Adding randomness to the rule in fact often makes them more tractable, as ergodic properties of random systems are much better understood than those of deterministic ones ([11] provides good examples).

Even though mathematical arguments are the ultimate objective, computer simulations are an indispensable tool for providing the all important initial clues for the subsequent analysis. However, as we explain in a few examples in the sequel, caution needs to be exercised while making predictions based on simulations. Despite the increased memory and speed of commercial computers, for some CA rules highly relevant events occur on spatial and temporal scales far beyond the present-day hardware. Simply put, mathematics and computers are both important, and one ignores each ingredient at one's own peril.

By nature, a subject in the middle of active research contains many exciting unresolved conjectures, vague ideas that need to be made precise, and intriguing examples in search of proper mathematical techniques. It is clear from what has already been accomplished that, often in sharp contrast with the simplicity of the initially posed problem, such techniques may be surprising, sophisticated, and drawn from such diverse areas as combinatorics, geometry, probability, number theory, and PDE. This is a field to which mathematicians of all stripes should feel invited.

2 Introduction

Let us begin with the general set-up. We will consider exclusively binary CA. Accordingly, a *configuration* will be a member of $\{0,1\}^{\mathbb{Z}^d}$, that is, an assignment of 0 or 1 to every site in the d-dimensional lattice $\mathbb{Z}^d$. This divides the lattice into two sets, those that are assigned state 1, called the *occupied* sites, and those in state 0, which are the *empty* sites. A configuration is thus represented by its occupied set A. This set will change in discrete time, its evolution given by $A_0, A_1, A_2, \cdots \subset \mathbb{Z}^d$.

The configuration changes subject to a CA rule, which is, in general, specified by the following two ingredients. The first is a finite *neighborhood* $\mathcal{N} \subset \mathbb{Z}^d$ of the origin, its translate $x + \mathcal{N}$ then being the neighborhood of point x. By convention, we assume that $\mathcal{N}$ contains the origin. Typically, $\mathcal{N} = B_\nu(0, \rho) \cap \mathbb{Z}^d$, where $B_\nu(0, \rho) = \{x \in \mathbb{R}^d : ||x||_\nu \leq \rho\}$ is the ball in the ℓ^ν–norm $|| \cdot ||_\nu$ and ρ is the *range*. When $\nu = 1$ the resulting $\mathcal{N}$ is called the Diamond neighborhood, while if $\nu = \infty$ it is referred to as the Box neighborhood. (In particular, when $d = 2$, range 1 Diamond and Box neighborhoods are also known as von Neumann and Moore neighborhoods, respectively.) The second ingredient is a map $\pi : 2^{\mathcal{N}} \to \{0, 1\}$, which flags the *sufficient*

configurations for occupancy. More precisely, for a set $A \subset \mathbb{Z}^d$, we let $\mathcal{T}(A) \subset \mathbb{Z}^d$ consist of every $x \in \mathbb{Z}^d$ for which $\pi((A_t - x) \cap \mathcal{N}) = 1$. Then, for a given *initial subset* $A_0 \subset \mathbb{Z}^d$ of occupied points, we define $A_1, A_2, \ldots$ recursively by $A_{t+1} = \mathcal{T}(A_t)$.

To explain this notation on arguably the most famous CA of all time, the Game of Life [26, 8] has $d = 2$, Moore neighborhood $\mathcal{N}$ consisting of the origin and nearest eight sites, so that the neighbor of x is

$$x + \mathcal{N} = \begin{matrix} \bullet & \bullet & \bullet \\ \bullet & x & \bullet \\ \bullet & \bullet & \bullet \end{matrix} \quad ,$$

and $\pi(S) = 1$ precisely when either $0 \in S$ and $|S| \in \{3, 4\}$, or $0 \notin S$ and $|S| = 3$. Here, $|S|$ is the size (cardinality) of $S \subset \mathcal{N}$ and note that the center x of the neighborhood itself is counted in the occupation number.

Usually, our starting set A_0 will consist of a possibly large, but finite set of 1's surrounded by 0's. However, other initial states are worthy of consideration, for example, half–spaces, wedges, and sets with finite complements, called *holes*. Finally, for understanding self-organizational abilities and statistical tendencies of the CA rule, the most natural starting set is the random "soup" $\Pi(p)$ to which every site is adjoined independently with probability p.

As already mentioned, we need to consider special classes if we hope to formulate a general theorem. Mathematically, the most significant restriction is to the class of *monotone* (or *attractive*) CA rules, for which $S_1 \subset S_2$ implies $\pi(S_1) \leq \pi(S_2)$. To avoid the trivial case we will also assume that monotone CA have $\pi(\mathcal{N}) = 1$.

Another important notion is that of solidification: we say that the CA *solidifies* if $\pi(S) = 1$ whenever $0 \in S$. In words, this means that once a site becomes occupied, it cannot be removed. To *every CA* on $\mathbb{Z}^d$ given by the rule $(\mathcal{N}, \pi)$ one can associate "space-time" solidification CA on $\mathbb{Z}^d \times \mathbb{Z}$, with unique solidification rule given by the neighborhood $\mathcal{N}' = (\mathcal{N} \times \{-1\}) \cup \{0^{d+1}\}$, and π' such that $\pi'(S \times \{-1\}) = \pi(S)$ for $S \subset \mathcal{N}$. This construction is useful particularly for one-dimensional CA whose space-time version interprets their evolution as a two-dimensional object [66], but we prefer to focus on the growth phenomena in the rule's "native" space.

A more restrictive, but still quite rich, class of rules is the *Threshold Growth (TG)* CA, which is a general *totalistic* monotone solidification CA rule. For such rules, $\pi(S)$ depends only on the cardinality $|S|$ of S whenever $0 \notin S$; therefore, for such S there exists a *threshold* $\theta \geq 0$ such that $\pi(S) = 0$ whenever $|S| < \theta$ and $\pi(S) = 1$ whenever $|S| \geq \theta$.

We will universally assume that a 1 cannot spontaneously appear in a sea of 0's, that is, that 1's only grow *by contact*: $\pi(\emptyset) = 0$. We also find it convenient to assume that π is *symmetric*: $-\mathcal{N} = \mathcal{N}$ and $\pi(-S) = \pi(S)$. This is not a necessary assumption in many contexts, but its absence makes many statements unnecessarily awkward.

Next is a very brief historical lesson. The first paper in CA modeling is surely [69], a precursor to the research into nucleation and self-organization in CA. The follow-up to this pioneering work had to wait until the 70's, when the influential work [43] appeared. The earliest work on CA growth is that of S. Willson [65, 66, 67], which still stands today as one of the notable achievements of mathematical theory. The importance of growth properties of CA, from theoretical and modeling perspectives, was more widely recognized in the mid-80's [54, 55, 60]. At about the same time, statistical physicists recognized the value of mathematical arguments in studying nucleation and metastability and hence the need to build tractable models [63, 64]. Bootstrap percolation ([61, 4, 2], and references therein), one of the most studied CA, which we

discuss in some detail in Section 5, originates from that period. Since the beginning of the 90's there has been a great expansion in the popularity of CA modeling [18, 51], while mathematical theory, which we review in the next three sections, proceeds at a much more measured pace.

The rest of the article is organized as follows. In Section 3 we consider properties of the set which the CA rule generates "at the end of time." In particular, we discuss when the CA eventually occupies the entire available space and, when it fails to do so, what proportion of space it does fill. Section 4 then focuses on the occupation mechanism, in particular on shapes attained from finite seeds. The main theme of Section 5 is nucleation, with the focus on sparse randomly populated initializations. We conclude with Section 6, a summary of issues in need of further research.

3   Final Set

Perhaps the most basic question that one may ask is: what proportion of space does a CA rule ultimately fill? Clearly we need to specify more precisely what is meant by this, but it should be immediately suspected that the answer in general depends on the initial state, even if we only restrict to finite ones. Indeed, consider the TG CA with Moore neighborhood and $\theta = 3$. It is easy to construct an initial set which stops growing, say, one containing fewer than 3 sites. It is not much harder to convince oneself that there exist finite sets (even some with only 3 sites) which eventually make every site occupied. It is a combinatorial exercise to show that these two are the only possibilities in this example. Is this dichotomy valid in any generality? This is one of the questions we address in this section.

Assume a fixed CA rule and the associated transformation $\mathcal{T}$, and fix an initial state A_0. If every $x \in \mathbb{Z}^d$ *fixates*, that is, changes state only finitely many times, then the *final set* $A_\infty = \mathcal{T}^\infty(A_0)$ exists. Notice that this is automatically true for every solidification rule, in which no site can change state more than once.

We say that A_0 *fills space* if $\mathcal{T}^\infty(A_0) = \mathbb{Z}^d$. One cannot imagine a greater ability of a CA rule to "conquer" the environment than if a finite set is able to fill space. Thus it is natural to ask whether there exist general conditions that assure this property, and indeed they do for monotone CA.

Induced by $\mathcal{T}$ is a growth transformation $\bar{\mathcal{T}}$ on closed subsets of $\mathbb{R}^d$, given by

$$\bar{\mathcal{T}}(B) = \{x \in \mathbb{R}^d : 0 \in \mathcal{T}((B - x) \cap \mathbb{Z}^d)\}.$$

In words, one translates the lattice so that $x \in \mathbb{R}^d$ is at the origin, and applies $\mathcal{T}$ to the intersection of Euclidean set B with the translated lattice. It is easy to verify that the two transformations are *conjugate*,

$$\mathcal{T}(B \cap \mathbb{Z}^d) = \bar{\mathcal{T}}(B) \cap \mathbb{Z}^d.$$

It will become immediately apparent why $\bar{\mathcal{T}}$ is convenient. Let S^{d-1} be the set of unit vectors $\mathbb{R}^d$ in and let

$$H_u^- = \{x \in \mathbb{R}^d : \langle x, u \rangle \leq 0\}$$

be the closed half–space with unit outward normal $u \in S^{d-1}$. Then, *provided that the CA rule is monotone*, there exists a $w(u) \in \mathbb{R}$ so that

$$\bar{\mathcal{T}}(H_u^-) = H_u^- + w(u) \cdot u$$

and consequently
$$\mathcal{T}^t(H_u^- \cap \mathbb{Z}^d) = (H_u^- + tw(u) \cdot u) \cap \mathbb{Z}^d.$$

Monotone CA with $w(u) > 0$ for every u are called *supercritical*. A supercritical CA hence enlarges every half-space.

Theorem 1. *Assume a monotone CA rule. A finite set A_0 which fills space exists if and only if $w(u) > 0$ for every direction $u \in S^{d-1}$.*

See [65, 31] for a proof. Before we proceed, a few remarks are in order. First, we should note that one direction of the above theorem has a one-line proof: if $w(u) \le 0$ for some u, then monotonicity prevents the CA from ever occupying a point outside a suitable translate of H_u^-. The other direction is proved by constructing a sufficiently "smooth" initial set. Moreover, supercriticality can be checked on a finite number of directions, in particular one can prove that a two-dimensional TG CA is supercritical if and only if $\theta \le \frac{1}{2}(|\mathcal{N}| - \max\{|\mathcal{N} \cap \ell| : \ell \text{ a line}$ through $0\})$ [31]. Thus, among the TG CA with Moore neighborhood, exactly those with $\theta \le 3$ are supercritical, while this is true for range 2 Box neighborhood when $\theta \le 10$.

A finite set A_0 for which $\cup_t A_t$ is infinite is said to *generate persistent growth*. Further, a CA for which any set that generates persistent growth has $A_\infty = \mathbb{Z}^d$ is called *omnivorous* [31]. For an omnivorous rule a finite seed has either a bounded effect or it fills space.

Is every supercritical TG CA omnivorous? The answer is no, and a counterexample in $d = 2$ is obtained by taking the neighborhood to be the cross of radius 2: $\mathcal{N} = \{(0,0), (0,\pm 1), (0,\pm 2), (\pm 1, 0), (\pm 2, 0)\}$, and $\theta = 2$. It is easy to check that for $A_0 = \{(0,0), (1,0)\}$ the final set A_∞ consists of the x-axis, while initialization with a 2×2 box results in $A_\infty = \mathbb{Z}^2$. On the other hand, the following theorem holds.

Theorem 2. *The two-dimensional TG CA is omnivorous provided either of the two conditions are satisfied:*

(1) $\mathcal{N}$ is box neighborhood of arbitrary range.

(2) $\mathcal{N} = \bar{\mathcal{N}} \cap \mathbb{Z}^2$, where $\bar{\mathcal{N}}$ is a convex set with the same symmetries as $\mathbb{Z}^2$, and $\theta \le \sigma^2/2$, where σ is the range of the largest box neighborhood contained in $\mathcal{N}$.

The theorem is proved in [9] and [10] by rather delicate combinatorial arguments involving analysis of invariant, or nearly invariant, quantities. The lack of robust methods makes conditions in the theorem far from necessary. In particular, proving a general analogue of Theorem 2 without solidification (while keeping monotonicity) is an intriguing open problem.

For non-monotone solidification rules, any general theory appears impossible, but one can analyze specific examples, and we list some recent results below. All are two-dimensional, therefore we assume $d = 2$ for the rest of this section.

In many interesting cases, it is immediately clear from computer simulations that $A_\infty \ne \mathbb{Z}^d$, but at least A_∞ is spread out fairly evenly. This motivates the following definition. Pick a set $A \subset \mathbb{Z}^2$. Let μ_ϵ be ϵ^2 times the counting measure on $\epsilon \cdot A$. We say that A has *asymptotic density* ρ if μ_ϵ converges to $\rho \cdot \lambda$ as $\epsilon \to 0$. Here λ is Lebesgue measure on $\mathbb{R}^2$ and the convergence holds in the usual sense:

(1)
$$\int f \, d\mu_\epsilon \to \rho \cdot \int f \, d\lambda$$

for any $f \in \mathcal{C}_c(\mathbb{R}^2)$. Equivalently, for any square $R \subset \mathbb{R}^2$, the quantity $\epsilon^2 \cdot |R \cap (\epsilon \cdot A)|$ converges to the area of R as $\epsilon \to 0$.

For totalistic solidification CA, the rule is determined by the neighborhood and a *solidification list* of neighborhood counts which result in occupation at the next time. Three neighborhoods have been studied so far: *Diamond* rules with von Neumann neighborhood, *Box* rules with Moore neighborhood and *Hex* rules with the neighborhood $\mathcal{N}$ consisting of $(0,0)$ and the six sites $(\pm 1, 0)$, $(0, \pm 1)$, and $\pm(1,1)$. (We note that this last neighborhood is a convenient way to represent the triangular lattice [60].) These rules are often referred to as *Packard snowflakes* [54, 38, 12]. As an example, in *Hex 135* rule, a 0 turns into a 1 exactly when it "sees" an odd number of already occupied neighbors in its hexagonal neighborhood.

We will assume that 1 is on the solidification list, for otherwise the analysis quickly becomes too difficult (see however [34] and [46] for some results on *Box 2* and *Box 3*) rules. Further, for *Hex* and *Diamond* cases, we will assume 2 is not on this list (or else the dynamics is too similar to a TG CA). We now summarize the main results of [38] and [12].

Theorem 3. *To each of the four Diamond and 16 Hex Packard snowflakes there corresponds a number $\rho \in (0,1]$, the asymptotic density of A_∞, which is independent of the finite seed A_0. The densities in Diamond cases are*

$$\rho_1 = 2/3, \; \rho_{13} = 2/3, \; \rho_{14} = 1, \; \rho_{134} = 29/36.$$

The Hex densities are exactly computable in 8 cases:

$$\rho_{13} = \rho_{135} = 5/6, \; \rho_{134} = \rho_{1345} = 21/22, \; \rho_{136} = \rho_{1356} = \rho_{1346} = \rho_{13456} = 1.$$

In six other Hex cases, one can estimate, within ± 0.0008,

$$\rho_1 \approx 0.6353, \; \rho_{14}, \rho_{145} \approx 0.9689, \; \rho_{15} \approx 0.8026, \; \rho_{16} \approx 0.7396, \; \rho_{156} \approx 0.9378.$$

The final two Hex rules have densities very close to 1:

$$\rho_{146} \in (0.995, 1), \; \rho_{1456} \in (0.9999994, 1).$$

The indices in the densities of course refer to the respective rule.

Note that, in each of the two cases, $\rho_{14} > \rho_{134}$, testimony to the fundamentally nonmonotone nature of these rules. It is also shown in [38] that observing *Hex 1456* from $A_0 = \{0\}$ on even the world's most extensive graphics array, with millions of sites on a side, one would still be led to the conclusion that $A_\infty = \mathbb{Z}^2$. In fact, the site in A_∞^c closest to the origin is still at distance of the order 10^9. Nevertheless, A_∞^c has a positive density and contains arbitrarily large islands of 0's. This is one illustration of limitations in making empirical conclusions based on simulations.

The fundamental tool to prove the above theorem is the fact that these dynamics have an additive component which creates an impenetrable *web* of occupied sites [38]. This web consists of sites at the edge of the light cone, or, to be more precise, the sites which are occupied at the same time at which TG CA with the same neighborhood and $\theta = 1$ would occupy them.

The web makes at least an approximate recursion possible, and the basic renewal theory applies. The delicacy of such results is conveyed effectively by comparison to *Box* solidification. There are 128 such rules with 1 on the solidification list. Although snowflake-like recursive carpets emerge in a great many cases, and exact computations are sometimes feasible, there

is no hope of a complete analysis as in the *Hex* and *Diamond* settings, and many fascinating problems remain. For instance, the density of *Box 1*, provided it exists at all, can depend on the initial seed. Namely, it is shown in [34] that *Box 1* solidification yields density 4/9 starting from a singleton. Later, D. Hickerson (private communication) engineered finite initial seeds with asymptotic densities 29/64 and 61/128. The latter is achieved by an ingenious arrangement of 180 carefully placed occupied cells around the boundary of an 83×83 grid. The highest density with which Box 1 solidification can fill the plane is not known, and neither is whether any seed fills with density less than 4/9. Most initial seeds generate what seems to be a chaotic growth with density about 0.51.

Many other *Box* rules have known asymptotic densities started from a singleton. Here is a sample (D. Griffeath, private communication):

rule	density
12	2/3
13	29/45
15	43/72
16	385/576
17	35/72
18	4/9

Table 1: Densities of A_∞ from $A_0 = \{0\}$ for some *Box* rules.

All exact density computations presented in this section are based on explicit recursions, made possible by an additive web. These recursions are in some cases far from simple, for example, D. Griffeath has shown that in the *Box 12* case, the following formula holds for $a_n = |A_\infty \cap B_\infty(0, 2^n - 1)|$, $n \geq 12$:

$$a_n = \frac{8}{3} \cdot 4^n + r_1 \gamma_3^n - \frac{8}{3} \cdot 3^n - \frac{16}{15} \cdot 2^n + \frac{2}{51} \cdot (-2)^n + 4n - 3 + \frac{8}{3} \cdot (-1)^n + r_2 \gamma_1^n + r_3 \cdot \gamma_2^n,$$

where

$$\gamma_1 \approx -.675 \,, \ \gamma_2 \approx .461 \,, \ \gamma_3 \approx 3.214$$

are the three real roots of the equation $\gamma^3 - 3\gamma^2 - \gamma + 1 = 0$, while

$$r_1 \approx .977 \,, \ r_2 \approx 2.099 \,, \ r_3 \approx .065$$

solve $3145r^3 + 19832r^2 - 22688r - 107648 = 0$.

Apparently very similar rules to those in the above table seem unsolvable, such as *Box 14*, and the "odd" rule *Box 1357* which does have an additive component, but the resulting web from $A_0 = \{0\}$ "leaks" and growth is apparently chaotic. The same problem plagues almost all *Box* rules started from general initial set. The sole exception seems to be the *12* rule, the best candidate for a general general theorem among the 128 rules, due to its quasiadditive web [50]. We should also mention that embedded additive dynamics have been used to study other models [22].

In all considered cases, the web consist of several copies of the final set generated by the space-time solidification associated to a one-dimensional CA. When this CA is linear, the web's
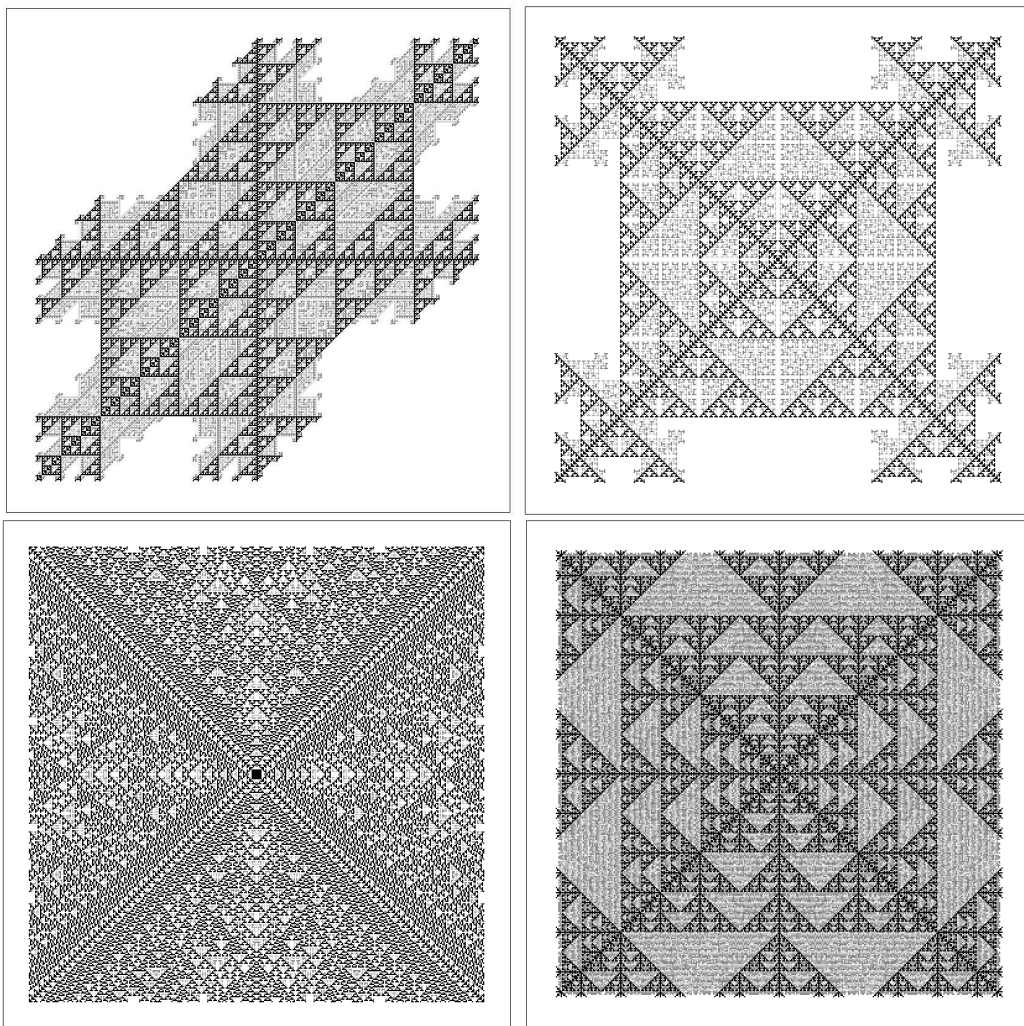
Figure 1: Some Packard snowflakes. Clockwise from top left: *Hex 1*; *Box 1*; *Box 1357*; and again *Box 1*. First three are started from $\{0\}$, and the last from an 8×8 box. The web is black, otherwise the updates are periodically shaded. Note that the chaotic growth can result from a chaotic web (bottom left) or from a leaky web (bottom right).

fractal dimension can be computed using the method from [67]. For example, the properly scaled webs in the top two frames on Figure 2 approach a set with Hausdorff dimension $\log 3/\log 2$, while for the bottom right web this dimension is $\log(1 + \sqrt{5})/\log 2$.

Given that all exactly given densities so far are rational, a natural question is whether there is an example of A_∞ with irrational density. Such example was given by Griffeath and Hickerson in [45], where an initial state for the *Game of Life* is provided for which the set $t^{-1}A_t$ converges to an asymptotic density $(3 - \sqrt{5})/90$ on an appropriate finite set L. This formulation masks the fact that every site x eventually periodically changes its state, so A_∞ does not exist. However, a closer look at the construction shows that the final periods are uniformly bounded. Therefore, if p is the lowest common multiple of all final periods, the p'th iterate of the *Game of Life* rule will generate A_∞ from the same A_0 and with the same density.

This is the only known example of a computable irrational density, and there is a good reason, which we now explain, why such examples are difficult to come by.

By analogy with statistical physics, we would call a set $A \subset \mathbb{Z}^2$ *exactly solvable*, if there exists a formula which decides whether a given x is an element of A. More formally, we require that there exists a finite automaton which, upon encountering x as input, decides whether $x \in A$. Representation of x as input is given as $(\pm i_1^1, \pm i_1^2, i_2^1, i_2^2, \dots)$, where i_1^1, i_1^2 are the most significant binary digits of the first and second coordinate of x; i_2^1, i_2^2 the next most significant, etc. (Some initial i_k^1's or i_k^2's may be 0, and the representation is finite but of arbitrary length.) This means that A is *automatic* [6], or equivalently a *uniform tag system* [16]. With a slight abuse of terminology we call a solidification CA *exactly solvable* (from A_0) if A_∞ is exactly solvable.

To our knowledge, the simplest nontrivial example of an exactly solvable CA is *Diamond 1* solidification, for which it can be shown by induction that $x \notin A_\infty$ iff $\max\{k : i_k^1 = 1\} = \max\{k : i_k^2 = 1\}$. It is easy to construct a (two-state) finite automaton that checks this condition, and the density ρ of A_∞ evidently must satisfy the equation $\rho = 1/2 + \rho/4$, so that $\rho = 2/3$ as stated in Theorem 3. In fact, all of the CA in Theorem 3 with exactly given densities are exactly solvable, and then, by [16, Theorem 6], these densities must be rational. Therefore, the Griffeath-Hickerson example given above is *not* exactly solvable, and the mechanism that forms A_∞ must be more complex in this precise sense. We note that none of the other examples from Theorem 3 are exactly solvable either, but for a different reason [38].

This section's final example, like many other fascinating CA rules, is due to D. Hickerson (private communication). His *Diamoeba* is a rule with the Moore neighborhood and $\pi(S) = 1$ whenever one of the following two conditions is satisfied:

$$0 \notin S, \text{ and } |S| \in \{3, 5, 6, 7, 8\}, \text{ or}$$
$$0 \in S, \text{ and } |S| \in \{6, 7, 8, 9\}.$$

This would be an easily analyzed monotone rule if the 3 were replaced by a 9, with $A_\infty = \emptyset$ for every finite A_0. At first, it seems that the Diamoeba shares this fate. In fact, D. Hickerson has demonstrated that, starting from $A_0 = B_\infty(0, r) \cap \mathbb{Z}^2$, $A_t = \emptyset$ at the smallest t given by

$$12r - 8 - 4r_1 + r_{11} + (r \mod 2),$$

where r_1 and r_{11} are, respectively, the number of 1's and the number of 11's in the binary representation of r. This interesting formula only gives a small taste of things to come (see [34] for a detailed discussion). One of the most intriguing examples is when A_0 is a 2×59 rectangle with a corner cell removed. This grows to a fairly large set in about a million updates, then apparently stops for several million more, after which another growth spurt is possible. The question whether $A_\infty = \mathbb{Z}^2$ for this A_0 is tantalizingly left open. However, there does exist an A_0 for which $A_\infty = \mathbb{Z}^2$. This initialization was discovered by D. Bell, and is an adaptation of a spaceship found by a search algorithm designed by D. Eppstein [21]. This startling object attests to the remarkable design expertise that *Game of Life* researchers have developed over the years.

4   Asymptotic Shapes

After addressing a rule's *ability* to grow in the previous section, we now turn to the *geometry* of growth: is it possible to predict the shape that the set of 1's attains as it spreads? It turns out that the complete answer is known in the monotone case.
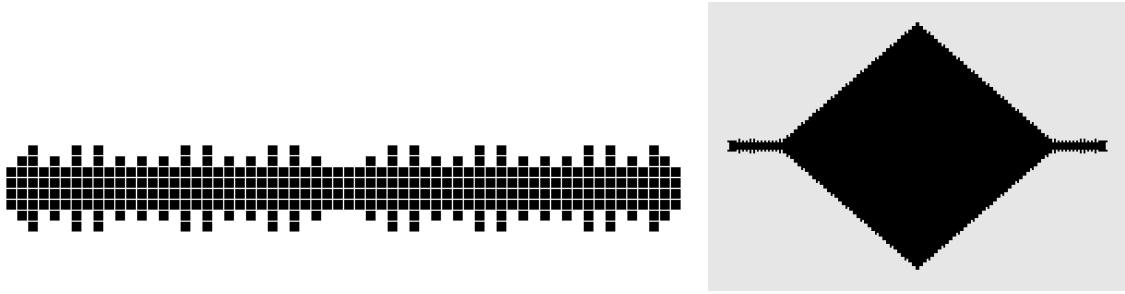
Figure 2: The Bell-Eppstein initial set (left) that results in $A_\infty = \mathbb{Z}^2$ for the *Diamoeba* rule. The set A_t, whose linear asymptotic shape is a rhombus with vertices $(\pm 1/7, 0)$ and $(0, \pm 1/8)$, is shown at $t = 500$.

Naturally, we need a notion of convergence of sets, and the most natural definition is due to Hausdorff (see [30, 32] for an introduction to such issues). We say that a sequence of compact sets $K_n \subset \mathbb{R}^d$ converges to a compact set $K \subset \mathbb{R}^d$ (in short, $K_n \to K$) if, for every $\epsilon > 0$, $K_n \subset K + B_2(0, \epsilon)$ and $K \subset K_n + B_2(0, \epsilon)$, for n large enough. Then we say that a CA has a *linear asymptotic shape* L from a finite initial seed A_0 if

$$\frac{1}{t} A_t \to L$$

as $t \to \infty$.

Turning to monotone CA, we recall the definition of half-space velocities w, and set

$$K_{1/w} = \cup \{[0, 1/w(u)] \cdot u : u \in S^{d-1}\}$$

and let L be the polar transform of $K_{1/w}$, that is,

$$L = K_{1/w}^* = \{x \in \mathbb{R}^d : \langle x, u \rangle \le w(u), \text{ for every } u \in S^{d-1}\}.$$

In general, the polar of a set $K \subset \mathbb{R}^d$ is given by $K^* = \{y \in \mathbb{R}^d : \langle x, y \rangle \le 1 \text{ for every } x \in K\}$. The set L is known as a *Wulff shape*, and is a very important notion in crystallography and statistical physics [56]. The next theorem was proved in the classic paper [66]. The core methods in its proof, as well as proofs of similar results [30], are those of convex and discrete geometry.

Theorem 4. *Assume a monotone CA rule with all $w(u) \ge 0$. Then there exists a large enough r so that for every finite initial set A_0, which contains $B_2(0, r) \cap \mathbb{Z}^d$, the linear asymptotic shape from A_0 equals the Wulff shape L. Even more, the difference between A_t and tL is bounded: there exists a constant C, which depends on the rule and on A_0, so that $A_t \subset tL + B_2(0, C)$ and $tL \subset A_t + B_2(0, C)$ for every $t \ge 0$.*

Note that supercriticality is not assumed here. If $w(u) = 0$ for some u, then $K_{1/w(u)}$ is an infinite object and L has dimension less than d. (The one trivial case is when $w \equiv 0$ and $L = \{0^d\}$.) Finally, note that if there exists a u so that $w(u) < 0$ (and hence $w(-u) < 0$, by symmetry), then A_t is sandwiched between two hyperplanes which approach each other and so eventually $A_t = \emptyset$.

It is also important to point out that $K_{1/w}$ is always a polytope, L is always a convex polytope and both are, for small neighborhoods, readily computable by hand or by computer
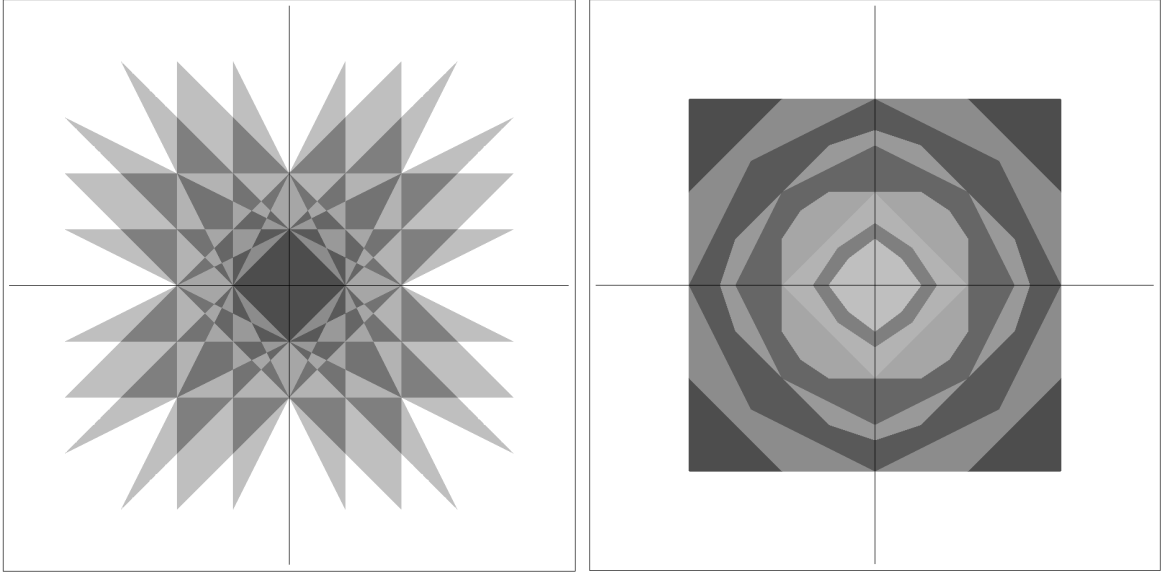
Figure 3: The sets $K_{1/w}$ (left) and the asymptotic shapes for all 10 supercritical range 2 TG CA. Note that there are only 9 shapes, as those with $\theta = 7$ and $\theta = 8$ coincide.

[31, 32, 33]. For example, the Moore neighborhood TG CA with $\theta = 3$ has $K_{1/w}$ with 16 vertices, of which three successive ones are $(0,1)$, $(1,2)$, $(1,1)$, and the remaining 13 are then continued by symmetry. It then follows that the limiting shape L is the convex hull of $(\pm 1/2, 0), (0, \pm 1/2), (\pm 1/3, \pm 1/3)$.

Matters become much murkier when the monotonicity assumption is dropped. We discuss a few interesting two-dimensional solidification examples next. They all hinge on *recursive specification* of iterates A_t for every t (see [34] for a definition). This is far from a general approach (and appears to fail even for simple monotone cases), but is the primary technique available.

We begin with the *Box 25* solidification, starting from $A_0 = B_2(0, r + 1/2) \cap \mathbb{Z}^2$. As was observed in [34], and can be quickly checked by computer, the linear asymptotic shape exists for $r = 2$, $r = 9$ and $r = 13$, but is in each case different, in fact it is convex in the first case, and nonconvex in the other two cases. This demonstrates that such shapes may depend on the initial seed.

A very interesting example was discovered by D. Hickerson (private communication). Consider *Box 37* solidification, with $A_0 = B_2(0, 7/2) \cap \mathbb{Z}^2$. Then $t^{-1/2} A_t$ converges to $B_\infty(0, 2\sqrt{2/3})$ as $t \to \infty$. This demonstrates the possibility of nontrivial *sublinear* asymptotic shapes.

We turn next to the *Hex* rules [38]. These exhibit subsequential limiting shapes, which are not always polygons, as we explain next.

Theorem 5. *Take any of the 16 Hex rules as in Theorem 3, and fix a finite A_0. There exists a one-parameter family of sets $\mathcal{S}_a$, $a \in [0,1]$, so that the following holds: for $t_n = a \cdot 2^n$,*

$$2^{-n} A_{t_n} \to \mathcal{S}_a,$$

as $n \to \infty$.

Furthermore, when 3 and 4 are not both on the solidification list, the family $\mathcal{S}_a$ is called simple and is independent of the initial set. In the opposite, diverse case, initial sets are divided into two classes, distinguished by two different families $\mathcal{S}_a$.

For rational a, it can be shown that the Hausdorff dimension of $\partial \mathcal{S}_a$ always exists, and is in principle computable. For example, for the simple $\mathcal{S}_a$ this dimension equals 5/4 for $a = 14/15$, evidently producing a non-polygonal subsequential shape.

This discussion brings forth the following question, which is probably the most interesting open problem on CA growth. For a *prescribed* set L, can we find a CA with linear asymptotic shape L, attained from a "generic" collection of initial sets? In particular, can L be a circle, thereby giving rise to asymptotic isotropy?

We note that the isotropic construction is possible for probabilistic CA [40], so it seems likely that the answer is yes for a properly constructed chaotic growth. However, techniques for such an approach are completely lacking at present. We should also remark that computational universality should allow for a construction of a CA and a *carefully engineered* initial state with circular (or any other) shape — although this has never been explicitly done. This would, however, violate the requirement of generic initialization.

We conclude this section by a short review of *reverse* shapes [35]. The question here is, if the initial set A_0 is a large hole, and evolves until shortly before the entire lattice is occupied, what is the resulting shape? The initial state has a large and persistent effect on the dynamics and thus the reverse shape geometry will depend on it. The detailed analysis depends on technical convexity arguments, but the cleanest instance is given by the following result.
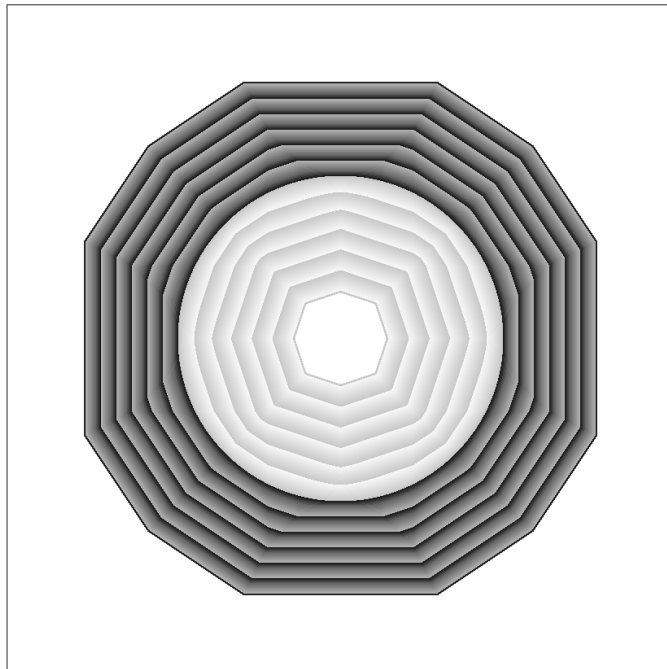

Figure 4: Superimposed convergence to the linear asymptotic shape and to the reverse shape, from, respectively, the interior and the exterior, of a large lattice circle. The rule is TG CA with range 2 and $\theta = 6$. Iterates are periodically shaded.

Theorem 6. *Assume a monotone CA, with $w \geq 0$ but not identically 0 on S^{d-1}. Assume also that its rule $\mathcal{T}$ preserves all symmetries of the lattice $\mathbb{Z}^d$. Pick a closed convex set $H \subset \mathbb{R}^d$, which has all symmetries of $\mathbb{Z}^d$, and let $A_0 = (mH)^c \cap \mathbb{Z}^d$ for some large m. Moreover, let*

$$T = \inf\{t : 0 \in A_t\}.$$

There is a nonempty bounded convex subset $\mathcal{R}(H) \subset \mathbb{R}^d$ such that

$$\lim_{M \to \infty} \lim_{m \to \infty} \frac{1}{M} \cdot A_{T-M} = \mathcal{R}(H)^c,$$

in the Hausdorff sense. Moreover, if

$$h_0 = \max\{h > 0 : h \cdot H^* \subset K_{1/w}\},$$

then

$$\mathcal{R}(H) = (h_0 \cdot H^* \cap \partial K_{1/w})^*$$

In words, one scales the polar H^* so that it touches the boundary of $K^{1/w}$; at this point, the intersection determines the reverse shape. (The shape does not change if H is multiplied by a constant, so h_0 determines its natural scale.) The paper [35] has many more details and examples.

5   Nucleation

In this section we assume that the initial state A_0 is the product measure $\Pi(p)$, with density $p > 0$ that is typically very small. Initially, then, there will be no significant activity on most of the space. Certainly this is no surprise as most of the space is empty, but isolated 1's or small islands of them are often not able to accomplish much either. Most of the lattice is thus in a *metastable* state. However, at certain rare locations there may, by chance, occur local configurations, which are able to spread their influence over large distances until they statistically dominate the lattice. These locations are called *nuclei*, and their frequency and mechanism of growth are the main self-organizational aspects of the CA rule. The majority of results are confined to two-dimensions, so we will assume $d = 2$ for the rest of this section and relegate higher dimensions to remarks.

We start with a simple example, for which we give a few details to introduce the basic ideas and demonstrate that a CA can go through more than one metastable state. For this example we do not specify the map π, but instead give a more informal description. In a configuration A, we call an *insurance* five sites in a cross formation in the state 1, or, more formally, a translate of von Neumann neighborhood which is inside A. The map $\mathcal{T}$ changes any 0 with a 1 in its von Neumann neighborhood to 1. Moreover, it automatically changes any 1 to 0, except that any 1 whose von Neumann neighborhood intersects with an insurance remains 1. Then, for every $\epsilon > 0$, as $p \to 0$,

$$P(0 \in A_t^c \text{ for all } t \leq p^{-1/2+\epsilon}) \to 1,$$
$$P(0 \in (A_t \text{ xor } A_{t+1}) \text{ for all } p^{-1/2-\epsilon} \leq t \leq p^{-5/2+\epsilon}) \to 1,$$
$$P(0 \in A_t \text{ for all } p^{-5/2-\epsilon} \leq t) \to 1.$$

(Here, xor is the exclusive union.) Roughly, most sites are 0 up to time $p^{-1/2}$, then periodic with period 2 up to time $p^{-5/2}$, and 1 afterwards. (In fact, stronger statements, along the lines of Theorem 7 below, are possible.)

The proof has two phases: the first deterministic and the second probabilistic. For the deterministic one, let $d_1(x)$ be the ℓ^1 distance from x to A_0, and assume that A_0 contains no insurance. Then one can prove by induction that, first, none of the A_t contain an insurance, and second, that for every x and $t \geq d_1(x)$, $x \in A_t$ precisely when $(t - d_1(x)) \mod 2 = 0$. On the other hand, an insurance in A_0 centered at the origin will result in $x \in A_t$ for every $t \geq d_1(x)-1$. The probabilistic part consists of noting that, with overwhelming probability when p is small, $B_1(0, p^{-1/2+\epsilon})$ (resp. $B_1(0, p^{-5/2+\epsilon})$) contains no 1 (resp. insurance) in $A_0 = \Pi(p)$, while $B_1(0, p^{-1/2-\epsilon})$ (resp. $B_1(0, p^{-5/2-\epsilon})$) does.

The bulk of the mathematical theory of nucleation and metastability addresses monotone CA, although some work has been done on the *Game of Life* [28] and its generalizations [23, 1], excitable media dynamics [43, 24, 25, 19], and artificial life models [52].

Our first general class are supercritical monotone solidification CA. (In fact, the solidification assumption is not necessary, but reduces technical details so much that it is assumed in most published works.) Such rules have two *nucleation parameters*. Let γ be the smallest i for which there exists an A_0 with $|A_0| = i$ that generates persistent growth. Moreover, let ν be the number of sets A_0 of size γ that generate persistent growth and have the leftmost among their lowest sites at the origin. (The last requirement ensures that ν counts the number of distinct smallest "shapes" that grow.) We call the rule *voracious* if, started from any of the ν initial sets A_0 described above, $A_\infty = \mathbb{Z}^2$. Voracity is a weak condition, which assures a minimal regularity of growth and can, for any fixed rule, be checked on finitely many cases (which is not true for the more restrictive omnivorous property).

For illustration, we briefly discuss these for range ρ Box neighborhood TG CA. For relatively small θ, $\gamma = \theta$; for example, when $\rho = 1$, $\gamma = \theta$ for all three supercritical rules, while when $\rho = 2$, γ exceeds θ only for $\theta = 10$, when it equals 11. For large ρ, and $\theta \sim \alpha\rho^2$, γ is asymptotically the smallest possible (that is, $\gamma \sim \alpha\rho^2$) when $\alpha < \alpha_c$ for some $\alpha_c \in (1.61, 1.66)$ [33]. One can also compute some ν, before they become too large.

	$\theta = 2$	$\theta = 3$	$\theta = 4$	$\theta = 5$	$\theta = 6$	$\theta = 7$
$\rho = 1$	12	42				
$\rho = 2$	40	578	4683	24938	94050	259308

Table 2. Nucleation parameter ν for small Box neighborhood TG CA.

Returning to $A_0 = \Pi(p)$, the most natural statistics to study is

$$T = \inf\{t : 0 \in A_t\},$$

the first time the CA occupies the origin.

Theorem 7. *Assume a monotone, supercritical, and voracious CA, with nucleation parameters* γ *and* ν. *Then, as* $p \to 0$,

$$\sqrt{\nu p^\gamma} \cdot T$$

15

converges in distribution to a nontrivial random variable τ, which is a functional of a Poisson point location $\mathcal{P}$ with unit intensity.

That $T \approx p^{-\gamma/2}$ can be easily guessed (and proved), but the more precise asymptotics described above require a considerable argument [31], as interaction between growing droplets is nontrivial. In particular, the higher dimensional version has not been proved, and the description of the limiting "movie" from $\mathcal{P}$ probably cannot avoid viscosity methods from PDE [59].

The most exciting nucleation results have been proved about *critical* models, for which $w(u)$ vanishes for some direction u but is positive for others. Although a general framework is presented in [36], we will instead focus on the most studied examples. Of these the most popular has been the *bootstrap percolation (BP)*, which is TG CA with von Neumann neighborhood and $\theta = 2$ [2, 4, 3, 61]. Its *modified* version *(MBP)* has the same neighborhood, still solidifies, but when $0 \notin S$, $\pi(S) = 1$ precisely when $\{\pm e_1\} \cap S \neq \emptyset$ and $\{\pm e_2\} \cap S \neq \emptyset$. (Here e_1 and e_2 are the basis vectors.)

Now $w(\pm e_1) = w(\pm e_2) = 0$, so no finite set can generate persistent growth, and it is not immediately clear that $P(T < \infty) = 1$. This is true [61], as very large sets are able to use sparse but helpful smattering of 1's around them and so are unlikely to be stopped. To determine the size of T, one needs more information about the necessary size of these nuclei, and the likelihood of their formation. This was started in [4] and culminated by the following theorem by A. Holroyd [47], which is arguably the crowning achievement of CA nucleation theory to date.

Theorem 8. *For BP let $\lambda = \pi^2/18$, and for MBP let $\lambda = \pi^2/6$. Then, for every $\epsilon > 0$,*

$$P(p \log T \in [\lambda - \epsilon, \lambda + \epsilon]) \to 1$$

as $p \to 0$.

To summarize, $T \approx \exp(\lambda/p)$, which is for small p a long time indeed and amply justifies the description of the almost empty lattice as metastable.

The most common formulation of the theorem above involves finite $L \times L$ squares with periodic boundary instead of infinite lattices. Then

$$I(L, p) = P(\text{the entire square is eventually occupied})$$

and, as $p \to 0$,

$$I(L, p) \to 1 \quad \text{if} \quad p \log L \geq \lambda + \epsilon,$$
$$I(L, p) \to 0 \quad \text{if} \quad p \log L \leq \lambda - \epsilon.$$

Here L is of course assumed to increase with p. Before the value of λ was known, this second formulation was used to estimate it by simulation. For example, [3] used L close to $30,000$ and obtained $\lambda \approx 0.245$ for BP, about a factor of two smaller than the true value $0.548\dots$. Other simulations of BP, MBP, and related models all exhibit a similar discrepancy. The reason apparently is that nuclei are, for realistic values of p, quite a bit more frequent than the asymptotics would suggest. Indeed, the following result from [39] confirms this.

Theorem 9. *For BP and MBP,*

$$I(L, p) \to 1 \quad \text{if} \quad p \log L \geq \lambda - c(\log L)^{-1/2},$$

for an appropriate constant c.

This alone indicates that to halve the error in approximating λ on an $L \times L$ system it is necessary to replace L by L^4. In addition, [39] shows that for the more tractable MBP one can do explicit calculations to conclude that to get an estimate of λ within 2%, one would need L at least 10^{500}, a non-achievable size.

For BP, the quantity $p \log L$ is the "order parameter," the quantity that, when varied, causes a phase transition (which, in addition, is sharp by Theorem 8). We will list now some other models with known order parameters — we also indicate the status of phase transition, when known:

- CA with von Neumann neighborhood and $\pi(S) = 1$ when $|S \setminus \{0\}| \geq 2$: $p^2 \log L$ [57];

- TG CA with range ρ Box neighborhood, $\theta \in [2\rho^2 + \rho + 1, 2\rho^2 + 2\rho]$: $p^{\theta - 2\rho^2 - \rho} \log L$ [31];

- TG CA with $\mathcal{N} = \{(0,0), (0, \pm 1), (\pm 1, 0), (\pm 2, 0)\}$, $\theta = 2$: $p^{3/2} L$, not sharp [31];

- TG CA with $\mathcal{N} = \{(0,0), (0, \pm 1), (\pm 1, 0), (\pm 2, 0)\}$, $\theta = 3$: $(-\log p)^{-2} p \log L$ [31, 62];

- TG CA with range ρ cross neighborhood $\mathcal{N} = \{(x, y) : |x| \leq \rho, |y| \leq \rho, xy = 0\}$ and $\theta = \rho + 1$: $p \log L$, sharp at $\lambda = \pi^2 / (3(\rho + 1)(\rho + 2))$ [49];

- TG CA on $\mathbb{Z}^d$ with $\mathcal{N} = B_1(0, 1) \cap \mathbb{Z}^d$ and $\theta \in [3, d]$, $p^{1/(d - \theta + 1)} \log_{\theta - 1} L$ (where $\log_k$ is the k'th iterate of log) [14], sharp at $\lambda = \pi^2 / 6$ for the modified version when $\theta = d$ [48].

Note that when $\theta = d = 3$ the last example gives the metastable scale $\exp(\exp(\lambda/p))$ [58, 13], making an even modest experimental approximation of λ impossible.

There are other interesting issues about critical growth models, which do not have to do with nucleation. One is decay rate for the first passage time T [5], which is connected to the properties of the very last holes to be filled in. Another is its ability to overtake random obstacles [41].

Apart from sending $p \to 0$, one could vary other parameters to get the metastability phenomena, and one natural example is the range. We explain this scenario on a non-monotone CA known as the *Threshold Voter Automaton (TVA)* [20, 32]. For simplicity, assume $\mathcal{N}$ is range ρ Box neighborhood, and fix a threshold θ. This rule makes a site change its "opinion" by contact with at least θ of the opposite opinions:

$$\pi(S) = 1 \quad \text{iff} \quad (0 \in S \text{ and } |S^c| < \theta) \text{ or } (0 \notin S \text{ and } |S| \geq \theta).$$

As the two opinions are symmetric, the most natural initial state of TVA is $\Pi(1/2)$. We also assume that ρ is large and the scaling $\theta = a|\mathcal{N}|$, for some $a \in (0, 1)$. It is proved in [20] that when $a > 3/4$, any fixed $x \in \mathbb{Z}^2$ changes its opinion only finitely many times with probability approaching 1 as $\rho \to \infty$ — and the rigorous results end there. The most interesting rare nucleation questions arise when $a \in (1/4, 3/4) \setminus \{1/2\}$. According to simulations, under this assumption the nuclei are rare and eventually tessellate the lattice into regions of consensus with either stable or periodic boundaries [32]. However, the definition of a nucleus is unclear and consequently their density cannot be estimated. Two torus simulations are given in Figure 5; it is important to point out that, for such *finite* systems, Lyapunov methods of [27] imply that every site eventually fixates or becomes periodic with period 2.

The *majority* TVA, when $a = 1/2$, is perhaps the most appealing of all [44]. The nucleation is now not rare; instead, this CA quickly self-organizes into visually attractive curvature driven
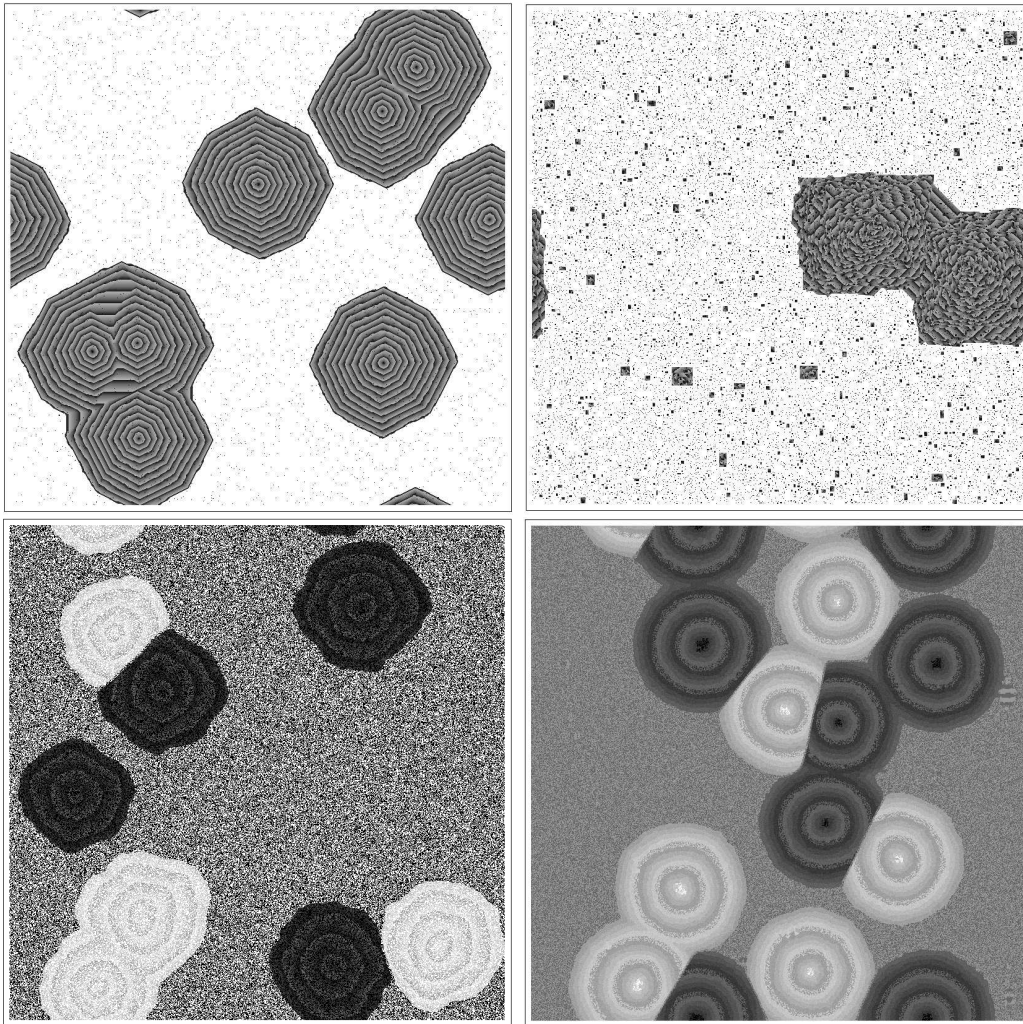
Figure 5: Four nucleation examples, each on an 800×800 array with periodic boundary. Clockwise from top left: TGM CA with Moore neighborhood, $\theta = 3$, and $p = 0.006$; bootstrap percolation with $p = 0.041$; TVA with $\rho = 10$ and $\theta = 194$; TVA with $\rho = 10$ and $\theta = 260$. The iterates are periodically colored to indicate growth and, in the TVA frames, the lighter shades indicate 0's.

dynamics. (Note that flat interfaces between the two opinions are now stable, so one opinion can advance only when it forms a concavity.) For any fixed ρ this must eventually stop, as finite islands with uniformly small enough curvature of either opinion are stable. However, when ρ is increased this effect is with large probability not felt on any fixed finite portion of space. (A similar effect is achieved by the Vichniac "twist" [64].) Many fascinating questions remain about this case, especially on the initial nucleating phase, whose analysis depends on delicate properties of random fields and remains an open problem.
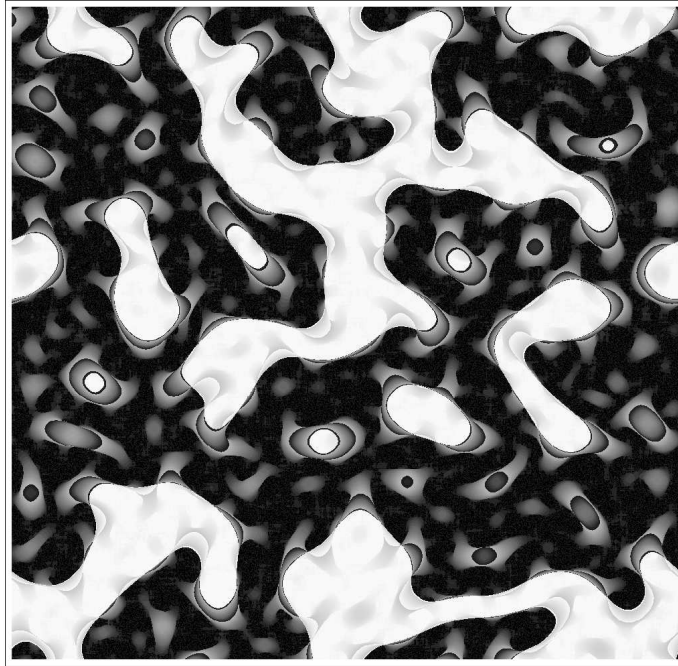
Figure 6: Majority vote: TGM with $\rho = 10$, $\theta = 221$ on a 1000×1000 array with periodic boundary. Again, iterates are periodically colored with the lighter shades reserved for 0's.

6 Future Directions

We will identify seven themes, which connect to open problems discussed in previous sections. Progress on each is bound to be a challenge, but also a significant advance in understanding CA growth.

6.1 Regularity of Growth

It is often important, and of independent interest, to be able to conclude that a cellular automaton rule generates growth without arbitrarily large tentacles, holes, or other undesirable features. An omnivorous CA, for example, has this property. The natural goal would be to develop techniques to establish such regularity for much more general monotone and non-monotone CA, and for arbitrary dimension. Many rules give the impression that regular growth is a generic trait, i.e., holds for a majority of initial sets.

6.2 Oscillatory Growth

Does there exist a class of CA with growing sets that oscillate on different scales? Hickerson's *Diamoeba* might be able to accomplish this from some initial sets, but perhaps there are other, more tractable, examples with identifiable mechanisms.

6.3 Analysis of Chaotic Growth

One look at the growth of *Box 1* solidification from a 8×8 initial box (bottom left frame in Figure 1) would convince most observers that it has a square asymptotic shape. However, there are no tools to prove, or disprove, this statement. A fully rigorous theory of chaotic CA, tailored to address such asymptotic issues, is almost nonexistent and constitutes perhaps the most important challenge for mathematicians in this area.

6.4 Nucleation Theory for Non-Monotone Models

Once nucleation centers are established, growth most often proceeds in a random environment, which consists of debris left over from the nucleation phase. This may help the analysis, as it adds a random perturbation to what may otherwise be intractable dynamics, but on the other hand random environment processes are notoriously tricky to analyze [42].

6.5 Robust Exact Constants and Sharp Transitions

The nucleation phase transition has been proved sharp for a few critical models, by rather delicate arguments. A more robust approach would extend them, and would perhaps provide further insights into the error terms, for which only one sided estimates are now known. The apparent *crossover* [3] phenomenon would also be interesting to understand rigorously.

6.6 Three-Dimensional Nucleation and Growth

With advances in computer power, extensive three-dimensional CA simulations have become viable on commercial hardware. Therefore, it may be possible to investigate nucleation, droplet interaction, clustering mechanisms, and other staples of two-dimensional CA research, at least experimentally. Proper visualization tools of complex three-dimensional phenomena may well require some novel ideas in computer graphics.

6.7 Generic Properties of CA with Large Range

A TG CA with range ρ, say, has on the order of ρ^2 possible thresholds θ. When can it be established that some property holds for the majority of relevant choices? One such property (sensitivity of shapes to random perturbations in the rule) was analyzed from this perspective in [37], but it would be interesting to provide further examples for TG or other classes of CA.

References

[1] A. Adamatzky, G. J. Martnez, J. C. S. T. Mora, *Phenomenology of reaction-diffusion binary-state cellular automata*, Internat. J. Bifur. Chaos Appl. Sci. Engrg. 16 (2006), 2985–3005.

[2] J. Adler, *Bootstrap percolation*, Physica A 171 (1991), 453–4170.

[3] J. Adler, D. Staufer, A. Aharony, *Comparison of bootstrap percolation models*, J. Phys. A: Math. Gen. 22 (1989), L279–L301.

[4] M. Aizenman, J. Lebowitz, *Metastability effects in bootstrap percolation*, J. Phys. A: Math. Gen. 21 (1988), 3801–3813.

[5] E. Andjel, T. S. Mountford, R. H. Schonmann, *Equivalence of decay rates for bootstrap percolation like cellular automata*, Ann. Inst. H. Poincaré 31 (1995), 13–25.

[6] J.-P. Allouche, J. Shallit, "Automatic Sequences: Theory, Applications, Generalizations," Cambridge University Press, 2003.

[7] T. Bäck, H. Dörnemann, U. Hammel, P. Frankhauser, *Modeling urban growth by cellular automata*, Lecture Notes In Computer Science 1141, Proceedings of the 4th International Conference on Parallel Problem Solving from Nature, pgs. 636–645. Springer-Verlag, 1996.

[8] E. R. Berlekamp, J. H. Conway, R. K. Guy, "Winning Ways for Your Mathematical Plays," Second Edition, Vol. 4, Chapter 25, Peters 2004.

[9] T. Bohman, *Discrete threshold growth dynamics are omnivorous for box neighborhoods*, Trans. Amer. Math. Soc. 351 (1999), 947–983.

[10] T. Bohman, J. Gravner, *Random threshold growth dynamics*. Random Structures and Algorithms 15 (1999), 93–111.

[11] M. Bramson, C. Neuhauser, *Survival of one-dimensional cellular automata under random perturbations*, Ann. Probab. 22 (1994), 244–263.

[12] C. D. Brummitt, H. Delventhal, M. Retzlaff, *Packard snowflakes on the von Neumann neighborhood*, to appear in Journal of Cellular Automata.

[13] R. Cerf, E. N. M. Cirillo, *Finite size scaling in three-dimensional bootstrap percolation*, Ann. Probab. 27 (1999), 1837–1850.

[14] R. Cerf, F. Manzo, *The threshold regime of finite volume bootstrap percolation*, Stochastic Process. Appl. 101 (2002), 69–82.

[15] B. Chopard, M. Droz, "Cellular Automata Modeling of Physical Systems." Cambridge University Press, 1998.

[16] A. Cobham, *Uniform tag sequences*, Math. Systems Theory 6 (1972), 164–192.

[17] M. Cook, *Universality in elementary cellular automata*, Complex Systems 15 (2005), 1–40.

[18] A. Deutsch, S. Dormann, "Cellular Automata Modeling of Biological Pattern Formation." Birkhäuser, 2005.

[19] R. Durrett, J. E. Steif, *Some rigorous results for the Greenberg-Hastings model*, J. Theoret. Probab. 4 (1991), 669–690.

[20] R. Durrett, J. E. Steif, *Fixation results for threshold voter systems*, Ann. Probab. 21 (1993), 232–247.

[21] D. Eppstein, *Searching for spaceships,* in "More games of no chance (Berkeley, CA, 2000)," pgs. 351–360. Cambridge University Press, 2002.

[22] K. M. Evans, *Replicators and Larger than Life examples*. In "New Constructions in Cellular Automata," eds. D. Griffeath and C. Moore, pgs. 119–159. Oxford University Press, 2003.

[23] K. M. Evans, *Larger than Life: digital creatures in a family of two-dimensional cellular automata*. in "Discrete Mathematics and Theoretical Computer Science," Vol. AA, eds. R. Cori, J. Mazoyer, M. Morvan and R. Mosseri, 2001, 177–192.

[24] R. Fisch, J. Gravner, D. Griffeath, *Threshold-range scaling for the excitable cellular automata.* Statistic and Computing 1 (1991), 23–39.

[25] R. Fisch, J. Gravner, D. Griffeath, *Metastability in the Greenberg-Hastings model.* Ann. Appl. Probab. 3 (1993), 935–967.

[26] M. Gardner, *Mathematical games,* Scientific American 133 (December 1976), 124–128.

[27] E. Goles, S. Martinez, "Neural and Automata Networks," Kluwer, 1990.

[28] N. M. Gotts, *Self-organized construction in sparse random arrays of Conway's Game of Life*. In "New Constructions in Cellular Automata," eds. D. Griffeath and C. Moore, pgs. 1–53. Oxford University Press, 2003.

[29] J. Gravner, *Growth phenomena in cellular automata*. In "New Constructions in Cellular Automata," eds. D. Griffeath and C. Moore, pgs. 161–181. Oxford University Press, 2003.

[30] J. Gravner, D. Griffeath, *Threshold growth dynamics.* Trans. Amer. Math. Soc. 340 (1993), 837–870.

[31] J. Gravner, D. Griffeath, *First passage times for the threshold growth dynamics on $\mathbf{Z}^2$.* Ann. Probab. 24 (1996), 1752–1778.

[32] J. Gravner, D. Griffeath, *Multitype threshold voter model and convergence to Poisson–Voronoi tessellation*. Ann. Appl. Probab. 7 (1997), 615–647.

[33] J. Gravner, D. Griffeath, *Nucleation parameters in discrete threshold growth dynamics*. Experimental Math. 6 (1997), 207–220.

[34] J. Gravner, D. Griffeath, *Cellular automaton growth on $\mathbb{Z}^2$: theorems, examples and problems*. Advances in Applied Mathematics 21 (1998), 241–304.

[35] J. Gravner, D. Griffeath, *Reverse shapes in first–passage percolation and related growth models*. In "Perplexing Problems in Probability. Festshrift in Honor of Harry Kesten," eds. M. Bramson and R. Durrett, pgs. 121–142. Birkhäuser, 1999.

[36] J. Gravner, D. Griffeath, *Scaling laws for a class of critical cellular automaton growth rules*. In "Random Walks," eds. P. Révész and B. Tóth, pgs. 167–186. János Bolyai Mathematical Society, Budapest 1999.

[37] J. Gravner, D. Griffeath, *Random growth models with polygonal shapes*. Ann. Probab. 34 (2006), 181–218.

[38] J. Gravner, D. Griffeath, *Modeling snow crystal growth. I. Rigorous results for Packard's digital snowflakes*. Experiment. Math. 15 (2006), 421–444.

[39] J. Gravner, A. E. Holroyd, *Slow Convergence in Bootstrap Percolation*. To appear in Ann. Appl. Probab.

[40] J. Gravner, N. Mastronarde, *Shapes in deterministic and random growth models*, in preparation.

[41] J. Gravner, E. McDonald, *Bootstrap percolation in a polluted environment*, J. Stat. Phys. 87 (1997), 915–927.

[42] J. Gravner, C. Tracy, H. Widom, *A growth model in a random environment.* Ann. Probab. 30 (2002), 1340–1368.

[43] J. Greenberg, S. Hastings, *Spatial patterns for discrete models of diffusion in excitable media.* SIAM J. Appl. Math. 4 (1978), 515–523.

[44] D. Griffeath, *Self-organization of random cellular automata: four snapshots.* In "Probability and phase transition," ed. G. Grimmett, pgs. 49–67. Kluwer, 1994.

[45] D. Griffeath, D. Hickerson, *A two-dimensional cellular automaton with irrational density.* In "New Constructions in Cellular Automata," eds. D. Griffeath and C. Moore, pgs. 119–159. Oxford University Press, 2003.

[46] D. Griffeath, C. Moore, *Life without Death is* **P**-*complete*, Complex Systems 10 (1996), 437–447.

[47] A. E. Holroyd, *Sharp metastability threshold for two-dimensional bootstrap percolation*, Probab. Theory Related Fields 125 (2003), 195–224.

[48] A. E. Holroyd, *The metastability threshold for modified bootstrap percolation in d dimensions*, Electron. J. Probab. 11 (2006), 418–433.

[49] A. E. Holroyd, T. M. Liggett, D. Romik, *Integrals, partitions, and cellular automata.* Trans. Amer. Math. Soc. 356 (2004), 3349–3368.

[50] E. Jen, *Exact solvability and quasiperiodicity of one-dimensional cellular automata*, Nonlinearity 4 (1991), 251–276.

[51] L. B. Kier, P. G. Seybold, C.-K. Cheng, "Cellular Automata Modeling of Chemical Systems." Springer, 2005.

[52] K. Lindgren, M. G. Nordahl, *Evolutionary dynamics of spatial games*, Physica D 75 (1994), 292–309.

[53] P. Meakin, "Fractals, Scaling and Growth Far from Equilibrium," Cambridge University Press, 1998.

[54] N. H. Packard, *Lattice models for solidification and aggregation*, Institute for Advanced Study preprint, 1984. Reprinted in "Theory and Application of Cellular Automata," ed. S. Wolfram, pgs. 305–310. World Scientific, 1986. pp. 305–310.

[55] N. H. Packard, S. Wolfram, *Two-dimensional cellular automata*, J. Statist. Phys. 38 (1985), 901–946.

[56] A. Pimpinelli, J. Villain,"Physics of Crystal Growth," Cambridge University Press, 1999.

[57] R. H. Schonmann, *Finite size scaling behavior of a biased majority rule cellular automaton*, Phys. A 167 (1990), 619–627.

[58] R. H. Schonmann, *On the behavior of some cellular automata related to bootstrap percolation*, Ann. Probab. 20 (1992), 174–193.

[59] M. Song, *Geometric evolutions driven by threshold dynamics*, Interfaces Free Bound. 7 (2005), 303–318.

[60] T. Toffoli, N. Margolus, "Cellular Automata Machines." MIT Press, 1997.

[61] A. C. D. van Enter, *Proof of Straley's argument for bootstrap percolation*, J. Statist. Phys. 48 (1987), 943–945.

[62] A. C. D. van Enter, T. Hulshof, *Finite-size effects for anisotropic bootstrap percolation: logarithmic corrections*, J. Statist. Phys. 128 (2007), 1383–1389.

[63] G. Y. Vichniac, *Simulating physics with cellular automata*, Physica D 10 (1984), 96–116.

[64] G. Y. Vichniac, *Cellular automata models of disorder and organization*. In "Disordered Systems and Biological Organization," eds. E. Bienenstock, F. Fogelman-Soulie, and G. Weisbuch, pgs. 1–20. Springer-Verlag, 1986.

[65] S. J. Willson, *On convergence of configurations*. Discrete Math. 23 (1978), 279–300.

[66] S. J. Willson, *Cellular automata can generate fractals*, Discrete Appl. Math. 8 (1984), 91–99.

[67] S. J. Willson, *Computing fractal dimensions for additive cellular automata,* Physica D (1987), 190-206.

[68] M. Wójtowicz, *Mirek's Cellebration: a 1D and 2D Cellular Automata explorer*. http://www.mirwoj.opus.chelm.pl/ca/

[69] N. Wiener, A. Rosenblueth, *The mathematical foundation of the problem of conduction of impulses in a network of connected excitable elements, specifically in cardiac muscle*, Arch. Inst. Cardiol. Mexico 16 (1946), 205–265.

Books and Reviews for Further Information

[70] A. Adamatzky, "Identification of Cellular Automata." Taylor & Francis, 1995.

[71] J.-P. Allouche, M. Courbage, J. Kung, G. Skordev, *Cellular automata*, in Encyclopedia of Physical Science and Technology, Third Edition, Vol. 2, pgs. 555–567. Academic Press, 2001.

[72] J.-P. Allouche, M. Courbage, G. Skordev, *Notes on cellular automata,* Cubo, Matemática Educacional 3 (2001), 213–244.

[73] R. Durrett, "Lecture Notes on particle Systems and Percolation." Wadsworth & Brooks/Cole, 1988.

[74] R. Durrett, *Stochastic spatial models*, SIAM Rev. 41 (1999), 677–718.

[75] A. E. Holroyd, *Astonishing cellular automata*, Bulletin du Centre de Recherches Mathématiques 10 (2007), 10–13.

[76] A. Ilachinsky, "Cellular Automata: A Discrete Universe." World Scientific, 2001.

[77] T. M. Liggett, "Interacting Particle Systems." Springer-Verlag, 1985.

[78] T. M. Liggett, "Stochastic Interacting Systems: contact, voter and exclusion processes." Springer-Verlag, 1999.

[79] A. Toom, *Cellular automata with errors: problems for students of probability*, in "Topics in Contemporary Probability and Its Applications," ed. J. L. Snell, pgs. 117–157. CRC Press, 1995.

[80] D. H. Rothman, S. Zaleski, "Lattice-Gas Cellular Automata." Cambridge University Press, 1997.