
Threshold-range scaling of excitable cellular automata

ROBERT FISCH,¹ JANKO GRAVNER² and DAVID GRIFFEATH²

¹Department of Mathematics, Colby College, Waterville, ME 04901, USA

²Department of Mathematics, University of Wisconsin, Madison, WI 53706, USA

Received December 1990 and accepted May 1991

Each cell of a two-dimensional lattice is painted one of κ colors, arranged in a ‘color wheel’. The colors advance (k to $k + 1 \bmod \kappa$) either automatically or by contact with at least a threshold number of successor colors in a prescribed local neighborhood. Discrete-time parallel systems of this sort in which color 0 updates by contact and the rest update automatically are called Greenberg–Hastings (GH) rules. A system in which all colors update by contact is called a cyclic cellular automaton (CCA). Started from appropriate initial conditions, these models generate periodic traveling waves. Started from random configurations the same rules exhibit complex self-organization, typically characterized by nucleation of locally periodic ‘ram’s horns’ or spirals. Corresponding random processes give rise to a variety of ‘forest fire’ equilibria that display large-scale stochastic wave fronts. This paper describes a framework, theoretically based, but relying on extensive interactive computer graphics experimentation, for investigation of the complex dynamics shared by excitable media in a broad spectrum of scientific contexts. By focusing on simple mathematical prototypes we hope to obtain a better understanding of the basic organizational principles underlying spatially distributed oscillating systems.

Keywords: Cellular automation, excitable medium, locally periodic, spiral, stable periodic object, Greenberg–Hastings model, cyclic CA, self-organization, phase transition, turbulence, bug, macaroni

1. The rules

A cellular automaton (CA) is a dynamic configuration on a lattice of sites that updates in parallel according to a local homogeneous deterministic rule (cf. Toffoli and Margolus, 1987). This paper studies two families of *excitable* cellular automata that we call *Greenberg–Hastings* (GH) models and *cyclic cellular automata* (CCA). Each family is indexed by three positive integer-valued parameters: the range, ρ , of interaction; the threshold, θ , of sites needed for excitation; and the number κ of available colors. We are interested in these CA rules because they emulate the behavior of a wide range of complex, coherent, periodic wave phenomena in space. Color Plates A–H illustrate some of the patterns generated by GH and CCA dynamics for different choices of the parameters ρ , θ and κ .

We will focus almost exclusively on the description of excitable CA rules in two dimensions, although many features of our analysis apply more generally. The most

expedient planar lattice from a mathematical point of view is \mathbb{Z}^2 , the two-dimensional integers. For graphical visualization, on the other hand, one uses a large finite lattice $\{0, 1, \dots, L - 1\}^2$, typically with wrap-around at the boundary. An especially intriguing aspect of our study is the interplay between computer experimentation and deductive reasoning; without such ‘symbiosis’, many subtle aspects of these complex systems would elude our understanding.

Let us now explain GH and CCA rules in detail. Each cell of the lattice (or pixel on the graphics screen) is painted with one of κ colors, arranged in a ‘color wheel’, and labeled $0, 1, \dots, \kappa - 1$. Colors can only advance in one direction around the wheel: $k \rightarrow k + 1 \bmod \kappa$. We assume throughout this paper that $\kappa \geq 3$, so ‘positive’ and ‘negative’ directions are distinguishable, although two-color variants also exhibit interesting behavior. Depending on the rule, colors either advance automatically or by *contact* during each discrete time unit. Advance by contact of color k at site x means that $k \rightarrow k + 1 \bmod \kappa$ if and only

if at least θ neighboring sites within range ρ of x have value $k + 1 \bmod \kappa$. Convenient interpretations of ‘within range ρ ’ use the *diamond* (l^1) or *box* (l^∞) neighbor set of radius ρ . For instance, $\rho = 1D$ yields the usual nearest neighbors: N, S, E, W; whereas $\rho = 1B$ adds four more neighbors: NE, NW, SE, SW. We call systems ξ_t in which color 0 updates by contact while all others advance automatically *Greenberg–Hastings* models, in honor of the authors of influential papers (Greenberg *et al.*, 1978; Greenberg and Hastings, 1978) that studied the case $\rho = 1D$, $\theta = 1$, $\kappa = 3$. Corresponding systems ζ_t in which *all* colors update by contact are called *cyclic cellular automata*.

More formally, let $\|x\|_p$ denote the l^p -norm of $x = (x_1, x_2) \in \mathbb{Z}^2$, $p \in [1, \infty]$, and introduce the neighborhoods $\mathcal{N}_\rho(x) = \mathcal{N}_\rho^p(x) = \{y: \|y - x\|_p \leq \rho\}$. For given parameters ρ , θ and κ , the update algorithms on $\{0, \dots, \kappa - 1\}$ -valued configurations are as follows:

$$\xi_{t+1}(x) \begin{cases} = (\xi_t(x) + 1) \bmod \kappa & \text{if } \xi_t(x) \geq 1 \text{ or} \\ & \# \{y \in \mathcal{N}_\rho(x): \\ & \xi_t(y) = 1\} \geq \theta \\ = \xi_t(x) & \text{otherwise} \end{cases} \quad (\text{GH})$$

$$\zeta_{t+1}(x) \begin{cases} = (\zeta_t(x) + 1) \bmod \kappa & \text{if } \# \{y \in \mathcal{N}_\rho(x): \\ & \zeta_t(y) = (\zeta_t(x) + 1) \bmod \kappa\} \\ & \geq \theta \\ = \zeta_t(x) & \text{otherwise} \end{cases} \quad (\text{CCA})$$

where $\#\Lambda$ is the cardinality of set Λ .

We are particularly interested in the evolution of excitable CA rules from disordered initial states, so our processes typically start in a *random* configuration that independently paints each site with a random color (probability $1/\kappa$ for each color). This paper has two primary objectives: to describe the rich ‘terrain’ of large-scale self-organization exhibited by excitable CAs depending on parameters ρ , θ and κ ; and to begin to identify mathematical principles underlying the evolution of GH and CCA dynamics from ‘chaos’ to ‘order’. We hope that this largely empirical investigation will help lay the foundation for a mathematical theory of excitable media with substantive rigorous underpinnings.

2. Basic phenomenology

How do GH and CCA dynamics evolve? Assume for simplicity that $p = \infty$, that is, the rules use box neighborhoods. Consider first an initial configuration comprised of contiguous horizontal ‘bands’ of colors $1, \dots, \kappa - 1$ from top to bottom, on a background of color 0. Let each band have width ρ in the vertical direction and stretch indefin-

itely in the horizontal direction (or wrap around the computer screen). It is easy to verify that such a band, subject to either update scheme, generates an upward-traveling wave if the threshold θ is sufficiently small. But the band *dies out* (vanishes) in GH or *fixates* (stops changing) in CCA when θ is large. Readers may want to accept these assertions on faith for now—they will be easier to understand in the context of Section 6—but precise statements are as follows. In the GH case, if $\theta \leq (2\rho + 1)[(\rho + 1)/2]$, then a wave propagates upward; each band in the wave has width $w = \rho + 1 - [\theta/(2\rho + 1)]$. If $\theta > (2\rho + 1)[(\rho + 1)/2]$, then the configuration dies out. In the CCA case, if $\theta \leq (2\rho + 1)\rho$, then a wave propagates upward; each band in the wave has width $w = \rho$. If $\theta > (2\rho + 1)\rho$, then the initial configuration never changes. (Here $[i]$ denotes the greatest integer less than or equal to i .) Surely waves can propagate through the lattice in any direction provided θ is suitably small. The horizontal and vertical orientations simply make computation of the cutoff easiest because they are especially suited to the box neighborhood.

One might suspect, then, that an initially disordered system tries to self-organize into spatially-distributed waves. But do the dynamics have enough ‘energy’ to organize, and, if so, what happens when waves moving in different directions collide? The same topological principle that guarantees a cowlick on a hairy billiard ball suggests the formation of vortices of some kind. Such vortices may then give rise to stable structures that serve as ‘centers’ for coherent growing structures. For instance, GH and CCA rules with $\rho = 3B$ and $\theta = 6$ support stable periodic spiral pairs, sometimes called *ram’s horns*, on a background of color 0. Figure 1, a gray-scale rendition of the eight-color case, illustrates the ability of ram’s horns to reproduce themselves exactly while generating concentric periodic waves.

It is tempting at this point to speculate on the role of parameters ρ , θ and κ . Larger range ρ involves a kind of averaging over lattice sites, so the corresponding wave shapes should become smoother, perhaps approaching some continuum limit. For fixed ρ , as we have seen, the

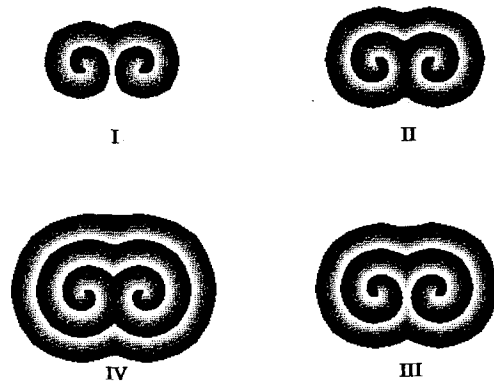


Fig. 1. *Ram’s horn dynamics* ($\rho = 3B$, $\theta = 6$, $\kappa = 8$)

threshold θ regulates the spread of wave fronts. Moreover, θ influences the ability of the system to remain active and organize out of initial disorder. The number of colors κ clearly dictates the period of waves, and also plays a key role in the evolution from randomness since only neighbors with adjacent colors promote updating by contact. Before beginning a detailed analysis of GH and CCA dynamics that will illuminate the parameter dependence, let us briefly review the existing literature on CA models of excitable media, and mention some examples of real-world spatial phenomena that have served as a primary motivation for this kind of modeling.

3. Background and connections

Spatially distributed periodic waves occur in scientific contexts as diverse as: oscillating chemical reactions, such as the Belousov–Zhabotinski reaction (Winfree, 1974), and carbon monoxide oxidation (Gerhardt and Schuster, 1989); atrial fibrillation (Winfree, 1987); the spread of epidemics (Murray *et al.*, 1986); electrical transmission in neural networks (Hodgkin and Huxley, 1952); self-organization of slime mold (Tomchik and Devreotes, 1981); algorithms for causal coherence in parallel computation (Bennett *et al.*, 1990; Toffoli and Margolus, 1987, pp. 90–95); and a theory of spiral galaxy formation (Freedman and Madore, 1983). Kapral (1991) has written a survey of excitable media and discrete reactions that includes about 200 references to journal articles and books in biology, computer science, chemistry, physics and physiology. Much of this literature is devoted to computer modeling of the listed phenomena, using CAs or more complicated coupled lattice maps. The methodology is overwhelmingly empirical; few findings are mathematically rigorous. CAs were pioneered by von Neumann and Ulam in the late 1940s (Toffoli and Margolus, 1987). At about the same time, Wiener and Rosenbluth (1946) analyzed discrete dynamics for excitable heart muscle. CA dynamics of the type we call GH, with $\rho = 1D$, $\theta = 1$, and $\kappa = 3$, were formalized and studied by Greenberg *et al.* (1987) and Greenberg and Hastings (1987). Qualitatively similar CA rules with more general threshold and range have been considered recently by Toffoli and Margolus (1987, pp. 82–84), Gerhardt *et al.* (1990) (for a popular account, see Dewdney, 1988), and Markus *et al.* (1991). We should also mention that there is widespread current interest in self-organization as a guiding principle of biological evolution; see Waldrop (1990) for a lively report on an ‘artificial life’ workshop.

Our research on excitable CAs evolved from the study of closely related processes with random dynamics. Bramson and Griffeath (1980) considered a continuous-time asynchronous Markov process on \mathbb{Z} similar to the one-dimensional CCA with $\rho = 1$, $\theta = 1$, $\kappa \geq 3$. They proved

clustering for $\kappa \leq 4$ and *fixation* for $\kappa \geq 5$. In that context, clustering means that every site changes color infinitely often but that any finite subset of the lattice is overwhelmingly likely to be all one color eventually; fixation means that every site is painted a final color with probability one. Fisch (1990; 1991) then obtained an analogous result for the CCA on \mathbb{Z} and calculated exact asymptotics concerning the rate of clustering for $\kappa = 3$. The behavior of random cyclic systems in higher dimensions remained a mystery until Fisch and Griffeath implemented a real-time simulation on the Cellular Automation Machine (CAM). An expository article (Griffeath, 1988) describes the computer-aided discovery of an exotic evolution on \mathbb{Z}^2 characterized by nucleation of wave droplets that interact to form a remarkably stable spiral-laden equilibrium. It came as a surprise that this scenario apparently occurs for *any* number of colors $\kappa \geq 3$.

Unfortunately, mathematical results for the random cyclic system in two or more dimensions seem very hard to come by. So Fisch *et al.* (1991a) turned to the basic CCA on \mathbb{Z}^2 ($\rho = 1D$, $\theta = 1$), hoping that deterministic dynamics would simplify matters. Indeed, this process ζ_t enjoys a number of regularity properties, essentially topological, that lead to a simple proof of local periodicity for any κ starting from a random initial state. We call a cellular automaton ζ_t *locally periodic* (LP) if

$$P((\zeta_t(x); t = 1, 2, \dots) \text{ is eventually periodic}) = 1, \quad \text{for each } x \in \mathbb{Z}^2 \quad (1a)$$

with the understanding that some sites x may fixate, *and*

$$\inf_{t, x \neq y, i, j} P(\zeta_t(x) = i, \zeta_t(y) = j) > 0 \quad (1b)$$

In the basic CCA, the existence of certain stable configurations on loops, known as *clocks*, implies that each x changes color at every update from some time on. Hence all sites have period κ (see Fisch *et al.*, 1991a; or Dewdney, 1981). Despite the contrast between statistical equilibrium in the random cyclic system and deterministic limiting behavior in ζ_t , on the computer screen both dynamics display nucleation of wave droplets and dynamic formation of large spirals. (Griffeath (1988) contains pictures of a random evolution; Plate E of this paper shows self-organization in a basic CCA.) Thus it makes sense to view the random cyclic system as a perturbation of the corresponding cellular automaton.

About a year ago, Rick Durrett (private communication) pointed out to us that the basic Greenberg–Hastings model ($\rho = 1D$, $\theta = 1$, $\kappa = 3$) enjoys the same regularity properties as the basic CCA but has additional simplifying features that make it even more amenable to mathematical analysis. (See Durrett and Steif, 1991, for several interesting rigorous results in one and two dimensions.) Durrett also noticed that the corresponding GH rules with larger numbers of colors display nucleation and spiral formula-

tion similar to that of the CCA. Moreover, he identified stable structures in the $\rho = 1$, $\theta = 2$, $\kappa = 3$ GH model that seemed to generalize the notion of a clock to higher threshold, functioning as centers for self-organization. His observations prompted our attempts to incorporate GH and CCA dynamics into a unified framework, and were the direct inspiration for our investigation of rules with higher threshold and range.

While carrying out the CA experiments described in this paper, we also simulated some random GH models analogous to the random cyclic ones that we had encountered first of all. Not surprisingly, simulations again showed nucleation leading to an equilibrium of coherent large-scale wave fronts. In this case, the steady-state indicates a close connection to the epidemic with regrowth, or ‘forest-fire model’, that is studied in a beautiful paper by Durrett and Neuhauser (1991). Evidently certain fundamental multitype particle systems may also be viewed as perturbations of excitable CAs. Consequently our interest in GH and CCA dynamics is fueled not only by their intrinsic significance but also by their potential to shed light on self-organization in random interactions.

4. Ergodic classification

Perhaps the most basic problem concerning GH and CCA dynamics is the classification of ergodic behavior, depending on ρ , θ and κ . One of the main objectives of this paper is to identify several qualitatively distinct possible limiting ‘phases’, as evidenced by Plates A–H, and suggest how the three parameters determine the phase of a given rule.

First let us mention a simple extension of the main argument in Fisch *et al.* (1991a). Namely, if $\theta = 1$, then for any ρ and $\kappa \geq 3$ the excitable systems of both types on \mathbb{Z}^2 are *locally periodic* with period κ . In other words, every site updates every time eventually, but sites sufficiently far apart are typically out of phase. The gist of the proof goes as follows. Somewhere in the random initial state there is a *clock*: a loop of sites on which all κ colors are arranged cyclically (perhaps cycling more than once). Since $\theta = 1$, the color at every site of the clock advances every time, regardless of the configuration at sites that do not belong to the clock. Thus the set \mathcal{A} of sites that eventually change color every time is nonempty. Now suppose that $x \notin \mathcal{A}$, $y \in \mathcal{A}$, and $\|y - x\|_\infty \leq \rho$. It follows that the color at y exceeds the color at x by 1 (mod κ) at some time t . But then, again because $\theta = 1$, y forces x to advance every time after t , a contradiction. Therefore $\mathcal{A} = \mathbb{Z}^2$ as claimed. With at least a certain positive probability that depends only on κ , distinct sites x and y are ‘slaved’ to clocks that are out of phase, so property 1b holds as well.

To see that not all GH and CCA rules on \mathbb{Z}^2 are locally periodic, let us now demonstrate different behavior when the threshold θ is sufficiently large. For simplicity consider

box (l^∞) neighborhoods. First, for GH dynamics, we claim that

$$\text{if } \theta > 2\rho^2 + 5\rho + 1, \text{ then } \xi_t \text{ dies out strongly} \quad (2)$$

(that is, the color at every site x is eventually 0 with probability one). The key observation is that, in this (ρ, θ) regime, a large lattice *disk* of 0s expands deterministically with at least a certain positive speed. To see this, set

$$\begin{aligned} D_R &= \{y \in \mathbb{Z}^2: \|y\|_2 \leq R\} \\ \partial D_R &= \{x \in \mathbb{Z}^2: R < \|x\|_2 \leq R + 1\} \end{aligned} \quad (3)$$

If R is sufficiently large ($R > 30\rho^2$, say), then in the vicinity of $x \in \partial D_R$ the edge of D_R is almost linear. As a consequence at least $(\rho - 1)(2\rho + 1)$ neighbors of x must lie within D_R ; for a proof of this fact see the Appendix. Even if all remaining neighbors of y are excited ($= 1$), their total is at most $(2\rho + 1)^2 - (\rho - 1)(2\rho + 1) - 1 = 2\rho^2 + 5\rho + 1$, which is below threshold. The same principle clearly applies to $x \in D_R$. Hence a 0 at $x \in D_{R+1}$ cannot possibly change to 1 if D_R has all 0s. Therefore, starting from some translate of D_R that has all 0s in the initial random configuration, with $R > 30\rho^2$, successive disks are deterministically covered by 0s within at most κ updates. We conclude that $P(\xi_t(x) = 0 \text{ eventually in } t) = 1$ for each x , as claimed.

Turning to the CCA ζ_t , a variation on the same argument shows that

$$\text{if } \theta > 2\rho^2 + 5\rho + 1 \text{ and } \kappa > 2\rho(\rho + 1), \text{ then } \zeta_t \text{ fixates} \quad (4)$$

(that is, every site x has a final color with probability one). As before, an initial configuration of all 0s on a sufficiently large D_{R_0} cannot possibly change. Moreover, for any R_0 , $x \in \partial D_R$ has at most $2\rho(\rho + 1)$ neighbors within D_R . Again, see the Appendix for a proof. Suppose that D_R fixates at some time t . Then, by the hypothesis on κ , some color k is missing from $\mathcal{N}_x \cap D_R$ at all times after t . Even if all remaining neighbors of x have color k at some later time, x cannot make a transition to color k for the same reason as in the GH case. Hence D_{R+1} must also fixate eventually. By induction, any $x \in \mathbb{Z}^2$ fixates with probability one.

Thus we see that GH and CCA models each exhibit at least two distinct ergodic behaviors, depending on ρ , θ and κ . But are there other possibilities? And is it possible to distinguish different kinds of local periodicity? In light of the complex self-organization characteristic of excitable media, it makes sense to enlist the aid of the computer as we attempt to answer these fundamental questions.

5. Modeling environments

Two computer graphics facilities were used for the extensive CA experimentation that will be described throughout the remainder of this paper. Before proceeding, let us describe them briefly.

5.1. EXCITE!

We have written interactive experimental software (Fisch and Griffeath, 1991) expressly for the study of GH and CCA rules, and their randomized counterparts, with $\rho \leq 10$. The program, called EXCITE!, runs on IBM-compatible PCs equipped with either EGA or preferably VGA color graphics. Plates A–H and all the figures in this paper were generated by our software. Since parallel updating of large arrays is very computation-intensive we recommend a 386 machine, but comparable images can be ‘grown from seed’ (for example, overnight) on slower computers. EXCITE! is being distributed as freeware, available by request from David Griffeath, Dept of Mathematics, University of Wisconsin, Madison, WI 53706, USA (griffeat@math.wisc.edu). Anyone who wishes to confirm findings reported here or experiment on their own is encouraged to request a copy. We have tried to make this paper readable without help from a PC, but strongly urge readers to try at least a few experiments. The distribution disk contains two slightly different executable versions of the main program, a ‘greatest hits’ file that identifies many of the most interesting parameter choices, a collection of windows designed for ‘planting’ in certain rules, the ram’s horn picture that generates Fig. 1, and a READ.ME file of documentation.

Our software is designed so that choice of rules and interaction with the dynamics takes place on a control panel text screen. A key press switches between the control panel and graphics screen. By choosing from menus and entering data, users are able to:

- (1) Pick GH or CCA, either random or deterministic, and specify parameters ρ (D or B), θ and κ ;
- (2) Specify an array size from 25×25 to 640×480 ;
- (3) Set the boundary condition to wrap-around, free, or all 1s;
- (4) Run a dynamic either until a key is pressed or until a predetermined final time;
- (5) Save the current time and configuration to disk (as a .PIC file);
- (6) Initialize the dynamic from a random or ‘special’ configuration, solid color, or pre-saved .PIC file;
- (7) Cut or paste a central window in a configuration, up to size 23×23 (as a .WND file);
- (8) Zoom the current window to display numerical values and edit these manually;
- (9) Test for the period of the process at prescribed sites;
- (10) Shift the configuration horizontally or vertically by any offset;
- (11) Change the update probability in randomized systems;
- (12) Insert or delete rules in the greatest hits file and edit accompanying descriptions.

Several more menu options are also available. We began writing EXCITE! in order to carry out the experiments

described in this paper. By now the program has evolved into a rather flexible modeling environment that we hope will provide a useful template for other developers of complex systems software.

5.2. Three CAMs plus wires

The Cellular Automaton Machine (CAM) is a pseudo-parallel processor invented by Toffoli and Margolus (1987) and implemented as a plug-in board for the IBM PC. CAM has the ability to update nearest neighbor CA rules on a 256×256 array at speeds up to 60 updates per second. If the rule is simple enough (for instance, $\rho = 1D$, $\kappa \leq 4$) this remarkable device can display ‘real-time’ dynamics, often yielding insights into system behavior that are difficult to glean from still frames. Unfortunately most of the rich phenomenology exhibited by GH and CCA rules exceeds the limitations of a single CAM. However Toffoli and Margolus wisely designed their hardware so that several CAMs could be connected in parallel, and included a ‘user connector’ that allows dedicated custom devices to be interfaced as well. By piggybacking three CAMs and designing an external circuit board to handle most of the CA logic, we have constructed a device that computes GH and CCA rules for box neighborhoods with range $\rho \leq 4$, thresholds $\theta \leq 15$, and $\kappa \leq 16$ colors. Performance deteriorates as ρ increases, but a 256×256 array typically updates several times per second. We have also devised a ‘scooping’ procedure that uses the CAMs to drive arrays of *arbitrary* size, subject only to memory constraints of the host computer. We expect that the CAMs will be particularly useful as we begin to study subtle varieties of self-organization which take place on very large length scales.

6. Stable periodic objects

To motivate a key concept in the analysis of excitable CA rules, we now describe our first experiment (which the reader can perform either using EXCITE! or by writing a little computer program). Consider the GH rule with parameters $\rho = 2B$, $\theta = 3$, $\kappa = 8$. Start from a (pseudo-) random array of eight colors over a lattice of size 100×100 , say, and let the dynamic evolve. Within the first few time steps the system will almost die out. But, except on a bad day, a few isolated clusters of excited states (= 1s) will persist. These rare pockets of activity generate ram’s horns that emit concentric cyclic waves. The waves spread out until they cover the entire lattice. Wave fronts emanating from different centers annihilate upon collision, and a final locally periodic pattern is formed. Plate A shows an intermediate stage in the evolution of such an experiment on a 512×400 grid.

Upon closer inspection of a ram’s horns’ structure, we find that the spiral pairs have centers comprised of period-8 cyclic patterns. For instance, in one of our experiments

collide with something:

$$\begin{matrix} 3 & 2 & 1 \\ 3 & 2 & 1 \end{matrix} \quad (0\text{'s}) \quad (8a)$$

Elsewhere clocks emerge; for instance:

$$\begin{matrix} 0 & 2 \\ 1 & 3 & 1 & 3 \\ 2 & 2 & 0 & 0 \\ 1 & 3 & 1 & 3 \\ 0 & 2 \end{matrix} \quad (8b)$$

Other more complicated growing structures abound. But in the end, every site would appear to cycle with period κ . As soon as there are five or more colors, the only surviving structures visible on a graphics array of any real-world size are bugs analogous to structure 8a. These bugs wander about in the four horizontal and vertical directions until they smash into one another, an event that usually but not always has dire consequences. One would guess that the system dies out slowly. But rather incredibly, after several months of false impressions, D. Pritikin (private communication) discovered the SPO in Fig. 3. EXCITE! comes with a ready-made window of Pritikin's bow tie. Plant it and watch it grow!

The last few examples dramatically illustrate the limits of computer visualization when dealing with excitable automata; several more stories with the same moral will follow. They also suggest that a rich set of possible ergodic behaviors may arise for intermediate threshold values. Spirals are the rule when θ is small, and for large θ the systems grind to a halt, but in-between and for relatively small κ there may exist SPOs that are extremely rare statistically, with an architecture based on complex and exotic organizational principles. As further evidence of this scenario, Pritikin has now produced a 'butterfly' SPO for the $\kappa = 6$ case of $\rho = 1B$, $\theta = 2$. His example, also on the EXCITE! disk, has more than 200 sites; we invite the reader to find it. Do SPOs exist for *all* κ in this simple case of eight nearest neighbors and threshold 2? We suspect not, but without any compelling rationale. If our hunch is right, then what is the highest number of possible colors?

The next section will suggest that a phase transition in the ergodic behavior of GH and CCA rules is governed by




Fig. 3. Pritikin's bow tie: a stable periodic object with complex architecture

the question of existence of SPOs. Even without this compelling connection to the dynamics, we find stable periodic objects to be intriguing combinatorial structures in their own right.

7. Engineering

Computer-generated dynamics from random initial states evidently do not always tell the whole story, so for additional insight we carry out various experiments specifically designed to display aspects of wave propagation. Let us first describe the *band test* for GH rules that is a 'special' initialization option in EXCITE! The starting state is an adjacent arrangement of horizontal bands, each of width $w \geq \rho$, with 1s at the top and $(\kappa - 1)$ s at the bottom, on a background of 0s. This is like the traveling wave described in Section 2, except that bands have a *finite* (horizontal) length l of our choosing. The evolution until $t = 20$ for the 'macaroni' rule of Fig. 2 is shown in Fig. 4 ($w = 6$, $l = 25$).

The first thing to notice is that the *ends* of the wave bend inwards, albeit slowly in this instance. They are trying to make ram's horns. One can imagine that if the threshold were lower—say, $\theta = 3$ —then the evolution would produce a tight spiral pair with SPOs at its two centers. This is indeed the case: the band test gives rise to stable centers for octagonal waves like those formed from a random initial state in Plate A. The exactly reproducing structure I of Fig. 1 was also obtained by means of a band test. However the excitation in Fig. 4 does not manage to reproduce: its ends collide prematurely, yielding only a bug that does not grow and escapes with the same velocity as the ring. To make a structure that generates periodic waves on a background of 0s there is a simple remedy: increase the length l of the band test. (In this case one could also reduce the initial widths w to 2.) However, if the resulting object has a period exceeding κ , due to wide bands of 0s, then it might be overrun by metastable macaroni (the environment of Plate B). The strongest case

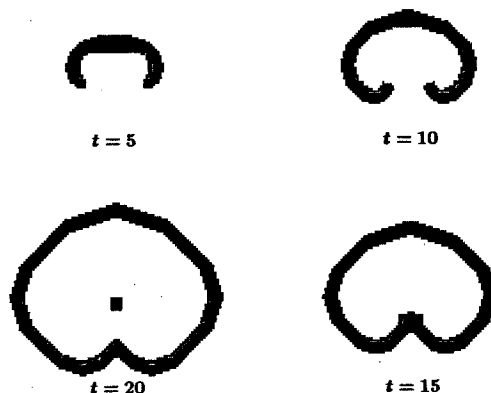


Fig. 4. A band test (GH, $\rho = 2B$, $\theta = 4$, $\kappa = 5$)

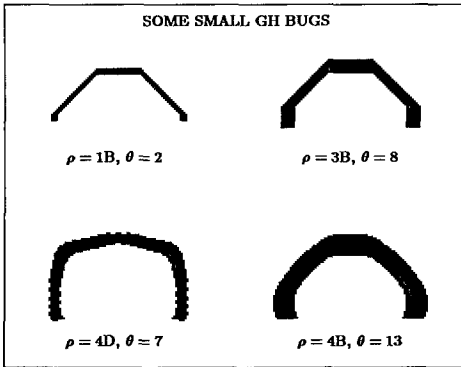


Fig. 5. Some small GH bugs

for local periodicity can be made if we manage to engineer a bona fide SPO. We will return to this issue later.

For a given range ρ , how does the band test behave as the threshold θ increases? Beyond a certain cutoff, wave fragments are unable to bend completely inward, and so cannot generate cyclic rings. Just above this critical value, for all small neighborhoods except 1D, there are two possibilities. In some cases wave ends are vulnerable, so the length of any band gradually decreases to 0. But in other cases a sufficiently long wave piece manages to survive without making rings: instead the band test creates a viable bug. Throughout this discussion we will use the term *bug* rather generally to describe waves with permanent ends that survive under GH dynamics in an environment of all 0s. Figure 5 shows examples of horizontally and vertically moving bugs for some rules with small parameter values. We have encountered bugs moving at 45° to the horizontal in some cases; perhaps other orientations are possible as well. It is tempting to make a distinction between structures that simply move along with a constant or periodic shape like (8a) and those that grow in size. Informally, we sometimes refer to the latter kind of bugs as *worms*.

Even if bugs with ends die, it may still be possible for suitably shaped rings of excitation to propagate. The simplest example of this rather subtle phenomenon has parameters $\rho = 3D$, $\theta = 4$. No band test survives for these values. However, if we start from all 1s on the lattice octagon $\mathcal{N}_5^1(0) \cap \mathcal{N}_4^\infty(0)$, then a stable octagonal ring spreads outward. Intriguing and delicate nonmonotonicity comes into play in this situation: excitation dies out starting from all 1s on *either* of the larger sets $\mathcal{N}_5^1(0)$ or $\mathcal{N}_4^\infty(0)$. Roughly speaking, the corners of a diamond or box are too vulnerable, so started from those configurations the ring breaks into four disjoint, short-lived fragments.

Increasing the threshold still further, we know from Section 2 that, for box neighborhoods, even infinite horizontal bands die out as soon as $\theta > (2\rho + 1)[(\rho + 1)/2]$. In this regime wave fronts self-destruct globally, so rings of any size or shape die out. A final transition occurs for

$\theta > \rho^2$ in the diamond case and $\theta > \rho(2\rho + 1)$ in the box case. The reader can easily verify that once the threshold is this large, horizontal and vertical wave fronts cannot advance at all. Thus the dynamics are *convex-confined*, and an isolated box of excitation disappears in one time step.

Let us conclude our overview of band tests with some remarks about the role of κ . As long as one is studying bug or ring propagation, the number of colors is not particularly important: directed spread of leading 1s onto a background of 0s is unaffected by trailing colors. (There is also no reason to start from a band width w greater than ρ .) The only exception to this principle is $\kappa = 3$, in which case ‘feedback’ excitation sometimes occurs behind a bug or within a ring. The feedback can cause bugs to stabilize as periodic structures similar to ram’s horns, and reinforce rings with internal ‘snowflake’ patterns. Particularly amusing examples of these effects, for $\rho = 4B$ and $\kappa = 3$, are exhibited by a band test when $\theta = 13$, and by an initial box $\mathcal{N}_3^\infty(0)$ of 1s when $\theta = 14$. To avoid such complicating features we typically use band tests with $\kappa = 4$ for bugs and rings, or $\kappa = 8$ if a longer history of the wave front is desired. In cases where bands can bend completely the choice of κ is much more important: too few colors can lead to a periodic wave that contains wide expanses of 0s, whereas too many colors can interfere with the excitation as it attempts to fold in on itself.

We now turn to a parallel exploration of wave propagation in cyclic cellular automata. If θ is small enough that the band test makes ram’s horns for GH, then essentially the same behavior occurs for CCA, that is, stable spiral pairs arise. However, for threshold values that produce bugs or die out under GH, the corresponding CCA fixates from the band test. Another initial condition—more tailored to cyclic dynamics—consists of a periodic arrangement of large squares (on a background of color 0, say). For example, if $\kappa = 5$ we might try a configuration of $w \times w$ squares arranged thus:

1s	2s	3s	(0s)
	0s	4s	

Such a *loop test*, with $w = 4\rho$, is the default ‘special’ initialization for CCA rules in EXCITE! The rough idea is to induce a vortex as the colors chase each other around, and then see whether a stable center for a spiral is formed. Not surprisingly, the loop test makes a spiral with ease whenever the band test yields a spiral pair. But for many θ such that the band test fixates, we encounter an unexpected phenomenon: loop-test behavior depends in a fundamental way on whether $\kappa \in \{3, 4\}$, or $\kappa \geq 5$.

With three or four colors the loop test makes spirals as long as the rule is not convex-confined. These structures, symmetric in the colors, have a period *exceeding* κ . The


```

5 5 0 0 1 1 2 2 2 2
5 0 0 1 1 2 2 3 3 2
5 0 0 1 1 2 3 3 3 3
5 5 0 3 1 1 0 4 4 3
4 5 5 4 3 0 4 4 4 4
4 4 4 4 0 3 4 5 5 4
3 4 4 0 1 1 3 0 5 5
3 3 3 3 2 1 1 0 0 5
2 3 3 2 2 1 1 0 0 5
2 2 2 2 1 1 0 0 5 5
    
```

Fig. 6. An organizing center with two strands (CCA, $\rho = 1B$, $\theta = 2$, $\kappa = 6$)

spirals are large if $\kappa = 4$ and enormous if $\kappa = 3$. Especially with three colors, it may be necessary to choose a very big w . Some nice examples are (ρ, θ, κ) values of $(1B, 3, 3)$, $(1B, 3, 4)$, $(3B, 9, 3)$, and $(3B, 8, 4)$.

In the same θ -regime, for $\kappa \geq 5$ the loop test tries to form a spiral but fails. No matter how big we make the initial squares, somehow there is not enough room at the center of the vortex to support a cohesive structure. After a few time steps, some pair of adjacent colors is pulled apart and the vortex self-destructs, leaving behind a rather curious assortment of debris. Try, for example, $(1B, 3, 5)$ with $w = 20$.

Let us end this section by discussing the fascinating CCA family with $\rho = 1B$, $\theta = 2$. Loop tests for $\kappa = 3$ ($w = 20$) and $\kappa = 4$ ($w = 4$) make large spirals. For $\kappa \geq 5$, as explained above, loop tests fail to produce periodic structures. But are there other spiral-like arrangements that might thrive in a chaotic or fixated environment? After some experimentation, we have discovered certain delicately arranged organizing centers that give rise to interwoven spirial strands. Figure 6 shows such a center, with two strands, for the six-color rule; the resulting dynamic is displayed in the title screen of *EXCITE!* Another arrangement with four strands is on the disk. Either pattern generates a creature that propagates with ease. Similar constructions have been carried out with $\kappa = 7, \dots, 12$; apparently there is at least one viable multistranded center for all $\kappa \geq 6$. An extensive search has failed to produce such a creature for the five-color rule, so we suspect that none exists. We would love to be proved wrong!

8. The data and phase portraits

Recall that our primary objective is to propose a classification of ergodic behavior in GH and CCA dynamics as a function of range, threshold, and number of colors. This section begins with a formulation of five cutoffs in θ for each fixed ρ , based on wave-front characteristics and existence of SPOs. Then we present experimental data enumerating each of the cutoff values for all diamond and box neighborhoods with $\rho \leq 6$. After that we describe the

behavior of excitable CA rules started from random configurations, as observed on large finite arrays using *EXCITE!* and 3 CAMs plus wires. We will see that the given cutoffs coincide with apparent phase transitions in the evolutions of infinite systems from initial noise. Finally, threshold-range scaling suggested by our data will be explained in terms of a limiting Euclidean wave dynamics as $\theta \rightarrow \infty$ and $\rho \rightarrow \infty$, with $\theta/\rho^2 \rightarrow c \in (0, 4)$.

First, for a given range ρ neighborhood (diamond or box), let us identify the following threshold cutoffs in increasing order. Since θ is integer-valued, each cutoff takes the form $n | n + 1$, meaning that one state of affairs occurs for $\theta \leq n$ while another occurs for $\theta \geq n + 1$.

- *bend*(ρ). Below *bend* the GH band test, with $\kappa = 4$, $w = \rho$, and l sufficiently large, makes stable periodic structures that generate waves forever; above *bend* the band is unable to fold in on itself completely, and so cannot reproduce indefinitely.

- *bug*(ρ). Below *bug* the GH band test, with $\kappa = 4$, $w = \rho$, and l sufficiently large, generates a wave front that survives forever; above *bug* the band has vulnerable ends and so dies.

- *ball*(ρ). Below *ball* the GH rule, with $\kappa = 4$, starting from all 1s on a suitable finite set of sites, produces a ring of excitation that survives forever; above *ball* any such ring is vulnerable somewhere and eventually breaks into fragments that die out. The set of 1s that produces a viable ring may need to be quite large and have the right shape.

- *spo₃*(ρ). Below *spo* the GH rule with $\kappa = 3$ has stable periodic objects; above *spo* it does not. (We defer until Section 10 a discussion of how this cutoff is evaluated.)

- *boot*(ρ). Below *boot* any lattice half-plane of excitation (1s) advances for at least one update; above *boot* excitable CA dynamics are convex-confined.

Table 1 shows empirically derived values of these five cutoffs for $\rho = nD$ and $\rho = nB$, $1 \leq n \leq 6$. We invite the reader to confirm our findings using *EXCITE!*

Having tabulated the cutoff values, we proceed to summarize their implications for the ergodic behavior of excitable CA rules started from *random* initial states. To keep matters as simple as possible, let us assume for now that the number of colors κ is fairly large. Just *how* large is not clear, but at the very least let $\kappa \geq 5$. We have already seen that three- and four-color rules exhibit various idiosyncratic phenomena; those cases will be addressed in Section 10. It would also simplify matters if one knew that SPOs like Pritikin's bow tie and butterfly did not exist. For given ρ and θ , we conjecture that such organizing centers with complex architecture cannot occur once κ is big enough. A rigorous result to this effect in the context of one-dimensional excitable systems has been obtained recently by Fisch and Gravner (to appear). It is quite clear that naive computer modeling will *never* address the existence of exotic SPOs. Rather, a theory based

Table 1. *Cutoffs*

ρ	Diamond	Box
	<i>bend</i> (ρ)	
1	1 2	1 2
2	2 3	4 5
3	4 5	7 8
4	6 7	12 13
5	9 10	18 19
6	12 13	26 27
	<i>bug</i> (ρ); <i>ball</i> (ρ)	
1	1 2	2 3
2	2 3	4 5
3	4 5; 5 6	8 9; 9 10
4	7 8	13 14; 14 15
5	10 11; 11 12	20 21; 21 22
6	14 15; 15 16	27 28; 29 30
	<i>spo</i> ₃ (ρ)	
1	1 2	2 3
2	3 4	5 6
3	5 6	10 11
4	8 9	16 17
5	12 13	23 24
6	16 17	31 32
	<i>boot</i> (ρ)	
	$\rho^2 \rho^2 + 1$	$\rho(2\rho + 1) \rho(2\rho + 1) + 1$

on combinatorial principles is needed; the current work of Fisch and Gravner is a first step in that direction.

First, let us consider the GH phase portrait. Our experimental conclusions for large κ are summarized thus:

$$\begin{aligned}
 \xi_t \text{ is locally periodic if and only if } \theta < \mathit{bend}(\rho) \\
 \text{dies weakly or clusters if and only if} \\
 \quad \mathit{bend}(\rho) < \theta < \mathit{ball}(\rho) \\
 \text{dies strongly if and only if } \theta > \mathit{ball}(\rho)
 \end{aligned} \tag{9}$$

Recall that these asymptotic behaviors, except clustering, were defined in Section 4. Each regime of Property 9 requires discussion. Our identification of the locally periodic regime is based on three bold claims:

- GH rules die weakly for large κ unless their wave fronts can reproduce by bending.
- If a wave front can bend in on itself, then the rule has SPOs.
- If a rule has SPOs, then the system is locally periodic.

The first claim is borne out by computer experiments, and seems ‘physical’, but we do not know a more compelling justification. The second claim is based on the principle

that easily bending wave fronts make ram’s horns, and the centers of perfect ram’s horns are SPOs. We verified this principle in one very amusing engineering feat. The goal was to produce an SPO for the macaroni-type rule (3B, 7, 5). We discovered that a band test with the given threshold and range, but with 30 colors, wrapped in on itself quite tightly and came very close to making a period of 30 ram’s horn. The period of a site would fluctuate unpredictably but remain close to minimal. After two or three hours of EXCITE! time, quite unexpectedly, the centers of the ram’s horns locked into a perfect period 30 cycle. Upon closer inspection each center contained an SPO. Recoloring mod 5, we had achieved our goal. As a reward we were able to paste our object in (3B, 7, 5) and watch it attack the macaroni. The third claim is based on our empirical observation that SPOs ‘enslave’ everything in their environment except terrain already enslaved by another SPO. That is to say, *demon* growth in $\theta = 1$ excitable CA rules generalized as SPO growth for $\theta > 1$. This property can be proved in a few special cases, but there are difficult problems concerning the manner in which SPOs spread when θ is large. Our claims include a conjecture that makes no reference to dynamics: A (ρ, θ, κ) -SPO exists for large κ if and only if $\theta < \mathit{bend}(\rho)$.

In the second regime of Property 9, *weak death* of the infinite system means that for each $x \in \mathbb{Z}^2$,

$$P(\xi_t(x) = 0) \rightarrow 1 \text{ as } t \rightarrow \infty,$$

$$\text{but } P(\xi_t(x) = 1 \text{ for arbitrarily large } t) = 1 \tag{11}$$

Between *bend* and *bug* the system supports wave fragments that annihilate upon collision but are otherwise capable of traveling indefinitely, perhaps growing all the while. In some cases only fragments with carefully arranged ends can survive and the ends caused by typical interactions are vulnerable. These models should satisfy property (11) since a few stable bugs will survive long enough to visit x after any prescribed time. We suspect that rules (1B, 2, κ) are of

Plate A. *Ram’s horns generate octagonal waves*
(GH, $\rho = 2B$, $\theta = 3$, $\kappa = 8$)

Plate B. *Three SPOs planted in the random initial state of GH*,
 $\rho = 2B$, $\theta = 4$, $\kappa = 5$

Plate C. *Bugs and SPOs generated by clocks*
(GH, $\rho = 1B$, $\theta = 2$, $\kappa = 4$)

Plate D. *Clustering from an engineered initial state of GH*,
 $\rho = 3B$, $\theta = 8$, $\kappa = 5$

Plate E. *Nucleation and spiral formation in the basic CCA*
($\rho = 1D$, $\theta = 1$, $\kappa = 16$)

Plate F. *An instance of fluid turbulence*
(CCA, $\rho = 2B$, $\theta = 5$, $\kappa = 8$)

Plate G. *Fixation of a CCA in the bootstrap regime*
($\rho = 2D$, $\theta = 5$, $\kappa = 3$)


Plate H. *Perfect spirals pop up within a metastable phase*
(CCA, $\rho = 1B$, $\theta = 3$, $\kappa = 4$)




A




B




C




D




E



F



G



H

this variety for κ sufficiently large. Another more exotic possibility, suggested by rule (3B, 8, 5) for instance, occurs when typical wave fragments grow at their ends. In this situation annihilation seems to induce increasing alignment of parades of ‘worms’ traveling in the same direction. Aligned worms also seem to be able to join at their ends and repair internal disturbances. The background state of Plate B, left on its own, exhibits this clustering behavior on the computer screen as time goes on even though that rule admits very rare SPOs. To illustrate the same phenomenon in rule (3B, 8, 5) required some invasive engineering since viable bugs are too rare to appear on the screen. Plate D was achieved by planting a smattering of suitable bands with random orientations on a background of all 0s. Although more experiments with very large arrays are needed, we suspect that alignment occurs on *arbitrary* length scales in this situation. Thus we conjecture that there is a qualitatively distinct *clustering* phase for some GH rules between *bend* and *bug*. At present our findings are so speculative that we will not even formulate a mathematical characterization of worm clustering, although we feel confident that such a formulation is possible. There are various simpler CA prototypes that exhibit large-scale clustered alignment. In a system we call *linebugs*, for instance, individual horizontal and vertical line segments move in random directions orthogonal to their orientations while growing on their ends and annihilating their overlap. Much larger configurations akin to Plate D are obtained.

The ergodic behavior of GH rules between *bug* and *ball* is virtually impossible to glean from computer visualization. In this regime wave fragments die out quickly so the graphics screen gives the impression of strong death. But we have already seen that certain large rings with the right shape can spread forever in a sea of 0s. Moreover, annihilating interaction causes two such rings to *join* when they interact. Thus, after a transient period in which virtually all fragments disappear, ξ_t should be governed by its *ring dynamics*. Observe that all *nested* rings of arbitrary multiplicities are preserved; these astronomically rare structures excite each site of the lattice at arbitrarily large times. (For rigorous results on ring dynamics of GH rules started from configurations of 0s and 1s, and for a proof of weak death for a ‘synchronized’ GH prototype with nested rings, see Gravner, 1991.)

Finally, the experiments indicate that ξ_t dies out strongly above *ball*—in the sense of display (2)—and we can conceive of no mechanism preventing rapid convergence to all 0s. However, we do not yet know a rigorous argument that improves display (2) substantially (for instance, a proof of strong death for some $\theta < boot$).

We now turn to a description of the CCA phase portrait. Our experimental findings for large κ are summarized thus:

$$\begin{array}{ll} \xi_t \text{ is locally periodic} & \text{if and only if } \theta < bend(\rho) \\ \text{is turbulent} & \text{if and only if} \\ & bend(\rho) < \theta < boot(\rho) \\ \text{fixates} & \text{if and only if } \theta > boot(\rho) \end{array} \quad (12)$$

Again, we will discuss the regimes of display (12) in order. Motivation for the locally periodic regime is very much the same as for display (10), except that the first claim should be replaced by:

- CCA rules can only make stable spiral centers for large κ if their wave fronts can bend completely, or if they belong to the multistrand family (1B, 2, $\kappa \geq 6$).

Of course, exceptional multistrand structures such as the one shown in Fig. 6 raise the specter of similar anomalies for larger ρ . We can only report that, after extensive engineering, we have not managed to find any.

So what happens to CCA rules above *bend*? Judging from a great many experiments with large arrays, there appears to be a new ‘fluid’ phase between *bend* and *boot* that we call *turbulence*. (This name is only suggestive; we do not claim any connection to classical turbulent dynamics.) Representative examples are (2D, 4, 5) and (2B, 5, 8), the latter shown in Plate F. Roughly speaking, ξ_t tries to make stable spiral centers but, as indicated by the loop test, these would-be centers are unstable. As a result, when vortices unravel they leave behind a smattering of debris that can act as a seed for the formation of new wave fronts. This debris is a permanent and crucial part of the dynamics, more like the fine-grained structure in the cyclic particle system (Griffeath, 1988) than the remnants of initial noise described in Fisch *et al.* (1991a) or in Section 10 of this paper. The turbulent fluid phase is characterized by very large length scales so even rule (2B, 5, 8) typically fixates on a 256×256 (CAM) array, due to finite-lattice effects. But CAM scooping indicates that the same rule on a 512×512 array has stable steady-state behavior. As in the case of worm clustering, we think it is premature to propose a precise mathematical formulation of turbulence. Suffice it to say that this regime entails a remarkably exotic self-organized equilibrium (stationary distribution) *not* concentrated on periodic orbits.

Once the threshold exceeds *boot*, ξ_t seems to fixate for any $\kappa \geq 3$. Plate G shows the final fixated state in a representative three-color case. Result (4) comes close to proving fixation above *boot* for ρ and κ large. A nice prototype for more in-depth analysis is the family of CCA rules (1D, 2, $\kappa \geq 3$). Perhaps it is possible to prove fixation in this case for κ sufficiently large, but the three-color situation is surely very delicate. Indeed, the name *boot* comes from *bootstrap percolation*, one variant of which is the (1D, 2, 3) CCA starting from only 0s and a random configuration with a small density p of 1s. That process, arguably the simplest prototype for *metastability effects*, is known *not* to fixate on the infinite lattice for any positive p , no matter how small (see, for example, Aizenman and Lebowitz, 1988).

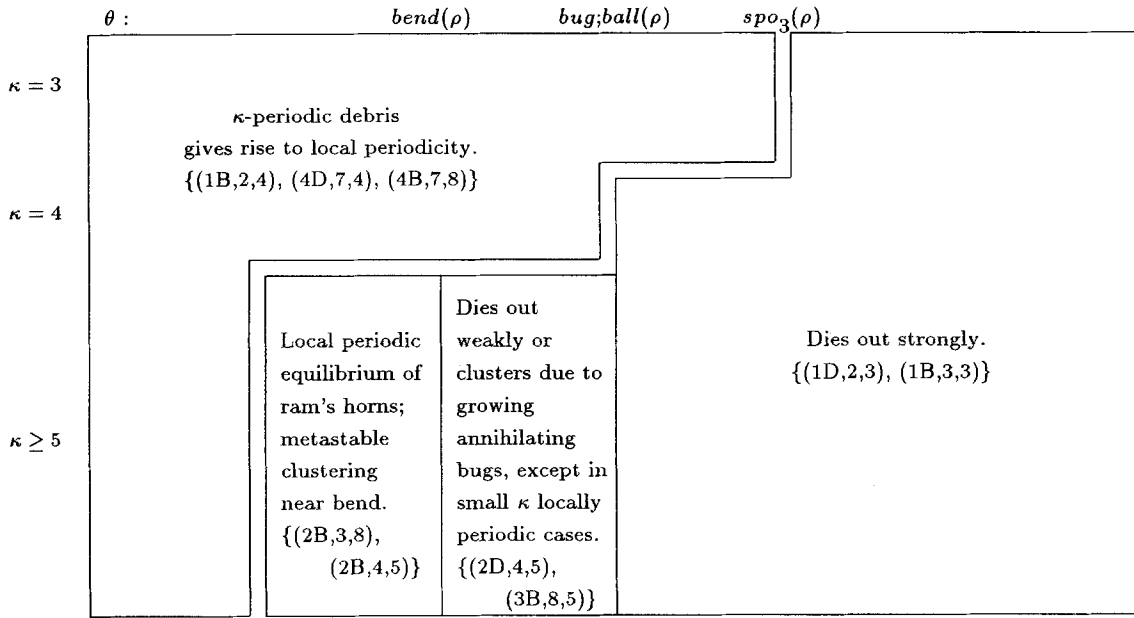


Fig. 7. Greenberg-Hastings models phase portrait

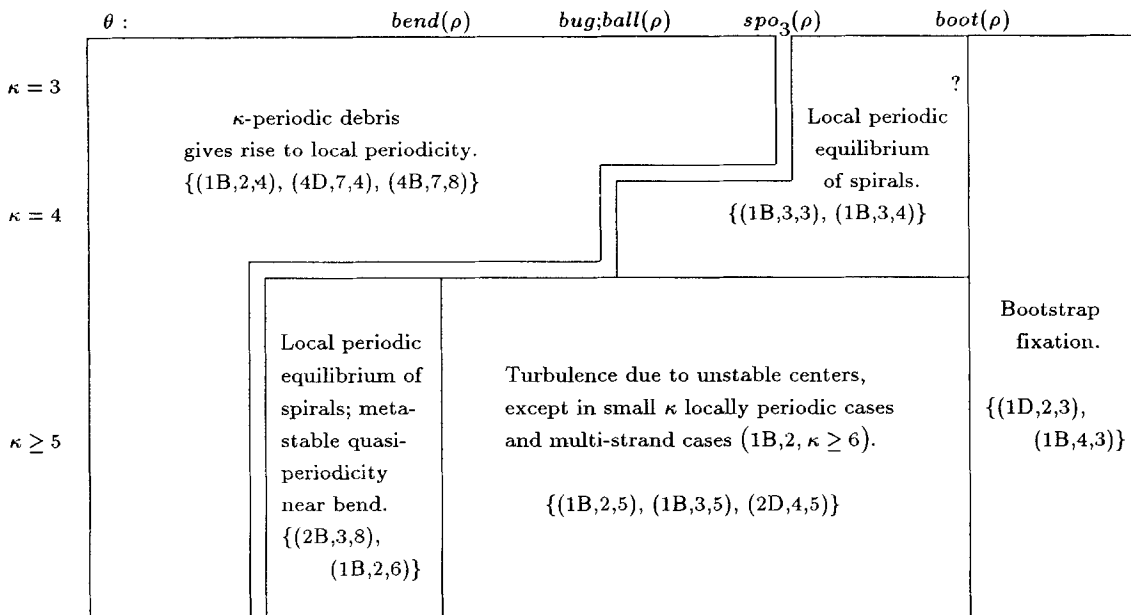


Fig. 8. Cyclic models phase portrait

Our empirical appraisals of the phase portraits for GH and CCA dynamics are summarized in Figs. 7 and 8. The periodic debris regime, its boundary, and $\kappa = 3, 4$, will be discussed in Section 10.

9. Threshold-range scaling

Our numerical values for $\rho \leq 6$ reveal an intriguing threshold-range scaling. The data suggest that $bend(\rho)/\rho^2$

converges to a limit as $\rho \rightarrow \infty$, and similarly for all the other cutoffs. Consequently the behavior of any GH or CCA rule is effectively determined by the ratio $\lambda = \theta/\rho^2$, at least for κ sufficiently large. Even models with the smallest threshold and range are remarkably consistent with the asymptotic picture.

There is a simple explanation for the scaling that is of fundamental significance for the theory of excitable cellular automata. Let us consider the case of GH rules on box neighborhoods to be concrete. Then, once a rule has

self-organized to a length scale of order ρ and consists of reasonably coherent waves, in order to update by contact site x asks the question: Is at least a proportion $\theta/(2\rho + 1)^2 \simeq \lambda/4$ of my neighbor set excited? Rescaling in the manner of Riemann, so that $\mathcal{N}_\rho(x)$ corresponds to the Euclidean neighbor set $\mathcal{N}_1(x)$ in the limit as $\rho \rightarrow \infty$, we obtain a limiting function form of contact updating in which x asks the question:

$$\text{Is } |\mathcal{N}_1(x) \cap \{\text{excited points } y\}| \geq \lambda? \quad (13)$$

Here $|\Lambda|$ denotes the area (Lebesgue measure) of Euclidean set Λ . In other words, for each $\lambda \in (0, 4)$ there is a corresponding GH dynamic defined on ‘nice’ colorings in \mathbb{R}^2 with κ colors. Any point with color 0 changes to 1 if and only if the answer to Question 13 is ‘yes’; all remaining colors update automatically. An analogous continuous space version of the CCA rule can also be defined. Moreover, the dynamics of lattice systems starting from organized lattice configurations of length scale ρ are well approximated by the dynamics of the continuous space rule started from a corresponding continuous initial state. To understand the observed consistent scaling behavior for small parameter values, note that contact updating with $\rho = 4B$ involves averaging over 80 neighbors; it is not so surprising that the Riemann sum with 80 terms agrees fairly well with the limiting integral.

The Euclidean rules have a number of mathematical advantages. Smooth wave fronts are preserved by the dynamics, so certain integral equations can be used to study key aspects of wave propagation such as the evolution of rings and bugs. Whereas lattice rules are subject to ‘Diophantine’ complexities, the threshold-range limit is surprisingly amenable to exact calculations. In addition, if we use an l^2 -neighborhood for our rules on \mathbb{R}^2 , then the models are *isotropic*. In particular the asymptotic shapes of spirals and rings are precisely circular, a desirable property according to Markus *et al.* (1991). In fact, we suspect that the recipe for isotropy proposed by Markus *et al.* (1991) is a complicated version of l^2 -ball threshold-range scaling.

Unfortunately if one wants to visualize continuous space dynamics, the most computationally efficient method is presumably to approximate by a CA with large ρ and θ . And one must also come to grips with the mysterious transition from initial randomness to waves, the topic of our next section.

10. The plot thickens

Upon closer inspection, the phase diagrams for GH and CCA rules are more complicated than the simplified description of Section 8. Let us now mention several anomalies that arise because the threshold is very low, or the number of colors is very small, or the finite lattice approx-

imation (for example, the graphics screen) is almost too small to handle the number of colors.

10.1. Periodic debris

If θ is quite small compared to ρ and κ is not too big, then SPOs will abound in the initial random configuration for *statistical* reasons. Durrett (1991) has observed that the limiting cutoff for this effect under threshold-range scaling is

$$\theta \simeq 2\rho^2/\kappa \quad (14)$$

In analogous interacting random models of either GH or CCA type, he proves that systems below this point have a stationary distribution which is asymptotically a product measure as $\rho \rightarrow \infty$. Experiments indicate that for θ just above $2\rho^2/\kappa$, random GH dynamics still have equilibria with short range correlations but the ‘stochastic spiral’ random CCA equilibria exhibit large-scale self-organization. In the cyclic case, at least, it appears that the asymptotic critical value has been identified precisely. For our excitable CA models one can prove local periodicity if $\theta < (2 - \epsilon)\rho^2/\kappa$ and ρ is sufficiently large ($\epsilon > 0$ arbitrary). The basic idea is that very large SPOs will be found in the initial state due to large deviation considerations, that this SPO region will percolate and have only bounded clusters in its complement, and finally that the dynamics will ‘fill in the holes’. This regime is characterized by a (rather boring) locally periodic limiting state of κ -*periodic debris* that is almost indistinguishable from random noise. Just above the cutoff of display (14), however, the deterministic systems seem to behave differently than the random ones. There is a curious *percolation transition* from the debris phase to the self-organized phase in which first debris predominates but pockets of wave activity are formed, and then for slightly higher θ the waves predominate but there are pockets of residual debris. The phase diagrams in Section 8 list interesting rules with mixed phases for both GH and CCA. As a rough empirical rule of thumb, substantial pockets of residual debris finally disappear once $\theta > 3\rho^2/\kappa$, although we do not claim that this value is exact in any sense. The ‘chasms’ in our phase pictures are meant to indicate the percolation transition, an effect that is not at all well understood and warrants further investigation.

10.2. Three-color phenomena

We conjecture that GH rules with $\kappa = 3$ undergo a direct transition from periodic debris to strong death at a cutoff $spo_3(\rho) \sim 2\rho^2/3$, in accordance with display (14). The name comes from the fact that SPOs abound below this cutoff, and our belief that they do not exist (for $\kappa = 3$, and hence for any κ) above it. The numerical data on spo_3 for $\rho \leq 6$ are derived by starting a large finite GH system from three-color noise with free (= all 0s) boundary conditions. If locally periodic pockets of excitation survive, then these

pockets contain SPO's and the rule is below the cutoff. If the large array dies out then θ exceeds spo_3 . Of course, this last claim is suspect if the only stable period objects are statistically very rare. But in three-color excitable CA rules there seems to be a sharp transition, so we have reason to hope that our numbers are accurate.

CCA rules with $\kappa = 3$ make periodic debris for $\theta < spo_3$ and seem to have locally periodic limits consisting of very large spirals for θ between spo_3 and $boot$, as suggested by the loop test. But we do not yet understand the detailed characteristics of the spirals, and have doubts about their stability for θ near $bend$. The very intriguing CCA rule (2B, 10, 3) even seems to cluster like some GH bug examples, and may indicate yet another phase. CAM real-time dynamics for this rule exhibit a nice surface tension effect so we sometimes call it the 'lava light' example.

10.3. Four-color phenomena

Four-color GH rules make periodic debris for $\theta < \rho^2/2$ by display (14). The percolation transition seems to last precisely to bug ; there is a strange effect whereby bugs emanating from debris pockets can stabilize the system even though the residual debris does not contain SPOs. Above bug ξ_i dies out (weakly up to $ball$, and then strongly), unless, of course, there are SPOs with strange architecture.

CCA rules with $\kappa = 4$ make nice spirals with period $n\kappa$ for some $n > 1$. In the (1B, 2, 4) case these spirals can dominate a large array for a long time before a structure, such as structure 8b, is suddenly created by the dynamics; then, quite rapidly, the spatial structure of the locally periodicity state changes completely as spirals are overtaken by SPOs. Another fascinating CCA rule is (1B, 3, 4), shown in Plate H. While most spirals that organize their environment contain occasional 'glitches', as they work their way around particularly troublesome boundary conditions, this and other four-color systems below $boot$ always manage to make very large *perfect* spirals. This ability to organize without any 'errors' whatsoever seems well worth further study.

10.4. Metastability and nucleation effects

Near the cutoffs of the phase diagrams computer visualization can also be very misleading. Due to metastability, dynamics on large finite arrays may indicate that the corresponding infinite system belongs to the regime on the other side of the phase boundary. We have mentioned that GH macaroni-type rules may appear to cluster when they are, in fact, locally periodic. Similarly, CCA rules may seem turbulent when they are locally periodic, or to fixate when they belong to the turbulent fluid phase.

Finally, for any model with ρ and θ values capable of sustaining wave activity, the distance between regions of

initial random noise that contain the 'right stuff' for self-organization increases very rapidly, with κ . GH models die out virtually everywhere, but widely separated clusters of excitation are then free to spread. CCA models fixate except for very rare wave *droplets* that manage to feed off their boundaries as they grow. These phenomena are instances of *nucleation*.

11. Rigorous results

For the most part, our methodology in this overview of excitable cellular automata has been *empirical*, relying on computer experiments, heuristics and intuition as much as precise deductive reasoning. We feel that the complex phenomenology under discussion warrants and indeed necessitates such an experimental approach. However, as mathematicians, we are also motivated by a sense that the theory of excitable CA rules constitutes a fertile terrain for rigorous results, that is, substantive theorems with illuminating proofs. Let us mention briefly some ongoing mathematical research and promising directions for additional work. Already underway are three projects that deal with GH and CCA rules on \mathbb{Z}^2 ;

1. *Nucleation scaling*. For fixed ρ and θ in the locally periodic regime, and for κ large, how far apart are the supercritical regions of initial wave activity that give rise to spirals? Closely related are the questions how large a finite ($L \times L$) system we need to ensure that GH does not die out or that CCA does not fixate. In Fisch *et al.* (1991b) nucleation of the basic GH models ($\rho = 1D$, $\theta = 1$, $\kappa \rightarrow \infty$) is studied. Using percolation ideas, it is proven that defects are formed from a random initial state on a scale that grows *exponentially* with κ , and the constant in the exponent is rigorously determined within a factor of 2. Some partial results for the analogous CCA problem will be found in Gravner (1991a); an exponential lower bound is proved there, but an analysis of droplet growth suggests that the scaling is in fact super-exponential.

2. *Existence and evaluation of critical values*. In joint work with R. Durrett, one of us (DG) will attempt to formulate rigorous definitions of limiting threshold-range scaling cutoffs, to compute as many of the cutoffs as possible, and to derive good numerical bounds for those that defy exact evaluation. We expect that some of the engineering principles described here will prove useful for both the proofs and the numerics.

3. *Asymptotic shapes*. Two of us (JG, DG), with undergraduate assistant D. Perry, will study the geometry of rings, ram's horns and spirals in the threshold-range limit. Asymptotic shapes can be computed exactly as a duality transform applied to an explicit width function. We will also try to characterize the polygonal shapes that arise for prescribed values of ρ and θ .

Two other large classes of rules exhibit ergodic behavior with intriguing connections to GH and CCA:

1. *Two-color models*. Two-color analogs of GH and CCA are called the *threshold contact automaton* and the *threshold voter automaton*, respectively (cf. Durrett, 1991). Interesting open problems abound.

2. *Random excitable systems*. A systematic overview of random excitable systems—the stochastic counterparts of GH and CCA rules, and various generalizations—would require another paper at least as lengthy as this one. The cyclic particle system in several dimensions (Griffeath, 1988) clearly indicates that large-scale self-organization is possible even in the presence of noisy dynamics. Pioneering papers such as Mollison and Kuulasmaa (1985), Durrett (1991) and Durrett and Neuhauser (1991) give a first glimpse of other exotic ergodic behaviors. Whereas SPOs are unstable under random perturbations, many of the phenomena described in this paper should have stochastic counterparts. For instance, we have seen ‘noisy ram’s horns’ in certain random GH equilibria, and we would not be surprised if clustering of aligned random worms were also possible. Perhaps our present work can serve as a useful template for future research in this exciting new area of stochastic processes.

Acknowledgements

Over the course of this project we have benefited from the help of many colleagues, but we would like to acknowledge two whose contributions have been indispensable. First, we thank Charles Bennett for introducing us to CAM modeling and the fascinating world of ram’s horns. Second, we thank Rick Durrett for his faith that these models would give rise to honest mathematics and for his guidance in showing the way on several occasions. This research was partially supported by research grants to each of the authors from the National Science Foundation.

References

- Aizenman, M. and Lebowitz, J. (1988) Metastability effects in bootstrap percolation. *J. Phys. A: Math. Gen.*, **21**, 3801–3813.
- Bennett, C., Grinstein, G., He, Y., Jayaprakash, C. and Mukamel, D. (1990) Stability of temporally-periodic states of classical many-body systems. *Phys. Rev. A*, **41**, 1932–1935.
- Bramson, M. and Griffeath, D. (1980) Flux and fixation in cyclic particle systems. *Ann. Probability*, **17**, 26–45.
- Dewdney, A. K. (1988) Computer recreations: the hodgepodge machine makes waves. *Scientific American*, August, 104–107.
- Dewdney, A. K. (1989) Computer recreations: a cellular universe of debris, droplets, defects and demons. *Scientific American*, August, 102–105.
- Durrett, R. (1991) Multicolor particle systems with large threshold and range. Submitted.
- Durrett, R. and Neuhauser, C. (1991) Epidemics with regrowth in $d = 2$. *Ann. Appl. Probability*, to appear.
- Durrett, R. and Steif, J. (1991) Some rigorous results for the Greenberg–Hastings model. *J. Theoretical Prob.*, to appear.
- Fisch, R. (1990) The one-dimensional cyclic cellular automaton: a system with deterministic dynamics which emulates an interacting particle system with stochastic dynamics. *J. Theoretical Prob.*, **3**, 311–338.
- Fisch, R. (1991) Clustering in the one-dimensional 3-color cyclic cellular automaton. *Ann. Probability*, to appear.
- Fisch, R. and Griffeath, D. (1991) EXCITE!: a periodic wave modeling environment. (1 disk for IBM 286/386 and compatibles). Freeware. Write: D. Griffeath, Mathematics Dept, University of Wisconsin, Madison, WI 53706, USA.
- Fisch, R., Gravner, J. and Griffeath, D. (1991a) Cyclic cellular automata in two dimensions, in *Spatial Stochastic Processes. A festschrift in honor of the seventieth birthday of T. E. Harris*, Birkhäuser, Boston.
- Fisch, R., Gravner, J. and Griffeath, D. (1991b) Metastability in the Greenberg–Hastings model. In preparation.
- Freedman, W. and Madore, B. (1983) Time evolution of disk galaxies undergoing stochastic self-propagating star formation. *Astrophysical J.* **265**, 140–147.
- Gerhardt, M. and Schuster, H. (1989) A cellular automaton describing the formation of spatially ordered structures in chemical systems. *Physica D*, **36**, 209–221.
- Gerhardt, M., Schuster, H. and Tyson, J. (1990) A cellular automaton model of excitable media including curvature and dispersion. *Science*, **247**, 1563–1566.
- Gravner, J. (1991) Mathematical aspects of excitable media. Ph.D thesis, University of Wisconsin.
- Greenberg, J., Hassard, B. and Hastings, S. (1978) Pattern formation and periodic structures in systems modeled by reaction-diffusion equations. *Bull. AMS*, **84**, 1296–1327.
- Greenberg, J. and Hastings, S. (1978) Spatial patterns for discrete models of diffusion in excitable media. *SIAM J. Appl. Math.*, **34**, 515–523.
- Griffeath, D. (1988) Cyclic random competition: a case history in experimental mathematics, in ‘Computers and Mathematics’, *AMS Notices*, 1472–1480.
- Hodgkin, A. and Huxley, A. (1952) A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiology*, **117**, 500–544.
- Kapral, R. (1991) Discrete models for chemically reacting systems. *J. Math. Chem.*, **6**, 113–163.
- Markus, M., Krafczyk, M. and Hess, B. (1991) Randomized automata for isotropic modelling of two- and three-dimensional waves and spatiotemporal chaos in excitable media, in *Nonlinear Wave Processes in Excitable Media*, A. Holden, M. Markus and H. Othmer (eds), Plenum Press New York. Forthcoming.
- Mollison, D. and Kuulasmaa, K. (1985) Spatial epidemic models: theory and simulations, in *Population Dynamics of Rabies in Wildlife*, P. J. Bacon (ed), Academic Press, London, pp. 291–309.

- Murray, J., Stanley, E. and Brown, D. (1986) On the spatial spread of rabies among foxes. *Proc. Roy. Soc. London B*, **229**, 111–150.
- Toffoli, T. and Margolus, N. (1987) *Cellular Automata Machines*. MIT Press, Cambridge, Mass.
- Tomchik, K. and Devreotes, P. (1981) Adenosine 3',5'-monophosphate waves in *Dictyostelium discoideum*: a demonstration by isotope dilution-fluorography. *Science*, **212**, 443–446.
- Waldrop, M. (1990) Spontaneous order, evolution, and life. *Science*, **247**, 1543–1545.
- Weiner, N. and Rosenbluth, A. (1946) The mathematical formulation of the problem of conduction of impulses in a network of connected excitable elements, specifically in cardiac muscle. *Arch. Inst. Cardiol. Mexico*, **16**, 205–265.
- Winfree, A. (1974) Rotating chemical reactions. *Scientific American*, June, 82–95.
- Winfree, A. (1987) *When Time Breaks Down: The Three Dimensional Dynamics of Electrochemical Waves and Cardiac Arrhythmias*. Princeton University Press, Princeton, N.J.

Appendix

For $R = 1, 2, \dots$, let $\mathcal{N}_x = \{y: \|y - x\|_\infty \leq \rho\}$, $D_R = \{y \in \mathbb{Z}^2: \|y\|_2 \leq R\}$, $\partial D_R = \{y \in \mathbb{Z}^2: R < \|y\|_2 \leq R + 1\}$.

Proposition. If $x \in \partial D_R$, and R is large, then

$$(\rho - 1)(2\rho + 1) \leq \#(\mathcal{N}_x \cap D_R) \leq 2\rho(\rho + 1)$$

Proof. Let us abbreviate $\|\cdot\|_2$ as $\|\cdot\|$ throughout. Given $y \in \mathcal{N}_x$, set $y = x + v$ ($v \in \mathcal{N}_0$); then

$$\|y\|^2 - \|x\|^2 = 2(x \cdot v) + \|v\|^2$$

Thus a necessary condition for $y \in D_R$ is $x \cdot v < 0$. By symmetry, the number of $v \in \mathcal{N}_0$ that satisfy this inequality is at most $\frac{1}{2}((2\rho + 1)^2 - 1) = 2\rho(\rho + 1)$. So the upper bound of the Proposition holds. Corresponding sufficient conditions to ensure that $y \in D_R$ are

$$x \cdot v \leq -(1 + \epsilon)R \quad (\text{A1})$$

and

$$\|v\|^2 \leq 2\epsilon R - 1 \quad (\text{A2})$$

for a suitably chosen $\epsilon > 0$. Note that inequality (A2) holds provided R is sufficiently large; for example, if

$$R \geq \frac{3\rho^2}{2\epsilon}$$

Dividing Inequality A1 through by $\|x\|$, it therefore suffices to obtain a lower bound on

$$n = \inf_{u: \|u\|=1} \# \{v \in \mathcal{N}_0: u \cdot v \leq -(1 + \epsilon)\}$$

For the infimal u , consider the lines in \mathbb{R}^2 given by $\ell_0 = \{z: u \cdot z = 0\}$, $\ell_1 = \{z: u \cdot z = -(1 + \epsilon)\}$, and the regions $\mathcal{R} = \{z: u \cdot z \leq -(1 + \epsilon)\}$, $\mathcal{S} = \{z: -(1 + \epsilon) < u \cdot z < 0\}$, $\mathcal{S}' = \{z: 0 < u \cdot z < (1 + \epsilon)\}$. Then $n = \#(\mathcal{N}_0 \cap \mathcal{R})$. Denote $n_0 = \#(\mathcal{N}_0 \cap \ell_0)$, $n_* = \#(\mathcal{N}_0 \cap \mathcal{S})$. Observe that ℓ_0 is a line through the origin orthogonal to u , whereas ℓ_1 is a line parallel to ℓ_0 and separated by distance $(1 + \epsilon)$. Symmetry dictates that

$$n + n_* = \frac{1}{2}((2\rho + 1)^2 - n_0)$$

so it suffices to obtain an upper bound on n_* . To this end, we consider the contribution to n_* from each ‘horizontal’ slice of \mathcal{N}_0 , that is, each section with fixed second coordinate. Since the width of the strip \mathcal{S} in the horizontal direction is at most $(1 + \epsilon)\sqrt{2}$, choosing $\epsilon = \frac{3}{4}\sqrt{2} - 1$ ensures that each slice of \mathcal{N}_0 intersects $\mathcal{S} \cup \ell_0 \cup \mathcal{S}'$ in at most three lattice points. Let n_1 be the number of such slices with no neighbors v on ℓ_0 and exactly one v in \mathcal{S} . Let n_2 be the number of such slices with no v on ℓ_0 and two in \mathcal{S} . Note that if a slice has a neighbor on ℓ_0 then it must have exactly one v in \mathcal{S} . Evidently

$$n_* = n_0 + n_1 + 2n_2 \quad \text{and} \quad n_0 + n_1 + n_2 = 2\rho + 1$$

Finally, by symmetry of \mathcal{S} and \mathcal{S}' , if a slice has two neighbors in \mathcal{S} then the same slice has one neighbor in \mathcal{S}' , and hence the corresponding slice of opposite sign has only one v in \mathcal{S} . Thus $n_2 \leq n_1$. We conclude that

$$\begin{aligned} n &= \frac{1}{2}((2\rho + 1)^2 - n_0) - n_* \\ &= \frac{1}{2}((2\rho + 1)^2 - n_0) - n_0 - \frac{3}{2}(n_1 + n_2) + \frac{1}{2}(n_1 - n_2) \\ &\geq \frac{1}{2}((2\rho + 1)^2 - n_0) - n_0 - \frac{3}{2}(2\rho + 1 - n_0) \\ &= (2\rho + 1)(\rho + 1) \end{aligned}$$

as desired. Our choice of ϵ guarantees inequality (A2) for $R \geq 30\rho^2$. \square