Robust periodic solutions and evolution from seeds
in one-dimensional edge cellular automata

Janko Gravner[1]
Mathematics Department
University of California
Davis, CA 95616
e-mail: gravner@math.ucdavis.edu

David Griffeath
Department of Mathematics
University of Wisconsin
Madison, WI 53706
e-mail: griffeat@math.wisc.edu

(September 15, 2012)

Abstract. We study one-dimensional cellular automata (CA) with values 0 and 1. We assume that such CA are started from semi-infinite configurations (those that have 0's to the left of some site), and we focus on the identification of *robust periodic solutions* (RPS), which, when observed from the left edge of the light cone, advance into any environment with positive velocity. We then utilize RPS and related concepts to analyze CA dynamics from *seeds*, i.e., initial configurations with finitely many 1's .

2000 *Mathematics Subject Classification*. 37B15, 68Q80.

Keywords: cellular automaton, robust periodic solution, regularity expansion, replication.

[1]Corresponding author, tel. 530-752-0825, fax 530-752-6635.

1 Introduction

A one-dimensional binary ***cellular automaton (CA)***[2] is an evolving configuration $\xi_t \in \{0,1\}^{\mathbb{Z}}$, given by a local, translation-invariant, deterministic rule that produces ξ_{t+1} given ξ_t. The local nature is specified by a finite interval $\mathcal{N} \subset \mathbb{Z}$ of sites, the *neighborhood*. The configuration in the neighborhood $x + \mathcal{N}$ of x, i.e., an element of $\{0,1\}^{x+\mathcal{N}}$, will be interpreted as a binary word. The *update rule* is a function $f : \{0,1\}^{\mathcal{N}} \to \{0,1\}$ such that $\xi_{t+1}(x) = f\left(\xi_t|_{x+\mathcal{N}}\right)$; that is, $\xi_{t+1}(x)$ is given by the application of f to the configuration ξ_t restricted to the neighborhood $x+\mathcal{N}$ of x, which is viewed as a configuration in $\mathcal{N}$. An assignment by f will be denoted by $\mapsto$. For example, the *Exactly 1* CA [**GG3**] is given by $\mathcal{N} = \{-1,0,1\}$, $001 \mapsto 1$, $010 \mapsto 1$, $100 \mapsto 1$, with the remaining five configurations mapped to 0.

The objects of this paper are ***range*** r ***edge CA***, which have $\mathcal{N} = [-r, 0]$, and update rules that satisfy

(1.1) $$[(r+1)\ 0\text{'s}] \mapsto 0$$

and

(1.2) $$[r\ 0\text{'s}]1 \mapsto 1$$

We are principally interested in the behavior of such CA from general ***seeds***, initial configurations with only finitely many 1's; by default, the leftmost 1 of a seed is placed at the origin. Moreover, one of our aims is to understand how CA develop nontrivial periodic solutions near the edge of the light cone, which we always choose to be the *left* edge, i.e., the furthest site to the left that can be influenced by the rule. Assumptions (1.1) and (1.2) reflect the perspective of the evolution viewed from this edge. For example, the edge version of *Exactly 1*, which we call *EDD0*, is an equivalent rule that has the neighborhood changed to $\mathcal{N} = \{-2,-1,0\}$ and the same assignments. The one-sided nature of these rules makes it convenient to consider more general ***semi-infinite initial states***: those that have only 0's to the left of the origin and, again by default, a 1 at the origin. We often call sites in state 1 *occupied* and those in state 0 *empty*.

Unless stated otherwise, all rules in this paper are edge CA. To give a famous example, *Rule 110* is a cellular automaton celebrated as the simplest supporting universal computation [**Coo**]. Instead of careful constructions such as the one in [**Coo**], this paper focuses on self-organizing properties of its edge version *E1ED*. Fig. 1 displays the evolution from two simple initializations that clearly evolve, on any interval $[0, N]$, into configurations that are temporally periodic, and also spatially periodic after an initial segment. We call such configurations *periodic solutions* of an edge CA. The initial segment, called a *handle*, does not belong to the spatially periodic part, but is instrumental in its emergence and maintenance. Both emerging configurations in the figure have the same temporal period 16 (which turns out to be the smallest possible for a periodic solution), but different spatial periods (16 and 160). The lengths of their handles are also quite different, which relates to different likelihoods of emergence from random seeds.

[2] We use bold italics for a key term when its formal mathematical definition is provided in context.

There has been considerable CA research into ***doubly periodic*** solutions (see [**BK**, **BL**] and references therein). Such solutions are spatially periodic configurations that extend over the entire lattice $\mathbb{Z}$ and are also temporally periodic for the CA rule under study. The paper [**BL**] contains an introduction and a computational approach, based on systematic case checking, conceptually related to ours. However, we must stress the distinction between periodic solutions in the semi-infinite setting of the present paper and doubly periodic solutions: we will see that many CA have none of the former but an infinite number of the latter (see Proposition 3.4 and subsequent remarks).

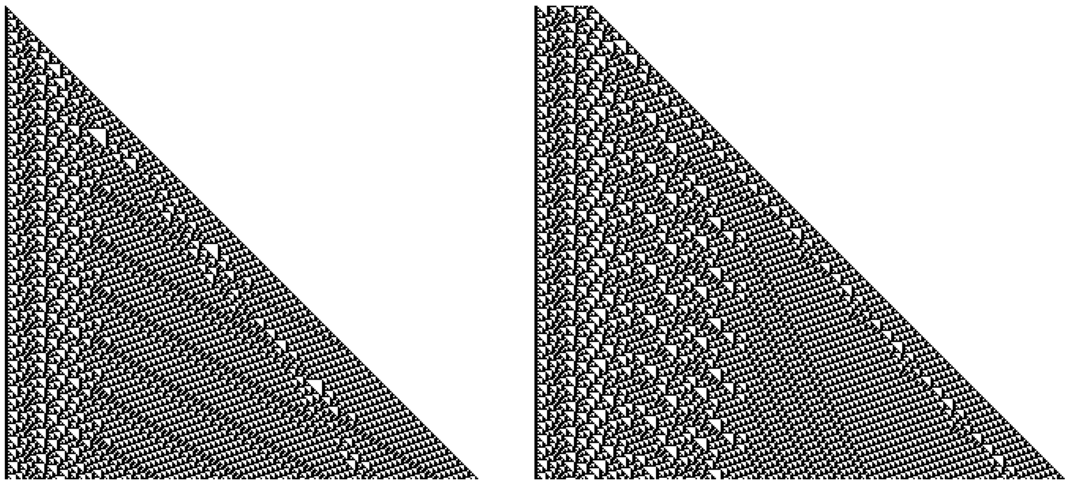

Fig. 1. Evolution of *E1ED* started from a single occupied site (left) and from the seed 11[6 0's][6 1's]000[8 1's]011010111011.

Of particular interest are periodic solutions that advance into any environment with a minimal velocity $v > 0$. We call these *robust periodic solutions (RPS)*. If an RPS η is replaced with a configuration η' which agrees with η within a long enough interval $[0, k]$ but is completely arbitrary outside $[0, k]$, then the CA dynamics from η and η' agree on a linearly growing portion of space $[0, vt]$. An RPS is thus stable: all nearby configurations converge to it at exponential rate (see Sections 2.5 and 5 for more details). Direct analogues in differential equations are asymptotically stable periodic orbits [**Ver**]. By contrast, arbitrary long segments of non-robust periodic solutions *stop* growing unless the environment to their right is very special.

In the next section we more precisely define RPS and discuss some of their properties, but we immediately mention that the two examples in Fig. 1 both have this property. It turns out that the smallest neighborhood, $r = 1$, admits only a trivial RPS. The 64 range 2 edge CA, however, feature many interesting examples; in Sections 3–8 we discuss general methods, but focus our applications on $r = 2$ cases. These correspond to the two (left and right) edge CA for *elementary* CA [**Wol1**, **Wol2**], that is, binary CA with neighborhood $\{-1, 0, 1\}$. Nearly 30 years after Wolfram's introduction of elementary CA, a systematic framework for studying these rules from seeds has still not been developed. In Sections 8 we will demonstrate that edge CA, RPS, and related concepts, are very useful tools for such investigations: we give a complete description of evolution from every seed for all but six symmetry classes of elementary CA, and explain the why the six exceptions are inherently much harder. These conclusions are based on results of the previous five sections. First, in Sections 3 and 4 we discuss general methods for proving (or disproving) existence and uniqueness of RPS and apply those methods to the $r = 2$ setting. Sections 5 and 6 study expansion velocities of the RPS to establish wedges filled by resulting periodic patterns. Section 7 addresses the edge CA with the opposite behavior, that is, those whose growth from seeds always stops. We conclude with Section 9 in which we discuss a few edge CA with $r > 2$.

As a motivation for the reader to slog through technical details we now mention a few highlights. We have managed to resolve the existence and uniqueness questions for all 64 range 2 edge CA, proving that 42 have no RPS, 19 a unique one, and 3 multiple ones. In the last category there is an example (*E1DD*) for which we can completely describe its infinite set of RPS. We also have a variety of results for expansion velocities, the best one a complex example (*ED1D*, see [**SCJ**]), which has a unique RPS that, for every seed, expands faster than is its minimum expansion velocity in half-infinite environments. Among our techniques, we single out regularity expansion (Section 2.7), which is a very effective tool for proving a variety of results.

Our present work is motivated primarily by earlier investigations of CA growth in two dimensions [**GG1**, **GG2**, **Hic1**], and by construction of periodic solutions in the *Exactly 1* CA [**GG3**, **Hic2**]. Many of our results were first conjectured through computer experimentation, using *MCell* [**Woj**], *Golly* [**TR**], and additional custom software. The Wolfram Atlas of Elementary Cellular Automata [**Wol5**] is another useful resource. We will maintain a web archive [**Gra**], which contains data used in this paper and additional examples.

2 Basic definitions and methods

2.1 Some notation

Contiguous intervals created by k concatenated copies of a configuration A will be indicated by [k A's] (e.g., [3 01's] means 010101); we also use A^∞ to denote a half-infinite concatenation of copies of A. The symbol $*$ will denote an arbitrary or unknown state (e.g., *Exactly 1* has the property that $11* \mapsto 0$). Moreover, a *block* in a configuration is a longest contiguous interval of one state, a 0-*block* of state 0 and a 1-*block* of state 1. The *length* of a block is the number of its sites, and the terms *odd* and *even* blocks refer to their lengths. For example, a configuration with no 1's on odd integers has only odd or infinite 0-blocks.

The configuration at time t of a CA under study will always be denoted by ξ_t. When we want to consider dependence of the evolution on the initial configuration A_0, we put it in the superscript: $\xi_t^{A_0}$.

2.2 Nomenclature

We introduce an edge CA naming convention. Wolfram's serial numbers [**Wol1**, **Wol2**] are in wide use, but do not make it easy to discern properties of the rule or its label dynamics defined below. For this reason, we prefer to use a base 4 notation. Every edge CA is given by a sequence of 2^r symbols, each chosen from the set $\{E, D, 0, 1\}$. All possible binary configurations in $[-r,-1]$ are assumed to be in lexicographic order and to each we associate one of the four symbols, which are then listed in the same order. For a configuration A on $[-r,-1]$, symbol E (for *equal*) signifies that $Ac \mapsto c$, for both $c \in \{0,1\}$, and in this case we refer to A as an *E-inducer*. Similarly, D (for *different*) is associated with A when $Ac \mapsto 1-c$ for both $c \in \{0,1\}$, *0* when $A* \mapsto 0$, and *1* when $A* \mapsto 1$; in these cases A is respectively a *D*-, a *0*-, and a *1-inducer*. Any configuration which is either a *0*-inducer or a *1*-inducer is called a *decider*. For example, *Exactly 1* is given by *EDD0*, while Wolfram's *Rule 214* is given *by EDD1* and for both of these rules the only decider is 11. Edge CA rules (i.e., those that satisfy (1.1) and (1.2)) are characterized by the initial E.

Throughout the remainder of the paper, we will only refer parenthetically to Wolfram numbers when discussing the left edge dynamics of the most widely studied $r=2$ CA. We will also sometimes give other natural names.

2.3 Label trees

A ***cycle*** with (temporal) period π for an edge CA is a periodic orbit on an interval $[0,N]$: an evolution starting from a configuration ξ_0 such that $\xi_\pi = \xi_0$ on $[0,N]$. Observe that, for any fixed N, and any initial configuration, the evolution eventually enters into a cycle of some period π. We now describe a convenient representation of the set of all possible cycles with fixed π; we naturally do not distinguish between temporal shifts of the same cycle.

Fix a cycle on $[0,N]$. For $k=0,1,\dots,N$, and $m=0,\dots,\pi-1$, let $\lambda_k[m] \in \{0,\dots,2^r-1\}$ be the number whose binary representation is the configuration in $[k-r+1,k]$ at time m; we will give $\lambda_k[m]$ in dyadic form (when viewing it as a configuration) or decimal form (to save space), as convenient. The ***label at generation*** k is the vector $\lambda_k = (\lambda_k[m]; m=0,\dots,\pi-1)$. Note that square brackets are used to denote coordinates of labels. Also observe that the state at site k at time m is given by $\lambda_k[m] \bmod 2$.

For a simple example, we consider the $r=2$ rule *ED1D* on $[-1,2]$ and $\pi=2$. Below is a cycle and its three labels. Recall that all states to the left of the origin are 0's; by convention, we write a decimal label horizontally as a word.

$$\begin{matrix}0100\\0111\end{matrix} \qquad \lambda_0 = \begin{bmatrix}01\\01\end{bmatrix} = 11 \quad \lambda_1 = \begin{bmatrix}10\\11\end{bmatrix} = 23 \quad \lambda_2 = \begin{bmatrix}00\\11\end{bmatrix} = 03$$

Assume that one of the coordinates of the label $\lambda \in \{0,\dots,2^r-1\}^\pi$ at generation k is a decider. Such a label uniquely determines the label at the next generation, that is, any cycle on $[0,k]$ whose last label is λ extends in a unique way to a cycle on $[0,k+1]$, with the same temporal period π. If none of the coordinates of λ is a decider and the number of its D-inducers is even, then there are two possibilities for the next-generation label: the 0-successor λ^0 and the 1-successor λ^1. The 0-successor (resp. 1-successor), is determined by assuming its first coordinate $\lambda^0[0]$ is a configuration built by the rightmost $r-1$ states of $\lambda[0]$ followed by a 0 (resp. 1). These determine the two possible extensions of the cycle to $[0,k+1]$, again with the same period π. Finally, if none of the coordinates of λ is a decider and the number of its D-inducers is odd, then there is no next label possible with temporal period π and there is, at this period, no extension of the cycle to $[0,k+1]$. However, if we embed period π labels in period 2π ones in the natural way, then we do get two successor labels. It easily follows that all periods π of cycles on $[0, N]$ are powers of 2 [**GG3**].

In the *ED1D* example above, the given cycle restricted to $[-1, 0]$ can be continued in two ways to $[-1, 1]$ as λ_0 has no deciders or D-inducers; we have chosen the 0-successor $\lambda_0^0 = \lambda_1$. Then λ_1 has a decider, as its first coordinate is $10 = 2$. Thus this cycle restricted to $[-1, 1]$ has a unique extension to $[-1, 2]$ and λ_2 is uniquely determined. In turn, λ_2 has no deciders and a single D-inducer, thus there is no way to extend the cycle to $[-1, 3]$ with temporal period $\pi = 2$.

This process naturally generates the ***label tree*** associated with a given period π, in which every node has a label that determines whether the number of its successors is 0, 1, or 2. By (1.1)–(1.2), the root label is the vector 1 with all coordinates 1 (similarly, we will use the shorthand 0 for a label with all coordinates 0), and successor labels are given by the rules in the previous paragraph. Thus any possible label at generation k as defined above appears as a label at tree generation k (i.e., at graph distance k from the root). A label tree may be finite, e.g., for *EDD0* at $\pi = 4$ [**GG3**] or for *ED1D* at $\pi = 2$, which happens exactly when long enough intervals have no cycle of period π. Call the label λ' *a rotation* by $i \in \{0, \ldots, \pi - 1\}$ of label λ if $\lambda'[m] = \lambda[(m - i) \bmod \pi]$ for all $m \in \{0, \ldots, \pi - 1\}$. Observe that the label tree is infinite if and only if a label at a node is a rotation of a label at a node in its ancestral line.

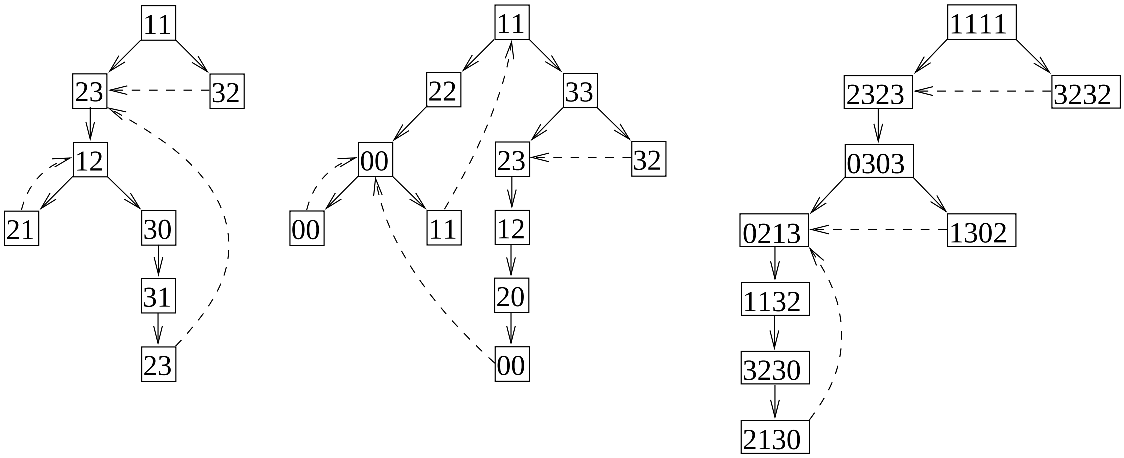

Fig. 2. The final Stage I trees, at their respective π_0, for *EDD1*, *EE0D*, and *ED1D*; dyadic labels are given as horizontal words. Backward edges, from a label to its rotation, are dashed.

Instead of dealing with the complete label tree, a commonly infinite object, we construct the ***Stage I tree*** as follows. Order nodes first so that previous generations are ahead of later generations, and within the same generation so that 0-successors (resp. their progeny) are ahead of 1-successors (resp. their progeny). A node is without successors, i.e., a leaf, if it either

- dies by period doubling: its label has no decider and an odd number of D-inducers; or
- dies by repetition: its label is a rotation of the label of a node ahead of it in the ordering.

Clearly, a Stage I tree is always finite. (See Fig. 5 in [**GG3**] for two *EDD0* examples. The $\pi = 2$ *ED1D* tree is a subtree of the $\pi = 4$ one in Fig. 2 comprised of the four nodes whose labels have period at most 2.) We augment the Stage I tree with oriented *backward* edges, which originate from any leaf x that dies by repetition and point to the first node in the ordering that duplicates x's label up to rotation. Note that a backward edge never points to a leaf, and that at most one backward edge can originate at a node (but many can end at the same node; see Fig. 2). The original edges, oriented from a predecessor to a successor, are called *forward*.

If, for some temporal period π_0, the Stage I tree has no deaths due to period doubling, then the tree is ***final***, i.e., the same is true for all $\pi \geq \pi_0$, with labels that are periodic extensions of the labels for π_0. When the rule has no D-inducers, this occurs immediately for $\pi_0 = 1$, and there can be no more than 2^{r+1} vertices in the tree. See Fig. 2 for examples with $\pi_0 \geq 2$ and [**Gra**] for pictures of all 44 finite label trees for $r = 2$ edge CA.

When there is a final Stage I tree, let the set of ***final configurations*** consist of all configurations on $[0, \infty)$ that are temporarily periodic with period π_0 (but not necessarily spatially periodic). These are all possible limit points: on any fixed interval $[0, N]$, ξ_t is eventually periodic with period π_0, and thus agrees with a final configuration there.

2.4 Periodic solutions and RPS

Let L be a finite configuration of length σ. Take another finite configuration H, and assume it has length h, with a 1 at the left end. Form the configuration HL^∞ by appending infinitely many concatenated copies of L to the right of H. Starting from HL^∞, run the edge CA until time π. Assume that there exist configurations $H_0, H_1, \ldots, H_{\pi-1}$ of length h, and $L_0, L_1, \ldots L_{\pi-1}$ of length σ, so that $H_0 = H$, $L_0 = L$ and at time t the configuration generated by the edge dynamics is $H_{t \bmod \pi} L^\infty_{t \bmod \pi}$. Then we call H a ***handle*** and L a ***link***, and the handle-link pair, which contains all information about the resulting ***periodic solution*** HL^∞, is denoted by $H + L$, with temporal period π and spatial period σ. For a doubly periodic solution, we also call a configuration of length σ, which periodically repeats in space, a link.

We assume that the spatial and temporal periods are minimal, i.e., the link L cannot be divided into two or more identical pieces and HL^∞ does not repeat at some time which is a proper divisor of π.

We call the $\pi \times \sigma$ matrix with rows L_i, $0 \leq i < \pi$, the ***tile*** of the periodic solution. Such a tile is a discrete torus, i.e., for any $i, k \in \mathbb{Z}$, the state at position (i, k) is the one at $(i \bmod \pi, k \bmod \sigma)$. Horizontal and vertical rotations give $\sigma \cdot \pi$ possible first rows (not necessarily all different). Interpret the first row as a binary representation of a number and choose the sequence of 0's and 1's with the *smallest* such number among all possibilities. This gives the ***signature*** of a periodic solution. We do not generally distinguish between periodic solutions with the same signature (i.e., all solutions with the same tile up to a rotation are considered identical, even though their handles may not be equal up to a rotation). Without loss of generality, we will assume that the handle is long enough so that the final handle label generated by HL^∞ is the same as the final link label: $\lambda_h = \lambda_{h+\sigma}$; a handle thus contains a "start" of the corresponding link. It is, however, quite possible that the handle in a periodic solution has a higher temporal period than the link; see Example 2 in [**GG3**] and Proposition 9.2 below.

The handle-link pair $H + L$ is *robust* if for each label λ_k; $k = h, \ldots, h + \sigma - 1$, there exists at least one m so that $\lambda_k[m]$ is a decider. Equivalently, every one of the σ $\pi \times r$ submatrices of the tile should contain at least one deciding row. Such a pair results in a ***robust periodic solution (RPS)***. Note that robustness depends only on the link L, so we sometimes also refer to L as robust, implicitly implying that a handle with such a link exists.

For example, *EDD0* (*Exactly 1*) has a single decider 11 and many known RPS. On the other hand, the rules with only E and D in their names cannot have an RPS since they have no decider.

We call an RPS with signature link L ***globally attracting*** if for any half-infinite initialization ξ_0 there exists a time t_0 so that ξ_{t_0} begins with a handle H for L. Thus, at periodically spaced times t, ξ_t agrees with HL^∞ on a linearly growing portion of space. Clearly, an RPS that is globally attracting is unique, but we will see that the reverse implication may not hold.

2.5 The R-algorithm

If a Stage I tree (final or not) can be generated for a given π, the following ***R-algorithm*** finds all possible links in a robust handle-link pair for period π [**GG3**].

A leaf in the Stage I tree with a label that includes a decider will be called a *decider leaf*. For each decider leaf, carry out the following procedure:

(R1) Generate the successors of the leaf, adding them to the Stage I tree, until one of them is an R-repetition or has a label with no decider.

The criterion for a node x to be an *R-repetition* is as follows. Trace back along the chain of ancestors of x with labels that include at least one decider. Call the resulting set of labels the *R-set* of x. Then x is an R-repetition if its label is a rotation of a label in its R-set.

Every decider leaf x for which (R1) ends in an R-repetition gives a robust handle-link pair, by the following formulas that construct the handle and link from the relevant labels:

(R2) Denote the label of node x by $\lambda(x)$. Consider the lineage starting from the root x_0, and continuing as $x_1, \dots, x_{i_0}, \dots, x_g = x$, where $\lambda(x_{i_0})$ is in the R-set of x_g, and $\lambda(x_g)$ is the rotation by r of $\lambda(x_{i_0})$. Let $n_2 = \pi/\gcd(\pi, r)$, with the convention that $n_2 = 1$ if $r = 0$. Then $h = i_0 + 1$, $\sigma = (g - i_0) \cdot n_2$, $H_i = x_i \bmod 2$, $i = 0, \dots, i_0$, and L the vector with σ entries

$$\lambda(x_{i_0+1})[(jr) \bmod \pi] \mod 2 \;,\; \dots \;,\; \lambda(x_g)[(jr) \bmod \pi] \mod 2 \;;\; j = 0, \dots n_2 - 1,$$

specify a robust handle-link pair. Determine the spatial period σ and the temporal period π; recall that these are assumed to be minimal.

After the procedures in (R1) and (R2) are concluded, additions to the Stage I tree are erased and the algorithm proceeds to the next decider leaf.

Proposition 2.1. *Up to rotation, all possible links from robust handle-link pairs are obtained by the R-algorithm.*

Proof. For two nodes x and y in the Stage I tree, we write $x \downarrow y$ if x is connected to a y by a sequence of forward edges $x = x_0 \to x_1 \to \ldots \to x_n = y$ in which every label $\lambda(x_k)$, $k = 0, \ldots n$ contains at least one decider (so that there is no branching). We write $x \uparrow y$ if there is a backward edge from x to y; recall that in this case x is a leaf, but y is not.

Labels $\lambda_h, \ldots, \lambda_{h+\sigma} = \lambda_h$, generated by a robust handle-link pair (with handle size h and temporal period σ), are labels of a chain of successors in the (complete) label tree. Thus they generate a sequence of nodes $y_0, \ldots, y_k$ in the Stage I tree with $\lambda_h = \lambda(y_0)$ and $y_0 \downarrow y_1 \uparrow y_2 \downarrow y_3 \ldots y_k \uparrow y_0$. It is easy to see that y_1 is a decider leaf detected by the R-algorithm. □

For example, in the *EDD1* Stage I tree of Fig. 2, the leaf with label 23 is a decider leaf. Its first successor generated by (R1), however, has label 12 and (R1) thus terminates without finding a robust pair. The other decider leaf with label 32, a rotation of 32, suffers the same fate, so there are no RPS. The same is true for *EE0D*, because the only decider link, the one with label 32, causes (R1) to terminate at label 00. Finally, the *ED1D* Stage I tree has three decider leaves. The one with label 3232 results in termination of (R1) at 3030. However, for the one with label 1302, (R1) generates, in order, three successors 3211, 3032, and 3201, terminates in an R-repetition, and thus yields a robust pair 1001+111000110101, with $\sigma = 3\pi_0 = 12$. The last decider leaf has label 2130 which is a rotation of the one just considered, so it yields the same RPS; thus there is a possibility of duplication which our implementation of the R-algorithm detects and eliminates.

2.6 Expansion velocity

Let $H+L$ be an arbitrary handle-link pair, and let A_0 be any half-infinite initialization for which ξ_t eventually agrees with H on $[0, h-1]$. We call such initializations ***proper*** for the periodic solution given by H and L. For any $t \geq 0$, let s_t be $x+1$, where x is the maximal integer for which the configurations started from A_0 and from HL^∞ agree on $(-\infty, x]$ at time t. Define the ***expansion velocity in environment*** A_0 as

$$v(A_0) = \liminf_{t\to\infty} \frac{s_t}{t},$$

and the ***expansion velocity*** of the periodic solution $H+L$ by

$$v = \inf\{v(A_0) : A_0 \text{ is a proper initialization for } H+L\}.$$

The reader should note that both $v(A_0)$ and v depend only on the signature of the periodic solution and not on its handle (although a handle is necessary to "start" a solution).

Proposition 2.2. *A handle-link pair is robust if and only if its expansion velocity is nonzero.*

Proof. The "only if" part has the same proof as Proposition 5.1. in [**GG3**]. To prove the "if" part, assume that the pair is not robust. Then there exists a k with $h \leq k \leq h+\sigma-1$, so that for each m, $0 \leq m < \pi$, $\lambda_k[m]$ is either a D-inducer or an E-inducer. The number of D-inducers must be even, as λ_k has at least one successor, but then it must have *two* distinct successors. One of them must be different from the one required to continue the periodic solution, which can therefore be stopped by that different choice. □

It is easy to see that $v \geq 1/\pi$ for any RPS. This a priori bound can be much improved by exploring the positions of deciders in the tile, which we will not pursue since we develop even better methods in Section 5.

2.7 Replication

One of our tools for proving that RPS do not exist for a given edge CA is replication [**GG3**, **GGP**]. In this paper, we will use only the following edge version of the behavior, very close to the one in [**GG3**]. The exposition here is self-contained but brief; we refer the reader to [**GGP**] for much more. The *EEDD* rule is the edge version of *Rule 90*; we will denote by ℓ_t its evolution started with a single occupied site at the origin. Then, for $t \geq 1$,

$$\ell_t(x) = (\ell_{t-1}(x-2) + \ell_{t-1}(x)) \bmod 2 \ .$$

This is arguably the simplest additive rule and its evolution generates a linear deformation of Sierpinski gasket [**Wil1**]. Intuitively, replication occurs in an evolution ξ_t if, at an arithmetic progression of times t_k, ξ_{t_k} is obtained by replacing properly spaced 1's from ℓ_k by finite configurations from a finite collection.

Formally, a *replicator rule* is a finite set $\mathcal{K}$ of finite configurations (none of which consists of all 0's though some may begin with 0's), called *replicating elements*, together with functions $\texttt{left}, \texttt{right} : \mathcal{K} \to \mathcal{K}$. An initial seed A_0 for an edge CA ξ_t is a ***replicator*** if there exist a replicator rule ($\mathcal{K}$, $\texttt{left}, \texttt{right}$), and $t_0 \geq 0, n_0 \geq 0$, so that the configurations ξ_t at times $t_k = t_0 + 2^{n_0}(k-1)$, $k = 1, 2, \ldots$, satisfy the following:

- for every k and x such that $\ell_k(x) = 1$, there is a replicating element, i.e., a configuration $K_{k,x} \in \mathcal{K}$, placed so that its leftmost state is at $2^{n_0}x$;
- all other states are 0's; and
- $K_{k,x} = \texttt{right}(K_{k-1,x-2})$ if $\ell_{k-1}(x-2) = 1$ and $\ell_{k-1}(x) = 0$, and $K_{k,x} = \texttt{left}(K_{k-1,x})$ if $\ell_{k-1}(x-2) = 0$ and $\ell_{k-1}(x) = 1$.

Thus all $K_{k,x}$ are determined by the initial pair $(K_{1,0}, K_{1,2})$ and successive applications of `left` and `right`. It is assumed that 2^{n_0} is large enough so that placed configurations do not overlap. As an illustration, Fig. 3 depicts *EE1D* started at 1100110001. In this case, we have the onset time $t_0 = 2$, step size $2^{n_0} = 4$, $\mathcal{K} = \{11, 110101\}$, and `left` and `right` identities except that $\texttt{left}(110101) = 11$.

We call an edge CA a ***replicator*** if every seed is a replicator.

The paper [**GGP**] considers more general replicators, with *EEDD* replaced by an arbitrary additive rule. In the present paper, all instances of non-additive replication are based on *EEDD*, so we have restricted the definition to this case. Furthermore, all our replicators (in Section 3) agree with additive or quasiadditive [**GGP**] dynamics outside a bounded neighborhood of the line $x = 0$ or the line $x = 2t$. Therefore, $\mathcal{K}$ can be made to consists of two configurations with easily identifiable succession rules `left` and `right`, thus we do not exhibit these objects explicitly in the arguments.

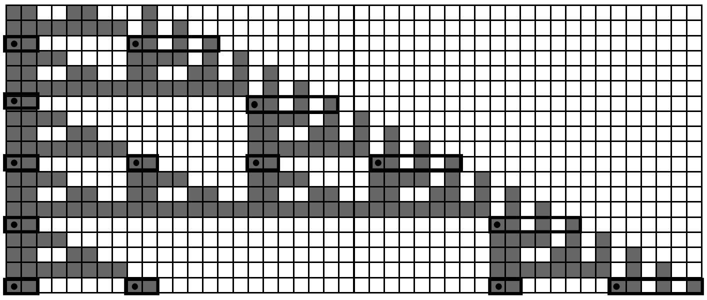

Fig. 3. *EE1D* from a small seed, run until time 18. Copies of replicating elements are emphasized by thick frames, with bullets placed at their left endpoints. After the coordinates of bullets are divided by 2, their positions match a translation of $\{(x,t) : t > 0 \text{ and } \ell_t(x) = 1\}$. Note that, `left` decides which configuration (if any) is placed below a replicating element, while `right` decides which one (if any) goes below and to its right.

The reason we consider this concept here is the following relation between replication and periodicity, which in fact holds for any type of replication considered in [**GGP**].

Proposition 2.3. *A replicator CA cannot have an RPS. Moreover, if a rule has at least one replicating seed, it cannot have a final Stage I tree.*

Proof. Assume throughout that ξ_t is started from an arbitrary replicating seed.

Let M be the longest length of a configuration in $\mathcal{K}$. It is well-known (and easy to see by induction) that at times $k = 2^m$, ℓ_k only contains two occupied sites, at 0 and at 2^{m+1}. Thus, at corresponding times t_k, ξ_{t_k} generates a longer and longer interval of 0's staring at position M. As a periodic solution with link 0 is not an RPS for any edge CA, ξ_t can not generate an RPS, which proves the first statement.

Now fix an $m_0 \geq 0$ and let A_i be the configurations on $[0, 2^{m_0+n_0+1} + M]$ generated by ξ_t at times $t = t_0 + 2^{n_0}(i \cdot 2^{m_0} - 1)$. When i is even (resp. odd), $\ell_{i \cdot 2^{m_0}}$ has 0 (resp. 1) at 2^{m_0+1} and 0's on $[1, 2^{m_0+1} - 1]$. Therefore, A_i at any even time differs from all A_i at odd times. It follows that the period of the resulting cycle on this interval is at least $2^{m_0+n_0}$. Thus there are cycles with arbitrary large periods and there cannot be a final Stage I tree. □

2.8 Regularity expansion

This is one of our main techniques to analyze behavior from seeds. Typically, a regular region, with desirable features that simplify further analysis, appears at one of the edges of the seed and proceeds to expand with sufficient velocity to make almost the entire configuration regular. To make the computational approach possible, regularity must be checkable by a finite automaton. This is therefore a generalization of the RPS concept and we now give a formal definition.

Assume that $\mathcal{R}$ is a set of half-infinite configurations (recall that we are assuming that the leftmost 1 is at the origin). We call $\mathcal{R}$ ***attracting*** for an edge CA ξ_t if

- $\mathcal{R}$ is ***invariant***, i. e., $\xi_0 \in \mathcal{R}$ implies that $\xi_1 \in \mathcal{R}$;

- ∘ there exists a finite automaton $\mathcal{F}$, which starts at the origin, stops at some finite position $x \in \mathbb{Z}$ when the configuration does not belong to $\mathcal{R}$, and never stops otherwise;
- ∘ the ***expansion velocity*** $v_{\mathcal{R}}$ of $\mathcal{R}$, as defined below, is nonzero.

We call $\mathcal{F}$ is *efficient* if, for every configuration $A \notin \mathcal{R}$, $\mathcal{F}$ stops on the first (i.e., leftmost) site x such that no configuration that agrees with A on $[0, x]$ is in $\mathcal{R}$ (and thus it is unnecessary to examine A after x). It is easy to see that, as soon as a finite automaton that recognizes $\mathcal{R}$ in the above sense exists, an efficient one exists as well. We will thus assume that $\mathcal{F}$ is efficient throughout the paper.

To define the expansion velocity, let $s_{\mathcal{R}}(A) \in [1, \infty)$ be the $x \in \mathbb{Z}$ at which the finite automaton stops when run on the configuration A. Then, for an initial configuration A_0,

$$v_{\mathcal{R}}(A_0) = \liminf_{t\to\infty} \frac{s_{\mathcal{R}}(\xi_t^{A_0})}{t}$$

and

$$v_{\mathcal{R}} = \inf_{A_0} v_{\mathcal{R}}(A_0).$$

In applications, $\mathcal{R}$ is most often the set of final configurations generated by a final Stage I tree. In that case, all three conditions for attractiveness are automatically satisfied. Other times, the desired $\mathcal{R}$ is identified through experimentation. In every instance, we need a good lower bound on $v_{\mathcal{R}}$, a task that depends heavily on case checking and computational analysis. We also note that if the distance between half-infinite configurations $A = a_0 a_1 \dots$ and $B = b_0 b_1 \dots$ is defined as $\text{dist}(A, B) = 2^{-\inf\{i: a_i \neq b_i\}}$, then $\text{dist}(\xi_t, \mathcal{R}) \le 2^{-v_{\mathcal{R}} t}$ and so $-v_{\mathcal{R}}$ is a kind of Lyapunov exponent [**BS**].

2.9 Growth velocity

In some instances (see Sections 7 and 8) we are also interested in how rapidly the occupied set grows. *Only in this context*, we consider configurations A that have no occupied sites to the right of some site, but may have occupied sites at arbitrary negative sites; we define $s_g(A)$, the ***(right) boundary*** of A, to be the position of the rightmost 1 in A. If an initial configuration A_0 has $s_g(A_0) < \infty$, we define its ***growth velocity*** to be

$$v_g(A_0) = \limsup_{t\to\infty} \frac{s_g(\xi_t^{A_0})}{t},$$

and the ***growth velocity of an edge CA*** is

$$v_g = \sup_{A_0} v_g(A_0).$$

A recent work by Brummit and Rowland [**BR**] studies asymptotic properties of $s_g(\xi_t^{A_0})$ for all range 3 rules and simple initial conditions A_0.

2.10 Mirror rules

The ***mirror*** of an edge CA with update rule f is the CA with the same neighborhood $[-r, 0]$ and update rule f^m given by $f^m(A) = f(A^m)$, where A^m is the reflection of the configuration A.

Interplay between a CA and its mirror is very useful when studying evolution from seeds. For instance, it is easy to see that the mirror of any replicator CA is also a replicator. A much more interesting application utilizes the possibility that a regular region spreads from the *right* edge fast enough so that it eventually reaches an RPS which expands from the left edge. To establish this rigorously, one needs an a priori lower bound on the expansion velocity from each edge, and the two bounds must add up to strictly more than r. The RPS now advances into the new environment, which may have a lower light speed (to be discussed below), and this fact alone may cause an increase in expansion velocity. Such mirror arguments are restricted to seeds and there is no possibility to extend them to semi-infinite configurations.

Assume that $\mathcal{R}$ is any set of configurations. We call a real number $c \geq 0$ a ***light speed*** for $\mathcal{R}$ if there exists a constant $d > 0$ such that the following holds for every $b, t \geq 0$: if $A_0 \in \mathcal{R}$ and $B_0 \in \mathcal{R}$ agree on $[b, \infty)$, then $\xi_t^{A_0}$ and $\xi_t^{B_0}$ agree on all integers in $[b + ct + d, \infty)$. (We will only be concerned with upper bounds on information propagation, so we will not bother taking the infimum over such c.)

For example, if all configurations in $\mathcal{R}$ are periodic with common period π, then its light speed is 0, as one can take $d = (\pi - 1)r$. Trivially, $c = r$ is a light speed for any $\mathcal{R}$. For a more interesting glider case, see Lemma 6.2. The following two observations, whose simple proofs are omitted, specify how to use properties of mirror rules to get lower and upper bounds on expansion velocities in seed environments.

Lemma 2.1. *Assume an edge CA given by the rule f. Assume that its mirror rule f^m has an attracting set $\mathcal{R}'$ with expansion velocity $v_{\mathcal{R}'}$ and a light speed c. If an RPS for f has expansion velocity that satisfies $v > r - v_{\mathcal{R}'}$, then $v(A_0) \geq r - c$ for every proper seed A_0. Similarly, if an attracting set $\mathcal{R}$ for f satisfies $v_{\mathcal{R}} > r - v_{\mathcal{R}'}$, then $v_{\mathcal{R}}(A_0) \geq r - c$ for every seed A_0.*

Lemma 2.2. *Assume that the edge CA rule f has a periodic solution $H + L$ with velocity v. Assume also that the mirror rule f^m has a globally attracting RPS, with expansion velocity v' and such that L^m is not one of its links. Then, for every proper seed for $H + L$, $v(A_0) \leq r - v'$.*

3 Range 2 edge CA: non-existence of RPS

In this section we present four methods to prove that RPS do not exist for an edge CA: lack of deciders, replication, inspection of the final Stage I tree, and left permutativity. Of the 64 edge CA with $r = 2$, we can thereby prove that 42 do not admit an RPS. We begin with a trivial necessary condition for RPS already mentioned in the previous section.

Proposition 3.1. *An edge CA with no deciders has no RPS.*

An edge CA with no deciders is also called ***right permutative*** in the literature, as the map $Ac \mapsto d$ is bijective (i.e., a permutation of $\{0,1\}$) for every configuration A of length r.

The 8 such edge CA with $r = 2$ are:

- *EDDD*, *EDDE* (*150*), *EDED* (*102*), *EDEE*, *EEDD* (*90*), *EEDE* (*154*), *EEED*, *EEEE*

The last of the above rules is the identity CA. The four with their Wolfram numbers in parentheses are replicators and cannot have RPS also by Proposition 2.3. We will now address the 12 edge CA for which we can use replication arguments; 8 of these are thus newly proved to be without an RPS.

- *EDDE*, *EDED*, and *EEDD* are *additive* automata: if $a_1a_2a_3$ is the configuration in $[-2, 0]$, they are given, respectively, by the following sums modulo 2: $a_1 + a_2 + a_3$, $a_2 + a_3$, and $a_1 + a_3$. See [**Lin**, **Wil1**, **Wil2**] for discussion of their replication properties.

- *E0D0* (*18*), *E0DE* (*146*), and *E11D* (*126*) were first analyzed rigorously by E. Jen [**Jen2**, **Jen3**]; they are also all replicators [**GG2**, **GG3**, **GGP**].

- *ED1E* is equivalent to Jen's replicator *E0DE* [**GGP**]. To see this (and expand on a remark in [**GGP**]), denote the *ED1E* dynamics by ξ_t and *E0DE* by η_t. For an arbitrary seed A (with leftmost 1 at 0), let its complement A^c have all states of A flipped, while A^n has states flipped only on $[-1, s_g(A) + 1]$. Note that A^c is not a seed and A^n is, albeit with a leftmost 1 at -1. It is easy to prove by induction that for any initial seed A_0, $(\xi_t^{A_0})^c = \eta_t^{A_0^c}$, and that $\eta_t^{A_0^c} = \eta_t^{A_0^n}$ on $[-1, s_g(\eta_t)]$. As $s_g(\eta_t) = s_g(\xi_t) + 1$, ξ_t thus replicates 0's on a background of 1's and cannot have an RPS.

- *EEDE* and *E0D1*: These two are a mirror pair, so it is enough to show that *E0D1* is a replicator. Let $\mathcal{R}$ be the set of configurations which have all 1's isolated and all 0-blocks odd. Clearly, *E0D1* agrees with *EEDD* on configurations in $\mathcal{R}$. Note also that the 1 at the origin is isolated from time 1 on. For $t \geq 1$, let R_t be the location, underlined in what follows, of the leftmost 1 that is followed either by another 1 or by an even 0-block. Then the configuration at R_t cannot be $0\underline{1}01$, so it is one of $0\underline{1}00$, $0\underline{1}10$, or $0\underline{1}11$. In each of these three cases, the state of ξ_{t+1} is either $\underline{0}01$ or $\underline{1}01$, and $R_{t+1} \geq R_t + 2$. It follows that $2t - s_{\mathcal{R}}(\xi_t)$ eventually stabilizes and replication follows.

- *EED0* and *E0DD*: Again, these are a mirror pair, and we consider *E0DD* first. The proof is the same as for *E0D1* up to the case $0\underline{1}11$, which leads to either $\underline{0}00$ or $\underline{1}00$. The last possibility may give $R_{t+1} = R_t$, but then we have to have $R_{t+2} \geq R_t + 2$. This is enough to prove that $v_{\mathcal{R}} \geq 1$ and that there are no RPS, but the lemma below in fact establishes replication.

Lemma 3.1. *Assume the EED0 rule. For every initial seed A_0, the configuration ξ_t contains no 111's for large enough t.*

Proof. Let $\mathcal{R}'$ comprise configurations in which there is no 1, followed by an even number $k \geq 0$ of 0's, and then followed by 11. If there is such a 1, let R'_t be the location of the leftmost one, again underlined in the sequel.

When the configuration at R'_t (or elsewhere) is $\underline{1}11$, the only possibility at time $t-1$ is $10\underline{0}11$, as is easy to check. This shows that, from time 1 on, there are no 1-blocks of length 4 or larger, and so we will assume from now on that all configurations have this property. Further, if the configuration at R'_t is $\underline{1}[k$ 0's$]11$, for some $k \geq 2$, then we have two possibilities at the preceding time. One is simply $10\underline{0}[k$ 0's$]11$, and the other is $01\underline{1}1[(k-2)$ 1's$]001$. This follows because the 11 can only result from 0011 or 1001, and because $0{*}1 \mapsto 1$. Thus $\mathcal{R}'$ is invariant and R'_t increases in t.

We claim that $v_{\mathcal{R}'} \geq 3/2$. Together with the $v_{\mathcal{R}} \geq 1$ estimate for *E0DD*, this will end the proof. We will show that either $R'_{t+1} \geq R'_t + 2$ or $R'_{t+2} \geq R'_t + 3$, which could be easily checked by computer, but for once we give full details by hand. See Fig. 4 for an illustrative example.

To demonstrate the claim, observe that $R'_{t+1} = R'_t + 2$ unless the configuration at R'_t is $\underline{1}11$, so this is the only case we need to analyze further. There are four possibilities for the three states to the right: $\underline{1}11000$, $\underline{1}11001$, $\underline{1}11010$, and $\underline{1}11011$. To the left of $\underline{1}$, we can either have 10 or at least three 0's; as ${*}10 \mapsto 0$, the four cases update, in order, to $0\underline{*}10010$, $0\underline{*}10011$, $0\underline{*}10000$, and $0\underline{*}10001$. In the first case $R'_{t+1} \geq R'_t + 4$, and in the fourth $R'_{t+1} \geq R'_t + 5$. The second and third case only have $R'_{t+1} \geq R_t + 1$, but in another time step they update, respectively, to $\underline{*}10111$ with $R'_{t+2} = R'_t + 3$, and to $\underline{*}10100$ with $R'_{t+2} \geq R'_t + 3$. This establishes the desired lower bound on $v_{\mathcal{R}'}$ and ends the proof. □

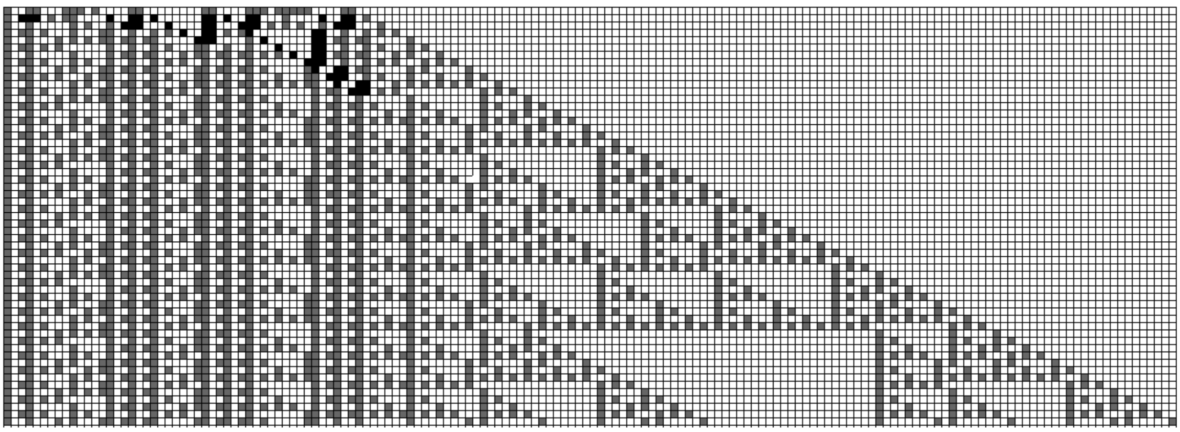

Fig. 4. *EED0* started from a short random seed. The 1's that are part of a configuration 1[even number of 0's]11 are highlighted after the initial time and disappear by time 12.

- *EE1D*: First assume that the initial configuration has 0's on all odd integers (recall that we assume a 1 at the origin). Any such configuration is static and clearly cannot be an RPS. To establish that *EE1D* has no RPS, we assume that a seed has at least one 1 at an odd integer and prove that it is a replicator. To this end, observe that a 11 is eventually created and let L_t be the leftmost non-isolated 1. As 011 leads to ${*}11$ directly below, L_t

decreases, and thus it eventually stabilizes (and all 1's to the left of L_t, if any, are on even integers). We will, without loss of generality, assume that $L_0 = 0$.

Now let $\mathcal{R}$ be the set of configurations that contain no isolated 1. It is clear that *EE1D* agrees with Jen's replicator *E11D* on $\mathcal{R}$. Let R_t be the (underlined) location of the rightmost non-isolated 1 to the left of the site where the corresponding finite automaton $\mathcal{F}$ stops. As $1\underline{1}0*$ updates to $**\underline{*}11$, $R_{t+1} \geq R_t + 2$, $2t - s_{\mathcal{R}}(\xi_t)$ eventually stabilizes and replication follows. See Fig. 3 for an example.

When an edge CA has a final Stage I tree, one could see whether there are RPS simply by inspection, that is, by running the R-algorithm. We will indeed do so, but first we state a simple condition that enables us to preclude RPS in advance.

Proposition 3.2. *An edge CA with no D-inducers and no 1-inducers has no RPS.*

Proof. For such a CA, $**0 \mapsto 0$, so a seed cannot generate a 1 to the right of its rightmost initial 1. Thus the ($\pi = 1$) configuration of all 0's is the only candidate for an RPS, but all 0's is not an RPS for any edge CA (as the 0 configuration is not a decider, by (1.1-1.2)). □

This condition eliminates RPS from an additional 7 range 2 edge CA:

- *E000*, *E00E*, *E0E0*, *E0EE*, *EE00*, *EE0E*, and *EEE0*.

More generally, a edge CA of bounded growth also cannot have an RPS; we investigate 13 such cases in Section 7, but it turns out that every one of them admits an easier argument given in this section: one (*EDEE*) has no deciders, and the others (12 cases in the first item below) have final Stage I trees.

The R-algorithm applied to the final Stage I tree eliminates RPS for the following 20 rules (parenthetically, we give corresponding π_0 and the number ν of vertices):

- *E00D* ($\pi_0 = 1$, $\nu = 5$), *E1EE* ($\pi_0 = 1$, $\nu = 8$), *E1E0* ($\pi_0 = 1$, $\nu = 7$), *EDE0* ($\pi_0 = 1$, $\nu = 8$), *E0ED* ($\pi_0 = 1$, $\nu = 6$), *E0E1* ($\pi_0 = 1$, $\nu = 6$), *E001* ($\pi_0 = 1$, $\nu = 5$), *E10E* ($\pi_0 = 1$, $\nu = 7$), *E100* ($\pi_0 = 1$, $\nu = 6$), *ED0E* ($\pi_0 = 1$, $\nu = 7$), *ED00* ($\pi_0 = 1$, $\nu = 7$), *EE0D* ($\pi_0 = 2$, $\nu = 11$);
- *EE1E* ($\pi_0 = 1$, $\nu = 6$), *EE10* ($\pi_0 = 1$, $\nu = 5$), *E11E* ($\pi_0 = 1$, $\nu = 5$);
- *ED10* ($\pi_0 = 2$, $\nu = 6$), *EDD1* ($\pi_0 = 2$, $\nu = 8$), *E1DE* ($\pi_0 = 2$, $\nu = 8$).

For instance, *EE10* has the final Stage I tree in which the root 1 branches to 2 and 3, and then $2 \to 1$, and $3 \to 2$; for *ED10*, the root 11 has two succession lines: $11 \to 23 \to 03 \to 02 \to 11$ and $11 \to 32$; for *EDD1*, see Fig. 2. In each of these three cases, there are infinitely many distinct periodic solutions, but no RPS.

This brings our count of $r = 2$ edge CA without RPS to 41. The one still left is *E1D0*, the famous *Rule 30* [**Wol3**, **Jen1**], which is used to generate pseudorandom numbers. We now introduce the key concept that handles this case.

An edge CA is ***left permutative*** if the map $cA \mapsto d$ is bijective for every configuration A of length r.

Proposition 3.3. *A left permutative edge CA has no periodic solutions, thus no RPS.*

Proof. Assume there is a periodic solution for a time period π and consider the (complete) label tree at that period. By left permutativity, every label in this tree has a unique predecessor label. As a periodic solution exists, there is a label λ, say at a node x, which is a repetition of a label on its ancestral line, say at a node x'. The same is true for predecessors of x and x', and by iteration we may assume that x' is the root, which has label 1. The unique predecessor of this label is the label 0, which is in turn its own unique predecessor, and so all labels on the ancestral line of x are 0 labels, a contradiction. □

It easily follows from Proposition 3.4 that an additive edge CA cannot have a periodic solution. By dramatic contrast, doubly periodic configurations are ***dense*** for a left permutative CA, in the sense that any finite configuration can be made a part of the link [**BK**].

- *E1D0* (*30*): It is very easy to show directly that this CA is left permutative [**Wol3**, **Jen1**], but a handy tool is its polynomial representation (see [**BL**] and references therein): modulo 2, $f(ca_1a_2) = c + a_1 + a_2 + a_1a_2$, and so $c = f(ca_1a_2) + a_1 + a_2 + a_1a_2$. Lack of RPS thus follows from Proposition 3.4.

4 Range 2 edge CA: existence and uniqueness of RPS

There are 22 range 2 edge CA for which the previous section does not disprove existence of RPS. In this section we will demonstrate that indeed all of these have at least one RPS: 19 rules a unique one, and 3 rules multiple ones. Let us begin with a general uniqueness criterion, valid for all r, based on analysis of the final Stage I tree.

Proposition 4.1. *If there is a final Stage I tree, then there are finitely many RPS, possibly none. When there is a final Stage I tree and a unique RPS, it is globally attracting if and only if there are no other periodic solutions.*

Proof. For the first claim, observe that the R-algorithm produces at most one RPS per decider leaf. Thus for any fixed π there are finitely many RPS, although there may be infinitely many periodic solutions. Clearly, existence of other periodic solutions precludes global attractiveness of an RPS. Conversely, assume that there are no other periodic solutions. Every path from the root following the forward and backward edges must eventually visit the same node twice. The loop part of this path must include a label from the RPS, or else another periodic solution would be constructed. If one of its label is reached on any such path, the RPS is globally attracting. □

The simplest situation is when an edge CA is ***filling***: it has a unique globally attracting RPS with link 1. That is, apart from an initial handle (which may have a higher period than 1), the half-lattice gets progressively filled by 1's for any semi-infinite initial state.

For most $r = 2$ rules we inspect the label tree [**Gra**] and appeal to Proposition 4.1; for some we give simple specialized arguments.

Three rules are based on the logical `or` operation:

- *E1E1*, *E111*, and *EE11* map $a_1a_2a_3$, respectively, to a_2 `or` a_3, a_1 `or` a_2 `or` a_3, and a_1 `or` a_3. All three *solidify*, i.e., occupied sites remain occupied forever. It is easy to see that the first two CA are filling. The last, *EE11* is *reducible*, i.e., the dynamics evolve as two independent CA, one on the even integers, the other on the odd ones, and each equivalent to *E1E1* after suitable relabeling of the sublattice. Again, "all 1's" is the unique RPS, but in this case it is not globally attracting since 1010 . . . is a static periodic solution. Indeed, any seed concentrated on the even integers leads to that solution.

There are 7 other rules with a $\pi_0 = 1$ final tree and a unique RPS that is globally attracting:

- *E010*, *E01D*, *E01E*, *E011*, *E101*, *E1D1*, *E110*: For all the rules whose names begin with *E01*, the final Stage I tree ($\pi_0 = 1$) is obviously $1 \to 2 \to 1$; so there is a unique periodic solution with signature 01 which is an RPS, thus globally attracting. For *E101*, the final Stage I tree ($\pi_0 = 1$) is $1 \to 3 \to 3$, thus this rule is filling. *E1D1* and *E110* have final Stage I trees ($\pi_0 = 1$) with 3 and 4 vertices, respectively; the first is filling and the second has a globally attracting RPS with signature 011.

In addition, there are 6 rules with a $\pi_0 > 1$ final trees and a globally attracting RPS:

- *ED0D*, *E10D*, *ED01*: The final Stage I trees ($\pi_0 = 2$) have 5, 6, and 5 vertices, respectively and a unique periodic solution, the 'checkerboard," with $\pi = \sigma = 2$ and signature 01, which is an RPS and hence globally attracting.
- *EDE1*, *ED11*: These rules have the same final Stage I tree ($\pi_0 = 2$) with 5 vertices and a unique periodic solution with signature link 0111 (which updates to 1101), which is an RPS, thus globally attracting.
- *ED1D*: There is a unique RPS, generated by 1110, with $\pi_0 = 4$, $\sigma = 12$, and signature 000110101111. As one can see from Fig. 2, the final tree has $\pi_0 = 4$, 9 vertices, and no other periodic solutions, so this RPS is globally attracting.

Each of the next 3 rules has "all 1's" as its unique RPS, which is not globally attracting, so none of these CA are filling:

- *EEE1*, *EE01*: The final label trees have $\pi_0 = 1$, with 8 and 7 vertices, respectively. By inspection, it is easy to see that "all 1's" is the unique RPS for these two rules. Also, any initial configuration with no 11 pair is fixed for *EEE1*, and any with all 1's separated by at least two 0's is fixed for *EE01*. Thus "all 1's" cannot be globally attracting.

- *EED1*: If the initial 1's are all at even sites, then the dynamics agrees with the additive rule *EEDD*. Otherwise, it is easy to check that a 11 pair is generated eventually, resulting in every site to the right of the pair eventually becoming a 1 (since $*11 \mapsto 1$ and $11* \mapsto 1$). By Proposition 3.4, the additive configurations with only even sites occupied have no periodic solutions, thus no RPS. We conclude that the RPS is unique, but not globally attracting. Also note that there is no final tree by Proposition 2.3.

Finally, there are 3 edge CA with multiple RPS:

- *EDD0* (*22*): The edge analysis of the present paper was developed in [**GG3**] to study this *Exactly 1* rule. In our large collection of RPS, there is one with $\pi = 8$, $\sigma = 104$, and shortest known handle of 35 sites, which will be studied further in the next section. It turns out that there are no RPS with $\pi = 16$, exactly 26 with $\pi = 32$, and many more with $\pi = 64$. See [**GG3**] for an extensive account.
- *E1ED* (*110*): The R-algorithm determines that there are no RPS for $\pi \leq 16$, but there are exactly three with $\pi = 32$: one with $\sigma = 16$ (generated by the seed 11[6 0's][6 1's]000[8 1's]011010111011), and two with $\sigma = 160$ (generated by a single 1, and 111[5 0's]11). (See Fig. 1.) Higher period RPS are also observed from small seeds; for instance, an example with $\pi = 64$ and $\sigma = 32$ emerges from 1101001000111.

Problem 1. Do *EDD0* and *E1ED* admit infinitely many RPS?

- *E1DD* (*94*): This is an example with provably infinitely many RPS, and we give details below.

We begin our analysis of *E1DD* with the observation that the doubly infinite configurations built from words 10 and 110 in any order are translated by 1 at every time step. It follows that any such periodic configuration with spatial period σ is also temporally periodic, with $\pi = \sigma$. These are not the only doubly periodic solutions (or even the only periodic solutions), but are the only ones generated by RPS.

Theorem 1. *Links for E1DD RPS are exactly configurations with a 10's and b 110's, of length $\sigma = 2a + 3b$ that is a power of 2.*

Proof. Represent a link of the desired type as a concatenation of a words 10 and b words 110. Reverse the order of these $a + b$ words, then replace each 10 with 12 and each 110 with 312 to create a label λ_0 of the tile. Each of the $\sigma = 2a + 3b$ labels is then a rotation of λ_0. Note that the RPS property is satisfied and that we may assume that λ_0 begins with 12. Create another label λ_3 by the following recipe: a 12 that precedes (in the cyclic order) 312 is replaced by 30; a 12 that precedes 12 is replaced by 33; a 312 that precedes 312 is replaced by 330; and a 312 that precedes 12 is replaced by 333. The number of D-inducers in λ_3 is $\sigma - b$, thus even; let λ_2 be its 1-successor. It is easy to check that 30, 33, 330, 333 in λ_3 result, respectively, in 30, 32, 230, 232 in λ_2. The number of D-inducers in λ_2 is again even, so we choose λ_1 to be its 0-successor. Now, λ_1 has 21 (resp. 121) in place of every 12 (resp. 312) in λ_0, and its unique successor is λ_0.

The key fact that *E1DD* coincides with the additive CA *EEDD* on configurations consisting entirely of even blocks is easy to check. It follows from the reverse label dynamics method of Lemma 5.1 in [**GG3**] that the label λ_3 is in the Stage I tree for $\pi = \sigma$, thus so is λ_0 and each of the claimed links indeed appears as part of an RPS.

We now proceed to proving that the described links are the only ones possible. We let $\mathcal{R}$ be the set of "additive" configurations that begin with 11 and have only even blocks. As 11 appears at the origin at time 1, it makes sense to consider the boundary of such a configuration, that is, the site $L_t = s_{\mathcal{R}}(\xi_t)$ at which the corresponding finite automaton $\mathcal{F}$ stops. That site will be underlined in the rest of this proof, which is a careful analysis of what might happen at this interface.

Our first step is to examine conditions under which 10$\underline{1}$ appears. At any time when this configuration is not on the interface, we have one of the following four possibilities:

(i) 11[k 0's]$\underline{1}$, $k \geq 3$ odd;

(ii) 11[k 0's]1$\underline{0}$01, $k \geq 2$ even;

(iii) 11[k 0's]1$\underline{0}$[ℓ 0's]1, $k \geq 2$ even, $2 \leq \ell \leq \infty$ ($\ell = \infty$ means the 0's extend indefinitely); or

(iv) 00[k 1's]$\underline{0}$, $k \geq 3$ odd.

Case (i) leads to 10$\underline{1}$ in $(k-1)/2$ steps, case (iv) updates to either 10$\underline{1}$ (when $k = 3$) or to case (i) in a single step, and (iii) updates to (iv), with $k = 3$, in a single step. Finally, (ii) results in advancement of L_t by at least 4. Therefore, unless L_t advances by at least 4 in each time step, 10$\underline{1}$ must appear. The latter must in particular happen when an RPS appears, as in that L_t must stabilize — we already know that an additive CA cannot have an RPS.

Now assume that 10$\underline{1}$ appears at time t_0, with $\underline{1}$ at the position x_0. We assume from now on that $t \geq t_0$, and denote by R_t the position $x_0 + t - t_0$. Observe that the configuration in $[R_t - 2, R_t]$ is 101.

We will now show that one of the below nine configurations (I)–(IX) occurs on the interface for $t \geq t_0$. This will be simply proved by computing possible one-step updates and checking that every result is on the list (I)–(IX). In each case, we specify the numerals of the updated configurations and indicate whether L_t advances. We specify all values which are determined by the given information. (As in other proofs, we underline L_t, and not L_{t+1}, at time $t + 1$.)

(I) 110$\underline{1}$0 updates to 001$\underline{0}$1 (V) or to 111$\underline{0}$1 (VI), no advance;

(II) 110$\underline{1}$10 updates to 001$\underline{0}$11 (V) or to 111$\underline{0}$11 (VI), no advance;

(III) 110$\underline{1}$110 updates to 001$\underline{0}$101 (V) or to 111$\underline{0}$101 (VI), no advance;

(IV) 110$\underline{1}$111 updates to 001$\underline{0}$100 (V) or to 111$\underline{0}$100 (VI), no advance;

(V) 001$\underline{0}$1 updates to 001$\underline{1}$01 or to 111$\underline{1}$01, both one of (I)–(IV), advance by exactly 2;

(VI) 111$\underline{0}$1 updates to 000$\underline{1}$01 (VII) or to 110$\underline{1}$01 (I), no advance;

(VII) 000$\underline{1}$01 updates to 000$\underline{1}$101 (VIII) or to 110$\underline{1}$101 (II), no advance;

(VIII) 000$\underline{1}$101 updates to 000$\underline{1}$1101 (IX) or to 110$\underline{1}$1101 (III), no advance;

(IX) 000$\underline{1}$1101 updates to 000$\underline{1}$10101 (VIII) or to 110$\underline{1}$10101 (II), no advance.

As L_t must stabilize, we could now eliminate (V), and consequently also (IV), for the purposes of this proof. However, in Section 8 we will refer to the general version of the claim we are about to state, so we keep the two on the list. Observe also that (V) cannot occur twice in a row, so $L_t \leq R_t$ and $v_{\mathcal{R}} \leq 1$.

Call a configuration in an interval $[a, b]$ *good* if it contains neither a 00 interval nor a 111 interval. We interpret $[a, b] = \emptyset$ when $b < a$. Our last claim is that $[L_t + 1, R_t]$ is good at time t.

We prove this claim by induction. Clearly, we can assume that $L_t < R_t$ or else the induction step is a triviality (as it is at t_0). Now let L'_t be any position at time t so that the configuration in $[L'_t, R_t]$ begins with 10 and is good. Then, the configuration in $[L'_t + 1, R_t + 1]$ at time $t + 1$ has the same properties. So, we need to check that the configuration at time $t + 1$ in $[L_t + 1, L'_t]$ is good and does not end with 11.

The induction hypothesis implies that case (IV) must have $L_t = R_t$, so it need not be considered further. The hypothesis also implies that in cases (I) and (II) we may take $L'_t = L_t - 2$, in cases (V) and (VI) $L'_t = L_t - 1$, in case (VII) $L'_t = L_t$, and in case (VIII) $L'_t = L_t + 1$; the induction step is now trivial in these cases as well. For (III) and (IX) we may let $L'_t = L_t + 2$, and the updates show that configuration at time $t + 1$ in $[L_t + 1, L'_t]$ is 10 in both of these cases. This establishes the claim. Therefore, the configurations to the right of L_t are of the desired form and the proof is concluded. □

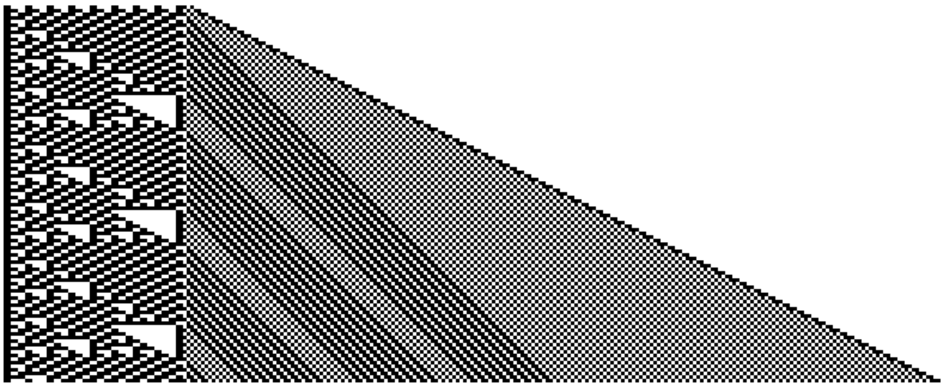

Fig. 5. Evolution of *E1DD* from [8 110000's]11101.

We should mention that we do not have the explicit form for the handle for all possible links, so we give here without proof simple examples of two infinite families: for $k \geq 4$, the link $[(2^{k-1} - 3)$ 01's]011011 is generated by its handle 11[2^k 0's][$(2^{k-2} - 2)$ 1100's]11011; and for $k \geq 1$, the link [2^k 011's][2^{k-1} 01's] is generated by its handle [2^k 110000's]11101 (see Fig. 5).

5 Range 2 edge CA: expansion velocities

The expansion velocity v for an RPS can often be determined by exhaustive checking of all possible configurations at the right boundary of the expanding periodic state, the basis of our methods in this section.

Assume $H+L$ is an RPS with periods π and σ, and that h is the length of H. Fix an integer $a \geq 1$. Assume A is one of the 2^a configurations of length a, and compare two dynamics ξ_t and ξ_t' on $[0, h+\sigma+a-1]$, where ξ_0 is given by HLA and ξ_0' by HL^∞. Suppose that $T(a, A)$ is the smallest time t for which $\xi_t = \xi_t'$ on $[0, h+\sigma+a-1]$ and ξ_t has the last σ states equal to L on the same interval. Then let

$$v_a = \frac{a}{\max\{T(a, A) : A\}}.$$

The following result follows easily.

Proposition 5.1. *For every a, $v \geq v_a$.*

It is common that $T(a, A) = \infty$ unless a satisfies a divisibility condition, resulting in the useless bound $v_a = 0$. In addition, we emphasize that v_a depends on the choice of the link L. Thus, care in the choice of a and L is necessary in most interesting cases.

A priori it might seem that good upper bounds should be easier, since $v \leq v(A_0)$ for any environment A_0. However, rigorous results may be elusive for most test environments. Therefore, we introduce two "finite" conditions suitable for computer search and verification. Both methods depend on configurations that stop the RPS expansion and repeat regularly.

For integers a and τ, we call a configuration B of length b a ***universal*** (a,τ)***-blocker*** if its first state is different from the first state of L and in addition B has the following property. Choose an arbitrary configuration A of length a and compare two evolutions on $[0, h+\sigma+b+a-1]$: ξ_t starts from $HLBA$, while ξ_t' starts from HL^∞. Then ξ_τ and ξ_τ' agree on $[0, h+\sigma+a-1]$, and the last $\sigma+b$ states of ξ_τ on $[0, h+\sigma+b+a-1]$ are equal to LB.

While finding a universal blocker is challenging, the upper bound it gives is completely independent of the rest of the environment, making this method suitable for an RPS that competes against a chaotic evolution [**GG3**]. Our second method, which works best for the glider rules [**Coo**, **SCJ**], is based on an RPS expanding into a periodic environment while maintaining a bounded interface.

Assume C is a configuration of length c that is a link for a doubly periodic solution — or equivalently, running ξ_t on c sites with periodic boundary, starting from C, eventually reproduces C. Assume also that B is a (possibly empty) configuration such that the configuration in the first σ states of BC^∞ differs (in at least one state) from L. Finally, assume that the following holds for a translation number $a > 0$ and a time $\tau > 0$. When ξ_0 and ξ_0' are respectively given on $[0, \infty)$ by $HLBC^\infty$ and HL^∞, $\xi_\tau = \xi_\tau'$ on $[0, h+a-1]$ and ξ_τ agrees with LBC^∞ on $[h+a, \infty)$. Then we call the pair (B, C) a ***periodic*** (a,τ)***-blocker***.

The next proposition follows immediately from the definitions.

Proposition 5.2. *If an (a,τ)-blocker of either type exists, then $v \leq a/\tau$.*

Using Propositions 5.1 and 5.2 we are able to determine exact expansion velocities for at least one RPS for each of the 22 range 2 edge CA identified in Section 4 that admit an RPS. Dealing quickly with the simplest cases, we provide additional details for the more interesting ones. The link L used to compute v_a and any other essential information are given in parentheses. Blockers and the corresponding bounds are independent of the handles, which we therefore do not provide (see [Gra] for explicit examples of handles for all rules discussed in this section).

The unique RPS for 4 edge CA have expansion velocities equal to the speed of light, $v = 2$, as a trivial check shows that $v_1 = 2$:

- *EE11* ($L = 1$); *E111* ($L = 1$); *E011* ($L = 01$); *ED11* ($L = 0111$).

There are 8 more rules with a unique RPS and $v = 1$, established by $100 \mapsto 0$, the assignment that makes any expansion velocity for any seed at most 1, and by $v_1 = 1$ except in the last case listed here:

- *E1E1* ($L = 1$); *ED0D* ($L = 01$); *E10D* ($L = 01$); *ED01* ($L = 01$); *E101* ($L = 1$); *EE01* ($L = 1$); *EEE1* ($L = 1$); *EDE1* ($L = 0111$, $v_4 = 1$).

An additional 5 rules have a unique RPS and $v = 1$, but need slightly more involved verification:

- *E010* ($L = 01$, $v_2 = 1$, with $C = L$, $B = 0$ a periodic $(2, 2)$-blocker);
- *E01E* (the same as *E010* except that $B = 1$);
- *E110* ($L = 011$, $v_3 = 1$, with $C = L$, $B = 000$ a periodic $(3, 3)$-blocker);
- *E1D1* and *EED1* ($L = 1$, $v_1 = 1$, with $C = L$, $B = 0$ a periodic $(1, 1)$-blocker).

The last 2 rules with a unique RPS have unusual velocities:

- *E01D* has $v = 4/3$. This is established by $L = 10$, $v_4 = 4/3$, and the periodic $(4, 3)$-blocker obtained by $B = \emptyset$, $C = 11100$.
- *ED1D* has $v = 3/4$. If we choose $L = 110001101011$, a computer computation gives $v_{12} = 3/4$. Moreover, a periodic $(12, 16)$-blocker is given by $B = \emptyset$, $C = 1001$ (see Fig. 6).

Finally, we address the three rules with multiple RPS, dealing only with the simplest examples in two cases.

- *E1DD* has $v = 1$ for every RPS constructed in Section 4. To see this, recall that the rule translates a copy of 101 by 1 in a single time step independently of the state of the rest of the lattice. Assume a link L ends in 101. Then we have $v_1 = 1$. Furthermore, one can get a universal $(1, 1)$-blocker by appending after L either 01 (if L begins with a 1) or 101 (if L begins with a 0).

- *E1ED* (*110*): despite its long handles, we consider the simplest RPS to be the one with $\pi = 32$ and $\sigma = 16$. This has $v = 1/2$. To check, we take $L = 1111101011100110$ and verify that $v_2 = 1/2$. Then, we verify that a periodic $(2, 4)$-blocker is constructed by the spatial period 8 environment with $B = 0$ and $C = 01111101$. (Although the verifications can be done by hand in this case, L and C were obtained by a computer search that found no C of length less than 8.)

- *EDD0* (*Exactly 1*): the simplest RPS has $\pi = 8$, $\sigma = 104$ [**GG3**, **Gri**] and commonly emerges from small seeds, e.g., from 1011011. In this case all possible links are rotations of one another and we choose for L the rotation that ends with [6 1's] (the largest possible such block). An exhaustive computer check then shows that $v_{13} = 13/43$. Moreover, this configuration has a universal $(13, 43)$-blocker, given by $B = 11$, as another computer verification of all 2^{13} possible configurations to the immediate right of HLB demonstrates. (We have been unable to find any other universal blockers for this rule.) We have thus proved that $v = 13/43$, as first conjectured by Dean Hickerson [**Hic2**]. Clearly, $v(A_0) = v$ whenever the described blocker appears, and we know of no seed from which the RPS emerges but the blocker does not; hence the open problem stated below. To add mystery to this puzzle we remark that there exist RPS which have a matching link [**GG3**]; thus there exists a seed A_0 with $v(A_0) = 2$. A nontrivial universal blocker therefore does *not* emerge from every seed that generates such an RPS.

Problem 2. For *EDD0*, does *every* seed that produces the $\pi = 8$ RPS also develop the universal blocker described above?

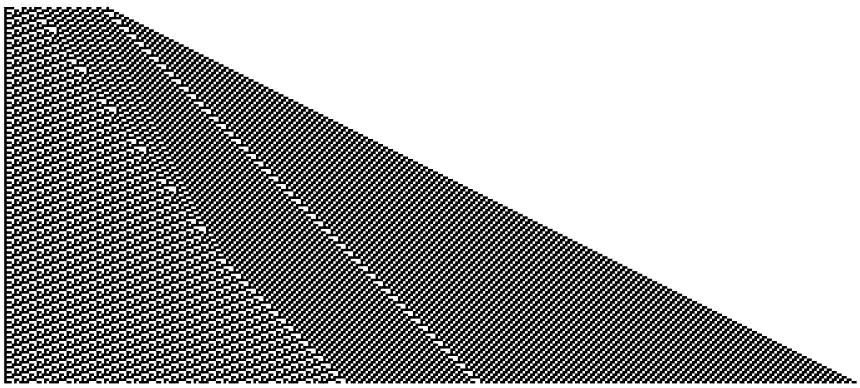

Fig. 6. *ED1D* started from a portion of its periodic blocker described in the text. Observe that, as required by Theorem 2, the RPS expands at velocity 1 once the "regularization wave" reaches the space-time line of slope 1.

6 Range 2 edge CA: expansion velocities in seed environments

This section is devoted to rules whose RPS have $v(A_0) > v$ for all seeds A_0. We begin with our best example of regularity expansion technique.

Theorem 2. *Assume the ED1D rule. For its globally attractive RPS, and any seed A_0, $v(A_0) = 1 > 3/4 = v$.*

It is well-known that *ED1D* supports a variety of gliders with quite complex interactions [**SCJ**]. One could perhaps prove the above theorem by proving that the glider phase eventually emerges and then carefully cataloging all possible glider interactions. We will see that such detailed knowledge is not necessary for our purposes.

We find it more convenient to consider the mirror rule *E110*, so this is our CA for a series of preparatory lemmas before the proof of Theorem 2 below. We declare that a configuration belongs to $\mathcal{R}$ if it begins with a 1-block of length at least 2, every 1-block has length 1, 2, 4, or 5; every 0-block has length 1, 2, 3, or 4; and satisfies the following transition rules on blocks from left to right. Call an isolated 0 a $0B$-block if it follows an isolated 1, and a $0A$-block otherwise. Then each of 11, 1111, and 11111 blocks may be followed by any of $0A$, 000, or 0000 and vice versa; each of $0A$, 000, and 0000 may also be followed by an isolated 1; an isolated 1 may be followed by $0B$ or 00; and each of $0B$ or 00 may be followed by any of 11, 1111, and 11111. The resulting finite automaton $\mathcal{F}$ has 12 $\mathcal{F}$-states and transitions

$$\begin{aligned}
&1A_1 \to 1A_2 \to 1A_3 \to 1A_4 \to 1A_5\\
&1A_2 \to 0A_1,\ 1A_4 \to 0A_1,\ 1A_5 \to 0A_1\\
&0A_1 \to 0A_2 \to 0A_3 \to 0A_4\\
&0A_1 \to 1AB_1,\ 0A_3 \to 1AB_1,\ 0A_4 \to 1AB_1\\
&1AB_1 \to 0B_1,\ 1AB_1 \to 1A_2\\
&0B_1 \to 0B_2 \to 1A_1,\ 0B_1 \to 1A_1
\end{aligned}$$

Here, the first letter of an $\mathcal{F}$-state indicates the binary state in a configuration on which $\mathcal{F}$ is allowed a transition. For example, if $\mathcal{F}$ is at x in the $\mathcal{F}$-state $1A_1$ and the state at $x+1$ is 1, $\mathcal{F}$ makes the transition to $1A_2$ and moves to $x+1$, but if the state at $x+1$ is 0 it stops at $x+1$. If $\mathcal{F}$ is instead in the $\mathcal{F}$-state $1A_2$, it transitions to $1A_3$, resp. $0A_1$, if the state at $x+1$ is 1, resp. 0, and it always moves to $x+1$. Also, indices give the current position within a block.

We should also note that a the configuration 11 appears at positions 0 and 1 at time 1 at the latest, providing a starting point for the expansion of $\mathcal{R}$.

Lemma 6.1. *So defined, $\mathcal{R}$ is invariant.*

Proof. Assume $\xi_0 \in \mathcal{R}$ and fix an $x \geq 3$. Observe that we may extend ξ_0 indefinitely to the left by $\ldots 011011 0\underline{1}1 \ldots$ (where the underlined state is at the origin), without affecting subsequent evolution on nonnegative sites; thus we do not need to consider x's close to the origin separately.

Assume that $\mathcal{F}$ stops at x for ξ_1. Let $y \leq x-2$ be the largest y so that a pair 11 occurs at y and $y+1$. It is easy to check that $y \geq x-9$. We may replace all states to the left of y by 0's.

Now we may simply check all 2^{y-x-1} configurations in $[y, x]$ which start with 11 (with all 0's to the left): there must be one for which $\mathcal{F}$ does not stop before or at x for ξ_0, but stops at x for ξ_1. A straightforward exhaustive (computer) check demonstrates that no such configuration exists and this contradiction ends the proof. □

Next, we show that $\mathcal{R}$ has light speed 1.

Lemma 6.2. *Assume that A_0 and B_0 are both in $\mathcal{R}$ and agree on $[b, \infty)$. Then $\xi_t^{A_0}$ and $\xi_t^{B_0}$ agree on $[b+t+2, \infty)$ for all $t \geq 0$.*

Proof. We claim that for any interval $[x-6, x] \subset [0, \infty)$, which contains any possible segment of a configuration in $\mathcal{R}$, the state at x three time units later only depends on states in $[x-3, x]$. Analogously to the previous proof, we need to consider all $\mathcal{R}$-configurations in $[y, x]$, for $x-13 \leq y \leq x-6$, that begin with 11 and have only 0's to their left. This is again accomplished by an exhaustive check.

Now, $\xi_t^{A_0}$ and $\xi_t^{B_0}$ trivially agree on $[b+2, \infty)$ for $t=1$ and on $[b+4, \infty)$ for $t=2$, and by the claim on $[b+3, \infty)$ for $t=3$. The proof is concluded by iteration. □

We suspect, but we are unable to prove, that $v_{\mathcal{R}} = 3/2$. Fortunately, we already know that the *ED1D* RPS has expansion velocity 3/4, so we only need $v_{\mathcal{R}} > 5/4$ to apply Lemma 2.1.

Lemma 6.3. *The expansion velocity satisfies $v_{\mathcal{R}} > 1.26$.*

Proof. Assume that $\mathcal{F}$ does not stop on ξ_0 at or before $b \geq 2$. Let B be ξ_0 restricted to $(-\infty, b]$. We will say that an integer a *works* for B if there exists an integer $a \geq 1$ such that, for any configuration A of length a the following is true: starting from BA, there exist a time $t_0 \geq 0$ and an x_0, $b + \frac{5}{4}t_0 < x_0 \leq b+a$, so that $\mathcal{F}$ does not stop on ξ_{t_0} at or before x_0. (Note that, for algorithmic efficiency, once one find a t_0 and x_0 for A, one does not need to check any configuration A' which agrees with A up to x_0.) If a fixed a works for every configuration B, then it is easy to see by iteration that $v_{\mathcal{R}} > 1.25$.

Assume that the initial set for ξ_t is BA as above. Assume that there is a $z \leq y$ with 11 at z and $z+1$. Let R_t be the first site at which $\mathcal{F}$ stops on ξ_t. (For this purpose, we assume $\mathcal{F}$ automatically stops at $b+a+1$ if it gets there.) Assume that the *influence condition* holds for z, that is, $z+t+2 \leq R_t - 2$ for $t \leq t_0$. Then, by Lemma 6.2, one may replace all sites to the left by 0's and obtain the same x_0 and t_0.

Assume that we find a and b_0 so that a works for all $\mathcal{R}$-configurations B with $b \in [b_0, b_0+7]$, and in addition the influence condition with $z=0$ can be verified for all relevant B and A. Then, by the same argument as in Lemma 6.1 a works for all configurations B, and the proof is concluded.

Verifying (or falsifying) the above property for a fixed a and b_0 is a finite task, albeit possibly a very time-consuming one. We were successful with $b_0 = 6$ and $a = 32$, close to the limit of feasible computation, and a computer implementation produced the explicit bound given in the statement. □

Proof. (Theorem 2) As $v = 3/4$ for the *ED1D* RPS (see Section 5), Lemmas 6.2, 6.3, and 2.1 imply that $v(A_0) \geq 1$ for any seed A_0. Conversely, the mirror CA *E110* has a globally attracting RPS with expansion velocity 1 (again, see Section 5) and therefore Lemma 2.2 implies that $v \leq 1$. □

We now address a much easier case.

Proposition 6.1. *For the unique E01D RPS, and any seed A_0, $v(A_0) = 2 > 4/3 = v$.*

Proof. Again, we initially prefer to deal with the mirror rule *EE10*. In this case, $\mathcal{R}$ comprises all final configurations given by the final Stage I tree with $\pi_0 = 1$; these have every 0 isolated and all 1-blocks of length 1 or 2. Such a configuration is fixed, and therefore it has light speed 0. Let R_t be the position, underlined in the sequel, at which the corresponding automaton stops for ξ_t. There are two possible configurations at R_t, $10\underline{0}$ and $11\underline{1}$, which update to $10\underline{1}$ and $11\underline{0}$, proving that $v_{\mathcal{R}} \geq 1$.

As it has $v = 4/3$, the *E01D* RPS therefore eventually expands into a $\mathcal{R}$-environment with light speed 0. By Lemma 2.1 it has the maximal possible expansion velocity 2. □

Next, we consider *EDD1*, a CA with no RPS, although one might be tempted to conclude otherwise by experimenting only with seeds. Namely, the periodic solution with the tile

0011
1110

behaves, in a seed environment, as if it were a globally attracting and fastest possible RPS.

Proposition 6.2. *Consider EDD1 edge CA. For every initial seed A_0, a handle for the above periodic solution emerges, and $v(A_0) = 2$.*

Proof. This time we need to consider both *EDD1* and its mirror rule *E1DE*. In each case, the set of regular configurations, $\mathcal{R}$ and $\mathcal{R}'$ respectively, will consist of final configurations determined by the respective final Stage I tree.

For *EDD1*, the final tree with $\pi_0 = 2$ is depicted in Fig. 2. It turns out that the resulting configurations in $\mathcal{R}$ have only blocks 0, 00, 1, 11 and 111, with left to right transitions, $0 \to 1$, $0 \to 111$, $00 \to 11$, $1 \to 0$, $1 \to 00$, $11 \to 0$, $11 \to 00$, and $111 \to 0$. Let R_t be the (underlined) position at which the corresponding finite automaton stops. There are six possibilities for the configuration at R_t: $101110\underline{0}$, $10111\underline{1}$, $100\underline{0}$, $1011\underline{0}$, $10011\underline{1}$, and $1001\underline{0}$. The second through fifth cases yield $R_{t+1} \geq R_t + 1$. The first case gives $R_{t+2} \geq R_t + 2$ regardless of the state of additional site to the right of $\underline{0}$. For the sixth case one needs to consider the four possibilities for states of two additional sites, and check that always $R_{t+3} \geq R_t + 3$. We conclude that $v_{\mathcal{R}} \geq 1$.

For *E1DE*, the final Stage I tree [**Gra**] establishes that configurations in its set $\mathcal{R}'$ are built by starting with intervals of 1's of length at least 2, followed by either 00 or 01, followed by another 1-interval of length at least 2, etc. Again, these are periodic with time period 2. Moreover, we claim that $v_{\mathcal{R}'} \geq 4/3$.

We begin by describing the resulting finite automaton $\mathcal{F}'$. It has five $\mathcal{F}'$-states and the following transitions (with the same naming convention as in the *ED1D* case above, and $1A_i$ the states when the automaton is within a requisite interval of 1's with length at least 2):

$$1A_1 \to 1A_2 \to 1A_2,\ 1A_2 \to 0_1$$
$$0_1 \to 0_2,\ 0_1 \to 1B_1$$
$$0_2 \to 1A_1,\ 1B_1 \to 1A_1$$

Assume again that $\mathcal{F}'$ stops at R_t at time t. Observe that $\xi_t(R_t) = 0$. Let $S_t < R_t$ be the rightmost position at which $\mathcal{F}'$ is in the $\mathcal{F}'$-state $1A_2$. Note that $S_t \geq R_t - 3$. The possible configurations at 6 sites starting with $S_t - 1$ are 110000, 110001, 110010, 110100, 110101, and 110110. For each of these, one can check that, on ξ_{t+3}, $\mathcal{F}'$ does not stop at or before $S_t + 4$, and furthermore, is in the $\mathcal{F}'$-state $1A_2$ at $S_t + 4$. Iteration gives the desired lower bound on $v_{\mathcal{R}'}$.

We now return to *EDD1*, and conclude that eventually the domains of expansion of $\mathcal{R}$ and $\mathcal{R}'$ intersect. It is easy to check that the intersection consists exactly of the claimed spatially and temporally periodic configuration. Finally, as $\mathcal{R}'$ is periodic and thus has light speed 0, Lemma 2.1 implies that $v_{\mathcal{R}} = 2$. □

7 Range 2 edge CA: bounded growth

An edge CA has ***bounded growth*** if, for every initial seed A_0, there exists an integer $K = K(A_0)$ so that the evolution never generates a 1 to the right of K.

We now discuss bounded growth for $r = 2$. A trivial necessary condition is that $100 \mapsto 0$, so we only need to consider 32 rules; of these, we know from Section 4 that 9 have RPS, one (*EDED*) is additive, and one (*EEED*) is proved not to have bounded growth in Section 8. An equally trivial sufficient condition for bounded growth is that there are no *D*-inducers and no *1*-inducers (see Proposition 3.3), which covers 8 cases. Below we prove bounded growth for the remaining 13 cases. For more challenging rules, we exploit the connection between regularity expansion and final Stage I trees: $\mathcal{R}$ will comprise exactly the final configurations. These are periodic (with periods 1 or 2 in these examples) and thus have zero speed of light. The strategy, then, is to show that the growth of the regular region outpaces the right boundary (which moves at velocity at most 1) and then causes the CA to stop growing. Complete descriptions of all Stage I trees are archived at [**Gra**]; see also Section 3.

The first four cases are handled by considering the configuration at the position $s_g(\xi_t)$ of the rightmost 1 of ξ_t, which will be underlined.

- *E00D*: As only 001 and 110 map to 1, all 1's at times $t \geq 1$ must be isolated. Moreover, $010 \mapsto 0$ and $100 \mapsto 0$, so the $s_g(\xi_t)$ cannot advance after the first update.
- *E1EE*, *E1E0*: By checking all cases, one finds that neither of these rules can have an isolated 1 at time 1. From that time on, the configuration at $s_g(\xi_t)$ must be $1\underline{1}$. As for both rules $1{*}0 \mapsto 0$, $s_g(\xi_t)$ thereafter cannot advance.

- *EDE0*: We have $s_g(\xi_{t+1}) \le s_g(\xi_t)$ unless the state at $s_g(\xi_t)$ is $0\underline{1}$. In this case the configuration at time $t+1$ is is $\underline{1}10$ and then $110 \mapsto 0$ and $*11 \mapsto 0$ imply that $s_g(\xi_{t+2}) \le s_g(\xi_t)$.
- *EDEE*: This is the only case among these 13 with no final Stage I tree. To see this, introduce the map F that flips all states on $[-1, \infty)$. Start *EDEE* (denoted by ξ_t) from $111\dots$ and *EEDD* (denoted by λ_t) from a single initially occupied site placed at -1; then $F(\xi_t) = \lambda_t$.

  To demonstrate bounded growth, note first that, again, $s_g(\xi_{t+1}) \le s_g(\xi_t)$ unless the configuration at $s_g(\xi_t)$ is $0\underline{1}$. In this case, let $s_g(\xi_t) - k$ be the rightmost position of a nonisolated 0 or a nonisolated 1 to the left of $s_g(\xi_t)$. In the case of a nonisolated 0, k is odd and it is easy to check that $s_g(\xi_{t+(k+3)/2}) = s_g(\xi_t)$, while in the other case k is even and $s_g(\xi_{t+(k+2)/2)}) = s_g(\xi_t)$.

For the remaining rules we use the following proposition.

Proposition 7.1. *Assume that an edge CA has final Stage I tree and $\mathcal{R}$ is the set of final configurations. If*

$$v_{\mathcal{R}} > v_g \tag{7.1}$$

then the CA has bounded growth.

Proof. Recall that π_0 is the temporal period for the final Stage I tree. It follows from (7.1) that there exists a time t_0 at which $s_{\mathcal{R}}(\xi_{t_0}) > s_g(\xi_{t_0}) + r \cdot \pi_0$. Assume $t \ge t_0$ for the rest of the argument. As ξ_t is periodic with period π_0 on $[0, s_{\mathcal{R}}(\xi_{t_0})]$, it has no 1 on $[s_g(\xi_{t_0}) + (\pi_0 - 1)r + 1, s_{\mathcal{R}}(\xi_{t_0})]$. As the last interval contains at least r sites, $\xi_t(x) \ne 1$ for $x > s_{\mathcal{R}}(\xi_{t_0})$. □

To verify (7.1), we will always use the trivial upper bound $v_g \le 1$ and prove $v_{\mathcal{R}} > 1$. We will also let $R_t = s_{\mathcal{R}}(\xi_t)$ be the (underlined) locations at which the corresponding finite automaton $\mathcal{F}$ stops.

- *E0ED*, *E0E1*, *E001*: The configurations in $\mathcal{R}$ are those that have all 1's isolated; by inspection of the final Stage I tree with $\pi_0 = 1$ (or a direct verification using $010 \mapsto 0$ and $100 \mapsto 0$), such configurations are fixed. All these rules have $011 \mapsto 0$, so $R_{t+1} \ge R_t + 2$ and $v_{\mathcal{R}} \ge 2$, verifying (7.1).
- *E10E*: The configurations in $\mathcal{R}$ are those that have all blocks of length at least 2, and again are fixed. Possible configurations at R_t are $001\underline{0}$ and $110\underline{1}$; the first updates to $1\underline{1}0$ and the second to $0\underline{0}1$. Therefore, $R_{t+1} \ge R_t + 2$, $v_{\mathcal{R}} \ge 2$, and the desired conclusion follows.
- *E100*: The definition of $\mathcal{R}$ is the same as for *E10E*, except that we add the requirement that 1-blocks have length exactly 2. The additional possibility at R_t, $0011\underline{1}$, updates to $11\underline{0}0$, the other two possible configuration behave the same as for *E10E*, and the same conclusion follows.

- *ED0E*, *ED00*: These rules have identical final trees with $\pi_0 = 2$ and 7 vertices. By inspection of the tree, it is easy to check that $\mathcal{R}$ consists of configurations with no 1-blocks or length greater than 2, which have every isolated 1 followed by at least 3 0's and a 11 followed by at least 2 0's.

  Observe that the configuration at R_t must have one of the following four forms: $0100\underline{1}$, $0010\underline{1}$, $0110\underline{1}$, or $0011\underline{1}$. For *ED00*, one can check two updates, in each of the four cases and for all four possible configurations on two sites to the right of $\underline{1}$, to see that $R_{t+2} \geq R_t + 3$ and hence $v_{\mathcal{R}} \geq 3/2$.

  The same is true for *ED0E*, except that the boundary configurations $0011\underline{1}00$ and $0011\underline{1}01$ do not yield $R_{t+2} \geq R_t + 3$. For each of these two, we check that either of the two states at $R_t + 3$ at time t leads to $R_{t+3} \geq 4$. Therefore, $v_{\mathcal{R}} \geq 4/3$, which suffices to establish (7.1).

- *EE0D*: The final tree has $\pi_0 = 2$ and 11 vertices (see Fig. 2). Here the description of configurations in $\mathcal{R}$ is somewhat more complicated: there are no 1-blocks of length greater than 3; isolated 1's are followed by at least 2 0's; 11 pairs are followed by 01 and then at least 2 0's; and 111 triples are followed by at least 3 0's.

  To show that $v_{\mathcal{R}} > 1$, observe that the configuration at R_t has one of the following 7 forms: $0010\underline{1}$, $00110\underline{0}$, $0011010\underline{1}$, $001101\underline{1}$, $0011100\underline{1}$, $001110\underline{1}$, or $00111\underline{1}$. Again, by computing the evolution from these configurations for up to 3 updates, with arbitrary states on $[R_t+1, R_t+3]$, one can verify that $R_{t+3} \geq R_t + 4$, except when the configuration is $00111\underline{1}10$ (which yields $R_{t+3} = R_t + 2$). In that case, however, one can consider all four possible configurations on $[R_t + 3, R_t + 4]$ at time t to check that $R_{t+4} \geq R_t + 5$. Hence $v_{\mathcal{R}} \geq 5/4$ and (7.1) follows.

Assume that an edge CA has a final Stage I tree. It is easy to see that a necessary condition for bounded growth is that each node on this graph is connected, through following a sequence of forward and backward edges, to a node with the 0 label. We have proved in this section that this condition is also sufficient for range 2 edge CA — something that fails to hold for larger ranges, as we now demonstrate by a simple range 3 counterexample. Consider the rule whose assignment has $a_1a_2a_3a_4 \mapsto 0$ exactly when either all four a_i are 0 or exactly one of a_1, a_2, a_3 is 1. This edge CA has a very simple final Stage I tree with $\pi_0 = 1$: $1 \to 2 \to 4 \to 0$, and the 0 branches into 0 and 1. Any seed with only isolated 1's fixates. However, any seed A_0 that contains a pair 11 has $v_g(A_0) = 2$ and, for the set of final configurations $\mathcal{R}$, $v_{\mathcal{R}}(A_0) = 5/3$. Thus the regular region may indeed lag behind the furthest occupied site and, at least when using our technique to establish bounded growth, a possibly tedious proof of (7.1) cannot be avoided. We also remark that equality in (7.1) does not suffice: the single change $1110 \mapsto 0$ in the described assignment yields an example without bounded growth, but with $v_{\mathcal{R}} = v_g = 2$.

8 Range 2 edge CA: asymptotic behavior from seeds

Our conclusions are summarized in the subsequent division of CA into four classes. As both members of a mirror pair have the same seed behavior, we consider them equivalent, which results in 56 range 2 rules. For the 50 rules in classes R, B and P, our results from previous sections rigorously describe the asymptotic behavior from every seed. For the remaining, and much more challenging, six class C rules we present partial results, empirical evidence and interesting open problems. First, we need a few additional definitions to clarify the notion of tractable behavior.

Denote the space-time wedge with slopes $\alpha, \beta \in \mathbb{R}$ by

$$W_{\alpha,\beta} = \{(x,t) \in \mathbb{Z} \times \mathbb{Z}_+ : \alpha t \leq x \leq \beta t\}$$

and for $M > 0$ let $W^M_{\alpha,\beta}$ consist of points at distance at most M from $W_{\alpha,\beta}$, and $W^{-M}_{\alpha,\beta}$ of points at distance at least M from the complement $(W_{\alpha,\beta})^c$. We say that a seed A_0 is ***completely periodic*** if there exist real numbers $\alpha_0, \dots, \alpha_k$, doubly periodic solutions with links $L_1, \dots, L_k$, and a finite constant M so that:

- $\xi^{A_0}_t$ is 0 outside $W^M_{\alpha_0,\alpha_k}$; and
- for each $i = 1, \dots, k$, $\xi^{A_0}_t$ agrees with the doubly periodic evolution given by L_i in $W^{-M}_{\alpha_{i-1},\alpha_i}$.

The wedge slopes α_i and the links L_i determine the *type* of a completely periodic seed. Similarly, the governing additive dynamics and ethers [**GGP**] determine the *type* of a replicating seed. We say that a property $\mathcal{P}$ of a seed holds ***almost surely*** if the proportion of seeds with $\mathcal{P}$ converges to 1 with increasing length. We also say that $\mathcal{P}$ is ***exactly resolved*** if there is a finite automaton (depending only on the rule) which decides whether $\mathcal{P}$ holds for a seed. We say that a rule is ***simple*** if every seed either results in bounded growth, one of r types of replication, or one of the p types of complete periodicity; and each of the $1+r+p$ scenarios is exactly resolved. With these definitions, we can make our classification formal.

- ***Class R*** comprises CA which are simple and almost surely replicators. When $r = 2$, there are 10 such rules: additive CA *EDDE*, *EDED*, *EEDD*, other replicator rules *E0D0*, *E0DE*, *E11D*, *ED1E*, *EED0-E0DD*, *EEDE-E0D1*, and *EE1D* which is a replicator unless all 1's are on odd integers in which case it fixates (see Section 3).
- ***Class B*** comprises simple rules with almost sure bounded growth. Here is the list of 21 such range 2 rules, for which bounded growth occurs from every seed (see Section 7): *E000*, *E001*, *E00D*, *E00E*, *E0E0*, *E0E1*, *E0ED*, *E0EE*, *E100*, *E10E*, *E1E0*, *E1EE*, *ED00*, *ED0E*, *EDE0*, *EDEE*, *EE00*, *EE0D*, *EE0E*, *EEE0*, and *EEEE*.
- ***Class P*** comprises simple rules which are almost surely completely periodic. There are 19 range 2 rules of this type. First, 11 rules have a single wedge (with a single periodic pattern, up to time shifts) and no exceptions: *E101*, *E10D*, *E111*, *E1E1*, *ED01*, *ED0D*, *EDE1*, *E011-EE1E*, *ED11-E11E*, *E01D-EE10* (Proposition 6.1), and *EDD1-E1DE* (Proposition 6.2). Then, 3 rules have a single wedge and one pattern, but with parenthetical exceptions: *EE11* (seeds that only occupy even sites grow at maximum speed to fill the even lattice), *EEE1* (seeds with only isolated 1's fixate), and *EE01* (seeds with 1's separated by at least two 0's fixate). Then, 4 rules have two wedges but no exception: *E010* (one pattern),

E01E (one pattern), *E1D1* (one pattern), and *ED1D-E110* (two patterns, Theorem 2). Finally, the single rule with two wedges (but one pattern) and exceptions is *EED1* (seeds on the even lattice are replicators).

- ***Class C*** rules belong to none of the above three categories. This is apparently true for 6 range 2 rules: *E1DD*, *E1ED*, *ED10*, *EDD0*, *EEED*, *E1D0-EDDD*. Intuitively, these rules either have a multitude of complex behaviors which are hard to predict by examining the seed, or have chaotic dynamics [**GG3**] from some seeds — or, very often, both.

It may be difficult to *prove* that a rule belongs to class C, despite convincing experimental evidence. For example, we know of no argument for *EDD0* (*Exactly 1*), which however will not be discussed further, as this was the subject of [**GG3**]. The rest of this section is devoted to the remaining 5 purportedly class C rules.

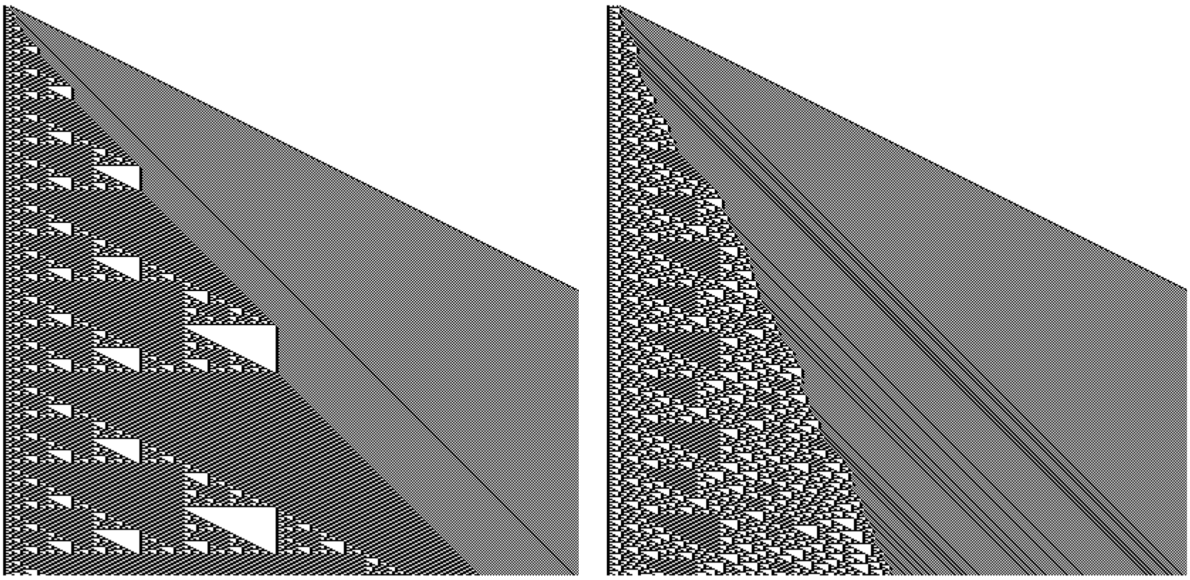

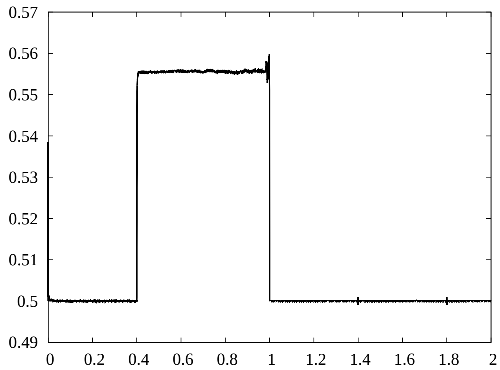

Fig. 7. *E1DD* from 110001 and 1100100101 (at time 400, cut at about $x = 400$). For the second seed, we also depict its density profile at time 10^6. This profile computes the average density on space-time rays $x = \alpha t$, $\alpha \in [0, 2]$ (see [**GG3**]).

We begin with *E1DD*. From Section 4 we know that this CA has infinitely many RPS. As in the proof of Theorem 1, let the set $\mathcal{R}$ consist of additive configurations; i.e., those that have only even blocks. We know that there are many initial configurations A_0 for which $s_{\mathcal{R}}(\xi_t)$ is bounded: exactly those attracted to one of the RPS. On the other hand, unless the entire configuration ξ_t started from A_0 eventually belongs to $\mathcal{R}$, $v_{\mathcal{R}}(A_0) \le 1$; this follows from the proof of Theorem 1. It also follows from that proof that $\mathcal{R}$ always expands into configurations with only blocks 0, 1, and 11. Another class of initial configurations is amenable to rigorous examination: those that behave similarly to mixed replicators [**GGP**], e.g., $A_0 = 110001$ (see Fig. 7). In such cases, an inductive analysis, carried out for several similar cases in [**GGP**] and omitted here, can be utilized to conclude that $s_{\mathcal{R}}(\xi_t)/t$ fluctuates between 6/7 and 3/4. Thus $v_{\mathcal{R}}(A_0) = 3/4$. Finally, there are initial configurations for which the additive region is apparently chaotic, e.g., $A_0 = 1100100101$. (See [**GG3**] for a general discussion on chaotic seeds.) Computer experiments (see Fig. 7) suggest that $v_{\mathcal{R}}(A_0)$ is close to 0.40 in such cases; however, we have neither an upper bound (smaller than 1) nor a lower bound. As the evolution contains both long time intervals during which $s_{\mathcal{R}}(\xi_t)$ is constant and those during which it advances with velocity 1, it is clear that no local method from the present paper applies.

Problem 3. Prove that $0 < v_{\mathcal{R}}(A_0) < 1$ if A_0 is a chaotic seed for rule *E1DD*.

Fig. 7 suggests that the chaotic wedge has density equal to 1/2; this wedge transitions at slope close to 2/5 to a "high correlations" state with density close to 5/9, and then at slope 1 to the checkerboard RPS that expands from the right edge. The density 1/2 raises the possibility that the chaotic configuration approximates, in the appropriate sense, the Bernoulli measure with density 1/2, which is well-known to be invariant for any right permutative dynamics, in this case *EEDD* [**Lin**]. We have no explanation for the (possible) appearance of 2/5 and 5/9.

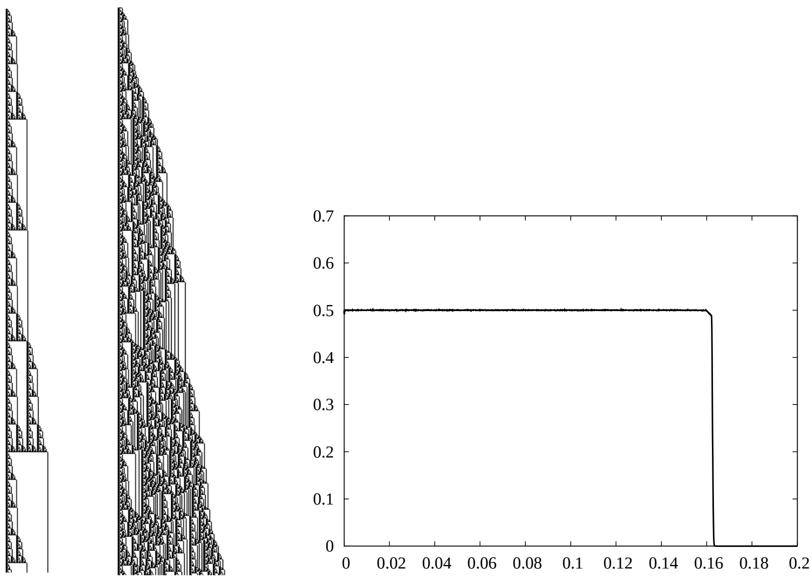

Fig. 8. *EEED* from 11, 11111, and the density profile [**GG3**] at time 10^6 in the latter case. Note that, again, the chaotic state appears to have density 1/2, presumably a consequence of right permutativity.

Our next rule is *EEED*. We know it has no RPS, but its behavior from seeds is quite interesting. Our brief discussion mostly concerns the growth velocity v_g, and $v_g(A_0)$ for seeds A_0. Recall that $s_g(\xi_t)$ is the location of the rightmost 1 in the configuration ξ_t and note that $s_g(\xi_t) \le s_g(\xi_{t+1}) \le s_g(\xi_t) + 1$ for *EEED*.

To get an upper bound on v_g, fix a configuration A of a sites with rightmost state 1, and place it so that its leftmost site (not necessarily in state 1) is at the origin, with 0's elsewhere. Use this as initial state ξ_0. Define L_t recursively as $L_0 = 0$ and

$$L_{t+1} = \begin{cases} L_t + 1 & \text{if } \xi_t(L_t) = 0, \\ L_t + 2 & \text{if } \xi_t(L_t) = 1. \end{cases}$$

That is, L_t is the leftmost "reliable" site, which is not influenced by the sites to the left of 0. The slower advance on 0 is justified by the transitions $*00 \mapsto 0$, $*01 \mapsto 1$. Note that we need only compute ξ_{t+1} on $[L_t, s_g(\xi_t) + 1]$. Further, let τ_A be the largest time t at which $L_t \le s_g(\xi_t) + 1$. The following proposition is immediate.

Proposition 8.1. *For any a,*

$$v_g \le \max_A \min_{t \le \tau_A} \frac{s_g(\xi_t) - a + 1}{t}.$$

The choice of $a = 21$ gives the upper bound $v_g \le 1/2$ (which, curiously, does not improve for a up to 32), and thus $v_g(A_0) \le 1/2$ for any seed A_0. What about lower bounds?

Any seed with only isolated 1's remains fixed. On the other hand, a seed with a least one 11 has $s_g(\xi_t) \to \infty$. To see this, assume that there are seeds for which $s_g(\xi_t)$ is eventually stationary. We need only consider seeds whose leftmost 1 is non-isolated and positioned at the origin. Among these, choose A_0 with the minimal $S = \lim_{t\to\infty} s_g(\xi_t^{A_0})$. Clearly $S \ge 2$. Let t_0 be a time at which $s_g(\xi_{t_0}^{A_0}) = S$. The 1 at S must be isolated in $\xi_t^{A_0}$ for $t \ge t_0$. The initial configuration A_0' that agrees with $\xi_{t_0}^{A_0}$, except at S where it has a 0, generates growth that never reaches $S - 1$, a contradiction with minimality of S.

There are seeds that generate sublinear growth. The simplest is $A_0 = 11$ (see Fig. 8), which has the property that at times $t_n = 2 \cdot 4^n$, $n = 0, 1, \dots$, $\xi_{t_n}^{A_0}$ equals $11[(3 \cdot 2^n - 2)\ 0\text{'s}]1$. This is easy to prove by induction and implies that

$$\limsup_t \frac{s_g(\xi_t)}{\sqrt{t}} = 3/\sqrt{2}.$$

Empirical evidence suggests that seeds grow linearly exactly when they are chaotic. The simplest such initialization is $A_0 = 11111$, and this behavior apparently becomes prevalent once the length of initial seed is large. Experiments (see Fig. 8) suggest that $v(A_0)$ is a bit larger than 0.16 in all such cases.

Problem 4. For *EEED*, prove that $v_g(A_0) > 0$ for at least one seed A_0.

The next rule we address is *E1ED*, for which we first state an open question that may be more tractable than others. We were unable to find seeds that would contradict a positive answer.

Problem 5. For *E1ED*, and its $\pi = 32$, $\sigma = 16$ RPS, is $v(A_0) = 1/2$ for every proper seed A_0?

The next question addresses the size of the domain of attraction of an RPS. Although it can be posed generally, we find it particularly compelling for this CA, as the RPS that arises from a single 1 is very common.

Problem 6. For each of the three $\pi = 32$ *E1ED* RPS, let p_k be the proportion of seeds of length k which are proper. Estimate p_k for large k.

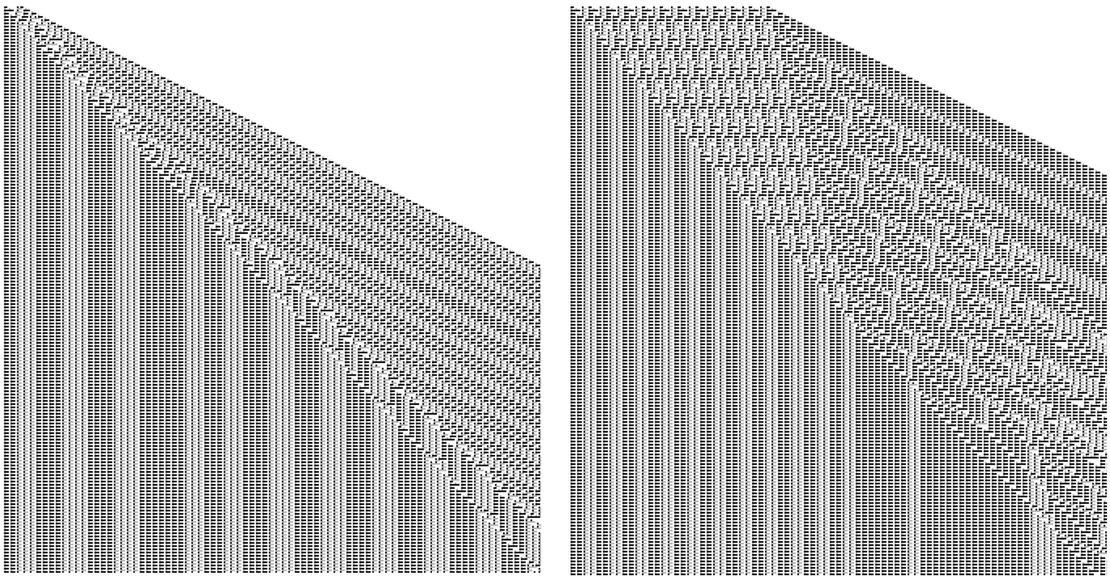

Fig. 9. *ED10* from 11[6 0's]1011 and from a portion of the "periodic blocker" defined in the text. Occupied sites are differently shaded at odd and even times or else no pattern would be visible. Observe that the left evolution does develop periodic "interfaces" and the distance between them is widening, albeit almost imperceptibly; they move at velocities 16/17 and 1.

Our next rule is *ED10*, which is known to support a variety of gliders with complex interactions [**MAM**]. The evolution from any seed, however, at first glance looks completely periodic and upon closer inspection almost so. We proceed to explain why. The final Stage I tree has $\pi_0 = 2$; one succession line of the root is $11 \to 23 \to 03 \to 02 \to 11$, and the other is $11 \to 32$. Thus the limiting configurations of leftmost sites have period 2: a 01 is followed by a 00 or by a 11, each of these two pairs is in turn followed by 01, and the updates flip between 00 and 11. Space-time images of all these look very similar. Let the described configurations comprise the set $\mathcal{R}$. A few experiments make it clear that the velocity $v_{\mathcal{R}}(A_0)$ is quite close to 1 for all seeds A_0.

It is very easy to verify that $v_{\mathcal{R}} \geq 1/2$ by computing two updates of six possible configurations such that the corresponding automaton $\mathcal{F}$ stops at the rightmost site (underlined): $011\underline{0}$, $010\underline{1}$, $0100\underline{1}$, $0111\underline{1}$, $01000\underline{0}$, $01110\underline{0}$. A computer verification involving arbitrary additional 7 sites to the right gives $v_{\mathcal{R}} \geq 8/9$. This is the best we can do with this local method, as a "periodic blocker" A_0 given by 11 followed by 00000100011^∞ has $v_\mathcal{R}(A_0) = 8/9$ and thus we prove that $v_\mathcal{R} = 8/9$. The seeds A_0 with the slowest expansion we know have $v_\mathcal{R}(A_0) = 16/17$, the simplest being A_0 =11[6 0's]1011. Fig. 9 depicts the evolution from these examples.

Problem 7. Do there exist seeds for either *E1ED* or *ED10* that have a wedge of nonperiodicity? By definition, such a wedge includes no smaller wedge on which the dynamics is periodic (i.e., agrees with the evolution from a doubly periodic solution).

We conclude this section with a discussion of remarkable behavior of *E1D0* that once more underscores the utility of Stage I label trees. Rowland [**Row**] has noted that all seeds result in the same period 16 cycle on the leftmost 53,208 sites. (Recall that we do not distinguish between time shifts of the same cycle.) Moreover, Rowland conjectures that there are cycles of infinitely many distinct time periods if the number of leftmost sites is allowed to be arbitrarily large. We will show that this claim follows from the left permutativity of this rule, and also explain why the number of cycles appears to be small.

We begin by observing that, for any edge CA, a Stage I tree for a period π is embedded in the tree of period 2π, with periodic extension of the labels. Any death by period doubling at period π is transformed into branching at period 2π, but the two resulting successors are each other's rotations by π. Thus one immediately dies by repetition; we call this *period doubling branching*. It increases the period of a limit cycle, but the *number* of distinct limit cycles is only increased by a ***genuine*** Stage I tree branching, i.e., one within a fixed period π.

Returning to *E1D0*, recall from Section 3 that periodic solutions are impossible as no label is a repetition of a label on its ancestral line on a Stage I tree. In fact, in this case, the only label repetition is among the two siblings at a period doubling branching. To see this, assume that two labels are equal up to rotation, then follow their unique ancestral lines generation by generation to a common ancestor; if this ancestor is not reached in the same number of generations, one label is a rotation of another on its ancestral line, a contradiction. Therefore, no period π Stage I tree can be final. Hence there is eventual death by period doubling for every π, confirming Rowland's conjecture.

As the deciders for *E1D0* are 1's and 3's, branching or death by period doubling occurs at a label consisting of only 0's and 2's, with branching exactly when the number of 2's is even. As a rough approximation, assume that labels are chosen independently and uniformly at random. The time along a path of period π labels before there is a label with no deciders should be of order 2^π, and then death or branching each occur with about equal probability. Thus each label has a unique successor for very long stretches once π is large, and then approximates a critical branching process [**AN**]. The size of the Stage I tree therefore should approximately square with each doubling of π, and asymptotic theory on critical branching process [**AN**] suggests that the same is true for the number of genuine branching events. This intuitive sketch suggests that the number of distinct limit cycles on the leftmost k sites grows as a small power of k — we have far too little data for a reliable estimate but would guess about $k^{0.1}$.

Turning to exact computations, the Stage I trees for $\pi = 1$, 2, 4, and 8 die at generations 2, 8, 28, and 399, with all branching of period doubling type. The first genuine branching occurs for period 16 at generation $53,207$, as observed in [**Row**]. For this reason, no matter what the seed, there is only one limit cycle on $[0, 53207]$. To illustrate this phenomenon even more dramatically, we have computed the full Stage I tree for $\pi = 16$ (which has $2,159,030$ vertices in $894,234$ generations [**Gra**]). We then constructed a partial $\pi = 32$ Stage I tree by continuing the succession line along each $\pi = 16$ leaf which at $\pi = 16$ died by period doubling, until we stopped at the first branching or death by period doubling. It turns out that there are 16 such succession lines, and the earliest stopping event was a death at generation $65,154,361$. Therefore, all seeds have only 16 distinct limit cycles on the leftmost 65 million cells. We wrap up this section with a natural open question that can be viewed as a strengthening of Rowland's conjecture.

Problem 8. Prove that the number of distinct cycles on $[0, N]$ goes to infinity as $N \to \infty$ or, equivalently, that the number of genuine branching events in the Stage I tree of period π goes to infinity as $\pi \to \infty$.

9 Totalistic rules with larger range

Fix a range r. We call an edge CA ***totalistic*** with the set of *occupation numbers* $S \subset \{1, \dots, r+1\}$ if $A \mapsto 1$ if and only if the number of 1's in A is in S. The name we give to such rules is *Tot* followed by the list of occupation numbers. Note that 1 has to be on the list due to (1.2), thus there are 2^r totalistic rules. The following result is easy to check.

Proposition 9.1. *Among the range r totalistic edge CA, the only rule without deciders is the additive one, i.e., the one whose occupation numbers are all odd integers in* $[0, r+1]$.

Totalistic rules can be replicators without being additive, as demonstrated for $r = 2$ by the *Tot 12* (*E11D*) rule [**Jen1**, **GG2**], and for general r by the *Quota* CA, which has $S = [1, r]$ [**GGP**].

We next give a result for general r, which also serves as an illustration how a handle can by necessity have a larger period than a link. As the proof suggests, the filling property can be established under more general conditions by considering temporal periods $\pi > 4$, but we do not do so here.

Proposition 9.2. *Assume that a totalistic edge CA with range r has $S = [1, \theta'] \cup [\theta, r+1]$, with $1 \le \theta' \le \theta - 2$. Assume also that $r \ge \min(3\theta - 2\theta' - 1, 2\theta - 2)$. Then the CA is filling, with a handle of period $\pi = 4$.*

Proof. Modulo rotations, any periodic orbit with period 4 starting from the origin begins with θ' columns $[1\,1\,1\,1]^T$, followed by $\theta - \theta' - 1$ columns $[1\,0\,1\,0]^T$, followed by $\theta - \theta' - 1$ columns $[1\,1\,0\,0]^T$, followed by a number $\theta - \theta' - 1 + k$, $1 \le k \le \theta - 1$, of columns $[0\,1\,1\,1]^T$. The number k is determined so that the next column is $[1\,1\,1\,1]^T$, and then so are all subsequent ones. Namely, $k = 1$ if $r \ge \theta' + 3(\theta - \theta' - 1) + 1 = 3\theta - 2\theta' - 1$ and $k = \theta'$ if $r < 3\theta - 2\theta' - 1$ but $r \ge \theta' + 2(\theta - \theta' - 1) + \theta' = 2\theta - 2$. ☐

Problem 9. Is it true that a totalistic edge CA is either a replicator or has at least one RPS?

Empirical evidence so far supports a positive answer to the above question; as an illustration, we will look at examples for $r = 3$. This range has 8 rules. Two have no RPS and one is an easy solidification:

- *Tot 13* is additive and *Tot 123* a replicator (the proof is similar to that of Theorem 2 in [**GGP**]).
- *Tot 1234* is a filling rule.

We now report the simplest RPS we found for the remaining five $r = 3$ rules.

- As for its $r = 2$ cousin *EDD0* (*Exactly 1*), RPS for *Tot 1* are not easy to come by for small π, but their number rapidly increases once they are common. In this case, the shortest RPS period is $\pi = 32$, when there are 47 different RPS, with six different σ's: 94 (3 representatives), 704 (1), 1440 (24), 1472 (2), 1488 (15), and 4416 (2). *Tot 12* has an RPS with $\sigma = 60$ for $\pi = 8$ and two with σ's 380 and 4504 for $\pi = 16$; simple seeds for these three are, respectively, 10101, 10101[5 0's]110111, and 10101[5 0's]110111001. *Tot 14* has a single RPS with $\sigma = 312$ for $\pi = 16$, and *Tot 124* has an RPS with the same π and σ. Finally, *Tot 134* has its final Stage I tree at $\pi_0 = 8$ from which one can read off its unique RPS with $\sigma = 16$. For this last RPS, can the reader determine the expansion velocity v, and $v(A_0)$ for all seed environments A_0?

The $r = 2$ *EDD0* rule is unique among the rules of that range in that each RPS from its large collection apparently advances against a chaotic state with a characteristic expansion velocity (see [**GG3**] for empirical evidence). This CA is also totalistic and we believe that this is not a coincidence. The $r = 3$ observations just described, as well as our preliminary experiments with $r = 4$ rules, suggest that indeed totalistic rules canonically feature the interplay between periodicity and chaos. The rare exceptions include: replicators; filling rules; and borderline cases like $r = 3$ *Tot 134* or $r = 6$ *Tot 12567*. One might say that totalistic CA are discrete counterparts to iterates of piecewise monotone maps [**BG**, **BS**] and the analogy carries through to their behavior. One phenomenon with no obvious counterpart in the continuous world is that totalistic CA exhibit vestiges of replication in two ways: some seeds lead to replication; and many RPS generate large triangles of 0's reminiscent of additive evolution. We began exploring such connections in [**GG3**], but they are worth further study.

Acknowledgments and support

Many initial insights for this work arose from our correspondence with Dean Hickerson on the *Exactly 1* cellular automaton. JG was partially supported by NSF grant DMS–0204376 and the Republic of Slovenia's Ministry of Science program P1-285.

References

[**AN**] K. B. Athreya, P. E. Ney, “Branching Processes.” Springer-Verlag, 1972.

[**BG**] A. Boyarsky, P. Gora, “Laws of Chaos: Invariant Measures and Dynamical Systems in One Dimension.” Birkhäuser, 1997.

[**BK**] M. Boyle, B. Kitchens, *Periodic points for onto cellular automata*, Indag. Mathem. N.S. 16 (1999), 483–493.

[**BL**] M. Boyle, B. Lee, *Jointly periodic points in cellular automata: computer exploration and conjectures*, Experimental Mathematics 16 (2007), 293–302.

[**BR**] C. D. Brummit, E. Rowland, *Boundary growth in one-dimensional cellular automata*, Complex Systems, to appear.

[**BS**] M. Brin, G. Stuck, “Introduction To Dynamical Systems.” Cambridge, 2002.

[**Coo**] M. Cook, *Universality in elementary cellular automata*, Complex Systems 15 (2004), 1–40.

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

[**GG1**] J. Gravner, D. Griffeath, *Cellular automaton growth on* $\mathbb{Z}^2$*: theorems, examples and problems*, Advances in Applied Mathematics 21 (1998), 241–304.

[**GG2**] J. Gravner, D. Griffeath, *Asymptotic densities for Packard Box rules*, Nonlinearity 22 (2009), 1817–1846.

[**GG3**] J. Gravner, D. Griffeath, *The one-dimensional Exactly 1 cellular automaton: replication, periodicity, and chaos from finite seeds*, Journal of Statistical Physics 142 (2011), 168-200.

[**GGP**] J. Gravner, G. Gliner, M. Pelfrey, *Replication in one-dimensional cellular automata*, Physica D 240 (2011), 1460–1474.

[**Gra**] J. Gravner, *Edge Cellular Automata*.
`http://www.math.ucdavis.edu/~gravner/edge/`

[**Gri**] D. Griffeath, *Primordial Soup Kitchen*.
`http://psoup.math.wisc.edu/exactly1/exactly1.html`

[**Hic1**] D. Hickerson, private communication (1998).

[**Hic2**] D. Hickerson, private communication (2009).

[**Jen1**] E. Jen, *Global properties of cellular automata*, Nonlinearity 4 (1991), 251–276.

[**Jen2**] E. Jen, *Exact solvability and quasiperiodicity of one-dimensional cellular automata*, Journal of Statistical Physics 43 (1986), 219–242.

[**Jen3**] E. Jen, *Transience and dislocations in one-dimensional cellular automata*. In "Cellular automata and cooperative systems (Les Houches, 1992)," NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci., 396, Kluwer Acad. Publ., 1993, 299–310.

[**Lin**] D. A. Lind, *Applications of ergodic theory and sofic systems to cellular automata*, Physica D 10 (1984), 36–44.

[**MAM**] G. J. Martínez, A. Adamatzky, H. V. McIntosh, *Phenomenology of glider collisions in cellular automaton Rule 54 and associated logical gates*, Chaos, Solitons and Fractals 28 (2006) 100-111.

[**Row**] E. S. Rowland, *Local Nested Structure in Rule 30*, Complex Systems 16 (2006), 239–258.

[**SCJ**] L. Shi, F. Chen, W. Jin, *Gliders, Collisions and Chaos of Cellular Automata Rule 62*, 2009 International Workshop on Chaos-Fractals Theories and Applications, IEEE (2009), 221–225.

[**TR**] A. Trevorrow, T. Rokicki, *Golly*. `http://golly.sourceforge.net/`

[**Ver**] F. Verhulst, "Nonlinear differential equations and dynamical systems." Springer, 1996.

[**Wil1**] S. J. Willson, *Cellular automata can generate fractals,* Discrete Appl. Math. 8 (1984), 91–99.

[**Wil2**] S. J. Willson, *Computing fractal dimensions for additive cellular automata,* Physica D 24 (1987), 190–206.

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

[**Wol1**] S. Wolfram, *Statistical Mechanics of Cellular Automata*, Rev. Mod. Phys. 55 (1983) 601–644.

[**Wol2**] S. Wolfram, *Computational Theory of Cellular Automata*, Communications in Mathematical Physics 96 (1984), 15–57.

[**Wol3**] S. Wolfram, *Random Sequence Generation by Cellular Automata*, Advances in Applied Mathematics 7 (1986), 123–169.

[**Wol4**] S. Wolfram, "A New Kind of Science." Wolfram Media, 2002.

[**Wol5**] S. Wolfram, "The Wolfram Atlas: Elementary Cellular Automata." `http://atlas.wolfram.com/01/01/`, 2010.