

Trace Complexity of Chaotic Reversible Cellular Automata^{*}

Jarkko Kari, Ville Salo, and Ilkka Törmä

TUCS – Turku Centre for Computer Science, Finland,
University of Turku, Finland
{jkari, vosalo, iatorm}@utu.fi

Abstract. Delvenne, Kůrka and Blondel have defined new notions of computational complexity for arbitrary symbolic systems, and shown examples of effective systems that are computationally universal in this sense. The notion is defined in terms of the trace function of the system, and aims to capture its dynamics. We present a Devaney-chaotic reversible cellular automaton that is universal in their sense, answering a question that they explicitly left open. We also discuss some implications and limitations of the construction.

Keywords: cellular automaton, reversible, chaos, computational complexity, trace, symbolic system

1 Introduction

A significant branch of dynamical systems research is the study of computability and computational complexity of finitely presented systems. In the literature, there are usually multiple incomparable notions of computability and computational universality for a sufficiently popular model, like cellular automata [16,18,4]. Traditionally, for a model to be considered computationally universal, it is sufficient for it to be able to simulate the computation process of any Turing machine in a suitably transparent way. However, seemingly minor variations to the formal definition (if one is presented) may reduce a universal system into a trivial one. A related notion is that of *intrinsic* universality, which refers to the ability of simulating any other instance of the same model in some formally defined way. In cellular automata, intrinsic universality is usually defined with respect to block simulations, although in earlier research this notion had usually also been left undefined. See [8] for a discussion on the implications of not defining these notions rigorously. In the context of reversible computation, intrinsic universality of reversible Turing machines has been discussed in [1].

In [5], a new definition of computational universality was proposed that can be applied to a wide range of discrete dynamical systems, including cellular automata, shift spaces, tag systems, and Turing machines, which can be viewed as dynamical systems in more than one way [9]. It is an update of the definition

^{*} Research supported by the Academy of Finland Grant 131558

given in [6], and aims to capture the dynamical complexity of the system, so that systems that are dynamically too simple (like the identity map on a set) or allow too much freedom (like the shift map) would not be universal.

The computational universality presented in [5] intuitively means the hardness of deciding prediction problems like ‘for subsets U, V, W of the state space, is there a point in U that is mapped by the dynamics to V , and stays there until it enters W .’ For example, one would show that Turing machines are computationally universal by defining U as the singleton set containing the initial configuration, V as the set of all configurations, and W as the set of final configurations. Of course, for the definition to be sensible, the subsets need to be restricted in some way. In symbolic systems, whose elements are infinite sequences of symbols, we require that the sets are clopen, that is, they are defined by the contents of finitely many coordinates.

One of the main observations in [5] is that universal systems tend to be ‘at the edge of chaos’: the dynamics appears chaotic, but has underlying structure that gives rise to the universality. They give examples of effective systems that are both universal and chaotic in the sense of Devaney [7], but the existence of a universal chaotic cellular automaton is explicitly left open. In this article, we construct a reversible universal chaotic cellular automaton, answering this question in the positive. Note that although reversible cellular automata were shown to be able to simulate arbitrary computation already in [11], and their construction seems to be universal in the sense of [5], it is not chaotic.

2 Definitions

Let \mathbb{M} be either \mathbb{N} or \mathbb{Z} , and let S be a finite alphabet. The set $S^{\mathbb{M}}$, equipped with the product topology, is called the *full \mathbb{M} -shift on S* , and its elements are called *configurations*. The monoid $(\mathbb{M}, +)$ acts on $S^{\mathbb{M}}$ by the *shift maps* $\sigma^m : S^{\mathbb{M}} \rightarrow S^{\mathbb{M}}$ for $m \in \mathbb{M}$, defined by $\sigma^m(x)_n = x_{n+m}$. We denote $\sigma^1 = \sigma$. For a word $w \in S^n$ and $x \in S^{\mathbb{M}}$, we say that w *occurs in x* , denoted $w \sqsubset x$, if there exists $m \in \mathbb{M}$ such that $w = x_{[m, m+n-1]}$. This notation is extended to sets of configurations in the obvious way. An *\mathbb{M} -shift space* is a topologically closed set $X \subset S^{\mathbb{M}}$ satisfying $\sigma(X) \subset X$. Equivalently, a shift space is defined by a set $F \subset S^*$ of *forbidden words* as $\mathcal{X}_F = \{x \in S^{\mathbb{M}} \mid \forall w \in F : w \not\sqsubset x\}$. If F is finite, \mathcal{X}_F is a *shift of finite type* (SFT for short). We denote $\mathcal{B}_n(X) = \{w \in S^n \mid w \sqsubset X\}$ and $\mathcal{B}(X) = \bigcup_{n \in \mathbb{N}} \mathcal{B}_n(X)$.

A *block map* is a continuous function $f : X \rightarrow Y$ between shift spaces $X, Y \subset S^{\mathbb{M}}$ that satisfies $f \circ \sigma|_X = \sigma|_Y \circ f$. Alternatively, a block map is defined by a *local rule* $\hat{f} : \mathcal{B}_{a+m+1}(X) \rightarrow \mathcal{B}_1(Y)$, where $a, m \in \mathbb{N}$ are the *anticipation* and *memory* of \hat{f} , by $f(x)_n = \hat{f}(x_{[n-m, n+a]})$. The interval $\{-m, \dots, a\} \subset \mathbb{M}$ is called the *neighborhood* of \hat{f} . In the case $\mathbb{M} = \mathbb{N}$, we must have $m = 0$. If $X = Y = S^{\mathbb{M}}$, then f is called a *cellular automaton* (CA for short), and a bijective CA is called *reversible*, since its inverse function is also a CA. We sometimes identify a CA and its local rule, but this should always be clear from the context.

A *symbolic system* is a tuple (X, f) , where X is a compact metric space with countable clopen basis (equivalently, homeomorphic to a closed subset of a full shift), and $f : X \rightarrow X$ is continuous. The system is *effective* if the clopen basis of X can be enumerated so that complementation, intersection and f -preimage are computable operations. It is *chaotic* (in the sense of Devaney [7]) if

- it is sensitive (there exists $\epsilon > 0$ such that for all $x \in X$ and $\delta > 0$ there exist $y \in X$ and $n \in \mathbb{N}$ with $d(x, y) < \delta$ and $d(f^n(x), f^n(y)) \geq \epsilon$),
- it is transitive (for all nonempty open sets $U, V \subset X$, there exists $n \in \mathbb{N}$ with $U \cap f^n(V) \neq \emptyset$), and
- the f -periodic points (those $x \in X$ for which $f^n(x) = x$ for some $n \in \mathbb{N}$) are dense in X .

In particular, every \mathbb{Z} -shift space (X, σ) with the left shift is a symbolic system, as is (X, f) for every block map $f : X \rightarrow X$. These are the only kinds of symbolic systems we use in this article; the full definition is given only for completeness, and to state the definitions of universality given in [5].

Example 1. For a cellular automaton $f : S^{\mathbb{Z}} \rightarrow S^{\mathbb{Z}}$, most dynamical notions have combinatorial characterizations. For example, f is transitive if and only if for all words $u, v \in S^{2\ell+1}$ of the same odd length, there exists a configuration $x \in S^{\mathbb{Z}}$ and $n \in \mathbb{N}$ such that $x_{[-\ell, \ell]} = u$ and $f^n(x)_{[-\ell, \ell]} = v$.

A *Muller automaton* is a quintuple $A = (Q, q_0, \Sigma, \delta, F)$, where Q is a finite state set, $q_0 \in Q$ an initial state, Σ a finite input alphabet, $\delta : Q \times \Sigma \rightarrow Q$ a transition function and $F \subset 2^Q$ a set of accepting subsets of states. A Muller automaton runs deterministically on infinite words $w \in \Sigma^{\mathbb{N}}$ analogously to a standard finite automaton, and accepts if the set of states that are visited infinitely often during the computation is in the set F . The language accepted by A is denoted L_A . For a language $L \subset A^*$, we denote by $L^\omega \subset A^{\mathbb{N}}$ the set of infinite concatenations of the words of L . In particular, if L is regular, then L^ω is accepted by a Muller automaton. See [17] for a reference on Muller automata, and other types of finite automata on infinite words.

In this article, a *Turing machine* is a sextuple $M = (Q, \Sigma, q_0, q_f, B, \delta)$, where Q is a finite state set, Σ a finite input alphabet, $q_0, q_f \in Q$ are the initial and final states, $B \in \Sigma$ is the blank letter and $\delta \subset Q \times (\Sigma \times \Sigma \cup \{/\} \times \{+, 0, -\}) \times Q$ a transition relation. Turing machines are run on two-way infinite tapes, and the initial input is placed immediately to the right of the head. The interpretation of a quadruple $[q_1, a, b, q_2] \in \delta$ in the case $a, b \in \Sigma$ is that if M is in state q_1 and reading the letter a , it may rewrite it to b and go to state q_2 . In the case $a = /$ and $b \in \{+, 0, -\}$, if M is in state q_1 , it may go to state q_2 and move one step in the direction indicated by b . Two quadruples $[q_1, a, b, q_2]$ and $[q'_1, a', b', q'_2]$ overlap in domain if $q_1 = q'_1$ and $a, a' \in \Sigma \implies a = a'$. They overlap in range if $q_2 = q'_2$ and $a, a' \in \Sigma \implies b = b'$. If no distinct quadruples overlap in domain (range), then M is *deterministic (reversible, respectively)*. As usual, the language of a Turing machine is the set of input words on which it eventually halts.

We use the following terminology for certain classes in the arithmetical and analytical hierarchies. A set $N \subset \mathbb{N}$ is called Σ_1^0 if it is recursively enumerable,

and Π_1^0 if its complement is. The set is called Σ_1^1 , if there exists an oracle Turing machine M such that

$$N = \{n \in \mathbb{N} \mid \exists f : \mathbb{N} \rightarrow \mathbb{N} : M \text{ never halts on input } n \text{ with oracle } f\}.$$

These are not the standard definitions of the classes, but characterizations whose proofs can be found, for example, in [15, Theorem 1.3]. Hardness and completeness of a set with respect to these classes is defined using Turing reductions. When classifying subsets of other countable sets than \mathbb{N} , for example $\{0, 1\}^*$, we assume that they are in some natural and computable bijection with \mathbb{N} .

In [12], it was proved that deterministic reversible Turing machines are capable of simulating any deterministic Turing machine (first proved in [2] for *multi-tape* Turing machines). We will not go into the details of the notion of simulation, but it is easy to see that it implies the following lemmas.

Lemma 1. *There exists a deterministic reversible Turing Machine M whose language is Σ_1^0 -complete.*

Lemma 2. *There exists a deterministic reversible Turing Machine M , whose tape alphabet contains 0, 1 and #, such that the set*

$$L = \{w \in \{0, 1\}^* \mid \exists u \in (0^*1)^\omega : M \text{ never halts on } w\#u\}$$

is Σ_1^1 -complete, and the head never steps left of the origin on right-infinite inputs.

3 Traces and Computational Universality

In this section, we recall the definition of computational universality for an effective symbolic system (X, f) , as given in [5]. First, a *clopen partition* of X is a finite collection $\mathcal{C} = (C_s)_{s \in \Sigma}$ of mutually disjoint clopen subsets of X such that $X = \bigcup_{s \in \Sigma} C_s$, labeled by a finite set Σ . The partition can be seen as an observation or experiment, with input $x \in X$ resulting in the unique label $\pi_{\mathcal{C}}(x) = s_0 \in \Sigma$ such that $x \in C_{s_0}$. More information can be extracted from x if we apply the dynamics function f to it and repeat the experiment, obtaining the result $\pi_{\mathcal{C}}(f(x)) \in \Sigma$. Iterating the idea leads to the following definition.

Definition 1. *Let (X, f) be a symbolic system, and let $\mathcal{C} = (C_s)_{s \in \Sigma}$ be a clopen partition of X . For $x \in X$, the f -itinerary of x via \mathcal{C} is the infinite sequence $\pi_{\mathcal{C}}^f(x) \in \Sigma^{\mathbb{N}}$ defined by $\pi_{\mathcal{C}}^f(x)_n = \pi_{\mathcal{C}}(f^n(x))$ for all $n \in \mathbb{N}$. The \mathcal{C} -trace shift of f is the \mathbb{N} -shift space $\tau_{f, \mathcal{C}} = \{\pi_{\mathcal{C}}^f(x) \mid x \in X\}$. In the case $X \subset S^{\mathbb{Z}}$, we denote by $\tau_{f, n}$ the trace with respect to the partition $\mathcal{C}_n = (C_w)_{w \in S^{2n+1}}$, where $C_w = \{x \in X \mid x_{[-n, n]} = w\}$ for all $w \in S^{2n+1}$.*

Example 2. Let $f : S^{\mathbb{Z}} \rightarrow S^{\mathbb{Z}}$ be a cellular automaton, and let $n \in \mathbb{N}$. Then the trace shift $\tau_{f, n}$ is obtained by taking, for each $x \in S^{\mathbb{Z}}$, the sequence

$$x_{[-n, n]}, f(x)_{[-n, n]}, f^2(x)_{[-n, n]}, f^3(x)_{[-n, n]}, \dots$$

of the central words occurring in the evolution of the initial state x under f .

The article [5] defines the following two decision problems.

Definition 2. *Let (X, f) be an effective symbolic system. The infinite time prediction problem asks whether we have $\tau_{f, \mathcal{C}} \cap L_A \neq \emptyset$ for a given partition $\mathcal{C} = (C_s)_{s \in \Sigma}$ and Muller automaton A on the alphabet Σ . The finite time prediction problem asks whether we have $\mathcal{B}(\tau_{f, \mathcal{C}}) \cap L \neq \emptyset$ for a given partition $\mathcal{C} = (C_s)_{s \in \Sigma}$ and regular language $L \subset \Sigma^*$. We say that (X, f) is computationally universal if its finite time prediction problem is Σ_1^0 -complete.*

As hinted in Section 1, the most obvious way of showing that an effective symbolic system (X, f) is computationally universal in the above sense is to construct a simulation of a Turing machine M (or some other universal computational device) by f , and identify three disjoint clopen sets $C_0, C_1, C_2 \subset X$ that correspond to the classes of initial, intermediate, and halting configurations of M . Then the instance 01*2 of the finite time prediction problem is positive if and only if M halts on one of the the initial configurations in C_0 .

Example 3. We continue Example 2. In the finite time prediction problem for $(S^{\mathbb{Z}}, f)$, we are given a clopen partition of $S^{\mathbb{Z}}$, which we (for now) assume to be \mathcal{C}_n for some $n \in \mathbb{N}$, and a regular language L over the alphabet S^{2n+1} . For example, if $S = \{0, 1\}$ and $n = 1$, then L may be given as the regular expression $([000][010])^*[111] + [100]^*[111]$ (note that the ‘letters’ of this regular expression are binary words of length 3). If there exists $x \in \{0, 1\}^{\mathbb{Z}}$ such that, for example, $x_{[-1,1]} = 000$, $f(x)_{[-1,1]} = 010$ and $f^2(x)_{[-1,1]} = 111$, then the answer to the finite time prediction problem with these inputs is ‘yes’, since the language of the \mathcal{C}_1 -trace shift of f contains the word $[000][010][111] \in L$.

Examples of chaotic universal effective symbolic systems and universal cellular automata were provided in [5], but it was explicitly left open whether a universal cellular automaton can be chaotic.

The notion of universality given in the earlier work [6] is also equivalent to the Σ_1^0 -completeness of a prediction problem, but instead of Muller automata, the definition uses a temporal logic that specifies subsets of the state space. We will not digress into this subject, as it is not necessary for stating and proving the main results of this article.

4 Main Results

In this section, we prove that a chaotic reversible cellular automaton can be computationally universal and even have a maximally hard infinite time prediction problem. In the proof, we use the following well-known lemma, found explicitly in [10].

Lemma 3. *A reversible cellular automaton is chaotic (in the sense of Devaney) if and only if it is transitive.*

We are now ready to state and prove our main theorem.

Theorem 1. *There exists a chaotic reversible cellular automaton whose finite-time prediction problem is Σ_1^0 -complete, even when restricted to the radius-one partition \mathcal{C}_1 . In particular, the automaton is computationally universal in the sense of Definition 2.*

Proof. We first describe the general idea of the construction. The configurations of the reversible CA we construct are divided into ‘compartments’, each of which may contain one read-write head of a reversible Turing machine. The automaton simulates the Turing machines separately in each compartment, changing the direction of the simulation if they halt, and no information can be passed between the compartments. The compartments and the machines are also constantly shifted to the left. When one of the machines halts, it sends a signal to the right.

Now, in the trace shift we wish to see the following pattern: the left wall of a compartment, then some empty space, then the Turing machine head in its initial state followed by an input word, then another empty stretch, and finally the right-moving signal emitted by the halting machine. If we see such a pattern, the machine must have halted, since the signal cannot have come from the other side of the wall, and it cannot be the result of a left-moving signal bouncing off the wall, for no such signal was seen earlier. Finally, the transitivity of the CA follows from the constant shifting of the compartments and the fact that they never communicate, so that every central pattern of a configuration can eventually be replaced by arbitrary data.

Let $M = (\Sigma, Q, \delta, B, q_0, q_f)$ be the reversible Turing Machine of Lemma 1. We may assume that the head of M makes a move only on every third step, that it does not make a move in state q_0 for two steps when run backwards or forwards, that it always takes at least $2|w| + 2$ steps for M to halt on an input word w , and that it always halts after an even number of steps or runs forever. Denote $S = \Sigma \times (Q \cup \tilde{Q} \cup \{\leftarrow, \rightarrow\})$, where $\tilde{Q} = \{\tilde{q} \mid q \in Q\}$ is a disjoint copy of Q , and define a reversible cellular automaton f_M on $S^{\mathbb{Z}}$ as follows.

We partition each configuration $x \in S^{\mathbb{Z}}$ into segments whose second track is of the form $\leftarrow^m q \rightarrow^n$, where $q \in Q \cup \tilde{Q}$, or $\leftarrow^{m+1} \rightarrow^{n+1}$, for some (possibly infinite) $m, n \geq 0$. Namely, each cell of x is contained in at least one pattern of this form, and when we take the maximal ones, the partition is uniquely determined. On these segments, f_M simulates a computation of M (backwards in time if $q \in \tilde{Q}$) in a standard way. Namely, consider a cell in state $(a, q) \in \Sigma \times Q$. If $[q, a, b, r] \in \delta$ for some $b \in \Sigma$ and $r \in Q$, the cell will update to (b, r) . Each two-cell pattern $(a, q)(b, \rightarrow)$ such that $[q, /, +, r] \in \delta$ for some $b \in \Sigma$ will update to $(a, \leftarrow)(b, r)$, and analogously for a 0- or --move. In all other cases (no applicable quadruple exists, or the segment ends), the cell becomes (a, \tilde{q}) . To such cells, the quadruples and the time-reversal rule are applied in the reverse direction: $[r, b, a, q] \in \delta$ results in (b, \tilde{r}) and so on. All cells not mentioned here retain their state. Thus the first track acts as the tape, the endpoints of the segments never move, and if the simulation cannot be carried on then it changes direction. For $x \in S^{\mathbb{Z}}$, denote by $r(x) \in S^{\mathbb{Z}}$ the configuration obtained from x by changing every $q \in Q$ to \tilde{q} , and vice versa; it is easy to see that $r \circ f_M \circ r = f_M^{-1}$, which implies that f_M is reversible. See Figure 1 for a visualization of the dynamics of f_M .

Now, let $D = \{\square, \boxminus, \boxplus, \boxtimes\}$, and denote $R = S \times D$. The cells in D represent particles traveling to the left or to the right, with the fourth one containing one of each. We define a CA g on $R^{\mathbb{Z}}$ that functions as follows:

1. Apply f_M to the first track.
2. Shift each particle to its direction, unless it would cross the barrier between two segments, in which case change its direction.
3. If a cell contains the final state of M in its first track, apply the bijection $\square \leftrightarrow \boxplus, \boxminus \leftrightarrow \boxtimes$ to the second track.

See Figure 2 for a visualization. Since each of the three steps is clearly reversible, so is g . Finally, define $h = \sigma \circ g^2$, which is likewise reversible, and denote by $\pi_S : R \rightarrow S$ and $\pi_D : R \rightarrow D$ the obvious projections from R .

Now, let $w \in \Sigma^*$ be arbitrary, and denote by $L(w) \subset R^*$ the regular language

$$(B, \rightarrow, \square)(B, \leftarrow, \square)^*(B, q_0, \square)(w \times (\rightarrow, \square)^{|w|})(B, \rightarrow, \square)^*(B, \rightarrow, \boxplus) \quad (1)$$

The words of the language $L(w)$ consist of the left border of a segment, then some ‘empty’ cells, followed by the initial state of M and its input word, then more empty cells, and finally a right-moving particle. Let $U(w) \subset (R^3)^{\mathbb{N}}$ be the open set of all configurations that have a prefix $v \in (R^3)^*$ such that the middle components of the triples in v form a word in $L(w)$, and none of the letters of v contain a left-moving particle. We claim that $\tau_{h,1} \cap U(w)$ is nonempty if and only if M halts on w , which is Σ_1^0 -complete.

First, assume that M halts on w after exactly $2n + 2$ steps, where $n \geq |w|$, and define $x \in R^{\mathbb{Z}}$ by

$$x_i = \begin{cases} (B, \rightarrow, \square), & \text{if } i \leq -n, \\ (B, \leftarrow, \square), & \text{if } -n < i < 0, \\ (B, q_0, \square), & \text{if } i = 0, \\ (w_{i-1}, \rightarrow, \square), & \text{if } 1 \leq i \leq |w|, \\ (B, \rightarrow, \square), & \text{if } i > |w|. \end{cases}$$

Denote $y = h^{-n}(x)$. We claim that $(h^i(y)_{[-1,1]})_{i \in \mathbb{N}} \in U(w)$. To prove that, we first remark that for all $k, \ell \in \mathbb{Z}$ such that $3|\ell| \geq |k|$ we have $\pi_S(g^k(x)_\ell) = \pi_S(x_\ell)$, since only the single Turing Machine head can introduce changes to the S -component of the configuration, and it only moves every third step by assumption. This means that for all $i \in \mathbb{Z}$ we have $\pi_S(h^i(y)_0) = \pi_S(x_{i-n})$.

Since M does not halt in $2n$ steps, the second track of every configuration in $\{h^i(y) \mid i \in \{0, \dots, 2n\}\}$ contains no particles. But since M halts at step $2n + 2$, there exists $k \in \mathbb{Z}$ with $|k| \leq \frac{n}{3}$ such that $\pi_D(h^{2n+1}(y)_{k-n-1}) = \pi_D(g^{2n+2}(x)_k) = \boxplus$. Since the single Turing Machine head will not enter the final state for another $4n + 4$ steps, we have $\pi_D(h^i(y)_0) = \square$ for every $i \in \{2n + 1, \dots, 3n - k + 1\}$, and $\pi_D(h^{3n-k+2}(y)_0) = \boxplus$. Furthermore, no letter in $(h^i(y)_{[-1,1]})_{i \in \mathbb{N}}$ contains a left-moving particle, and together with the previous paragraph, this shows that $(h^i(y)_{[-1,1]})_{i \in \mathbb{N}} \in U(w)$, and thus $\tau_{h,1} \cap U(w) \neq \emptyset$. See Figure 3 for a visualization that will also be helpful in the converse direction.

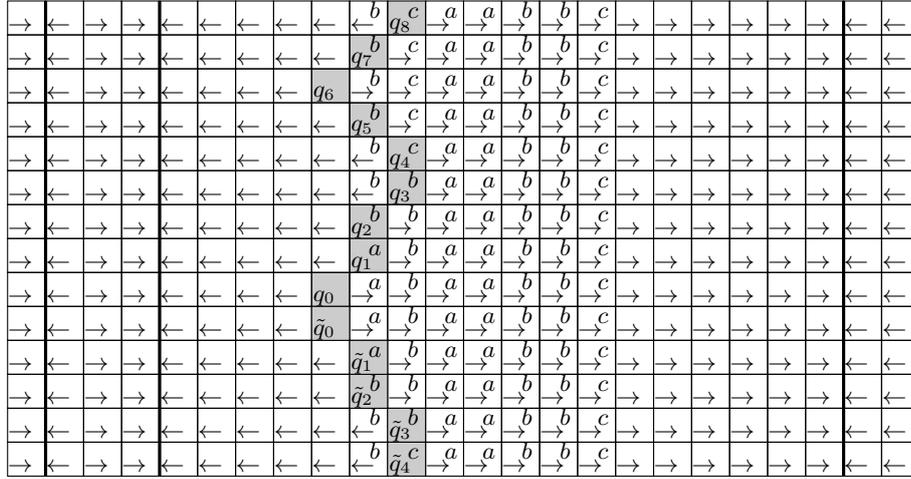


Fig. 1. A spacetime diagram of the reversible cellular automaton f_M simulating the reversible Turing machine M . Each row is (the central pattern of) a configuration of $S^{\mathbb{Z}}$, and time increases upwards. The shaded cells contain a Turing machine head, and the thick vertical lines mark the borders of segments. The initial state of M is q_0 , which is preceded in the simulation by the ‘backward’ state \tilde{q}_0 , and a, b, c are elements of the tape alphabet. To save space, M does not move only every third step in this figure.

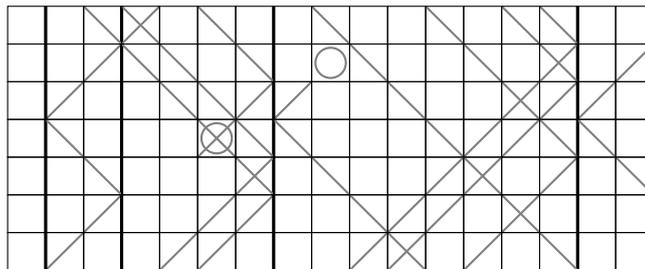


Fig. 2. A spacetime diagram of g , showing the dynamics of the particles. The thick lines mark the borders of segments, and the circles denote the final state of M . Other information from the first track is not shown. The particles are drawn in gray for clarity.

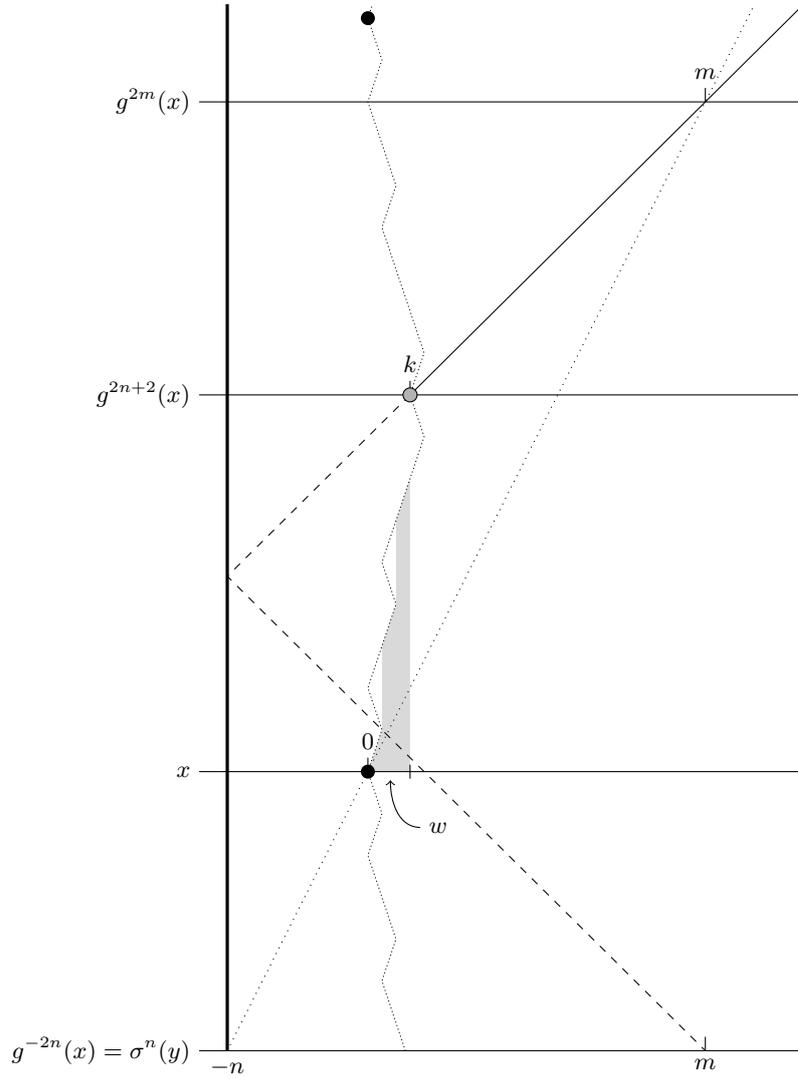


Fig. 3. A schematic spacetime diagram of g . The black (gray) circles represent the initial (final) states of M (at coordinates 0 and k , respectively), and the densely dotted line traces the read-write head of M as it carries out its computation. The sparsely dotted line corresponds to the central coordinates of the configurations $h^{i-n}(x)$ for $i \in \mathbb{N}$. The thick vertical line is a border of segments, the diagonal line is a particle, and the dashed continuation is its hypothetical path assuming that M never halted. The particle meets the sparsely dotted line at coordinate m . Each horizontal line represents a configuration. The shaded area represents the letters of the input word $w \in \Sigma^*$ in x that the head of M has not yet read (which the sparsely dotted line goes through).

Second, assume that $y \in R^{\mathbb{Z}}$ is such that $(h^i(y))_{i \in \mathbb{N}} \in U(w)$, with the $*$ -symbols in the definition of $L(w)$ being replaced by numbers $n-1, m-|w|-1 \in \mathbb{N}$, so that we have $h^n(y)_0 = (B, q_0, \square)$ and $h^{n+m}(y)_0 = (B, \rightarrow, \boxplus)$. Let $x = h^n(y)$. Since the segments used in the simulation of M never move under the action of g , the interval $I = [-n, m]$ is contained in a single segment of y , and contains its left endpoint. The segment contains a Turing Machine head, and as above, for all $k, \ell \in \mathbb{Z}$ such that $3|\ell| \geq 2|k|$, we have $\pi_S(h^{k+n}(y)_{\ell-k}) = \pi_S(g^{2k}(x)_\ell) = \pi_S(x_\ell)$. In particular, the word $\pi_S(x_I)$ contains a Turing Machine head in state q_0 followed by the input word w , surrounded by the blank symbols.

Consider now the two remaining tracks of x and y . We know from the above that $g^{2m+2n}(y)_{m+n} = h^{m+n}(y)_0 = (B, \rightarrow, \boxplus)$ (the intersection of the particle and the densely dotted line in Figure 3). This implies that either the coordinate $g^{2m+2n-i}(y)_{m+n-i}$ contains the final state of M for some $i \in \{0, \dots, m+n\}$, or each of them contains a right-moving particle. In the latter case, $g^{m+n-1}(y)_0$ contains a left-moving particle, since it is next to a segment border, and so does $g^{m+n-1-j}(y)_j$ for all $j \in \{0, \dots, m+n\}$ (see the dashed line in the figure). Now, at $j = 0$ we have $m+n-1-j > 2j$, while at $j = m+n$, the opposite holds. Thus we have $|(m+n-1-j_0)/2 - j_0| \leq 1$ for some j_0 such that $m+n-1-j_0$ is even; denote $m+n-1-j_0 = 2\ell$. Then $0 \leq \ell < m+n$, and the cell

$$h^\ell(y)_b = g^{2\ell}(y)_{\ell+b} = g^{m+n-1-j_0}(y)_{j_0}$$

contains a left-moving particle for some $b \in \{-1, 0, 1\}$ (the intersection of the dashed line with the densely dotted line in the figure). But this is impossible since $(h^i(y))_{i \in \mathbb{N}} \in U(w)$, so the choice that none of the cells $g^{2m+2n-i}(y)_{m+n-i} = g^{2m-i}(x)_{m-i}$ contain the final state of M was incorrect. Thus one of them does, implying that M eventually halts, since the initial and final states lie in the same segment (at different times). This finishes the proof of $\tau_{h,1} \cap U(w) \neq \emptyset$ being equivalent to M halting on w .

Finally, we show that h is chaotic, and by Lemma 3, it suffices to prove transitivity. The proof is standard for reversible CA that have ‘shifting barriers’, in our case borders of segments. Let thus $u, v \in R^n$ be two words of the same length n . Define $w = (B, \leftarrow, \square)(B, \rightarrow, \square)$. Then for all $x, y \in R^{\mathbb{Z}}$ with $x_{[0,1]} = y_{[0,1]} = x_{[n+2, n+3]} = y_{[n+2, n+3]} = w$ and $x_{[2, n+1]} = y_{[2, n+1]}$, we have $g^i(x)_{[2, n+1]} = g^i(y)_{[2, n+1]}$ for all $i \in \mathbb{Z}$. This is because the evolution of a segment under g is independent of other segments.

Now, let $x \in R^{\mathbb{Z}}$ be such that $x_{[0, n+3]} = wuw$, and let $y \in R^{\mathbb{Z}}$ be defined by

$$y_i = \begin{cases} (wvw)_i, & \text{if } i \in [0, n+3], \\ h^{n+4}(x)_i, & \text{otherwise.} \end{cases}$$

By applying the above argument to y , which satisfies $y_{[n+4, n+5]} = y_{[2n+6, 2n+7]} = w$, we have that $h^{-n-4}(y)_{[2, n+1]} = u$, and by definition $y_{[2, n+1]} = v$. This shows that h is transitive, and thus chaotic. See Figure 4 for a visualization of this argument. \square

The regular expression (1) used in this construction is *local* in the sense that it can be recognized by a DFA whose state only depends on the n symbols it

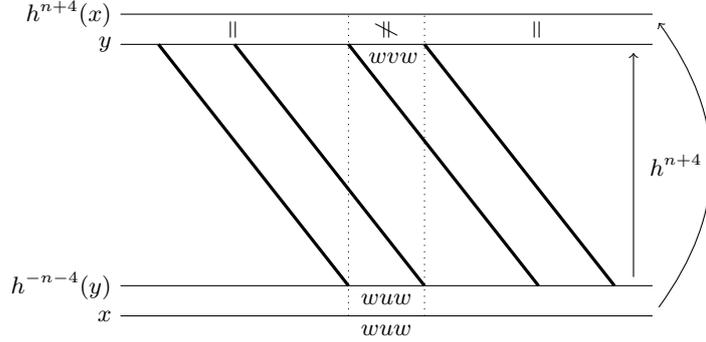


Fig. 4. A schematic spacetime diagram of h . The vertical lines represent configurations, and each thick line is a border of segments. The dotted lines mark the interval $[0, n + 3]$.

last read. This also applies to the regular language used in the definition of $U(w)$, where the small extra condition of not having left-moving particles in the neighboring coordinates was added.

The reversible cellular automaton we constructed above also has a maximally hard infinite time prediction problem, provided that we choose the machine M correctly.

Theorem 2. *There exists a chaotic reversible cellular automaton whose infinite-time prediction problem is Σ_1^1 -complete, even when restricted to the radius-one partition \mathcal{C}_1 .*

5 Further Discussion

The automaton h constructed in Theorem 1 and Theorem 2 is chaotic, but it is not *expansive*. A reversible cellular automaton $h : S^{\mathbb{Z}} \rightarrow S^{\mathbb{Z}}$ is expansive, if there exists $r \in \mathbb{N}$ such that for all distinct pairs $x \neq y \in S^{\mathbb{Z}}$, there exists $k \in \mathbb{Z}$ with $h^k(x)_{[-r,r]} \neq h^k(y)_{[-r,r]}$. Intuitively, this means that all discrepancies between two configurations are propagated to the left and to the right, if we consider both their past and future itineraries. Expansivity can thus be seen as an extreme form of sensitivity to initial conditions, and it makes sense to ask whether an expansive CA can be computationally universal.

Unfortunately, the trace shifts of expansive cellular automata are a deep and mysterious subject, and not much is known about them. In [3], it was shown that if h is an expansive CA, then all of its wide enough traces ($\tau_{h,n}$ for large enough $n \in \mathbb{N}$) have a property called *total chain transitivity*, which makes it difficult to find much structure in them. Also, in [14], it was shown that if h has memory or anticipation 0, then all wide enough traces are actually SFTs. It is currently unknown whether this holds for all expansive CA (as conjectured in [13]), but it would directly imply that their prediction problems are decidable, at least for a fixed clopen partition.

6 Acknowledgments

The authors are thankful to the anonymous referees for their valuable comments that helped to improve the quality and readability of this article.

References

1. Holger Bock Axelsen and Robert Glück. What do reversible programs compute? In *Foundations of software science and computational structures*, volume 6604 of *Lecture Notes in Comput. Sci.*, pages 42–56. Springer, Heidelberg, 2011.
2. C. H. Bennett. Logical reversibility of computation. *IBM J. Res. Develop.*, 17:525–532, 1973.
3. Mike Boyle. Some sofic shifts cannot commute with nonwandering shifts of finite type. *Illinois J. Math.*, 48(4):1267–1277, 2004.
4. Matthew Cook. Universality in elementary cellular automata. *Complex Systems*, 15(1):1–40, 2004.
5. Jean-Charles Delvenne, Petr Kůrka, and Vincent Blondel. Decidability and universality in symbolic dynamical systems. *Fund. Inform.*, 74(4):463–490, 2006.
6. Jean-Charles Delvenne, Petr Kůrka, and Vincent D. Blondel. Computational universality in symbolic dynamical systems. In *Machines, computations, and universality*, volume 3354 of *Lecture Notes in Comput. Sci.*, pages 104–115. Springer, Berlin, 2005.
7. Robert L. Devaney. *An introduction to chaotic dynamical systems*. Addison-Wesley Studies in Nonlinearity. Addison-Wesley Publishing Company Advanced Book Program, Redwood City, CA, second edition, 1989.
8. Bruno Durand and Zsuzsanna Róka. The game of life: universality revisited. In *Cellular automata (Saissac, 1996)*, volume 460 of *Math. Appl.*, pages 51–74. Kluwer Acad. Publ., Dordrecht, 1999.
9. Petr Kůrka. On topological dynamics of Turing machines. *Theoret. Comput. Sci.*, 174(1-2):203–216, 1997.
10. Ville Lukkarila. Sensitivity and topological mixing are undecidable for reversible one-dimensional cellular automata. Technical Report 927, TUCS, 2009.
11. Kenichi Morita and Masateru Harao. Computation universality of one-dimensional reversible (injective) cellular automata. *The Transactions of The IEICE*, E72-E(6):758–762, 1989.
12. Kenichi Morita, Akihiko Shirasaki, and Yoshifumi Gono. A 1-tape 2-symbol reversible turing machine. *The Transactions of the IEICE*, E72:223–228, 1989.
13. Masakazu Nasu. Textile systems for endomorphisms and automorphisms of the shift. *Mem. Amer. Math. Soc.*, 114(546):viii+215, 1995.
14. Masakazu Nasu. Textile systems and one-sided resolving automorphisms and endomorphisms of the shift. *Ergodic Theory and Dynamical Systems*, 28:167–209, 2008.
15. G.E. Sacks. *Higher recursion theory*. Perspectives in mathematical logic. Springer-Verlag, 1990.
16. Alvy Ray Smith, III. Simple computation-universal cellular spaces. *J. Assoc. Comput. Mach.*, 18:339–353, 1971.
17. Wolfgang Thomas. Automata on infinite objects. In *Handbook of theoretical computer science, Vol. B*, pages 133–191. Elsevier, Amsterdam, 1990.
18. Stephen Wolfram. Universality and complexity in cellular automata. *Phys. D*, 10(1-2):1–35, 1984. Cellular automata (Los Alamos, N.M., 1983).

Appendix

Proof of Theorem 2:

Proof. The CA $h : R^{\mathbb{Z}} \rightarrow R^{\mathbb{Z}}$ and the proof idea are exactly the same as in Theorem 1, but the machine M now needs to satisfy the claim of Lemma 2.

Now, let $w \in \{0, 1\}^*$, and define the Muller-recognizable language $L(w) \subset R^{\mathbb{N}}$ by the infinite regular expression

$$(B, \rightarrow, \square)(B, q_0, \square)(w \times (\rightarrow, \square)^{|w|})(\#, \rightarrow, \square)((0, \rightarrow, \square)^*(1, \rightarrow, \square))^\omega \quad (2)$$

It is similar to (1), except that the head of M is situated right next to the end of the segment, and after the input word w we may have an infinite tail of 0s and 1s, but no particles at all. As before, let $U(w) \subset (R^3)^{\mathbb{N}}$ be the set of configurations whose middle letters form a configuration in $L(w)$, and none of the letters of which contains a left-moving particle.

With a proof mimicking that of Theorem 1, we can now show that $\tau_{f,1} \cap U(w) \neq \emptyset$ if and only if there exists $u \in \{0, 1\}^{\mathbb{N}}$ such that M never halts on $w\#u$, and this is Σ_1^1 -complete by the choice of M . Namely, if M never halts, then the preimage of a configuration $x \in R^{\mathbb{Z}}$ containing the initial configuration of M with input $w\#u$ for some $u \in (0^*1)^\omega$ has its trace in $U(w)$ since no particles are ever introduced. Conversely, if such a configuration exists, then it necessarily simulates a non-halting computation of M , since a particle must be either created or destroyed at the time of halting, both of which are impossible. Furthermore, h is chaotic by the same argument as before. \square