

Cellular Automata Reversible over Limit Set^{*}

SIAMAK TAATI[†]

*Turku Centre for Computer Science, and
Department of Mathematics, University of Turku, Finland*

Reversibility of dynamics is a fundamental feature of nature, as it is currently believed that all physical processes are reversible in the ultimate microscopic scale. In this paper, we consider cellular automata (CA) whose dynamics are reversible when restricted to the limit set; i.e., those that obey reversibility in equilibrium.

We exploit standard topological and combinatorial arguments to show that the limit set, in this case, is a mixing subshift of finite type (SFT), and is reached in finite time. In one dimensional case, any mixing SFT which contains at least one homogeneous configuration, may arise this way. We also discuss the decidability of two related algorithmic questions.

Key words: Cellular Automata, Reversibility, Limit Set, Subshifts of Finite Type, Conservation Laws, Undecidability

1 INTRODUCTION

Reversibility is a widely accepted principle in physics, according to which, the microscopic laws governing the dynamics of the physical nature can hypothetically be reversed to let the system run backward in time. Reversible cellular automata (CA) have been extensively studied as convenient tools for modeling physical systems, whenever capturing the microscopic reversibility is desirable (see e.g. [16]).

^{*} This research was partially supported by the Academy of Finland grant 54102.

[†] email: staati@utu.fi

Although various elegant techniques have been devised (see e.g. [14, 15]), the task of designing a reversible CA with certain behavior still remains far from trivial. The fundamental reason behind this difficulty is that no algorithm could exist to answer whether a given (two-dimensional) CA is reversible or not [7].

Here we study a broader class of CA which can be used for modeling reversible processes. Namely, we let the CA undergo an irreversible transient, and require reversibility only when settled in the limit set. This is a natural relaxation of the original property, and as we are to show, it meets another course of generalizing the notion of reversible CA, i.e., the study of automorphisms of subshifts of finite type.

To further clarify the relevance of this notion, let us consider the following abstract model of particles moving in a one-dimensional lattice. Each cell in the lattice may take one of the states \blacktriangleleft , \blacktriangleright , or $-$, representing respectively, a particle moving to the left, a particle moving to the right, and an empty space. It is easy to verify that the rewriting rules in Figure 1(a) (together with their mirror images) uniquely define a CA with neighborhood radius 2. This CA




FIGURE 1
An abstract model of moving particles in 1d. Over its limit set, this CA is reversible and conserves the number of particles. (a) Transition rules (mirror images are omitted). (b) A sample space-time snapshot.

is not reversible, neither does it even conserve the number of particles, which seems disappointing for a model of physical particles. Yet, after one step, it settles in its limit set (i.e., the set of configurations with no occurrence of blocks $\blacktriangleleft\blacktriangleleft$ and $\blacktriangleright\blacktriangleright$), over which it is reversible, and conserves the number of particles (Figure 1(b)).

The two ends of the spectrum of the CA that are reversible over limit set,

are when the limit set contains only a single point, and when it includes the whole space. The former, gives a nilpotent CA; the latter an ordinary reversible one.

In Section 2 we fix the notations and review some required preliminaries. The reader may consult [11, 12] for a thorough treatment of the topic. The main results are given in Section 3, where we show that for the CA in question, the limit set is reached in finite time. This was known for the nilpotent CA, and is trivial for the reversible ones. An immediate consequence is that the limit set must be a mixing subshift of finite type. This gives rise to a characterization of the limit behavior in Section 4, which is based on some known results in symbolic dynamics. Alas, this characterization only works for one-dimensional CA. The higher dimensional case remains open.

In Section 5 we consider two algorithmic questions important in designing CA whose limit behavior are supposed to model a physical phenomenon. The first question is whether a given CA is reversible over its limit set, and the second asks whether it obeys a given conservation law over its limit set. Both questions are seen to be undecidable (even in the one-dimensional case) using the proof of the undecidability of nilpotency. Some further remarks and open problems are noted in Section 6.

2 NOTATIONS AND PRELIMINARIES

Let S be a finite set. An assignment $c : \mathbb{Z}^d \rightarrow S$ ($d \geq 1$) is called a *configuration* of the *lattice* \mathbb{Z}^d . Each point $i \in \mathbb{Z}^d$ is referred to as a *cell* of the lattice, and $c(i)$ is called the *state* of the cell i . d is the *dimension* of the lattice. For the sake of succinctness, let $\mathbb{T} = \mathbb{Z}^d$.

For each $a \in \mathbb{T}$, we define a *shift* operator $\sigma^a : S^{\mathbb{T}} \rightarrow S^{\mathbb{T}}$, which simply translates any configuration by a :

$$(\sigma^a c)(i) \triangleq c(a + i) \quad \text{for all } i \in \mathbb{T}$$

A partial configuration is a *pattern*. A *finite* pattern is one with finite domain. Whenever it is not ambiguous, we may consider the patterns modulo shifts, and loosely refer to the class of all patterns that can be obtained from $p : D \rightarrow S$ by some shift, as the pattern p .

Let A and B be finite sets. Given a finite set $N \subseteq \mathbb{T}$ and a function $\varphi : A^N \rightarrow B$, we can define a mapping $f : A^{\mathbb{T}} \rightarrow B^{\mathbb{T}}$ as follows: A configuration $x \in A^{\mathbb{T}}$ is mapped to a configuration $y = f(x) \in B^{\mathbb{T}}$, where

$$y(i) \triangleq \varphi((\sigma^i x)|_N) \quad \text{for all } i \in \mathbb{T}$$

(Here $g|_X$ denotes the restriction of a function g to a subset X of its domain.) The function φ is called a *local rule*, and N its *neighborhood*. When a mapping f is induced by a local rule, like above, we say that it is *local* (or, it is locally defined).

A *cellular automaton* (CA) consists of a finite state set S , a lattice \mathbb{T} , a finite neighborhood $N \subseteq \mathbb{T}$, and a local rule $\varphi : S^N \rightarrow S$. The set $S^{\mathbb{T}}$ of configurations together with the induced map $f : S^{\mathbb{T}} \rightarrow S^{\mathbb{T}}$ (the *global map* of the CA) form a dynamical system. We often identify the CA with its global map, and speak of the CA f .

The space of configurations $S^{\mathbb{T}}$ is naturally equipped with the product topology if we provide S with discrete topology. This space is compact (e.g., by Tychonoff's Theorem) and metrizable. Given a finite pattern $w : A \rightarrow S$, the set $[w]_A \triangleq \{x \in S^{\mathbb{T}} \mid x|_A = w\}$ is called a *cylinder*. Cylinders are both open and closed, and form a basis for the topology.

A compact and translation invariant subspace of $S^{\mathbb{T}}$ is called a *subshift*. Subshifts are exactly those subsets of $S^{\mathbb{T}}$ that can be defined by forbidding a collection of finite patterns: Given a collection F of finite patterns, we define the subshift

$$\mathcal{X}_F \triangleq \{x \in S^{\mathbb{T}} \mid (\sigma^a x)|_{\text{dom}(p)} \neq p, \text{ for all } a \in \mathbb{T} \text{ and } p \in F\}$$

Any subshift, has a representation of this form. For a subshift \mathcal{X} , let us define $L(\mathcal{X})$ as the set of all finite patterns that occur in configurations in \mathcal{X} . Because of compactness, we have $\mathcal{X}_{L^c(\mathcal{X})} = \mathcal{X}$. If F is finite, \mathcal{X}_F is called a *subshift of finite type* (SFT).

Continuous translation invariant maps between subshifts coincide with those that are induced by local rules. This is the Curtis-Hedlund-Lyndon Theorem [5]. Put it more clearly, for subshifts $\mathcal{X} \subseteq S^{\mathbb{T}}$ and $\mathcal{Y} \subseteq T^{\mathbb{T}}$, a mapping $f : \mathcal{X} \rightarrow \mathcal{Y}$ which commutes with all shift operators, is continuous, if and only if, there is a finite neighborhood N , and a local rule $\varphi : S^N \rightarrow T$ that generates f . (Here φ does not need to be total.) This again follows from a simple compactness argument. As a consequence, we get that, any bijective local map is a homeomorphism and hence is (locally) *reversible*; i.e., its inverse is also local. In particular, the inverse of a bijective CA is a CA.

Let (\mathcal{X}, Φ) be a dynamical system, with \mathcal{X} a compact topological space, and Φ a finitely generated semigroup of continuous transformations over \mathcal{X} . We say (\mathcal{X}, Φ) is (topologically) *mixing*, if for any two open sets $U, V \subseteq \mathcal{X}$, we have $\alpha U \cap V \neq \emptyset$ for all but finitely many $\alpha \in \Phi$.

Mixing property is invariant under epimorphisms. Namely, let (\mathcal{X}, Φ) and (\mathcal{Y}, Φ) be dynamical systems, and $h : \mathcal{X} \rightarrow \mathcal{Y}$ a surjective continuous map

that commutes with the action of Φ (i.e., an epimorphism). It is easy to see that if (\mathcal{X}, Φ) is mixing, so is (\mathcal{Y}, Φ) .

If we think of a subshift and the shift operators as a dynamical system, the mixing property can be expressed in terms of patterns: \mathcal{X} is mixing, if and only if, for any two patterns $p, q \in L(\mathcal{X})$ and all but finitely many shifts σ^a , there can be found a pattern $u \in L(\mathcal{X})$ that agrees with $\sigma^a p$ and q over their domains.

An element $p \in \mathbb{T}$ is called a *period* of a configuration $x \in S^{\mathbb{T}}$, if $\sigma^p(x) = x$. A configuration x is said to be (spatially) *periodic* if the set $\{\sigma^a x \mid a \in \mathbb{T}\}$ is finite. It is *homogeneous* if it is constant over \mathbb{T} .

Given a CA f , a configuration c is *temporally periodic* if $f^i(c) = c$ for some $i > 0$. It is *eventually* temporally periodic if $f^k(c)$ is temporally periodic for some $k \geq 0$.

Let $\varphi : S^N \rightarrow S$ be the local rule of a CA. For a state $s \in S$, let $q_s : N \rightarrow S$ be a pattern with $q_s(i) = s$ (for all $i \in N$). A state $o \in S$ with $\varphi(q_o) = o$ is called a *quiescent* state.

The *limit set* of a CA $f : S^{\mathbb{T}} \rightarrow S^{\mathbb{T}}$ is the intersection of all forward images of the space of configurations $\Lambda \triangleq \bigcap_{i \geq 0} f^i(S^{\mathbb{T}})$. The limit set is translation-invariant and compact, and so it is a subshift. A CA with singleton limit set is called *nilpotent*. The single element of Λ in this case is a homogeneous configuration in which all cells are quiescent.

One way to formulate the conservation laws in CA is in terms of relative invariance. An *additive quantity* is a function $\mu : S \rightarrow \mathbb{R}$ that assigns a real number to each state. For any finite set $A \subseteq \mathbb{T}$ of cells, define $M_A(x) \triangleq \sum_{i \in A} \mu(x(i))$, for all $x \in S^{\mathbb{T}}$. Given a subshift $\Gamma \subseteq S^{\mathbb{T}}$, we say an endomorphism $f : \Gamma \rightarrow \Gamma$ *conserves* the additive quantity μ , if for any two configurations $x, x' \in \Gamma$ that differ in at most a finite number of cells, there is a finite set $F \subseteq \mathbb{T}$, such that

$$M_A(f(x')) - M_A(f(x)) = M_A(x') - M_A(x)$$

for all $A \supseteq F$.

3 INJECTIVITY OVER LIMIT SET

It is known [3] that for any nilpotent CA there is a finite time when all cells go quiescent, regardless of the initial configuration; i.e., the CA reaches its singleton limit set after a certain finite time. We show that the same is true for every CA whose global map is injective restricted to the limit set. In what follows, let $\mathbb{T} = \mathbb{Z}^d$ where $d \geq 1$.

Let $f : S^{\mathbb{T}} \rightarrow S^{\mathbb{T}}$ be a CA, and Λ its limit set. Consider the chain

$$\Lambda \subseteq f^{-1}(\Lambda) \subseteq f^{-2}(\Lambda) \subseteq \dots \subseteq S^{\mathbb{T}}.$$

If for some k , $f^{-k}(\Lambda) = f^{-(k+1)}(\Lambda)$, then for all $m > k$ we will have $f^{-k}(\Lambda) = f^{-m}(\Lambda)$. In this case, we claim, we must also have $f^{-k}(\Lambda) = S^{\mathbb{T}}$. For, suppose on the contrary, that $\mathcal{X} \triangleq S^{\mathbb{T}} - f^{-k}(\Lambda)$ is non-empty. \mathcal{X} is the set of all configurations that, under iteration of f , never get inside the limit set. On the other hand, \mathcal{X} is an open set, and hence contains at least one spatially periodic configuration. But every spatially periodic configuration is eventually temporally periodic, and so cannot be in \mathcal{X} . This is a contradiction. Therefore, we have

Lemma 1. *Either $f^m(S^{\mathbb{T}}) = \Lambda$ for some $m \geq 0$, or for every $m \geq 0$ there is a configuration $c_{-m} \in S^{\mathbb{T}}$ that enters Λ after exactly m steps.*

Consider now the case that f is one-to-one over Λ . We show that there is an $m \geq 0$ such that $f^m(S^{\mathbb{T}}) = \Lambda$, i.e., the limit set is reached in finite time. If this is not true, according to Lemma 1, for any $m \geq 0$ there is a sequence $c_{-m}^{(m)}, c_{-m+1}^{(m)}, \dots, c_0^{(m)}$ of configurations in $S^{\mathbb{T}}$, such that $c_{t+1}^{(m)} = f(c_t^{(m)})$ for each $-m \leq t < 0$, and $c_0^{(m)} \in \Lambda$, but $c_t^{(m)} \notin \Lambda$ for $t < 0$. Since $c_0^{(m)}$ is in Λ , it also has a pre-image $e^{(m)}$ in Λ . Clearly $c_{-1}^{(m)} \neq e^{(m)}$. Without loss of generality (possibly using a proper translation), we can assume that they are different at position 0, i.e., $c_{-1}^{(m)}(0) \neq e^{(m)}(0)$. Consider the sequences

$$\{e^{(i)}\}_i \quad \text{and} \quad \{c_0^{(i)}\}_i, \{c_{-1}^{(i)}\}_i, \{c_{-2}^{(i)}\}_i, \dots$$

Since $S^{\mathbb{T}}$ is compact, we can choose $\mu : \mathbb{N} \rightarrow \mathbb{N}$, such that, all the subsequences

$$\{e^{(\mu(i))}\}_i \quad \text{and} \quad \{c_0^{(\mu(i))}\}_i, \{c_{-1}^{(\mu(i))}\}_i, \{c_{-2}^{(\mu(i))}\}_i, \dots$$

converge. Denote the limits by $e, c_0, c_{-1}, c_{-2} \dots$, respectively. Since f is continuous, it preserves the limits, i.e., $f(e) = c_0$ and $f(c_t) = c_{t+1}$ for any $t < 0$. It follows that e and c_{-1} are both in Λ . Furthermore, $f(e) = f(c_{-1})$, but $e \neq c_{-1}$, which means $f|_{\Lambda}$ is not one-to-one.

Lemma 2. *If the restriction of f to Λ is injective, there must be an $m \geq 0$ such that $f^m(S^{\mathbb{T}}) = \Lambda$.*

The restriction $f|_{\Lambda}$ is continuous, shift-invariant and onto. In case it is one-to-one, it turns out, from Curtis-Hedlund-Lyndon theorem, that it must

also be *reversible* over Λ ; i.e., $f|_{\Lambda}$ has a local inverse $g_0 : \Lambda \rightarrow \Lambda$. The local map g_0 can now be extended to a CA, by completing the transition table of its local rule, arbitrarily. In particular, the extension $g : S^{\mathbb{T}} \rightarrow S^{\mathbb{T}}$ defined by $g(x) = g_0^{m+1} \circ f^m(x)$ (for all $x \in S^{\mathbb{T}}$) has the convenient property that it has the same limit set as f . (In fact, $g(S^{\mathbb{T}}) = \Lambda$.)

Lemma 2 also implies that the limit set Λ in this case is a subshift of finite type. To see this, note that Λ is exactly the set of configurations $x \in S^{\mathbb{T}}$ that satisfy $g^m(f^m(x)) = x$.

The following theorem summarizes the above discussion.

Theorem 1. *If a CA $f : S^{\mathbb{T}} \rightarrow S^{\mathbb{T}}$ is injective on its limit set Λ ,*

- a) *there is a CA $g : S^{\mathbb{T}} \rightarrow S^{\mathbb{T}}$ with limit set Λ , such that $f \circ g(x) = g \circ f(x) = x$, for any $x \in \Lambda$.*
- b) *Λ is a subshift of finite type, and under the evolution of f , is reached in finite time.*

A closely related theorem [6] states that if the limit set of a CA is an SFT, it is necessarily reached in finite time: using compactness, one may show that each forbidden pattern of the limit set must be excluded at some particular time.

4 THE LIMIT BEHAVIOR

Let $f : S^{\mathbb{T}} \rightarrow S^{\mathbb{T}}$ be a CA, and Λ its limit set. Clearly, Λ must contain at least one homogeneous configuration. On the other hand, if f is reversible over Λ , we have $f^m(S^{\mathbb{T}}) = \Lambda$ for some $m > 0$. The full shift $S^{\mathbb{T}}$ is, trivially, mixing. Therefore, Λ should also be mixing.

In one-dimensional case, any surjective endomorphism f_s of a mixing SFT $\Lambda \subseteq S^{\mathbb{Z}}$ which has at least one homogenous configuration is in fact the restriction of a CA to its limit set [13]. The idea is simple: given an arbitrary configuration x , one first erases all the Λ -forbidden blocks from x . Large erased blocks can then be filled (up to a margin) by blocks from the homogeneous configuration. Now each two consecutive non-erased blocks can be glued together via connectors as in the definition of the mixing property for the subshifts. Since Λ is of finite type, the connector depends only on a bounded-length suffix of the left block and a bounded-length prefix of the right block. For any suffix and prefix, some canonical connectors should be fixed beforehand to make the process deterministic. It is straightforward to

realize these operations using a local mapping. The endomorphism f_s later may be applied to obtain a CA f which maps every configuration in $S^{\mathbb{Z}}$ to one in Λ , and which over Λ agrees with f_s .

Theorem 2. *Let $f : S^{\mathbb{Z}} \rightarrow S^{\mathbb{Z}}$ be a 1d CA which is reversible over its limit set Λ . Then, Λ must be a mixing SFT having at least one homogeneous configuration. Conversely, let $\Lambda \subseteq S^{\mathbb{Z}}$ be a 1d mixing SFT with at least one homogeneous configuration, and $f_s : \Lambda \rightarrow \Lambda$ any automorphism of Λ . Then, there is a CA f with limit set Λ , such that $f|_{\Lambda} = f_s$.*

In higher dimensions, one needs a stronger mixing property to ensure that an SFT is the limit set of a cellular automaton which is reversible over it. The following is an example of a 2d mixing SFT having a homogeneous configuration which cannot be the limit set of any such CA.

Example 1. Consider the binary SFT \mathcal{X} , defined over the lattice \mathbb{Z}^2 by forbidding the block $\begin{smallmatrix} \square & \blacksquare \\ \square & \blacksquare \end{smallmatrix}$. This subshift contains two homogeneous configurations. It is also mixing. To see this, note that any finite pattern $p : A \rightarrow \{\square, \blacksquare\}$ in $L(\mathcal{X})$ can be extended to one which contains only \square 's in its border. Let us denote this new pattern by \tilde{p} . Now, for any two patterns $p, q \in L(\mathcal{X})$, we can clearly glue \tilde{p} and \tilde{q} together to obtain a new pattern $w \in L(\mathcal{X})$, provided that we shift one of them far enough apart from the other so that their domains do not intersect.

Yet \mathcal{X} could not be the limit set of any CA which is reversible over its limit set. In particular, no matter how a CA maps the configuration c , defined by

$$c(i, j) \triangleq \begin{cases} \blacksquare & \text{if } i \geq 0, \\ \square & \text{if } i < 0, \end{cases}$$

it keeps its vertical period unchanged. Therefore, it can never get rid of the $\begin{smallmatrix} \square & \blacksquare \\ \square & \blacksquare \end{smallmatrix}$ blocks in finite time, unless it takes c to a homogeneous configuration, which leads to a violation of injectivity over \mathcal{X} . (In fact, using a slightly more complicated argument, one can show that \mathcal{X} cannot be the limit set of *any* CA.)

5 ALGORITHMIC ASPECTS

We now discuss two algorithmic questions concerning the global behavior of CA over their limit sets. Both questions turn out to be undecidable, taking advantage of the well-known proof of the undecidability of nilpotency. Again, let $\mathbb{T} = \mathbb{Z}^d$ ($d \geq 1$) be the lattice.

First, consider the algorithmic question, given a CA, whether it is reversible over its own limit set. Reversibility in general (over the whole space) is decidable in 1d [1], but undecidable in higher dimensions [7, 9]. On the other hand, the structure of the limit set is undetectable. In particular, for every $d \geq 1$, any non-trivial property of the limit set is undecidable [10].

Any nilpotent CA is trivially reversible over its limit set. Nilpotency of the CA is known to be undecidable for any $d \geq 1$ [8]. Not quite surprisingly, the proof is invariably effective here, to show that our problem is undecidable, too.

The original proof of the undecidability of nilpotency (for $d = 1$; note this implies the undecidability for $d > 1$), uses a reduction from a version of the tiling problem, called *NW-deterministic* tiling, which is in turn, shown to be undecidable [8]. The latter can be viewed as the emptiness question of a class of \mathbb{Z}^2 -SFTs: Given a partial function $\varphi : S \times S \rightarrow S$, define the subshift Γ_φ , as the set of all configurations $x \in S^{\mathbb{Z}^2}$ that satisfy $\varphi(x(i, j), x(i + 1, j)) = x(i, j + 1)$ for all $i, j \in \mathbb{Z}$. The problem asks if Γ_φ is empty.

A partial function φ can be extended to a total function

$$\psi : S \cup \{q\} \times S \cup \{q\} \rightarrow S \cup \{q\}$$

by $\psi(a, b) \triangleq \varphi(a, b)$, whenever $\varphi(a, b)$ is defined, and $\psi(a, b) \triangleq q$, otherwise. Now, it is easy to see that the 1d CA with local rule ψ is nilpotent, if and only if, Γ_φ is empty. On the other hand, whenever Γ_φ is not empty, the homogeneous configuration x_q with $x_q(i) \equiv q$ has infinitely many pre-images in the limit set of the CA. (Note that if y is another configuration in the limit set, any configuration that agrees with y on a finite interval and with x_q everywhere else must also be in the limit set.) Therefore, we have

Theorem 3. *Given a CA over $\mathbb{T} = \mathbb{Z}^d$ ($d \geq 1$), it is undecidable whether it is reversible over its limit set.*

Note that using Theorem 1, it is easy to see that our question is semi-decidable. Namely, in case a CA is reversible over its limit set, the limit set can be detected in finite time, and one can also recognize the existence of an inverse rule, simply by enumerating all local maps as candidates.

For the next question, we are given a CA, and a proposed conservation law. We are, then, asked whether that conservation law holds for the CA in equilibrium. It is known that, given a CA f and an additive quantity μ , it is decidable whether f conserves μ over the whole space [4]: Note that one only needs to verify the relative invariance criteria for the cases that the two configurations x and x' differ in only one cell (cf. [11]).

Here, however, we are concerned about the question, whether f conserves μ over its limit set. This is undecidable, again following the same line as in the proof of the undecidability of nilpotency. Note, a nilpotent CA conserves anything on its limit set. For the 1d CA with local rule ψ as constructed above, the additive quantity μ defined with $\mu(q) \triangleq 0$ and $\mu(a) \triangleq 1$ (for $a \neq q$), is conserved on the limit set, if and only if the CA is nilpotent.

Theorem 4. *Given a CA f over $\mathbb{T} = \mathbb{Z}^d$ ($d \geq 1$) and an additive quantity μ , it is undecidable whether f conserves μ over its limit set.*

6 CONCLUSION AND REMARKS

We studied the class of CA that are reversible over their limit sets. Since they inevitably reach their limit sets in finite time, they are well-justified to be called *eventually reversible*. The concept is comparable to that of CA that eventually obey some conservation law [2].

An important open problem here is to find a characterization of the class of subshifts that arise as the limit set of such CA. The limit dynamics, then, can be determined by any automorphism of such a subshift. We know that the limit set must be a mixing SFT, and contains at least one homogeneous configuration. This provides the required characterization in 1d, but is shown to be insufficient in higher dimensions.

Another open issue is to understand the transient behavior. Nilpotent CA may be found useful in this respect, as their non-transient behavior is trivial. Can the transient of any eventually reversible CA be somehow reduced to that of a nilpotent CA?

As a final remark, note that although we presented our main result for the ordinary checker-board lattices, it is valid for a large class of interesting lattices, including the 2d triangular and hexagonal lattices, and all kinds of crystalline lattice structures of the 3d space.

The definitions in Section 2 do not need to be changed if we substitute \mathbb{T} with an arbitrary finitely generated group. The Curtis-Hedlund-Lyndon Theorem likewise holds on any group \mathbb{T} . The bottleneck in Section 3 is when we argue that any open set in $S^{\mathbb{T}}$ has a spatially periodic point. This is valid whenever for any finite set $F \subseteq \mathbb{T}$ the group \mathbb{T} has a finite factor \mathbb{G} with homomorphism $h : \mathbb{T} \rightarrow \mathbb{G}$, such that the restriction $h|_F$ is one-to-one. Not all groups have such a property, as there exist infinite simple groups. Yet, the counter-examples are weird structures that do not naturally appear in CA theory.

The undecidability results in Section 5 can be extended to any finitely generated group that has a copy of \mathbb{Z} as subgroup. Again, the negative answer to Burnside's problem shows that this is not the case for all finitely generated groups, though the counter-examples are obscure.

7 ACKNOWLEDGMENTS

I am indebted to Jarkko Kari, for his encouragement and support, and for educative discussions. Eugen Czeizler kindly accepted to present the work in AUTOMATA'05 workshop, for which I am grateful.

REFERENCES

- [1] S. Amoroso and Y. N. Patt. (1972). Decision procedures for surjectivity and injectivity of parallel maps for tessellation structures. *Journal of Computer and System Sciences*, 6:448–464.
- [2] Nino Boccara. (2005). Eventually number-conserving cellular automata. To be published.
- [3] Karel Culik II, Jan Pachl, and Sheng Yu. (1989). On the limit sets of cellular automata. *SIAM Journal on Computing*, 18(4):831–842.
- [4] Tetsuya Hattori and Shinji Takesue. (1991). Additive conserved quantities in discrete-time lattice dynamical systems. *Physica D*, 49:295–322.
- [5] G. A. Hedlund. (1969). Endomorphisms and automorphisms of the shift dynamical system. *Mathematical Systems Theory*, 3(4):320–375.
- [6] Lyman P. Hurd. (1990). Recursive cellular automata invariant sets. *Complex Systems*, 4:119–129.
- [7] Jarkko Kari. (1990). Reversibility of 2D cellular automata is undecidable. *Physica D*, 45:379–385.
- [8] Jarkko Kari. (1992). The nilpotency problem of one-dimensional cellular automata. *SIAM Journal of Computing*, 21(3):571–586.
- [9] Jarkko Kari. (1994). Reversibility and surjectivity problems of cellular automata. *Journal of Computer and System Sciences*, 48(1):149–182.
- [10] Jarkko Kari. (1994). Rice's theorem for the limit sets of cellular automata. *Theoretical Computer Science*, 127:229–254.
- [11] Jarkko Kari. (2005). Theory of cellular automata: A survey. *Theoretical Computer Science*, 334:3–33.
- [12] Douglas Lind and Brian Marcus. (1995). *An Introduction to Symbolic Dynamics and Coding*. Cambridge University Press.
- [13] Alejandro Maass. (1995). On the sofic limit sets of cellular automata. *Ergodic Theory and Dynamical Systems*, 15:663–684.
- [14] Norman Margolus. (1984). Physics-like models of computation. *Physica D*, 10:81–95.
- [15] Kenichi Morita and Masateru Harao. (1989). Computation universality of one-dimensional reversible (injective) cellular automata. *Transactions of IEICE*, E 72:758–762.
- [16] T. Toffoli and N. Margolus. (1990). Invertible cellular automata: A review. *Physica D*, 45:229–253.