Complex Systems **9** (1995) 235–250

Some Results on Invertible Cellular Automata

A. Clementi*
Department of Computer Science, Università di Roma "La Sapienza", Italy

P. Mentrasti†
Department of Mathematics, Università di Roma "La Sapienza', Italy

P. Pierini‡
MIT Laboratory for Computer Science, Cambridge, MA, USA

Abstract. We address certain questions concerning invertible cellular automata (ICA), and present new results in this area. Specifically, we explicitly construct a cellular automaton (CA) in a class (*residual class*) previously known not to be empty only via a nonconstructive existence proof. This class contains CAs that are invertible on every finite support but not on an infinite lattice. Moreover, we show a class that contains ICAs having bounded neighborhood, but whose inverses constitute a class of CAs for which there is no recursive function bounding all the neighborhoods.

1. Introduction

Computational models satisfying physical laws are the subject of several recent studies [3, 8]; of particular interest are *invertible* models [5, 8]. Cellular automata (CAs) represent one of the best models of parallel computation; the study of invertibility in CAs is of great interest in modelling physics.

Several theoretical results concerning invertibility in CAs have been presented [2, 9, 10, 12, 13, 15, and 18], some leading to open questions.

- In [18], the existence of a peculiar class (*residual class*) of CAs had been predicted but, until now, no such CAs had been exhibited. Here we explicitly construct a CA in this class, that is, a CA that is invertible on every finite support but is not invertible on an infinite support.

*Electronic mail address: `clementi@sci.uniroma1.it`.
†Electronic mail address: `mentrasti@sci.uniroma1.it`.
‡Electronic mail address: `pp@im.lcs.mit.edu`, `pierini@sci.uniroma1.it` (supported by the Italian National Research Council).

- It is known [18] that for the class of all ICAs, an upper bound to the radius of the inverses cannot be found. We investigate the meaning of this constraint, exhibiting a class of ICAs with inverse local maps having neighborhoods that cannot be bounded by any recursive function.

We construct these CAs starting from a space tiling technique introduced in [14]. More precisely, using a variant of their technique, we discuss a particular set of local maps that had first been presented in [9, 10] and we prove that this set has the above mentioned properties.

2. Infinite cellular automata

2.1 Cellular automata

A *cellular automaton* is a set of identical finite automata (also called *cells*) locally connected to each other in a uniform way. In this paper, we consider two kinds of CAs, depending on their *support*, that is, the grid containing the cells.

- If the cells are located in the infinite d-dimensional lattice (i.e., $\mathcal{Z}^d$), we have a *proper* CA.
- If the support is a d-dimensional toroidal array (*torus*) of period (or *size*) n along all dimensions (i.e., $\mathcal{Z}_n^d$), we have a *toroidal* CA.

The state q of each cell varies according to a uniform, deterministic *local function* defined on the set of neighborhood states. The *neighborhood*, a set N of displacements, specifies the relative positions (with respect to the cell to be updated) of the cells used by the local function. By *radius* of a CA, we mean the radius of its neighborhood. In this paper, we use the *Moore neighborhood*, consisting of a cell and the eight adjacent cells in a two-dimensional grid.

Hence, a complete description of a CA (both for the infinite and the finite case) can be given by defining:

- the support space,
- the state set Q of the cells,
- the neighborhood N, and
- the local function $f : Q^{|N|} \to Q$.

The pair $\lambda = \langle$neighborhood, local function$\rangle$ will be called a *local map* or *rule*. The cells change their states in a parallel, synchronous way. The local function determines a global function F acting on the space Σ of all possible configurations.

2.2 Invertibility

A CA is *invertible* if its global function is bijective. The invertibility of a CA is an important issue in modelling reversible physical phenomena. Here we give a brief summary of the main results about the invertibility of CAs.

The invertibility of a CA is a property of its global function, while the CA itself is described in terms of a local map; in [13] it is proved that the bijectivity (i.e., invertibility) of a CA's global function implies the existence of an *inverse local map*. Note that even if every global state is the successor of another state, a CA is not invertible if there is a global state having more than one possible predecessor.

Theorem 2.1. [13] *If the global map of a CA is injective, then it is invertible, and its inverse is the global map of a CA as well.*

In other words, if the global map is injective then it is also surjective and the inverse global process can be described in local terms. The proof in [13] is not constructive: it does not give any procedure for finding the inverse local map. However, given two CAs it is possible to decide whether they are inverses of each other by using Lemma 2.1.

Lemma 2.1. [18] *There is an effective procedure for deciding, for any two local maps λ and λ' defined on the same set of configurations, whether the corresponding global maps F and F' are inverses of one another.*

Early investigators conjectured that ICA could not be computational universal [1, 4]. In [15] the existence of universal invertible d-dimensional CAs when $d > 1$ is proved; in [12] the existence of computational universal CAs in the one-dimensional case is proved.

For many years a major challenge has been deciding whether or not a given CA is invertible. For the one-dimensional case Theorem 2.2 is proved in [2].

Theorem 2.2. [2] *There is an effective procedure for deciding whether or not an arbitrary one-dimensional CA, given in terms of a local map, is invertible.*

In other words, the class of invertible one-dimensional CAs is recursive. The class of multidimensional ICAs is recursively enumerable (see [18] for a proof). However, it was recently proved [9, 10] that, for d greater than one, the class of invertible d-dimensional CAs is not recursive.

Theorem 2.3. [9, 10] *There is no effective procedure for deciding whether or not an arbitrary two-dimensional CA, given in terms of its local map, is invertible. Thus, in general, the invertibility of a CA is undecidable.*

The proof of Theorem 2.3 is based on transformation from another undecidable problem: the tiling problem on the infinite two-dimensional lattice [14].

The invertibility of a CA considered on toroidal finite supports has been proved to be co-NP-complete [6]; the same result of completeness arises also in average-case complexity theory [7].

Along with these theoretical results, there have been technological, architectural, and algorithmic developments (e.g., [11, 16, 17, and 19]); as a result, CAs have become a very productive tool for modelling and executing parallel computation.

3. A tiling technique

In this section we summarize some results and definitions used in [10] to prove Theorem 2.3. In particular, since it is used for proving our results, we briefly describe a finite set of tiles having the properties that they cover an infinite two-dimensional grid and that they define a path through all the elements of this grid. Moreover, we prove that this set of tiles cannot be used to tile a finite toroidal support of any size.

A *tile* is a square with colored edges. Formally, given a finite set C of colors, a *tile set* is a subset $\tau \subset C^4$ and a τ-*tile* is any ordered quadruple t of C^4.

A *tiling* (denoted as τ-tiling) of a fixed grid (*support space*) using the set τ, is a mapping from the sites of the grid to the set of tiles.

By *correct* tiling we mean a tiling such that edges of adjacent tiles have the same color.

Using colors, labels, or numbers to distinguish different kinds of tile edges is just a matter of convention. In what follows we adapt the concept of edge color as in [10].

We replace each color with one or more *arrows* pointing inwards or outwards from an edge. In other words, each edge is tagged by the *heads* and *tails* of different arrows (see Figure 1). Two adjacent edges match correctly if each head meets, in the adjacent tile, the tail of an identical arrow.

Let us further generalize the concept of color by assigning labels to the corners as well. If we call the four *corners* of a tile NE, NW, SW, and SE (see Figure 2(a)), a *passage* is a pair $\langle a_{\text{in}}, a_{\text{out}}\rangle$, where both a_{in} and a_{out} (which denote, respectively an *inward direction* and an *outward direction*) belong to the set {NE,NW,SW,SE} (see Figure 2(b)). The corner NW is naturally called *opposite* to SE and NE is opposite to SW.

With this formulation, the concept of "tile color" has been generalized to one of "arrows and passages." According to these definitions, a tiling is correct if both of the following are true.

- Each arrow head meets an equal arrow tail.
- For each passage $\langle a_{\text{in}}, a_{\text{out}}\rangle$ the neighbor tile in the a_{out} direction is "colored" with a passage whose inward direction is the opposite of a_{out}.

If one of these two conditions does not arise, we say that a *tiling error* occurs.

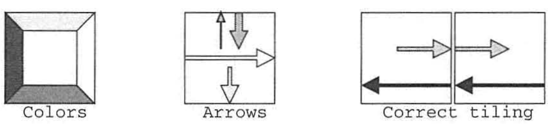

Figure 1: Colors are replaced by arrows.

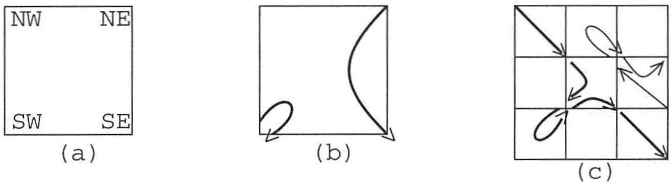

Figure 2: (a) Corners. (b) ⟨SW,SW⟩ and ⟨NE,SE⟩ passages. (c) Paths.

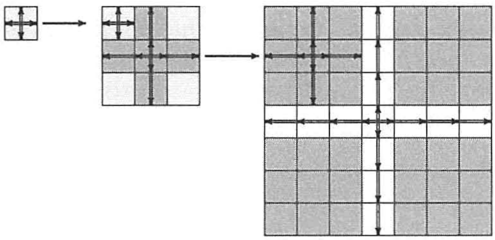

Figure 3: Construction of a correct tiling for a square of side 7.

We define a *path* as a sequence (possibly cyclic) of consecutive passages $p_1, \ldots, p_i, \ldots$ associated with neighboring cells, such that the outward direction of p_i is equal to the opposite of the inward direction of p_{i+1} (see Figure 2(c)).

A complete description of the set τ_{K} of 160 tiles defined in [10] is given in appendix A. This set has three very interesting properties expressed by Lemmas 3.1 through 3.3.

Lemma 3.1. [10] *An arbitrarily large square grid can be correctly tiled with the set τ_{K}. From Koenig's infinity lemma, one can then correctly tile the entire plane $\mathcal{Z}^2$.*

The correct tiling of arbitrarily large square grids (*squares* for simplicity) is obtained by a recursive construction that, given the correct tiling of a square of side $2^n - 1$, determines the correct tiling of a square of side $2^{n+1} - 1$. A schematic representation of this construction is shown in Figure 3.

Lemma 3.2. [10] *In every tiling of the plane from the set τ_{K}, only two types of path may arise:*

- *either the path has a tile for which there is a tiling error in its neighborhood, or*
- *the path visits all the tiles of an arbitrarily large square.*

The property of tiling arbitrarily large squares is trivial to achieve; it is the property expressed in Lemma 3.2 that makes this set of tiles really interesting. In other words, whatever square we tile with the set τ_K, either this tiling induces a continuous path through all the tiles, or the fact that the path will not cover the entire square is locally recognizable through a tiling error. The path induced by the passages has the shape of the Hilbert curve (see Figure 11). We remark that the set of tiles τ_K is independent of the size of the square that ones wants to tile.

The set τ_K resembles that presented in [14] as an example of tiles that permit only nonperiodic tiling: neither the tiles in [14] nor the set τ_K can be used to tile correctly a torus. In fact, referring to Figure 3, a tiling error must occur when the space is wrapped around by joining together opposite edges since some arrow heads will then meet other heads instead of tails. In Lemma 3.3 we formalize this result, using the construction of $(2^n-1)\times(2^n-1)$ squares with the set τ_K given in appendix A.

Lemma 3.3. *The set of tiles τ_K does not permit correct tiling for tori of any size.*

Proof. (By contradiction.) Suppose that there is a correct tiling for a torus of size n. This tiling is equivalent to a correct *periodic* tiling of period n for the infinite lattice. Let t be a single cross of the torus (which always exists, since in every correct tiling each 2×2 square must have a single cross). From Lemma A.1, t is in a XY-$(2^m - 1)$-square constructed as explained in appendix A; since the entire plane is tiled correctly, m is as large as we want. Thus we can find the smaller k such that there exists a XY-$(2^k - 1)$-square including our planar representation of the torus (see Figure 4). Since $n > 2^{k-1} - 1$, the torus must include the central cross of the square (A in figure) as well as part of the arms leaving from it; thus, when the space is wrapped around, a tiling error is encountered. ■

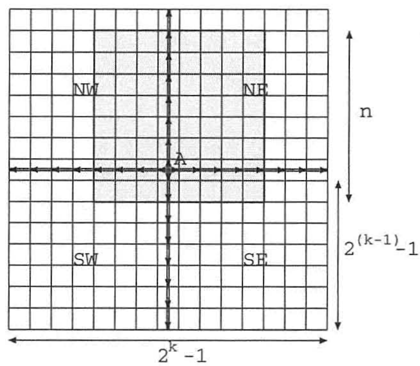

Figure 4: A correct tiling for a $2^{k-1}\times 2^{k-1}$ square including the planar representation of a torus (shaded).

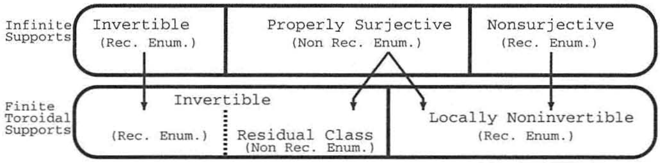

Figure 5: Invertibility for CA.

4. A cellular automaton in the residual class

Here we exhibit the first constructive example of CA in the residual class.

As described in [18], a CA on an infinite support can be:

- injective and surjective (*invertible*);
- not injective and not surjective (*nonsurjective*); or
- surjective but not injective (*properly surjective*).

It is impossible for a CA to be injective but not surjective [13]. When the local map is considered on a finite support, the three classes listed above have the following behaviors (see Figure 5):

- ICAs remain invertible;
- nonsurjective CAs remain nonsurjective; and
- properly surjective CAs can yield either invertible or nonsurjective finite CAs.

The class of CAs that are invertible on every finite support but noninvertible on an infinite support is called the *residual class*. The residual class cannot be empty [18] but no examples of CAs in this class have been shown until now.

4.1 A specific noninvertible cellular automaton

Using the set of tiles defined in section 3, we construct a CA in the residual class.

We tag each passage of the tiles τ_K with a binary digit; let us consider the CA having as states these modified tiles and performing, at each step, the XOR between bits of adjacent cells along the path induced by the passages. Formally we have Definition 4.1.

Definition 4.1. *The two-dimensional CA A_K is defined by the following.*
<u>States</u>: Each state consists of two components: a tile $t \in \tau^K$ (tile component) and one bit (0 or 1) for each passage of t (bit components).
<u>Neighborhood</u>: Moore.
<u>Local function</u>: The local function does not change the tile component.

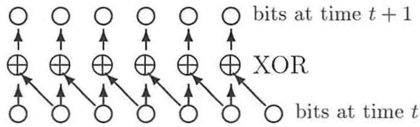

Figure 6: Bits along a path.

- *If there is no tiling error in the neighborhood of a cell, then the function performs the XOR between the bit of each passage of the cell and the bit of the corresponding next passage in the path (see Figure 6).*
- *If there is a tiling error, the bits remain unchanged.*

A_K is noninvertible and has an interesting property that is stated in Lemma 4.1.

Lemma 4.1. [10] *The CA* A_K

1) *is noninvertible on the infinite support* $\mathcal{Z}^2$*; and*
2) *for each pair of distinct configurations* c_1 *and* c_2 *having the same image under the update function (i.e.,* $F(c_1) = F(c_2)$*), the tile component of both* c_1 *and* c_2 *constitutes a correct tiling of the plane (this property has been called* almost injectivity *in [10]).*

Proof. From Lemma 3.1 we can choose two configurations c_1 and c_2 with the tile components constituting a correct tiling of the plane; thus the local function performs the XOR between each bit and the next bit in the space covering path. If the bit components are set to 0 in c_0 and to 1 in c_1, the images of both configurations under the update function coincide with c_0 itself; hence A_K is not invertible.

If $c_1 \neq c_2$ and $F(c_1) = F(c_2)$, there must be a cell $\vec{x}$ whose states q_1 and q_2, respectively in c_1 and c_2, are different. The local function f must change at least one of q_1 and q_2; however, the tile components of these states must be the same, since they are not changed by f. According to Definition 4.1, there cannot be a tiling error in the neighborhood of $\vec{x}$ (otherwise the bits would remain unchanged), f computes the XOR between bits of adjacents passages, thus also the bits in the cell that follows $\vec{x}$ along the path, must have different states. By iterating such a process, since a correct path visits every cell (Lemma 3.2), we see that there cannot be any tiling errors anywhere. Thus the tile component of c_1 and c_2 must constitute a correct tiling of the plane. ∎

4.2 From infinite to finite support

Here we prove that the noninvertible CA A_K in Definition 4.1 becomes invertible when considered on finite toroidal supports. Thus A_K is a constructive example of a CA in the *residual class*.

The noninvertibility of A_{K} on an infinite support comes from the noninjectivity of the XOR operator; if the path, defined by the passages, is an infinite one, we cannot know the predecessor of a configuration (see Figure 6). Nonetheless, if we could know the predecessor of at least one cell then we could also determine the predecessor of the cell that follows in the path, and iteratively determine all the other cell predecessors. The CA then becomes invertible.

The invertibility of A_{K} on a finite support follows from the fact that, on a torus, the set of tiles τ_{K} always leads to a tiling error (Lemma 3.3): from a given configuration, we can reconstruct the previous one starting from a cell (there is at least one) where the local function does not change the state since the tiling is not correct. Formally we have Theorem 4.1.

Theorem 4.1. *The two-dimensional CA A_{K} is in the residual class.*

Proof. We still call A_{K} the toroidal CA obtained from the local map of A_{K} in Definition 4.1. From Lemma 3.3, for any n, for any configuration of A_K on the toroidal support $\mathcal{Z}_n{}^2$ there must be a tiling error due to the tile component of the states. From the second property of Lemma 4.1, given two configurations c_1 and c_2 such that $F(c_1) = F(c_2)$, it must be that $c_1 = c_2$, otherwise the tiling would be correct. Thus A_{K} is injective on any toroidal support. ■

The proper surjectivity of A_{K} (see Figure 5) follows from the the fact that A_{K} is noninvertible on infinite support and becomes invertible on finite supports. Independent of this result, Theorem 4.2 can be proved.

Theorem 4.2. *The two-dimensional CA A_{K} is properly surjective.*

Proof. By Lemma 4.1, A_{K} is not injective on the infinite support $\mathcal{Z}^2$. However, A_{K} is surjective; indeed, let c be a configuration of A_{K} and n be a positive integer; on any path $p = \langle p_1, \ldots, p_n \rangle$ of length n (passing through the cells $i = 1, \ldots, n$) we denote as $\langle x_1^t, \ldots, x_n^t \rangle$ the bits associated with the passages in p at time t; then, all possible situations can be easily reduced to the following two cases.

1. The cells $i = 1, \ldots, n-1$ have no tiling errors and cell n has a tiling error; then, in the predecessor of c, x_n must have the same value of that in c (see Definition 4.1). Moreover, for any $t > 0$ and $i = 1, \ldots, n-1$, in the XOR function

   $$x_i^t = x_i^{t-1} \oplus x_{i+1}^{t-1}$$

   the value x_i^{t-1} is uniquely determined by x_i^t and x_{i+1}^{t-1}. From these facts, by a backward-iterative procedure, it is easy to correctly define all the bit values of p in the predecessor of c.

2. None of the cells $i = 1, \ldots, n$ have tiling errors; then, by Theorem 3.2, the path p continues into a cell $n+1$ having a bit x_{n+1} associated with the passage p_{n+1}. Then, for any value of x_{n+1}, it is not hard to determine all bit values of p in the predecessor of c by making use of the same procedure mentioned in the first case.

The proof is completed by observing that the local function f^K (and thence the global one) is different from the identity only on the bits of passages. ■

5. Unbounded neighborhood

Here we define a class of ICAs for which the neighborhood of the inverse is not bounded by any recursive function (*nonreciprocal* property). From [18] we have Theorem 5.1.

Theorem 5.1. [18] *There cannot exist a recursive function $f(\lambda)$ defined on the local maps of the two-dimensional CAs and bounding the neighborhoods of all the ICAs.*

Proof. If this function existed, given a local map λ we could sequentially generate all the local maps with a radius bounded by $f(\lambda)$, and by Lemma 2.1 we could check if one of these maps is the inverse of λ. On reaching $f(\lambda)$, either we found an inverse or we can conclude that λ is not invertible. But this contradicts the undecidability of CA invertibility (Theorem 2.3). ■

This result can be easily extended to any class of CAs for which it is undecidable whether a CA is invertible. Thus, the class of CAs used in [10] to prove Theorem 2.3 has the *nonreciprocal* property.

Showing a class that is "small" and still has the nonreciprocal property is useful in understanding the nature of the theoretical result of Theorem 5.1. We give a different proof, without using the result in Theorem 2.3, of the nonreciprocal property for the class introduced in [10], emphasizing the reason of this property. Let us now define such a class in a formal way.

Definition 5.1. *For any tile set τ, we consider the CA $A_K{}^{\tau}$ defined in the following way.*

States: Each state is a pair $\langle t, q\rangle$ where $t \in \tau$ and q is an element of the state set defined in Definition 4.1.

Neighborhood: Moore.

Local function: The local function operates as the function of A_K except that, when checking for tiling errors, it also considers the state component given by the set τ.

It can be easily proved, as done in [10] for different goals, that if a tile set τ_{error} does not admit a correct tiling for the plane $\mathcal{Z}^2$ (that is, it is a NO-instance of the tiling problem), then the corresponding CA $A_K{}^{\tau_{\text{error}}}$ is invertible.

Lemma 5.1. *For any NO-instance τ_{error} of the tiling problem, the corresponding CA ${A_K}^{\tau_{\text{error}}}$ is invertible.*

Proof. If we suppose, by contradiction, that ${A_K}^{\tau_{\text{error}}}$ is noninvertible, by the same reasoning used for proving the second statement of Lemma 4.1, we can prove that τ_{error} admits a correct tiling of the plane. But this is false, and thus the lemma is proved. ■

If we consider any possible NO-instance τ_{error} of the tiling problem and we construct the corresponding CA ${A_K}^{\tau_{\text{error}}}$, we then obtain the following class:

$$\text{INV} = \{{A_K}^{\tau_{\text{error}}} \ : \ \tau_{\text{error}} \text{ is a NO-instance of the tiling problem}\}.$$

From Lemma 5.1, the class INV consists of CAs that are invertible. However, in order to explicitly obtain the predecessor of a cell $\vec{x}$ we must follow the path originating from $\vec{x}$ until a tiling error is encountered. When we find a tiling error in a cell $\vec{y}$, since the local function in $\vec{y}$ is the identity, we know the predecessor of $\vec{y}$; by going backwards along the path, we can find the predecessor of $\vec{x}$. This is the only way to construct the inverse; it follows that the radius of the neighborhood must be large enough to recognize the nearest tiling error (this last property will be formally proved in Theorem 5.2). But the tiling error generated by an arbitrary NO-instance τ_{error} of the tiling problem cannot be bounded by a recursive function. Indeed, for the tiling problem we have a result similar to that of Theorem 5.1.

Lemma 5.2. *There cannot exist a recursive function $g(\tau)$ defined on the set of instances of the tiling problem and bounding the maximum distance between two errors (or, equivalently, bounding the distance between one point and its nearest tiling error) in a τ_{error}-tiling of any possible NO-instances τ_{error}.*

Proof. The proof is by contradiction. if this function existed, given any tile set τ we would then be able to decide the tiling problem by generating all possible τ-tiling of size at most $g(\tau)$. ■

Since, by Lemma 5.2, we cannot predict where a tiling error eventually occurs, the radii of the inverses of the CAs in INV cannot be bounded by any recursive function. Theorem 5.2 formalizes this result.

Theorem 5.2. *The radius $|N^{-1}|$ of the inverses of the CAs in INV cannot be bounded by any recursive function.*

Proof. Let us denote by f^{-1} the local function of the inverse of a CA in INV. By Lemma 5.2, it is sufficient to prove that N^{-1} must always contain at least one tiling error; that is, every cell, in the backward evolution, must have at least one tiling error in its neighborhood in order to compute its next state. Indeed, let us assume by contradiction that N^{-1} does not always contain a tiling error (with respect to the τ_K component or the τ_{error} component). Thus, there exists a configuration in which a cell $\vec{x}$ has a correct tiling (for both τ_K and τ_{error} components) in its neighborhood. For simplicity we call N^{-1} the

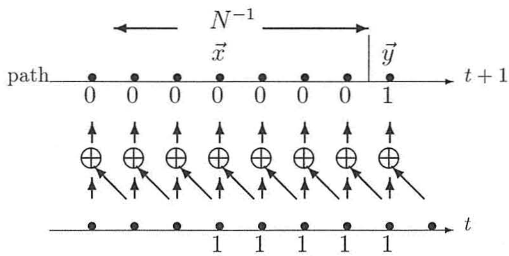

Figure 7: A tiling error in the neighborhood of the inverse.

neighborhood of $\vec{x}$. By Lemma 3.2, the path passing on $\vec{x}$ must leave this neighborhood and reach a cell $\vec{y}$ which is adjacent to N^{-1}. Without loss of generality, let us suppose that there is a tiling error (there must always exist one) in the cell $\vec{y}$ (more precisely, in the part of its neighborhood which is outside N^{-1}); thus the bits in $\vec{y}$ are not changed by the local function. Let us consider the configuration in which all bits in N^{-1} are 0. The predecessor of $\vec{x}$ in the forward evolution of the CA is $f^{-1}(N^{-1})$; let us suppose that this value is 0 and consider the configuration in which the bit in $\vec{y}$ is 1 (see Figure 7). Under these conditions, all the bits that preceed $\vec{y}$ in the path that goes from $\vec{y}$ to $\vec{x}$ must be 1, also, bits in $\vec{x}$ must be 1, but this is a contradiction. Similar arguments can be applied if we suppose $f^{-1}(N^{-1}) = 1$. ■

6. Conclusions

The existence of families of invertible toroidal CA having an inverse local map with large and complex interactions could determine a set of one-way functions having practical applications in cryptography. Indeed, knowledge of the direct local map (the cryptor) does not give sufficient information (to the cryptoanalyst) on the inverse local map (the decryptor) (e.g., [7]). In terms of dynamical system theory, the results shown in this paper imply the existence of reversible dynamical systems having local and simple interactions but whose inverses have almost-global interactions.

Acknowledgments

We thank Tommaso Toffoli for several comments and helpful suggestions on this paper.

Appendix A. Kari's tiling technique

Here we describe the basic structure of the tiles in the set τ_K defined by Kari and used in our results. We will thus follow the terminology and notations adopted in [10].

<u>*Arrows*</u> Different kinds of arrows (see section 3) are distinguished by drawing them in different ways and by labelling them with different tags. Thus

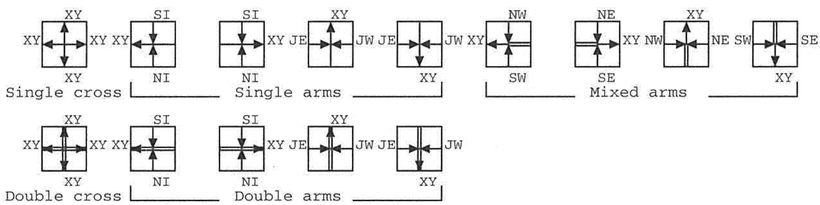

Figure 8: Labelled arrows.

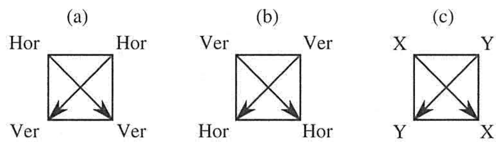

Figure 9: The diagonal arrows on (a) horizontal arms, (b) vertical arms, and (c) crosses ($X, Y \in$ {Hor,Ver}).

we have the set of *labelled arrows* shown in Figure 8. Each of these tiles also has two *diagonal arrows* as shown in Figure 9. This set permits the recursive construction of correct tilings for arbitrarily large squares (see Figure 10). The diagonal arrows force the horizontal and vertical arms to alternate on each diagonal row of tiles.

Passages If we denote a tile by the label of its arrow; we have the following.

- Double crosses must have the passages {(NW,NE),(NE,SE), (SE,SW),(SW,NW)}.
- Single crosses must have one of the following six passage sets: {(XY,XY)} such that XY ∈ {NE,NW,SE,SW}; {(NW,SE),(SE,NW)}; or {(NE,SW), (SW,NE)}.
- No passage for any type of arms.

In Figure 11 the path induced by a correct tiling on a square of dimension 7 is shown.

In a $(2^n - 1) \times (2^n - 1)$ square correctly tiled by the recursive costruction sketched in Figure 3, the tile in the middle is always a *double cross* (see Figures 3 and 10); we denote this square by $(2^n - 1)$-XY-square where XY is the label of the central double cross.

The set τ_K is such that, given a $(2^n - 1)$-square tiled correctly, the tiles immediately outside the square are the ones that allow the correct tiling to be extended to a $(2^{n+1} - 1)$-square.

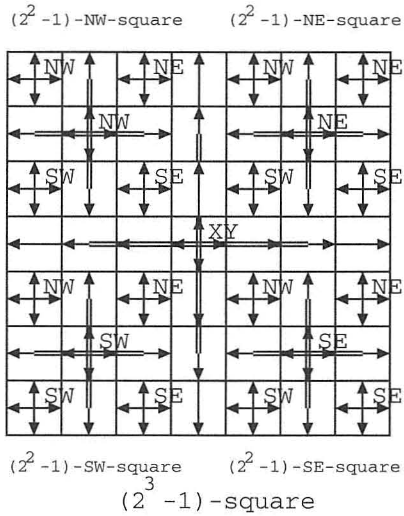

Figure 10: Correct tiling. Arms are drawn without secondary arrows and without labels.

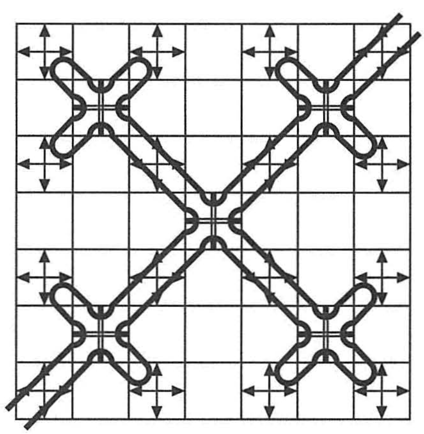

7-NW-square

Figure 11: A correct path. The path is drawn without direction.

In a correct tiling every 2×2 block of tiles contains a single cross. Thus, if we consider each *final* tile consisting of one of all possible correct 2×2 blocks of elementary tiles, we obtain a path visiting the entire plane (Lemma 3.1); the following technical lemma proves this result.

Lemma A.1. [10] *Let t be a single cross on the plane. Consider the path that goes via t. Suppose that there are no tiling errors in any of the 4^n tiles that precede and the 4^n tiles that follow t on this path. Then t belongs to a XY-$(2^n - 1)$-square (XY can be as usual, NE, NW, SE, or SW) whose single crosses are all visited by the path.*

Note that Lemma A.1 implies that if there are no tiling errors in the plane, there is an unique, infinite covering path. A formal proof of this consequence can be found in [10].

References

[1] Aladyev, V., "Computability in Homogeneous Structures," *Izvestija Akademii Nauk Estonskoi SSR, serija Fiziko-Matematicheskaja*, **21** (1972) 80–83.

[2] Amoroso S., and Patt, Y. N., "Decision Procedures for Surjectivity and Injectivity of Parallel Maps for Tesselation Structures," *Journal of Computer & System Sciences* **6** (1972) 448–464.

[3] Bernstein, E. and Vazirani, U., "Quantum Complexity Theory," *Proceedings of the 25^{th} ACM Symposium on Theory of Computing*, (1993) 11–20.

[4] Burks, A., *Essays on Cellular Automata* (University of Illinois Press, Urbana, 1970).

[5] Bennett, C.H., Gacs, P., Li, M., Vitanyi, P., and Zurek, W., "Thermodynamics of Computation and Information Distance," *Proceedings of the 25^{th} ACM Symposium on Theory of Computing*, (1993) 21–30.

[6] Clementi, A. and Pierini, P., "Computational Complexity of the Finite Cellular Automata Reversibility Problem," in *Proceedings of the 4^{th} Italian Conference on Theoretical Computer Science*, L'Aquila, Italy, (1992) 165–178.

[7] Clementi, A., Pierini, P., and Impagliazzo, R., "On the Average-case Complexity of the Reversibility Problem for Finite Cellular Automata," *Proceedings of IEEE—PhysComp '94*, Dallas, Texas, (1994) 151–155.

[8] Jacopini, G. and Sontacchi, G., "Reversible Parallel Computation: An Evolving Space Model," *Theoretical Computer Science* **73** (1993).

[9] Kari, J., "Reversibility of 2D Cellular Automata is Undecidable," *Physica D* **45** (1990) 379–385.

[10] Kari, J., "Reversibility and Surjectivity Problems of Cellular Automata," *Journal of Computer & System Sciences* **48** (1994) 149–182.

[11] Margolus, N., "Cellular Automata Machines: A New Kind of Computer," to appear in the *Proceedings of the* 6^{th} *International Conference on Computational Physics*, Bejing, China, 1993.

[12] Morita, K. and Harao, M., "Computation Universality of One-Dimensional Reversible CA," *Transactions of IEICE* **72** (1989) 758–762.

[13] Richardson, D., "Tesselation with Local Transformations," *Journal of Computer & System Sciences* **6** (1972) 373–388.

[14] Robinson, R.M., "Undecidability and Nonperiodicity for Tilings of the Plane," *Inventionae Mathematicae* **12** 177–209.

[15] Toffoli, T., "Computation and Construction Universality of Reversible Cellular Automata," *Journal of Computer & System Sciences* **15** (1977) 213–231

[16] Toffoli, T., "A High Performance Cellular Automata Machine," *Physica D*, **10** (1984) 195–204.

[17] Toffoli, T. and Margolus N., *Cellular Automata Machines – A New Environment for Modeling* (MIT Press, Cambridge, MA, 1987).

[18] Toffoli, T. and Margolus N., "Invertible Cellular Automata: A Review," *Physica D* **666** (1990) 229–253.

[19] Toffoli, T. and Margolus N., "Programmable Matter: Concepts and Realization," *Physica D* **47** (1991) 263–272.