

Self-stabilisation of Cellular Automata on Tilings

Nazim Fatès

*Université de Lorraine
CNRS, Inria, LORIA
F-54000 Nancy, France
nazim.fates@loria.fr*

Irène Marcovici

*Université de Lorraine
CNRS, Inria, IECL
F-54000 Nancy, France
irene.marcovici@univ-lorraine.fr*

Siamak Taati^{*†}

*Department of Mathematics
American University of Beirut
Beirut, Lebanon
siamak.taati@gmail.com*

Abstract. Given a finite set of local constraints, we seek a cellular automaton (i.e., a local and uniform algorithm) that self-stabilises on the configurations that satisfy these constraints. More precisely, starting from a finite perturbation of a valid configuration, the cellular automaton must eventually fall back into the space of valid configurations where it remains still. We allow the cellular automaton to use extra symbols, but in that case, the extra symbols can also appear in the initial finite perturbation. For several classes of local constraints (e.g., k -colourings with $k \neq 3$, and North-East deterministic constraints), we provide efficient self-stabilising cellular automata with or without additional symbols that wash out finite perturbations in linear or quadratic time, but also show that there are examples of local constraints for which the self-stabilisation problem is inherently hard. We note that the optimal self-stabilisation speed is the same for all local constraints that are isomorphic to one another. We also consider probabilistic cellular automata rules and show that in some cases, the use of randomness simplifies the problem. In the deterministic case, we show that if finite perturbations are corrected in linear time, then the cellular automaton self-stabilises even starting from a random perturbation of a valid configuration, that is, when errors in the initial configuration occur independently with a sufficiently low density.

Keywords: tilings, shifts of finite type, cellular automata, self-stabilisation, noise, fault-tolerance, reliable computing, symbolic dynamics.

^{*}The work of ST was partially supported by NWO grant 612.001.409.

[†]Address for correspondence: Department of Mathematics, American University of Beirut, Beirut, Lebanon.

Introduction

While all living organisms possess some ability to stabilise or repair themselves when subjected to perturbations or attacks, artificial systems rarely have such an ability. In particular, in systems designed in engineering and computer science, a small local perturbation (e.g., due to noise or tampering by an adversary) can often propagate throughout the system leading to a total or partial devastation of the behaviour of the system. The inevitability of such perturbations has led to the study of systems which, in addition to their normal functionality, have the *self-stabilisation* property. A self-stabilising system has the capacity to re-enter a set of “legal” or “desirable” states once the system has been taken out of its normal behaviour by an external perturbation.

The concept of self-stabilisation in computational processes was first introduced in 1970s by Dijkstra, who presented examples of networks of finite-state automata with a non-trivial self-stabilisation property [1]. Since then, self-stabilisation has been widely studied in the context of distributed computing (see e.g. [2, 3]). In the current paper, we explore the question of self-stabilisation in the context of cellular automata.

In a cellular automaton (CA), the components of the system, the *cells*, are arranged regularly on an infinite d -dimensional lattice. The cells are identical finite-state automata that interact locally and change their states synchronously. The overall state of the cells is referred to as a *configuration* of the CA. We will assume that the set of legal configurations of the system is specified with a finite number of local constraints, which must be satisfied at every position. As a prototypical example, one may consider the *colouring* constraints: each cell can have any of a finite number of colours, and the legal configurations are those in which every two adjacent cells have different colours. More generally, we think of the legal states as *tilings* of the lattice with a finite number of tile types (identified with the states of the cells) satisfying local matching constraints. In the language of symbolic dynamics, the set of legal configurations is simply a shift space of finite type (SFT). We will clarify the terminology further in the following section.

We require our CA to have the following form of self-stabilisation:

- (1) Starting from a configuration that deviates from a legal configuration only on a finite region, the CA must evolve back to a legal configuration in a finite number of steps.
- (2) Starting from a legal configuration, the CA must remain unchanged.

We note that in our definition, the CA has no functionality other than to keep the constraints satisfied. Depending on the local constraints, even this simplified notion of self-stabilisation can be quite challenging to achieve. The difficulty is that the cells are indistinguishable and the information available to each cell is limited to the state of its close neighbours. Since the cells are not aware of their own absolute position or the position and extent of the error region, it is thus for example not possible to correct the error region by starting from the upper-left corner and then proceeding sequentially.

Self-stabilisation can be understood as a weak form of fault-tolerance, in which the perturbations occur only at the beginning. Stronger forms of fault-tolerance have been studied in the setting of cellular automata, although aside from a few strong proof-of-concept constructions, the field remains wide open. Around the same time as Dijkstra, Toom found a class of CA which self-stabilise even in

presence of sufficiently weak *temporal* noise [4, 5] (see Examples 3.1 and 6.2 below). Gács and Reif exploited Toom’s simplest example (i.e., the NEC-majority rule) to construct a three-dimensional CA which, in presence of noise, can perform universal computation reliably [6]. Subsequently, Gács was able to construct a sophisticated one-dimensional CA capable of reliable universal computation [7, 8]. In our setting, Toom’s NEC-majority CA solves the self-stabilisation problem for the constraint that adjacent cells must have the same colour.

Various other problems studied in the setting of cellular automata can be related to self-stabilisation. For instance, the *density classification* problem [9, 10] can be formulated as a problem of self-stabilisation where the system needs to return to a homogeneous configuration (all-zero or all-one) with the additional requirement that the colour which appears less frequently in the initial configuration is the one which has to be wiped off. It is known that Toom’s NEC-majority CA solves this problem, at least when the system starts from a biased Bernoulli random configuration [9]. Another example is the *global synchronisation* problem, which can again be understood as a self-stabilisation problem with the homogeneous configurations as the legal states, with the requirement that, in its legal state, the system must oscillate rather than remain unchanged [11, 12].

This article has grown out of a conference paper in which some of our results were presented [13]. The scope of the current paper is however more general and contains various new results. The structure of the paper is as follows:

- In Section 1, we introduce the terminology and notation.
- In Section 2, we present a general construction for self-stabilising one-dimensional SFTs.
- As is the case for many other problems regarding cellular automata and tilings, the self-stabilisation problem in two and higher dimensions is significantly more complex than in one dimension. Section 3 is dedicated to the two-dimensional case, where we present several constructions of self-stabilising CA depending on the type of the constraints. Here, the example of k -colourings serves as a running example, as it has different levels of difficulty depending on the parameter k . While the cases of $k = 2$ and $k \geq 5$ admit relatively simple solutions that self-stabilise in linear time, our solution for $k = 4$ self-stabilises in quadratic time, and we could not find any solution whatsoever for the case $k = 3$. We also provide a linear-time solution for the case of deterministic SFTs. Deterministic SFTs encompass a relatively rich family of constraints, including some which admit only non-periodic legal configurations.
- In Section 4, we investigate the self-stabilisation of probabilistic CA, and show that, in some cases, access to randomness simplifies the self-stabilisation problem. For instance, our probabilistic solution for k -colourings with $k \geq 5$ works in logarithmic time rather than linear time, and furthermore, has the same symmetries as the colouring constraints.
- After having explored the “algorithmic” aspects of self-stabilisation, we turn to the question of “complexity” in Section 5. We show that for some choices of the legal constraints, the self-stabilisation problem is inherently hard (i.e., requires super-polynomial stabilisation time, unless $\mathbf{P} = \mathbf{NP}$). We also show that “isomorphic” constraints (i.e., isomorphic SFTs) admit solutions with roughly the same stabilisation times.

- Section 6 is about a different notion of self-stabilisation in which the initial perturbations are random rather than finite. We show that if a (deterministic) CA self-stabilises from finite perturbations in linear time, then it also self-stabilises from sufficiently weak Bernoulli random perturbations. The more interesting question of self-stabilisation in presence of temporal noise (as in the case of Toom's CA) is left open.
- The article ends with some remarks and open questions in Section 7.

1. Terminology and notations

1.1. Tilings

Configurations and patterns. Let Σ be an *alphabet*, that is, a finite set of *symbols*, and let $d \geq 1$. An assignment $x: \mathbb{Z}^d \rightarrow \Sigma$ is called a *configuration* of the *lattice* \mathbb{Z}^d . The symbol x_i is the *state* (or *colour*) of *cell* i . A configuration is said to be *homogeneous* if all the cells are in the same state. Given $c \in \Sigma$, we denote by \underline{c} the homogeneous configuration in which all cells have symbol c .

The restriction of the configuration $x \in \Sigma^{\mathbb{Z}^d}$ to a set $A \subseteq \mathbb{Z}^d$ is denoted by x_A . A *pattern* is an assignment $p: A \rightarrow \Sigma$ with finite *shape* $A \subseteq \mathbb{Z}^d$, i.e., a partial configuration with a finite domain. We denote by $\Sigma^\#$ the set of all patterns.

A sequence $x^{(1)}, x^{(2)}, \dots$ of configurations is said to *converge* to another configuration x if the state of each cell in $x^{(n)}$ eventually fixates at the value of the same cell in x , that is, if for every $i \in \mathbb{Z}^d$ there exists an n_i such that $x_i^{(n)} = x_i$ for all $n \geq n_i$. This is the notion of convergence in the product topology on $\Sigma^{\mathbb{Z}^d}$. The space $\Sigma^{\mathbb{Z}^d}$ with the product topology is compact and metrizable.

The *shift* by $k \in \mathbb{Z}^d$ is the map $\sigma^k: \Sigma^{\mathbb{Z}^d} \rightarrow \Sigma^{\mathbb{Z}^d}$ defined by $\forall i \in \mathbb{Z}^d, \sigma^k(x)_i \triangleq x_{k+i}$. Every shift is continuous in the product topology.

Shift spaces of finite type. A *shift space of finite type (SFT)* is a set $X \subseteq \Sigma^{\mathbb{Z}^d}$ of configurations identified by a finite number of local constraints. More specifically, let $\mathcal{F} \subseteq \Sigma^\#$ be a finite set of finite patterns, which we refer to as the *forbidden patterns*. The set of configurations $x \in \Sigma^{\mathbb{Z}^d}$ that avoid the patterns in \mathcal{F} (i.e. $\sigma^k(x)_A \notin \mathcal{F}$ for every pattern $p: A \rightarrow \Sigma$ of \mathcal{F} and every $k \in \mathbb{Z}^d$) is called an SFT and is denoted by $X_{\mathcal{F}}$. Every SFT is closed (hence compact) in the product topology, and is invariant under every shift.

Observe that the choice of the defining forbidden sets \mathcal{F} is not unique, and in our discussion, we occasionally need to consider distinct collections defining the same SFT. The smallest integer m for which X can be identified by a collection of forbidden patterns with shape $S_m \triangleq \{0, 1, \dots, m-1\}^d$ is referred to as the *interaction range* of X .

A pattern (or partial configuration) $p: A \rightarrow \Sigma$ is said to be *globally admissible* in X if $p = x_A$ for some $x \in X$, and is said to be *locally admissible* with respect to \mathcal{F} if it has no occurrence of the patterns from \mathcal{F} , that is, if $\sigma^k(p)_B \notin \mathcal{F}$ for all $k \in \mathbb{Z}^d$ and finite $k+B \subseteq A$. Note that in general, a locally admissible pattern does not need to be globally admissible.

Tiling spaces (or nearest-neighbour SFTs). A *tiling space* (or *nearest-neighbour SFT*) is an SFT defined by a collection of *nearest-neighbour* forbidden patterns, that is to say, patterns whose shapes

consist of exactly two adjacent cells. Formally, let e_1, e_2, \dots, e_d denote the standard basis vectors in \mathbb{R}^d . A nonempty set $X \subseteq \Sigma^{\mathbb{Z}^d}$ is a (d -dimensional) tiling space if there exist functions $v_1, v_2, \dots, v_d: \Sigma^2 \rightarrow \{0, 1\}$ such that

$$X = \left\{ x \in \Sigma^{\mathbb{Z}^d} : \forall c \in \mathbb{Z}^d, \forall i \in \{1, 2, \dots, d\}, v_i(x_c, x_{c+e_i}) = 1 \right\}.$$

Example 1.1. (Homogeneous space)

We denote by $\mathcal{H}_2 = \{\underline{0}, \underline{1}\}$ the d -dimensional SFT containing only the two homogeneous configurations $\underline{0}, \underline{1} \in \{0, 1\}^{\mathbb{Z}^d}$. This can be seen as the tiling space defined by the functions $v_1 \triangleq \dots \triangleq v_d \triangleq v$ where $v(a, b) \triangleq 1$ if $a = b$, and 0 if $a \neq b$. ○

Example 1.2. (k -colourings)

A k -colouring of the lattice \mathbb{Z}^d is an assignment of colours from $\Sigma \triangleq \{0, 1, \dots, k-1\}$ to each position in such a way that the adjacent positions have different colours. The set of all k -colourings \mathcal{C}_k is a tiling space identified by the functions $v_1 \triangleq \dots \triangleq v_d \triangleq v$ where $v(a, b) \triangleq 1$ if $a \neq b$, and 0 if $a = b$. ○

Example 1.3. (Hard-core)

The d -dimensional *hardcore* tiling space on the set of symbols $\Sigma = \{0, 1\}$ is defined by the function $v_1 \triangleq \dots \triangleq v_d \triangleq v$ where $v(a, b) \triangleq 1$ if and only if $(a, b) \neq (1, 1)$. ○

Example 1.4. (Wang tiles)

A general family of tiling spaces are those defined by Wang tiles. A *Wang tile* is a unit square with coloured edges (see Figure 14). The colours indicate the matching rules for tiling: two tiles placed next to each other must have the same colour on their touching edges. A finite collection Θ of Wang tiles identifies a two-dimensional tiling space $X \subseteq \Theta^{\mathbb{Z}^2}$, consisting of all *valid tilings* (i.e., configurations that respect the matching rule). In other words, X is defined by the functions $v_1, v_2: \Theta^2 \rightarrow \{0, 1\}$ where $v_1(a, b) \triangleq 1$ if and only if the right edge of a has the same colour as the left edge of b , and $v_2(a, b) \triangleq 1$ if and only if the top edge of a has the same colour as the bottom edge of b . ○

1.2. Perturbations of configurations

Finite perturbations. For two configurations $x, y \in \Sigma^{\mathbb{Z}^d}$, we denote by $\Delta(x, y) \triangleq \{i \in \mathbb{Z}^d : x_i \neq y_i\}$ the set of cells at which x and y disagree. A *finite perturbation* of a configuration $x \in \Sigma^{\mathbb{Z}^d}$ in $\Sigma^{\mathbb{Z}^d}$ is a configuration $\tilde{x} \in \Sigma^{\mathbb{Z}^d}$ such that $\Delta(x, \tilde{x})$ is finite.

The *diameter* of a finite set $A \subseteq \mathbb{Z}^d$, denoted by $\text{diam}(A)$, is the smallest $m \in \mathbb{N}$ such that A fits in a hypercube of size m , that is, $A \subseteq i + [0, m]^d$ for some $i \in \mathbb{Z}^d$. For two configurations $x, y \in \Sigma^{\mathbb{Z}^d}$, the diameter of $\Delta(x, y)$ is denoted by $\delta(x, y)$.

Given an SFT $X \subseteq \Sigma^{\mathbb{Z}^d}$, we denote by $\tilde{X}(\Sigma)$ the set of finite perturbations of the elements of X in $\Sigma^{\mathbb{Z}^d}$, that is $\tilde{X}(\Sigma) \triangleq \{y \in \Sigma^{\mathbb{Z}^d} : \exists x \in X, \delta(x, y) < \infty\}$.

Let us stress that the set $\tilde{X}(\Sigma)$ depends on the choice of the alphabet Σ . A larger alphabet $\Sigma' \supseteq \Sigma$ would lead to a larger set $\tilde{X}(\Sigma')$ of finite perturbations. When the choice of the alphabet is clear from the context, we will simply use the notation \tilde{X} as a shortcut.

Case of tiling spaces. When considering an element $x \in \tilde{X}$, we will often examine the set of cells where the constraints of the SFT are not respected. We will say that such cells are *defective*, or *have defects*. In the specific case of tiling spaces (or nearest-neighbour SFT), we introduce different notions of defects.

For a configuration $x \in \Sigma^{\mathbb{Z}^d}$, a cell $c \in \mathbb{Z}^d$ is said to have a *defect in direction* e_i (with respect to v_i) if $v_i(x_c, x_{c+e_i}) = 0$. It has a *defect in direction* $-e_i$ if $v_i(x_{c-e_i}, x_c) = 0$. In the two-dimensional case, we will also use the terminology *E-defect*, *W-defect*, *N-defect*, *S-defect* instead of respectively defect in direction $e_1, -e_1, e_2 - e_2$. The set of cells having a defect is then defined by

$$\mathcal{D}(x) \triangleq \{c \in \mathbb{Z}^d : \exists e \in \{\pm e_1, \dots, \pm e_d\}, c \text{ has a defect in direction } e\}.$$

A cell $c \in \mathbb{Z}^d$ is said to be *defect-free* if it does not belong to $\mathcal{D}(x)$, meaning that it obeys the local constraints in the $2d$ directions.

Note that in some cases, even if a configuration contains only very few defects, it is necessary to modify a much larger set of cells in order to reach a valid configuration. More precisely, for some tiling spaces X , neither the cardinality nor the diameter of $\mathcal{D}(\tilde{x})$ gives much information about $\delta(\tilde{x}, X)$.

Example 1.5. (3-colourings)

Let \mathcal{C}_3 be the set of two-dimensional 3-colourings. For any integer $n \geq 1$, there exists a configuration $\tilde{x} \in \tilde{\mathcal{C}}_3$ such that $\mathcal{D}(x)$ contains only two adjacent cells, and $\delta(\tilde{x}, \mathcal{C}_3) \geq n$. In other words, in order




Figure 1. A configuration in $\tilde{\mathcal{C}}_3$ that has only two defects but is nonetheless far from every element of \mathcal{C}_3 . The arrows pointing South and East are emphasized.

to correct a single defect, one may have to modify the state of cells that are arbitrarily far. The construction of such a configuration is illustrated in Figure 1, using the connection between the set of 3-colourings and the six-vertex model. This connection and a further discussion of 3-colourings will be presented in Section 7.1. \circ

A symbol $\alpha \in \Sigma$ is called a *safe* symbol for an SFT $X \subseteq \Sigma^{\mathbb{Z}^d}$ if for every $x \in X$ and each $k \in \mathbb{Z}^d$, the configuration \tilde{x} obtained from x by replacing x_k with α is again in X . In case of a tiling space, this means that $v_i(\alpha, \sigma) = v_i(\sigma, \alpha) = 1$ for all $\sigma \in \Sigma$ and $i \in \{1, 2, \dots, d\}$. For example, for the hardcore tiling space of Example 1.3, the symbol 0 is a safe symbol. If a tiling space has a safe symbol, then the phenomenon discussed above for 3-colourings cannot occur, since one can always update the cells having a defect with a safe symbol in order to recover a valid configuration.

1.3. Self-stabilising cellular automata

Cellular automata. A *cellular automaton* (CA) is a dynamical system on $\Sigma^{\mathbb{Z}^d}$ obtained by repeated parallel updating of the symbols on the lattice using a local rule. More specifically, given a finite set $\mathcal{N} \subseteq \mathbb{Z}^d$ and a map $f: \Sigma^{\mathcal{N}} \rightarrow \Sigma$, we can define a mapping $F: \Sigma^{\mathbb{Z}^d} \rightarrow \Sigma^{\mathbb{Z}^d}$ by $F(x)_i \triangleq f(\sigma^i(x)_{\mathcal{N}})$. This is the *global map* of the CA defined with *local rule* f and *neighbourhood* \mathcal{N} .

The neighbourhood can always be chosen to have the form $\mathcal{N} = \{-r, \dots, r\}^d$, in which case we say that the CA has *neighbourhood radius* r . The neighbourhood $\mathcal{M} \triangleq \{-1, 0, 1\}^d$ is called the *Moore* neighbourhood, and the neighbourhood $\mathcal{N} = \{0, \pm e_1, \dots, \pm e_d\}$ is referred to as the *von Neumann* neighbourhood.

For a neighbourhood \mathcal{N} and a set $A \subseteq \mathbb{Z}^d$, we write $\mathcal{N}(A)$ and $A + \mathcal{N}$ interchangeably. For an integer $t \geq 1$, we also introduce the set $\mathcal{N}^t = \underbrace{\mathcal{N} + \dots + \mathcal{N}}_{t \text{ times}}$. If the CA F has neighbourhood \mathcal{N} , then F^t is a CA of neighbourhood \mathcal{N}^t .

Self-stabilisation. We say that a CA $F: \Sigma^{\mathbb{Z}^d} \rightarrow \Sigma^{\mathbb{Z}^d}$ *stabilises* an SFT $X \subseteq \Sigma^{\mathbb{Z}^d}$ from finite perturbations if

- i) (*consistency*) the configurations of X are fixed points, that is, $F(x) = x$ for every $x \in X$,
- ii) (*attraction*) finite perturbations of the elements of X evolve to X in finitely many steps, that is, for every $\tilde{x} \in \tilde{X}(\Sigma)$, there exists a time $t \in \mathbb{N}$ such that $F^t(\tilde{x}) \in X$.

The first such t is called the *stabilisation time* (or the *recovery time*) starting from \tilde{x} . We say that F stabilises X from finite perturbations *in time* $\tau(n)$ if for each $n \in \mathbb{N}$, the largest stabilisation time among all the configurations \tilde{x} with $\delta(\tilde{x}, X) = n$ is $\tau(n)$. We remark that the above notion of stabilisation makes sense even if X is an arbitrary set of configurations, and not only when it is an SFT. We restrict ourselves to the scenario in which X is an SFT, because local and uniform constraints appear more natural in the present context.

Note that the above definition does not rule out the possibility that Σ has extra symbols in addition to those appearing in the elements of X . In other words, it is possible that $X \subseteq \Gamma^{\mathbb{Z}^d}$ for some $\Gamma \subsetneq \Sigma$. Let us emphasize that, according to the above definition, in order for F to stabilise X , it is necessary that all finite perturbations of X in $\Sigma^{\mathbb{Z}^d}$ (not just those in $\Gamma^{\mathbb{Z}^d}$) evolve to X . Even

without this requirement, the problem of finding a cellular automaton that corrects finite perturbation on a given SFT remains non-trivial. We will come back to this in Section 7.1 and in Problem 7.6 of Section 7.3.

Following an observation made earlier, note that if a tiling space $X \subseteq \Sigma^{\mathbb{Z}^d}$ has a safe symbol, then one can easily construct a CA that stabilises X from finite perturbations, without extra symbols. Indeed, suppose that α is a safe symbol of X . Then, the map $F: \Sigma^{\mathbb{Z}^d} \rightarrow \Sigma^{\mathbb{Z}^d}$ defined by

$$\forall c \in \mathbb{Z}^d, \quad F(x)_c \triangleq \begin{cases} \alpha & \text{if } c \in \mathcal{D}(x), \\ x_c & \text{otherwise,} \end{cases} \quad (1)$$

is a CA with neighbourhood $\mathcal{N} = \{0, \pm e_i\}$, and it stabilises X in one step.

The aim of the current article is to present self-stabilising CAs for other families of SFTs, for which finding a CA achieving the stabilisation from finite perturbations is a non-trivial problem.

2. One-dimensional case

In this section, we focus on the problem of self-stabilisation for a one-dimensional SFT $X \subseteq \Sigma^{\mathbb{Z}}$. As a specific case, let us present a well-known example of a CA stabilising the one-dimensional tiling space $\mathcal{H}_2 = \{0, 1\}$ in linear time, without extra symbols.

Example 2.1. (GKL)

The Gács-Kurdyumov-Levin (GKL) cellular automaton is the CA $\text{GKL}: \{0, 1\}^{\mathbb{Z}} \rightarrow \{0, 1\}^{\mathbb{Z}}$ with neighbourhood $\mathcal{N} = \{-3, -1, 0, 1, 3\}$ defined for any $x \in \{0, 1\}^{\mathbb{Z}}$ and $k \in \mathbb{Z}$ by

$$\text{GKL}(x)_k \triangleq \begin{cases} \text{maj}(x_k, x_{k+1}, x_{k+3}) & \text{if } x_k = 1, \\ \text{maj}(x_k, x_{k-1}, x_{k-3}) & \text{if } x_k = 0. \end{cases}$$

It is known that this CA is both a 0-eroder and a 1-eroder, which precisely means that it stabilises \mathcal{H}_2 from finite perturbations [14]. Furthermore, the stabilisation occurs in linear time [15]. Another slightly simpler CA having the same property was proposed by Kari and Le Gloanec, under the name of *modified traffic* [16]. \circ

More generally, we will prove that (essentially) every one-dimensional SFT has a CA that stabilises it from finite perturbations in linear time.

Let $X \subseteq \Sigma^{\mathbb{Z}}$ be a one-dimensional SFT. Then, there exists an integer $k \geq 1$ such that X can be described by a set \mathcal{F} of forbidden words of length $k + 1$. In this case, we say that X is a k -step SFT. Indeed, the constraints can be represented by a transition matrix A indexed by $\Sigma^k \times \Sigma^k$ such that for two words $v = v_1 \cdots v_k$ and $w = w_2 \cdots w_{k+1}$ of length k ,

$$A(v_1 \cdots v_k, w_2 \cdots w_{k+1}) \triangleq \begin{cases} 1 & \text{if } v_2 = w_2, \dots, v_k = w_k \text{ and } v_1 v_2 \cdots v_k w_{k+1} \notin \mathcal{F}, \\ 0 & \text{otherwise.} \end{cases}$$

Note that if $k = 1$, then X is a tiling space, that is, a nearest-neighbour SFT.

For an integer $n \geq 1$, we denote by $L_n(X)$ the set of words of length n occurring in X , that is,

$$L_n(X) \triangleq \{w \in \Sigma^n : \exists x \in X, x_1 \cdots x_n = w_1 \cdots w_n\}.$$

The set $L(X) \triangleq \bigcup_{n \in \mathbb{N}} L_n(X)$ is called the *language* of X .

Remark 2.2. In the specific case when X is a *mixing* one-dimensional SFT (meaning that there exists $n_0 \geq 0$ such that for every two words $u, v \in L(X)$ and every $n \geq n_0$, there exists a word $w \in L_n(X)$ such that $uwv \in L(X)$) and has at least one homogeneous configuration, then a result of Maass implies that the self-stabilisation problem is trivial for X [17, Theorem 3.2]. Indeed, this result states that there exists a CA $F: \Sigma^{\mathbb{Z}} \rightarrow \Sigma^{\mathbb{Z}}$ (i.e., without extra symbols) such that

1. F is the identity map on X ,
2. $F(\Sigma^{\mathbb{Z}}) = X$.

In particular, we achieve self-stabilisation in one step. ◇

For $u, v \in L_k(X)$, we write $u \sim v$ if there is a word $w \in \Sigma^*$ such that $uwv \in L(X)$. We say that X is *non-wandering* if \sim is an equivalence relation, meaning that the transition graph defined by the adjacency matrix A consists of strongly connected components, with no connections between them.

Example 2.3. (Wandering vs. non-wandering)

The SFT $\{0^{\mathbb{Z}}, 1^{\mathbb{Z}}\} \subseteq \{0, 1\}^{\mathbb{Z}}$ is non-wandering, but the SFT consisting in $0^{\mathbb{Z}}, 1^{\mathbb{Z}}$ and the translations of $\cdots 000111 \cdots$ is not. ○

Example 2.4. (A non-wandering SFT)

The 1-step SFT on the alphabet $\Sigma = \{0, 1, 2, 3, 4\}$ with the transition matrix and graph illustrated in Figure 2 is non-wandering. We will use this as a running example to illustrate the constructions of this section. ○




Figure 2. An example of a non-wandering SFT. (left) transition matrix (right) transition graph.

Theorem 2.5. (Self-stabilisation in one dimension)

For every non-wandering one-dimensional SFT $X \subseteq \Sigma^{\mathbb{Z}}$, there exists a CA $F: \Sigma'^{\mathbb{Z}} \rightarrow \Sigma'^{\mathbb{Z}}$ with $\Sigma' \supseteq \Sigma$ that stabilises X from finite perturbations in linear time.

Remark 2.6. Ilkka Törmä has communicated with us an alternative construction showing that *every* one-dimensional SFT is stabilised by a CA with *no extra symbol* [18]. His construction uses long markers made of forbidden patterns, allowing to avoid the introduction of extra symbols. ◇

The rest of this section is devoted to the proof of Theorem 2.5. Thus, for the rest of this section, we will assume that X is a k -step non-wandering SFT.

We denote by m the smallest non-negative integer such that for every $u, v \in L_k(X)$ with $u \sim v$, we have $uvw \in L(X)$ for some word w of length at most m .

In the case of the tiling space of Example 2.4, recall that $k = 1$, and one can check that $m = 2$. Indeed, for $u = v = 1$ or $u = v = 2$, one needs a word w of length at least 2 in order to have $uvw \in L(X)$, while in all the other cases, a word of length 0 or 1 is sufficient.

For a configuration $y \in \Sigma^{\mathbb{Z}}$, we denote by

$$\mathcal{D}(y) \triangleq \{i \in \mathbb{Z} : y_{[i-k, i]} \notin L_{k+1}(X)\}$$

the set of cells at which a *defect* occurs. Note that if y is a (finite) perturbation of a configuration $x \in X$, then $\mathcal{D}(y) \subseteq \Delta(x, y) + \{0, 1, \dots, k\}$, (i.e., every defect on y is on the right within distance k from an element of $\Delta(x, y)$), but $\mathcal{D}(y)$ could be much smaller than $\Delta(x, y)$.

Let us consider again Example 2.4. Let $y \in \Sigma^{\mathbb{Z}}$ be the configuration below. The set $\mathcal{D}(y)$ contains only two elements, which correspond to the positions marked by a cross. But in order to recover a valid configuration, we need to change the values of at least seven cells (the ones taking values 3 and 4).

$$\begin{array}{cccccccccccccccccccc} \dots & 0 & 1 & 2 & 0 & 1 & 3 & 4 & 3 & 4 & 3 & 4 & 3 & 2 & 0 & 1 & 2 & 0 & 0 & 0 \dots \\ \times & \times \end{array}$$

A sequential correction process. We first describe a particular *sequential* procedure for correcting finite “islands” of defects on X . This procedure will not be a cellular automaton itself, but will be used in the construction of the cellular automaton. Applied on a configuration y , the procedure involves updating the symbols on y one by one, from left to right, starting from a cell on the left of the leftmost element of $\mathcal{D}(y)$. Each update is performed according to the same local rule, which we call the *patching rule*. If y is in X , then the procedure does not modify any symbol in y . If y is a finite perturbation of a configuration $x \in X$, then the procedure eventually turns y into a configuration $z \in X$, before reaching few cells to the right of the rightmost element of $\mathcal{D}(y)$. The existence of an appropriate patching rule relies on the fact that X is non-wandering.

More specifically, the patching rule will be a function $g: \Sigma^{2k+m} \rightarrow \Sigma$. In order to update the symbol at position i on y , we replace it with $g(y_{[i-k, i+m+k]})$. The sequential updating of y from a cell i to a cell $j \geq i$ proceeds by first updating the symbol at cell i , then updating the symbol at cell $i + 1$, and so forth until we update the symbol at cell j .

The patching rule g is constructed as follows.

- For $u \in L_k(X)$ and $q \in \Sigma^{m+k}$, we let $r \in [0, m]$ be the smallest index (if exists) such that $uwq_{[r+1, m+k]} \in L(X)$ for some $w \in \Sigma^r$. If no such r exists, we choose an arbitrary $a \in \Sigma$ such that $ua \in L(X)$ and set $g(uq) \triangleq a$. If r exists, we choose a corresponding $w \in \Sigma^r$ and set $g(uq) \triangleq w_1$ if $r > 0$ and $g(uq) \triangleq q_1$ if $r = 0$.

- For $u \in \Sigma^k \setminus L_k(X)$ and $q \in \Sigma^{m+k}$, we simply set $g(uq) \triangleq q_1$. (This case is not used during a sequential correction.)

Let us see on some examples how we can construct a patching rule in the context of the tiling space of Example 2.4, where $k = 1$ and $m = 2$. If $u \in \{0, 1, 2\}$ and $q_3 \in \{3, 4\}$, then, there is no index $r \in [0, 2]$ as above. So, we can simply set $g(uq_1q_2q_3) \triangleq (u + 1) \bmod 3$. Consider now for example $g(0432)$. Then $r = 2$, and for $w = 01$, we have $0012 \in L(X)$. So, we can set $g(0432) = 0$, which is in fact the only possible choice here. For $g(0320)$, we have $r = 1$, and again there is no other choice than to set $g(0320) = 1$.

Intuitively, when applied at the leftmost defect $i \in \mathcal{D}(y)$, the patching rule updates y_i based on the (roughly) shortest patch that would remove the defects on y . The following simple lemma formulates the main property of the patching rule.

Lemma 2.7. (Sequential process)

Let y be a finite perturbation of a configuration $x \in X$. Let $[a, b]$ be an interval containing $\mathcal{D}(y)$. Then, the sequential updating of y from a to $b + m$ using the above-constructed patching rule g ends with an element of X .

Proof:

For $i \in [a, b + m]$, let $y^{(i)}$ denote the configuration obtained during the sequential updating procedure right after updating cell i , and set $y^{(a-1)} \triangleq y$ for consistency.

Observe that $y' \triangleq y^{(b-1)}$ has the following property: the finite words in $y'_{(-\infty, b-1]}$ and $y'_{[b, \infty)}$ are all in $L(X)$. Since $y'_{[b-k, b-1]}$ and $y'_{[i, i+k)}$ for $i \geq b$ are in the same transitive component of X (here we are using the non-wandering property), there must be an integer $r \in [0, m]$ such that $y'_{[b-k, b-1]} w y'_{[b+r, b+r+k)} \in L(X)$ for some $w \in \Sigma^r$. Now, the construction of g ensures that $y^{(b+r-1)}$ is in X . □

Let us observe on Figure 3 how the patching rule described for Example 2.4 corrects the configuration y represented above, when applied from left to right. The smallest interval $[a, b]$ containing the defects is marked on the first configuration. In order to update the symbol at position i , we replace it by $g(y_{i-1}y_iy_{i+1}y_{i+2})$. The successive cells that are updated are underlined. The last configuration that is shown belongs to the SFT, so that afterwards, the patching rule does not introduce any new change.

Let us observe that the above procedure may fail if the SFT is wandering. For instance, if we modify the SFT of Example 2.4 so as to allow transition from 1 to 3, then the configuration

$$\begin{array}{cccccccccccccccccccc} \dots & 0 & 1 & 2 & 0 & 1 & 3 & 4 & 3 & 4 & 3 & 4 & 3 & 2 & 0 & 1 & 2 & 0 & 0 & 0 \dots \end{array}$$

×

will contain only one defect. As before, in order to correct this configuration, we have no choice but to modify every 3 and 4. However, in this case, starting from any position, the sequential process described above can only modify the cells to the right of the rightmost 3, hence it will not succeed.




Figure 3. An example of evolution of the sequential correcting process in the context of the SFT of Example 2.4. Time goes downwards.

From the sequential rule to the stabilising CA. We now construct a CA $T: \Sigma^{\mathbb{Z}} \rightarrow \Sigma^{\mathbb{Z}}$ with an extended alphabet $\Sigma' \supseteq \Sigma$ that corrects finite islands of defects on X in linear time. The CA uses the patching rule g to correct defects sequentially from left to right. Since the CA cannot a priori identify the leftmost defect, it instead applies the patching rule simultaneously everywhere that locally looks like the leftmost defect. Therefore, if there are several far apart defect regions, there will be a correction trail initiated from the left of each of them. These correction trails need not be consistent with one another. We use suitable signals to make sure that the leftmost trail of correction is “dominant”, and the other ones do not continue forever to the right, corrupting the original configuration. To this end, each correction trail leaves a trace (using extra symbols) that is slowly faded away on its own. If a correction trail coming from the left meets a trace in front of it, it sends a fast signal ahead (again using extra symbols) to stop the correction trail that has left that trace. So, the leftmost trail eventually stops all the trails in front of it and goes on to correct the entire island of defects. We must of course make sure that this scenario works even if the defects in the initial configuration involve symbols from the extended alphabet.

The extended alphabet will be $\Sigma' \triangleq \Sigma \times \{\circ, \bullet, \otimes\}$, in which $\{\circ\} \times \Sigma$ is identified with Σ . The symbol \bullet represents the *trace* and \otimes signifies the *stop* signal. For convenience, we identify the configurations with alphabet Σ' with pairs (y, α) where $y \in \Sigma^{\mathbb{Z}}$ and $\alpha \in \{\circ, \bullet, \otimes\}^{\mathbb{Z}}$. The configuration (y, α) has a *defect* at cell i if either y has a defect at i , or i contains a signal symbol \bullet or \otimes . Extending the notation $\mathcal{D}(y)$, we denote the set of defects on (y, α) by

$$\mathcal{D}(y, \alpha) \triangleq \{i \in \mathbb{Z} : i \in \mathcal{D}(y) \text{ or } \alpha_i \neq \circ\}.$$

We also define the set

$$\mathcal{D}_0(y, \alpha) \triangleq \{i \in \mathcal{D}(y, \alpha) : y_{[i-k, i-1]} \in L_k(X)\}$$

whose elements we interpret as the cells that are the leftmost elements of a defect island. Observe that the patching rule g can only affect the cells in $\mathcal{D}_0(y, \alpha)$.

The CA T will be constructed as a composition $T_2(T_g T_1^2 T_0)^2$ of four CA maps $T_g, T_0, T_1, T_2 : \Sigma'^{\mathbb{Z}} \rightarrow \Sigma'^{\mathbb{Z}}$. This composition will ensure that the fading of the traces is half as slow as the correction speed, while the stop signals propagate twice as fast as the correction speed.

Patching. The map T_g is defined by $T_g(y, \alpha) \triangleq (y', \alpha')$, where

$$(y'_i, \alpha'_i) \triangleq \begin{cases} (g(y_{[i-k, i+m+k]}), \bullet) & \text{if } \alpha_{i-1} \neq \otimes \text{ and } i \in \mathcal{D}_0(y, \alpha), \\ (y_i, \alpha_i) & \text{otherwise.} \end{cases}$$

It simply applies the patching rule on the elements of $\mathcal{D}_0(y, \alpha)$ (if no stop signal on the left) and leaves a trace behind. It also erases any stop symbol which sees no stop symbol on its left neighbour, replacing it with a trace symbol.

Generation of stop signals. The map T_0 is responsible for generating stop signals, and is defined by $T_0(y, \alpha) \triangleq (y, \alpha')$ (i.e., no change on y), where

$$\alpha'_i \triangleq \begin{cases} \otimes & \text{if } \alpha_i = \bullet \text{ and } i \in \mathcal{D}(y), \\ \alpha_i & \text{otherwise.} \end{cases}$$

Propagation of stop signals. The propagation of the stop signals is governed by the map T_1 , which is defined by $T_1(y, \alpha) \triangleq (y, \alpha')$ (i.e., no change on y), where

$$\alpha'_i \triangleq \begin{cases} \otimes & \text{if } \alpha_{i-1} = \otimes \text{ and } \alpha_i = \bullet, \\ \alpha_i & \text{otherwise.} \end{cases}$$

Fading of the traces. Finally, the map T_2 handles the fading of the traces and is defined by $T_2(y, \alpha) \triangleq (y, \alpha')$ (i.e., no change on y), where

$$\alpha'_i \triangleq \begin{cases} \circ & \text{if } \alpha_{i-1} = \circ \text{ and } \alpha_i = \bullet, \\ \alpha_i & \text{otherwise.} \end{cases}$$

In Figure 4, we illustrate the operation of the cellular automaton in the context of Example 2.4. The defective cells are underlined. Note that in this example $\mathcal{D}_0 = \mathcal{D}$. Namely, since $k = 1$, we always have $y_{[i-k, i-1]} = y_{i-1} \in \Sigma = L_1(X)$. Two consecutive configurations correspond to one application of $T_g T_1^2 T_0$. Every two steps, the map T_2 is also applied, erasing one trace symbol on the left of each correction region. Figure 5 illustrates the evolution of the CA in a longer time span. When the front of a correction trail meets the fading trace of another correction trail in front of it, a stop signal is created. This stop signal travels faster than the correction trails, and hence quickly catches up with all the correction trails in front of it. As a result, every correction trail which is not initiated by the leftmost defect is eventually stopped.

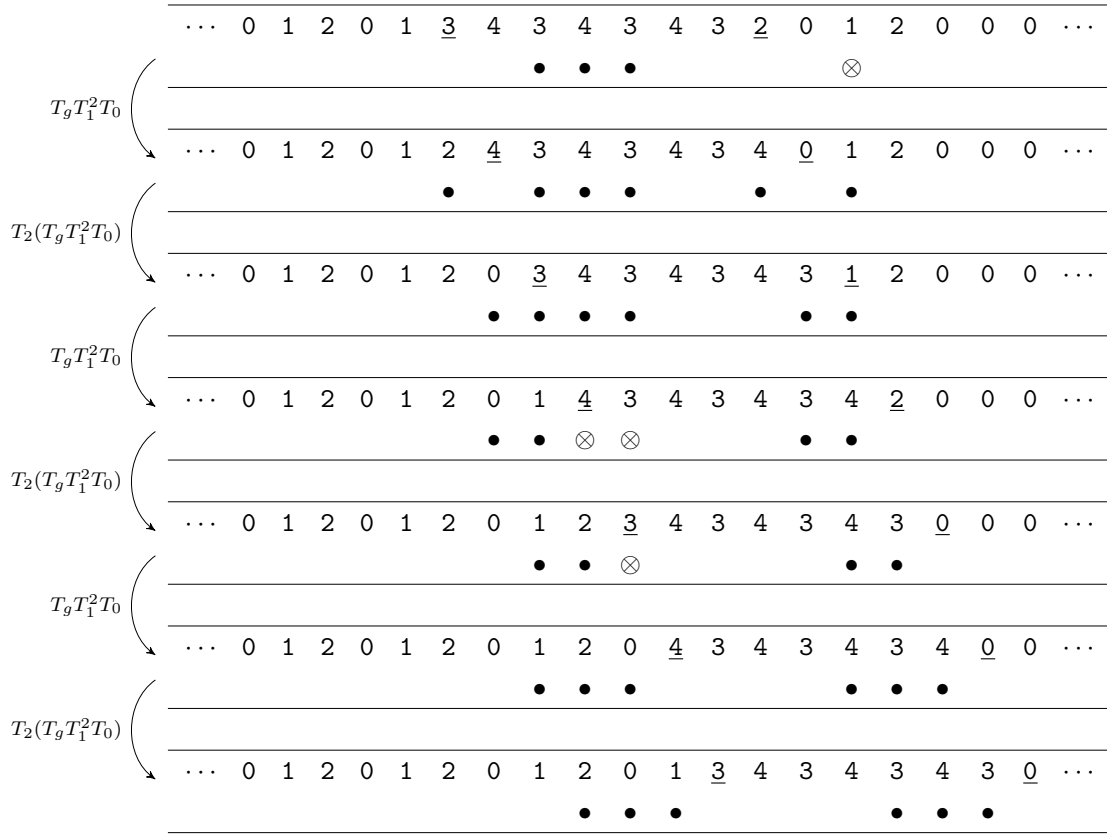


Figure 4. An example of evolution of the cellular automaton in the context of the SFT of Example 2.4. Time goes downwards. Note that in this example, the initial perturbation already contains symbols from the extended alphabet Σ' .

Proof of Theorem 2.5:

We verify that the CA T constructed above corrects finite islands of defects on X in linear time.

Let $(y, \alpha): \mathbb{Z} \rightarrow \Sigma'$ be a finite perturbation of a configuration $x \in X$ (or more explicitly, a finite perturbation of $(x, \circ\mathbb{Z})$). Let us call a cell $i \in \mathcal{D}(y, \alpha)$

- *active* on (y, α) if $i \in \mathcal{D}_0(y, \alpha)$ and $\alpha_{i-1} \neq \otimes$,
- *frozen* if $\alpha_{i-1} = \otimes$, and
- *fading* if $\alpha_i = \bullet$.

The *tail* of an active cell i is the longest (possibly empty) interval $[i - l, i - 1]$ of fading cells. A tail whose leftmost element is frozen is said to be *freezing*. An active cell together with its tail is a *correction trail*. Observe that the leftmost element of $\mathcal{D}(y, \alpha)$ is either active, or fading and non-frozen.




Figure 5. Illustration showing the behaviour of the different signals composing the cellular automaton. Time goes downwards.

The CA works intuitively as follows. Let $[a, b] \subseteq \mathbb{Z}$ be the smallest interval containing $\Delta((x, \circ^{\mathbb{Z}}), (y, \alpha))$. The leftmost trail moves with speed *at least* 2 to the right, while its tail is erased with speed 1 from the left. The rightmost trail moves with speed *at most* 2 to the right, occasionally becoming frozen when the following trail reaches its tail. In summary, at time $t > 0$,

- the leftmost active cell is on the right of $a + 2t$,
- the leftmost fading cell is on the right of $a + t$,
- the rightmost active cell is on the left of $b + 2t$,
- the rightmost non-freezing tail is on the left of $b + t$.

It follows that before time $t_0 \triangleq b - a$, the leftmost trail overpasses the rightmost non-freezing tail, and causes it to freeze. From that moment on, it takes at most $(b - a)/4$ steps before every cell in front of the leftmost trail is inactive. At time $t_1 \triangleq t_0 + \lfloor (b - a)/4 \rfloor < \frac{5}{4}(b - a)$, the CA is then in a configuration (y', α') with the following properties.

- The configuration (y', α') is a finite perturbation of $(x, \circ^{\mathbb{Z}})$ with $\Delta(x, y') \subseteq [a + 2t_1, b + 2t_1]$ and $\Delta((x, \circ^{\mathbb{Z}}), (y', \alpha')) \subseteq [a + t_1, b + 2t_1]$.
- The only active cell on (y', α') is on the right of $a + 2t_1$ and every element of $\mathcal{D}(y', \alpha')$ to the left of it is fading and non-frozen.

From time t_1 onward, the CA essentially applies the patching rule, removing the defects on y in no more than $b - a + m$ steps. At time $t_2 \triangleq t_1 + b - a + m$, the CA is in a configuration (y'', α'') in which

- y'' has no defect,
- α'' contains no stop symbol \otimes and its trace symbols \bullet are contained in region $[a + t_2, a + 2t_2]$.

Eventually, in at most $(a + 2t_2) - (a + t_2) = t_2$ more steps, every trace symbol fades away, and we arrive at a configuration in X before time $t_3 \triangleq t_2 + t_2 \leq \frac{9}{4}(b - a) + m$. \square

3. Two-dimensional case

We now consider the problem of stabilising a two-dimensional SFT. In this section, we will focus on some specific families of tiling spaces, for which different strategies can be implemented in order to achieve self-stabilisation efficiently. The first cases we consider are inspired from the study of self-stabilisation for k -colourings [13]. Namely, Sections 3.1, 3.2, and 3.3 treat and extend respectively the cases of 2-colourings, k -colourings with $k \geq 5$, and 4-colourings. Finally, Section 3.4 is based on the same construction as in Section 2, allowing us to handle the case of *deterministic* SFTs. Most of the constructions in this section can be easily generalized to higher dimensions, but for concreteness, we focus on the two-dimensional case.

As a first example, let us present a well-known CA stabilising the two-dimensional tiling space $\mathcal{H}_2 = \{\underline{0}, \underline{1}\}$ in linear time and without extra symbols.

Example 3.1. (Toom’s North-East-Center majority rule)

Toom’s (deterministic) majority cellular automaton is the CA NEC-Maj: $\{0, 1\}^{\mathbb{Z}^2} \rightarrow \{0, 1\}^{\mathbb{Z}^2}$ with neighbourhood $\mathcal{N} = \{0, e_1, e_2\}$ defined, for any $x \in \{0, 1\}^{\mathbb{Z}^2}$ and $k \in \mathbb{Z}^2$, by

$$\text{NEC-Maj}(x)_k \triangleq \text{maj}(x_k, x_{k+e_1}, x_{k+e_2}),$$

where maj denotes the majority function, outputting the symbol which is in majority among the input symbols.




Figure 6. Illustration of the evolution of Toom’s CA NEC-Maj.

It is known that this two-dimensional CA is both a *0-eroder* and a *1-eroder*, which is equivalent to the CA stabilising $\mathcal{H}_2 = \{\underline{0}, \underline{1}\}$ from finite perturbations. Furthermore, the stabilisation occurs in linear time [5, 19]. \circ

As we shall discuss in Section 5, it seems hopeless to be able to construct, for any two-dimensional SFT, a CA that stabilises it from finite perturbations in linear time, or even in polynomial time. However, in the following of this section, we will present several classes of SFTs for which we are able provide CA that stabilise in linear or quadratic time. We first focus on colourings and similar tiling spaces, then we treat the case of two-dimensional deterministic tiling spaces.

3.1. Finite SFTs

In this section, we treat the case where the SFT contains only a finite number of configurations. Observe that the configurations of such a finite SFT are necessarily spatially periodic.

Example 3.2. (Finite SFTs)

The set $\mathcal{H}_2 = \{\underline{0}, \underline{1}\} \subseteq \{0, 1\}^{\mathbb{Z}^2}$ is a finite tiling space. The set \mathcal{C}_2 of all 2-colourings of \mathbb{Z}^2 is also a finite tiling space, since it contains only two configurations, namely the odd and even checkerboard configurations. \circ

Let us consider an arbitrary finite SFT $X \subseteq \Sigma^{\mathbb{Z}^2}$. Then, for each configuration $x \in X$, there exist integers $n_1, n_2 \geq 1$ (the *horizontal* and *vertical* periods of x) such that $\sigma^{n_1 e_1}(x) = \sigma^{n_2 e_2}(x) = x$. Let N_1 (respectively, N_2) denote the least common multiple of the horizontal (resp., vertical) periods of all the configurations in X . Since X is finite, N_1 and N_2 are finite. Then, for all $x \in X$, we have $\sigma^{N_1 e_1}(x) = \sigma^{N_2 e_2}(x) = x$. We define a CA F on $\Sigma^{\mathbb{Z}^2}$ by

$$F(x)_k \triangleq \text{maj}(x_k, x_{k+N_1 e_1}, x_{k+N_2 e_2}),$$

where the majority function $\text{maj}(a, b, c)$ assigns to three symbols a, b, c the symbol which is most common among a, b, c , with the convention that when a, b, c are distinct, one can choose arbitrarily the value of the function. Observe that F simply consists in applying Toom's majority rule on each sub-lattice generated by $N_1 e_1$ and $N_2 e_2$.

Proposition 3.3. (Self-stabilisation of finite SFTs)

Let $X \subseteq \Sigma^{\mathbb{Z}^2}$ be a finite SFT. Then, the CA $F: \Sigma^{\mathbb{Z}^2} \rightarrow \Sigma^{\mathbb{Z}^2}$ defined above stabilises X from finite perturbations in linear time.

Proof:

It is clear from the definition that $F(x) = x$ for every $x \in X$. For each $n \in \mathbb{N}$, define the triangular set $T_n = \{k \in \mathbb{Z}^2 : k_1 + k_2 \leq n, k_1, k_2 \geq 0\}$. Let $x \in X$ and take y such that $\Delta(x, y)$ is finite. By translating x and y if needed, we can assume without loss of generality that the difference set $\Delta(x, y)$ is included in the triangle T_n for some n . It is then easy to verify that $\Delta(x, F(y)) \subseteq T_{n-1}$. Indeed, for every cell outside T_n , the local rule does not modify the state, whereas for the cells $k \in \mathbb{Z}^2$ which are inside T_n and satisfy $k_1 + k_2 = n$, we have $F(y)_k = x_k$, since $x_k = x_{k+N_1 e_1} = x_{k+N_2 e_2} = y_{k+N_1 e_1} = y_{k+N_2 e_2}$. Iterating F we obtain $\Delta(x, F^t(y)) \subseteq T_{n-t}$ for each $t \geq 0$. Consequently, as

time goes by, the set of disagreements becomes smaller. In particular, for $t = n + 1$, we obtain $\Delta(x, F^{n+1}(y)) \subseteq T_{-1} = \emptyset$, hence $F^{n+1}(y) = x \in X$. This means that the configuration y has been corrected in at most $n + 1$ steps. \square

Remark 3.4. The above result clearly extends to dimensions $d \geq 2$. One can simply apply Toom's majority rule on each (two-dimensional) sub-lattice generated by $N_1 e_1$ and $N_2 e_2$, and in the proof, replace T_n by the triangular prism $T_n = \{k \in \mathbb{Z}^d : k_1 + k_2 \leq n, k_1, k_2 \geq 0\}$. \diamond

3.2. Single-cell fillable tiling spaces

We say that a two-dimensional tiling space is *single-cell fillable* if there exists a map $\psi: \Sigma^4 \rightarrow \Sigma$ such that, for any possible choice $(a, b, c, d) \in \Sigma^4$ of symbols surrounding a cell (see Figure 7), assigning the value $\psi(a, b, c, d)$ to the central cell ensures that it is defect-free. More specifically, the value $\alpha = \psi(a, b, c, d)$ satisfies $v_1(a, \alpha) = v_1(\alpha, c) = 1$ and $v_2(\alpha, b) = v_2(d, \alpha) = 1$. One can observe that a tiling space is single-cell fillable if and only if the functions v_1 and v_2 can be chosen in such a way that every locally admissible pattern is globally admissible [20].




Figure 7. Illustration of the notion of single-cell fillability.

Note that any tiling space having a safe symbol is trivially single-cell fillable. The following is a more interesting family of single-cell fillable tiling spaces.

Example 3.5. (Single-cell fillable)

For $k \geq 5$, the space \mathcal{C}_k of two-dimensional k -colourings is single-cell fillable.

More generally, let $G = (V, E)$ be a finite undirected graph (possibly with self-loops). Let $X_G \subseteq V^{\mathbb{Z}^d}$ denote the set of all graph homomorphisms from \mathbb{Z}^d (with nearest-neighbour edges) to G . In other words, $x \in X_G$ if and only if $x_i x_j \in E$ for every two adjacent cells $i, j \in \mathbb{Z}^d$. The set X_G is clearly a tiling space, defined by the functions $v_1 \triangleq \dots \triangleq v_d \triangleq v$ where $v(a, b) \triangleq 1$ if and only if $ab \in E$. Note that if G is the complete graph on k vertices (without self-loops), then X_G coincides (up to renaming of the vertices) with the space of k -colourings \mathcal{C}_k . It is easy to see that X_G is single-cell fillable if and only if G has the following property: for any subset $A \subseteq V$ with $|A| = 2d$, there exists a vertex $b \in V$ such that $ab \in E$ for every $a \in A$. In particular, X_G is single-cell fillable if the minimum degree of the vertices in G is larger than $(1 - (2d)^{-1})|V|$. See Refs. [21, 22] for more on the homomorphism spaces. \circ

Let X be a single-cell fillable tiling space. We introduce the following terminology for defects. We say that a cell (i, j) has a *NE-defect* if it has a N-defect or an E-defect (or both). For $x \in \Sigma^{\mathbb{Z}^2}$, we denote by $\mathcal{D}_{\text{NE}}(x)$ the set of cells having a NE-defect, that is:

$$\mathcal{D}_{\text{NE}}(x) \triangleq \{(i, j) \in \mathbb{Z}^2 : v_1(x_c, x_{c+e_1}) = 0 \text{ or } v_2(x_c, x_{c+e_2}) = 0\}.$$

Let $\psi: \Sigma^4 \rightarrow \Sigma$ be a function which assigns, to a choice $(a, b, c, d) \in \Sigma^4$ of symbols surrounding a cell, a symbol $\psi(a, b, c, d)$ for the central cell ensuring that it is defect-free. We define a CA F on $\Sigma^{\mathbb{Z}^2}$ by

$$\forall c \in \mathbb{Z}^2, \quad F(x)_c \triangleq \begin{cases} \psi(x_{c-e_1}, x_{c-e_2}, x_{c+e_1}, x_{c+e_2}) & \text{if } c \in \mathcal{D}_{\text{NE}}(x), \\ x_c & \text{otherwise.} \end{cases}$$

Proposition 3.6. (Self-stabilisation of single-cell fillable tiling spaces)

Let $X \subseteq \Sigma^{\mathbb{Z}^2}$ be a single-cell fillable tiling space. Then, the CA $F: \Sigma^{\mathbb{Z}^2} \rightarrow \Sigma^{\mathbb{Z}^2}$ defined above stabilises X from finite perturbations in linear time.

Proof:

It is clear from the definition that $F(x) = x$ for every $x \in X$. Let us now take $x \in \tilde{X}$.

First, observe that if $c \notin \mathcal{D}_{\text{NE}}(x)$, then $c \notin \mathcal{D}_{\text{NE}}(F(x))$, so that the set of NE-defects can only decrease under the action of F . Let us indeed take $c \notin \mathcal{D}_{\text{NE}}(x)$. By definition of F , the value of cell c is not modified when applying F , and if the value of cell $c + e_1$ (resp. $c + e_2$) is modified, that is, if $F(x)_{c+e_1} \neq x_{c+e_1}$ (resp. $F(x)_{c+e_2} \neq x_{c+e_2}$), then the new value of $c + e_1$ (resp. $c + e_2$) is chosen in such a way that $v_1(F(x)_c, F(x)_{c+e_1}) = 1$ (resp. $v_2(F(x)_c, F(x)_{c+e_2}) = 1$), so that $c \notin \mathcal{D}_{\text{NE}}(F(x))$.

Second, if $c \in \mathcal{D}_{\text{NE}}(x)$ is such that $c + e_1, c + e_2 \notin \mathcal{D}_{\text{NE}}(x)$, then $c \notin \mathcal{D}_{\text{NE}}(F(x))$, so that the set of NE-defects is progressively eroded, from the NE to the SW. More formally, we can assume without loss of generality that there exists an integer $n \geq 0$ such that $\mathcal{D}_{\text{NE}}(x) \subseteq T_n$, where $T_n = \{(i, j) \in \mathbb{Z}^2 : i + j \leq n, i, j \geq 0\}$, and one can check that after t steps, we have $\mathcal{D}_{\text{NE}}(F^t(x)) \subseteq T_{n-t}$. Thus, after $n + 1$ steps, we have $\mathcal{D}_{\text{NE}}(F^{n+1}(x)) = \emptyset$, meaning that the configuration is fully corrected: $F^{n+1}(x) \in X$. □

Remark 3.7. The result extends naturally to d -dimensional single-cell fillable tiling spaces, with $d \geq 2$. ◇

3.3. Strongly ℓ -fillable tiling spaces (with $\ell \geq 2$)

We say that a tiling space is *strongly ℓ -fillable* if there exists a map $\psi: \Sigma^{4\ell} \rightarrow \Sigma^{\ell^2}$ such that, for any possible choice $(a_1, \dots, a_{4\ell}) \in \Sigma^{4\ell}$ of symbols surrounding an ℓ -square (that is, an $\ell \times \ell$ block of cells), assigning the values $(b_1, \dots, b_{\ell^2}) = \psi(a_1, \dots, a_{4\ell})$ to the inner cells of the ℓ -square ensures that each cell of the ℓ -square is defect-free (see Figure 8). Note that here, we do not assume any further condition on $(a_1, \dots, a_{4\ell}) \in \Sigma^{4\ell}$. We refer to the article by Alon et al. [23] for a similar but weaker condition of ℓ -fillability.




Figure 8. Illustration of the notion of strong ℓ -fillability.

Example 3.8. (Strongly 2-fillable)

The set \mathcal{C}_4 of two-dimensional 4-colourings is not single-cell fillable. We claim that \mathcal{C}_4 is strongly 2-fillable. To prove this, we need to show that for any possible choice $(a, b, c, d, e, f, g, h) \in \Sigma^8$ of symbols surrounding a 2-square (see Figure 9), there exists a choice $(\alpha, \beta, \gamma, \delta) \in \Sigma^4$ for the cells of the 2-square such that the four cells of the 2-square are defect-free.




Figure 9. Illustration of the notion of strong 2-fillability.

If $\{a, d, e, h\} \subsetneq \Sigma$, then we can choose a colour from $\Sigma \setminus \{a, d, e, h\}$ and assign it to both α and γ . We are then sure that we can find suitable colours for the two remaining cells, since each of these two cells is surrounded by at most three different colours. In the same way, if $\{b, c, f, g\} \subsetneq \Sigma$, we can find a valid pattern.

Let us now assume that $\{a, d, e, h\} = \{b, c, f, g\} = \Sigma$. Without loss of generality, we can assume that $a = 0, h = 1, d = 2, e = 3$. The set of allowed colours for α is then $\{2, 3\}$, and the set of allowed colours for γ is $\{0, 1\}$. If the allowed colours for β and δ are $\{0, 1\}$ and $\{2, 3\}$ respectively, then a valid pattern is given by $(\alpha, \beta, \gamma, \delta) = (2, 0, 1, 3)$. If the allowed colours for β and δ are $\{0, 2\}, \{1, 3\}$ respectively, then a valid pattern is given by $(\alpha, \beta, \gamma, \delta) = (2, 0, 1, 3)$. The other cases are analogous. ○

Example 3.9. (Strongly 2-fillable)

Consider the set Θ of all Wang tiles whose edges are coloured either black or white, except the one with four white edges. It is easy to see that the space of valid tilings with tiles from Θ is strongly 2-fillable but not single-cell fillable. \circ

Example 3.10. (Strongly 2-fillable)

Consider the set of Wang tiles depicted in Figure 10. The decorations symbolise the edge colours, hence there are three possible colours: “horizontal line”, “vertical line” and “none”. Let X denote the space of all valid tilings. It can be verified that X is strongly 2-fillable but not single-cell fillable. \circ




Figure 10. Illustration of Example 3.10. (a) The tile set. (b) Example of a valid tiling.

For $\ell = 1$, strong ℓ -fillability corresponds to single-cell fillability. Therefore, in what follows, we will assume that $\ell \geq 2$. The construction that we present below is slightly different from the one of the earlier article [13], allowing for a simpler proof. The proof from the latter article had some inaccuracies which are now avoided.

Let us design a CA that corrects the finite perturbations of a strongly ℓ -fillable tiling space X . Let $\psi: \Sigma^{4\ell} \rightarrow \Sigma^{\ell^2}$ be a function that maps some $(a_1, \dots, a_{4\ell}) \in \Sigma^{4\ell}$ to an element of Σ^{ℓ^2} , such that the pattern formed by these values is a defect-free pattern. The aim is to apply ψ on a collection of non-overlapping, non-adjacent ℓ -squares containing defects so as to reduce the number of defects. More specifically, the CA must (locally) select some ℓ -squares containing defects in such a way that the following two conditions are satisfied:

1. Every two selected ℓ -squares are at distance at least 1 (they do not overlap or touch each other).
2. If the configuration contains a defect, then at least one ℓ -square containing a defect is selected.

To this end, let us first identify a set of cells that will play the role of the top-right corners of the selected ℓ -squares. Given a configuration $x \in \Sigma^{\mathbb{Z}^2}$, let us denote again the set of cells having a NE-defect by $\mathcal{D}_{NE}(x) = \{c \in \mathbb{Z}^2 : v_1(x_c, x_{c+e_1}) = 0 \text{ or } v_2(x_c, x_{c+e_2}) = 0\}$. For the sake of clarity, we first define a stabilising CA in the case $\ell = 2$, and then treat the general case.

Case $\ell = 2$. We say that a cell $c \in \mathbb{Z}^2$ is a *NE-corner* if $c \in \mathcal{D}_{\text{NE}}(x)$ and

$$\begin{aligned} c - e_1 + e_2, c - 2e_1 + 2e_2 &\notin \mathcal{D}_{\text{NE}}(x), \\ c + 2e_1 - e_2, c + e_1, c + e_2, c - e_1 + 2e_2 &\notin \mathcal{D}_{\text{NE}}(x), \\ c + 2e_1, c + e_1 + e_2, c + 2e_2 &\notin \mathcal{D}_{\text{NE}}(x), \\ c + 2e_1 + e_2, c + e_1 + 2e_2 &\notin \mathcal{D}_{\text{NE}}(x). \end{aligned}$$

See Figure 11 for an illustration of the definition. We denote by $\mathcal{C}_{\text{NE}}(x)$ the set of NE-corners in a configuration $x \in \Sigma^{\mathbb{Z}^2}$.




Figure 11. Illustration of the notion of NE-corner in the case $\ell = 2$. The central cell is a NE-corner if at least one of the red lines (North or East or both) contains a defect and all the green lines are defect-free.

We define a CA F by the following rule: if a cell $c = (i, j) \in \mathbb{Z}^2$ is a NE-corner, then apply ψ to the 2-square whose NE-corner is c , that is, replace the symbols of the cells $(i-1, j-1), (i-1, j), (i, j), (i, j-1)$ by $\psi(a, b, \dots, h)$, where $a = x_{i-2, j-1}, b = x_{i-2, j}, \dots, h = x_{i-1, j-2}$ (see Figure 9). Let us first observe that the CA F given by this rule is well-defined. Indeed, by definition of a NE-corner, one can check that there are no two adjacent NE-corners, vertically or horizontally, or in diagonal. Consequently, at each time step, the 2-squares that are updated do not overlap. In addition, note that the definition of NE-corners ensures that the 2-squares that are updated cannot share a common edge, so that the property 1 above is satisfied.

General case ($\ell \geq 2$). In the general case, we modify the notion of NE-corner as follows. We say that a cell $c \in \mathbb{Z}^2$ is a *NE-corner* if $c \in \mathcal{D}_{\text{NE}}(x)$ and for any $-\ell \leq i, j \leq \ell$,

$$\text{if } 1 \leq i + j \leq \ell - 1 \text{ or } [i + j = 0 \text{ and } j > i], \text{ then } c + ie_1 + je_2 \notin \mathcal{D}_{\text{NE}}(x).$$

See Figure 12 for an illustration in the case $\ell = 3$. We denote again by $\mathcal{C}_{\text{NE}}(x)$ the set of NE-corners in a configuration $x \in \Sigma^{\mathbb{Z}^2}$.

Again, we define a CA F by the following rule: if a cell $c = (i, j) \in \mathbb{Z}^2$ is a NE-corner, then apply ψ to the ℓ -square whose NE-corner is c . By a similar argument as for $\ell = 2$, F is well-defined, and in addition, the ℓ -squares that are updated cannot share a common edge, so that the property 1 above is satisfied.

We are now able to state the result, whose proof now consists in proving that the property 2 above is also satisfied.




Figure 12. Illustration of the notion of NE-corner in the case $\ell = 3$. The central cell is a NE-corner if at least one of the red lines (North or East or both) contains a defect and all the green lines are defect-free.

Proposition 3.11. (Self-stabilisation of strongly fillable tiling spaces)

Let $X \subseteq \Sigma^{\mathbb{Z}^2}$ be a strongly ℓ -fillable tiling space, with $\ell \geq 2$. Then, the CA $F: \Sigma^{\mathbb{Z}^2} \rightarrow \Sigma^{\mathbb{Z}^2}$ defined above stabilises X from finite perturbations in quadratic time.

Proof:

Let $x \in \tilde{X}$. We prove that if $\mathcal{D}(x) \neq \emptyset$, then $\mathcal{C}_{NE}(x) \neq \emptyset$. Let us indeed sweep the configuration x by NW-SE diagonals, from the NE to the SW. Since $\mathcal{D}(x)$ is finite, we can consider the first diagonal which contains a NE-defect, and on this diagonal, we consider the leftmost NE-defect (which is also the uppermost). By definition of a NE-corner, this NE-defect is a NE-corner.

In order to end the proof, it is then sufficient to observe that while $\mathcal{D}(x) \neq \emptyset$, the number of defects decreases strictly when applying F . Indeed, at least all the NE-corners become defect-free, and since the updated ℓ -squares do not share any common edge, no new defect is created. \square




Figure 13. Illustration of the proof of Proposition 3.11 in the case $\ell = 2$. Defects are represented in red, and cells that are NE-corners are shaded. For the other NE-defects, the constraints represented in green on Figure 11 are not all satisfied. As represented, we know that by construction, there exists at least one NE-corner on the first NW-SE diagonal that contains a NE-defect.

3.4. Deterministic SFTs

We now extend the construction proposed in Section 2 for one-dimensional SFTs to a specific class of higher-dimensional SFTs: the *deterministic* SFTs. To begin with, let us define this notion. Again, we focus on the two-dimensional case, but the ideas can also be applied in higher dimensions.

We say that a two-dimensional SFT $X \subseteq \Sigma^{\mathbb{Z}^2}$ is *NE-deterministic* if it possible to describe it by a set \mathcal{F} of forbidden patterns in such a way that all the forbidden patterns of \mathcal{F} have shape $\{0, e_1, e_2\}$, and \mathcal{F} has the additional property that for every $a, b \in \Sigma$, there exists at most one element $c \in \Sigma$ such that the pattern $p: \{0, e_1, e_2\} \rightarrow \Sigma$ defined by $p(e_1) = a$, $p(e_2) = b$ and $p(0) = c$ does not belong to \mathcal{F} .

The constraints are therefore completely specified by the partial function $f: \Sigma \times \Sigma \rightarrow \Sigma$, that maps $(a, b) \in \Sigma^2$ to the unique symbol c that is allowed, if any.

Example 3.12. (Ledrappier SFT)

The SFT

$$X \triangleq \{x \in \{0, 1\}^{\mathbb{Z}^2} : x_k = x_{k+e_1} + x_{k+e_2} \pmod{2} \text{ for all } k \in \mathbb{Z}^2\}$$

known as the *Ledrappier* SFT, or as the *three-dot system*, is NE-deterministic. ○

In case of SFTs identified by Wang tiles, NE-determinism means that for every two tiles $a, b \in \Sigma$, there is at most one tile $c = f(a, b) \in \Sigma$ such that the right edge of c is compatible with the left edge of a and the upper edge of c matches the lower edge of b .

Example 3.13. (Aperiodic deterministic tile sets)

The SFTs in two and higher dimensions can be significantly more complex than the one-dimensional SFTs. For instance, in one dimension, every non-empty SFT contains periodic configurations, whereas in two dimensions, there are various examples of non-empty SFTs that do not contain periodic configurations. Moreover, it is algorithmically undecidable whether the SFT identified by a given finite set of forbidden patterns is non-empty [24, 25].

Figure 14 shows a set of 16 Wang tiles, discovered by R. Ammann, which is *aperiodic*, meaning that it admits valid tilings but no periodic valid tilings [26, Chapter 11]. Remarkably, Ammann's tile




Figure 14. Ammann's set of Wang tiles is NE-deterministic and SW-deterministic.

set is both NE-deterministic and SW-deterministic. The question of whether a given finite set of Wang tiles admits a valid tiling remains undecidable when restricted to NE-deterministic (or even four-way deterministic) tile sets [27, 28]. \circ

As illustrated by the two examples above, the class of NE-deterministic SFTs is thus rich and multi-faceted.

In order to stabilise a NE-deterministic SFT, a natural first idea is to exploit the partial rule f that defines the SFT. Namely, the CA can apply the partial rule f when f is applicable and leave the state of the cell unchanged when f is not applicable. Figure 15 illustrates why this approach fails on the Ledrappier SFT. Indeed, the CA F defined by $F(x)_k \triangleq x_{k+e_1} + x_{k+e_2} \pmod{2}$ cannot stabilise any finite perturbation of the all-0 configuration.

0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
0	0	0	1	1	1	1	0	0
0	0	0	0	1	0	1	0	0
0	0	0	0	0	1	1	0	0
0	0	0	0	0	0	1	0	0
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0

Figure 15. Illustration of how the naive approach may fail on a NE-deterministic SFT such as the Ledrappier SFT. In the depicted example, the initial configuration contains a single 1 (the red cell) in a sea of 0s. Although each 1 is eventually turned into 0, the two foremost 1s (the yellow cells) always lead to more 1s.

Theorem 3.14. (Self-stabilisation of NE-deterministic SFTs)

For every two-dimensional NE-deterministic SFT $X \subseteq \Sigma^{\mathbb{Z}^2}$, there exists a CA $F: \Sigma'^{\mathbb{Z}^2} \rightarrow \Sigma'^{\mathbb{Z}^2}$ with $\Sigma' \supseteq \Sigma$ that stabilises X from finite perturbations in linear time.

We construct a CA $T: \Sigma'^{\mathbb{Z}^2} \rightarrow \Sigma'^{\mathbb{Z}^2}$ with an extended alphabet $\Sigma' \supseteq \Sigma$ that corrects finite islands of defects on X in linear time. The construction is similar to the one-dimensional case (Section 2), except that the patching rule is replaced by the map f stemming from the fact that X is a deterministic tiling space.

Let us thus consider a NE-deterministic SFT $X \subseteq \Sigma^{\mathbb{Z}^2}$, identified by a partial map $f: \Sigma \times \Sigma \rightarrow \Sigma$. For a configuration $y \in \Sigma^{\mathbb{Z}^2}$, we denote by

$$\mathcal{D}(y) \triangleq \{(i, j) : y_{i,j} \neq f(y_{i+1,j}, y_{i,j+1})\}$$

the set of cells at which a *defect* occurs.

As before, the extended alphabet will be $\Sigma' \triangleq \Sigma \times \{\circ, \bullet, \otimes\}$, in which $\{\circ\} \times \Sigma$ is identified with Σ , and \bullet and \otimes represent, respectively, the *trace* and *stop* signals. Again, for convenience, we identify the configurations with alphabet Σ' with pairs (y, α) where $y \in \Sigma^{\mathbb{Z}^2}$ and $\alpha \in \{\circ, \bullet, \otimes\}^{\mathbb{Z}^2}$. The configuration (y, α) has a *defect* at cell (i, j) if either y has a defect at (i, j) , or (i, j) contains a signal symbol \bullet or \otimes . Extending the notation $\mathcal{D}(y)$, we denote the set of defects on (y, α) by

$$\mathcal{D}(y, \alpha) \triangleq \{(i, j) \in \mathbb{Z}^2 : (i, j) \in \mathcal{D}(y) \text{ or } \alpha_{i,j} \neq \circ\}.$$

The CA T will be, as before, constructed as a composition $T_2(T_g T_1^2 T_0)^2$ of four CA maps $T_g, T_0, T_1, T_2: \Sigma'^{\mathbb{Z}^2} \rightarrow \Sigma'^{\mathbb{Z}^2}$.

Patching. The map T_g is defined by $T_g(y, \alpha) \triangleq (y', \alpha')$, where

$$(y'_{i,j}, \alpha'_{i,j}) \triangleq \begin{cases} (f(y_{i+1,j}, y_{i,j+1}), \bullet) & \text{if } \alpha_{i+1,j} \neq \otimes, \alpha_{i,j+1} \neq \otimes, \\ & (i, j) \in \mathcal{D}(y, \alpha) \text{ and } f(y_{i+1,j}, y_{i,j+1}) \text{ exists,} \\ (y_{i,j}, \alpha_{i,j}) & \text{otherwise.} \end{cases}$$

It simply replaces the symbol $y_{i,j}$ with the one prescribed by its North and East neighbours (if none of them contains a stop mark) and leaves a trace mark behind.

Generation of stop signals. The map T_0 is responsible for generating stop signals, and is defined by $T_0(y, \alpha) \triangleq (y, \alpha')$ (i.e., no change on y), where

$$\alpha'_{i,j} \triangleq \begin{cases} \otimes & \text{if } \alpha_{i,j} = \bullet \text{ and } (i, j) \in \mathcal{D}(y), \\ \alpha_{i,j} & \text{otherwise.} \end{cases}$$

Propagation of stop signals. The propagation of the stop signals is governed by the map T_1 , which is defined by $T_1(y, \alpha) \triangleq (y, \alpha')$ (i.e., no change on y), where

$$\alpha'_{i,j} \triangleq \begin{cases} \otimes & \text{if } (\alpha_{i+1,j} = \otimes \text{ or } \alpha_{i,j+1} = \otimes) \text{ and } \alpha_{i,j} = \bullet, \\ \alpha_{i,j} & \text{otherwise.} \end{cases}$$

Fading of the traces. Finally, the map T_2 handles the fading of the traces and is defined by $T_2(y, \alpha) \triangleq (y, \alpha')$ (i.e., no change on y), where

$$\alpha'_{i,j} \triangleq \begin{cases} \circ & \text{if } \alpha_{i+1,j} = \alpha_{i,j+1} = \circ \text{ and } \alpha_{i,j} = \bullet, \\ \alpha_{i,j} & \text{otherwise.} \end{cases}$$

The composition $T_2(T_g T_1^2 T_0)^2$ ensures that the fading of the traces is half as slow as the correction speed, while the stop signals propagate twice as fast as the correction speed.

Proof of Theorem 3.14:

The proof is similar to the proof of Theorem 2.5. □

4. Self-stabilisation of isotropic probabilistic CA

All the cellular automata discussed above provide directional solutions: the cells need to distinguish between the four directions North, South, East, West. In general, finding self-stabilising CA that respect the symmetries of the tiling space appears to be a difficult problem, and not always possible. For instance, in case of the homogeneous space \mathcal{H}_2 , Pippenger has shown that self-stabilisation cannot be achieved by a monotone, self-dual (i.e., respecting the $0 \leftrightarrow 1$ symmetry), centrosymmetric rule [29].

In this section, we will examine to which extent the use of randomness in the evolution of the cellular automata may provide us with a means to design *simpler* solutions. More precisely, we now study probabilistic CA that achieve self-stabilisation with *nearest-neighbour isotropic* rules, that is, rules with von Neumann neighbourhood which treat the neighbours “equally”. Our aim is to show that the use of randomness can extend the range of possibilities for designing self-stabilising processes.

Since this is a broad topic, we will restrict our scope to two examples: finite SFTs and single-cell fillable tiling spaces. In Section 7.2, we will discuss some further questions related to other families of tiling spaces.

4.1. Setting

To begin with, let us recall the notion of *probabilistic CA*, and formulate the concept of self-stabilisation in this context.

Probabilistic cellular automata. The specificity of *probabilistic CA* is that the outcome of the local rule is now a probability distribution on Σ , and the cells of the lattice are updated simultaneously and independently at each time step, according to the distributions prescribed by the local rule.

Formally, the local rule in this case is thus given by a function $\varphi: \Sigma^{\mathcal{N}} \rightarrow \mathcal{P}(\Sigma)$, where $\mathcal{P}(\Sigma)$ denotes the set of probability distributions on Σ , and where we still denote by $\mathcal{N} \subseteq \mathbb{Z}^d$ the (finite) neighbourhood of the rule.

We describe the evolution of the system as a time-homogeneous Markov chain $(\mathbf{x}^t)_{t \in \mathbb{N}}$ with values in $\Sigma^{\mathbb{Z}^d}$, such that for any finite $C \subseteq \mathbb{Z}^d$, and for any $x^0, \dots, x^t \in \Sigma^{\mathbb{Z}^d}$, we have:

$$\begin{aligned} \mathbb{P}(\mathbf{x}_C^{t+1} = x_C^{t+1} \mid \mathbf{x}^0 = x^0, \dots, \mathbf{x}^t = x^t) &= \mathbb{P}(\mathbf{x}_C^{t+1} = x_C^{t+1} \mid \mathbf{x}_{C+\mathcal{N}}^t = x_{C+\mathcal{N}}^t) \\ &= \prod_{c \in C} \varphi((x_{c+i}^t)_{i \in \mathcal{N}})(\{x_c^{t+1}\}). \end{aligned}$$

Self-stabilisation. We say that a probabilistic CA *stabilises* a tiling space $X \subseteq \Sigma^{\mathbb{Z}^d}$ from finite perturbations if

- i) (*consistency*) the configurations of X are absorbing states, that is, if $\mathbf{x}^t \in X$, then $\mathbf{x}^{t+1} = \mathbf{x}^t$,
- ii) (*attraction*) finite perturbations of the elements of X evolve almost surely to X in finitely many steps, that is, if $\mathbf{x}^0 \in \tilde{X}$, then there exists almost surely a time $\mathbf{t} \in \mathbb{N}$ such that $\mathbf{x}^{\mathbf{t}} \in X$.

The first such \mathbf{t} is called the *stabilisation time* (or the *recovery time*) starting from \mathbf{x}^0 . Note that \mathbf{t} is now a random variable. We say that F stabilises X from finite perturbations *in time* $\tau(n)$ if for each $n \in \mathbb{N}$, the maximum of the expected stabilisation time $\mathbb{E}[\mathbf{t}]$ among all possible (deterministic) initial configurations $\tilde{x} \in \Sigma^{\mathbb{Z}^d}$ with $\delta(\tilde{x}, X) = n$ is $\tau(n)$.

4.2. Finite tiling spaces

Let us first examine the case where the tiling space is the set $\mathcal{H}_2 = \{0, \underline{1}\} \subseteq \{0, 1\}^{\mathbb{Z}^2}$ of Example 1.1. Toom's North-East-Center majority rule (see Example 3.1) is a deterministic CA that stabilises \mathcal{H}_2 . We now present an isotropic self-stabilising probabilistic CA for this tiling space.

Let $\Sigma \triangleq \{0, 1\}$, and consider the probabilistic CA *Maj-Random-If-Equal* on $\Sigma^{\mathbb{Z}^2}$, defined on the von Neumann neighbourhood $\mathcal{N} \triangleq \{0, e_1, -e_1, e_2, -e_2\}$ by the local rule $\varphi: \Sigma^{\mathcal{N}} \rightarrow \mathcal{P}(\Sigma)$ where

$$\varphi((x_i)_{i \in \mathcal{N}}) \triangleq \begin{cases} \delta_1 & \text{if } x_{e_1} + x_{e_2} + x_{-e_1} + x_{-e_2} > 2, \\ \delta_0 & \text{if } x_{e_1} + x_{e_2} + x_{-e_1} + x_{-e_2} < 2, \\ \mathcal{B}(1/2) & \text{otherwise,} \end{cases}$$

where δ_1 and δ_0 are the Dirac distributions on 1 and 0, respectively, and $\mathcal{B}(1/2)$ denotes the Bernoulli random variable with parameter 1/2. In words, at every step, the state of each cell is changed to the state which is in majority among its four adjacent cells, and in case of a tie, the tie is broken with a flip of a fair coin, independently of the other cells. A few sample snapshots from the evolution of *Maj-Random-If-Equal* are shown in Figure 16. The continuous-time version of *Maj-Random-If-Equal* was studied by Fontes, Schonmann and Sidoravicius [30] (see Examples 4.2 and 6.4 below).




Figure 16. Snapshots from the evolution of the *Maj-Random-If-Equal* rule.

Proposition 4.1. (Isotropic self-stabilisation of \mathcal{H}_2)

The probabilistic CA defined above stabilises \mathcal{H}_2 from finite perturbations in at most cubic time.

Proof:

Let $x \in \tilde{\mathcal{H}}_2$, and let us assume that the defects of x are initially included in a rectangle R . By symmetry, we can consider that the defects are 1s and that the system needs to return to the all-0 configuration. Observe that over time, the defects always stay within R .

We first determine an *upper bound* for the average time it takes for the 1s of the upper row of R to disappear. Let us number from left to right by $1, \dots, k$ the cells of the upper row of R . We also consider the cell 0 which is on the left of cell 1 and the cell $k + 1$ which is on the right of cell k .

We bound the evolution of the cells $1, \dots, k$ by a new process, designed by imagining that these cells evolve in an environment where for each cell i ($1 \leq i \leq k$), its North neighbour is in state 0 and its South neighbour in state 1 (see Figure 17). Because of the monotonicity of the local rule φ , the new process can be coupled with the original process in such a way that the state of the cells $1, \dots, k$ in the original process remain dominated by their states in the new process (i.e., wherever the former has a 1, so does the latter).




Figure 17. Study of the evolution of the upper row of a rectangle of defects under the Maj-Random-If-Equal probabilistic CA (see proof of Proposition 4.1).

Since the two North and South neighbours have their state fixed, the evolution of the cells $1, \dots, k$ can be modelled as a one-dimensional probabilistic CA $(y^t)_{t \in \mathbb{N}}$ with neighbourhood radius 1 and fixed boundary conditions $y_0^t = y_{k+1}^t = 0$.

Let us analyse the evolution of this one-dimensional probabilistic CA, whose behaviour can be observed in Figure 18. The local rule of this one-dimensional CA is given in the following table:

Value of the neighbourhood	000	001	010	011	100	101	110	111
Probability of symbol 1	0	1/2	0	1/2	1/2	1	1/2	1

Observe that this local rule does not depend on the value of the cell itself, but only of its left and right neighbours. Indeed, if the neighbourhood is in state (x, y, z) , the new state of the central cell is equal to x if $x = z$ and to a random value with distribution $\mathcal{B}(1/2)$ otherwise. This observation allows us to describe the evolution of this CA as the combination of two independent processes, one on the even space-time positions and the other on the odd space-time positions.

Consider first the process on the *odd* space-time positions. In this process, the position of the leftmost cell in state 1 is bounded from below by a symmetric random walk that is reflected on the




Figure 18. Space-time diagrams showing evolutions of the one-dimensional probabilistic CA appearing in the proof of Proposition 4.1, with fixed boundary conditions. Time goes from top to bottom. Blue and white squares respectively represent states 0 and 1. (left) evolution from the all-one configuration; (middle) evolution showing how the central homogeneous zone disappears; (right) evolution with various appearances and reappearances of the central homogeneous zone.

left boundary (cell 0) and vanishes when it reaches the right boundary (cell $k + 1$). One can show that the expected time this random walk needs to reach $k + 1$ is of order k^2 . Indeed, by a standard argument, if T_i denotes the expected time needed to reach $k + 1$ from cell i , then we have the recursion $T_i = 1 + (T_{i-1} + T_{i+1})/2$, with the boundary conditions $T_1 = 1 + T_2$ and $T_{k+1} = 0$. It follows that T_1 is quadratic in k .

The same result holds for the process on the *even* space-time positions. Since the time needed for the one-dimensional CA to reach the all-0 configuration is the maximum of the times of the two processes on the odd and even space-time positions, the former time is also quadratic. In particular, the time it takes for the original two-dimensional CA to wipe out the first row of defects from rectangle R is quadratic in the diameter of the rectangle R . Finally, since we have a linear number of lines to wipe out, the stabilisation time is at most cubic. \square

Experimental evidence suggests that the stabilisation is faster than cubic. Note that in the above argument, we did not use the fact that the different rows of the rectangle evolve simultaneously. We conjecture that the true stabilisation time is in fact quadratic. Our intuition is further supported by the following known result regarding the analogous model in continuous-time.

Example 4.2. (Continuous-time Maj-Random-If-Equal)

Fontes, Schonmann and Sidoravicious studied the continuous-time variant of Maj-Random-If-Equal in which the cells are updated asynchronously, triggered by independent Poisson clocks with rate 1[30].¹ Their Theorem 1.3 states that, in dimension $d \geq 2$, the system stabilises \mathcal{H}_2 from finite perturbations,

¹In fact, Fontes et al. considered a more general family of local rules in which, in case of a neighbourhood tie, the current state of the cell is flipped with probability $0 < \alpha \leq 1$. The Maj-Random-If-Equal rule corresponds to $\alpha = 1/2$.

and that the stabilisation occurs in time $O(n^d)$, in the sense that there exist constants $c, \gamma > 0$ such that the stabilisation time \mathbf{t} starting from a configuration $\tilde{x} \in \{0, 1\}^{\mathbb{Z}^d}$ with $\delta(\tilde{x}, \mathcal{H}_2) = n$ satisfies $\mathbb{P}(\mathbf{t} > cn^d) \leq e^{-\gamma n}$. \circ

Extension to finite SFTs. We can extend the above isotropic probabilistic rule to stabilise any finite SFT in at most cubic time. Indeed, if we define N_1 and N_2 as in Section 3.1, then the following rule is suitable: if a state appears strictly more than twice among $x_{a+N_1, b}, x_{a, b+N_2}, x_{a-N_1, b}, x_{a, b-N_2}$, then this state becomes the new value for $x_{a, b}$; otherwise, choose one of the latter values uniformly at random (taking into account the multiplicities). Again, all the defects stay within some enveloping rectangle and are eventually corrected, and a similar argument shows that the stabilisation time is at most cubic.

4.3. Single-cell fillable tiling spaces

The second case we examine is the one of single-cell fillable tiling spaces, such as k -colourings for $k \geq 5$ (see Section 3.2). Again, a simple isotropic probabilistic rule can then be designed. Interestingly, the new rule turns out to stabilise faster than the one in Section 3.2.

Let $X \subseteq \Sigma^{\mathbb{Z}^2}$ be a single-cell fillable tiling space, and let $\psi: \Sigma^4 \rightarrow \Sigma$ be a function as in Section 3.2, assigning, for any possible choice $(a, b, c, d) \in \Sigma^4$ of symbols surrounding a cell, a consistent value $\psi(a, b, c, d)$. For $x \in \Sigma^{\mathbb{Z}^2}$, recall that we denote by $\mathcal{D}(x)$ the set of cells having a defect, that is,

$$\mathcal{D}(x) \triangleq \{k \in \mathbb{Z}^2 : \exists e \in \{\pm e_1, \pm e_2\}, k \text{ has a defect in direction } e\} .$$

Given $\alpha \in (0, 1)$, we define a probabilistic CA on X which leaves the state of cell k unchanged if $k \notin \mathcal{D}(x)$ and changes it to $\psi(x_{k-e_1}, x_{k-e_2}, x_{k+e_1}, x_{k+e_2})$ with probability α if $k \in \mathcal{D}(x)$. In other words, the CA has a von Neumann neighbourhood $\mathcal{N} \triangleq \{0, e_1, -e_1, e_2, -e_2\}$ and local rule

$$\varphi((x_i)_{i \in \mathcal{N}}) \triangleq \begin{cases} \alpha \delta_{\psi(x_{-e_1}, x_{-e_2}, x_{+e_1}, x_{+e_2})} + (1 - \alpha) \delta_{x_0} & \text{if } (x_i)_{i \in \mathcal{N}} \text{ contains a defect,} \\ \delta_{x_0} & \text{otherwise.} \end{cases}$$

Note that if ψ is isotropic, then so is φ .

To estimate the stabilisation time of this CA, we will use the following standard lemma. Let us recall that a *geometric* random variable with parameter $\varepsilon \in [0, 1]$ is a discrete random variable \mathbf{r} with possible values $k = 1, 2, 3, \dots$ such that $\mathbb{P}(\mathbf{r} = k) = (1 - \varepsilon)^{k-1} \varepsilon$.

Lemma 4.3. (Expectation of maximum of i.i.d. geometric random variables)

For every $\varepsilon \in (0, 1)$, there exist constants $a, b > 0$ such that, whenever $\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_n$ are independent geometric random variables with parameter ε , we have $\mathbb{E}[\max\{\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_n\}] \leq a \log n + b$.

Proof sketch:

This can be shown via a simple comparison with the maximum of n independent exponential random variables, for which the expected value can be calculated explicitly (see e.g. [31]). \square

Proposition 4.4. (Isotropic self-stabilisation of single-cell fillable tiling spaces)

Let $X \subseteq \Sigma^{\mathbb{Z}^2}$ be a single-cell fillable tiling space. Then, for every $\alpha \in (0, 1)$, the probabilistic CA defined above stabilises X from finite perturbations in at most logarithmic time.

Proof:

Let $x \in \tilde{X}$ be an initial configuration. Let $\mathbf{x}^0, \mathbf{x}^1, \dots$ denote the Markov process described by the probabilistic CA starting from $\mathbf{x}^0 = x$. Note that if a cell k is non-defective at time t (i.e., $c \notin \mathcal{D}(\mathbf{x}^t)$), then it will remain non-defective at time $t + 1$, because according to the local rule, the state of k will not change, and by the property of ψ , changes in the state of the neighbours of k cannot create a defect at k . Therefore, for every t we have $\mathcal{D}(\mathbf{x}^{t+1}) \subseteq \mathcal{D}(\mathbf{x}^t)$. On the other hand, at every step, every defective cell has probability at least $\alpha(1 - \alpha)^4$ of becoming non-defective. Indeed, consider a cell $k \in \mathcal{D}(\mathbf{x}^t)$. If at time $t + 1$, cell k is updated according to ψ and none of its four neighbours are updated according to ψ , then k becomes non-defective, and this event occurs with probability at least $\alpha(1 - \alpha)^4$. It follows that inside $\mathcal{D}(x)$, the probabilistic CA behaves like an absorbing finite-state Markov chain that eventually reaches a configuration with no defects.

To analyse the stabilisation time, let us imagine that the process is constructed using a collection $(\mathbf{z}_k^t)_{c \in \mathbb{Z}^2, t \in \mathbb{N}}$ of independent Bernoulli random variables with parameter α , representing the random choices taken at every time step and each cell. Namely, for $t = 0, 1, 2, \dots$ and $k \in \mathbb{Z}^2$, we have

$$\mathbf{x}_k^{t+1} \triangleq \begin{cases} \psi(\mathbf{x}_{k-e_1}^t, \mathbf{x}_{k-e_2}^t, \mathbf{x}_{k+e_1}^t, \mathbf{x}_{k+e_2}^t) & \text{if } \mathbf{z}_k^{t+1} = 1 \text{ and } k \in \mathcal{D}(\mathbf{x}^t), \\ \mathbf{x}_k^t & \text{otherwise.} \end{cases}$$

For each $k \in \mathbb{Z}^2$, let

$$\mathbf{t}_k \triangleq \inf \{t > 0 : \mathbf{z}_k^t = 1 \text{ and } \mathbf{z}_{k+e_1}^t = \mathbf{z}_{k-e_1}^t = \mathbf{z}_{k+e_2}^t = \mathbf{z}_{k-e_2}^t = 0\}$$

denote the first time that k is updated according to ψ and none of its neighbours are updated according to ψ . This is a geometric random variable with parameter $\varepsilon \triangleq \alpha(1 - \alpha)^4$. As we observed above, $k \notin \mathcal{D}(\mathbf{x}^t)$ for all $t \geq \mathbf{t}_k$. In particular, defining $\mathbf{t}_A \triangleq \max\{\mathbf{t}_k : k \in A\}$ for $A \subseteq \mathbb{Z}^2$, we have $\mathbf{x}^t \in X$ for all $t \geq \mathbf{t}_{\mathcal{D}(x)}$. We show that $\mathbb{E}[\mathbf{t}_{\mathcal{D}(x)}]$ is logarithmic in $|\mathcal{D}(x)|$, and hence also in $\delta(\tilde{x}, X)$.

The random variables \mathbf{t}_k (for $k \in \mathbb{Z}^2$) are not independent. However, if k_1, k_2, \dots are cells that do not share neighbours, then $\mathbf{t}_{k_1}, \mathbf{t}_{k_2}, \dots$ are independent. Note that \mathbb{Z}^2 can be partitioned into five

0	1	2	3	4	0	1	2	3	4
3	4	0	1	2	3	4	0	1	2
1	2	3	4	0	1	2	3	4	0
4	0	1	2	3	4	0	1	2	3
2	3	4	0	1	2	3	4	0	1
0	1	2	3	4	0	1	2	3	4

Figure 19. A partitioning of \mathbb{Z}^2 into 5 parts in such a way that the cells in each part do not share neighbours with one another.

parts Q_0, \dots, Q_4 in such a way that the cells in each part do not share neighbours with one another (see Figure 19).

Therefore, setting $\mathcal{D}_i(x) \triangleq \mathcal{D}(x) \cap Q_i$, we have

$$\mathbf{t}_{\mathcal{D}(x)} = \max\{\mathbf{t}_{\mathcal{D}_i(x)} : i = 0, 1, \dots, 4\} \leq \sum_{i=0}^4 \mathbf{t}_{\mathcal{D}_i(x)}$$

and hence $\mathbb{E}[\mathbf{t}_{\mathcal{D}(x)}] \leq \sum_{i=0}^4 \mathbb{E}[\mathbf{t}_{\mathcal{D}_i(x)}]$. By Lemma 4.3, $\mathbb{E}[\mathbf{t}_{\mathcal{D}_i(x)}] = O(\log|\mathcal{D}_i(x)|)$. The claim follows.

An alternative argument for the logarithmic stabilisation time can be given by considering an appropriate martingale as in Ref. [32, Lemma 6]. \square

5. Complexity of self-stabilisation

In this section, we consider the complexity of self-stabilisation as a computational task. One can consider at least three different measures of complexity: the speed of stabilisation, the number of extra symbols, the size of the neighbourhood (Section 5.1). Our focus in this article is on the speed of stabilisation, with a preference for having no extra symbols. In Section 5.3, we show that there is an SFT whose self-stabilisation problem is inherently hard, in the sense that it requires super-polynomial time (unless $\mathbf{P} = \mathbf{NP}$). Interestingly, the optimal speed of stabilisation turns out to be a topological invariant: if two SFTs are topologically isomorphic, then their stabilisation requires roughly the same amount of time (Section 5.2). Here, we will be focusing on the deterministic setting, and leave it open whether the same results hold in the setting of probabilistic cellular automata.

5.1. Measures of complexity

The problem of designing a cellular automaton that stabilises an SFT X is an algorithmic problem, albeit a parallel one with extra requirements. The examples discussed so far suggest that the complexity of this algorithmic problem could drastically vary with X . The efficiency of a CA F in stabilising X can be judged based on the resources it uses:

Speed of stabilisation $\tau_F(n)$

How fast does $\tau_F(n)$ grow with n ? Recall that $\tau_F(n)$ denotes the maximum time it takes for F to correct a finite perturbation \tilde{x} with $\delta(\tilde{x}, X) = n$.

Number of extra symbols κ_F

How many extra states per cell does F have compared to the alphabet of X ?

Neighbourhood radius r_F

How far does the local rule of F need to look in order to update the state of one cell?

The complexity of self-stabilisation for X can be measured by the optimal values of $\tau_F(n)$, κ_F and r_F . We let $\kappa_* \triangleq \min_F \kappa_F$ and $r_* \triangleq \min_F r_F$, where the minimums are over all CA F that stabilise X .

We also allegorically use $\tau_*(n)$ to indicate the optimal (in order of magnitude) speed of stabilisation among all CA F that stabilise X .²

Let us make a couple of remarks about these measures:

1. The stabilisation can be linearly sped up at the cost of increasing the neighbourhood radius. Namely, if F has neighbourhood radius r_F and stabilises X in time $\tau_F(n)$, then F^k has neighbourhood radius kr_F and stabilises X in time $\lceil \frac{1}{k} \tau_F(n) \rceil$. In particular, the value of $\tau_*(n)$ is meaningful only up to a multiplicative constant.
2. As we will see in Proposition 5.3, the speed of stabilisation $\tau_*(n)$ is (almost) invariant under topological isomorphisms. Namely, if two SFTs X and Y are topologically isomorphic, then their minimum stabilisation speeds are roughly the same. We suspect that no such invariance holds for the minimum number of extra symbols κ_* or for the minimum neighbourhood radius r_* : for every SFT X , it should be possible to find an isomorphic SFT Y for which $\kappa_* = 0$ and $r_* = 1$.

In this paper, we have focused on the speed of stabilisation $\tau(n)$ with preference towards having no extra symbols.

5.2. Topologically isomorphic SFTs

In this section, we show that the optimal stabilisation time $\tau_*(n)$ for an SFT is an isomorphism invariant, meaning that topologically isomorphic SFTs have roughly the same optimal stabilisation times.

To prove this claim, let us first recall some terminology from symbolic dynamics (see the monograph by Lind and Marcus [33] for more details).

The space $\Sigma^{\mathbb{Z}^d}$ of all d -dimensional configurations with symbols from a finite alphabet Σ is a compact metrisable space with the product topology. The shift maps $\sigma^k: \Sigma^{\mathbb{Z}^d} \rightarrow \Sigma^{\mathbb{Z}^d}$ (for $k \in \mathbb{Z}^d$) are all continuous. A closed subset $X \subseteq \Sigma^{\mathbb{Z}^d}$ is called a *shift space* if it is invariant under all shifts, that is, $\sigma^k x \in X$ for all $x \in X$ and $k \in \mathbb{Z}^d$. Clearly, every SFT is a shift space. Given a shift space X and a finite set $M \subseteq \mathbb{Z}^d$, we define $L_M(X) \triangleq \{x_M : x \in X\}$ as the set of all patterns with shape M that appear in X .

A *homomorphism* between two shift spaces $X \subseteq \Sigma^{\mathbb{Z}^d}$ and $Y \subseteq \Gamma^{\mathbb{Z}^d}$ is a continuous map $\Phi: X \rightarrow Y$ that commutes with the shifts, that is, $\Phi \circ \sigma^k = \sigma^k \circ \Phi$ for every $k \in \mathbb{Z}^d$. It can be verified that a map $\Phi: X \rightarrow Y$ is a homomorphism if and only if it is realized by a local rule, that is, if and only if there exists a finite set $\mathcal{N} \subseteq \mathbb{Z}^d$ and a map $\varphi: L_{\mathcal{N}}(X) \rightarrow \Gamma$ such that $\Phi(x)_i = \varphi(\sigma^i(x)_{\mathcal{N}})$ for each $x \in X$ and $i \in \mathbb{Z}^d$. A map $\Phi: X \rightarrow Y$ is an *isomorphism* if it is bijective and both Φ and Φ^{-1} are homomorphisms. Since every shift space is compact and Hausdorff, every bijective homomorphism is in fact an isomorphism. Two shift spaces X and Y are *topologically isomorphic* (or *topologically conjugate*) if there is an isomorphism between them.

²Note that this is not a rigorous notation. For instance, it might be that for each $\varepsilon > 0$, there exists an F with $\tau_F(n) = O(n^{1+\varepsilon})$ but no F with $\tau_F(n) = O(n)$, in which case $\tau_*(n)$ is not well defined.

Example 5.1. (Two isomorphic SFTs)

Let $X \subseteq \{0, 1\}^{\mathbb{Z}}$ and $Y \subseteq \{0_0, 0_1, 1\}^{\mathbb{Z}}$ be the 1-step SFTs whose transition graphs are depicted in Figure 20. It is easy to verify that the map $\Phi : Y \rightarrow X$ given by

$$\Phi(y)_i \triangleq \begin{cases} 0 & \text{if } y_i \in \{0_0, 0_1\}, \\ 1 & \text{if } y_i = 1, \end{cases}$$

is an isomorphism. Observe that 0 is a safe symbol for X (see Example 1.3), whereas Y does not have a safe symbol. While the CA F defined in (1) (Section 1.3) stabilises X in one step, it is not immediately clear if F can be “translated” into a CA that stabilises Y . Indeed, Y has more symbols than X , and as a result, a configuration in Y can have many more finite perturbations than the corresponding configuration in X . Nevertheless, having the correspondence between X and Y in mind, one can re-implement the mechanism of stabilisation by F (i.e., replacing each defect with the safe symbol) to obtain a CA that stabilises Y . Namely, the CA $G : \{0_0, 0_1, 1\}^{\mathbb{Z}} \rightarrow \{0_0, 0_1, 1\}^{\mathbb{Z}}$ defined by

$$G(y)_i \triangleq \begin{cases} 0_0 & \text{if } y_{i-1}y_i = 11 \text{ or } y_iy_{i+1} \in \{11, 0_10_0, 0_10_1\} \text{ or } y_iy_{i+1}y_{i+2} = 0_111, \\ 0_1 & \text{if } y_iy_{i+1}y_{i+2} \in \{0_010_0, 0_010_1\}, \\ y_i & \text{otherwise,} \end{cases}$$

stabilises Y (in one step), and is in some sense the “translation” of F . ○



Figure 20. Graphs of allowed transitions for two isomorphic one-dimensional 1-step SFTs.

Example 5.2. (Higher block presentation)

Let $X \subseteq \Sigma^{\mathbb{Z}^d}$ be an SFT and $M \subseteq \mathbb{Z}^d$ be a finite set with at least two elements. Then, the map $\Phi : X \rightarrow (\Sigma^M)^{\mathbb{Z}^d}$ defined by $\Phi(x)_i \triangleq x_{i+M}$ is an isomorphism between X and $Y \triangleq \Phi(X)$. As in the previous example, the isomorphism between X and Y does not extend to a correspondence between the finite perturbations of X and the finite perturbations of Y , and it is not immediately clear how a CA stabilising X from finite perturbations can be translated into a CA stabilising Y from finite perturbations. ○

The following proposition shows that a CA stabilising an SFT X from finite perturbations can always be translated into a CA stabilising an isomorphic SFT Y from finite perturbations (if we appropriately extend the alphabet of Y) while keeping the stabilisation time roughly unchanged.

Proposition 5.3. (Stabilisation time for isomorphic SFTs)

Let X and Y be two topologically isomorphic SFTs. If there exists a CA that stabilises X from finite perturbations in time $\tau(n)$, then there also exists a CA that stabilises Y from finite perturbations in time $\tau(n + O(1))$.

Proof:

Let Σ and Γ be two finite alphabets such that $X \subseteq \Sigma^{\mathbb{Z}^d}$ and $Y \subseteq \Gamma^{\mathbb{Z}^d}$. Let $\Phi: X \rightarrow Y$ be the isomorphism between X and Y , and denote its inverse by Ψ . Let $F: \Sigma^{\mathbb{Z}^d} \rightarrow \Sigma^{\mathbb{Z}^d}$ be a CA that stabilises X from finite perturbations in time $\tau(n)$.

Let us first assume that Φ can be extended to a one-to-one shift-invariant continuous map $\hat{\Phi}: \hat{\Gamma}^{\mathbb{Z}^d} \rightarrow \hat{\Gamma}^{\mathbb{Z}^d}$, where $\hat{\Gamma}$ is a finite alphabet including Γ . Let $\hat{\Psi}: \hat{\Phi}(\Sigma^{\mathbb{Z}^d}) \rightarrow \Sigma^{\mathbb{Z}^d}$ denote the inverse of $\hat{\Phi}$, and observe that $\hat{\Psi}$ is an extension of Ψ , that is, $\hat{\Psi}|_Y = \Psi$. Let $\tilde{\Psi}: \hat{\Gamma}^{\mathbb{Z}^d} \rightarrow \Sigma^{\mathbb{Z}^d}$ be a shift-invariant continuous extension of $\hat{\Psi}$. Such an extension can be constructed by (arbitrarily) completing the local rule of $\hat{\Psi}$. Let us now define $G: \hat{\Gamma}^{\mathbb{Z}^d} \rightarrow \hat{\Gamma}^{\mathbb{Z}^d}$ by $Gy \triangleq \hat{\Phi}F\tilde{\Psi}y$. To see that G stabilises Y from finite perturbations, first note that for every $y \in \hat{\Gamma}^{\mathbb{Z}^d}$, the diagram

$$\begin{array}{ccccccc}
 x & \xrightarrow{F} & Fx & \xrightarrow{F} & F^2x & \xrightarrow{F} & \dots \\
 \uparrow \tilde{\Psi} & & \hat{\Phi} \left(\begin{array}{c} \uparrow \\ \downarrow \end{array} \right) \tilde{\Psi} & & \hat{\Phi} \left(\begin{array}{c} \uparrow \\ \downarrow \end{array} \right) \tilde{\Psi} & & \\
 y & \xrightarrow{G} & Gy & \xrightarrow{G} & G^2y & \xrightarrow{G} & \dots
 \end{array}$$

commutes. If $y \in Y$, then clearly $Gy = y$. Suppose that $y \in \tilde{Y} \setminus \langle \hat{\Gamma} \rangle$, that is, a finite perturbation of an element of Y in $\hat{\Gamma}^{\mathbb{Z}^d}$. Then, $x \triangleq \hat{\Phi}y \in \tilde{X} \setminus \langle \Sigma \rangle$. More specifically, let $\bar{y} \in Y$ be such that $\Delta(\bar{y}, y)$ is finite with diameter n . Then, $\bar{x} \triangleq \tilde{\Psi}\bar{y} = \Psi\bar{y}$ is in X and the diameter of $\Delta(\bar{x}, x)$ is at most $n + C$, where C is the diameter of the neighbourhood of the local rule of $\tilde{\Psi}$. Following the above diagram, the iterations of G on y correspond to the iterations of F on x . Since F stabilises X , we have $F^t x \in X$ for some $t \leq \tau(n + C)$. But as soon as $F^t x \in X$, we also obtain $G^t y = \hat{\Phi}F^t x = \Phi F^t x \in Y$. It follows that G stabilises Y from finite perturbations in time $\tau(n + C)$.

It remains to construct a one-to-one shift-invariant continuous extension $\hat{\Phi}: \Sigma^{\mathbb{Z}^d} \rightarrow \hat{\Gamma}^{\mathbb{Z}^d}$ of Φ . Let $\varphi: L_N(X) \rightarrow \Gamma$ and $\psi: L_M(Y) \rightarrow \Sigma$ be the local rules of Φ and Ψ respectively. Without loss of generality, we assume that Γ and Σ are disjoint. Let $\hat{\Gamma} \triangleq \Gamma \cup \Sigma \cup (\Gamma \times \Sigma)$. Define $\hat{\Phi}$ with the local rule $\hat{\varphi}: \Sigma^{M+N} \rightarrow \hat{\Gamma}$, given by

$$\hat{\varphi}(p) \triangleq \begin{cases} \varphi(p_N) & \text{if } p \in L_{M+N}(X), \\ (\varphi(p_N), p_0) & \text{if } p \notin L_{M+N}(X) \text{ but } p_N \in L_N(X), \\ p_0 & \text{otherwise.} \end{cases}$$

The first case in the definition ensures that $\hat{\Phi}|_X = \Phi$. Furthermore, p_0 can always be recovered from $\hat{\varphi}(p)$, either directly, or by applying ψ on $(\varphi(\sigma^i(p)_N))_{i \in M}$, which can be extracted from $\hat{\varphi}(p)$. This means that $\hat{\Phi}$ is indeed a one-to-one extension of Φ as required. \square

5.3. Super-polynomial hardness

In this section, we prove the following theorem:

Theorem 5.4. (SFT with slow stabilisation)

Let $d \geq 2$. Unless $\mathbf{P} = \mathbf{NP}$, there exists a d -dimensional SFT X which is not stabilised from finite perturbations by any CA in polynomial time.

Let us emphasize that we do not know whether every SFT can be stabilised by some CA or not (see below, Problem 7.6 of Section 7.3). According to the above theorem, assuming $\mathbf{P} \neq \mathbf{NP}$, there exists an SFT for which either there is no solution whatsoever or every solution requires super-polynomial stabilisation time. The proof of Proposition 5.6 below contains an explicit construction of such an SFT, but we do not prove that the constructed SFT is stabilised by some CA.

For concreteness, we present the proof for the two-dimensional case. The proof of the higher-dimensional case carries through similarly.

The *square tiling problem* of a finite set of Wang tiles Θ is the decision problem of whether a square of size n can be tiled admissibly using Θ in such a way as to achieve a prescribed colouring of its boundary. The square tiling problem of every tile set is clearly in class \mathbf{NP} . The existence of a set of Wang tiles for which the square tiling problem is \mathbf{NP} -complete is folklore.

The idea of the proof of Theorem 5.4 is that a CA that stabilises the SFT X associated to a set of Wang tiles Θ can be used to solve a variant of the square tiling problem for Θ . Namely, suppose that F is a CA that stabilises X in time $\tau(n)$, and let r be the neighbourhood radius of F . Then, given a configuration $x \in \tilde{X}$ and a finite set $A \subseteq \mathbb{Z}^2$ such that $x_{\mathbb{Z}^2 \setminus A}$ is globally admissible, the cellular automaton is able to “patch” the defects of x in $\tau(\text{diam}(A))$ steps and turn it into a valid tiling by only changing the states of the cells that are no farther than $r\tau(n)$ from A . Note that only the states of the cells within distance $r\tau(n)$ from A can possibly change during this computation, and only the states of the cells within distance $2r\tau(\text{diam}(A))$ from A are relevant for such changes. Thus, the relevant parts of the computation can be simulated by a Turing machine. As we will see, such a Turing machine can then be used to solve the standard square tiling problem.

To be specific, let us state the latter problem more explicitly. Let $\mathcal{M} \triangleq \{-1, 0, 1\}^2$ denote the Moore neighbourhood. The (*Moore*) *boundary* of a set $A \subseteq \mathbb{Z}^2$ is the set $\partial\mathcal{M}(A) \triangleq \mathcal{M}(A) \setminus A$. For each $n \in \mathbb{N}$, let $S_n \triangleq \{0, 1, \dots, n-1\}^2$. Thus, for $k \geq 0$, $\mathcal{M}^k(S_n)$ represents the square $\{-k, -k+1, \dots, n+k-1\}^2$.

Algorithmic Problem 5.1. (Global tiling patching)

- Parameters: A finite set of Wang tiles Θ and two functions $\alpha, \beta: \mathbb{N} \rightarrow \mathbb{N}$.
- Input: A globally admissible tiling q of $\mathcal{M}^{\alpha(n)+\beta(n)+1}(S_n) \setminus S_n$.
- Task: Find a globally admissible pattern \tilde{q} on $\mathcal{M}^{\alpha(n)+1}(S_n)$ that agrees with q on the band $\partial\mathcal{M}(\mathcal{M}^{\alpha(n)}(S_n))$.

See Figure 21 for an illustration.




Figure 21. Illustration of the global tiling patching problem. (a) The input is a globally admissible pattern on the shaded region. (b) The output is a globally admissible pattern on the shaded region which agrees with the input on the region which is shaded in blue.

Proposition 5.5. (Serial simulation)

Let X be the SFT of the valid tilings of a finite set of Wang tiles Θ . Suppose $X \neq \emptyset$ and that there exists a CA with neighbourhood radius r that stabilises X in time $\tau(n)$. Then, there exists a Turing machine with a two-dimensional tape that solves the global tiling patching problem for Θ and $\alpha(n) \triangleq \beta(n) \triangleq r\tau(n)$ in time $O([n + 4r\tau(n)]^2\tau(n))$.

Proof:

The Turing machine simulates the CA on its two-dimensional tape. Once $\tau(n)$ steps of the CA simulation have been carried out, the Turing machine stops and returns the current configuration of the CA.

A serial simulation of the CA can be done in a standard fashion. The tape has two layers (numbered 0 and 1) for storing the configurations at even and odd time steps. The input is initially provided on layer 0. In an initialization stage, the input pattern is extended by writing an arbitrary symbol from Θ at every position in S_n . When simulating time step t of the CA, the Turing machine reads the configuration at time $t - 1$ from layer $(t - 1) \bmod 2$ and writes the result on layer $t \bmod 2$. At positions in which not enough information is available (i.e., outside $\mathcal{M}^{2r\tau(n)+1-rt}(S_n) \setminus \mathcal{M}^{rt}(S_n)$) the blank symbol is written instead.

In order to keep of track of the number of simulated steps, the cells in $\mathcal{M}^{2r\tau(n)+1}(S_n) \setminus S_n$ are, in the initialization stage, marked with \star . At every stage of the simulation, the marks are also updated: a mark is kept if all its $(2r + 1)^2$ neighbours are marked and is erased otherwise. After $\tau(n)$ stages of the simulation, only the cells in the band $\mathcal{M}^{r\tau(n)+1}(S_n) \setminus \mathcal{M}^{r\tau(n)}(S_n)$ are marked and thus the Turing machine easily recognizes that it has to enter its final stage of computation. In the final stage, the marks and the distinction between the two layers are erased so that only the configuration at time $\tau(n)$ is left on the tape.

The simulation of each time step of the CA takes $O([n + 4r\tau(n)]^2)$ time steps of the Turing machine. Likewise, the initialization stage and the final stage take only $O([n + 4r\tau(n)]^2)$ time steps each. Thus, in overall, the Turing machine accepts or rejects its input in $O([n + 4r\tau(n)]^2\tau(n))$ steps. \square

Proposition 5.6. (NP-hardness)

There exists a finite set of Wang tiles Θ such that for every two functions $\alpha, \beta: \mathbb{N} \rightarrow \mathbb{N}$ with polynomial growth, the global tiling patching problem associated to Θ , α and β is **NP-hard**.

Proof:

Let Θ_0 be a finite set of Wang tiles for which the square tiling problem is **NP-complete**, and let C_0 denote the set of colours that appear in the tiles of Θ_0 . Let Θ_1 be the set of Wang tiles depicted in Figure 22 and their four-fold rotations. Without loss of generality, we assume

$$C_1 \triangleq \{\diamond, \rightarrow, \leftarrow, \uparrow, \downarrow, \dashrightarrow, \dashleftarrow, \updownarrow, \downuparrow\} \cup \{(c, \bullet) : c \in C_0\}$$

is disjoint from C_0 . Let $\Theta = \Theta_0 \cup \Theta_1$.




Figure 22. The extra tiles used in the proof of Proposition 5.6. The tile set Θ_1 consists of all these tiles (for each $c \in C$) and their rotations. The unlabeled edges are coloured with the blank symbol \diamond .

Let $\alpha, \beta: \mathbb{N} \rightarrow \mathbb{N}$ be any two functions with polynomial growth. We show that the square tiling problem for Θ_0 can be reduced in polynomial time to the global tiling patching problem for Θ , α and β . It will then follow that the latter problem is **NP-hard**.

The reduction is illustrated in Figure 23. An instance of the square tiling problem for Θ_0 is given by prescribing colours from C for the boundary of the square S_n for some n as in Figure 23a. This is transformed into a locally admissible tiling of the region $\mathcal{M}^{\alpha(n)+\beta(n)+1}(S_n) \setminus S_n$ with tiles from Θ as depicted in Figure 23b. Let us call the latter tiling q . Observe that the restriction of q to the band $\partial\mathcal{M}(\mathcal{M}^{\alpha(n)}(S_n))$ (the region shaded in blue) enforces the tiling of the band $\partial\mathcal{M}(S_n)$ (the region shaded in red), which in turn fixes the colouring of the boundary of the square S_n . Also observe that if \tilde{q} is a locally admissible extension of q to $\mathcal{M}^{\alpha(n)+\beta(n)+1}(S_n)$, then \tilde{q} is in fact globally admissible (can be extended to the entire plane).

Let $G(q)$ be an answer to the global tiling patching problem for Θ with the tiling of Figure 23b as an input.

First, suppose that there is a locally admissible tiling of S_n using the tiles from Θ_0 that achieves the prescribed colouring in Figure 23a. It is easy to see that in this case, the partial tiling given in

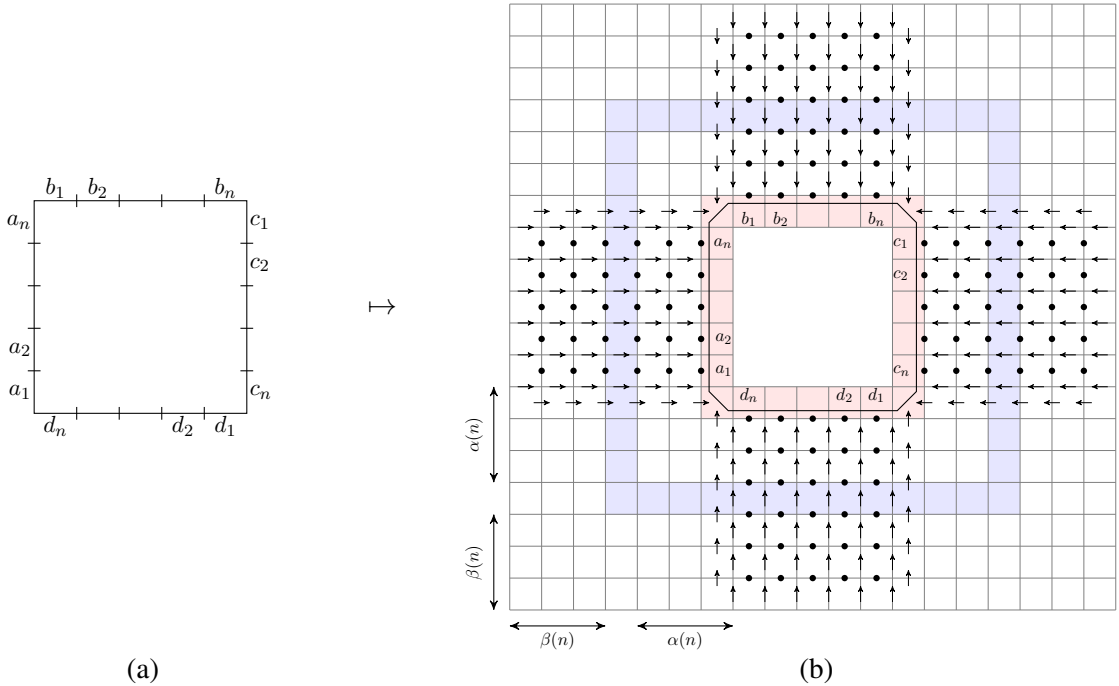


Figure 23. The reduction in the proof of Proposition 5.6. (a) An instance of the square tiling problem, where a_i, b_i, c_i, d_i are colours from C . (b) An instance of the global tiling patching problem. The colour of the edges marked with \bullet (not depicted) match accordingly.

Figure 23b is globally admissible. Thus, $G(q)$ will be globally admissible and will agree with q on $\partial\mathcal{M}(\mathcal{M}^{\alpha(n)}(S_n))$. In particular, by the above remark, $G(q)_{S_n}$ will be a locally admissible tiling of S_n with Θ_0 compatible with the prescribed colouring of Figure 23a. Conversely, assume that the prescribed colouring of Figure 23a cannot be achieved by a locally admissible tiling of S_n using Θ_0 . Then, either $G(q)$ is not locally admissible, or $G(q)$ and q disagree on the band $\partial\mathcal{M}(S_n)$.

We see that, in either case, the answer to the square tiling problem for Θ_0 on Figure 23a can easily (in polynomial time) be extracted from $G(q)$. \square

The proof of Theorem 5.4 is obtained by putting together Propositions 5.5 and 5.6.

Proof of Theorem 5.4.

Let $X \neq \emptyset$ be the SFT of the valid tilings of the tile set Θ in Proposition 5.6. Suppose there exists a CA F that stabilises X in time $\tau(n) = O(n^k)$ for some $k \in \mathbb{N}$. By Proposition 5.5, this implies that there exists a Turing machine solving the global tiling patching problem associated to Θ and $\alpha(n) \triangleq \beta(n) \triangleq r\tau(n)$ (where r is the neighbourhood radius of F) in polynomial time. Unless $\mathbf{P} = \mathbf{NP}$, this is a contradiction. \square

Remark 5.7. In principle, Proposition 5.5 may also lead to other hardness results. For instance, if we have a tile set for which solving the global tiling patching problem with a Turing machine (with a two-dimensional tape) requires more than $O(n^3)$ time steps, then it follows from Proposition 5.5 that there is no CA that stabilises the valid tilings of that tile set in linear time. \diamond

6. Stability against random noise

A true fault-tolerant system should involve efficient error-correction mechanisms so as to maintain its structure and functionality in presence of random noise. In this scenario, noise is present everywhere and throughout the evolution of the system. Such level of stability is achieved by Toom’s NEC-Maj CA and its variants [5] (Examples 3.1 and 6.2) and Gács’s (very sophisticated) reliable CA [7, 8], but appears very challenging in general.

As a modest intermediate step, in this section we consider the problem of self-stabilisation starting from tilings that are perturbed by random noise. (Hence, noise is present only in the initial configuration but not at every time step.) We show that if a CA stabilises from finite perturbations in linear time, then it also stabilises from Bernoulli random perturbations with a sufficiently low density of errors. The argument is based on the idea of sparseness due to Gács [6, 7, 8] and Durand, Romashchenko and Shen [34]. In this section, we restrict ourselves to the case of stabilisation by deterministic CA and leave the corresponding problem for probabilistic CA open.

6.1. Formulation and examples

Let us start by extending the notion of self-stabilisation to the case where the tiling is perturbed with random noise. This is not as straightforward as one might hope for, and there are indeed several variants depending on the requirements and the type of noise. Here we use one such possible notion.

Consider a configuration $x: \mathbb{Z}^d \rightarrow \Sigma$. Let $\varepsilon \geq 0$. An ε -perturbation of x in $\Sigma^{\mathbb{Z}^d}$ is a random configuration \tilde{x} in $\Sigma^{\mathbb{Z}^d}$ with the property that for each finite set $I \subseteq \mathbb{Z}^d$, we have

$$\mathbb{P}(\Delta(x, \tilde{x}) \supseteq I) = \mathbb{P}(\tilde{x}_i \neq x_i \text{ for each } i \in I) \leq \varepsilon^{|I|}.$$

We think of \tilde{x} as a “noisy version” of x where random errors have occurred leading to a change in the state of some cells. A special type of an ε -perturbation is a *Bernoulli ε -perturbation* for which the set $\Delta(x, \tilde{x})$ is a Bernoulli random set with parameter ε , that is, a random subset of \mathbb{Z}^d in which each cell $k \in \mathbb{Z}^d$ is included with probability ε independently of the other cells. The notion of ε -perturbation is however much more general. It turns out that the usual arguments regarding noise-resilience in computational models often work equally well with this more general notion of noise. This was first observed by Toom [5].

The (*Besicovitch*) density of a set $A \subseteq \mathbb{Z}^d$ is defined as

$$\bar{d}(A) \triangleq \limsup_{n \rightarrow \infty} \frac{|A \cap \{-n, -n+1, \dots, n\}^d|}{(2n+1)^d}.$$

Self-stabilisation from random perturbations. Let $\varepsilon, \delta \geq 0$. We say that a CA $F: \Sigma^{\mathbb{Z}^d} \rightarrow \Sigma^{\mathbb{Z}^d}$ δ -stabilises an SFT $X \subseteq \Sigma^{\mathbb{Z}^d}$ from ε -perturbations if

- (i) (*consistency*) the configurations of X are fixed points, that is, $F(x) = x$ for every $x \in X$,
- (ii) for every $x \in X$ and every ε -perturbation \tilde{x} of x ,
 - (ii.a) (*attraction*) $F^t(\tilde{x})$ converges almost surely in the product topology to a (random) configuration \mathbf{y} from X ,
 - (ii.b) (*stability*) The density of the disagreements between \mathbf{y} and x is almost surely less than δ , that is, $\bar{d}(\Delta(x, \mathbf{y})) < \delta$.

We say that F stabilises X from random perturbations if for every $\delta > 0$, there exists an $\varepsilon > 0$ such that F δ -stabilises X from ε -perturbations. The stability condition may sound unnecessary at first sight. The following example illustrates why it is a desired property.

Example 6.1. (Unstable attraction)

Let $F: \{0, 1\}^{\mathbb{Z}} \rightarrow \{0, 1\}^{\mathbb{Z}}$ be the one-dimensional CA with neighbourhood $\mathcal{N} = \{-1, 0, 1\}$ defined by

$$F(x)_k \triangleq \begin{cases} 0 & \text{if } x_{k-1} = x_k = x_{k+1} = 0, \\ 1 & \text{otherwise.} \end{cases}$$

As usual, let $\mathcal{H}_2 \triangleq \{\underline{0}, \underline{1}\}$ be the finite tiling space consisting only of the two homogeneous configurations in $\{0, 1\}^{\mathbb{Z}}$. Then, F satisfies the conditions of consistency and attraction in the definition of stabilisation from random perturbations. However, note that if \mathbf{x} is a non-trivial Bernoulli perturbation of $\underline{0}$, then $F^t(\mathbf{x})$ converges almost surely to $\underline{1}$ rather than to $\underline{0}$. In other words, the system recovers from the initial noise but the distinction between the elements of X is completely lost. \circ

Before stating our result, let us mention some examples.

Example 6.2. (Toom's CA; Continuation of Example 3.1)

Toom's CA has strong forms of stability against noise.

Bušić, Fatès, Mairesse and Marcovici have shown that Toom's CA *classifies* the Bernoulli random configurations according to their *density* [9]. Namely, if we start from a Bernoulli random configuration \mathbf{x} with parameter p (i.e., the state of different cells are chosen independently at random with probability p of choosing 1 and probability $1 - p$ of choosing 0), then

$$\begin{aligned} \text{NEC-Maj}^t(\mathbf{x}) &\rightarrow \underline{0} && \text{if } p < 1/2, \\ \text{NEC-Maj}^t(\mathbf{x}) &\rightarrow \underline{1} && \text{if } p > 1/2. \end{aligned}$$

In particular, for $\delta = 0$ and every $\varepsilon < 1/2$, the CA δ -stabilises the homogeneous tiling space $\mathcal{H}_2 \triangleq \{\underline{0}, \underline{1}\}$ from Bernoulli ε -perturbations.

Moreover, the two homogeneous configurations $\underline{0}$ and $\underline{1}$ remain stable under NEC-Maj even if there is a small independent noise at every time step. In his original paper [5], Toom showed that for

every $\delta > 0$, there exists an $\varepsilon > 0$ such that if $(\mathbf{x}_k^{(t)})_{k \in \mathbb{Z}^2, t \in \mathbb{N}}$ is an ε -perturbed trajectory of NEC-Maj starting from $\underline{0}$ in the sense that for every finite set $I \subseteq \mathbb{Z}^2 \times \mathbb{N}$ of space-time positions, we have

$$\mathbb{P} \{ \mathbf{x}_i^{(s)} \text{ deviates from the local rule of NEC-Maj at every } (i, s) \in I \} \leq \varepsilon^{|I|},$$

then this trajectory remains δ -close to its initial configuration $\underline{0}$, that is,

$$\mathbb{P} (\mathbf{x}_k^{(t)} \neq \underline{0}) \leq \delta \quad \text{for each } (k, t) \in \mathbb{Z}^2 \times \mathbb{N}.$$

In fact, Toom proved the same kind of stability for random perturbations of all monotonic eroders. Note that for NEC-Maj, a similar stability property holds by symmetry for the trajectories starting from $\underline{1}$. \circ

Example 6.3. (GKL and modified traffic; Continuation of Example 2.1)

Using the sparseness result of Durand, Romashchenko and Shen [34], Taati showed that GKL and modified traffic stabilise the homogeneous tiling space $\mathcal{H}_2 \triangleq \{\underline{0}, \underline{1}\}$ from random perturbations [35]. Below, we use the same sparseness result to obtain a more general stabilisation result. \circ

Example 6.4. (Continuous-time Maj-Random-If-Equal; Continuation of Example 4.2)

The notion of stabilisation from random perturbations extends naturally to probabilistic CA and to continuous-time models. Fontes, Schonmann and Sidoravicius [30] proved that, in dimension two and higher, the continuous-time Maj-Random-If-Equal of Example 4.2 stabilises the homogeneous tiling space $\mathcal{H}_2 \triangleq \{\underline{0}, \underline{1}\}$ from random perturbations. We conjecture that a similar result holds for the discrete-time version discussed in Section 4.2. \circ

6.2. Sparseness and the general result

Our objective is to prove the following result.

Theorem 6.5. (Self-stabilisation from random perturbations)

Let $F: \Sigma^{\mathbb{Z}^d} \rightarrow \Sigma^{\mathbb{Z}^d}$ be a CA and $X \subseteq \Sigma^{\mathbb{Z}^d}$ an SFT. If F stabilises X from finite perturbations in linear time, then F also stabilises X from random perturbations.

A proof of Theorem 6.5 will be provided in the next subsection. Here we discuss the idea of the proof and its main ingredient, namely the notion of sparseness. Let us remark that a recent result of Gács [36] would allow us to improve the above result by replacing the “linear time” condition with the “sub-quadratic time” condition. See Remark 6.8 below for more details.

A basic idea in the proof of Theorem 6.5 is that the speed of the propagation of information in a cellular automaton is bounded.

Observation 6.6. (Speed of light)

Let $F: \Sigma^{\mathbb{Z}^d} \rightarrow \Sigma^{\mathbb{Z}^d}$ be a CA with neighbourhood $\mathcal{N} \triangleq \{-r, \dots, r\}^d$. Then, for every $A \subseteq \mathbb{Z}^d$ and $t \geq 0$, and every two configurations $x, y \in \Sigma^{\mathbb{Z}^d}$, we have

- (a) If $x_{\mathcal{N}^t(A)} = y_{\mathcal{N}^t(A)}$, then $F^t(x)_A = F^t(y)_A$.
- (b) If $\Delta(x, y) \subseteq A$, then $\Delta(F^t(x), F^t(y)) \subseteq \mathcal{N}^t(A)$.

Suppose that a CA $F: \Sigma^{\mathbb{Z}^d} \rightarrow \Sigma^{\mathbb{Z}^d}$ stabilises an SFT $X \subseteq \Sigma^{\mathbb{Z}^d}$ from finite perturbations. Since information propagates at bounded speed, this immediately implies that F stabilises starting from any perturbation in which the errors are “sufficiently sparse”. Indeed, if the set of errors in a perturbation \tilde{x} can be decomposed into well-separated finite islands, then under the iterations of F , each island will disappear before sensing or affecting the other islands (see Figure 24).

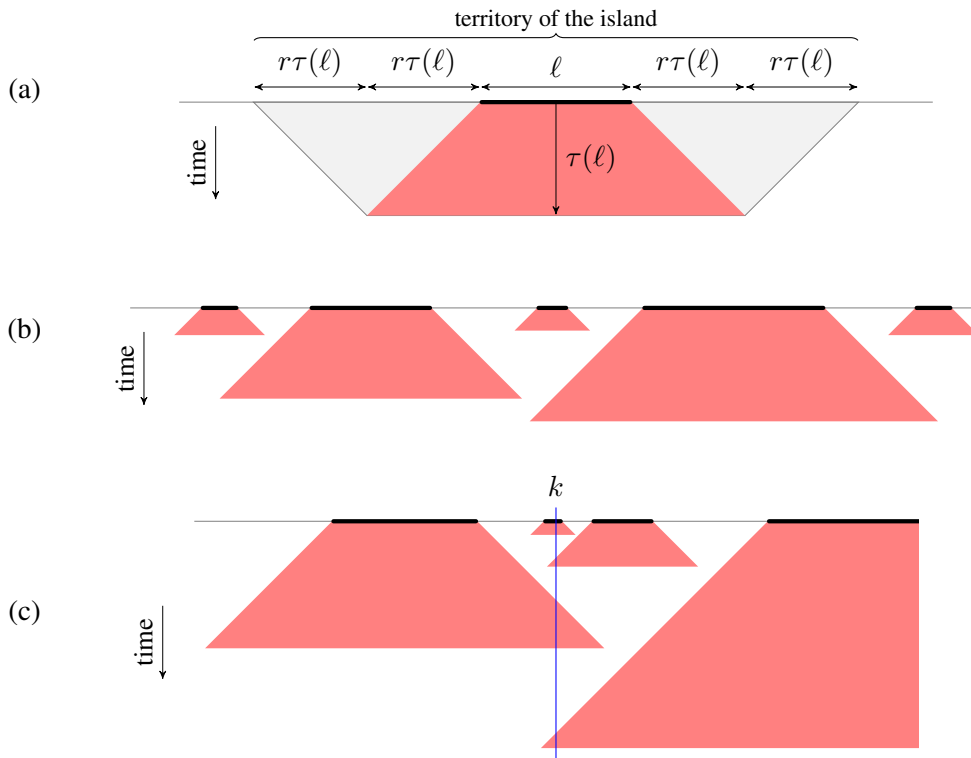


Figure 24. Stabilisation from a sparse set of errors. (a) An isolated island of errors and its territory. (b) Well-separated islands of errors disappear before sensing or affecting one another. (c) The disappearance of all the islands of error is not sufficient for stabilisation.

More specifically, suppose that F stabilises X in time $\tau(n)$. Suppose that F has a neighbourhood radius r and X has an interaction range m . Let $\tilde{x} \in \Sigma^{\mathbb{Z}^d}$ be a (possibly infinite) perturbation of a configuration $x \in X$, and let $S \triangleq \Delta(x, \tilde{x})$ denote the set of positions at which an error has occurred. Let us call a subset $A \subseteq S$ with diameter ℓ an *isolated island* if A is at distance more than $2r\tau(\ell) + m$ from $S \setminus A$, where r denotes the neighbourhood radius of F (Figure 24a). Note that under the iterations of F on \tilde{x} , every isolated island of S is corrected before interacting with the rest of S . Let $S' \subseteq S$ be the subset of S obtained by removing all its isolated islands. Then, again under the iterations of F on \tilde{x} , every isolated island of S' is corrected before interacting with the rest of S' . This reasoning can be repeated to identify an infinite hierarchy of islands in S , each of which is corrected without interacting with the rest of S . If every element of S is included in one of these islands, then this means that

every error on \tilde{x} is eventually patched (Figure 24b). However, note that this alone does not guarantee attraction towards an element of X , for it is possible that a cell $k \in \mathbb{Z}^d$ is within the interaction range of infinitely many islands of S (Figure 24c) and hence never stabilises.

To make this idea precise, we need to introduce the notion of sparseness.

Let $\rho: \mathbb{N} \rightarrow \mathbb{N}$ be an arbitrary function, and let $\mathcal{M} \triangleq \{-1, 0, 1\}^d$ be the Moore neighbourhood. We define the ρ -territory of a finite set $A \subseteq \mathbb{Z}^d$ as the set $\mathcal{M}_\rho(A) \triangleq \mathcal{M}^{\rho(\text{diam}(A))}$ of all cells k that are within $\rho(\text{diam}(A))$ from A . (We use the Moore neighbourhood and the ℓ^∞ distance on \mathbb{Z}^d , but this is an arbitrary choice.) We say that a set $S \subseteq \mathbb{Z}^d$ is ρ -sparse if there is a partitioning $\mathcal{C}(S)$ of S into finite sets, called the ρ -islands of S , such that

- (i) (separation) every two distinct islands $A, B \in \mathcal{C}(S)$ are *well separated* from each other, that is, either $A \cap \mathcal{M}_\rho(B) = \emptyset$ or $\mathcal{M}_\rho(A) \cap B = \emptyset$.
- (ii) (thinness) every cell $a \in \mathbb{Z}^d$ is in the territory of no more than finitely many islands, that is, $\{C \in \mathcal{C}(S) : \mathcal{M}_\rho(C) \ni a\}$ is finite.

The significance of the concept of sparseness in the context of fault-tolerant computation was noticed by Gács [6, 7, 8], who used more sophisticated variants of it in the scenario in which errors due to noise occur at every time step. The above definition of sparseness is equivalent to the one introduced by Durand, Romashchenko and Shen [34], who used it in a context very similar to (but different from) ours.

Let $\varepsilon > 0$. By an ε -random subset of \mathbb{Z}^d , we shall mean a random set $\mathbf{S} \subseteq \mathbb{Z}^d$ with the property that for each finite $I \subseteq \mathbb{Z}^d$,

$$\mathbb{P}(\mathbf{S} \supseteq I) \leq \varepsilon^{|I|}.$$

Durand, Romashchenko and Shen [34, Section 9.2] proved the following.

Theorem 6.7. (Linear sparseness of ε -random sets)

Let $\rho: \mathbb{N} \rightarrow \mathbb{N}$ be such that $\rho(\ell) = O(\ell)$ as $\ell \rightarrow \infty$. For every $\delta > 0$, there exists an $\varepsilon > 0$ such that

- (i) every ε -random set $\mathbf{S} \subseteq \mathbb{Z}^d$ is almost surely ρ -sparse,
- (ii) for every ε -random set $\mathbf{S} \subseteq \mathbb{Z}^d$, the density of the union of the ρ -territories of all the ρ -islands in \mathbf{S} is almost surely less than δ .

Let us clarify that Durand, Romashchenko and Shen proved their result for Bernoulli ε -random sets (for all sufficiently small ε). However, their proof goes through without modification for arbitrary ε -random sets.

Remark 6.8. The proof of Theorem 6.5 relies on the linear sparseness of ε -random sets (Theorem 6.7). Recently, Gács has strengthened the latter sparseness result by proving the sub-quadratic sparseness of ε -random sets [36]. More precisely, he has shown that if $\rho: \mathbb{N} \rightarrow \mathbb{N}$ is any sub-quadratic function (i.e., $\rho(\ell) = O(\ell^\beta)$ for some $\beta < 2$), then for $\varepsilon > 0$ sufficiently small, every ε -random set is almost surely ρ -sparse. Applying the result of Gács in place of Theorem 6.7, we immediately obtain that F stabilises X from random perturbations as long as it stabilises X from finite perturbations in sub-quadratic time. \diamond

6.3. Proof of the result

In this section, we prove Theorem 6.5

Suppose that F stabilises X from finite perturbations in time $\tau(\ell)$. There is no loss of generality to assume that $\tau(\ell)$ is an increasing function with $\tau(\ell) \rightarrow \infty$ as $\ell \rightarrow \infty$. If not, we can replace $\tau(\ell)$ with $\tau'(\ell) \triangleq \max\{\ell, \tau(k) : k \leq \ell\}$ in the argument below.

Suppose that F has neighbourhood radius r so that $\mathcal{M}^r = \{-r, \dots, r\}^d$ is a neighbourhood for the local rule of F . Suppose that X has interaction range m so that it can be identified by a set \mathcal{F} of forbidden patterns with shape $S_m \triangleq \{0, 1, \dots, m-1\}^d$. Define $\rho(\ell) \triangleq 3r\tau(\ell) + m$. (We use this choice rather than $\rho(\ell) \triangleq 2r\tau(\ell) + m$ used in the intuitive explanation above to simplify the proof.) By assumption, $\tau(\ell) = O(\ell)$, hence $\rho(\ell) = O(\ell)$. Let $\delta > 0$, and choose $\varepsilon > 0$ as in Theorem 6.7. We claim that F δ -stabilises X from ε -perturbations.

The consistency property of F is satisfied by assumption. Let $x \in X$, and let \tilde{x} be an ε -perturbation of x . Let $\mathbf{Q} \triangleq \Delta(x, \tilde{x})$ denote the set of positions where an error has occurred. Note that \mathbf{Q} is an ε -random set. Thus, by the choice of ε , we know that \mathbf{Q} is almost surely ρ -sparse. Let $\mathcal{C}(\mathbf{Q})$ be a partition of \mathbf{Q} into ρ -islands, witnessing the ρ -sparseness of \mathbf{Q} . For each $\ell \geq 0$, let us define

$$\begin{aligned} \mathcal{C}_\ell(\mathbf{Q}) &\triangleq \{C \in \mathcal{C}(\mathbf{Q}) : \text{diam}(C) = \ell\}, \\ \mathcal{C}_{>\ell}(\mathbf{Q}) &\triangleq \{C \in \mathcal{C}(\mathbf{Q}) : \text{diam}(C) > \ell\} = \bigcup_{k>\ell} \mathcal{C}_k(\mathbf{Q}), \\ \mathcal{C}_{\geq\ell}(\mathbf{Q}) &\triangleq \{C \in \mathcal{C}(\mathbf{Q}) : \text{diam}(C) \geq \ell\} = \mathcal{C}_\ell(\mathbf{Q}) \cup \mathcal{C}_{>\ell}(\mathbf{Q}). \end{aligned}$$

Note that by the separation property of \mathbf{Q} , each $C \in \mathcal{C}_\ell(\mathbf{Q})$ is at distance more than $\rho(\ell) = 3r\tau(\ell) + m$ from every $C' \in \mathcal{C}_{>\ell}(\mathbf{Q})$ distinct from C .

Lemma 6.9. (Recursive correction)

For all $\ell \geq 0$, we can construct configurations $\tilde{\mathbf{x}}^{(\ell)} \in \Sigma^{\mathbb{Z}^d}$ and $\mathbf{y}^{(\ell)} \in X$ such that the following conditions are satisfied:

- (a) (forgetting) $F^t(\tilde{\mathbf{x}}) = F^t(\tilde{\mathbf{x}}^{(\ell)})$ for all $t \geq \tau(\ell)$, that is, the distinction between $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{x}}^{(\ell)}$ is forgotten in $\tau(\ell)$ time steps.
- (b) (patching) $\mathbf{y}^{(0)} = x$ and $\Delta(\mathbf{y}^{(\ell-1)}, \mathbf{y}^{(\ell)}) \subseteq \bigcup_{C \in \mathcal{C}_\ell(\mathbf{Q})} \mathcal{M}^{r\tau(\ell)}(C) \subseteq \bigcup_{C \in \mathcal{C}_\ell(\mathbf{Q})} \mathcal{M}_\rho(C)$ for $\ell \geq 1$, that is, $\mathbf{y}^{(\ell-1)}$ and $\mathbf{y}^{(\ell)}$ differ only in the territory of the ρ -islands with diameter ℓ in $\mathcal{C}(\mathbf{Q})$.
- (c) (correction) $\Delta(\mathbf{y}^{(\ell)}, \tilde{\mathbf{x}}^{(\ell)}) \subseteq \bigcup_{C \in \mathcal{C}_{>\ell}(\mathbf{Q})} C$, that is, the ρ -islands of errors on $\tilde{\mathbf{x}}^{(\ell)}$ are simply the ρ -islands of diameter larger than ℓ in $\mathcal{C}(\mathbf{Q})$.

Proof:

We start with $\tilde{\mathbf{x}}^{(0)} \triangleq \tilde{\mathbf{x}}$ and $\mathbf{y}^{(0)} \triangleq x$.

In step $\ell \geq 1$, assuming we have constructed $\tilde{\mathbf{x}}^{(\ell-1)}$ and $\mathbf{y}^{(\ell-1)}$, we construct $\tilde{\mathbf{x}}^{(\ell)}$ as follows. Let C_1, C_2, \dots be an enumeration of the elements of $\mathcal{C}_\ell(\mathbf{Q})$. Let $\mathbf{z}^{(n)}$ be the configuration that agrees

with $\tilde{\mathbf{x}}^{(\ell-1)}$ on C_n and with $\mathbf{y}^{(\ell-1)}$ outside C_n . Note that $\mathbf{z}^{(n)}$ is a finite perturbation of $\mathbf{y}^{(\ell-1)}$ with $\Delta(\mathbf{y}^{(\ell-1)}, \mathbf{z}^{(n)}) \subseteq C_n$. Since F stabilises X from finite perturbations in time $\tau(\cdot)$, we know that $F^{\tau(\ell)}(\mathbf{z}^{(n)}) \in X$. By Observation 6.6, the two configurations $\mathbf{y}^{(\ell-1)}$ and $F^{\tau(\ell)}(\mathbf{z}^{(n)})$ agree everywhere except possibly in $\mathcal{M}^{r\tau(\ell)}(C_n)$, that is, the set of positions within distance $r\tau(\ell)$ from C_n . Define

$$\tilde{\mathbf{x}}_k^{(\ell)} \triangleq \begin{cases} F^{\tau(\ell)}(\mathbf{z}^{(n)})_k & \text{if } k \in \mathcal{M}^{r\tau(\ell)}(C_n) \text{ for some } n \in \mathbb{N}, \\ \tilde{\mathbf{x}}_k^{(\ell-1)} & \text{otherwise.} \end{cases}$$

$$\mathbf{y}_k^{(\ell)} \triangleq \begin{cases} F^{\tau(\ell)}(\mathbf{z}^{(n)})_k & \text{if } k \in \mathcal{M}^{r\tau(\ell)}(C_n) \text{ for some } n \in \mathbb{N}, \\ \mathbf{y}_k^{(\ell-1)} & \text{otherwise.} \end{cases}$$

We use induction to show that $\mathbf{y}^{(\ell)} \in X$ and conditions (a)–(c) are satisfied. The case $\ell = 0$ is trivial. Assuming that these conditions are satisfied in the first $\ell - 1$ steps of the construction, we show that they remain satisfied in step ℓ .

(–) Let us first verify that $\mathbf{y}^{(\ell)} \in X$. To see this, note that for each $n \in \mathbb{N}$, the configuration $\mathbf{y}^{(\ell)}$ agrees with $F^{\tau(\ell)}(\mathbf{z}^{(n)})$ not only on $\mathcal{M}^{r\tau(\ell)}(C_n)$ but on $\mathcal{M}^{r\tau(\ell)+m}(C_n)$. Namely, by Observation 6.6, $F^{\tau(\ell)}(\mathbf{z}^{(n)})$ agrees with $\mathbf{y}^{(\ell-1)}$ on $\mathcal{M}^{r\tau(\ell)+m}(C_n) \setminus \mathcal{M}^{r\tau(\ell)}(C_n)$, and on the same set, $\mathbf{y}^{(\ell-1)}$ agrees with $\mathbf{y}^{(\ell)}$ by construction. Now, for every $i \in \mathbb{Z}^d$, the set $i + S_m$ is either entirely in $\mathcal{M}^{r\tau(\ell)+m}(C_n)$ for some n or entirely in $\mathbb{Z}^d \setminus \bigcup_{n=1}^{\infty} \mathcal{M}^{r\tau(\ell)}(C_n)$. In the first case, $\mathbf{y}_{i+S_m}^{(\ell)} = F^{\tau(\ell)}(\mathbf{z}^{(n)})_{i+S_m} \notin \mathcal{F}$ and in the second case $\mathbf{y}_{i+S_m}^{(\ell)} = \mathbf{y}_{i+S_m}^{(\ell-1)} \notin \mathcal{F}$. Thus, $\mathbf{y}^{(\ell)} \in X$ as claimed.

(a) By the induction hypothesis, $F^t(\tilde{\mathbf{x}}) = F^t(\tilde{\mathbf{x}}^{(\ell-1)})$ for $t \geq \tau(\ell - 1)$. Since $\tau(\ell - 1) < \tau(\ell)$, it is enough to show that $F^t(\tilde{\mathbf{x}}^{(\ell)}) = F^t(\tilde{\mathbf{x}}^{(\ell-1)})$ for $t \geq \tau(\ell)$.

Note that for each $n \in \mathbb{N}$, the configuration $\mathbf{x}^{(\ell)}$ agrees with $F^{\tau(\ell)}(\mathbf{z}^{(n)})$ not only on $\mathcal{M}^{r\tau(\ell)}(C_n)$ but on $\mathcal{M}^{3r\tau(\ell)}(C_n)$. Namely, by construction, Observation 6.6, and the fact that the elements of X are fixed points, $F^{\tau(\ell)}(\mathbf{z}^{(n)})$ agrees with $\mathbf{y}^{(\ell-1)}$ outside $\mathcal{M}^{r\tau(\ell)}(C_n)$. Furthermore, by the induction hypothesis, $\Delta(\mathbf{y}^{(\ell-1)}, \tilde{\mathbf{x}}^{(\ell-1)}) \subseteq \bigcup_{C \in \mathcal{C}_{\geq \ell}(\mathbf{Q})} C$. Since C_n has distance more than $\rho(\ell) = 3r\tau(\ell) + m$ from the rest of $\mathcal{C}_{\geq \ell}(\mathbf{Q})$, we find that $\mathbf{y}^{(\ell-1)}$ agrees with $\tilde{\mathbf{x}}^{(\ell-1)}$ on $\mathcal{M}^{3r\tau(\ell)}(C_n) \setminus \mathcal{M}^{r\tau(\ell)}(C_n)$. Lastly, again by construction and the fact that C_n has distance more than $\rho(\ell) = 3r\tau(\ell) + m$ from the rest of $\mathcal{C}_{\ell}(\mathbf{Q})$, the two configurations $\tilde{\mathbf{x}}^{(\ell-1)}$ and $\tilde{\mathbf{x}}^{(\ell)}$ also agree on $\mathcal{M}^{3r\tau(\ell)}(C_n) \setminus \mathcal{M}^{r\tau(\ell)}(C_n)$.

Now, Observation 6.6 and the fact that $F^{\tau(\ell)}(\mathbf{z}^{(n)}) \in X$ is a fixed point imply that for each $n \in \mathbb{N}$, the configurations $F^{\tau(\ell)}(\tilde{\mathbf{x}}^{(\ell-1)})$, $F^{\tau(\ell)}(\mathbf{z}^{(n)})$ and $F^{\tau(\ell)}(\tilde{\mathbf{x}}^{(\ell)})$ agree on $\mathcal{M}^{2r\tau(\ell)}(C_n)$. On the other hand, by construction, $\Delta(\tilde{\mathbf{x}}^{(\ell-1)}, \tilde{\mathbf{x}}^{(\ell)}) \subseteq \bigcup_{n=1}^{\infty} \mathcal{M}^{r\tau(\ell)}(C_n)$. Therefore, according to Observation 6.6, the two configurations $F^{\tau(\ell)}(\tilde{\mathbf{x}}^{(\ell-1)})$ and $F^{\tau(\ell)}(\tilde{\mathbf{x}}^{(\ell)})$ agree outside $\bigcup_{n=1}^{\infty} \mathcal{M}^{2r\tau(\ell)}(C_n)$. We conclude that $F^{\tau(\ell)}(\tilde{\mathbf{x}}^{(\ell-1)})$ and $F^{\tau(\ell)}(\tilde{\mathbf{x}}^{(\ell)})$ agree everywhere. That $F^t(\tilde{\mathbf{x}}^{(\ell-1)}) = F^t(\tilde{\mathbf{x}}^{(\ell)})$ for all $t \geq \tau(\ell)$ follows immediately.

(b) That $\Delta(\mathbf{y}^{(\ell-1)}, \mathbf{y}^{(\ell)}) \subseteq \bigcup_{C \in \mathcal{C}_{\ell}(\mathbf{Q})} \mathcal{M}^{r\tau(\ell)}(C) \subseteq \bigcup_{C \in \mathcal{C}_{\ell}(\mathbf{Q})} \mathcal{M}_{\rho}(C)$ is immediate from the construction.

(c) By the induction hypothesis,

$$\Delta(\mathbf{y}^{(\ell-1)}, \tilde{\mathbf{x}}^{(\ell-1)}) \subseteq \bigcup_{C \in \mathcal{C}_{\geq \ell}(\mathbf{Q})} C, \quad \Delta(\mathbf{y}^{(\ell-1)}, \mathbf{y}^{(\ell)}) \subseteq \bigcup_{C \in \mathcal{C}_{\ell}(\mathbf{Q})} \mathcal{M}^{r\tau(\ell)}(C).$$

On the other hand, by construction, $\tilde{\mathbf{x}}^{(\ell)}$ and $\mathbf{y}^{(\ell)}$ agree on $\bigcup_{C \in \mathcal{C}_{\ell}(\mathbf{Q})} \mathcal{M}^{r\tau(\ell)}(C)$. It follows that $\Delta(\mathbf{y}^{(\ell)}, \tilde{\mathbf{x}}^{(\ell)}) \subseteq \bigcup_{C \in \mathcal{C}_{> \ell}(\mathbf{Q})} C$.

This completes the proof of the lemma. \square

Item (b) in Lemma 6.9 ensures that $\mathbf{y}^{(\ell)} \rightarrow \mathbf{y}$ for some $\mathbf{y} \in X$. Indeed, let $k \in \mathbb{Z}^d$. By the thinness condition in the definition of ρ -sparseness, we have $k \in \mathcal{M}_{\rho}(C)$ for at most finitely many $C \in \mathcal{C}(\mathbf{Q})$. It follows that $\mathbf{y}_k^{(\ell-1)} \neq \mathbf{y}_k^{(\ell)}$ for no more than finitely many values of ℓ . Define \mathbf{y}_k as the eventual value of $\mathbf{y}_k^{(\ell)}$. Then, $\mathbf{y}^{(\ell)} \rightarrow \mathbf{y}$ in the product topology. Furthermore, we have $\mathbf{y} \in X$ because X is closed.

Lemma 6.10. (Attraction)

$F^t(\tilde{\mathbf{x}}) \rightarrow \mathbf{y}$ as $t \rightarrow \infty$.

Proof:

Let $t \geq 0$, and choose ℓ such that $\tau(\ell) \leq t \leq \tau(\ell + 1)$. By item (a) in Lemma 6.9, $F^t(\tilde{\mathbf{x}}) = F^t(\tilde{\mathbf{x}}^{(\ell)})$. By item (c) in Lemma 6.9, $\Delta(\mathbf{y}^{(\ell)}, \tilde{\mathbf{x}}^{(\ell)}) \subseteq \bigcup_{C \in \mathcal{C}_{> \ell}(\mathbf{Q})} C$. This, together with Observation 6.6, implies that $\Delta(F^t(\mathbf{y}^{(\ell)}), F^t(\tilde{\mathbf{x}}^{(\ell)})) \subseteq \bigcup_{C \in \mathcal{C}_{> \ell}(\mathbf{Q})} \mathcal{M}^{rt}(C)$. Recall that $\mathbf{y}^{(\ell)}$ is an element of X and thus $F^t(\mathbf{y}^{(\ell)}) = \mathbf{y}^{(\ell)}$. Note further that for $C \in \mathcal{C}_{> \ell}(\mathbf{Q})$, we have $2rt \leq \rho(\ell + 1) \leq \rho(\text{diam}(C))$, and thus $\mathcal{M}^{rt}(C) \subseteq \mathcal{M}_{\rho}(C)$. It follows that

$$\begin{aligned} \Delta(\mathbf{y}^{(\ell)}, F^t(\tilde{\mathbf{x}})) &= \Delta(F^t(\mathbf{y}^{(\ell)}), F^t(\tilde{\mathbf{x}})) \\ &= \Delta(F^t(\mathbf{y}^{(\ell)}), F^t(\tilde{\mathbf{x}}^{(\ell)})) \\ &\subseteq \bigcup_{C \in \mathcal{C}_{> \ell}(\mathbf{Q})} \mathcal{M}^{rt}(C) \subseteq \bigcup_{C \in \mathcal{C}_{> \ell}(\mathbf{Q})} \mathcal{M}_{\rho}(C). \end{aligned}$$

From item (b) in Lemma 6.9, it also follows that $\Delta(\mathbf{y}, \mathbf{y}^{(\ell)}) \subseteq \bigcup_{C \in \mathcal{C}_{> \ell}(\mathbf{Q})} \mathcal{M}_{\rho}(C)$. Thus,

$$\Delta(\mathbf{y}, F^t(\tilde{\mathbf{x}})) \subseteq \bigcup_{C \in \mathcal{C}_{> \ell}(\mathbf{Q})} \mathcal{M}_{\rho}(C).$$

Let $k \in \mathbb{Z}^d$. By the thinness condition in the definition of ρ -sparseness, we know that $k \in \mathcal{M}_{\rho}(C)$ for no more than finitely many $C \in \mathcal{C}(\mathbf{Q})$. Thus, the value of $F^t(\tilde{\mathbf{x}})_k$ will eventually stabilise at \mathbf{y}_k . We conclude that $F^t(\tilde{\mathbf{x}}) \rightarrow \mathbf{y}$ as $t \rightarrow \infty$. \square

Lemma 6.11. (Stability)

$\bar{d}(\Delta(x, \mathbf{y})) < \delta$ almost surely.

Proof:

In order to have $y_k \neq x_k$, we must have $y_k^{(\ell)} \neq y_k^{(\ell-1)}$ for some $\ell \in \mathbb{Z}$. By item (b) in Lemma 6.9, this means that $k \in \bigcup_{C \in \mathcal{C}(\mathbf{Q})} \mathcal{M}_\rho(C)$. However, by property (ii) in Theorem 6.7,

$$\bar{d} \left(\bigcup_{C \in \mathcal{C}(\mathbf{Q})} \mathcal{M}_\rho(C) \right) < \delta .$$

The claim follows.

This concludes the proof of Theorem 6.5. □

7. Discussion and open problems

7.1. Stabilising 3-colourings

In Section 3, we presented self-stabilising CA for different families of two-dimensional tiling spaces, including k -colourings for $k \neq 3$. When $k \neq 3$, the k -colourings have the property that any configuration with a finite island of defects can be corrected in a purely local manner, by modifying the configuration only in a bounded region around the defects. In contrast, as remarked in Example 1.5, this is not always possible for two-dimensional 3-colourings. This makes the self-stabilisation problem more challenging.

Question 7.1. (Self-stabilisation of 3-colourings)

Is there a deterministic CA that stabilises 3-colourings from finite perturbations in polynomial time, ideally without additional symbols?

We conjecture that it is not possible to stabilise 3-colourings in linear time, but that it could be possible in quadratic time, at least if we allow additional symbols.

In order to shed some light on this problem, let us describe a representation of the two-dimensional 3-colourings based on the configurations of the so-called *six-vertex* model, and in the process, explain Figure 1. The correspondence between 3-colourings and the six-vertex model was first discovered by A. Lenard (see [37, Section 8.13]).

Connection with the six-vertex model. Let us start with a valid 3-colouring in two dimensions. For each pair of neighbouring cells (horizontal or vertical), let us draw an arrow on the boundary between the two cells according to the following rule. Let q and q' be the colours of the two neighbouring cells. Since $q' \neq q$, we either have $q' = q + 1 \pmod{3}$ or $q' = q - 1 \pmod{3}$. Depending on this, we draw the arrow in one direction or the other:

- The arrow on a vertical boundary is directed upwards if the colour on its right is one more than the colour on its left (modulo 3), and downwards otherwise.
- The arrow on a horizontal boundary is directed towards the right if the colour below it is one more than the colour above it (modulo 3), and towards the left otherwise.




Figure 25. The convention used for encoding 3-colouring configurations in the six-vertex model and an example of a configuration with its associated six-vertex image.

These conventions are depicted in Figure 25.

One can then check that starting from a 3-colouring, the resulting arrow configuration is such that at each vertex, there are exactly two incoming arrows and two outgoing arrows. Conversely, from a six-vertex configuration (i.e., a configuration of arrows with equal number of incoming and outgoing arrows at each vertex), there are exactly three 3-colourings giving rise to that arrow configuration. (Once we choose the colour of one cell, all the other colours can be deduced.)

Figure 25 shows an example of such an encoding of a valid 3-colouring. By contrast, Figure 1 depicts a finite perturbation of a 3-colouring. Notice that, in Figure 1, the arrows pointing to the South and the ones pointing to the West are drawn in bold. In a valid six-vertex configuration, the knowledge of these two types of arrows is sufficient to fully describe the configuration: indeed, the other horizontal or vertical arrows have to be East or North arrows, respectively. In Figure 1, we have also shaded the unperturbed portion of the configuration, thus the perturbed region consists in the inner (unshaded) square. Therefore, in order to correct the tiling, we need to fill in the inner square with an admissible configuration. One can verify that the only way to do so corresponds to a six-vertex configuration that has a direct downward vertical path. Indeed, there is only one bold incoming arrow and one bold outgoing arrow at the boundary of the square, and these two arrows have to be connected.

The example in Figure 1 shows that we can construct finite perturbations of a 3-colouring which contain only two defective cells (i.e., a single interface with same colour), but for which in order to obtain a valid configuration, one needs to modify an arbitrarily large domain. In fact, taking a limit of such configurations, one obtains a configuration with only two defective cells which is not a finite perturbation of any valid 3-colouring.

A sequential correction procedure. In order to tackle the problem of self-stabilisation for 3-colourings, it may be helpful to address the simpler question of how to correct a perturbed 3-colouring with a conventional, sequential, non-local algorithm.

Suppose we are given a 3-colouring with a finite number of defects and a square region S of size ℓ containing all the defective cells. Is there a simple (sequential, non-local) algorithm to decide if the tiling can be corrected by modifying only the colour of the cells inside S ? The 3-colouring problem on general graphs is **NP**-complete, and an exhaustive search takes an exponential amount of time in ℓ^2 . Nevertheless, in this case, the representation in terms of the six-vertex configurations allows us to solve the decision problem in linear time, and to find the actual correction (when it exists) in quadratic time.

More precisely, consider the pattern of incoming and outgoing arrows on the boundary of S . In order to decide if the colouring of S can be corrected, we just need to know if it is possible to pair the incoming and outgoing arrows on the boundary of S in an admissible way. This is easy to do sequentially in linear time. Starting from the NE-corner, let us enumerate the incoming arrows on the North and the West sides from 1 to n_{in} counter-clockwise, and the outgoing arrows on the East and the South sides from 1 to n_{out} clockwise. The square can be coloured if we can match each incoming arrow number k with the outgoing arrow number k by a SE-path of arrows. (In particular, this would imply $n_{in} = n_{out}$.) In order to know if this can be done, we try to match successively the incoming and outgoing arrows from 1 to $n_{in} = n_{out}$ with disjoint paths, by moving East if the edge has not already been selected, and South otherwise. As an additional condition, we need to ensure that at each step, the path does not go beyond the corresponding outgoing arrow or come across another path. This procedure succeeds if and only if there is at least one admissible matching. See Figure 26 for an illustration.




Figure 26. Example of valid matching of the arrows of the contour, with the procedure that match successively the incoming and outgoing arrows, by moving East if possible, and South otherwise.

At this point, we may want to try to turn the above sequential procedure into a self-stabilising CA. The CA would start by marking squares encompassing the defects. On each such square, the CA would then simulate a Turing machine that performs the above sequential procedure. If the procedure is successful, a signal is propagated throughout the square in order to erase the markings. If the procedure is unsuccessful, another signal is sent to increase the size of the square, and to repeat the simulation of the Turing machine on this larger square. If two squares collide, they merge and form a larger square. Such a CA can indeed be constructed. However, the main difficulty in turning such a construction into a self-stabilising CA is that the initial perturbation could now involve the extra states. The construction must ensure stabilisation starting from any such perturbation, not just the perturbations involving the three colours.

7.2. Probabilistic self-stabilisation

In this paper, we have barely touched the topic of self-stabilisation with probabilistic rules. There are several questions left to be answered, and much more to be explored.

An isotropic candidate for stabilising 4-colourings. In Section 4, we have presented isotropic probabilistic cellular automata achieving self-stabilisation on two-dimensional k -colourings, with $k = 2$ and $k \geq 5$. We now propose a rule which we believe does the job for $k = 4$, but for which we have no formal proof of convergence. The idea is to modify the method used for the case $k \geq 5$, and make an exception when there is no colour available to directly correct a defective cell. More specifically, for $\alpha \in (0, 1)$, we define a probabilistic CA with the following rule. If a cell k is not defective, its state is kept unchanged. Otherwise, the state of k is changed, with probability α , to a colour which is chosen at random from among all colours consistent with the current colours of its four neighbours. If no consistent colour exists, a colour is chosen at random from among all possible colours. As usual, different cells are updated independently. Experimentally, we have observed that this rule rapidly corrects the defects. However, unlike the case $k \geq 5$, for $k = 4$, we cannot ensure with the above rule that the defects stay in some bounded area. See Figure 27 for an example of stabilisation in this probabilistic CA.




Figure 27. Illustration of the evolution of the probabilistic CA proposed in Section 7.2 for stabilising 4-colourings. The dots indicate the defective cells.

Question 7.2. (Probabilistic self-stabilisation of 4-colourings)

Does the probabilistic CA defined above stabilise 4-colourings from finite perturbations, for any or for some values of $\alpha \in (0, 1)$?

We suspect that the answer is positive. To support this claim, one can try to look for configurations for which this rule could potentially fail to stabilise. Consider the configuration depicted in Figure 28. It has only two defective cells in the center. Furthermore, all cells (including the defective ones), see the three other colours in their neighbourhoods. Consequently, if one of the defective cells changes its state alone, it will remain defective. For this specific configuration, some kind of coordination is thus necessary, which cannot here occur by a specific mechanism as in the deterministic case. This at first might suggest that the defects may propagate arbitrary far from their origin. However, we have experimentally observed that this is not the case: defects have a tendency to stay in the same area,

and the correcting process is more rapid than the diffusion of the defects. Surprisingly enough, even when the cells are updated successively at random (i.e., the fully asynchronous case), we also noticed that the rule succeeds in correcting the defects. Indeed, when the defects propagate, they modify the configuration in such a way that the property of seeing three different colours in the neighbourhood is lost, which finally enables a correction to take place. This observation supports the idea that when we use parallel updates, as we do in our rule, we can only increase the possibilities of correction.

2	3	0	1	2	3	0	1	2	3	0	1
0	1	2	3	0	1	2	3	0	1	2	3
3	0	1	2	3	0	0	1	2	3	0	1
1	2	3	0	1	2	3	0	1	2	3	0
3	0	1	2	3	0	1	2	3	0	1	2

Figure 28. A 3-colouring with a single defect (i.e., two defective cells). The configuration has the property that every cell sees the other three colours in its neighbourhood.

No isotropic candidate for stabilising 3-colourings. In contrast with the previous cases, when we have three colours, we cannot use the method of taking an available colour or a random colour when no colour is available. We noticed experimentally that the errors diffuse and we could not find any rule that keeps them confined, even in statistical terms.

Question 7.3. (Probabilistic self-stabilisation of 3-colourings)

Is there a probabilistic CA that stabilises 3-colourings from finite perturbations, ideally rapidly and having the same symmetries (isotropy, symmetry of colours) as the tiling space?

Self-stabilisation of the Maj-Random-If-Equal rule. In Section 4.2, we showed that the Maj-Random-If-Equal probabilistic CA stabilises the homogeneous space \mathcal{H}_2 in no more than cubic time.

Question 7.4. (Speed of stabilisation in the Maj-Random-If-Equal CA)

What is the precise speed of stabilisation in the Maj-Random-If-Equal probabilistic CA?

We conjecture that the stabilisation occurs in quadratic time. See the discussion after the proof of Proposition 4.1.

Self-stabilisation from random perturbations. Does the result of Theorem 6.5 extend in some form to probabilistic CA?

Question 7.5. (Stabilisation from random perturbations)

Suppose that a probabilistic CA stabilises an SFT X from finite perturbations in linear (or sub-quadratic) time. Does the CA also stabilise X from random perturbations?

As a special case, it would be interesting to know if the `Maj-Random-If-Equal` rule stabilises from random perturbations, as in the case of its continuous-time counterpart (see Example 6.4).

7.3. Other open questions

Stabilising any tiling space. In the current paper, we have searched for *efficient* solutions to the self-stabilisation problem for specific classes of SFTs. However, even if we drop the efficiency ambition, it is not clear if one can always find a CA stabilising any given SFT.

Question 7.6. (General solution)

Is it true that for every SFT X , there exists a deterministic CA that stabilises X from finite perturbations, possibly in exponential time in the number of errors? What if we require the solution to have no extra symbols compared to the alphabet of the SFT?

We conjecture that an approach similar to the one suggested for 3-colourings in the last paragraph of Section 7.1 should be possible for a general SFT. In particular, one should be able to come up with a general construction for translating a Turing machine that solves the functional version of the square tiling problem for a finite set of Wang tiles (see Section 5.3) into a self-stabilising CA for the valid tilings. This approach, if successful, stabilises the SFT in at most exponential time by testing all the possible ways to fill in a given region. As before, the main difficulty would be to handle the appearance of the extra symbols in the initial perturbation.

Stabilisation in presence of temporal noise. Recall from Example 6.2 that Toom’s `NEC-Maj` CA maintains a form of stability on \mathcal{H}_2 even in presence of temporal noise. A comparison argument with Toom’s CA can be used to show that the two CA discussed in Sections 3.1 and 3.2 have the same stability property.

Question 7.7. (Self-stabilisation in presence of temporal noise)

Under which conditions does a CA stabilise an SFT in presence of temporal noise?

In particular, do the CA discussed in Sections 2 and 3.4 self-stabilise in presence of temporal noise?

Acknowledgments

We would like to thank Peter Gács and Ilkka Törmä for helpful discussions and for communicating their results with us. We are grateful to Daniel Martins Fiebich for pointing out some inaccuracies in an earlier version of the paper, in particular in the construction of Section 2. We would also like to thank the two anonymous referees for their careful reading of the text and for providing numerous helpful suggestions.

References

- [1] Dijkstra EW. Self-stabilizing systems in spite of distributed control. *Communications of the ACM*, 1974. 17(11):643–644. doi:10.1145/361179.361202.

- [2] Dolev S. Self-Stabilization. MIT Press, 2000. ISBN:9780262041782.
- [3] Altisen K, Devismes S, Dubois S, Petit F. Introduction to Distributed Self-Stabilizing Algorithms. Morgan & Claypool Publishers, 2019. doi:10.2200/S00908ED1V01Y201903DCT015.
- [4] Toom AL. Nonergodic Multidimensional System of Automata. *Problems of Information Transmission*, 1974. **10**(3):239–246.
- [5] Toom AL. Stable and attractive trajectories in multicomponent systems. In: Dobrushin RL, Sinai YG (eds.), *Multicomponent Random Systems*, pp. 549–575. Marcel Dekker. ISBN:9780824768317, 1980.
- [6] Gács P, Reif J. A simple three-dimensional real-time reliable cellular array. *Journal of Computer and System Sciences*, 1988. **36**(2):125–147. doi:10.1016/0022-0000(88)90024-4.
- [7] Gács P. Reliable computation with cellular automata. *Journal of Computer and System Sciences*, 1986. **32**(1):15–78. doi:10.1016/0022-0000(86)90002-4.
- [8] Gács P. Reliable cellular automata with self-organization. *Journal of Statistical Physics*, 2001. **103**(1–2):45–267. doi:10.1023/A:1004823720305.
- [9] Bušić A, Fatès N, Mairesse J, Marcovici I. Density classification on infinite lattices and trees. *Electronic Journal of Probability*, 2013. **18**:51. doi:10.1214/EJP.v18-2325.
- [10] de Oliveira PPB. On Density Determination With Cellular Automata: Results, Constructions and Directions. *Journal of Cellular Automata*, 2014. **9**(5–6):357–385.
- [11] Richard G. On the Synchronisation Problem over Cellular Automata. In: Vollmer H, BVallée (eds.), *Proceedings of STACS 2017*, volume 66 of *LIPICs*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017 pp. 54:1–54:13. doi:10.4230/LIPICs.STACS.2017.54.
- [12] Fatès N. Remarks on the cellular automaton global synchronisation problem: deterministic versus stochastic models. *Natural Computing*, 2019. **18**(3):429–444. URL <https://hal.inria.fr/hal-01653631>.
- [13] Fatès N, Marcovici I, Taati S. Cellular automata for the self-stabilisation of colourings and tilings. In: *Reachability Problems*, volume 11674 of *Lecture Notes in Computer Science*, pp. 121–136. Springer, 2019. doi:10.1007/978-3-030-30806-3_10.
- [14] Gács P, Kurdyumov GL, Levin LA. One-dimensional uniform arrays that wash out finite islands. *Problems of Information Transmission*, 1978. **14**(3):223–226.
- [15] Gonzaga de Sá P, Maes C. The Gacs–Kurdyumov–Levin Automaton Revisited. *Journal of Statistical Physics*, 1992. **67**(3/4):507–522. doi:10.1007/BF01049718.
- [16] Kari J, Le Gloanec B. Modified traffic cellular automaton for the density classification task. *Fundamenta Informaticae*, 2012. **116**(1–4):141–156. doi:10.3233/FI-2012-675.
- [17] Maass A. On the sofic limit sets of cellular automata. *Ergodic Theory and Dynamical Systems*, 1995. **15**(4):663–684. doi:10.1017/S0143385700008609.
- [18] Törmä I. Personal communication.
- [19] Toom AL, Vasilyev NB, Stavskaya ON, Mityushin LG, Kuryumov GL, Pirogov SA. Discrete local Markov systems. In: Dobrushin RL, Kryukov VI, Toom AL (eds.), *Stochastic cellular systems: ergodicity, memory, morphogenesis*. Manchester University Press. 1990. ISBN:9780719022067.
- [20] Marcus B, Pavlov R. An integral representation for topological pressure in terms of conditional probabilities. *Israel Journal of Mathematics*, 2015. **207**(1):395–433.

- [21] Briceño R, Pavlov R. Strong Spatial Mixing in Homomorphism Spaces. *SIAM Journal on Discrete Mathematics*, 2017. **31**(3):2110–2137. doi:10.1137/16M1066178.
- [22] Chandgotia N, Marcus B. Mixing properties for hom-shifts and the distance between walks on associated graphs. *Pacific Journal of Mathematics*, 2018. **294**(1):41–69. doi:10.2140/pjm.2018.294.41.
- [23] Alon N, Briceño R, Chandgotia N, Magazinov A, Spinka Y. Mixing properties of colorings of the \mathbb{Z}^d lattice. *Combinatorics, Probability and Computing*, 2021. **30**(3):360–373. doi:10.1017/S0963548320000395.
- [24] Berger R. The undecidability of the domino problem. *Memoirs of the American Mathematical Society*, 1966. **66**.
- [25] Robinson RM. Undecidability and nonperiodicity for tilings of the plane. *Inventiones Mathematicae*, 1971. **12**:177–209. doi:10.1007/BF01418780.
- [26] Grünbaum B, Shephard GC. *Tilings and Patterns*. W. H. Freeman and Co., 1987. ISBN:978-0-7167-1193-3.
- [27] Kari J. The Nilpotency Problem of One-Dimensional Cellular Automata. *SIAM Journal on Computing*, 1992. **21**(3):571–586. doi:10.1137/0221036.
- [28] Lukkarila V. The 4-way deterministic tiling problem is undecidable. *Theoretical Computer Science*, 2009. **410**(16):1516–1533. doi:10.1016/j.tcs.2008.12.006.
- [29] Pippenger N. Symmetry in self-correcting cellular automata. *Journal of Computer and System Sciences*, 1994. **49**(1):83–95.
- [30] Fontes LR, Schonmann RH, Sidoravicius V. Stretched exponential fixation in stochastic Ising models at zero temperature. *Communications in Mathematical Physics*, 2002. **228**:495–518. doi:10.1007/s002200200658.
- [31] Eisenberg B. On the expectation of the maximum of IID geometric random variables. *Statistics & Probability Letters*, 2008. **78**(2):135–143. doi:10.1016/j.spl.2007.05.011.
- [32] Fatès N. A tutorial on elementary cellular automata with fully asynchronous updating. *Natural Computing*, 2020. **19**(1):179–197. doi:10.1007/s11047-020-09782-7.
- [33] Lind D, Marcus B. *An Introduction to Symbolic Dynamics and Coding*. Cambridge University Press, 1995. ISBN:0-521-55124-2.
- [34] Durand B, Romashchenko A, Shen A. Fixed-point tile sets and their applications. *Journal of Computer and System Sciences*, 2012. **78**(3):731–764. doi:10.1016/j.jcss.2011.11.001.
- [35] Taati S. Restricted density classification in one dimension. In: Kari J (ed.), *Proceedings of AUTOMATA 2015*, volume 9099 of *LNCS*. Springer, 2015 pp. 238–250. doi:10.1007/978-3-662-47221-7_18.
- [36] Çapuni I, Gács P. A reliable Turing machine, August 9, 2021. URL <https://www.cs.bu.edu/fac/gacs/>.
- [37] Baxter R. *Exactly Solved Models in Statistical Mechanics*. Academic Press, 1982. ISBN:0-12-083180-5.