



HAL
open science

Cold Dynamics in Cellular Automata: a Tutorial

Guillaume Theyssier

► **To cite this version:**

Guillaume Theyssier. Cold Dynamics in Cellular Automata: a Tutorial. Natural Computing, 2022, <10.1007/s11047-022-09886-2>. <hal-03462040v2>

HAL Id: hal-03462040

<https://hal.science/hal-03462040v2>

Submitted on 14 Jun 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Cold Dynamics in Cellular Automata: a Tutorial

Guillaume Theyssier*

June 14, 2022

Abstract

This tutorial is about cellular automata that exhibit *cold dynamics*. By this we mean zero Kolmogorov-Sinai entropy, stabilization of all orbits, trivial asymptotic dynamics, or similar properties. These are essentially transient and irreversible dynamics, but they capture many examples from the literature, ranging from crystal growth to epidemic propagation and symmetric majority vote. A collection of properties is presented and discussed: nilpotency and asymptotic, generic or mu-variants, unique ergodicity, convergence, bounded-changeness, freezingness. They all correspond to the *cold dynamics* paradigm at various degrees, and we study their links and differences by key examples and results. Besides dynamical considerations, we also focus on computational aspects: we show how such *cold cellular automata* can still compute under their dynamical constraints, and what are their computational limitations. The purpose of this tutorial is to illustrate how the richness and complexity of the model of cellular automata are preserved under such strong constraints. By putting forward some open questions, it is also an invitation to look more closely at this *cold dynamics* territory, which is far from being completely understood.

1 Introduction

This tutorial is about cellular automata (CA for short), a well-known class of discrete dynamical systems, useful for modeling natural phenomena, and also a model of computation. Informally, a CA is a lattice of finite-state cells that evolve according to a uniform local rule. CA have been extensively studied as chaotic dynamical systems [44], as models of physical phenomena with positive (Kolmogorov-Sinai) entropy [16, 80], as a model of reversible computing [59, 28] or as groups when considering only reversible CA [9, 69].

None of these kinds of CA are considered here. On the contrary, we are interested in examples with zero (Kolmogorov-Sinai) entropy, that have a strong convergence property, whose asymptotic dynamics is essentially trivial, and that are strongly irreversible. These examples are essentially transient dynamically, they erase information from the initial configuration, and display a strong effect of the *arrow of time*.

We will refer to this class as *cold dynamics*, without trying to precisely define it. It turns out that literature on CA abounds in results and examples that fit into this class. Actually, examples of CA with such kinds of dynamics

*I2M, Université Aix-Marseille, CNRS




Figure 1: Evolution starting from the same initial configuration for the two examples of cold dynamics presented in this section. Positions belonging to sets D_i are represented in black.

were already pointed out by S. Ulam in his pioneering work on crystal (or pattern) growth more than 50 years ago [75]. The computational aspects of some *cold dynamics* CA were also considered more than 40 years ago [77] with the point of view of language recognition. Since then, many theoretical works have focused on some particular *cold* properties, the flagship result among them probably being the undecidability of nilpotency¹, established 30 years ago [42] (and actually much before but not explicitly in [1]).

The purpose of this tutorial is to present and organize this literature around key concepts like convergence, and to illustrate how the richness and complexity of the model of cellular automata is preserved despite the strong constraints behind the *cold dynamics* paradigm.

Concrete examples. Let us consider the discrete plane \mathbb{Z}^2 and say that two elements are neighbors if one is at distance 1 to the north, east, south or west of the other. Now consider any initial set $D_0 \subseteq \mathbb{Z}^2$ and define D_{i+1} for any $i \in \mathbb{N}$ as the union of D_i and all elements of \mathbb{Z}^2 that have exactly one neighbor in D_i . This actually defines a CA with a cold dynamics. Indeed, the sets D_i are only increasing and therefore each position of \mathbb{Z}^2 can only have one of two destinies: either it will never belong to any D_i , or it will appear in some D_i and then stay in all subsequent ones. This first example was introduced in [75] by S. Ulam and it is actually one of the very first CA ever defined. It was thought of as a toy model of crystal growth.

As our second example, consider again any initial set $D_0 \subseteq \mathbb{Z}^2$ and define D_{i+1} as set of positions from \mathbb{Z}^2 such that, among itself and its 4 neighbors, a majority of positions is in D_i (*i.e.* at least 3). In this second example, the destiny of a position is less clear as the local majority could change several times. It turns out that if we consider one step every two, this example also

¹This notion will be presented in Section 4.




Figure 2: A configuration reached by a naive CA model of forest fire: fire (yellow) propagates over neighboring trees (in green) and become ashes (gray) in one step. Shades of gray indicates the time since the fire left the position (the darker the sooner). The density of trees has a huge impact on the fire propagation (in this example the initial density of trees is 0.4).

has a strong convergence property (we will establish this for a large class of examples in Theorem 10). Note that majority dynamics like this one might well be the single most studied class of CA.

An example of evolution of both CA is show in Figure 1. Other examples of two states CA with a *cold dynamics* have been specifically studied [14, 40, 37]. In the context of modeling through CA some propagation phenomena (epidemic, rumor, fire, etc), many models have been proposed following the so-called compartmental paradigm [43]: the population is divided into categories (like Susceptible, Infected and Recovered for the SIR model) and the model describes the local conditions for an individual to change from a category to another. When the spatial aspects are considered, this kind of models naturally give rise to multi-state CA and many of them have a *cold dynamics* (see Figure 2).

One could start to analyze these specific examples in depth, but we will rather try to understand what are the properties that make them special, and work on entire classes of examples.

Main questions around cold dynamics. This tutorial is organized into five parts: we will first precise our formal settings in Section 2 (using the language of orbits, topology and probability measures) and introduce in Section 3 the pivotal notion for cold dynamics: convergence. Then, we will address three general questions as starting points to present many results and examples:

1. *What kind of behaviors can be obtained if we push the convergence property to an extreme and ask that the system collapse all initial configurations into a single one? What if we only specify that the asymptotic dynamics be trivial (a singleton) and how does it depend on the language used (configurations versus probability measures)?*

We will tackle these questions in Section 4 by studying the classical notion of nilpotency and its variants.

2. *How to establish that a given CA is convergent?*

In Section 5, we will both present proof techniques and introduce useful sub-classes of convergent CA (bounded-change and freezing CA) that put stronger constraints on the dynamics but already capture many natural examples, like the two presented above.

3. *How does the ability of general CA to make arbitrary computations survive under the constraint of convergence, bounded-change or freezingness?*

In Section 6, we will see that even under the strongest constraints, computational universality is preserved, but the effectiveness of this computational power greatly varies with the constraints considered.

A tutorial. This tutorial is an invitation to discover what we believe are nice examples, concepts and proof technics. It is not a survey, and we try to find a good compromise between accessibility and generality. We do include some proofs (or sketch of) as they give much insight in the topic. However, many results are stated without proof. This is mostly due to space constraint and not an indication that the result is less important: the reader is warmly invited to consult the corresponding references.

2 Definitions and Notations

A (discrete) dynamical system is a pair (X, F) where X is a compact metric space (seen as a set of possible global states of the system) and $F : X \rightarrow X$ a continuous map (seen as the action or global evolution rule of the system). Given an initial point $x \in X$, its *orbit* is the sequence $x, F(x), F^2(x), F^3(x)$, etc. It represents the evolution of the system at successive *time steps* when initialized in x .

We will consider only two kinds of dynamical systems in this tutorial: deterministic CA acting on *configurations*, and deterministic CA acting on *probability measures*. A CA is defined by some *uniform spatial structure* of cells, an *alphabet* and a *local evolution rule* that defines its dynamics. For the reader unfamiliar with these notions, let us start by a concrete example to fix ideas:

- choose the alphabet to be $\{0, 1\}$,
- choose the spatial structure of cells to be a bi-infinite line,
- configurations are then bi-infinite sequences of 0s and 1s,
- choose the following local rule that defines how the CA transforms a configuration into a new one at each time step: each cell takes as new state the sum modulo 2 of its own state and the one of its right neighbor. Denote by F this *global evolution rule* on configuration.

Starting from a given initial configuration and applying iteratively the global transformation F , we get a sequence of configurations, *i.e.* an *orbit* (see Figure 3). For instance, if c is the configuration where each cell is in state 0, then $F(c) = c$ and we say that c is a *fixed point* of F .

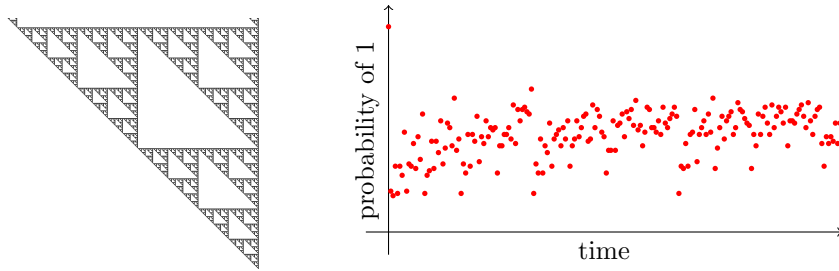


Figure 3: The sum modulo 2 CA seen as two dynamical systems. On the left: example of an orbit of a configuration represented as a space-time diagram (time goes from bottom to top, each horizontal line represents a configuration, white represents state 0 and black state 1). On the right: example of an orbit of a measure (only a partial information on measures is represented).

We can also see our example CA as a deterministic dynamical system acting on *probability measures*. To simplify exposition, let's think of a measure μ as a process that produces (random) configurations. Then one can consider a new measure $F(\mu)$ which is the process that first produces a random configuration by μ and then applies F to get a new random configuration. From there we can of course apply iteratively F and get a sequence of measures: F can thus also be seen as a dynamical system acting on measures. For instance, consider μ_p the random process that chooses the state of each cell independently to be 0 with probability p and 1 with probability $1 - p$ (such measures are called Bernoulli measures). It can be shown that $F(\mu_{1/2}) = \mu_{1/2}$, said differently $\mu_{1/2}$ is an *invariant measure* for F . However $F(\mu_{1/4}) \neq \mu_{1/4}$ because when the states of a given cell and its right neighbor are chosen independently to be 0 with probability $1/4$ and 1 with probability $3/4$, then after one step of F the considered cell has a probability $\frac{1}{4} \times \frac{3}{4} + \frac{3}{4} \times \frac{1}{4} = \frac{3}{8}$ to be in state 1. A measure can convey much more information than the probability for an individual cell to be in some state: for instance the reader can verify that $F(\mu_{1/4})$ is not of the form μ_p for any p (the probability of states in two adjacent cells are no longer independent).

In the remaining of this section we present standard definitions about CA and establish a precise formal framework as well as some notations. Most of the results presented in this tutorial do not require in-depth knowledge of the notions defined below. The reader expecting a more detailed presentation is invited to consult [13, 64, 48, 44, 78].

Spatial structure of cells. A natural and fairly general settings for the spatial cell structure is to consider a finitely generated group \mathbb{G} (without mentioning it each time, all groups \mathbb{G} below will be finitely generated). We will use additive notation for groups and denote by $0_{\mathbb{G}}$ the identity of \mathbb{G} . For $z \in \mathbb{G}$, we denote by $\|z\|_{\mathbb{G}}$ the length n of the smallest sequence of generators $(g_i)_{1 \leq i \leq n}$ such that $g_1 + \dots + g_n = z$, and we let $\|0_{\mathbb{G}}\|_{\mathbb{G}} = 0$. We denote by B_n the set of elements $z \in \mathbb{G}$ such that $\|z\|_{\mathbb{G}} \leq n$. Most of the time we will consider $\mathbb{G} = \mathbb{Z}^d$ and discuss about cases $d = 1$ or $d \geq 2$ (called 1D CA, 2D CA, etc). We will only scratch the surface of the deep and largely open question of how CA theory depends on \mathbb{G} .

We will sometimes use the notion of amenability which is an emblematic example of this dependence through the so-called *Garden of Eden* theorem (see [13]). A finitely additive probability measure on \mathbb{G} is a map μ from subsets of \mathbb{G} to $[0, 1]$ such that $\mu(\mathbb{G}) = 1$ and $\mu(A \cup B) = \mu(A) + \mu(B)$ whenever $A \cap B = \emptyset$. A finitely additive probability measure μ is said left-invariant if $\mu(A) = \mu(g + A)$ for any $g \in \mathbb{G}$ and any $A \subseteq \mathbb{G}$. Then we say that \mathbb{G} is *amenable* if admits a left-invariant finitely additive probability measure. Groups \mathbb{Z}^d are amenable and we will sometimes consider free groups with 2 or more generators as an example of non-amenable group.

Space of configurations and its topology. Given a finite alphabet Q and a group \mathbb{G} , a *configuration* is a map $c : \mathbb{G} \rightarrow Q$. In order to remove some parentheses in expressions, we will often denote $c(z)$ by c_z . For any $q \in Q$, we denote by \bar{q} the uniform configuration $c : z \mapsto q$. The set of configurations $Q^{\mathbb{G}}$ can be endowed with the prodiscrete topology (product of the discrete topology on Q) to form a compact space [44]. Given a finite set $D \subseteq \mathbb{G}$ and a partial configuration $u : D \rightarrow Q$, called a *pattern*, we denote by $[u]$ the *cylinder set* of configurations that coincide with u on domain D :

$$[u] = \{c \in Q^{\mathbb{G}} : c|_D = u\}.$$

By a slight abuse of notation, we denote by $[q]$ for $q \in Q$ the cylinder set $[u]$ of domain $\{0_{\mathbb{G}}\}$ where $u = 0_{\mathbb{G}} \mapsto q$. Cylinder sets are clopen sets and form a base of the prodiscrete topology. This topology is also metrizable as follows. For any pair of configurations $x, y \in Q^{\mathbb{G}}$ define their distance by

$$d(x, y) = \begin{cases} 0 & \text{if } x = y, \\ 2^{-n} & \text{where } n = \max\{i : x|_{B_i} = y|_{B_i}\}. \end{cases}$$

We will often consider sequences of configurations $(x_n)_n$ that converges to some limit, denoted $x_n \rightarrow_n x$ or $\lim_{n \rightarrow \infty} x_n = x$ in the sequel: it will always be in the sense of this prodiscrete topology. $x_n \rightarrow_n x$ is equivalent to the property that for each $g \in \mathbb{G}$ the sequence of states $x_n(g)$ converges to $x(g)$, *i.e.* is ultimately constant equal to $x(g)$. \mathbb{G} naturally acts on $Q^{\mathbb{G}}$ by *translation* as follows: for any $z \in \mathbb{G}$ we define the shift map $\sigma_z : Q^{\mathbb{G}} \rightarrow Q^{\mathbb{G}}$ by

$$\sigma_z(c)_{z'} = c_{z+z'}.$$

Shift maps are homeomorphisms. For $c \in Q^{\mathbb{G}}$, we denote by $\mathcal{O}_{\mathbb{G}}(c)$ its orbit under translations, that is: $\mathcal{O}_{\mathbb{G}}(c) = \{\sigma_z(c) : z \in \mathbb{G}\}$. A natural kind of subset of configurations are closed translation invariant subsets of $Q^{\mathbb{G}}$. Such a subset X is called a *subshift* and it is characterized by its *forbidden language*, *i.e.* the set of patterns it avoids: $\{u : [u] \cap X = \emptyset\}$ (see [48]).

Evolution rule. A CA on $Q^{\mathbb{G}}$ is defined by a local evolution rule $\lambda : Q^N \rightarrow Q$ where N is a finite subset of \mathbb{G} called *neighborhood*, and whose associated global evolution rule $F : Q^{\mathbb{G}} \rightarrow Q^{\mathbb{G}}$ is defined as follows:

$$F(c)_z = \lambda(\sigma_z(c)|_N)$$

for all $z \in \mathbb{G}$. We say that F has *radius* r if $N \subseteq B_r$ (note that the same map F could be defined by different local maps λ with different neighborhoods N , but

finding the minimal such N and hence the minimal radius is usually not important in the sequel). Any CA map F is continuous and commutes with translations: $\sigma_z \circ F = F \circ \sigma_z$. Then $(Q^{\mathbb{G}}, F)$ is a dynamical system. Conversely, by Curtis-Hedlund-Lyndon theorem [39, 13], any dynamical system $(Q^{\mathbb{G}}, \phi)$ where ϕ commutes with translations is actually a CA. A CA is said bijective (resp. surjective, resp. injective) if its corresponding global map on configuration is. A fundamental consequence of Curtis-Hedlund-Lyndon theorem is that if F is the global map of a bijective CA, then the inverse map F^{-1} is also the global map of a CA. For this reason, bijective CA are often referred to as *reversible CA*.

Given a CA $F : Q^{\mathbb{G}} \rightarrow Q^{\mathbb{G}}$, we are naturally interested in orbits, *i.e.* sequences of the form $c, F(c), \dots, F^t(c), \dots$ for some configuration c . A partial information on orbits called *canonical trace* will play an important role: it is the sequence $\mathcal{T}_c^F : \mathbb{N} \rightarrow Q$ defined as $c_{0_{\mathbb{G}}}, F(c)_{0_{\mathbb{G}}}, \dots, F^t(c)_{0_{\mathbb{G}}}, \dots$ for a given $c \in Q^{\mathbb{G}}$. The map $c \mapsto \mathcal{T}_c^F$ is a factor map from $Q^{\mathbb{G}} \rightarrow Q^{\mathbb{N}}$. It is generally neither surjective, nor injective.

Space of probability measures and action of a CA on it. The set of Borel probability measures on $Q^{\mathbb{G}}$ will be denoted by $\mathcal{M}(Q^{\mathbb{G}})$. A measure $\mu \in \mathcal{M}(Q^{\mathbb{G}})$ is a countably additive function from Borel sets to $[0, 1]$ such that $\mu(Q^{\mathbb{G}}) = 1$. The *support* of a measure is the smallest closed set X such that $\mu(X) = 1$. We say that a measure has *full-support* if its support is $Q^{\mathbb{G}}$. Concretely, a measure is characterized by its value on cylinders (by the Carathéodory-Fréchet extension theorem since cylinders form a semi-ring that generates the σ -algebra of Borel sets). It is also convenient to define a measure this way. For instance, a natural and important class of measures are the *Bernoulli measures* which are product measures defined by coefficients β_q for each $q \in Q$ such that $\sum_{q \in Q} \beta_q = 1$ as follows: for any cylinder $[u]$ with $u : D \rightarrow Q$,

$$\mu([u]) = \prod_{z \in D} \beta_{u_z}.$$

The uniform Bernoulli measure is the one where all coefficients β_q are equal to $1/\#Q$ where the notation $\#X$ represents the cardinality of the set X . A Bernoulli measure has full-support if and only if $\beta_q > 0$ for all $q \in Q$. A fundamental property of Bernoulli measures is that they are *translation-ergodic*, meaning that for any Borel set X which is invariant under translation, $\mu(X)$ is either 1 or 0 [78]. The somewhat opposite example of measures are dirac measures. The dirac measure associated to a configuration $c \in Q^{\mathbb{G}}$, denoted δ_c , is defined by

$$\delta_c(A) = \begin{cases} 1 & \text{if } c \in A, \\ 0 & \text{if } c \notin A \end{cases}$$

for any Borel set A .

The set $\mathcal{M}(Q^{\mathbb{G}})$ is naturally endowed with the weak-* topology which is the coarsest topology such that the map $\mu \mapsto \mu([u])$ is continuous for each cylinder set $[u]$. This topology is compact and metrizable by the distance

$$d(\mu, \nu) = \sum_{n \in \mathbb{N}} 2^{-n} \max\{|\mu([u]) - \nu([u])| : D \subseteq B_n, u : D \rightarrow Q\}.$$

In this topology, a sequence of measures $(\mu_n)_n$ converges to μ if and only if, for each cylinder set $[u]$, $\mu_n([u])$ converges to $\mu([u])$.

Any CA $F : Q^{\mathbb{G}} \rightarrow Q^{\mathbb{G}}$ acts naturally on the set of measures $\mathcal{M}(Q^{\mathbb{G}})$ as follows: $F(\mu) = \mu \circ F^{-1}$ or, said differently, $F(\mu)$ is the measure such that for each Borel set one has

$$F(\mu)(X) = \mu(F^{-1}(X)).$$

This way $(\mathcal{M}(Q^{\mathbb{G}}), F)$ is also a dynamical system. Note that for any $c \in Q^{\mathbb{G}}$ one has $F(\delta_c) = \delta_{F(c)}$ because $F(c) \in A \iff c \in F^{-1}(A)$. Hence the dynamical system $(Q^{\mathbb{G}}, F)$ actually embeds into the dynamical system $(\mathcal{M}(Q^{\mathbb{G}}), F)$.

Decision problems and undecidability. To conclude this definition section, let us briefly clarify the settings behind the decision problems that we will consider. Our problems deal with CA and/or configurations as input. A CA is always given by its finite description (alphabet and local evolution rule). A configuration in input will be given as a map $c : \mathbb{G} \rightarrow Q$, which is actually a map from A^* to Q where A is a fixed set of generators of \mathbb{G} . Depending on the context (computable, finite, periodic configuration) there are various ways to represent this map as an input. In the general case, it will be a Turing machine that computes the map.

Our undecidability results will sometimes refer to the arithmetical hierarchy, we refer to [66] for an in depth presentation. Intuitively, it is a hierarchy of levels of uncomputability. Contrary to the polynomial hierarchy in complexity theory, this one is known to be strict. We will only use the first levels of the hierarchy, for which we can give canonical complete problems about Turing machines to give some intuition:

- the halting problem is Σ_1^0 -complete,
- the set of Turing machine that halts on all inputs is Π_2^0 -complete, its complement is Σ_2^0 -complete,
- the set of Turing machine such that the set of inputs on which it doesn't halt is infinite is Π_3^0 -complete.

3 Convergent CA and their basic properties

The central definition considered in this tutorial is convergence.

Definition 1. A dynamical system (X, T) is convergent if for any $x \in X$ the orbit of x under T converges towards some limit denoted $T^\omega(x)$:

$$\forall x \in X : \lim_{n \rightarrow \infty} T^n(x) = T^\omega(x).$$

The convergence of a CA $(Q^{\mathbb{G}}, F)$ is equivalent to the property that for any $c \in Q^{\mathbb{G}}$, \mathcal{T}_c^F is eventually constant. When it is the case, we denote by $\zeta_F(c, z)$ the *freezing time* of cell z in the orbit of c , *i.e.* the first time after which the cell z becomes constant in this orbit:

$$\zeta_F(c, z) = \min\{t_0 \geq 0 : \forall t \geq t_0, F^{t+1}(c)_z = F^t(c)_z\}.$$

We will see various examples of convergent CA in the next sections, but let's start by one of the simplest to fix ideas.




Figure 4: Freezing time line in an orbit of Example 1 supposing that all the right part of the configuration not shown here is in state 1: in black state 1, in white state 0 and the red dots represent the freezing time of each cell as positions in space-time.

Example 1. Let $Q = \{0, 1\}$ and define the CA F on $Q^{\mathbb{G}}$ by

$$F(x)_z = \min_{a \in A} x_{z+a}$$

where A is a fixed finite neighborhood containing the identity of \mathbb{G} . Clearly, F can only turn 1s into 0s in any configuration, so it is a convergent CA. It turns out that this simple dynamical systems actually give much information on the group \mathbb{G} and its generating sets (see [22]). If $\mathbb{G} = \mathbb{Z}$ and $A = \{0, 1\}$ then the freezing time $\zeta_F(c, z)$ of cell z starting from configuration c is either 0 if $c_z = 0$ or if there is no occurrence of state 0 to the right of z in c , or it is the distance to the closest such occurrence otherwise (see Figure 4).

A first observation is that convergence property on configurations extends to the action on measures.

Lemma 1. A CA $(Q^{\mathbb{G}}, F)$ is convergent if and only if $(\mathcal{M}(Q^{\mathbb{G}}), F)$ is convergent.

Proof. Suppose first that $(\mathcal{M}(Q^{\mathbb{G}}), F)$ is convergent and consider any $c \in Q^{\mathbb{G}}$. Since $F^t(\delta_c) = \delta_{F^t(c)}$ and since $(F^t(\delta_c))_t$ converges to some measure $F^\omega(\delta_c)$ by hypothesis, we have that for any $q \in Q$, $\delta_{F^t(c)}([q])$ must converge to $F^\omega(\delta_c)([q])$. But $\delta_{F^t(c)}([q])$ is either 1 if $\mathcal{T}_c^F(t) = q$ or 0 otherwise. This shows that \mathcal{T}_c^F is eventually constant.

Suppose now for the other direction that $(Q^{\mathbb{G}}, F)$ is convergent. Consider any measure μ and any pattern $u \in Q^D$ with $D \subseteq B_n$ for some $n \in \mathbb{N}$. We want to show that $(F^t(\mu)([u]))_t$ converges. Consider the set X_t of initial configurations whose orbit is frozen for all cells in B_n after at most t steps:

$$X_t = \{c \in Q^{\mathbb{G}} : \forall z \in B_n, \zeta_F(c, z) \leq t\}.$$

$(X_t)_t$ is monotone increasing and $\bigcup_t X_t = Q^{\mathbb{G}}$ so the measure of these sets converges to 1. $(F^{-t}([u]) \cap X_t)$ is also monotone increasing (if $c \in F^{-t}([u]) \cap X_t$ then $F^{t'}(c) \in [u]$ for any $t' \geq t$), so the measure of these sets converges (from below) to some α . Then, for any $\epsilon > 0$, there is some $t_0 \in \mathbb{N}$ such that $\mu(X_{t_0}) \geq 1 - \epsilon$ and $\alpha - \epsilon \leq \mu(F^{-t_0}([u]) \cap X_{t_0}) \leq \alpha$. For any $t \geq t_0$, we can write $F^t(\mu)([u])$ as:

$$F^t(\mu)([u]) = \mu(F^{-t}([u]) \cap X_{t_0}) + \mu(F^{-t}([u]) \setminus X_{t_0}).$$

But it actually holds that $F^{-t}([u]) \cap X_{t_0} = F^{-t_0}([u]) \cap X_{t_0}$, so we have actually shown $|F^t(\mu)([u]) - \alpha| \leq \epsilon$, which concludes convergence of $(\mathcal{M}(Q^{\mathbb{G}}), F)$ because μ and u were chosen arbitrarily. \square

A second observation on convergent CA is that they cannot be too chaotic. Precisely, a dynamical system (X, T) is said sensitive to initial conditions if there is ϵ such that for all δ and all $x \in X$ there is some $y \in X$ which is δ -close to x but whose orbit will be ϵ -far from that of x at some future step (see [44] for more details on this notion in the context of CA).

In the following lemma, we use the notation $T_A(c)$ for some finite set $A \subseteq \mathbb{G}$ to denote the trace of configuration c on A , i.e. the map $t \mapsto F^t(c)|_A$.

Lemma 2. *No convergent CA is sensitive to initial conditions.*

Proof. Suppose by contradiction that F on $Q^{\mathbb{G}}$ is convergent and sensitive to initial conditions: there is $N \in \mathbb{N}$ such that for any $c \in Q^{\mathbb{G}}$ and any $p \in \mathbb{N}$ there exists $c' \in Q^{\mathbb{G}}$ such that $c|_{B_p} = c'|_{B_p}$ and $F^t(c)|_{B_N} \neq F^t(c')|_{B_N}$ for some $t \in \mathbb{N}$. Consider any configuration $c_0 \in Q^{\mathbb{G}}$ and $p_1 \geq N$. By sensitivity there is $c' \in Q^{\mathbb{G}}$ and $t_1 \in \mathbb{N}$ such that either $T_{B_N}(c_0)$ or $T_{B_N}(c')$ is non-constant on time interval $[0, t_1]$. Denote by c_1 the one among c_0 and c' that corresponds to the non-constant trace. Let $p_2 = p_1 + N + r(t_1 + 1)$ where r is the radius of F . Applying sensitivity again on c_1 and p_2 we know there exist c' and t_2 such that:

- c_1 and c' are identical on B_{p_2} ;
- therefore, by choice of p_2 , $T_{B_N}(c_1)$ and $T_{B_N}(c')$ coincide on interval $[0, t_1]$;
- $T_{B_N}(c_1)$ and $T_{B_N}(c')$ differ at time t_2 .

So one of c_1 or c' , denoted c_2 , is such that $T_{B_N}(c_2)$ is not constant on interval $[t_1, t_2]$. Going on with the same reasoning we construct a converging sequence $(c_n)_{n \in \mathbb{N}}$ of configurations such that $T_{B_N}(c_n)$ is not constant on each interval $[t_i, t_{i+1}]$ for $0 \leq i < n$. Taking $c = \lim_{n \rightarrow \infty} c_n$ we get a trace $T_{B_N}(c)$ which is not eventually constant contradicting convergence. \square

Our last observation on convergent CA is that they are always irreversible (meaning non-bijective, see Section 2) and even non-surjective (except the identity), at least if \mathbb{G} is amenable. The proof below essentially relies on Poincaré recurrence theorem.

Lemma 3. *Let \mathbb{G} be an amenable group. If a convergent CA on $Q^{\mathbb{G}}$ is surjective, then it is the identity map.*

Proof. Let F be surjective and convergent and let μ denote the uniform product measure on $Q^{\mathbb{G}}$. By [6], μ is preserved under F : $\mu(X) = \mu(F^{-1}(X))$ for any Borel set X (see [11] for an overview of properties of surjective CA on amenable groups). If we suppose that F is not the identity map, then there is a word $u \in Q^{B_n}$ such that for all $x \in [u]$ it holds $F(x)_0 \neq x_0$. We claim that there is a configuration x such that $F^t(x) \in [u]$ for infinitely many t . From this claim we deduce that F is not convergent because the orbit of x is not convergent, and the lemma follows. To prove the claim, let us denote $R_t = \bigcup_{t' \geq t} F^{-t'}([u])$ and $R = \bigcap_t R_t$. By definition R is the set of configurations whose orbit visits $[u]$ infinitely many times and we want to show that $R \neq \emptyset$. We actually show that

$\mu(R) > 0$. Since $R_{t+1} = F^{-1}(R_t)$ for all $t \geq 0$ we have $\mu(R_t) = \mu(R_0)$. Moreover $[u] \subseteq R_0$ so we deduce that $\mu([u] \setminus R_t) \leq \mu(R_0 \setminus R_t) = 0$ since $R_t \subseteq R_0$ and $\mu(R_t) = \mu(R_0)$. Finally, by Boole's inequality we have $\mu([u] \setminus R) = 0$ so $\mu(R) > 0$ since $\mu([u]) > 0$. \square

4 Nilpotency and its Variants

This section explores examples where the asymptotic dynamics is reduced to a singleton. We will consider various properties using either the language of orbits of configurations or measures, or the language of traces.


4.1 Extreme cases of convergence

Let us start by the strongest forms of nilpotency, which are strengthened forms of convergence.

Definition 2. Let F be a CA on $Q^{\mathbb{G}}$.

- F is said nilpotent if there is $x \in Q^{\mathbb{G}}$ and $t_0 \in \mathbb{N}$ such that for all $y \in Q^{\mathbb{G}}$ it holds $F^{t_0}(y) = x$.
- F is said asymptotically nilpotent if there is $x \in Q^{\mathbb{G}}$ such that for all $y \in Q^{\mathbb{G}}$ it holds $F^t(y) \rightarrow_t x$.

Example 2. A local evolution map which is a constant map yields a nilpotent CA, but let us give a non-trivial example. Consider the CA $F : Q^{\mathbb{Z}} \rightarrow Q^{\mathbb{Z}}$ of radius 1 where $Q = \{\square, \blacksquare, \blacksquare\}$ and defined by all the transitions appearing in the following space-time diagram (times goes from bottom to top) and such that any transition not appearing in it produces the state \square :



We claim that $F^{19}(c)_0 = \square$ for any configuration c so that F is nilpotent (we verified this claim by a computer program). However, the diagram above clearly shows that F^{18} is not a constant map. We challenge the reader to find a nilpotent CA with same alphabet and radius but such that F^{19} is not a constant map.

The attracting configuration x in the definitions above must be uniform, because it must attract in particular uniform configurations. Thus, these definitions can be expressed as properties of traces:

- F is nilpotent if and only if there is t_0 and $q \in Q$ such that $\mathcal{T}_x^F(t) = q$ for all $t \geq t_0$ and all $x \in Q^{\mathbb{G}}$.

- F is asymptotically nilpotent if and only if there is $q \in Q$ such that all $x \in Q^{\mathbb{G}}$ there is $t_x \in \mathbb{N}$ such that $\mathcal{T}_x^F(t) = q$ for all $t \geq t_x$.

Note that the state q in the above definitions must be *quiescent*, *i.e.* such that the uniform configuration \bar{q} is a fixed point of F . The property of nilpotency is also well-known as a property of the *limit set*. The limit set of a CA F is the set $\Omega_F = \bigcap_t F^t(Q^{\mathbb{G}})$ (see [17, 41]). Nilpotency is equivalent to Ω_F being a singleton.

Asymptotic nilpotency can be reformulated through the action of the CA on measures.

Lemma 4. *A CA F on $Q^{\mathbb{G}}$ is asymptotically nilpotent if and only there is $q \in Q$ such that for any $\mu \in \mathcal{M}(Q^{\mathbb{G}})$ it holds $F^t(\mu) \rightarrow_t \delta_{\bar{q}}$.*

Proof. The property on measures implies asymptotic nilpotency by the same argument as the first part of the proof of Lemma 1 with the additional fact that $F^\omega(\delta_c) = \delta_{\bar{q}}$ which implies that $F^\omega(c) = \bar{q}$.

Conversely, if we suppose that F is asymptotically nilpotent, we can use the same reasoning as in part two of the proof of Lemma 1. The additional fact in this case is that $F^\omega(c) = \bar{q}$ for all $c \in Q^{\mathbb{G}}$ so that $F^{-t}([u]) \cap X_t = \emptyset$ for all pattern u except the uniform one: $z \in B_n \mapsto q$. We deduce that $F^t(\mu)$ converges towards $\delta_{\bar{q}}$. \square

It turns out that, under some conditions on \mathbb{G} , nilpotency and asymptotic nilpotency are equivalent. One might believe that this can be easily proved by a standard compactness argument, but it is not the case.

Theorem 1 ([38, 67]). *If $\mathbb{G} = \mathbb{Z}^d$ then nilpotency is equivalent to asymptotic nilpotency.*

sketch. Let q be the quiescent state involved in the asymptotic nilpotency property. First, we remark that for any ball B_n there is a uniform bound T_n such that in the orbit of any configuration c there is some time step $t \leq T_n$ for which the cells inside B_n are all in state q : $\forall z \in B_n, F^t(c)_z = q$. Indeed, otherwise we would obtain by compactness a configuration whose orbit would not converge to \bar{q} . Call *finite* any configuration which is everywhere q except on a finite region, and call *mortal* any configuration c such that there is t with $F^t(c) = \bar{q}$.

Let us first establish the theorem for $\mathbb{G} = \mathbb{Z}$ as in [38]. There are two key arguments:

- first, we show nilpotency on finite mortal configurations, *i.e.* there is a uniform T such that for all finite mortal c it holds $F^T(c) = \bar{q}$. This is proved using the remark above about times T_n : if there is no uniform time bound on the mortality of finite mortal configuration, then for arbitrary large B_n and arbitrary large t we can find a finite mortal configuration c such that c is uniformly q on B_n but $F^t(c)_0 \neq q$. From there, we can produce an infinite *nested space-time diagram* as in Figure 5 by taking a well-chosen sequence (c_i) of such configurations and considering their superposition c : it is sufficient that the non- q portions of space-time diagrams are disjoint enough among the c_i (taking into account the radius of the CA). The orbit of this configuration c gives a contradiction with the convergence hypothesis.




Figure 5: The infinite nested space-time diagram contradicting asymptotic nilpotency in Theorem 1. The white part represents state q , the grayed parts represent states other than q , and each nuance of gray corresponds to a particular configuration c_i .

- second, it can be shown that if the CA is not nilpotent, then we can extract finite mortal configurations that die arbitrarily late: indeed, by enclosing a finite configuration alive at time t with blocking words (whose existence is granted by Lemma 2, see [44]), we obtain a mortal configuration still alive at time $t - b$ where b is a constant depending only on the choice of blocking words. This contradicts the previous item, so we deduce that the CA is nilpotent.

To extend the result to higher dimensions, it is enough to show that all finite configurations are mortal (this together with asymptotic nilpotency implies nilpotency, by the nested construction argument). To simplify we restrict to dimension 2 and consider any asymptotically nilpotent CA F . The idea is to reduce to the one-dimensional case by remarking that if a finite configuration c is periodized vertically to form a configuration c' , then c' is mortal because the action of F on c' can be seen as the action of a one-dimensional asymptotically nilpotent CA on a one-dimensional configuration. The key step of the proof is then to show that no finite configuration can spread too much in the horizontal direction. This step is proved by contradiction using again a nested construction of an infinite space-time diagram, but this time a vertical periodization is applied at each step to guarantee mortality. Of course, this argument on the bounded horizontal spreading of finite configurations is also valid vertically, so we deduce that the orbit of any finite configuration remains inside a finite support forever. This is enough to conclude mortality since F is supposed asymptotically nilpotent. \square

The above result has been greatly generalized by [71] in at least two ways: the class of dynamical systems considered (for instance, one can consider CA defined over a subshift rather than the whole space $Q^{\mathbb{G}}$) and the class of groups \mathbb{G} considered. To obtain these results, a notion of tiered dynamical systems was developed which goes way beyond the scope of this tutorial. The authors also coined the term *nil-rigidity* and used it with various classes of dynamical systems. Let us stick to the simplest setting and say that a group \mathbb{G} is nilrigid if any asymptotically nilpotent CA on $Q^{\mathbb{G}}$ is nilpotent. Despite the many results

obtained about nilrigidity in [71] they left some interesting open questions (see section 11), some of which were already asked in [68].

Question 1 ([71, 68]). *Is there any finitely generated group \mathbb{G} which is not nilrigid? In particular, is the free group with 2 or more generators nilrigid?*

4.2 Variations with partial convergence and measures

Let us now consider variations of these strong notions of nilpotency. We can first consider a similar property of orbits but restricted to a subset of initial configurations. Let us begin by two properties related to dense subsets of configurations (at least when $\mathbb{G} = \mathbb{Z}^2$).

A configuration is *q-finite* for some state q if it is everywhere q except on a finite region of \mathbb{G} . A configuration $x \in Q^{\mathbb{G}}$ is *totally periodic* if its orbit by translation $\mathcal{O}_{\mathbb{G}}(x)$ is finite. When $\mathbb{G} = \mathbb{Z}^d$ a totally periodic configuration is a configuration with d non-colinear periods.

Definition 3. • *F is nilpotent on periodic configurations if there is a configuration c such that for any totally periodic configuration x there is t with $F^t(x) = c$.*

- *F is nilpotent on finite configurations if there is a state q such that for any q -finite configuration x there is t with $F^t(x) = \bar{q}$.*

Note again that the attracting configuration c must be uniform in the first definition. Nilpotency on finite configurations is also sometimes called *eroder property* (see Section 7). An asymptotically nilpotent CA must be nilpotent on periodic configurations and nilpotent on finite configurations, but the converse is far from being true as we will see below.

As a second (seemingly unrelated) variation around nilpotency, we can change the point of view and switch from configurations to measures. A measure $\mu \in \mathcal{M}(Q^{\mathbb{G}})$ is invariant for a CA on $Q^{\mathbb{G}}$ if $F(\mu) = \mu$. The following definition is classical in dynamical system theory [78] and can be defined as a rigidity property of invariant measures.

Definition 4. *F is uniquely ergodic if it possesses a unique invariant measure.*

The invariant measures of F are exactly the limit points of Cesaro mean sequences $(u_n)_{n \in \mathbb{N}}$ of the form ([78, Theorem 6.9])

$$u_n = \frac{1}{n} \sum_{t=0}^{n-1} F^t \mu$$

for $\mu \in \mathcal{M}(Q^{\mathbb{G}})$.

Orbit-wise unique ergodicity is equivalent to the convergence of all such Cesaro mean sequences to the same limit. Therefore we could also name this property *Cesaro-nilpotency* to stress the dynamical side of it but it is best known as *unique ergodicity*. Note that by Lemma 4, any asymptotically nilpotent CA is also uniquely ergodic.

Nilpotency over periodic configurations and unique ergodicity are both implied by asymptotic nilpotency. The next lemma shows that the three properties actually form a hierarchy when $\mathbb{G} = \mathbb{Z}^d$.

Lemma 5. *A uniquely ergodic CA with $\mathbb{G} = \mathbb{Z}^d$ is nilpotent on periodic configurations.*

Proof. Consider any uniquely ergodic CA F on $Q^{\mathbb{G}}$. Suppose first that there are two disjoint temporal cycles of totally periodic configurations, *i.e.* two sequences of totally periodic configurations (c_0, \dots, c_{k-1}) and (d_0, \dots, d_{l-1}) with $F(c_i) = c_{i+1 \bmod k}$ and $F(d_j) = d_{j+1 \bmod l}$ for $0 \leq i < k$ and $0 \leq j < l$. Then we can define two measures $\mu_c = \frac{1}{k} \sum_{i=0}^{k-1} \delta_{c_i}$ and $\mu_d = \frac{1}{l} \sum_{j=0}^{l-1} \delta_{d_j}$ that are both invariant under F because (c_i) and (d_j) are cyclic orbits. Moreover since the two cycles are disjoint there must exist a cylinder $[u]$ such that $c_i \in [u]$ for some i but $d_j \notin [u]$ whatever j . Thus $\mu_c([u]) > 0$ and $\mu_d([u]) = 0$ which proves that μ_c and μ_d are two distinct invariant measures: a contradiction.

We have shown that F has the following synchronization property:

there is a cycle of totally periodic configurations (c_0, \dots, c_{k-1}) such that the orbit of any totally periodic configuration eventually enters this cycle (c_i) .

We claim that $k = 1$ which concludes the proof since it means that F is nilpotent on periodic configurations.

We are actually going to show the stronger claim that no CA F (not necessarily uniquely ergodic) with $\mathbb{G} = \mathbb{Z}^d$ can have the synchronization property above with $k \geq 2$. The argument we give is due to G. Richard (see [23] for more background on the synchronization problem). First note that if this synchronization property holds for some CA F with $\mathbb{G} = \mathbb{Z}^d$ and $d > 1$ then it must also hold for some CA with $\mathbb{G} = \mathbb{Z}$. Indeed, considering F on configurations which are constant in $d - 1$ directions and periodic in the last one, we actually define a one-dimensional CA which has the synchronization property as F does. We can thus suppose without loss of generality that $d = 1$. Clearly, each configuration c_i in the attracting cycle must be uniform because any uniform configuration has an orbit eventually entering this cycle. However, if the synchronization property holds with $k \geq 2$ then F has no quiescent state. Now consider $f : Q^{\mathbb{N}} \rightarrow Q$ the local map of F and let $w \in Q^{\mathbb{N}}$ be a semi-infinite word verifying

$$w_{n+m} = f(w_n, \dots, w_{n+m-1})$$

for all $n \in \mathbb{N}$. There must exist n_1 and n_2 with $n_1 + m < n_2$ such that

$$w_{[n_1, n_1+m-1]} = w_{[n_2, n_2+m-1]}.$$

Consider the word $w = w_{[n_1, n_2-1]}$ of length at least $m + 1$. By construction and since F has no quiescent state w contains at least two distinct letters from Q . Moreover, if c is a periodic \mathbb{Z} -configuration of period w then $F(c) = \sigma_p(c)$ for some p . Therefore the orbit of c is a cycle of non-uniform configurations, which contradicts the synchronization property. \square

Interestingly, we can also capture unique ergodicity as a property of traces of orbits on configurations: it is the property that there is some attracting state whose frequency goes to 1 when considering larger and larger prefixes of any trace. To be more precise, for any $q \in Q$ and any word $w \in Q^{\mathbb{N}}$ we denote by $d_q(w)$ the (inferior) asymptotic density of occurrences of q in w :

$$d_q(w) = \liminf_n \frac{\#\{0 \leq i < n : w_i = q\}}{n}.$$




Figure 6: Properties of any space-time diagram of a nilpotent, asymptotically nilpotent and uniquely ergodic CA respectively. White color represents the quiescent state, gray color any state and various shades of gray corresponds to various density of the quiescent state (the lighter, the higher the density). Time goes from bottom to top.

We then get this equivalent formulation of unique ergodicity [73, Proposition 3.2].

Lemma 6. *A CA F on $Q^{\mathbb{G}}$ with $\mathbb{G} = \mathbb{Z}^d$ is uniquely ergodic if and only if there is $q \in Q$ such that for any $x \in Q^{\mathbb{G}}$ it holds $d_q(\mathcal{T}_x^F) = 1$.*

Proof. Suppose first that F is uniquely ergodic, then by Lemma 5 it is nilpotent on periodic configurations and has a (unique) quiescent state q and its unique invariant measure is $\delta_{\bar{q}}$. Consider any $x \in Q^{\mathbb{G}}$ and any $\epsilon > 0$. If there were infinitely many $n \in \mathbb{N}$ such that $\frac{\#\{0 \leq i < n: \mathcal{T}_x^F(i) \neq q\}}{n} > \epsilon$, then we could extract a limit point μ from the sequence $(\frac{1}{n} \sum_{t=0}^{n-1} F^t(\delta_x))_n$ which would be an invariant measure [78, Theorem 6.9] such that $\mu([q]) < 1 - \epsilon$ (recall that $F^t(\delta_x) = \delta_{F^t(x)}$) hence different from $\delta_{\bar{q}}$: a contradiction. We deduce that $d_q(\mathcal{T}_x^F) = 1$.

Suppose now that $d_q(\mathcal{T}_x^F) = 1$ for all $x \in Q^{\mathbb{G}}$ for some $q \in Q$. Then q must be a quiescent state and thus the measure $\delta_{\bar{q}}$ is invariant and also ergodic. If it were not the unique invariant measure then there would be another invariant ergodic measure μ (see [78, Theorem 6.10]), which would necessarily verify $\mu([g \mapsto q]) < 1$ for some $g \in \mathbb{G}$ in order to differ from $\delta_{\bar{q}}$ (otherwise the Choquet's decomposition theorem on invariant measures would be contradicted). Note that μ is not necessarily translation-invariant, but translating it by g gives another F -invariant ergodic measure so we can actually suppose $g = 0_{\mathbb{G}}$ so $\mu([q]) < 1$. Then by the ergodic Theorem (see [78, Lemma 6.13]), we would have some $x \in Q^{\mathbb{G}}$ with

$$\frac{1}{n} \sum_{t=0}^{n-1} f(F^t(x)) \rightarrow_n \mu([q])$$

where f is the characteristic map of $[q]$. But then $d_q(\mathcal{T}_x^F) < 1$ which contradicts the hypothesis. \square

Uniquely ergodic CA form a strictly larger class than nilpotent CA. The example in the following result to show this separation is highly non-trivial and we will not present it here, but it is quite remarkable that such behaviors are possible with CA.

Theorem 2 ([73]). *There exists a uniquely ergodic CA on \mathbb{Z} which is not nilpotent.*

Remark 1. *The CA from the above theorem cannot be convergent because then it would be asymptotically nilpotent by Lemma 6 and this would contradict Theorem 1. From Lemma 1 it is also non-convergent when starting from some measures. However, there is convergence in Cesaro mean starting from any μ since otherwise, a Cesaro mean sequence would have several limit points and hence produce several invariant measure for the CA as already explained above.*

Theorem 3. *There exists a CA on \mathbb{Z} which is nilpotent over periodic configurations but not uniquely ergodic.*

Sketch of proof. Such a CA can be obtained using the classical construction of [42] using Wang tiles, which formalizes some tiling of the plane \mathbb{Z}^2 . Roughly, a Wang tileset is a finite set of tile types X together with vertical and horizontal constraints specifying which pairs of adjacent tile types is allowed horizontally (resp. vertically). A valid tiling is a configuration of $X^{\mathbb{Z}^2}$ that satisfies the constraints everywhere. A tileset is aperiodic if it admits valid tilings, but no periodic ones. A tileset is NE-deterministic when the constraints have the following property: when fixing a tile at the north of a position and another one at the east, then there is at most one possible tile type that can be placed at the north-east without violating the local horizontal and vertical constraints. Taking a NE-deterministic aperiodic Wang tileset T (whose existence is proven in [42]) one can transform it into a one-dimensional CA F with alphabet $T \cup \{s\}$, where s is a spreading state, and neighborhood $\{0, 1\}$ as follows: if a cell and its right neighbor hold a pair of state from T that represent two tiles (one at the north, one at the east) that can be uniquely completed to t according to the NE-deterministic rule, then the cell turns into state t , otherwise it turns into state s (configurations of the CA represent diagonals of a tiling). One obtains a CA F that admits $\delta_{\bar{s}}$ as invariant measure. Because the tileset has no valid periodic tiling, F is nilpotent on periodic configurations (the orbit of a periodic configuration without occurrence of s would give a periodic tiling of T , which is impossible, and any occurrence of s in a periodic configuration forces its orbit to converge towards \bar{s}). On the other hand, because the tileset admits a valid (aperiodic) tiling, it means that there is a configuration c whose orbit under F has no occurrence of s . Taking any limit point of the sequence of Cesaro means

$$\mu_n = \frac{1}{n} \sum_{t=0}^{n-1} F^t(\delta_c)$$

gives an invariant measure such that $\mu([s]) = 0$ since $\mu_n([s]) = 0$ for all $n \in \mathbb{N}$. We thus get a second invariant measure for F . \square

So far we considered properties concerning all initial configurations, or a meager set of initial configuration. Natural variants of nilpotency have been considered in the literature to capture the behavior on *most* or *a large set* of initial configurations. Two notions in particular emerged in literature that use different approach to define precisely *most* or *large set*: one is topological [57, 20], the other uses measure theory [45]. Recall that a set in a topological space is comeager (intuitively, large) if it is a countable intersection of sets with

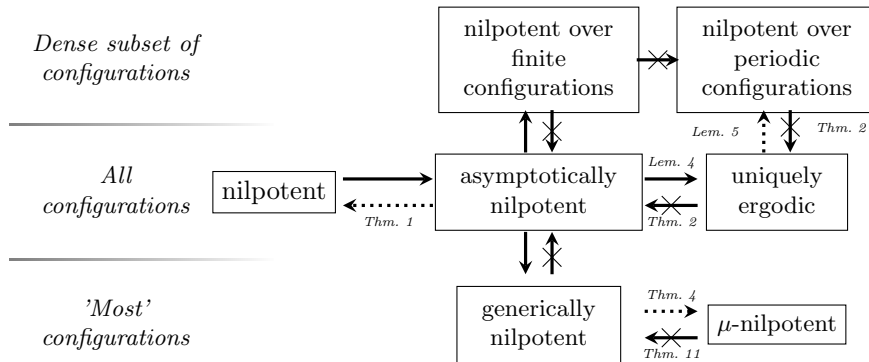


Figure 7: Nilpotency and its variants with (non-)implications between them as seen so far. Implications with dotted arrows refer to Theorems that require some hypothesis on either \mathbb{G} or μ . Counter-examples (or non-implications) are represented with a cross on the arrow.

dense topological interior (for instance a countable intersection of dense open sets).

Definition 5. Consider a CA F on $Q^{\mathbb{G}}$ and $\mu \in \mathcal{M}(Q^{\mathbb{G}})$. Then F is:

- *generically nilpotent* if there is $q \in Q$ and a comeager set $B \subseteq Q^{\mathbb{G}}$ such that $F^t(x) \rightarrow_t \bar{q}$ for all $x \in B$;
- *μ -nilpotent* if there is $q \in Q$ such that $F^t(\mu) \rightarrow_t \delta_{\bar{q}}$.

Generic nilpotency and μ -nilpotency both specify a behavior on *most configurations*, where *most* is to be understood either topologically or according to a measure. It turns out that for well-behaved measures the topological version is stronger.

Theorem 4 ([20]). Let F be a generically nilpotent CA on $Q^{\mathbb{G}}$ and $\mu \in \mathcal{M}(Q^{\mathbb{G}})$ a full-support translation-ergodic measure. Then F is μ -nilpotent.

There are ways to characterize the previous variants of nilpotency by defining variants of limit sets, *i.e.* sets of (typical) asymptotic configurations [45, 20]. Each time, the nilpotency variant is equivalent to having such asymptotic sets being singletons (see [74, 8]).

The implication diagram of Figure 7 is not complete, in particular concerning the link between unique ergodicity and μ -nilpotency.

Question 2. What are the other implications or counter-examples in the diagram of Figure 7? In particular, for what measure μ does unique ergodicity imply μ -nilpotency? Said differently, for what measure μ the convergence in Cesaro mean implied by unique ergodicity is actually a simple convergence?

4.3 The undecidability garden

We conclude this section by undecidability results related to the notions presented above. Of course, an undecidability result is a negative result and this might appear repulsive to some. However, behind almost any undecidability

result there is a positive one: a construction technique (a reduction) providing examples with rich and complex behaviors. Each of the results below relies on a non-trivial and nice construction idea that is of independent interest. A presentation of each technique would take too much space for this tutorial, but we invite the reader to dig into the corresponding references.

The following theorem follows from the undecidability of the domino problem on deterministic sets of Wang tiles. Several proof techniques are known and we suggest to read [2, Chapter 6] for a survey.

Theorem 5 ([42, 1]). *The set of 1D nilpotent CA is Σ_1^0 -complete.*

As for the existence of a uniquely ergodic CA which is not nilpotent (Theorem 2), the key tool in the following result is a self-simulating CA inspired from [25] in which the only orbits that can survive produce sparser and sparser computations.

Theorem 6 ([73]). *The set of 1D uniquely ergodic CA is Π_2^0 -complete.*

The two following results use a construction technique introduced in [18] which has proven to be useful when one wants to control the behavior of a CA on most initial configuration: it relies on a special precursor state that generates well-behaved zones that grow in a random context until they meet another well-behaved zone. Each zone can hold a properly initialized Turing computation. Then, a global process of merging or destruction of meeting zones is used to ensure that at least some Turing computations can go on forever.

Theorem 7 ([74]). *The set of 1D generically nilpotent CA is Σ_2^0 -complete.*

Recall from Theorem 4 that a generically nilpotent CA is necessarily μ -nilpotent for μ the uniform Bernoulli measure (which is ergodic and has full support). The converse is false and the following theorem illustrates how the weaker property of μ -nilpotence is strictly harder to decide.

Theorem 8 ([8]). *Let μ denote the uniform Bernoulli measure. The set of 1D μ -nilpotent CA is Π_3^0 -complete.*

5 Proving Convergence

5.1 Bounded change CA

Proving convergence of a CA F means proving that there exists a time bound for each trace \mathcal{T}_c^F after which the trace is constant. This has been defined above as the freezing time $\zeta_F(c, 0)$. In general this time bound depends on the initial configuration (see Example 3), but, even if there is a global bound that doesn't depend on initial configuration like in nilpotent CA, it is generally uncomputable. Instead of bounding directly the time of the last change in traces, a way to tackle the problem is to bound the number of changes that can occur in a trace. The following definition is useful [77, 46].

Definition 6. *A CA F is k -change for some $k \in \mathbb{N}$ if for any $x \in Q^{\mathbb{G}}$, there are at most k state changes in \mathcal{T}_x^F , i.e.*

$$\#\{t \in \mathbb{N} : \mathcal{T}_x^F(t+1) \neq \mathcal{T}_x^F(t)\} \leq k.$$

A CA is bounded-change if it is k -change for some k .

A bounded-change CA is necessarily convergent: a trace which is not ultimately constant cannot exist since it would contain infinitely many changes. However the converse is false as shown by the following example (it is used extensively as a basic building block in [62]).

Example 3. Let us consider the state set $Q = \{\square, \rightarrow, \leftarrow, \triangleleft, \triangleright, e\}$ where \triangleleft and \triangleright represent an active head moving left or right respectively, \rightarrow and \leftarrow are states indicating the position of the head (to the right and to the left respectively), \square is a blank state and e is an error state. Let Σ be the set of configurations whose patterns of length two all appear in the space-time diagram of Figure 8, i.e. configurations made of zones with a unique active head (possibly none if the zone is infinite) in a blank background. Define the CA $F : Q^{\mathbb{Z}} \rightarrow Q^{\mathbb{Z}}$ of radius 2, such that:

- e is a spreading state and any cell turns into state e if there is some pattern not in L_2 in its neighborhood,
- in the Σ -valid zones with a unique head, the only cells that change their states in one step are the one around the head, and they update in such a way that the head makes zigzag inside the zone and reduces the size of the zone by one unit at each bounce on a border. Precisely, all the transitions appear in the space-time diagram of Figure 8.

First, it is clear that F cannot be k -change whatever the value of k is, because in a Σ -valid finite zone of size n , the orbit of the central cells contains at least n changes. On the other hand, if we consider any initial configuration $c \in Q^{\mathbb{Z}}$ and any position $z \in \mathbb{Z}$ there are only 3 possible cases:

- either $c \notin \Sigma$ and in this case the spreading state e will appear somewhere and reach position z at some step,
- or $c \in \Sigma$ and z belongs to a finite valid zone, so after some number of zigzags of the head, the zone will shrink letting z outside in state \square ,
- or $c \in \Sigma$ and z belongs to an infinite valid zone, and after at most two passages of the head (at most one bounce on the boundary if any), the head will move towards infinity and cell z will remain unchanged forever.

In any case, we have that the trace at z starting from c is eventually constant. Hence F is convergent.

To use a metaphor inspired from physics, each cell of a bounded-change CA can be seen as a system that evolves according to information received from its neighbors, and for which each state change has a fixed energy cost. Then, in each orbit and at each cell, the instantaneous energy defined as the number of remaining potential state changes is non-increasing with time, and strictly decreasing at each state change. The number k in Definition 6 can be interpreted as the cell-wise capacity, i.e. the maximal energy an individual cell can hold initially.

As seen in the previous section, a nilpotent CA F is such that F^t is a constant map for some t , so it is always bounded-change. Nilpotency for CA with a spreading state is actually Turing reducible to the property of bounded-change as follows: consider F with a spreading state and add a $\{0, 1\}$ component to




Figure 8: A convergent but not bounded-change CA (time from bottom to top in the space-time diagram on the left).

states that constantly exchanges 0 and 1, except if there is a spreading state in the neighborhood in which case the additional component turns into 0. This new CA F' is bounded-change if and only if F is nilpotent: first, if F is nilpotent then F' also because the spreading state of F forces 0 everywhere on the second component; conversely, if F is not nilpotent then it must possess an orbit without any occurrence of the spreading state (by compactity), and then the corresponding orbit in F' is constantly changing the state on the second layer. By undecidability of nilpotency we deduce the following theorem.

Theorem 9. *It is undecidable to determine whether a 1D CA is bounded-change or not.*

Despite the general undecidability, this approach turns out to be very fruitful to analyze some majority CA. A majority cellular automaton consists in taking in each cell the state which has the majority of occurrences among the neighboring cells. The majority CA with neighborhood $V \subseteq \mathbb{G}$ is defined on alphabet $Q = \{0, 1\}$ by

$$F(x)_z = \begin{cases} 1 & \text{if } \sum_{i \in V} x_{z+i} > \#V/2, \text{ or } \sum_{i \in V} x_{z+i} = \#V/2 \text{ and } x_z = 1 \\ 0 & \text{otherwise.} \end{cases}$$

A subset $S \subseteq \mathbb{Z}^d$ is symmetric if $x \in S \implies -x \in S$. The following result concerning majority CA is a specific case of a much more general result from [29] which deals with infinite graphs with some growth condition. A similar qualitative behavior for majority evolution rules was established before in the settings of automata networks (the lattice of cells is finite, the evolution rule is possibly non-uniform), see [35]. It is based on a decreasing energy function in both case. Note however that [29] do not state their main theorem as a bounded change property but only as a convergence result, although their proof technique allows to establish bounded-changedness as shown below.




Figure 9: An orbit of F^2 where F is some majority CA on \mathbb{Z}^2 with symmetric neighborhood S with $(0,0) \in S$ as in the conditions of Theorem 10, so that F^2 is bounded-change and hence convergent.

Theorem 10 ([29]). *Let F be any majority CA with $\mathbb{G} = \mathbb{Z}^d$ and symmetric neighborhood S with $0 \in S$, then F^2 is bounded-change.*

Proof. Denote by r the radius of the CA, i.e. the maximum $\|u\|_\infty$ for $u \in S$. Let us fix any $x \in Q^{\mathbb{G}}$. The core of the argument is to consider the following “energy” function for any $x \in Q^{\mathbb{G}}$ and any $n, t \in \mathbb{N}$:

$$E_n^t = \sum_{u \in B_n} \sum_{v \in B_n} a(u, v) |F^{t+1}(x)_u - F^t(x)_v|$$

where $a(u, v) = 1$ if $v \in u + S$, and 0 else. Because S is symmetric $a(u, v) = a(v, u)$ so we can write E_n^{t-1} as follows:

$$E_n^{t-1} = \sum_{u \in B_n} \sum_{v \in B_n} a(u, v) |F^{t-1}(x)_u - F^t(x)_v|.$$

For any $t > 0$, we will first give an upper bound on $\Delta_n^t = E_n^t - E_n^{t-1}$ which intuitively means that the energy E_n^t essentially decreases with time, but only up to a “controlled perturbation” of the border of the ball B_n . Denoting

$$\begin{aligned} \Sigma_{n,t}^+(u) &= \sum_{v \in B_n} a(u, v) |F^{t+1}(x)_u - F^t(x)_v|, \text{ and} \\ \Sigma_{n,t}^-(u) &= \sum_{v \in B_n} a(u, v) |F^{t-1}(x)_u - F^t(x)_v| \end{aligned}$$

we can rewrite Δ_n^t as

$$\Delta_n^t = \sum_{0 \leq i \leq n-r} \sum_{u \in C_{i,t}} \Sigma_{n,t}^+(u) - \Sigma_{n,t}^-(u) + \sum_{n-r < i \leq n} \sum_{u \in C_{i,t}} \Sigma_{n,t}^+(u) - \Sigma_{n,t}^-(u) \quad (1)$$

where $C_{i,t}$ denote the cells that have different states between step $t+1$ and $t-1$ (i.e. $F^{t+1}(x)_u \neq F^{t-1}(x)_u$) and belong to the sphere of radius i (i.e. $\|u\|_\infty = i$). This expression for Δ_n^t holds because $\Sigma_{n,t}^+(u) = \Sigma_{n,t}^-(u)$ when $F^{t+1}(x)_u = F^{t-1}(x)_u$.

We first claim that $\Sigma_{n,t}^+(u) - \Sigma_{n,t}^-(u) \leq \frac{\#S-1}{2}$ for any $u \in B_n$: this holds because actually $\Sigma_{n,t}^+(u) \leq \frac{\#S-1}{2}$ must hold to ensure that a majority of neighbors of u are in state $F^{t+1}(x)_u$ at time t .

We now claim that $\Sigma_{n,t}^+(u) - \Sigma_{n,t}^-(u) \leq -1$ for any $u \in \bigcup_{0 \leq i \leq n-r} C_{i,t}$. Indeed, we can split the neighbors of any such u into two sets

$$S_u^+ = \{v \in u + S : F^t(x)_v = F^{t+1}(x)_u\}, \text{ and}$$

$$S_u^- = \{v \in u + S : F^t(x)_v = F^{t-1}(x)_u\},$$

and then $\#S_u^+ = \Sigma_{n,t}^-(u)$ because $u + S \subseteq B_n$, and similarly $\#S_u^- = \Sigma_{n,t}^+(u)$. By the majority rule we must have $\#S_u^+ > \#S_u^-$ (recall that S is of odd cardinality by hypothesis) which proves the claim.

Let $s = \frac{(\#S-1)}{2}$. From the two claims we can rewrite Equation (1) above as:

$$\Delta_n^t \leq -\# \bigcup_{0 \leq i \leq n-r} C_{i,t} + s\# \bigcup_{n-r < i \leq n} C_{i,t}. \quad (2)$$

We are now going to deduce a bound on the number of changes in $\mathcal{T}_x^{F^2}$ from the above inequality in two steps: using a weighted sum spatially and then summing over time. Spatially, the trick is to use a telescoping sum to attenuate the perturbation of the second term in the upper bound of Equation (2) and put more weight on the central cell. Choose $n = kr$ and consider the sum

$$TE_n^t = E_n^t + \frac{1}{s}E_{n-r}^t + \left(1 + \frac{1}{s}\right)\frac{1}{s}E_{n-2r}^t + \left(1 + \frac{1}{s}\right)^2\frac{1}{s}E_{n-3r}^t + \cdots + \left(1 + \frac{1}{s}\right)^{k-1}\frac{1}{s}E_0^t.$$

The weights are chosen so that $TE_n^t - TE_n^{t-1} = \Delta_n^t + \frac{1}{s}\Delta_{n-r}^t + \cdots$ can be nicely upper-bounded. Indeed Equation (2) gives

$$\Delta_n^t + \frac{1}{s}\Delta_{n-r}^t \leq -\left(1 + \frac{1}{s}\right)\# \bigcup_{0 \leq i \leq n-2r} C_{i,t} + s\# \bigcup_{n-r < i \leq n} C_{i,t}$$

so the changes in the annulus between spheres of radii $n-2r$ and $n-r$ are exactly compensated. Then, when adding $\left(1 + \frac{1}{s}\right)\frac{1}{s}\Delta_{n-2r}^t$ upper-bounded by Equation (2), the terms involving annulus from $n-3r$ to $n-2r$ are exactly compensated again, giving:

$$\Delta_n^t + \frac{1}{s}\Delta_{n-r}^t + \left(1 + \frac{1}{s}\right)\frac{1}{s}\Delta_{n-2r}^t \leq -\left(1 + \frac{1}{s}\right)^2\# \bigcup_{0 \leq i \leq n-3r} C_{i,t} + s\# \bigcup_{n-r < i \leq n} C_{i,t}.$$

Repeating this argument, and by the identity

$$1 + \frac{1}{s} + \left(1 + \frac{1}{s}\right)\frac{1}{s} + \cdots + \left(1 + \frac{1}{s}\right)^{k-1}\frac{1}{s} = \left(1 + \frac{1}{s}\right)^k,$$

we get

$$TE_n^t - TE_n^{t-1} \leq -\left(1 + \frac{1}{s}\right)^k C_{0,t} + s\# \bigcup_{n-r < i \leq n} C_{i,t} \leq -\left(1 + \frac{1}{s}\right)^k C_{0,t} + s\#(B_n \setminus B_{n-r}).$$

By summing over time the successive differences we then obtain

$$TE_n^t - TE_n^0 \leq -\left(1 + \frac{1}{s}\right)^k \sum_{i < t} C_{0,i} + (t-1)s\#(B_n \setminus B_{n-r})$$

so

$$\sum_{0 \leq i < t} C_{0,i} \leq \frac{TE_n^0}{\left(1 + \frac{1}{s}\right)^k} + \frac{(t-1)s\#(B_n \setminus B_{n-r})}{\left(1 + \frac{1}{s}\right)^k}.$$

When n grows, the first term of this upper bound converges to a constant independent of t and of x (from the definition of TE_n^0 because E_n^0 is at most the square of the size of B_n), and the second term vanishes. This concludes the proof because it shows that $\sum_{i=0}^{\infty} C_{0,i}$ is bounded independently of x , and it counts the number of state changes in the trace $\mathcal{T}_x^{F^2}$ and the trace $\mathcal{T}_{F(x)}^{F^2}$. \square

Let us show two counter-examples to Theorem 10 when hypothesis on either \mathbb{G} or the symmetry of neighborhood are removed:

1. let \mathbb{G} be the free group with 2 generators, V the symmetric neighborhood made of all generators and their inverse and element $0_{\mathbb{G}}$, and denote by F the associated majority CA. Define the configuration c^b for any $b : \mathbb{N} \rightarrow \{0, 1\}$ by

$$c_g^b = b(\|g\|_{\mathbb{G}}).$$

Because we are in the free group with 2 generators, if $g \in B_n \setminus B_{n-1}$ then $g + V$ contains 3 elements in $B_{n+1} \setminus B_n$ and V is of size 5. Therefore we have $F(c^b)_g = b(\|g\|_{\mathbb{G}} + 1)$, and thus $F(c^b) = c^{b'}$ with $b'(n) = b(n+1)$ for all $n \in \mathbb{N}$. By choosing b such that $b(i+1) \cdots b(i+2k) = 0^k 1^k$ for arbitrarily large k and for some i , we deduce that there is no $p > 0$ such that F^p is convergent.

2. for $\mathbb{G} = \mathbb{Z}^2$, let us consider the non-symmetric neighborhood $V = \{(0, 0), (0, 1), (1, 0)\}$ and let F be the associated majority CA. Define the configuration c^b for any $b : \mathbb{N} \rightarrow \{0, 1\}$ by

$$c_{(x,y)}^b = b(x+y).$$

Again, for any $(x, y) \in \mathbb{Z}^2$ there are two elements z' in $(x, y) + V$ such that $c_{z'}^b = b(x+y+1)$, so $F(c^b)_{(x,y)} = b(x+y+1)$. We can conclude as above that there is no $p > 0$ such that F^p is convergent.

5.2 Freezing CA

The argument of Theorem 10 to prove the bounded-change property is not immediate and the property is generally undecidable, even when $\mathbb{G} = \mathbb{Z}$ (Theorem 9).

In some cases, however, the structure of the local rule of a CA directly implies the bounded-change property. The so called *freezing CA*, introduced in [30], are an example.

Definition 7. A CA F on $Q^{\mathbb{G}}$ is freezing if there is some partial order (Q, \leq) on its state set such that for any $c \in Q^{\mathbb{G}}$ and any $z \in \mathbb{G}$ it holds:

$$F(c)_z \leq c_z.$$

Clearly, in a freezing CA the number of changes at any cell in any orbit is bounded by the depth of the partial order, so all freezing CA are bounded-change. The freezing property can be tested efficiently by looking at the local transition table. Indeed, given any CA F , we can compute the canonical state change graph (Q, \rightarrow) defined by $q_1 \rightarrow q_2$ if some transition changes state q_1 into q_2 . Then one can check that F is freezing if and only if this graph is acyclic.

Going back to our earlier metaphor, the partial order involved in the freezing property can be seen as a local non-increasing energy that serves as a certificate for the bounded change property. Of course, not all bounded-change CA are freezing: for instance F^2 , where F is some CA from the hypothesis of Theorem 10, is a bounded-change CA but it cannot be freezing since F and hence F^2 commute with the transformation that permutes states 0 and 1. Thus, if transition $0 \rightarrow 1$ exists in the canonical state change graph of F^2 , then transition $1 \rightarrow 0$ also exists, which prevent any freezing order to satisfy Definition 7.

There are many example of freezing CA studied in the literature [37, 75, 24, 4, 36, 7].

A freezing CA F for some order \leq can also be *monotone*:

$$x \leq y \Rightarrow F(x) \leq F(y)$$

where \leq here denotes the cellwise extension of \leq to configurations. Among these examples, the classical bootstrap percolation CA [14] has been the starting point of a rich branch of percolation theory, which can be viewed as the study of the qualitative behavior of monotone freezing CA initialized on random Bernoulli configurations [7, 40, 5]. This CA, denoted F_B in the sequel, is defined on $\{0, 1\}^{\mathbb{Z}^2}$ as follows:

$$F_B(c)_z = \begin{cases} 0 & \text{if } c_z = 0 \text{ or } \#\{z' \in z + N : c_{z'} = 0\} \geq 2, \\ 1 & \text{else.} \end{cases}$$

where $N = \{(1, 0), (0, 1), (-1, 0), (0, -1)\}$.

F_B is not nilpotent on periodic configurations (it admits both $\bar{0}$ and $\bar{1}$ as fixed points), and a classical result of [76] shows that it is μ -nilpotent for any full-support Bernoulli measure². Actually, much more is known on this CA started from random initial configurations [40]. A qualitatively very different behavior is obtained by a similar CA, denoted F_O , and defined on $\{0, 1\}^{\mathbb{Z}^2}$ as follows:

$$F_O(c)_z = \begin{cases} 0 & \text{if } c_z = 0 \text{ or } c_{z+(1,0)} = 0 \text{ or } c_{z+(1,1)} = 0, \\ 1 & \text{else.} \end{cases}$$

F_O , contrary to F_B , is μ -nilpotent for some Bernoulli measures but not all.

As a third example of monotone freezing CA with 2 states, let's consider F_M defined as follows:

$$F_M(c)_z = \begin{cases} 0 & \text{if } c_z = 0 \text{ or } \#\{z' \in z + N : c_{z'} = 0\} \geq 3, \\ 1 & \text{else.} \end{cases}$$

A classification of behaviors of all monotone freezing CA with 2 states started from random Bernoulli configurations is given in [5, 7]. The following theorem focuses on examples F_B and F_O to illustrate this classification.

²[76] doesn't use the notion of μ -nilpotence but the property $\mu(\{c : F^\omega(c) = \bar{0}\}) = 1$. It can be shown that the two are equivalent for freezing CA in general [70, Lemma 3.6].




Figure 10: Representation of fixed points reached under either F_B or F_O or F_M starting from μ -random initial configurations with Bernoulli measures μ . The color interpretation is as follows: red represents state 1 and shades of blue represent state 0 with the convention that the lighter the color, the earlier the freezing time. This illustrates Theorem 11: F_B converges almost surely to $\bar{0}$ no matter how small $\mu([0])$ is, while F_O needs a large enough $\mu([0])$ to erase almost surely all 1s, and F_M almost never converges to $\bar{0}$, no matter how large $\mu([0])$ is.

Theorem 11. F_B is generically nilpotent, and μ -nilpotent for any translation-ergodic measure μ of full support. However, there is a threshold $0 < \theta < 1$ such that F_O is μ -nilpotent for Bernoulli measure μ when $\mu([0]) > \theta$ and not μ -nilpotent if $\mu([0]) < \theta$. In particular F_O is not generically nilpotent. Finally, F_M is not μ -nilpotent as soon as μ has full support.

Proof. It is enough to show generic nilpotency of F_B and the μ -nilpotency follows by Theorem 4. For any $n \in \mathbb{N}$ consider the set of configurations having an annulus of 0 of radius n around position $(0, 0)$:

$$A_n = \{c \in \{0, 1\}^{\mathbb{Z}^2} : \forall z, \|z\| = n \implies c_z = 0\}.$$

It is easy to show by induction that if $\|z\| \leq n$ and $c \in A_n$ then $F_B^\omega(c)_z = 0$ (the annulus of radius $n - 1$ is first turned into 0 starting from the corner, and so on). A_n is open (it's a union of cylinders) and for any $m \in \mathbb{N}$ the union $U_m = \bigcup_{n \geq m} A_n$ is open and dense (any cylinder intersects A_n for large enough n). Therefore the intersection $I = \bigcap_{m \geq 0} U_m$ is a comeager set. From the remark above, it holds $F_B^\omega(c)_z = 0$ for any $z \in \mathbb{Z}^2$ and any $c \in I$. Therefore I is included in the realm of attraction of $\bar{0}$ and we deduce that F_B is generically nilpotent.

Concerning F_O , the claim on Bernoulli measures for which it is μ -nilpotent is enough to show that it is not generically nilpotent by Theorem 4. One can show that $F_O^\omega(c)_{(0,0)} = 1$ if and only if there is a directed path $(\pi_n)_{n \in \mathbb{N}}$ in \mathbb{N}^2 starting from $\pi_0 = (0, 0)$ and such that $\pi_{n+1} \in \pi_n + \{(0, 1), (1, 1)\}$ for all $n \in \mathbb{N}$, and verifying $c(\pi_n) = 1$ for all $n \in \mathbb{N}$. Indeed, if such path exists, it is easy to check that $F^t(c)_{\pi_n} = 1$ for all $t \geq 1$ and all $n \in \mathbb{N}$ by induction, so in particular $F_O^\omega(c)_{(0,0)} = 1$. On the contrary, if no such infinite path exists, then there is a finite bound B on the longest finite directed path with the same properties (otherwise we could extract an infinite one by compactity). One can then check that $F_O^{B+1}(c)_{(0,0)} = 0$. It is well-known from oriented percolation theory [21] that when c is a μ -random configuration for Bernoulli μ , there is a threshold $0 < \theta < 1$ such that an infinite path $(\pi_n)_{n \in \mathbb{N}}$ exists:

- with positive probability when $\mu([0]) < \theta$,
- with probability 0 when $\mu([0]) > \theta$.

We deduce that $F_O^t(\mu)([1]) \rightarrow 0$ so F_O is μ -nilpotent in the second case, while $F_O^t(\mu)([1]) \geq \epsilon > 0$ for all $t \geq 0$ so F_O is not μ -nilpotent in the first case.

Finally, F_M admits a finite obstacle: precisely, if a 2 by 2 square of 1s appear in some configuration, it will never disappear. Thus if μ gives a probability larger than $\epsilon > 0$ to a cylinder made of this pattern and also to the cylinder $[0]$, then so will $F^t(\mu)$ for any $t \geq 1$. We deduce that F_M is not μ -nilpotent. \square

Although μ -nilpotency is well understood for monotone freezing CA with 2 states by bootstrap percolation theory [61, 60], it remains hard for freezing CA in general.

Theorem 12 ([70]). *The set of 2D freezing CA with 2 states which are μ -nilpotent for all full-support Bernoulli measure μ is recursively inseparable from the set of 2D freezing CA with 2 states which are μ -nilpotent for no full-support Bernoulli measure μ .*

A measure μ is said limit-computable if there is a computable map ϕ such that $\lim_{n \rightarrow \infty} \phi(u, n) - \mu([u]) = 0$ for each cylinder $[u]$. One can check that if $\mu' = \lim_{t \rightarrow \infty} F^t(\mu)$ for some CA F and μ a computable measure, then μ' is limit computable. There is a remarkable converse to this observation.

Theorem 13 ([19, 54]). *Let $\mathbb{G} = \mathbb{Z}^d$ for some $d \geq 1$. Let μ be any translation-invariant and limit-computable measure on some alphabet. Then there exists a CA F (on a possibly different alphabet) such that $\mu = \lim_{t \rightarrow \infty} F^t(\mu_0)$ where μ_0 is the uniform Bernoulli measure on the alphabet of F .*

Interestingly, the construction in the above theorem relies on CA that are not convergent, although the particular orbit starting from μ_0 (and actually many other initial measures) is itself convergent. We have no idea about whether it is possible to realize the same limit measures with convergent CA.

Question 3. *What are the limit measures $F^\omega(\mu_0)$ where F is convergent (or bounded change, or freezing) and μ_0 is the uniform Bernoulli measure (or some computable measure)?*

6 Computational Power of Convergent Cellular Automata

The convergence property dramatically restricts the possible dynamics of CA as we have seen so far. One can therefore legitimately ask whether the obvious fact that general CA are computationally universal still holds for convergent, bounded-change and freezing CA. This section tackles this question and outlines an answer in three steps: yes, even freezing CA in dimension 1 are computationally universal, but there is some loss of complexity, and for a convergent CA F , it is interesting to study the map $c \mapsto F^\omega(c)$ from a computational point of view.




Figure 11: On the left, a Minsky machine with 1 counter and a few steps of the execution starting from counter value 0. On the right the corresponding space-time diagram of the freezing CA of Example 5 encoding it. Counter value is encoded as the blue part in each trace: value n is represented by state sequence $1^n \#_{-1} \#_0 \#_1$. Green cells in the space-time diagram indicate position where all the required information to compute a Minsky transition is available locally. Cells containing the result of Minsky transitions are represented in red.

6.1 Embeddings of computationally universal systems

Let's start by the obvious observation that the space-time diagrams of any 1D CA can be grown by a 2D freezing CA.

Example 4. Any 1D CA F with states Q and neighborhood V can be simulated by a 2D freezing CA F' with states $Q \cup \{*\}$ as follow. Let $V' = \{(v, -1) : v \in V\}$. A cell in a state from Q never changes. A cell in state $*$ looks at cells in its V' neighborhood: if they are all in a state from Q then it updates to the state given by applying F on them, otherwise it stays in $*$. Starting from a all- $*$ configuration except on one horizontal line where it is in a Q -configuration c_0 , this 2D freezing CA will compute step by step the space-time diagram of F on configuration c_0 .

In dimension 1, it is also possible to embed any Turing machine computation (with a polynomial slowdown) inside a convergent CA using the zigzag trick of Example 3: by adding more states, one can make the head actually compute like a Turing head and do one transition at each zigzag (see [62, Proposition 5] for details).

For 1D freezing CA, the embedding of computationally universal systems is much more constrained but still possible using Minsky machine [30, 12, 62].

Definition 8. A k -counter Minsky machine is a 4-tuple $M = (Q_M, q_0, h, \tau)$ where $q_0, h \in Q_M$ are the initial and halting states and

$$\tau : Q_M \times \{0, 1\}^k \rightarrow Q_M \times \{-1, 0, 1\}^k$$

is its transition map, which verifies $\tau(h, \cdot) = (h, (0, \dots, 0))$. A configuration of M is an element of $Q_M \times \mathbb{N}^k$. M transforms any configuration $c = (q, (\chi_i)_{1 \leq i \leq k})$ in one step into configuration

$$M(c) = (q', (\max(0, \chi_i + \delta_i))_{1 \leq i \leq k})$$

where $(q', (\delta_i)_{1 \leq i \leq k}) = \tau(q, (\min(1, \chi_i))_{1 \leq i \leq k})$. M halts on input $(\chi_i)_{1 \leq i \leq k} \in \mathbb{N}^k$ if there is a time t such that $M^t(q_0, (\chi_i)_{1 \leq i \leq k}) \in (h, \mathbb{N}^k)$.

Example 5. Any k -counter Minsky machine M can be embedded into a 1D freezing CA F in the following sense. To simplify the exposition, let's take $k = 1$. A configuration of M at some step of the evolution is a pair (q, n) where q is a finite state and $n \in \mathbb{N}$ the current value of the counter. This configuration will be encoded into the trace of some cell $z \in \mathbb{Z}$ of F : the state q will appear at cell z at some time, and after the value n will be encoded by the time interval between specific state changes occurring at z . Successive configurations of an orbit of M are then encoded into successive cells of F through the temporal traces. In order to allow F to correctly produce an orbit where the trace at position $z + 1$ encodes the configuration which is the image by M of the configuration encoded into the trace at position z , the key point is to shift temporally the time interval containing the encoded configuration from position to position $z + 1$. This technical condition allows to implement all the operations on a counter locally (zero test, increment and decrement). Concretely, if the configuration of M at time 0 is encoded into the time interval $[0, \Delta_0]$ at cell 0, then the configuration at step t of machine M is encoded into the time interval $[2t, 2t + \Delta_t]$ at cell t . See Figure 11 for an example with details.

From the embedding of Minsky machines as in the above example, one expects undecidability results to be transferred to 1D freezing CA. There are many ways to state such results, depending on the additional details implemented. Let us mention the following one which essentially relies on a Minsky machine embedding but requires a more technical construction than the above example. A freezing CA is always k -change for k no smaller than the cardinal of its alphabet, but determining the minimal k for which it is k -change is undecidable.

Theorem 14 ([62, Theorem 6]). *There exists a constant k such that the following problem is undecidable: given a freezing 1D CA F , decide whether F is k -change.*

The intuition for the construction in the proof of the above theorem is as follows:

- first, the construction is such that the largest sequence of state changes that can possibly occur in some orbit at a given cell is unique, no other sequence of state changes can be as long as this one;
- second, if this sequence appears in some orbit, it enforces the presence of a correct halting computation of the Minsky machine; conversely, if the simulated Minsky machine halts, there exists an orbit where this special sequence of state changes occurs at some cell;
- third, the length of this sequence, denoted L , is independent of the number of states of the Minsky machine being simulated (intuitively, only the initial and halting states of the machine appear in it).

With these ingredients, one has a reduction from the halting problem of Minsky machines to the problem of whether a freezing CA is $(L - 1)$ -change.

6.2 Dimension 1: complexity gap between bounded change and convergent CA

There is a strong information flow bottleneck in bounded change CA that becomes critical in 1D. It is not the case for convergent CA. This can be formulated using communication complexity or classical computational complexity of the prediction problem [77, 30, 62] (see also [33] which generalizes the result to other \mathbb{G}). The prediction problem associated to a CA asks for the state of a cell after a given time starting from a given input, a possible formalization is as follows: given an input pattern $u \in Q^{Btr}$ where r is the radius of the considered CA, the problem is to determine the value of cell $0_{\mathbb{G}}$ after t steps starting from a configuration $c \in [u]$ (note that by choice of the domain of u , this does not depend on c). The prediction problem has communication complexity at most $\log(n)$ and is in class NL for 1D bounded change CA, while it can be $\Omega(\sqrt{n})$ and P-complete for 1D convergent CA.

Instead of presenting the above results and their formal setting, we are going to illustrate this difference between bounded change and convergent CA through the problem of recognizing palindromes. Before stating the result, let's formalize the notion of language recognition by CA under time constraints (but without space constraint). Say a language $L \subseteq \{0, 1\}^*$ is recognized by a CA F on $Q^{\mathbb{Z}}$ in time $\tau : \mathbb{N} \rightarrow \mathbb{N}$ if $\{0, 1, B, A, R\} \subseteq Q$ (where B is a blank state, A an accepting state and R a rejecting state) and for any $n \in \mathbb{N}$ and any $u \in \{0, 1\}^n$, the orbit of configuration c^u defined by

$$c^u(z) = \begin{cases} u_z & \text{if } 0 \leq z < n, \\ B & \text{else,} \end{cases}$$

is such that:

- there is $t \leq \tau(n)$ such that $F^t(c^u)_0 \in \{A, R\}$,
- for the minimal such time t it holds

$$F^t(c^u)_0 = A \iff u \in L.$$

Finally, PAL is the language of palindromes on alphabet $\{0, 1\}$, *i.e.* words u such that $u_i = u_{n-1-i}$ for $0 \leq i < n = |u|$. The next theorem states that PAL is hard to recognize with bounded change CA, while it is easily recognized by convergent CA. Note that limitations of bounded change (or bounded communication) CA were established previously in different settings: in [77] language recognition is considered with a space limitation (only the cells initially containing the input can be modified), and in [46] the focus is on real-time recognition (*i.e.* $\tau(n) = n$).

Theorem 15. *Let $\mathbb{G} = \mathbb{Z}$. Suppose that a k -change CA can recognize PAL in time $\tau(n)$, then the recognition time is at least exponential, *i.e.* $\tau(n) \geq \alpha^n$ for some $\alpha > 1$ and large enough n . However, there exists a convergent CA that recognizes PAL in quadratic time.*

Proof. Suppose by contradiction that a k -change CA F on $Q^{\mathbb{Z}}$ and of radius r recognizes PAL in time $\tau(n) \leq \alpha^n$ with $\alpha < 2^{\frac{1}{4rk}}$. For $n \in \mathbb{N}$ and $u \in \{0, 1\}^n$ denote by $T_z^{u,n}$ the prefix of the trace of length $\tau(n)$ and of width $2r$ starting at position z in the orbit of c^u :

$$T_z^{u,n} : t \in \{0, \dots, \tau(n)\} \mapsto F^t(c^u)_z, \dots, F^t(c^u)_{z+2r}.$$

By the k -change property, there are at most $(|Q|\tau(n))^{2rk} \leq |Q|^{2rk} \alpha^{2rkn}$ such prefixes (for all choices of u and z) because for each of the $2r$ columns of the trace, there are at most k changes and each change can be described by a new state and a time step between 1 and $\tau(n)$. On the other hand, there are $2^{\frac{n}{2}}$ words in PAL of even length n , so for large enough even n and by choice of α there are $u \neq v$ in PAL such that

$$T_{n/2}^{u,n} = T_{n/2}^{v,n} \quad (3)$$

Then, consider the word $w = u_0 \cdots u_{n/2-1} v_{n/2} \cdots v_{n-1}$. By the Equality (3) above, we have that $\mathcal{T}_{c^u}^F(t) = \mathcal{T}_{c^w}^F(t)$ for all $0 \leq t \leq \tau(n)$. So in particular F must accept w in time at most $\tau(n)$ because it does for u . However, w is not a palindrome, which yields the desired contradiction.

It is not difficult to adapt Example 3 to make a convergent CA that recognizes PAL in quadratic time. Roughly, using extra states, any input word BuB is turned in one step into a segment of state with two components: the binary component that keeps the information of u and the zigzag component of the form $\boxed{B} \boxed{\triangleright} \boxed{\leftarrow} \cdots \boxed{\leftarrow} \boxed{B}$, where the head state $\boxed{\triangleright}$ holds an extra bit, initialized to u_0 . Then, with the zigzag movement of the head, this bit held inside the right-moving head $\boxed{\triangleright}$ is compared to the bit in the binary component when the head bounces on the right border (*i.e.* u_0 is compared to u_{n-1} at the end of the first zig). If the comparison fails, an error state e appears which is interpreted as a rejecting state. If not, the head turns into $\boxed{\leftarrow}$ (which does not hold any additional information) and goes back to the left boundary as in the rule of Example 3. Once the left boundary is reached again, the head becomes $\boxed{\triangleright}$ and copies the bit from the binary component at this position (*i.e.* u_1 at the start of the second zig). When the zone is shrunk down to a single position, an accepting state is generated that spreads to the left to reach the position which initially was holding value u_0 (the accepting state spreads over any other state except e , which is not present in the orbit of c^u if u is in PAL). The additional mechanism on top of the rule of Example 3 does not compromise the convergent property: the analysis is the same as in Example 3 with the presence of two spreading states instead of one (rejecting state e spreading over the accepting state which spreads over any other state). \square

Note that the above theorem does not say anything about the existence of a bounded-change CA recognizing PAL in super-exponential time. It would be surprising that such a CA exists because the information to compare between the two halves of a palindrome is too dense to be dealt with and transmitted within the bounded change constraint, however we have no proof of this impossibility.

Question 4. *Is there a bounded change CA recognizing PAL (with arbitrary time and space) or any language with linear communication complexity in the sense of [10]?*

We don't think that the quadratic time recognition of PAL by convergent CA is optimal in Theorem 15, however we have no idea on how to recognize PAL in real-time (*i.e.* $\tau(n) = n$) with a convergent CA. More generally, we don't know whether the recognition time of languages by convergent CA can always be as good as the recognition time for general CA.

6.3 Complexity of limit configurations

Any convergent CA F on $Q^{\mathbb{G}}$ induces a map $F^{\omega} : Q^{\mathbb{G}} \rightarrow Q^{\mathbb{G}}$. We say that a configuration $c \in Q^{\mathbb{G}}$ is computable if it is computable as a map. For any t , the map F^t transforms computable configurations into computable configurations, however there is no reason to expect that the same holds for F^{ω} .

Let us first make the following observation: the computability of the limit configuration is linked to the computability of freezing times.

Lemma 7. *Let F be a convergent CA over $Q^{\mathbb{G}}$ and $c \in Q^{\mathbb{G}}$ be any computable configuration. Then the map $z \mapsto F^{\omega}(c)_z$ is Turing reducible to the map $z \mapsto \zeta_F(c, z)$. Moreover, if F is freezing, both maps are Turing-equivalent.*

Proof. To compute $F^{\omega}(c)_z$ from $\zeta_F(c, z)$, it is sufficient to compute $F^{\zeta_F(c, z)}(c)_z$. Reciprocally, if F is supposed freezing, $\zeta_F(c, z)$ can be computed from $F^{\omega}(c)_z$ as it is the first time step t such that $F^t(c)_z = F^{\omega}(c)_z$. \square

For bounded change 1D CA, it is impossible to produce uncomputable limits from computable initial configurations. The argument of the following proof is due to G. Richard.

Theorem 16 ([62]). *For any 1D bounded-change CA F and any computable configuration c , $F^{\omega}(c)$ is computable.*

sketch. We restrict without loss of generality to F of radius 1 and fix a computable configuration c . If we know the number of state changes in the orbit of c at cells z_1 and z_2 , with $z_1 < z_2$, then we can compute the value $F^{\omega}(c)_z$ for all $z \in [z_1, z_2]$. Indeed, we can compute $F^t(c)$ on cells $[z_1, z_2]$ for increasing values of t until the correct number of changes is observed at z_1 and z_2 . From that time on, the evolution of cells in the segment $[z_1, z_2]$ is independent of the context since cells z_1 and z_2 no longer change and F has radius 1. So we can compute the evolution until this set of cells reaches a fixed point, which is then the value they have in $F^{\omega}(c)$.

The number of changes at any cell is bounded, and there is a maximal value L for which there are infinitely many positions $z < 0$ with exactly L changes. Moreover, there is a limit position z_L to the left of which no cell has more than L changes. The same is true for positions $z > 0$ giving the corresponding constants R and z_R .

Then the algorithm to compute $F^{\omega}(c)_z$ given z is the following: compute larger and larger portions of the space-time diagram around position z until finding $z_1 \leq z \leq z_2$:

1. $z_1 \leq z_L$ and $z_R \leq z_2$,
2. the state of z_1 has changed L times and the state of z_2 has changed R times.




Figure 12: Examples of configurations obtained after some time from a finite seed by randomly chosen 2D freezing CA. All CA have the same neighborhood and the same number of states, with the same freezing order. All finite seed have the same size (the non-uniform portion of the configuration is a small centered square), and all example configurations were obtained after the same number of steps. The difference in the uniform background comes from the fact that not all examples have the same quiescent states.

Then it is sufficient to apply the algorithm of the above to compute $F^\omega(c)_{[z_1, z_2]}$ and therefore obtain $F^\omega(c)_z$. \square

Interestingly, the above limitation for bounded change CA disappears when considering convergent CA.

Theorem 17 ([62]). *There exists a 1D convergent CA F and a computable configuration c such that $F^\omega(c)$ is uncomputable.*

In 2D, it is easy to find a freezing CA F and a computable configuration c such that $F^\omega(c)$ is uncomputable: let $c(i, j)$ be 1 if $i, j \in \mathbb{N}$ and Turing machine i halts in at most j steps on the empty tape, and 0 otherwise. Then, if F is the freezing CA that spreads state 1 vertically, we get that $F^\omega(c)_{(i, 0)} = 1$ for $i \in \mathbb{N}$ if and only if machine i halts.

A more interesting question is to ask whether $F^\omega(c)$ can be uncomputable for some finite or eventually periodic configuration c (*i.e.* a configuration which is totally periodic up to a finite region of \mathbb{Z}^2). The behavior of freezing CA from finite configuration is quite rich (see Figure 12) and it is actually possible to realize uncomputable limits with a freezing 1-change CA and slightly more with a 2-change CA.

Theorem 18. [47, 62] *For a configuration $c \in Q^{\mathbb{Z}^2}$ and $q \in Q$, define the set of cells of c in state q : $\chi_q(c) = \{z \in \mathbb{Z}^2 : c(z) = q\}$. The following holds:*

- *There exists a 1-change freezing CA F_1 , a state q of F_1 , and a finite configuration c^1 such that $\chi_q(F_1^\omega(c^1))$ is not computable.*




Figure 13: Configuration reached from a random finite initial configuration (center) by the CA game of life without death after a few hundreds of steps. Black represents state 0, white represents state 1.

- For any 1-change freezing CA F , any state q and any finite configuration c , $\chi_q(F^\omega(c))$ is either recursively enumerable or co-recursively enumerable.
- There exists a 2-change freezing CA F_2 , a state q of F_2 , and a finite configuration c^2 such that $\chi_q(F_2^\omega(c^2))$ is neither recursively enumerable, nor co-recursively enumerable.

Proof. The first item is the main result of [47]. The second item is [62, Proposition 9]. The third item is [62, Corollary 1]. \square

Recall from Lemma 7 that in order to produce an uncomputable limit configuration from a computable one by a freezing CA, the freezing times of cells must be uncomputable. We stress that for this reason it is not sufficient for a 2D freezing CA to be *computationally universal* in order to produce uncomputable limit configurations from eventually periodic initial ones. For instance, in the orbits described in the simple embedding of Example 4, the freezing times are computable by construction since cells are frozen progressively line by line. We don't have a general understanding of the ingredients behind the capacity of freezing CA to produce uncomputable limit configuration starting from simple (computable) initial ones.

Concretely, one can consider the *game of life without death* CA [37], which

is a CA on the space on configurations $\{0,1\}^{\mathbb{Z}^2}$ defined by³:

$$F(c)_z = \begin{cases} 0 & \text{if } c_z = 0 \text{ or } Z_8(c, z) = 3, \\ 1 & \text{else,} \end{cases}$$

where $Z_8(c, z)$ counts the number of occurrences of state 0 among the 8 neighbors of cell z in configuration c . The prediction problem of this CA was shown to be P-complete in [37], however its ability to produce complex limit configurations from simple initial ones has not been established to our knowledge (see Figure 13).

Question 5. *Let F be the game of life without death CA. Is there a finite or eventually periodic configuration c such that $F^\omega(c)$ is uncomputable? More generally, for which freezing CA is this possible? What are sufficient or necessary conditions to achieve this behavior?*

7 Some related works and further reading

What was presented so far is of course not an exhaustive view of cold dynamics in cellular automata, and it also does not tackle other dynamical systems. Let us finish this tutorial by pointing out some related research directions that does not exactly fit in the formalism or classes seen so far.

Randomization and measure rigidity in CA. There have been much work on CA considered as dynamical systems acting on probability measures [64]. Let us mention two (linked) aspects that extends what we presented above: recall that unique ergodicity (Definition 4) is both the property of having a unique invariant measure, and the property that any initial measure converges in Cesaro mean towards this unique invariant measure under the CA. If one considers a surjective CA on $\mathbb{G} = \mathbb{Z}^d$, then the uniform Bernoulli measure is an invariant measure [6]. It can never be the only one (a CA always admits some linear combination of dirac measures as an invariant measure), but a natural question is to ask under which additional conditions it becomes unique, a phenomenon called *measure rigidity* (see [64, section 2] and references therein). Note that the famous Furstenberg conjecture can be formulated as measure rigidity question on a particular CA [64, section 2F]. Another much studied phenomenon is called *randomization* and studied since the 80s [58, 49]: it is the property that a large class of initial measures (including Bernoulli measures of full support) converge in Cesaro mean towards the uniform Bernoulli measure (see [64, section 3A] and references therein). More recently, a characterization of randomization was obtained for Abelian CA and an example was shown where the convergence is not only in Cesaro mean but actually a direct convergence (in the topology on measures) [53]. Recall that this difference in the kind of convergence plays also a role in the links between μ -nilpotency and unique ergodicity as discussed in Section 4. These two phenomena are now rather well understood on CA with an algebraic structure, but the generalization to other kind of CA remains largely open.

³For the reader familiar with life-like CA, note that we use the somewhat unconventional representation of alive cells by 0 to stay consistent with Definition 7.

Probabilistic CA, noise and ergodicity. A probabilistic CA is a CA whose local evolution rule is probabilistic, thus randomness can be introduced at each time step (contrary to deterministic CA acting on probability measures where randomness only appears at the initial time step) [51]. In this context, the property of unique ergodicity has been much studied, in particular by asking if the unique invariant measure is also attractive which is then called *ergodicity* [15] and which can be seen as an analog of asymptotic nilpotency. Among the many examples of probabilistic CA, a particularly interesting family are those obtained from a deterministic one by adding a (small) noise at each transition. In that setting, ergodicity intuitively means that the noise will progressively erase any information from the initial configuration, and non-ergodicity is a form of resistance to noise. It appears that in many cases the resulting noisy CA is ergodic [52] as a probabilistic CA. In fact, proving that a particular noisy CA is *not* ergodic turns out to be difficult. A famous example in dimension 2 is a non-symmetric majority rule which was proven to be non-ergodic in [72]. The result in [72] is in fact more general: any monotone deterministic CA which is nilpotent over finite configurations yields a non-ergodic probabilistic CA when adding a small enough noise. Thus in dimension 2, very simple rules can produce CA resistant to noise. In dimension 1, the only such example known is extremely complicated [25] (recall that the construction technique involved in this result inspired the construction of Theorem 2).

Self-assembly tilings. In the field of DNA computing, several models have been proposed that describe how a spatially extended structure (an *assembly*) grows according to some local rules. For instance, the aTAM model [79], usually considered on the grid \mathbb{Z}^2 , uses the formalism of Wang tiles with glues to define the local process of growth: an assembly is a (usually finite) set of occupied positions of \mathbb{Z}^2 and the aTAM rules determine what tiles can be added on an empty position in the neighborhood of the assembly. A key parameter of the model is *temperature* which determines the amount of glues necessary to allow a tile to be added. We skip the details of the definition, but they are inspired by biochemistry and one of the strengths of the model is that it is both theoretically interesting and realistic enough to allow in vitro experiments. We refer to [63] for a global overview (including many other self-assembly models). Any aTAM system can be seen as a 1-change (asynchronous and non-deterministic) process acting on $Q^{\mathbb{Z}^2}$ where Q is the set of tiles plus a special symbol representing an empty position. Note that in the case of a deterministic aTAM system (called *directed*), one can even define a freezing (1-change) CA F that produces the same limit configurations (see proof of Theorem 18). Important progress have been accomplished recently in understanding the computational power of aTAM depending on the temperature parameter [55, 56].

Automata networks. We worked with \mathbb{G} infinite, but in some sense we already considered the case of a finite space of cells when introducing nilpotency on periodic configurations (Definition 3): indeed, the orbit of a totally periodic configuration on \mathbb{Z}^d can be seen as an orbit of configurations on a finite space of cell which is a torus (because since CA commute with translations, periods are preserved under iterations). One can go much further by considering automata networks which are non-uniform CA on a space of cells which can be an

arbitrary graph [50, 32]. Actually, Theorem 10 has its roots in the automata network settings [35]. In this context, a particularly fruitful approach is to analyze how the graph of cells influences the dynamics of the network (see [26] and references therein for an overview). Following this approach, many results focused on nilpotency or convergence towards a unique attractor [65, 3, 27]. More recently, freezing automata networks were also considered [33, 31, 34].

Acknowledgment

We thank the anonymous referees for their careful reading, their many suggestions and corrections, and in particular, for having corrected us on the fact that the bootstrap percolation CA F_B is generically nilpotent (Theorem 11).

References

- [1] Stål O. Aanderaa and Harry R. Lewis. Linear sampling and the aea case of the decision problem. *Journal of Symbolic Logic*, 39(3):519–548, 1974. doi:10.2307/2272894.
- [2] Shigeki Akiyama and Pierre Arnoux, editors. *Substitution and Tiling Dynamics: Introduction to Self-inducing Structures*. Springer International Publishing, 2020. doi:10.1007/978-3-030-57666-0.
- [3] Julio Aracena. Maximum number of fixed points in regulatory Boolean networks. *Bulletin of mathematical biology*, 70:1398–1409, 2008.
- [4] Per Bak, Kan Chen, and Chao Tang. A forest-fire model and some thoughts on turbulence. *Physics Letters A*, 147(5):297 – 300, 1990. doi:10.1016/0375-9601(90)90451-S.
- [5] Paul Balister, Béla Bollobás, Michał Przykucki, and Paul Smith. Subcritical \mathcal{U} -bootstrap percolation models have non-trivial phase transitions. *Transactions of the American Mathematical Society*, 368(10):7385–7411, 2016. doi:10.1090/tran/6586.
- [6] Laurent Bartholdi. Gardens of eden and amenability on cellular automata. *Journal of the European Mathematical Society*, pages 241–248, 2010. doi:10.4171/jems/196.
- [7] Béla Bollobás, Paul Smith, and Andrew Uzzell. Monotone cellular automata in a random environment. *Combinatorics, Probability and Computing*, 24(4):687–722, 2015. doi:10.1017/S0963548315000012.
- [8] L. Boyer, M. Delacourt, V. Poupet, M. Sablik, and G. Theyssier. mu-limit sets of cellular automata from a computational complexity perspective. *Journal of Computer and System Sciences*, 81(8):1623–1647, dec 2015. doi:10.1016/j.jcss.2015.05.004.
- [9] Mike Boyle, Douglas Lind, and Daniel Rudolph. The automorphism group of a shift of finite type. *Transactions of the American Mathematical Society*, 306(1):71–71, jan 1988. doi:10.1090/s0002-9947-1988-0927684-2.

- [10] Juraj Hromkovič and Georg Schnitger. Communication complexity and sequential computation. In *MFCS '97*, pages 71–84, London, UK, 1997. Springer-Verlag.
- [11] Silvio Capobianco, Pierre Guillon, and Jarkko Kari. Surjective cellular automata far from the Garden of Eden. *Discrete Mathematics and Theoretical Computer Science*, Vol. 15 no. 3(3):41–60, October 2013. Automata, Logic and Semantics. URL: <https://hal.inria.fr/hal-00966380>.
- [12] Olivier Carton, Bruno Guillon, and Fabian Reiter. Counter machines and distributed automata. In *Cellular Automata and Discrete Complex Systems*, pages 13–28. Springer International Publishing, 2018. doi:10.1007/978-3-319-92675-9_2.
- [13] Tullio Ceccherini-Silberstein and Michel Coornaert. *Cellular Automata and Groups*. Springer, 2010.
- [14] J. Chalupa, P. L. Leath, and G. R. Reich. Bootstrap percolation on a Bethe lattice. *Journal of Physics C: Solid State Physics*, 12(1):L31–L35, 1979. doi:10.1088/0022-3719/12/1/008.
- [15] Philippe Chassaing and Jean Mairesse. A non-ergodic probabilistic cellular automaton with a unique invariant measure. *Stochastic Processes and their Applications*, 121(11):2474–2487, nov 2011. doi:10.1016/j.spa.2011.06.009.
- [16] B. Chopard and M. Droz. *Cellular automata modeling of physical systems*. Cambridge University Press New York, 1998.
- [17] K. Čulik, II, J. Pachl, and S. Yu. On the limit sets of cellular automata. *SIAM Journal on Computing*, 18(4):831–842, August 1989.
- [18] M. Delacourt, V. Poupet, M. Sablik, and G. Theyssier. Directional dynamics along arbitrary curves in cellular automata. *Theoretical Computer Science*, 412(30):3800–3821, jul 2011. doi:10.1016/j.tcs.2011.02.019.
- [19] Martin Delacourt and Benjamin Hellouin de Menibus. Characterisation of limit measures of higher-dimensional cellular automata. *Theory of Computing Systems*, 61(4):1178–1213, 2017. doi:10.1007/s00224-017-9753-1.
- [20] Saliha Djenaoui and Pierre Guillon. The generic limit set of cellular automata. *Journal of Cellular Automata*, 14(5-6):435–477, 2019. URL: <https://www.oldcitypublishing.com/journals/jca-home/jca-issue-contents/jca-volume-14-number-5-6-2019/jca-14-5-6-p-435-477/>.
- [21] Richard Durrett. Oriented percolation in two dimensions. *The Annals of Probability*, 12(4), nov 1984. doi:10.1214/aop/1176993140.
- [22] Jeremias Epperlein and Tom Meyerovitch. Iterated minkowski sums, horoballs and north-south dynamics, 2020. arXiv:2009.09221.
- [23] Nazim Fatès. Remarks on the cellular automaton global synchronisation problem: deterministic versus stochastic models. *Natural Computing*, 18(3):429–444, 2019. doi:10.1007/s11047-018-9683-0.

- [24] M.A. Fuentes and M.N. Kuperman. Cellular automata and epidemiological models with spatial dependence. *Physica A: Statistical Mechanics and its Applications*, 267(3–4):471 – 486, 1999.
- [25] Peter Gács. Reliable cellular automata with self-organization. *Journal of Statistical Physics*, 103(1/2):45–267, 2001. doi:10.1023/a:1004823720305.
- [26] Maximilien Gadouleau. On the influence of the interaction graph on a finite dynamical system. *Natural Computing*, 19(1):15–28, feb 2019. doi:10.1007/s11047-019-09732-y.
- [27] Maximilien Gadouleau and Adrien Richard. Simple dynamics on graphs. *Theoretical Computer Science*, 628:62–77, may 2016. doi:10.1016/j.tcs.2016.03.013.
- [28] Anahí Gajardo, Jarkko Kari, and Andrés Moreira. On time-symmetry in cellular automata. *Journal of Computer and System Science*, 78(4):1115–1126, jul 2012. doi:10.1016/j.jcss.2012.01.006.
- [29] Yuval Ginosar and Ron Holzman. The majority action on infinite graphs: strings and puppets. *Discrete Mathematics*, 215(1-3):59–71, mar 2000. doi:10.1016/s0012-365x(99)00228-9.
- [30] E. Goles, N. Ollinger, and G. Theyssier. Introducing freezing cellular automata. In *Exploratory Papers of Cellular Automata and Discrete Complex Systems (AUTOMATA 2015)*, pages 65–73, 2015.
- [31] Eric Goles, Diego Maldonado, Pedro Montealegre-Barba, and Nicolas Ollinger. Fast-parallel algorithms for freezing totalistic asynchronous cellular automata. In *Cellular Automata - 13th International Conference on Cellular Automata for Research and Industry, ACRI 2018, Como, Italy, September 17-21, 2018, Proceedings*, volume 11115 of *Lecture Notes in Computer Science*, pages 406–415. Springer, 2018. doi:10.1007/978-3-319-99813-8_37.
- [32] Eric Goles and Servet Martínez. *Neural and Automata Networks: Dynamical Behavior and Applications*. Kluwer Academic Publishers, Norwell, MA, USA, 1990.
- [33] Eric Goles, Pedro Montealegre, Martín Ríos Wilson, and Guillaume Theyssier. On the impact of treewidth in the computational complexity of freezing dynamics. In *17th Conference on Computability in Europe, CiE 2021*, volume 12813 of *Lecture Notes in Computer Science*, pages 260–272. Springer, 2021. doi:10.1007/978-3-030-80049-9_24.
- [34] Eric Goles, Pedro Montealegre-Barba, and Ioan Todinca. The complexity of the bootstrapping percolation and other problems. *Theoretical Computer Science*, 504:73–82, sep 2013. doi:10.1016/j.tcs.2012.08.001.
- [35] E Goles-Chacc, F Fogelman-Soulie, and D Pellegrin. Decreasing energy functions as a tool for studying threshold networks. *Discrete Applied Mathematics*, 12(3):261–277, 1985.

- [36] Janko Gravner and David Griffeath. Cellular automaton growth on \mathbb{Z}^2 : Theorems, examples, and problems. *Advances in Applied Mathematics*, 21(2):241 – 304, 1998.
- [37] D. Griffeath and C. Moore. Life without death is P-complete. *Complex Systems*, 10, 1996.
- [38] Pierre Guillon and Gaétan Richard. Nilpotency and limit sets of cellular automata. In *Mathematical Foundations of Computer Science 2008, MFCS 2008*, volume 5162 of *Lecture Notes in Computer Science*, pages 375–386. Springer, 2008. doi:10.1007/978-3-540-85238-4_30.
- [39] G. A. Hedlund. Endomorphisms and Automorphisms of the Shift Dynamical Systems. *Mathematical Systems Theory*, 3(4):320–375, 1969.
- [40] Alexander E. Holroyd. Sharp metastability threshold for two-dimensional bootstrap percolation. *Probability Theory and Related Fields*, 125(2):195–224, 2003.
- [41] Lyman P. Hurd. Formal language characterizations of cellular automaton limit sets. *Complex Systems*, 1:69–80, 1987.
- [42] J. Kari. The Nilpotency Problem of One-dimensional Cellular Automata. *SIAM Journal on Computing*, 21:571–586, 1992.
- [43] William Ogilvy Kermack and A. G. McKendrick. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 115(772):700–721, aug 1927. doi:10.1098/rspa.1927.0118.
- [44] P. Kůrka. *Topological and symbolic dynamics*. Société Mathématique de France, 2003.
- [45] P. Kůrka and A. Maass. Limit Sets of Cellular Automata Associated to Probability Measures. *Journal of Statistical Physics*, 100(5-6):1031–1047, 2000.
- [46] Martin Kutrib and Andreas Malcher. Cellular automata with sparse communication. *Theoretical Computer Science*, 411(38-39):3516–3526, 2010. doi:10.1016/j.tcs.2010.05.024.
- [47] James I. Lathrop, Jack H. Lutz, Matthew J. Patitz, and Scott M. Summers. Computability and complexity in self-assembly. *Theory of Computing Systems*, 48(3):617–647, 2011.
- [48] D. Lind and B. Marcus. *An introduction to symbolic dynamics and coding*. Cambridge University Press, 1995.
- [49] D.A. Lind. Applications of ergodic theory and sofic systems to cellular automata. *Physica D: Nonlinear Phenomena*, 10(1-2):36 – 44, 1984. doi: [http://dx.doi.org/10.1016/0167-2789\(84\)90246-X](http://dx.doi.org/10.1016/0167-2789(84)90246-X).
- [50] W. S. Mac Culloch and W. S. Pitts. A logical calculus of the ideas immanent in nervous activity. *Bull. Math. Bio. Phys.*, 5:113–115, 1943.

- [51] Jean Mairesse and Irène Marcovici. Around probabilistic cellular automata. *Theoretical Computer Science*, 559:42–72, nov 2014. doi:10.1016/j.tcs.2014.09.009.
- [52] Irene Marcovici, Mathieu Sablik, and Siamak Taati. Ergodicity of some classes of cellular automata subject to noise. *Electronic journal of probability*, 24, 2019. doi:10.1214/19-EJP297.
- [53] B. Hellouin De Menibus, V. Salo, and G. Theyssier. Characterizing asymptotic randomization in abelian cellular automata. *Ergodic Theory and Dynamical Systems*, 40(4):923–952, sep 2018. doi:10.1017/etds.2018.75.
- [54] Benjamin Hellouin De Menibus and Mathieu Sablik. Characterization of sets of limit measures of a cellular automaton iterated on a random configuration. *Ergodic Theory and Dynamical Systems*, 38(2):601–650, sep 2016. doi:10.1017/etds.2016.46.
- [55] Pierre-Étienne Meunier, Damien Regnault, and Damien Woods. The program-size complexity of self-assembled paths. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pages 727–737. ACM, 2020. doi:10.1145/3357713.3384263.
- [56] Pierre-Étienne Meunier and Damien Woods. The non-cooperative tile assembly model is not intrinsically universal or capable of bounded turing machine simulation. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 328–341. ACM, 2017. doi:10.1145/3055399.3055446.
- [57] John Milnor. On the concept of attractor. *Communications in Mathematical Physics*, 99(2):177–195, jun 1985. doi:10.1007/bf01212280.
- [58] Munemi Miyamoto. An equilibrium state for a one-dimensional life game. *Journal of Mathematics of Kyoto University*, 19(3):525–540, 1979.
- [59] Kenichi Morita. Reversible computing and cellular automata—a survey. *Theoretical Computer Science*, 395(1):101–131, apr 2008. doi:10.1016/j.tcs.2008.01.041.
- [60] Robert Morris. Bootstrap percolation, and other automata. *European Journal of Combinatorics*, 66:250–263, dec 2017. doi:10.1016/j.ejc.2017.06.024.
- [61] Robert Morris. *Monotone cellular automata*, page 312–371. London Mathematical Society Lecture Note Series. Cambridge University Press, 2017. doi:10.1017/9781108332699.008.
- [62] Nicolas Ollinger and Guillaume Theyssier. Freezing, Bounded-Change and Convergent Cellular Automata. *Discrete Mathematics & Theoretical Computer Science*, vol. 24, no. 1, January 2022. URL: <https://dmtcs.episciences.org/9004>, doi:10.46298/dmtcs.5734.

- [63] Matthew J. Patitz. An introduction to tile-based self-assembly and a survey of recent results. *Natural Computing*, 13(2):195–224, jun 2013. doi:10.1007/s11047-013-9379-4.
- [64] Marcus Pivato. Ergodic theory of cellular automata. In *Encyclopedia of Complexity and Systems Science*, pages 2980–3015. Springer, New York, NY, 2009. doi:10.1007/978-0-387-30440-3_178.
- [65] F. Robert. *Discrete iterations: a metric study*, volume 6 of *Series in Computational Mathematics*. Springer, 1986.
- [66] H. Rogers. *Theory of Recursive Functions and Effective Computability*. MIT Press, 1967.
- [67] Ville Salo. On nilpotency and asymptotic nilpotency of cellular automata. In *AUTOMATA & JAC 2012*, volume 90 of *EPTCS*, pages 86–96, 2012. URL: <https://doi.org/10.4204/EPTCS.90.7>.
- [68] Ville Salo. Strict asymptotic nilpotency in cellular automata. In *Cellular Automata and Discrete Complex Systems*, pages 3–15, Cham, 2017. Springer International Publishing.
- [69] Ville Salo. No tits alternative for cellular automata. *Groups, Geometry, and Dynamics*, 13(4):1437–1455, sep 2019. doi:10.4171/ggd/529.
- [70] Ville Salo, Guillaume Theyssier, and Ilkka Törmä. Bootstrap percolation and cellular automata, 2021. arXiv:2110.00656.
- [71] Ville Salo and Ilkka Törmä. Nilpotent endomorphisms of expansive group actions. *International Journal of Algebra and Computation*, 31(03):393–452, feb 2021. doi:10.1142/s021819672150020x.
- [72] A. L. Toom. Stable and attractive trajectories in multicomponent systems. *Multicomponent random systems*, 6:549–575, 1980.
- [73] Ilkka Törmä. A uniquely ergodic cellular automaton. *Journal of Computer and System Sciences*, 81(2):415–442, 2015. URL: <https://www.sciencedirect.com/science/article/pii/S0022000014001366>, doi: <https://doi.org/10.1016/j.jcss.2014.10.001>.
- [74] Ilkka Törmä. Generically nilpotent cellular automata, 2021. arXiv:2108.12925.
- [75] S. M. Ulam. On some mathematical problems connected with patterns of growth of figures. In A. W. Burks, editor, *Essays on Cellular Automata*, pages 219–231. U. of Illinois Press, 1970.
- [76] Aernout C. D. van Enter. Proof of straley's argument for bootstrap percolation. *Journal of Statistical Physics*, 48(3-4):943–945, aug 1987. doi:10.1007/bf01019705.
- [77] R. Vollmar. On cellular automata with a finite number of state changes. In *Parallel Processes and Related Automata / Parallele Prozesse und damit zusammenhängende Automaten*, volume 3 of *Computing Supplementum*, pages 181–191. Springer Vienna, 1981. doi:10.1007/978-3-7091-8596-4_13.

- [78] P. Walters. *An Introduction to Ergodic Theory*. Graduate Texts in Mathematics. Springer New York, 1981.
- [79] Erik Winfree. *Algorithmic Self-Assembly of DNA*. PhD thesis, California Institute of Technology, 1998.
- [80] Dieter A. Wolf-Gladrow. *Lattice Gas Cellular Automata and Lattice Boltzmann Models*. Springer Berlin Heidelberg, 2000. doi:10.1007/b72010.