

PROBABILISTIC CELLULAR AUTOMATA, INVARIANT MEASURES, AND PERFECT SAMPLING

ANA BUŠIĆ,* INRIA - École Normale Supérieure

JEAN MAIRESSE ** *** AND

IRÈNE MARCOVICI,** **** Université Paris Diderot

Abstract

A probabilistic cellular automaton (PCA) can be viewed as a Markov chain. The cells are updated synchronously and independently, according to a distribution depending on a finite neighborhood. We investigate the ergodicity of this Markov chain. A classical cellular automaton is a particular case of PCA. For a one-dimensional cellular automaton, we prove that ergodicity is equivalent to nilpotency, and is therefore undecidable. We then propose an efficient perfect sampling algorithm for the invariant measure of an ergodic PCA. Our algorithm does not assume any monotonicity property of the local rule. It is based on a bounding process which is shown to also be a PCA. Last, we focus on the PCA majority, whose asymptotic behavior is unknown, and perform numerical experiments using the perfect sampling procedure.

Keywords: Probabilistic cellular automata; perfect sampling; ergodicity; invariant measure

2010 Mathematics Subject Classification: Primary 37B15; 60J05; 60J22

Secondary 37A25; 60K35; 68Q80

1. Introduction

A *deterministic cellular automaton* (DCA) consists of a lattice (e.g. \mathbb{Z} or \mathbb{Z}^2 or $\mathbb{Z}/n\mathbb{Z}$) divided into regular cells, each cell containing a letter of a finite alphabet. The cells evolve synchronously, each one evolving as a function of a finite number of cells in its neighborhood, according to a local rule.

DCAs form a natural mathematical object. By Hedlund's theorem [15], the mappings realized by DCAs are precisely the continuous functions (for the product topology) commuting with the shift. They also constitute a powerful model of computation, in particular, they can 'simulate' any Turing machine. Last, due to the considerable gap between the simplicity of the definition and the intricacy of the generated behaviors, DCAs are good candidates for modeling 'complex systems' appearing in physical and biological processes.

Probabilistic cellular automata. To take into account random events, one is led to consider probabilistic versions of DCAs. In one of them, at most one cell is updated at each time, this cell being randomly chosen according to a given distribution. For an infinite set of cells, one is led to consider continuous-time models, obtaining what is known in probability theory as an *interacting particle system* [20]. Another model is the *probabilistic cellular automaton* (PCA) [24], [27]. For the PCA, time is discrete, and all the cells evolve synchronously as for the DCA,

Received 22 March 2011; revision received 20 July 2012.

* Postal address: Laboratoire d'Informatique de l'École Normale Supérieure (UMR 8548), INRIA - École Normale Supérieure, 23 avenue d'Italie, CS 81321, 75214 Paris Cedex 13, France. Email address: ana.busic@inria.fr

** Postal address: CNRS, UMR 7089, LIAFA, Université Paris Diderot, Sorbonne Paris Cité, F-75205 Paris, France.

*** Email address: jean.mairesse@liafa.univ-paris-diderot.fr

**** Email address: irene.marcovic@liafa.univ-paris-diderot.fr

but the difference is that, for each cell, the new content is randomly chosen, independently of the others, according to a distribution depending only on a finite neighborhood of the cell.

Let us mention a couple of motivations. First, the investigation of fault-tolerant computational models was the motivation for [10] and [27]. Second, PCAs appear in combinatorial problems related to the enumeration of directed animals [1] [7] [19]. Third, in the context of the classification of DCAs (Wolfram's program), robustness to random errors can be used as a discriminating criterion [9] [23]. Recently, PCAs also proved to be pertinent for the density classification problem [8], that is, efficiently testing whether some sequence contains more occurrences of 0 or 1. Last, PCAs are used in statistical physics and in life sciences. They model various phenomena, ranging from the dynamical properties of neural tissue [18] to competition between species.

We focus our study on the equilibrium behavior of PCAs. Observe that a PCA may be viewed as a Markov chain over the state space \mathcal{A}^E , where \mathcal{A} is the alphabet and E is the set of cells. So the equilibrium is studied via the invariant measures of the Markov chain. Several questions are in order.

Ergodicity. A PCA is *ergodic* if it has a unique and attractive invariant measure. A challenging problem in this area is the *positive rates conjecture*. A PCA is said to have *positive rates* if, for any neighborhood, the updated content of a cell can be any letter with a strictly positive probability. The positive rates conjecture states that any one-dimensional ($E = \mathbb{Z}$) model with positive rates is ergodic. Gács exhibited in 2001 a very large and complex counterexample in a paper of more than 200 pages [10] which was published with an introductory article by Gray [11]. This counterexample is far from being completely understood and several questions remain. For instance, does the conjecture hold true for small alphabets and neighborhoods? Even for alphabets and neighborhoods of size 2, the question is not settled.

Performance evaluation. The second natural question is whether the invariant measures can be evaluated. A PCA with an alphabet and a neighborhood of size 2 is determined by four parameters. If the parameters satisfy a given polynomial equation, there exists an invariant measure with an explicit product form [27, Chapter 16]. Under another polynomial condition, there exists an invariant measure with an explicit Markovian form; see [27, Chapter 16] or [1]. What happens for generic values of the parameters, or for PCAs with a larger neighborhood or alphabet? When explicit computation is not possible, simulation becomes the alternative. Simulating PCAs is known to be a challenging task, costly both in time and space. Also, configurations cannot be tracked down one by one (there is an infinite number of them when $E = \mathbb{Z}$) and may only be observed through some measured parameters. So the crucial point is whether some guarantees can be given upon the results obtained from simulations.

The contributions of the present paper focus on (i) ergodicity and (ii) performance evaluation.

(i) We prove that the ergodicity of a DCA on \mathbb{Z} is undecidable. This was mentioned as Unsolved Problem 4.5 of [26]. Since a DCA is a special case of a PCA, it also provides a new proof of the undecidability of the ergodicity of a PCA (see [27, Chapter 14] and [25]).

(ii) Given an ergodic PCA, a *perfect sampling* procedure is a random algorithm which returns a configuration distributed according to the invariant measure. By applying the procedure repeatedly, we can estimate the invariant measure with arbitrary precision. We propose such an algorithm for PCAs by adapting the *coupling-from-the-past* method of Propp and Wilson [21]. When the set of cells is $E = \mathbb{Z}/n\mathbb{Z}$, a PCA is a finite state space Markov chain. Therefore, coupling from the past from all possible initial configurations provides a basic perfect sampling procedure. But it is a very inefficient sampling procedure since the number of configurations is exponential in n . Here, the contribution consists in simplifying

the procedure. We define a new PCA on an extended alphabet, called the *envelope PCA*. We obtain a perfect sampling procedure for the original PCA by running the envelope PCA on a single initial configuration. When the set of cells is $E = \mathbb{Z}^d$, a PCA is a Markov chain on an uncountable state space. So there is no basic perfect sampling procedure anymore. We prove that if the PCA is ergodic then the envelope PCA may or may not be ergodic. If it is ergodic then we can use the envelope PCA to design an efficient perfect sampling procedure (the result of the algorithm is the finite restriction of a configuration with the right invariant distribution). In the case $E = \mathbb{Z}$, we give a sufficient condition for the envelope PCA to be ergodic. The envelope PCA can be viewed as a systematic treatment of ideas already used by Toom for *percolation PCAs* (see, for instance, [26, Section 2]).

The perfect sampling procedure can also be run on a PCA whose ergodicity is unknown, with the purpose of testing it. We illustrate this approach on the majority PCA, a prototype of a PCA whose equilibrium behavior is not well understood. More precisely, we define a parametrized family of PCA, called $\text{Majority}(\alpha)$, $\alpha \in (0, 1)$. We conjecture the existence of a phase transition between two situations: (i) several invariant measures; (ii) a unique but nonattractive invariant measure. We provide some numerical evidence for the phase transition, which would be the first example of this kind. In fact, the mere existence of a PCA satisfying (ii) had been a long standing open problem which has recently been positively resolved [4].

In Section 2 we give the basic definitions. Section 3 is devoted to the ergodicity problem. In Section 4 we present the perfect sampling procedures. Last, Section 5 is devoted to the case study of the majority PCA.

A short version without proofs of this paper appears in the proceedings of the STACS'2011 conference [3].

2. Probabilistic cellular automata

Let \mathcal{A} be a finite set called the *alphabet*, and let E be a countable or finite set of *cells*. We denote by X the set \mathcal{A}^E of *configurations*.

We assume that E is equipped with a commutative semigroup structure, whose law is denoted by '+'. In examples, we consider mostly the cases $E = \mathbb{Z}$ or $E = \mathbb{Z}/n\mathbb{Z}$. Given $K \subset E$ and $V \subset E$, we define

$$V + K = \{v + k \in E \mid v \in V, k \in K\}.$$

A *cylinder* is a subset of X having the form $\{x \in X \mid \text{for all } k \in K, x_k = y_k\}$ for a given finite subset K of E and a given element $(y_k)_{k \in K} \in \mathcal{A}^K$. When there is no possible confusion, we shall denote by y_K the cylinder $\{x \in X \mid \text{for all } k \in K, x_k = y_k\}$. For a given finite subset K , we denote by $\mathcal{C}(K)$ the set of all cylinders of base K .

Let us equip $X = \mathcal{A}^E$ with the product topology, which can be described as the topology generated by cylinders. We denote by $\mathcal{M}(\mathcal{A})$ the set of probability measures on \mathcal{A} and by $\mathcal{M}(X)$ the set of probability measures on X for the σ -algebra generated by all cylinder sets, which corresponds to the Borelian σ -algebra. For $x \in X$, denote by δ_x the Dirac measure concentrated on the configuration x .

Definition 2.1. Given a finite set $V \subset E$, a *transition function of neighborhood V* is a function $f: \mathcal{A}^V \rightarrow \mathcal{M}(\mathcal{A})$. The PCA P of transition function f is the application

$$\begin{aligned} P: \mathcal{M}(X) &\rightarrow \mathcal{M}(X), \\ \mu &\mapsto \mu P, \end{aligned}$$

defined on cylinders by

$$\mu^P(y_K) = \sum_{x_{V+K} \in \mathcal{C}(V+K)} \mu(x_{V+K}) \prod_{k \in K} f((x_{k+v})_{v \in V})(y_k).$$

Let us look at how P acts on a Dirac measure δ_z . The content z_k of the k th cell is changed into the letter $a \in \mathcal{A}$ with probability $f((z_{k+v})_{v \in V})(a)$, independently of the evolution of the other cells. The real number $f((z_{k+v})_{v \in V})(a) \in [0, 1]$ is thus to be thought of as the conditional probability that, after application of P , the k th cell will be in the state a if, before its application, the neighborhood of k was in the state $(z_{k+v})_{v \in V}$.

Let u be the uniform measure on $[0, 1]$. We define the product measure $\tau = \otimes_{i \in E} u$ on $[0, 1]^E$.

Definition 2.2. An *update function* of the probabilistic cellular automaton P is a deterministic function $\phi: \mathcal{A}^E \times [0, 1]^E \rightarrow \mathcal{A}^E$ (the function ϕ takes as argument a configuration and a sample in $[0, 1]^E$, and returns a new configuration), satisfying, for each $x \in \mathcal{A}^E$ and each cylinder y_K ,

$$\tau(\{r \in [0, 1]^E; \phi(x, r) \in y_K\}) = \prod_{k \in K} f((x_{k+v})_{v \in V})(y_k).$$

In practice, it is always possible to define an update function ϕ for which the value of $\phi(x, r)_k$ depends only on $(x_{k+v})_{v \in V}$ and on r_k . For example, if the alphabet is $\mathcal{A} = \{a_1, \dots, a_n\}$, we can set

$$\phi(x, r)_k = \begin{cases} a_1 & \text{if } 0 \leq r_k < f((x_{k+v})_{v \in V})(\{a_1\}), \\ a_2 & \text{if } f((x_{k+v})_{v \in V})(\{a_1\}) \leq r_k < f((x_{k+v})_{v \in V})(\{a_1, a_2\}), \\ \vdots & \\ a_n & \text{if } f((x_{k+v})_{v \in V})(\{a_1, a_2, \dots, a_{n-1}\}) \leq r_k \leq 1. \end{cases} \tag{2.1}$$

For a given initial configuration $x^0 \in \mathcal{A}^E$, and samples $(r^t)_{t \in \mathbb{N}}$, $r^t \in [0, 1]^E$, let $(x^t)_{t \in \mathbb{N}} \in (\mathcal{A}^E)^\mathbb{N}$ be the sequence defined recursively by

$$x^{t+1} = \phi(x^t, r^t).$$

Such a sequence is called a *space–time diagram*. It can be viewed as a realization of the Markov chain. Examples of space–time diagrams appear in Figures 1 and 8.

Classical cellular automata are a specialization of PCAs.

Definition 2.3. A *deterministic cellular automaton* (DCA) is a PCA such that, for each sequence $(x_v)_{v \in V} \in \mathcal{A}^V$, the measure $f((x_v)_{v \in V})$ is concentrated on a single letter of the alphabet. A DCA can thus be seen as a deterministic function $F: \mathcal{A}^E \rightarrow \mathcal{A}^E$.

In the literature, the term *cellular automaton* denotes what we call here a DCA. DCAs have been widely studied, in particular on the set of cells $E = \mathbb{Z}$; see Section 3. For a DCA, any initial configuration defines a unique space–time diagram.

Example 2.1. Let $\mathcal{A} = \{0, 1\}$, $E = \mathbb{Z}$, and $V = \{0, 1\}$. Consider $0 < \varepsilon < 1$ and the local function

$$f(x, y) = (1 - \varepsilon)\delta_{x+y \bmod 2} + \varepsilon\delta_{x+y+1 \bmod 2}.$$

This defines a PCA that can be considered as a perturbation of the DCA $F: \mathcal{A}^E \rightarrow \mathcal{A}^E$ defined by $F(x)_i = x_i + x_{i+1} \bmod 2$, with errors occurring in each cell independently with probability ε .




FIGURE 1: Space–time diagrams of the PCA of Example 2.2 for $V = \{-1, 0, 1\}$.

Example 2.2. Let $\mathcal{A} = \{0, 1\}$, let $E = \mathbb{Z}^d$, and let V be a finite subset of E , $|V| \geq 2$. Consider $0 < \alpha < 1$ and the local function

$$f((x_v)_{v \in V}) = \alpha \delta_{\max(x_v, v \in V)} + (1 - \alpha) \delta_0.$$

The corresponding PCA is called the *percolation PCA* associated with V and α . The particular case of the space $E = \mathbb{Z}$ and the neighborhood $V = \{0, 1\}$ is called the *Stavskaya PCA*. In Figure 1, we represent two space–time diagrams of the percolation PCA for $V = \{-1, 0, 1\}$.

2.1. Invariant measures and ergodicity

A PCA can be seen as a Markov chain on the state space \mathcal{A}^E . We use the classical terminology for Markov chains that we now recall.

Definition 2.4. A probability measure $\pi \in \mathcal{M}(X)$ is said to be an *invariant measure* of the PCA P if $\pi P = \pi$. The PCA is *ergodic* if it has exactly one invariant measure π which is *attractive*, that is, for any measure $\mu \in \mathcal{M}(X)$, the sequence μP^n converges weakly to π (i.e. for any cylinder C , $\lim_{n \rightarrow +\infty} \mu P^n(C) = \pi(C)$).

A PCA has at least one invariant measure, and the set of invariant measures is convex and compact. This is a standard fact, based on the observation that the set $\mathcal{M}(X)$ of measures on X is compact for the weak topology; see, for instance, [27]. Therefore, there are three possible situations. A PCA has

- (i) several invariant measures;
- (ii) a unique invariant measure which is not attractive;
- (iii) a unique invariant measure which is attractive (ergodic case).

Example 2.3. Consider the PCA of Example 2.1. Using the results in [27, Chapters 16 and 17], we can prove that the PCA is ergodic and that its unique invariant measure is the uniform measure, i.e. it is the product of Bernoulli measures of parameter $\frac{1}{2}$.

Example 2.4. Consider the percolation PCA of Example 2.2. Observe that the Dirac measure δ_{0^E} is an invariant measure. Using a coupling with a site percolation model, we can prove the following; see, for instance, [26, Section 2]. There exists $\alpha^* \in (0, 1)$ such that

$$\begin{aligned} \alpha < \alpha^* &\implies \text{(iii) ergodicity,} \\ \alpha > \alpha^* &\implies \text{(i) several invariant measures.} \end{aligned}$$

The exact value of α^* is not known but it satisfies $1/|V| \leq \alpha^* \leq \frac{53}{54}$.

The existence of a PCA corresponding to situation (ii) had been a longstanding conjecture (see Unsolved problem 5.7 of [26]). In [4], it was proved that situation (ii) occurs for the PCA on $\{0, 1\}^{\mathbb{Z}}$ with neighborhood $V = \{0, 1\}$, and local function a defined by $a(00)(1) = \frac{1}{2}$, $a(01)(1) = 0$, $a(10)(1) = 1$, and $a(11)(1) = \frac{1}{2}$.

The PCA of Example 2.2 exhibits a phase transition between the situations (i) and (iii). In Section 5 we study a PCA that may have a phase transition between situations (ii) and (iii). It would provide the first example of this type.

3. Ergodicity of the DCA

DCAs form the simplest class of PCAs; it is therefore natural to study the ergodicity of DCAs. In this section we prove the undecidability of ergodicity for DCAs (Theorem 3.1). This also gives a new proof of the undecidability of the ergodicity for PCAs.

Remark. In the context of DCAs, the terminology of Definition 2.4 might be confusing. Indeed, a DCA P can be viewed in two different ways:

- a (degenerated) Markov chain;
- a symbolic dynamical system.

In the dynamical system terminology, P is *uniquely ergodic* if [there exists a unique μ , $\mu P = \mu$]. In the Markov chain terminology (that we adopt), P is *ergodic* if [there exists a unique μ , $\mu P = \mu$] and [for all ν , $\nu P^n \xrightarrow{w} \mu$], where ‘ \xrightarrow{w} ’ stands for the weak convergence. Knowing whether the unique ergodicity (of symbolic dynamics) implies the ergodicity (of the Markovian theory) is an open question for DCAs.

The *limit set* of P is defined by

$$LS = \bigcap_{n \in \mathbb{N}} P^n(\mathcal{A}^E).$$

In words, a configuration belongs to LS if it may occur after an arbitrarily long evolution of the cellular automaton.

Observe that LS is nonempty since it is the decreasing limit of nonempty closed sets. A constructive way to show that LS is nonempty is as follows. A *monochromatic* configuration is a configuration of the type x^E , $x \in \mathcal{A}$. The image under P of a monochromatic configuration is monochromatic. In particular, there exists a monochromatic periodic orbit for P , and we have

$$x_0^E \rightarrow x_1^E \rightarrow \dots \rightarrow x_{k-1}^E \rightarrow x_0^E \implies \{x_0^E, x_1^E, \dots, x_{k-1}^E\} \subset LS. \tag{3.1}$$

Recall that δ_u denotes the probability measure concentrated on the configuration u . The periodic orbit $(x_0^E, \dots, x_{k-1}^E)$ provides an invariant measure given by $(\delta_{x_0^E} + \dots + \delta_{x_{k-1}^E})/k$. More generally, the support of any invariant measure is included in the limit set.

Definition 3.1. A DCA is *nilpotent* if its limit set is a singleton.

Using (3.1), we see that a DCA is nilpotent if and only if $LS = \{x^E\}$ for some $x \in \mathcal{A}$. The following stronger statement is proved in [6], using a compactness argument:

$$[P \text{ nilpotent}] \iff [\text{there exist } x \in \mathcal{A} \text{ and } N \in \mathbb{N}, P^N(\mathcal{A}^E) = \{x^E\}].$$

We obtain the next proposition as a corollary.

Proposition 3.1. *Consider a DCA P . We have*

$$[P \text{ nilpotent}] \implies [P \text{ ergodic}].$$

Proof. Let $x \in \mathcal{A}$ and $N \in \mathbb{N}$ be such that $P^N(\mathcal{A}^E) = \{x^E\}$. For any probability measure μ on \mathcal{A}^E , we have $\mu P^N = \delta_{x^E}$. Therefore, P is ergodic with unique invariant measure δ_{x^E} . This completes the proof.

If we restrict ourselves to DCAs on \mathbb{Z} , we obtain the converse statement.

Theorem 3.1. *Consider a DCA P on the set of cells \mathbb{Z} . We have*

$$[P \text{ nilpotent}] \iff [P \text{ ergodic}].$$

Proof. Let P be an ergodic DCA. Assume that there exists a monochromatic periodic orbit $(x_0^{\mathbb{Z}}, \dots, x_{k-1}^{\mathbb{Z}})$ with $k \geq 2$. Then $\mu = (\delta_{x_0^{\mathbb{Z}}} + \dots + \delta_{x_{k-1}^{\mathbb{Z}}})/k$ is the unique invariant measure. The sequence $\delta_{x_0^{\mathbb{Z}}} P^n$ does not converge weakly to μ , which is a contradiction. Therefore, there exists a monochromatic fixed point $P(x^{\mathbb{Z}}) = x^{\mathbb{Z}}$, and $\delta_{x^{\mathbb{Z}}}$ is the unique invariant measure.

Define the cylinder $C = x_K^E = \{v \in \mathcal{A}^{\mathbb{Z}} \mid \text{for all } i \in K, v_i = x\}$, where K is some finite subset of \mathbb{Z} . For any initial configuration $u \in \mathcal{A}^{\mathbb{Z}}$, using the ergodicity of P , we have

$$\delta_u P^n(C) \rightarrow \delta_{x^{\mathbb{Z}}}(C) = 1.$$

But $\delta_u P^n$ is a Dirac measure, so $\delta_u P^n(C)$ is equal to 0 or 1. Consequently, we have $\delta_u P^n(C) = 1$ for large enough n , that is,

$$\text{there exists } N \in \mathbb{N} \text{ such that, for all } n \geq N \text{ and } i \in K, P^n(u)_i = x.$$

In words, in any space–time diagram of P , any column becomes eventually equal to $xxx \dots$. Using the terminology of Guillon and Richard [12], the DCA P has a *weakly nilpotent trace*. It was proved in [12] that the weak nilpotency of the trace implies the nilpotency of the DCA. (The result is proved for cellular automata on \mathbb{Z} and left open in larger dimensions.) This completes the proof.

Kari proved in [17] that the nilpotency of a DCA on \mathbb{Z} is undecidable. (For DCAs on \mathbb{Z}^d , $d \geq 2$, the proof appears in [6].) By coupling Kari’s result with Theorem 3.1, we obtain the following.

Corollary 3.1. *Consider a DCA P on the set of cells \mathbb{Z} . The ergodicity of P is undecidable.*

The undecidability of the ergodicity of a PCA was a known result, proved by Kurdyumov; see Toom [27] and also Toom [25]. Kurdyumov’s and Toom’s proofs use a nondeterministic PCA of dimension 1 and a reduction of the halting problem of a Turing machine.

Corollary 3.1 is a stronger statement. In fact, the (un)decidability of the ergodicity of a DCA was mentioned as Unsolved Problem 4.5 of [26]. We point out that Corollary 3.1 can also be obtained without Theorem 3.1, by directly adapting Kari’s proof to show the undecidability of the ergodicity of the DCA associated with an NW-deterministic tile set.

4. Sampling the invariant measure of an ergodic PCA

Generally, the invariant measure(s) of a PCA cannot be described explicitly. Numerical simulations are consequently very useful to get an idea of the behavior of a PCA. Given an ergodic PCA, we propose a *perfect sampling* algorithm which generates configurations *exactly* according to the invariant measure.

A perfect sampling procedure for finite Markov chains has been proposed by Propp and Wilson [21] using a coupling-from-the-past scheme. Perfect sampling procedures have been developed since in various contexts. We mention below some works directly linked to the present article. For more information see the annotated bibliography *Perfectly Random Sampling with Markov Chains* (available online at <http://dbwilson.com/exact/>).

The complexity of the algorithm depends on the number of all possible initial conditions, which is prohibitive for PCAs. A first crucial observation already appears in [21]; for a monotone Markov chain, we have to consider two trajectories corresponding to minimal and maximal states of the system. For antimonotone systems, an analogous technique has been developed by Häggström and Nelander [14] that also considers only extremal initial conditions. To cope with more general situations, Huber [16] introduced the idea of a bounding chain for determining when coupling has occurred. The construction of these bounding chains is model dependent and in general not straightforward. In the case of a Markov chain on a lattice, Bušić *et al.* [2] proposed an algorithm to construct bounding chains.

Our contribution is to show that the bounding chain ideas can be given in a particularly simple and convenient form in the context of a PCA via the introduction of the *envelope PCA*.

4.1. Basic coupling from the past for PCAs

4.1.1. *Finite set of cells.* Consider an ergodic PCA P on the alphabet \mathcal{A} and on a finite set of cells E (for example, $\mathbb{Z}_m = \mathbb{Z}/m\mathbb{Z}$). Let π be the invariant measure on $X = \mathcal{A}^E$. A *perfect sampling* procedure is a random algorithm which returns a state $x \in X$ with probability $\pi(x)$. Let us present the Propp and Wilson, or coupling-from-the-past (CFTP), perfect sampling procedure.

Algorithm 1. (*Basic CFTP algorithm for a finite set of cells.*) *Data:* an update function $\phi: X \times [0, 1]^E \rightarrow X$ of a PCA. A family $(r_k^{-n})_{(k,n) \in E \times \mathbb{N}}$ of independent and identically distributed (i.i.d.) random variables (RVs) with uniform distribution in $[0, 1]$.

```

begin
   $t = 1$ ;
  repeat
     $R_{-t} = X$ ;
    for  $j = -t$  to  $-1$  do
       $R_{j+1} = \{\phi(x, (r_i^j)_{i \in E}); x \in R_j\}$ 
    end
     $t = t + 1$ 
  until  $|R_0| = 1$ ;
  return the unique element of  $R_0$ 
end

```

The ‘good’ way to implement Algorithm 1 is to keep track of the partial couplings of trajectories. This allows us to consider only one-step transitions.

Proposition 4.1. ([21].) *If the procedure stops almost surely then the PCA is ergodic and the output is distributed according to the invariant measure.*




FIGURE 2: Coupling from the past.

The converse statement is not true in general; even for ergodic PCAs, there exist choices of ϕ for which the procedure does not stop. Nevertheless, for PCAs having positive rates (see the introduction), the algorithm stops almost surely in finite time if the update function is chosen according to (2.1).

In Figure 2, we illustrate the algorithm for the toy example of a PCA on the alphabet $\{0, 1\}$ and the set of cells \mathbb{Z}_2 . The state space is thus $X = \{x_1 = 00, x_2 = 01, x_3 = 10, x_4 = 11\}$. On this sample, the algorithm returns x_2 .

A sketch of the proof of Proposition 4.1 can be given using Figure 2. In the bottom-right picture, the Markov chain is run from time -4 onwards and its value is x_2 at time 0. If we had run the Markov chain from time $-\infty$ to 0, then the result would obviously still be x_2 . But if we started from time $-\infty$, then the Markov chain would have reached equilibrium by time 0.

4.1.2. *Infinite set of cells.* Assume that the set of cells E is infinite. Then a PCA defines a Markov chain on the infinite state space $X = \mathcal{A}^E$, so the above procedure is not effective anymore. However, it is possible to use the locality of the updating rule of a PCA to still define a perfect sampling procedure. (This observation already appears in [28].)

Let P be an ergodic PCA, and denote by π its invariant distribution. In this context, a perfect sampling procedure is a random algorithm taking as input a finite subset K of E and returning a cylinder $x_K \in \mathcal{C}(K)$ with probability $\pi(x_K)$.

To get such a procedure, we use the following fact. If the PCA is run from time $-k$ onwards then to compute the content of the cells in K at time 0, it is enough to consider the cells in the finite dependence cone of K . This is illustrated in Figure 3 for the set of cells $E = \mathbb{Z}$ and the neighborhood $V = \{-1, 0, 1\}$, with the choice $K = \{0\}$.

Observe that the orientation has changed with respect to Figure 2 in order to be consistent with the convention of Figures 1 and 8 for space–time diagrams.

Let us define this more formally. Let V be the neighborhood of the PCA. Given a subset K of E , the *dependence cone* of K is the family $(V_{-t}(K))_{t \geq 0}$ of subsets of E defined recursively by $V_0(K) = K$ and $V_{-t}(K) = V + V_{-t+1}(K)$. Let $\phi: X \times [0, 1]^E \rightarrow X$ be an update function, for instance, that defined according to (2.1). For a given subset K of E , we denote by $\phi_{-t}: \mathcal{A}^{V_{-t}(K)} \times [0, 1]^{V_{-t}(K)} \rightarrow \mathcal{A}^{V_{-t+1}(K)}$ the corresponding restriction of ϕ .

With this notation, the algorithm can be written as follows.




FIGURE 3: Dependence cone of a cell.

Algorithm 2. (Basic CFTP algorithm for an infinite set of cells.) *Data:* an update function $\phi: X \times [0, 1]^E \rightarrow X$ of a PCA. A family $(r_k^{-n})_{(k,n) \in E \times \mathbb{N}}$ of i.i.d. RVs with uniform distribution in $[0, 1]$. A finite subset K of E .

begin

$V_0(K) = K;$

$t = 1;$

repeat

$V_{-t}(K) = V + V_{-t+1}(K);$

$R_{-t} = \mathcal{A}^{V_{-t}(K)};$

for $j = -t$ **to** -1 **do**

$R_{j+1} = \{\phi_j(x, (r_i^j)_{i \in V_j(K)}); x \in R_j\} \subset \mathcal{A}^{V_{j+1}(K)}$

end

$t = t + 1$

until $|R_0| = 1;$

return the unique element of R_0

end

The next proposition is an easy extension of Proposition 4.1.

Proposition 4.2. *If the procedure stops almost surely then the PCA is ergodic and the output is distributed according to the marginal of the invariant measure.*

The converse statement is not true in general. It would be interesting to know whether it holds true for the update function (2.1) and for PCAs having positive rates (possibly under an additional hypothesis).

4.2. Envelope probabilistic cellular automata

The CFTP algorithm is inefficient when the state space is large. This is the case for PCAs. When E is finite, the set \mathcal{A}^E is very large, and when E is infinite, it is the number of configurations living in the dependence cone described above which is very large. We cope with this difficulty by introducing the envelope PCA.

To begin with, let us assume that P is a PCA on the alphabet $\mathcal{A} = \{0, 1\}$ (as previously, the set of cells is denoted by E , the neighborhood by $V \subset E$, and the local function by f). The case of a general alphabet is treated in Section 4.5.

4.2.1. *Definition of the envelope probabilistic cellular automaton.* Let us introduce a new alphabet

$$\mathcal{B} = \{\mathbf{0}, \mathbf{1}, \mathbf{?}\}.$$

A word over \mathcal{B} is to be thought of as a word over \mathcal{A} in which the letters corresponding to some positions are not known, and are thus replaced by the symbol ‘?’ . Formally, we identify \mathcal{B} with $2^{\mathcal{A}} - \emptyset$ as follows: $\mathbf{0} = \{0\}$, $\mathbf{1} = \{1\}$, and $\mathbf{?} = \{0, 1\}$. So each letter of \mathcal{B} is a set of possible letters of \mathcal{A} . With this interpretation, we view a word over \mathcal{B} as a set of words over \mathcal{A} . For instance,

$$\mathbf{?1?} = \{010, 011, 110, 111\}.$$

We associate to the PCA P a new PCA on the alphabet \mathcal{B} , that we call the *envelope probabilistic cellular automaton* (EPCA) of P .

Definition 4.1. The EPCA of P is the PCA $\text{env}(P)$ of alphabet \mathcal{B} , defined on the set of cells E , with the same neighborhood V as for P , and a local function $\text{env}(f): \mathcal{B}^V \rightarrow \mathcal{M}(\mathcal{B})$ defined for each $y \in \mathcal{B}^V$ by

$$\begin{aligned} \text{env}(f)(y)(\mathbf{0}) &= \min_{x \in \mathcal{A}^V, x \in y} f(x)(0), & \text{env}(f)(y)(\mathbf{1}) &= \min_{x \in \mathcal{A}^V, x \in y} f(x)(1), \\ \text{env}(f)(y)(\mathbf{?}) &= 1 - \min_{x \in \mathcal{A}^V, x \in y} f(x)(0) - \min_{x \in \mathcal{A}^V, x \in y} f(x)(1). \end{aligned}$$

We point out that $\min_{x \in \mathcal{A}^V, x \in y} f(x)(1) + \max_{x \in \mathcal{A}^V, x \in y} f(x)(0) = 1$, so that the last quantity $\text{env}(f)(y)(\mathbf{?})$ is nonnegative.

Moreover, $\text{env}(P)$ acts like P on configurations which do not contain the letter ‘?’ . More precisely, for all $y \in \mathcal{A}^V$,

$$\text{env}(f)(y)(\mathbf{0}) = f(y)(0), \quad \text{env}(f)(y)(\mathbf{1}) = f(y)(1), \quad \text{env}(f)(y)(\mathbf{?}) = 0. \tag{4.1}$$

In particular, we obtain the following.

Proposition 4.3. *If the EPCA $\text{env}(P)$ is ergodic then the PCA P is ergodic.*

Proof. According to (4.1), any invariant measure of P corresponds to an invariant measure of $\text{env}(P)$. Therefore, if P has several invariant measures, so does $\text{env}(P)$. Assume that P has a unique invariant measure μ which is nonergodic. Let μ_0 be such that $\mu_0 P^n$ does not converge to μ . Then $\mu_0 \text{env}(P)^n$ does not converge either; see (4.1). To summarize, we have proved that P nonergodic implies that $\text{env}(P)$ is nonergodic. This completes the proof.

The converse of Proposition 4.3 is not true and counterexamples will be given in Section 4.3.3.

4.2.2. *Construction of an update function for the EPCA.* Let us define the update function

$$\tilde{\phi}: \mathcal{B}^E \times [0, 1]^E \rightarrow \mathcal{B}^E$$

of the PCA $\text{env}(P)$ by

$$\tilde{\phi}(y, r)_k = \begin{cases} \mathbf{0} & \text{if } 0 \leq r_k < \text{env}(f)((y_{k+v})_{v \in V})(\mathbf{0}), \\ \mathbf{1} & \text{if } 1 - \text{env}(f)((y_{k+v})_{v \in V})(\mathbf{1}) \leq r_k \leq 1, \\ \mathbf{?} & \text{otherwise.} \end{cases} \tag{4.2}$$

The value of $\tilde{\phi}(y, r)_k$ as a function of r_k can thus be represented as in Figure 4.

For a PCA of neighborhood $V = \{0, 1\}$, we represent in Figure 5 the construction of the updates of the EPCA when the value of the neighborhood is $\mathbf{0?}$.




FIGURE 4.




FIGURE 5.

Let ϕ be the natural update function for the PCA P defined as in (2.1). Observe that the function $\tilde{\phi}$ coincides with ϕ on configurations which do not contain the letter ‘?’ . Furthermore, we have for all $r \in [0, 1]^E$, $x \in \mathcal{A}^E$, and $y \in \mathcal{B}^E$,

$$x \in y \implies \phi(x, r) \in \tilde{\phi}(y, r). \tag{4.3}$$

4.3. Perfect sampling using EPCAs

We propose two perfect sampling algorithms: one for a finite and another for an infinite number of cells. We show that in both cases, the algorithms stop almost surely if and only if the EPCA is ergodic (Theorem 4.1). The ergodicity of the EPCA implies the ergodicity of the PCA but the converse is not true. We provide a counterexample for each finite and infinite case (Section 4.3.3). We also give conditions of ergodicity of the EPCA (Propositions 4.4 and 4.5).

4.3.1. *Algorithms. Finite set of cells.* The idea is to consider only one trajectory of the EPCA—the trajectory that starts from the initial configuration $?^E$ (coding the set of all configurations of the PCA). The algorithm stops when at time 0 this trajectory hits the set \mathcal{A}^E

Algorithm 3. (*Perfect sampling using the EPCA for a finite set of cells.*) *Data:* The pre-computed update function ϕ . A family $(r_k^{-n})_{(k,n) \in E \times \mathbb{N}}$ of i.i.d. RVs with uniform distribution in $[0, 1]$.

```

begin
  t = 1;
  repeat
    c =  $?^E$ ;
    for j = -t to -1 do
      c =  $\tilde{\phi}(c, (r_i^j)_{i \in E})$ 
    end
    t = t + 1
  
```

```

until  $c \in \mathcal{A}^E$ ;
return  $c$ 
end
    
```

Infinite set of cells. Once again, we consider only one trajectory of the EPCA.

Algorithm 4. (*Perfect sampling using the EPCA for an infinite set of cells.*) *Data:* the pre-computed update function ϕ . A family $(r_k^{-n})_{(k,n) \in E \times \mathbb{N}}$ of i.i.d. RVs with uniform distribution in $[0, 1]$. A finite subset K of E .

```

begin
   $V_0(K) = K$ ;
   $t = 1$ ;
  repeat
     $V_{-t}(K) = V + V_{-t+1}(K)$ ;
     $c = \mathcal{P}^{V_{-t}(K)}$ ;
    for  $j = -t$  to  $-1$  do
       $c = \tilde{\phi}_j(c, (r_i^j)_{i \in V_j(K)}) \in \mathcal{B}^{V_{j+1}(K)}$ 
    end
     $t = t + 1$ 
  until  $c \in \mathcal{A}^K$ ;
  return  $c$ 
end
    
```

Theorem 4.1. *Algorithms 3 and 4 stop almost surely if and only if the EPCA is ergodic. In that case, the output of the algorithm is distributed according to the unique invariant measure of the PCA.*

Proof. The argument is the same in the finite and infinite cases. We give it for the finite case. Assume first that Algorithm 3 stops almost surely. By construction, it implies that, for all μ_0 , the measure $\mu_0 \text{env}(P)^n$ is asymptotically supported by \mathcal{A}^E . Therefore, we can strengthen the result in Proposition 4.3: the invariant measures of $\text{env}(P)$ coincide with the invariant measures of P . In that case, $\text{env}(P)$ is ergodic if and only if P is ergodic. Using (4.3), the halting of Algorithm 3 implies the halting of Algorithm 1. Furthermore, if we use the same samples $(r_k^{-n})_{(k,n) \in E \times \mathbb{N}}$, Algorithms 1 and 3 and will have the same output. According to Proposition 4.1, this output is distributed according to the unique invariant measure of P . In particular, P is ergodic. So $\text{env}(P)$ is ergodic.

Assume now that the EPCA is ergodic. The unique invariant measure π of $\text{env}(P)$ has to be supported by \mathcal{A}^E . Also, by ergodicity, we have $\delta_{\mathcal{P}^E} \text{env}(P)^n \xrightarrow{w} \pi$. This means precisely that Algorithm 3 stops almost surely.

4.3.2. *Criteria of ergodicity for the EPCA. Finite set of cells.* In the next proposition, we give a necessary and sufficient condition for the EPCA to be ergodic. In particular, this condition is satisfied if the PCA has positive rates (see the introduction).

Proposition 4.4. *The EPCA $\text{env}(P)$ is ergodic if and only if $\text{env}(f)(\mathcal{P}^V)(\mathcal{P}) < 1$. This condition can also be written as*

$$\min_{x \in \mathcal{A}^V} f(x)(0) + \min_{x \in \mathcal{A}^V} f(x)(1) > 0. \tag{4.4}$$

Proof. If $\text{env}(f)(\mathcal{P}^V)(\mathcal{P}) = 1$ then, for almost any $r \in [0, 1]^E$, we have $\tilde{\phi}(\mathcal{P}^E, r) = \mathcal{P}^E$, so that at each step of the algorithm, the value of c is \mathcal{P}^E with probability 1.

Conversely, if we assume for example that $p = \min_{x \in \mathcal{A}^V} f(x)(0) > 0$, then, for any configuration $d \in \mathcal{B}^E$, the probability of having $\tilde{\phi}(x, r) = \mathbf{0}^E$ is greater than $p^{|E|}$, so that the algorithm stops almost surely, and the expectation of the running time can be roughly bounded by $1/p^{|E|}$. This completes the proof.

Infinite set of cells. For an infinite set of cells, the situation is more complex. The condition of Proposition 4.4 is not sufficient to ensure the ergodicity of the EPCA. A counterexample is given in Section 4.3.3. First, we propose a rough sufficient condition of convergence for Algorithm 4.

Proposition 4.5. *Let $\alpha^* \in (0, 1)$ be the critical probability of the percolation PCA with neighborhood V ; see Examples 2.2 and 2.4. The EPCA $\text{env}(P)$ is ergodic if*

$$\text{env}(f)(?)^V(?) < \alpha^* \tag{4.5}$$

and nonergodic if

$$\min_{x \in \mathcal{B}^V - \mathcal{A}^V} \text{env}(f)(x)(?) > \alpha^*. \tag{4.6}$$

Proof. Recall that $\mathcal{B} = \{\mathbf{0}, \mathbf{1}, ?\}$. Define $\mathcal{C} = \{d, ?\}$, with $d = \{\mathbf{0}, \mathbf{1}\}$. A word over \mathcal{C} is interpreted as a set of words over \mathcal{B} , for instance, $d? = \{\mathbf{0?}, \mathbf{1?}\}$. The symbol d stands for the determined letter, as opposed to $?$ which represents an unknown letter.

We define a new PCA Q on the alphabet \mathcal{C} , with the same neighborhood V as P and $\text{env}(P)$, and with the transition function $g: \mathcal{C}^V \rightarrow \mathcal{M}(\mathcal{C})$ defined by

$$g(d^V) = \delta_d, \quad g(u) = \alpha \delta_? + (1 - \alpha) \delta_d \quad \text{for all } u \in \mathcal{C}^V - \{d^V\},$$

where $\alpha = \max_{x \in \mathcal{B}^V} \text{env}(f)(x)(?) = \text{env}(f)(?)^V(?)$.

Observe that δ_{d^E} is an invariant measure of Q . Recall that $\tilde{\phi}$ is an update function of $\text{env}(P)$; see (4.2). Given the way Q is defined, we can construct an update function ϕ_Q of Q such that, for all $x \in \mathcal{B}^E$, all $y \in \mathcal{C}^E$, and all $r \in [0, 1]^E$,

$$x \in y \implies \tilde{\phi}(x, r) \in \phi_Q(y, r). \tag{4.7}$$

In particular, assume that Q is ergodic. Then $\delta_{?^E} Q^n \xrightarrow{w} \delta_{d^E}$. Using (4.7) implies that Algorithm 4 stops almost surely, and $\text{env}(P)$ is ergodic according to Theorem 4.1. To summarize, the ergodicity of Q implies the ergodicity of $\text{env}(P)$.

Observe that the PCA Q is a percolation PCA as defined in Example 2.2 (here, d plays the role of 0 and $?$ plays the role of 1). Let $\alpha^* \in (0, 1)$ be the critical probability of the percolation PCA with neighborhood V ; see Example 2.4. For $\alpha < \alpha^*$, the percolation PCA Q is ergodic. This completes the proof of (4.5).

Define a PCA R on the alphabet \mathcal{C} , with neighborhood V , and with the transition function

$$h(d^V) = \delta_d, \quad h(u) = \beta \delta_? + (1 - \beta) \delta_d \quad \text{for all } u \in \mathcal{C}^V - \{d^V\},$$

where $\beta = \min_{x \in \mathcal{B}^V - \mathcal{A}^V} \text{env}(f)(x)(?)$. Given the way R is defined, we can construct an update function ϕ_R of R such that for all $x \in \mathcal{B}^E$, $y \in \mathcal{C}^E$, $r \in [0, 1]^E$, and $k \in E$,

$$[x \in y, \phi_R(y, r)_k = ?] \implies \tilde{\phi}(x, r)_k = ?.$$

Therefore, the ergodicity of $\text{env}(P)$ implies the ergodicity of R . Equivalently, the nonergodicity of R implies the nonergodicity of $\text{env}(P)$. Observe that the PCA R is a percolation PCA. Therefore, for $\beta > \alpha^*$, the percolation PCA R is nonergodic. This completes the proof of (4.6).

4.3.3. *Counterexamples.* Recall Proposition 4.3: [EPCA ergodic] implies that [PCA ergodic]. We now show that the converse is not true.

Example 4.1. Consider the PCA with alphabet $\mathcal{A} = \{0, 1\}$, neighborhood $V = \{-1, 0, 1\}$, set of cells $E = \mathbb{Z}/n\mathbb{Z}$, and transition function

$$f(x, y, z) = \begin{cases} \delta_{1-y} & \text{if } xyz \in \{101, 010\}, \\ \alpha\delta_y + (1 - \alpha)\delta_{1-y} & \text{otherwise,} \end{cases}$$

for a parameter $\alpha \in (0, 1)$. This is the PCA majority studied in Section 5. For odd n , it is easy to check that the PCA is ergodic. However, the associated EPCA satisfies $\text{env}(f)(\delta_\varphi) = \delta_\varphi$. According to Proposition 4.4, the EPCA is not ergodic.

Example 4.2. Consider the PCA of Example 2.1. This PCA has positive rates, in particular, it satisfies (4.4). So the EPCA is ergodic on a finite set of cells. Now let the set of cells be \mathbb{Z} . The PCA is ergodic for $\varepsilon \in (0, 1)$; see Example 2.3. Consider the associated EPCA $\text{env}(P)$. Assume for instance that $\varepsilon \in (0, \frac{1}{2})$. We have

$$\text{env}(f)(u) = \begin{cases} f(u) & \text{if } u \in \{\mathbf{0}, \mathbf{1}\}^V, \\ \varepsilon\delta_{\mathbf{0}} + \varepsilon\delta_{\mathbf{1}} + (1 - 2\varepsilon)\delta_\varphi & \text{otherwise.} \end{cases}$$

By applying Proposition 4.5, $\text{env}(P)$ is nonergodic if $1 - 2\varepsilon > \alpha^*$.

4.4. Decay of correlations

In what follows, the set of cells is $E = \mathbb{Z}^d$, $d \geq 1$. It is easy to prove that the invariant measure of an ergodic PCA is shift invariant. Using the coupling-from-the-past tool, we give conditions for the invariant measure of an ergodic PCA to be shift mixing.

Definition 4.2. A measure μ on $X = \mathcal{A}^{\mathbb{Z}^d}$ is *shift mixing* if, for any nontrivial translation shift τ of \mathbb{Z}^d and any cylinders U, V of X ,

$$\lim_{n \rightarrow +\infty} \mu(U \cap \tau^{-n}(V)) = \mu(U)\mu(V).$$

The proof of the following proposition is inspired by the proof of the validity of the coupling-from-the-past method (see [21] or [13]).

Proposition 4.6. *If Algorithm 2 stops almost surely then the unique invariant measure of the PCA is shift mixing. It is in particular the case under condition (4.5).*

Proof. Assume that P is an ergodic PCA, and denote by π its unique invariant measure. Let K and L be two finite subsets of E , and denote by x_K and y_L some cylinders corresponding to these subsets. Since the perfect sampling algorithm stops almost surely for each $\varepsilon > 0$, there exists an integer t_ε for which, with probability greater than $1 - \varepsilon$, the algorithm stops before reaching the time $-t_\varepsilon$ when it is run for the set of cells K or for the set of cells L . If $n \in \mathbb{N}^d$ is large enough, we have $V_{-t_\varepsilon}(K) \cap V_{-t_\varepsilon}(\tau^{-n}(L)) = \emptyset$.

Let Z be the output of the algorithm if it is asked to sample the marginals of π corresponding to the cells of $K \cup \tau^{-n}(L)$. As illustrated in Figure 6, imagine running the PCA from time $-t_\varepsilon$ and consider the set of cells $V_{-t_\varepsilon}(K) \cup V_{-t_\varepsilon}(\tau^{-n}(L))$ up to time 0, using the same update variables as those used to get Z . Choose the initial condition at time $-t_\varepsilon$ independently on $V_{-t_\varepsilon}(K)$ and $V_{-t_\varepsilon}(\tau^{-n}(L))$, and according to the relevant marginals of π . Let X and Y respectively be




FIGURE 6: Illustration of the proof of Proposition 4.6.

the output at time 0 on the set of cells K and $\tau^{-n}(L)$. Observe that X and Y are distributed according to the marginals of π . Furthermore, X and Y are independent since the dependence cones of K and $\tau^{-n}(L)$ originating at time $-t_\epsilon$ are disjoint. We therefore obtain

$$\begin{aligned} \pi(x_K \cap \tau^{-n}(y_L)) - \pi(x_K)\pi(y_L) &= \mathbb{P}(Z_K = x_K, Z_{\tau^{-n}(L)} = y_L) - \mathbb{P}(X = x_K)\mathbb{P}(Y = y_L) \\ &= \mathbb{P}(Z_K = x_K, Z_{\tau^{-n}(L)} = y_L) - \mathbb{P}(X = x_K, Y = y_L) \\ &\leq \mathbb{P}((Z_K, Z_{\tau^{-n}(L)}) = (x_K, y_L) \text{ and } (X, Y) \neq (x_K, y_L)) \\ &\leq \mathbb{P}((Z_K, Z_{\tau^{-n}(L)}) \neq (X, Y)) \\ &\leq 2\epsilon. \end{aligned}$$

In the same way, we get $\pi(x_K)\pi(y_L) - \pi(x_K \cap \tau^{-n}(y_L)) \leq 2\epsilon$. This completes the proof.

In Proposition 4.6, the coupling-from-the-past method is not used as a sampling tool but as a way to get theoretical results. Knowing if there exists an ergodic PCA having an invariant measure which is not shift mixing is an open question (see [5] for details).

4.5. Extensions

In a PCA, the dynamic is homogeneous in space. It is possible to get rid of this characteristic by defining nonhomogeneous PCAs, for which the neighborhood and the transition function depend on the position of the cell. The definition below is to be compared with Definition 2.1. The configuration space $X = \mathcal{A}^E$ is unchanged.

Definition 4.3. For each $k \in E$, denote by $V_k \subset E$ the (finite) neighborhood of the cell k , and by $f_k: \mathcal{A}^{V_k} \rightarrow \mathcal{M}(\mathcal{A})$ the transition function associated to k . Set $\mathcal{V}(K) = \bigcup_{k \in K} V_k$. The *non-homogeneous PCA (NH-PCA)* of transition functions $(f_k)_{k \in E}$ is the application $P: \mathcal{M}(X) \rightarrow \mathcal{M}(X)$, $\mu \mapsto \mu P$, defined on cylinders by

$$\mu P(y_K) = \sum_{x_{\mathcal{V}(K)} \in \mathcal{C}(\mathcal{V}(K))} \mu(x_{\mathcal{V}(K)}) \prod_{k \in K} f_k((x_v)_{v \in V_k})(y_k).$$

Observe that it is not necessary for E to be equipped with a semigroup structure anymore. We use this below to define the finite restriction of a PCA.

It is quite straightforward to adapt the coupling-from-the-past algorithms to NH-PCAs. More precisely, given an NH-PCA, we define the associated NH-EPCA by considering Definition 4.1 and replacing V and $\text{env}(f)$ by V_k and $\text{env}(f)_k$ for each $k \in E$. The algorithms of Section 4.1 and 4.3.1 are then unchanged, and Proposition 4.3 and Theorem 4.1 still hold in the non-homogeneous setting.

In Section 5 we study the PCA majority by approximating it by a sequence of NH-PCAs. Let us explain the construction in a general setting.

Let P be a PCA on the infinite set of cells E , with neighborhood V and transition function $f: \mathcal{A}^V \rightarrow \mathcal{M}(\mathcal{A})$. Let D be a finite subset of E . Define

$$\bar{V}(D) = (V + D) \cup D, \quad B(D) = \bar{V}(D) - D.$$

The set $B(D)$ is the *boundary* of the domain D . Fix a probability measure ν on \mathcal{A} . The *restriction* of P associated with ν and D is the NH-PCA $P(\nu, D)$ with set of cells $\bar{V}(D)$ and neighborhoods

$$V_u = \{u\} + V \quad \text{for all } u \in D, \quad V_u = \emptyset \quad \text{for all } u \in B(D),$$

and transition functions

$$f_u = f \quad \text{for all } u \in D, \quad f_u(\cdot) = \nu \quad \text{for all } u \in B(D).$$

In words, the boundary cells are i.i.d. of law ν and the cells of D are updated according to P .

If μ is a probability measure on \mathcal{A}^S , where S is a finite subset of E , we define its extension $\tilde{\mu}$ on \mathcal{A}^E by setting, for a fixed letter $a \in \mathcal{A}$ and for all $x \in \mathcal{A}^E$,

$$\tilde{\mu}(x) = \begin{cases} \mu((x_k)_{k \in S}) & \text{if, for all } i \in E - S, x_i = a, \\ 0 & \text{otherwise.} \end{cases}$$

Lemma 4.1. *Let $(D_i)_{i \in \mathbb{N}}$ be an increasing sequence of finite domains $D_i \subset E$ for which $\bigcup_{i \in \mathbb{N}} D_i = E$. Let $(\nu_i)_{i \in \mathbb{N}}$ be a sequence of probability measures on \mathcal{A} . For each i , let μ_i be an invariant measure of $P(\nu_i, D_i)$. Any accumulation point of the sequence $(\tilde{\mu}_i)_{i \in \mathbb{N}}$ is an invariant measure of the original PCA P defined on E .*

Proof. Upon extracting a subsequence, we may assume that $(\tilde{\mu}_j)_{j \in \mathbb{N}}$ converges to $\tilde{\mu} \in \mathcal{M}(X)$. We need to prove that, for any cylinder $y_K \in \mathcal{C}(K)$, we have $\tilde{\mu}P(y_K) = \tilde{\mu}(y_K)$.

By definition, $\mu_j P(\nu_j, D_j) = \mu_j$. Let the subset K of E and the cylinder $y_K \in \mathcal{C}(K)$ be fixed. If j is large enough, we have $K \subset D_j$ and $\bar{V}(K) \subset D_j$. So that $\mu_j(y_K) = \tilde{\mu}_j(y_K)$ and $P(\nu_j, D_j)$ and P coincide on K . We deduce that $\tilde{\mu}_j P(y_K) = \tilde{\mu}_j(y_K)$. By taking the limit on both sides, we get $\tilde{\mu}P(y_K) = \tilde{\mu}(y_K)$. This completes the proof.

Alphabet with more than two elements. The EPCA and the associated algorithms have been defined on a two-letter alphabet. It is possible to extend the approach to a general finite alphabet.

Let \mathcal{A} be the finite alphabet. Let P be a PCA with set of cells E , neighborhood V , and transition function $f: \mathcal{A}^V \rightarrow \mathcal{M}(\mathcal{A})$.

Consider the alphabet $\mathcal{B} = 2^{\mathcal{A}} - \{\emptyset\}$, that is, the set of nonempty subsets of \mathcal{A} . A word over \mathcal{B} is viewed as a set of words over \mathcal{A} .

The EPCA $\text{env}(P)$ associated with P is a PCA on the alphabet \mathcal{B} with neighborhood V and transition function $\text{env}(f)$ that we now determine. Let us fix some $v \in \mathcal{B}^V$ and define $\rho_S = \min_{u \in v} f(u)(\{S\})$. For a single letter $a \in \mathcal{A}$, we still want to have $\text{env}(f)(v)(\{a\}) = \min_{u \in v} f(u)(a) = \rho_a$. Now, let us consider some $b \in \mathcal{A}$, $b \neq a$. We set $\text{env}(f)(v)(\{a, b\}) = \rho_{a,b} - \rho_a - \rho_b$, and so on.

By the inclusion–exclusion principle, we finally obtain the following formula for the transition function $\text{env}(f)$: for all $v \in \mathcal{B}^V$ and all $y \in \mathcal{B}$,

$$\text{env}(f)(v)(y) = \sum_{x \subset y} (-1)^{|y|-|x|} \min_{u \in v} f(u)(x).$$

For instance, $\text{env}(f)(v)(\{0, 1, 2\}) = \rho_{0,1,2} - \rho_{1,2} - \rho_{0,2} - \rho_{0,1} + \rho_0 + \rho_1 + \rho_2$.

The algorithms of Section 4.3 are unchanged. Observe however that the construction of an update function is not as natural as in the two-letter alphabet case.

5. The majority-flip PCA: a case study

The majority-flip PCA, which we call for simplicity the majority PCA, is one of the simplest examples of a PCA whose behaviour is not well understood. Therefore, it provides a good case study for the sampling algorithms of Section 4.

5.1. Definition of the majority PCA

Given $0 < \alpha < 1$, the PCA Majority(α), or simply Majority, is the PCA on the alphabet $\mathcal{A}=\{0, 1\}$, with set of cells $E=\mathbb{Z}$ (or $\mathbb{Z}/n\mathbb{Z}$), neighborhood $V=\{-1, 0, 1\}$, and transition function

$$f(x, y, z) = \alpha\delta_{\text{maj}(x,y,z)} + (1 - \alpha)\delta_{1-y},$$

where $\text{maj}: \mathcal{A}^3 \rightarrow \mathcal{A}$ is the *majority function*: the value of $\text{maj}(x, y, z)$ is 0 or 1 depending on whether there are two or three 0s, or two or three 1s, respectively, in the sequence x, y, z . The transition function of PCA Majority(α) can thus be represented as in Figure 7. It consists of choosing independently for each cell to apply the elementary rule 232 (with probability α) or to flip the value of the cell.

The PCA Minority(α) has also been studied (see [23]). It is defined by the transition function $g(x, y, z) = f(1 - x, 1 - y, 1 - z)$.

Let $x = (01)^\mathbb{Z} \in \{0, 1\}^\mathbb{Z}$ be defined by $x_{2n} = 0$ and $x_{2n+1} = 1$ for all $n \in \mathbb{Z}$. The configuration $(10)^\mathbb{Z}$ is defined similarly. Consider the probability measure

$$\mu = \frac{\delta_{(01)^\mathbb{Z}} + \delta_{(10)^\mathbb{Z}}}{2}. \tag{5.1}$$

Clearly, μ is an invariant measure for the PCA majority. The question is whether there exists other invariant measures.

To get some insight into this question, consider the PCA majority on the set of cells $\mathbb{Z}_n = \mathbb{Z}/n\mathbb{Z}$. This PCA has two completely different behaviors depending on the parity of n . Indeed, a simple analysis of the structure of the transition matrix shows that the Markov chain has a unique invariant measure which is $(\delta_{(01)^{\mathbb{Z}_n}} + \delta_{(10)^{\mathbb{Z}_n}})/2$ if n is even, and which is supported on $\{0, 1\}^{\mathbb{Z}_n}$ if n is odd.

Let us return to the PCA majority on \mathbb{Z} . The invariant measure μ in (5.1) can be viewed as the ‘limit’ over n of the invariant measures of the PCA on \mathbb{Z}_{2n} . What about the ‘limits’ of the invariant measures of the PCA on \mathbb{Z}_{2n+1} ? Do they define other invariant measures for the PCA on \mathbb{Z} ?




FIGURE 7: The transition function of the PCA majority.

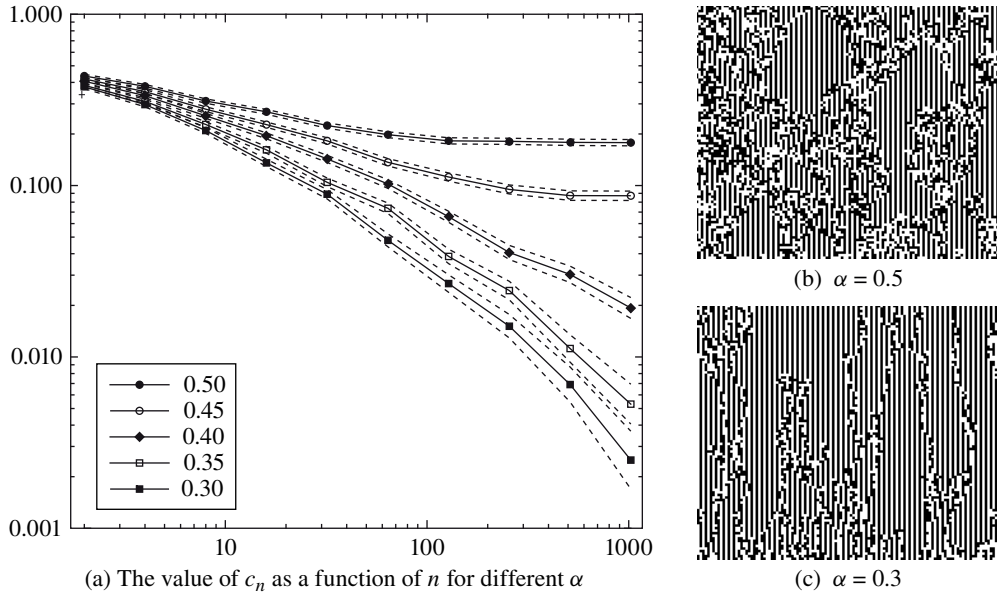


FIGURE 8: Experimental study of Majority(α) (the configurations at odd times only are represented on the space–time diagrams).

One of the motivations of our work on perfect sampling algorithms for PCAs was to test the following conjecture, which is inspired by the observations made in [22] and [23] on a PCA equivalent to majority. This conjecture concerns the existence of a ‘phase transition’ phenomenon for the PCA majority.

Conjecture 5.1. *There exists $\alpha_c \in (0, 1)$ such that Majority(α) has a unique invariant measure if $\alpha < \alpha_c$, and several invariant measures if $\alpha > \alpha_c$.*

We tried to get some numerical evidence for Conjecture 5.1 using the perfect sampling tools developed in the previous section. To study the PCA majority experimentally, a first idea would be to consider the same PCA on the set of cells \mathbb{Z}_n , n odd. This does not work well. First, due to the state space explosion, computing exactly the invariant measure is possible only for small values (we did it up to $n = 9$ using MAPLE®). Second, the algorithms of Section 4 cannot be applied since the EPCA is not ergodic.

Instead, we use approximations of the PCA by NH-PCAs on a finite subset of cells, the methodology sketched in Section 4.5. Again, computing exactly the invariant measure is impossible except for very small windows. But now the sampling algorithms become effective.

Let P be the PCA majority. Set $D_n = \{-n, \dots, n\}$, and let ν be the uniform measure on $\{0, 1\}$. Consider the NH-PCA $P(\nu, D_n)$. Let μ_n be the unique invariant measure of $P(\nu, D_n)$. We are interested in the quantity

$$c_n = \mu_n\{x \in X \mid x_0 = x_1 = 0\} + \mu_n\{x \in X \mid x_0 = x_1 = 1\}.$$

Indeed, by application of Lemma 4.1, if $\limsup_n c_n > 0$ then there exists a nontrivial invariant measure for the PCA majority on \mathbb{Z} .

Now the NH-EPCA is ergodic, so the sampling algorithms of Section 4 can be used. We were able to run the algorithms up to a window size of $n = 1024$ before running into a timeout. The experimental results appear in Figure 8, with a logarithmic scale. We ran the sampling

algorithms 10^4 times. We show in the figure the confidence intervals calculated with the Wilson score test at 95%.

It is reasonable to believe that the top two curves in Figure 8 do not converge to 0 while the bottom three converge to 0. This is consistent with the visual impression of space–time diagrams. It reinforces Conjecture 5.1 with a possible phase transition between 0.4 and 0.45.

Acknowledgements

We used the applet FiatLux developed by N. Fatès and available on his website (LORIA, INRIA Lorraine) to draw the space–time diagrams of Figures 1 and 8.

This work was partially supported by the ANR project MAGNUM (ANR-2010-BLAN-0204).

References

- [1] BOUSQUET-MÉLOU, M. (1998). New enumerative results on two-dimensional directed animals. *Discrete Math.* **180**, 73–106.
- [2] BUŠIĆ, A., GAUJAL, B. AND VINCENT, J.-M. (2008). Perfect simulation and non-monotone Markovian systems. In *Proc. 3rd Internat. Conf. Performance Evaluation Methodol. Tools*, Athens, 10pp.
- [3] BUŠIĆ, A., MAIRESSE, J. AND MARCOVICI, I. (2011). Probabilistic cellular automata, invariant measures, and perfect sampling. In *28th Internat. Symp. Theoret. Aspects Comput. Sci. (STACS'11)*, Leibniz-Zentrum, Dagstuhl, pp. 296–307.
- [4] CHASSAING, P. AND MAIRESSE, J. (2011). A non-ergodic probabilistic cellular automaton with a unique invariant measure. *Stoch. Process. Appl.* **121**, 2474–2487.
- [5] COLETTI, C. F. AND TISSEUR, P. (2010). Invariant measures and decay of correlations for a class of ergodic probabilistic cellular automata. *J. Statist. Phys.* **140**, 103–121.
- [6] CULIK, K., II, PACHL, J. AND YU, S. (1989). On the limit sets of cellular automata. *SIAM J. Comput.* **18**, 831–842.
- [7] DHAR, D. (1983). Exact solution of a directed-site animals-enumeration problem in three dimensions. *Phys. Rev. Lett.* **51**, 853–856.
- [8] FATÈS, N. (2011). Stochastic cellular automata solve the density classification problem with an arbitrary precision. In *28th Internat. Symp. Theoret. Aspects Comput. Sci. (STACS'11)*, Leibniz-Zentrum, Dagstuhl, pp. 284–295.
- [9] FATÈS, N., REGNAULT, D., SCHABANEL, N. AND THIERRY, E. (2006). Asynchronous behavior of double-quiescent elementary cellular automata. In *LATIN 2006: Theoretical Informatics (Lecture Notes Comput. Sci. 3887)*, Springer, Berlin, pp. 455–466.
- [10] GÁCS, P. (2001). Reliable cellular automata with self-organization. *J. Statist. Phys.* **103**, 45–267.
- [11] GRAY, L. F. (2001). A reader's guide to P. Gács's 'positive rates' paper: 'Reliable cellular automata with self-organization'. *J. Statist. Phys.* **103**, 1–44.
- [12] GUILLON, P. AND RICHARD, G. (2008). Nilpotency and limit sets of cellular automata. In *Mathematical Foundations of Computer Science (Lecture Notes Comput. Sci. 5162)*, Springer, Berlin, pp. 375–386.
- [13] HÄGGSTRÖM, O. (2002). *Finite Markov Chains and Algorithmic Applications* (London Math. Soc. Student Texts **52**). Cambridge University Press.
- [14] HÄGGSTRÖM, O. AND NELANDER, K. (1998). Exact sampling from anti-monotone systems. *Statist. Neerlandica* **52**, 360–380.
- [15] HEDLUND, G. (1969). Endomorphisms and automorphisms of the shift dynamical system. *Math. Systems Theory* **3**, 320–375.
- [16] HUBER, M. (2004). Perfect sampling using bounding chains. *Ann. Appl. Probab.* **14**, 734–753.
- [17] KARI, J. (1992). The nilpotency problem of one-dimensional cellular automata. *SIAM J. Comput.* **21**, 571–586.
- [18] KOZMA, R. *et al.* (2004). Neuropercolation: a random cellular automata approach to spatio-temporal neurodynamics. In *Cellular Automata (Lecture Notes Comput. Sci. 3305)*, Springer, Berlin, pp. 435–443.
- [19] LE BORGNE, Y. AND MARCKERT, J.-F. (2007). Directed animals and gas models revisited. *Electron. J. Combin.* **14**, R71.
- [20] LIGGETT, T. (1985). *Interacting Particle Systems*. Springer, New York.
- [21] PROPP, J. G. AND WILSON, D. B. (1996). Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures Algorithms* **9**, 223–252.
- [22] REGNAULT, D. (2008). Directed percolation arising in stochastic cellular automata analysis. In *Mathematical Foundations of Computer Science 2008 (Lecture Notes Comput. Sci. 5162)*, Springer, Berlin, pp. 563–574.

- [23] SCHABANEL, N. (2009). *Habilitation à diriger des recherches: systèmes complexes & algorithmes*. Res. Rep. Université Paris Diderot - Paris 7.
- [24] TOOM, A. (1995). Cellular automata with errors: problems for students of probability. In *Topics in Contemporary Probability and Its Applications*, CRC, Boca Raton, FL, pp. 117–157.
- [25] TOOM, A. (2000). Algorithmical unsolvability of the ergodicity problem for binary cellular automata. *Markov Process. Relat. Fields* **6**, 569–577.
- [26] TOOM, A. (2001). *Contours, Convex Sets, and Cellular Automata*. IMPA, Rio de Janeiro.
- [27] TOOM, A. *et al.* (1990). Discrete local Markov systems. In *Stochastic Cellular Systems: Ergodicity, Memory, Morphogenesis*, eds R. Dobrushin, V. Kryukov and A. Toom. Manchester University Press, pp. 1–182.
- [28] VAN DEN BERG, J. AND STEIF, J. E. (1999). On the existence and nonexistence of finitary codings for a class of random fields. *Ann. Prob.* **27**, 1501–1522.