Probing robustness of cellular automata through variations of asynchronous updating

Olivier Bouré, Nazim A. Fatès, Vincent Chevrier

Probing robustness of cellular automata through variations of asynchronous updating

Olivier Bouré, Nazim Fatès, Vincent Chevrier

Université de Lorraine – INRIA Nancy–Grand-Est – LORIA
`{olivier.boure,nazim.fates,vincent.chevrier}@loria.fr`

Abstract

Typically viewed as a deterministic model of spatial computing, cellular automata are here considered as a collective system subject to the noise inherent to natural computing. The classical updating scheme is replaced by stochastic versions which either randomly update cells or disrupt the cell-to-cell transmission of information. We then use the novel updating schemes to probe the behaviour of Elementary Cellular Automata, and observe a wide variety of results. We study these behaviours in the scope of macroscopic statistical phenomena and microscopic analysis. Finally, we discuss the possibility to use updating schemes to probe the robustness of complex systems.

Keywords: Asynchronous cellular automata, robustness, discrete dynamical systems, phase transitions, directed percolation.

1   Introduction

Cellular automata (CA) are known as a model of computation based on the parallel evolution of cells according to their neighbourhood state. Originating from the study of self-replicating systems as introduced by von Neumann (1966), they are considered as an alternative to sequential computing models. Indeed, their spatially-extended and yet simple structure make them a suitable model for parallel computing and simulations of natural phenomena.

In their original form they are deterministic and synchronously evolved, that is, all components of the system are updated simultaneously. However, as natural computing supposes the use of non-classical methods involving a number of interacting components, the determinism of the model is somewhat incompatible with the presence of noise. This raises the question of the robustness of cellular automata, that is, the stability of the behaviour despite external disturbances, such as an asynchronous update or errors in transitions or interactions between cells.

This question was initially discussed by Ingerson and Buvel (1984) who "wanted to estimate how much of the behaviour of cellular automata comes from synchronous modeling and how much is intrinsic to the iteration process". For that purpose, they questioned the perfect synchrony hypothesis by randomising the transition function of Elementary Cellular Automata, and observed behavioural changes. From their observations they concluded that the updating scheme plays a fundamental role in the behaviour of the system. Since then, authors have tackled the question of whether cellular automata are robust to non-synchronous updating, either by keeping the same individual rule or adding adapted constructs (Peper et al. 2002). For instance, Grilo and Correia (2011) investigated behaviour changes for asynchronous symmetric 2-player evolutionary games, such as the Prisoner's Dilemma. Similarly, models of cellular automata have been studied under asynchronous updating, such as Elementary Cellular Automata (Fatès and Morvan 2005, Regnault 2008) or the Game of Life (Blok and Bergersen 1999, Fatès 2010) and revealed qualitatively different behaviours.

The definitions of asynchronous cellular automata used so far have mainly focused on either randomising the transition function or updating cells sequentially (Bandini et al. 2010a). But as asynchronism is first defined as a negation of synchronicity, several interpretations remain possible. Our approach, which shares the main ideas from the fields of spatially-extended computing (Gouaïch et al. 2005, Ackley and Williams 2011) and autonomous agents (Bandini et al. 2010b), intends to consider cellular automata as a collective system where cells are autonomous components and their interactions are subject to perturbations. Instead of questioning whether cells are updated at each time step, we represent explicitly the perception by the cells of their neighbourhood, and consider failures in the cell-to-cell transmission of information.

Questioning the perfect synchrony hypothesis. Let us first consider an illustrative example: two planes are flying in opposing directions and at successive time steps must decide to go *up* or *down* in order to avoid collision. The decision function, the same for both planes, is deterministic: if the opposing plane goes down, then go up and vice versa. Here three types of updating schemes are considered (see Fig. 1):

(a) The perfect synchrony assumes that planes update their state simultaneously at each time step. An initial configuration with both planes going in the same direction inevitably results in a periodic cycle and a crash.

(b) However, when planes update their state with a given probability, the system finds a stable situation after a few time steps. This updating scheme is called *asynchronous update*.

(c) This time, planes do update their state at the same time, but the updated state of one plane is not always transmitted to the other. When one plane fails to transmit its new state, the system stabilises. We call this novel type of updating scheme *asynchronous information transmission*.

Figure 1: Different updating schemes applied to a rudimentary collision avoidance system. Black planes represent actual states, and in the case (c), white planes denotes how they are perceived by the opposite plane.

As one can deduce from this simple example, both asynchronous updating schemes alter the long-term behaviour of the system, as they directly prevent the stability of the crash scenario of the synchronous case. However, their intrinsic difference also suggests that there may exist models which would react differently to each of them. We intend to use these updating schemes to *probe* cellular automata, that is, to study how each of the different updating schemes influences the behaviour of a set of cellular automata rules.

Asynchronous information transmission in cellular automata. To study the effects of asynchronous information transmission on cellular automata, we propose a novel updating scheme, extending two features of the classical definition: (1) we extend the state space to take into account the internal state of the cell as well as how it is perceived by the neighbourhood and (2) we describe the updating process of the system in the frame of a cellular cycle, that is, the sequential application of two steps, (a) the local computation of the new state, and (b) the transmission of the updated state to the cells of the neighbourhood.

Extending a previous work (Bouré et al. 2011), we present a more general definition of asynchronous information transmission and a more complete study of Elementary Cellular Automata. In Sec. 2 we give a general formal definition of asynchronous information transmission for cellular automata. Then, an overview of our observations on asynchronous Elementary Cellular Automata is given in Sec. 3 and in Sec. 4, we present an analysis of phase transitions. Section 5 presents two case studies of a microscopic approach to the study of the effects on two models, which present different reaction depending on the asynchronism. Finally, we discuss our results in Sec. 6.

3

2 Asynchronous Cellular Automata

2.1 Definition of Cellular Automata

A cellular automaton is a discrete dynamical system defined by $(\mathcal{L}, Q, \mathcal{N}, f)$ where :

- $\mathcal{L} \subset \mathbb{Z}^d$ is the cellular space, where elements of $\mathcal{L}$ represent *cells*.

- Q is a finite set of states.

- $\mathcal{N} \subset \mathcal{L}$ is a finite set of vectors called the *neighbourhood*, which associates to a cell the set of its neighbouring cells. $\mathcal{N}$ and $\mathcal{L}$ are such that for all $c \in \mathcal{L}$ and for all $n \in \mathcal{N}$, the *neighbour* $c + n$ is in $\mathcal{L}$.

- f is the *local transition rule*, which defines the next state of a cell according to the state of this cell and the ones of its neighbours.

A *configuration* x^t represents the state of the automaton at time t; it is defined as a function $x^t : \mathcal{L} \to Q$ which maps each cell to a state. Classically, cellular automata are synchronously updated, meaning that at each time the local transition rule is applied simultaneously to all cells. The *global transition function* is therefore defined as $x^{t+1} = F(x^t)$, so that, for $\mathcal{N} = \{n_1, ..., n_k\}$:

$$\forall c \in \mathcal{L}, \ x^{t+1}(c) = f\left(x^t(c), x^t(c + n_1), ..., x^t(c + n_k)\right).$$

Without loss of generality, we assume that the neighbourhood $\mathcal{N}$ does *not* contain the cell itself. This hypothesis is necessary to explicitly represent the flow of information between a cell and its neighbours. Note that this does not restrict the expressiveness of f since the current state of a cell $x^t(c)$ is always a parameter of f, possibly not taken into account in the transition function.

To give a proper definition of asynchronous information transmission, we need to extend the classical framework of cellular automata.

1. The cell update is represented by a two-step *cell cycle* (see Fig. 2):

   - the *state update* step, where a cell changes its state according to the local transition function and its perception of the neighbourhood.

   - the *information transmission* step, where the cell transmits the updated state to its neighbours.

2. The state space Q must be extended to distinguish:

   - the *eigenstate*, that is, the actual, internal state of a cell in Q,

   - a vector of m *observation states*, that is, m elements of Q that represent how a cell is perceived by its neighbours.

In the frame of this extension of cellular automata, we define three types of updating schemes (see Fig. 3 for schematic representations):

4

Figure 2: Left: schematic representation of the classical update. Right: representation of the update as a cellular cycle, for a cell with two neighbours. A triplet $\{y_L, y, y_R\}$ represents a cell with eigenstate y, and its observable states, for each of its neighbours (y_L for left cell and y_R for right). The classical transition corresponds to $y' \leftarrow f(x, y, z)$. However, for the cellular cycle, the transition is divided into the update step $y' \leftarrow f(x_R, y, z_L)$ and the transmission step, which performs $y'_L \leftarrow y'$ and $y'_R \leftarrow y'$.

α-synchronism ($m = 0$): at each time step, each cell is updated with a fixed probability α, or else left unchanged (no representation of perception).

β-synchronism ($m = 1$): each cell is updated at each time step, but the transmission of the new state to the whole neighbourhood is realized with a probability β.

γ-synchronism ($m = k$): each cell is updated at each time step, but the transmission to each neighbour is realized independently with a probability γ.

As a result, perturbations consist in applying one of the steps of the cell cycle with a probability defined as the *synchrony rate*. Remark that for the sake of simplicity, α, β and γ denote both the type of asynchronism and the associated synchrony rates.

2.2 Formal definition of the updating schemes

α-synchronism. We introduce a *selection function* $\Delta_\alpha : \mathbb{N} \to \mathcal{P}(\mathcal{L})$ which gives for time t the subset of cells to be updated, where each cell has a probability α to be selected. Note that when $\alpha = 1$ the updating is fully synchronous and the system is deterministic. The global transition function becomes $\forall t \in \mathbb{N}$, $\forall c \in \mathcal{L}$, $\mathcal{N} = \{n_1, ..., n_k\}$:

$$F_\Delta(x^t(c)) = \begin{cases} f(x^t(c), x^t(c + n_1), ..., x^t(c + n_k)) & \text{if } c \in \Delta(t) \\ x^t(c) & \text{otherwise.} \end{cases}$$

By contrast with asynchronous updating, asynchronous information transmission aims at modelling the disruption of the exchange of information about states between cells. We now give a framework for this new type of asynchronism, modifying the classical definition of cellular automata with adequate extensions.

5

Figure 3: Example of the time cycle for a 3-cell sample for three perturbations of the updating scheme: α-synchronous (left), β-synchronous (middle) and γ-synchronous (right) (0 is represented in white, 1 in gray). The applied rule is ECA **50** (see Sec. 3).

Consider a CA defined by $(\mathcal{L}, Q, \mathcal{N}, f)$. First, the state space Q is extended to $Q' = Q \times Q^m$ in order to differentiate for a cell configuration $x^t(c)$, its *eigenstate* $e^t(c) \in Q$ and a vector of *m observation states* $\mathbf{o}^t(c) \in Q^m$. Secondly, the transition function f is replaced by $f' = f_{\mathrm{tr}} \circ f_{\mathrm{u}}$, that is, two functions applied sequentially to represent the actual update and the transmission of the state:

- $f_{\mathrm{u}} : Q'^{k+1} \to Q'$ is the *update* function; it computes the new eigenstate of a cell based on its eigenstate and the observation state of the neighbours, leaving the observation states unchanged:

$$f_{\mathrm{u}}\left(\begin{pmatrix} e \\ \mathbf{o} \end{pmatrix}, \begin{pmatrix} \cdot \\ \mathbf{o}_1 \end{pmatrix}, \cdots, \begin{pmatrix} \cdot \\ \mathbf{o}_k \end{pmatrix} \right) = \begin{pmatrix} \phi(e, \mathbf{o}_1, \cdots, \mathbf{o}_k) \\ \mathbf{o} \end{pmatrix} \qquad (1)$$

  where $\phi : Q \times Q'^k \to Q$ computes the new eigenstate based on the result of the local transition function f and the observation states for each neighbour.

- $f_{\mathrm{tr}} : Q' \to Q'$ is the *transmission* function; it carries out the "propagation" of the eigenstate $e^t(c)$ to the neighbourhood: this operation consists of replacing the values of observation states $\mathbf{o}^t(c)$ of the cell by the eigenstate. To this end, we introduce a set of random variables $(\theta_i^{t,c})_{t \in \mathbb{N}, c \in \mathcal{L}, i \in [\![1,m]\!]}$ which represent the atomic transmissions of the eigenstate $e^t(c)$ to the i-th element of $\mathbf{o}^t(c)$, such that:

$$\forall t \in \mathbb{N}, \ \forall c \in \mathcal{L}, \ \forall i \in [\![1, m]\!], \ \theta_i^{t,c}(e, o) = \begin{cases} e & \text{with a probability } P, \\ o & \text{with a probability } 1 - P, \end{cases} \qquad (2)$$

  where P is the synchrony rate corresponding to the asynchronisms (α, β, γ). The transmission function f_{tr} is thus defined, for the configuration $(e, \mathbf{o})$

of a cell c at time t, as a stochastic function such that:

$$f_{\mathrm{tr}}^{t,c}\begin{pmatrix} e \\ \mathbf{o} \end{pmatrix} = \begin{pmatrix} e \\ \theta_1^{t,c}(e, o_1) \\ \vdots \\ \theta_m^{t,c}(e, o_m) \end{pmatrix}. \tag{3}$$

We now detail the two implementations of this definition:

β-**synchronism.** This asynchronism consists in considering perturbations as disruptions of the emission of the eigenstate by a cell. We represent the observation state by a single element of Q ($m = 1$), which is perceived by all neighbouring cells:

- the state of a cell, for a given configuration x^t, is given by $x^t(c) = (e^t(c), \mathbf{o}^t(c))$, where $\mathbf{o}^t(c) \in Q$.

- $\phi(e, o_1, \cdots, o_k) = f(e, o_1, \cdots, o_k)$.

γ-**synchronism.** This asynchronism consists in considering the disruption of the transmission of a cell's eigenstate to each neighbour independently. As a consequence, the vector of observation states associates an element of Q for each element of the neighbourhood ($m = k$):

- the state of a cell, for a given configuration x^t, is given by $x^t(c) = (e^t(c), \mathbf{o}^t(c))$, where $\mathbf{o}^t(c) \in Q^k$.

- The transition function ϕ is given by:

$$\begin{aligned} \phi(e, \mathbf{o}_1, \cdots, \mathbf{o}_k) &= \phi\left(e, \begin{pmatrix} o_{1,1} \\ \vdots \\ o_{1,k} \end{pmatrix}, \cdots, \begin{pmatrix} o_{k,1} \\ \vdots \\ o_{k,k} \end{pmatrix}\right) \\ &= f(e, o_{1,1}, \cdots, o_{k,k}), \end{aligned}$$

where $o_{i,i}(c)$ corresponds to the independent transmitted state to cell c from its i-th neighbour.

As we defined them, β- and γ-synchronism differ only in the representation of the observation states. Readers should note that, like α-synchronism, β- and γ-synchronism can be understood as classical probabilistic cellular automata with an extended state space (Bouré et al. 2011).

Let us now observe how these updating schemes affect the behaviours of cellular automata using a well-known set of rules.

3 Preliminary observations

We now apply β- and γ-synchronism to one of the simplest classes of cellular automata, the Elementary Cellular Automata, and compare the resulting behaviours[1] to the dynamics of α-synchronism.

Elementary Cellular Automata (ECA). An *ECA* is a one-dimensional binary cellular automaton with nearest-cell neighbourhood. It is defined with $(\mathcal{L}, \mathcal{N}, \{0,1\}, f)$ where:

- $\mathcal{L} = \mathbb{Z}/L\mathbb{Z}$ is a 1-dimensional ring of length L,

- $\mathcal{N} = \{-1, +1\}$, *i.e.* the left and right neighbours,

- the local transition function $f : Q^3 \to Q$.

It is common to designate each ECA by the decimal code $W = f(0,0,0) \cdot 2^0 + f(0,0,1) \cdot 2^1 + ... + f(1,1,1) \cdot 2^7$. Through the use of left/right reflexion and 0/1 complementarity, it is possible to narrow down the study of the ECA space to 88 classes, each represented by the rule of smallest number.

In the following, initial conditions are constructed by setting the value of the eigenstates randomly and uniformly in $\{0,1\}$. As we consider the disruption of information transmission as a perturbation, we set the values of the observable states equal to the eigenstate, thus assuming that no error exists at initial time.

Experimental protocol. It is well known that the classification of ECA rules is a difficult problem, due to the richness and the complexity of the encountered behaviours (see *e.g.* Wolfram classes in (Culik II and Yu 1988)). Similarly, studies on α-synchronism have suggested that there is no straightforward correlation between Wolfram classes and their robustness to α-synchronism (see Fatès 2009). As β- and γ-synchronism define new updating mechanisms, they raise the question of whether they have the same effects than α-synchronism. As an analytical prediction of behaviour is a difficult problem (Fukś and Skelton 2012), we first decide to probe the entire ECA rule space experimentally[2].

In order to facilitate the comparison of the effects between different asynchronism, we will classify ECA rules according to their response to asynchronism, that is, an observation of the behaviour as a function of the perturbation applied. To estimate the response of ECA, we here use two methods (see Fig. 4 for visualisations):

1. We visually compare the space-time diagrams obtained by displaying the values of eigenstates over a few time steps. This approach provides a good idea on the changes of behaviour, but cannot hold as a formal classification criterion.

---

[1] Because of the proximity of their definitions, the effects of β- and γ-synchronism on CA behaviours are expected to be similar in most cases. Therefore γ-synchronism will only be mentioned when a difference with β-synchronism is observed.

[2] The complete set of results can be found at: `http://www.loria.fr/~boure/eca`.
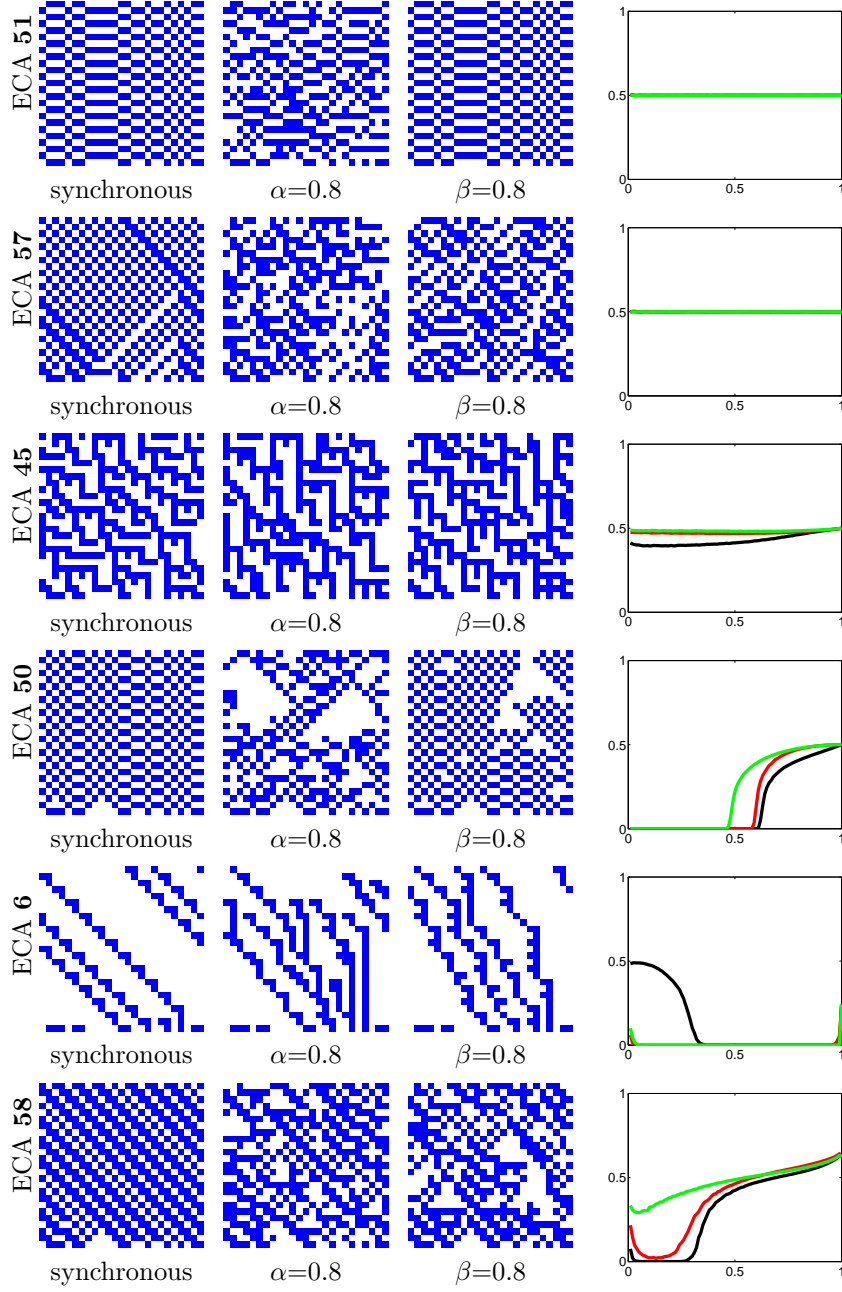
Figure 4: A sample of the observations made on ECA rules on asynchronous updating schemes. The plot show the profile of the density parameter as a function of the synchrony rate (black line for α, red/dark gray for β, and green/light gray line for γ). All diagrams start from the same initial configuration. Time goes from bottom to top.

Figure 5: Samples of space-time diagrams of ECA **50** for different synchrony rates and different types of asynchronism, starting from the same initial configuration.

2. We quantify the behaviour using an order parameter, the density, defined as the ratio of cells with eigenstate 1. We give an estimation of the asymptotic density, by letting a system of size $L = 2000$ evolve for 5000 steps, and then averaging the parameter value for 100 steps.

Results. For the vast majority of rules, neither the space-time diagram nor the density profile allow to qualify any change of behaviour between α-synchronism and the two other types of asynchronism. However, some rules show remarkable changes, as illustrated by the examples of Fig. 4:

- The space-time diagrams of rule ECA **51** (**NOT** rule) shows a trivial example of a qualitative difference between α- and β-synchronism. Indeed, as the transition function does not actually depend on the state of the neighbourhood, ECA **51** is perfectly robust to asynchronous information transmission.

- For rules like ECA **51** and ECA **57**, the initial value of the density parameter is conserved for all values of the synchrony rate. However, this observation is not representative of the behaviour as a quick comparison of their space-time diagrams reveals that it is radically different.

- For rules with a behaviour similar to ECA **45**, the density shows a slight difference between α-, β- and γ-synchronism.

- Finally, some rules show a discontinuity in the plot profile of the density, such as for ECA **6, 50, 58**, suggesting that a qualitative change of behaviour occurs for a critical value of the synchrony rate. This change is confirmed by the observation of the space-time diagrams (see Fig. 5).

In Tab. 1, we collect rules according to the type of response obtained for the density parameter and divide them into three categories: (a) a brutal change of

Table 1: Classification of the minimal ECA per observed difference of behaviour, using the density parameter.

(a) Discontinuity	6, 18, 26, 38, 50, 58, 106, 134, 146
(b) Other different response	3, 4, 9, 12, 19, 22, 33, 35, 36, 44, 45, 46, 54, 62, 72, 76, 104, 108, 110, 122, 132, 152, 164, 172
(c) No observable difference	0, 1, 2, 5, 7, 8, 10, 11, 13, 14, 15, 23, 24, 25, 27, 28, 29, 30, 32, 34, 37, 40, 41, 42, 43, 51, 56, 57, 60, 73, 74, 77, 78, 90, 94, 105, 126, 128, 130, 136, 138, 140, 142, 150, 154, 156, 160, 162, 168, 170, 178, 184, 200, 204, 232

behaviour occurs for at least one of the asynchronisms, (b) a difference of level or slope can be observed and (c) no difference can be observed. The observation of a qualitative change of behaviour is a strong evidence of non-robustness, and conveniently, it is quantifiable. To clarify further how updating schemes affects cellular automata, we choose to focus on this category and study these discontinuities in the context of statistical physics.

4 Study of phase transitions

The occurrence of phase transitions is probably one of the most remarkable phenomena that arises in asynchronous cellular automata: in such models, there exists a non-trivial value of the synchrony rate, the critical threshold, which separates two distinct "phases": a passive phase for which the system converges to an homogeneous fixed point of 0s, and an active phase where the system the system does not converge to a fixed point but nevertheless oscillates around a fixed density. In our case, different discontinuities appear depending upon the updating scheme used:

(a) for the rules ECA **18, 26, 50, 106, 146**, a phase transition occurs for all three types of asynchronism,

(b) for the rules ECA **6, 38, 134**, a phase transitions is observed only for α-synchronism,

(c) a phase transition can be observed for α- and β-synchronous ECA **58**, but not for γ-synchronism.

A first observation is that all these rules show a discontinuity for α-synchronism, but not necessarily for the two others. For the α-synchronous case, these abrupt changes were identified by Fatès (2009) as phase transitions, which belong to the directed percolation (DP) universality class. This means that these models follow the same characteristic power laws at criticality (see Hinrichsen 2000). These rules were then classified into two categories[3]:

---

[3]For the sake of conciseness, we choose to put the third class DP_2 aside, as it displays a phase transition in ECA **178**, but for an order parameter different than density.

Figure 6: Phase transition analysis for β-synchronous ECA **50** on a log-log plot. The straight line follows a power law $f(t) = K \times t^{-\delta}$ with the critical exponent $\delta = 0.1595$. Measures are averaged over 32 automata of size 10000.

- the DP_{hi} class comprises all ECA with a "positive" phase transition: the active phase occurs for high values of the synchrony rate, and the passive phases for lower values. This class is made of ECA **18, 26, 50, 58, 106, 146**.

- the DP_{low} class gathers ECA with a "negative" phase transition: active and passive phases are reversed in order. It is made of ECA **6, 38, 134**.

The existence of a correlation between these classes and the rules' response in β-synchronism is an interesting question, which remains as an open problem.

Quantification of the critical behaviour. Because all drastic changes in β-synchronism also belong to the DP_{hi} class, we conjectured that all observed discontinuities are caused by phase transitions which belong to the directed percolation universality class. To verify this hypothesis for β- and γ-synchronism, we reused the same experimental protocol described by Fatès (2009) for α-synchronism:

1. We fix the synchrony rate, start from a random initial configuration and let it evolve for a given number of steps.

2. We monitor the evolution of the density for a long simulation time until we observe a sub-critical or super-critical behavior. In a log-log plot such as Fig. 6, a concave curve occurs for the passive phase and a convex curve for the active phase. As we expect the order parameter to follow a power law $K.t^{-\delta}$ near criticality, the curve at critical value should appear as a straight line of slope $-\delta$.

3. We repeat the experiment with a value of the synchrony rate closer to the critical point until a satisfactory precision is reached.

Figure 6 presents the results of this experiment for β-synchronous ECA **50**. For this particular example, the asymptotic behaviour at critical threshold follows a power law $K.t^{-\delta}$ with $\delta \sim 0.1595$, which is the characteristic exponent for the directed percolation universality class in $1 + 1$ dimensions (Hinrichsen 2000). Using successive experiments with different synchrony rates, we measured the

Table 2: Measurements of the critical synchrony rates for different rules and asynchronisms.

Class	ECA	α_{c}	β_{c}	γ_{c}
$\mathrm{DP_{hi}}$	18	0.714	0.749	0.655
	26	0.475	0.526	0.406
	50	0.628	0.601	0.487
	58	0.340	0.289	–
	106	0.815	0.848	0.813
	146	0.675	0.732	0.635
$\mathrm{DP_{low}}$	6	0.283	–	–
	38	0.041	–	–
	134	0.082	–	–

critical value for each asynchronism with a precision of 10^{-3}. Table 2 lists the values of the critical synchrony rates; all measured values of the critical exponent δ show good evidence of directed percolation.

To sum up, these experiments reveal a strong connection between α- and β-synchronism for the $\mathrm{DP_{hi}}$ class, which contrasts with the results obtained for the $\mathrm{DP_{low}}$ class and the exception of γ-synchronous ECA **58**. This is already an interesting result, as we have identified phenomena that are exclusive to each asynchronism. A potential use of these differences is that they can be used as a "signature" of the updating scheme. In order to understand the origin of these differences, we now propose to study these models at a lower level.

5   Microscopic approach

In this section, we focus on the study of the local evolution of cells under α- and β-synchronous updating schemes, using three examples that represent different responses to asynchronism: ECA **50** for the $\mathrm{DP_{hi}}$ class, ECA **6** for the $\mathrm{DP_{low}}$ class, and ECA **58** as an exception to the $\mathrm{DP_{hi}}$ class. Note that these approaches are not aimed at giving a full proof of the observed behaviour, but rather to give us an insight on how different types of asynchronism may produce different responses.

5.1   First case study: ECA 50

ECA **50** is defined by the following transition table:

000	001	010	011	100	101	110	111
0	1	0	0	1	1	0	0

The study of this rule reveals a paradoxal response to α- and β-synchronism: if the evolution of the density parameter suggests that directed percolation phenomena occurs in both cases, a quick inspection of the space-time diagrams over a few time steps clearly shows a radical difference in the local behaviours

| t (a) checkerboard | (b) anomalies | (c) particle shift | (d) particle |
| and particles | | | annihilation |

Figure 7: Patterns appearing in space-time diagrams of β-synchronous ECA **50**.

(see Fig. 4 page 9). This suggests that the local behaviour differs from one asynchronism to the other.

Describing patterns.   By looking at the β-synchronous space-time diagrams, we perceive a kind of periodicity in the organisation of states that does not appear for α-synchronism. In particular, configurations divide up in large patterns of cells with consecutive or alternated eigenstates 0 and 1, which appear or disappear with random events:

- The *checkerboard region* are defined as a group of spatially alternated eigenstates. Under a synchronous scheme they will flip all their states, thus drawing a "checkerboard" over a few time steps (see Fig. 7-a).

- The *particles* separate checkerboard regions; they are formed by pairs of cells with identical eigenstates (*i.e.* 00 and 11), which both flip when updated (see Fig. 7-a).

- Occasionally, random *anomalies* appear, that is, groups of three or more cells with eigenstates 0 that perturbate the regularity of the two previous patterns. Typically, this results in the appearance of "white holes" in the patterns above (Fig. 7-b). When these anomalies include particles, they can displace them spatially (Fig. 7-c) or even annihilate them by pairs (Fig. 7-d). Remark that as the synchrony rate β gets smaller, the size of anomalies and their frequency of appearance increase.

Consequence on the global behaviour.   As a result of the observed microscopic behaviour, we can deduce that the asymptotic behaviour will converge to one of the following configurations:

- if β is below the critical value, anomalies take over other patterns until they cover the entire lattice. This scenario corresponds to the convergent phase as observed by the density parameter in the previous section.

- if β is high enough and strictly inferior to 1, anomalies continue to appear but checkerboard regions will "repair" them in a few time steps. Instead, particles follow a spatial non-biased random walk until they meet and

14

Figure 8: Samples of space-time diagrams of ECA **6** under different values of synchrony rates α and β. Note the active phase occurring for $\alpha = 0.3$.

annihilate. After long simulation times, the checkerboard pattern will cover the entire lattice.

To sum up, this behaviour, derived from the stability of checkerboard patterns under β-synchronism, is a novel property exclusively observed for this updating scheme, and suggests that the behaviour of ECA **50** is indeed qualitatively different in α- and β-synchronism.

5.2   Second case study: ECA 6

ECA **6** shows a peculiar behaviour under different asynchronisms: a phase transition occurs for the α-synchronous version but it is not observed for β-synchronism (see Fig. 8). The transition table of the rule is:

000	001	010	011	100	101	110	111
0	1	1	0	0	0	0	0

In order to understand the mechanism involved in the dynamics of this rule, one needs to analyse the evolution of cells states at a microscopic level. To simplify the description of configurations of ECA **6**, we propose to consider *branches* as contiguous groups of 1-cells surrounded by 0-cells. Under synchronous updating, these branches shift regularly but when noise is introduced, variations can occur in the form of delays, creation, destruction of branches.

Behaviour in α-synchronism.   In order to describe the qualitative behaviour of α-synchronous ECA **6**, we analyse two "events": the destruction or the creation of a branch, and determine which event is more likely to occur depending on the value of the synchrony rate α. Two mechanisms need to be highlighted:

1. We first consider the case of a single branch and start to build the graph of the possible evolutions, with their associated probabilities (see Fig. 9). We see that for all values of the synchrony rate α, the creation of a new branch ("birth") is more probable than its destruction ("death").

15

Figure 9: A sample of transition graph which illustrates the possible evolutions from a single branch in α-synchronous ECA **6**. Here, by considering only groups of four cells, we find that the probability of "Birth" events is higher than "Death" events.

2. On contrary, when two branches are separated by only one cell, the right branch can be destroyed (see the synchronous case of Fig. 10). As the distance between branches varies due to random updates, multiple possibilities to destruct branches exist through this process.

For values of the synchrony rate α close to 1, the probabilities to create or destroy branches are almost equal. On the other hand, the probability to destruct branches through their interactions is high as, for high α, this phenomenon happens almost surely as soon as the distance between branches becomes critical. However, for low values of α, the opposite behaviour occurs: as bigger clusters of continuous 1s appear more often, it becomes more probable to create new branches than to destroy them, thus overcoming the effect of interactions branches. This behaviour is confirmed by experiments, for which the behaviour remains active and the density parameter converges to a fixed value (see Fig. 4 page 9).

To sum up, studying the local dynamics of ECA **6** allows us to understand better why two different phases may occur for different values of α. We now intend to consider the β-synchronous case in order to investigate why the phase transition does not exist.

Novel mechanisms under β-synchronism. Because of the extension of the state space, to study the dynamics of β-synchronous updating using the previous approach is difficult. However, comparing the mechanism involved behind the creation of branches gives an insight on why the non-converging phase does not exist for β-synchronism.

Let us consider the examples given in Fig. 10, where creation of branches are visualized under two updating schemes. In the α-synchronous case, as soon as an isolated 1 is created, it remains stable (case a). However, in the β-synchronous case, branches are created when a cell changes its eigenstate to 0 but keeps its observable as 1 (light gray cells). These branches have a short life expectancy, as (1) they cannot activate neighbouring cells until some space is created between branches and (2) they disappear as soon as they finally manage

16

Figure 10: Branching phenomenon in ECA **6** under α- and β-synchronism. Light gray cells represent state $(0,1)$ and green/darker gray state $(1,0)$.

to update their observable state to `0` (case b). The only means to create branches in β-synchronous ECA **6** occurs for low values of the synchrony rate β (case c): in that case, the branch will survive sufficiently long time for some space to be created and then it will survive like a normal branch.

The arguments above do not formally constitute a proof, but this first analysis may provide a possible explanation for a difference of behaviour between α- and β-synchronism.

5.3   The case of ECA 58

The observations made on ECA **58** are probably one of the most surprising result of this study: among all ECA, it is the only rule that displays a qualitative difference of behaviour between β- and γ-synchronism (see Fig. 11). To this date, it is not understood why the behaviour of ECA **58** changes between the two updating schemes. This issue is justified by several factors:

1. First, proving that the absence of a passive phase is a much more difficult problem than proving its existence. The complexity of establishing such a proof is also greatly increased by the fact that the state space is larger (with observation states).

2. Second, β- and γ-synchronism follow similar mechanisms. Indeed, they differ only in the sense that in β-synchronism, cell-to-cell transmission of the same neighbourhood are correlated.

3. Finally, it should be noted that although the rule ECA **58** differs only by one bit from ECA **50**, it is sufficient to introduce novel patterns that limits the use of previous techniques:

000	001	010	011	100	101	110	111
0	1	0	1	1	1	0	0

Displaying the mechanisms behind this phenomenon and explaining why the phase transitions does not occur in γ-synchronism is still an open problem which could provide another relevant example of how sensitive to the updating method cellular automata behaviours can be. For now, deeper statistical experiments
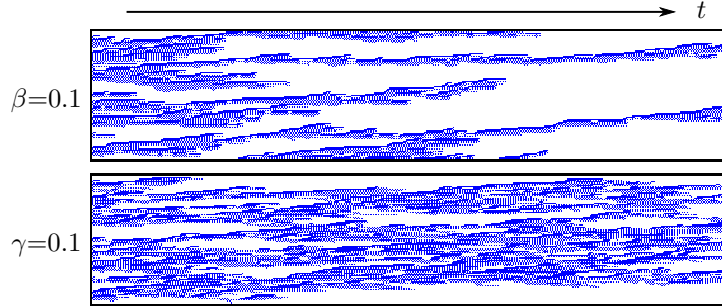
17

Figure 11: Comparative observation of the behaviour of ECA **58** under β- and γ-synchronism. According to quantitative experiments, the sub-critical behaviour for β occurs for $\beta < 0.289$ whereas the γ-synchronous version does not seem to converge to an all 0 configuration.

must be conducted in order to confirm this difference of behaviour, and novel approaches must be elaborated in order to find possible divergences between β- and γ-synchronism.

6 Discussion

This paper presented a formalism for a new type of asynchronism in cellular automata, called asynchronous information transmission. We compared this type of perturbation to α-synchronism for Elementary Cellular Automata. We then highlighted specific examples for which directed percolation phenomena occur, and observed different responses to each asynchronism:

- For ECA **50**, the macroscopic behaviour is conserved, but a microscopic analysis of the dynamics reveals novel properties for β-synchronism.

- For ECA **6**, β-synchronism introduces modifications in the microscopic dynamics, which alter qualitatively the macroscopic behaviour.

- For ECA **58**, a qualitatively different response is observed for β- and γ-synchronism.

The implications of these observations seem promising: in spite of the resemblance of the responses of most ECA, for each updating scheme we could find at least one example where the behaviour differs from the other updating schemes. This observation supports the idea that the behaviour of cellular automata can sometimes be largely determined by the updating scheme. This echoes recent results obtained for a biological model with a stochastic transition rule (Bouré et al. 2012).

However, it is still an open question to predict how the updating scheme influences the behaviour of cellular automata and why phase transitions appear

in some automata and not in others (*e.g.* Fatès 2010). By resuming the exploration of the field of updating schemes, we bring to light new examples of behaviour alterations appearing under different perturbations.

Probing a model to reveal its specificities. The approach we followed used the updating scheme as a "probing tool" to explore the behaviour of cellular automata. By subjecting models to different kinds of updating scheme and observing the resulting behaviour, we gained insight on (1) the models' properties – by highlighting to which perturbation they are robust to – and (2) the updating schemes themselves, by underlining the properties they affect, or the novel phenomena that they introduce.

Updating schemes as a component. To this date, studies on complex systems have mainly focused on a variation of the individual rules considering a single updating scheme. Our conclusion is that, in studies of complex systems, the updating scheme should hold a modular position in order to allow a scrutiny of the behaviour that is *orthogonal* to the exploration of the transition rule space.

References

Dave H. Ackley and Lance R. Williams. Homeostatic architecture for robust spatial computing. In *Proceedings of the 5th IEEE International Conference on Self-Adaptive and Self-Organizing Systems*, 2011.

Stefania Bandini, Andrea Bonomi, and Giuseppe Vizzari. What do we mean by asynchronous CA? A reflection on types and effects of asynchronicity. In *Proceedings of ACRI 2010*, volume 6350 of *Lecture Notes in Computer Science*, pages 385–394, 2010a.

Stefania Bandini, Andrea Bonomi, and Giuseppe Vizzari. Modeling and programming asynchronous automata networks: The MOCA approach. In *Proceedings of ACRI 2010*, volume 6350 of *Lecture Notes in Computer Science*, pages 345–355, 2010b.

Hendrik J. Blok and Birger Bergersen. Synchronous versus asynchronous updating in the "Game of Life". *Physical Review E*, 59(4):3876–3879, 1999.

Olivier Bouré, Nazim Fatès, and Vincent Chevrier. Robustness of Cellular Automata in the Light of Asynchronous Information Transmission. In *Unconventional Computation*, volume 6714 of *Lecture Notes in Computer Science*, pages 52–63, 2011.

Olivier Bouré, Nazim Fatès, and Vincent Chevrier. First steps on asynchronous lattice-gas models with an application to a swarming rule. In *Proceedings of ACRI 2012*, Lecture Notes in Computer Science, 2012. URL `http://hal.inria.fr/hal-00687987/`. (to appear).

Karel Culik II and Sheng Yu. Undecidability of CA classification schemes. *Complex Systems*, 2:177–190, 1988.

Nazim Fatès. Asynchronism induces second order phase transitions in elementary cellular automata. *Journal of Cellular Automata*, 4(1):21–38, 2009.

Nazim Fatès. Does *Life* resist asynchrony? In *Game of Life Cellular Automata*, pages 257–274. Springer, 2010.

Nazim Fatès and Michel Morvan. An experimental study of robustness to asynchronism for elementary cellular automata. *Complex Systems*, 16:1–27, 2005.

Henryk Fukś and Andrew Skelton. Orbits of bernoulli measure in asynchronous cellular automata. *DMTCS*, AP:95–112, 2012. Proceedings of Automata 2011.

Abdelkader Gouaïch, Fabien Michel, and Yves Guiraud. MIC*: a deployment environment for autonomous agents. In *Post-Proceedings of the 1st International Workshop in Environments for Multiagent Systems*, pages 109–126. Springer-Verlag, 2005.

Carlos Grilo and Luís Correia. Effects of asynchronism on evolutionary games. *Journal of Theoretical Biology*, 269(1):109–122, 2011.

Haye Hinrichsen. Nonequilibrium critical phenomena and phase transitions into absorbing states. *Advances in Physics*, 49:815–958, 2000.

Thomas E. Ingerson and Raymond L. Buvel. Structure in asynchronous cellular automata. *Physica D: Nonlinear Phenomena*, 10(1-2):59 – 68, 1984.

Ferdinand Peper, Teijiro Isokawa, Yousuke Takada, and Nobuyuki Matsui. Self-timed cellular automata and their computational ability. *Future Generation Computer Systems*, 18(7):893 – 904, 2002.

Damien Regnault. Directed percolation arising in stochastic cellular automata analysis. In *Mathematical Foundations of Computer Science*, volume 5162 of *Lecture Notes in Computer Science*, pages 563–574. Springer-Verlag, 2008.

John von Neumann. *Theory of Self-Reproducing Automata*. University of Illinois Press, Champaign, IL, USA, 1966.