

A decentralised diagnosis method with probabilistic cellular automata

Nazim Fatès*
nazim.fates@loria.fr

Régine Marchand†
regine.marchand@univ-lorraine.fr

Irène Marcovici‡
irene.marcovici@univ-lorraine.fr

May 12, 2023

Abstract

The decentralised diagnosis problem consists in the detection of a certain amount of defects in a distributed network. Here, we tackle this problem in the context of two-dimensional cellular automata with three states : neutral, alert and defect. When the density of defects is below a given threshold, we want the alert state to coexist with the neutral state while when this density is above the threshold, we want the alert state to invade the whole grid.

We present two probabilistic rules to answer this problem. The first one is isotropic and is studied with numerical simulations. The second one is defined on Toom's neighbourhood and is examined with an analytical point of view. These solutions constitute a first step towards a broader study of the decentralised diagnosis problem on more general networks.

1 Introduction

Consider the following problem. We have a system of interconnected components that interact in discrete time steps according to a local rule. At each time step, each component can fail with a given probability, in which case it becomes *defective*. The non-defective components see their defective neighbours and they may switch from their normal state, or *neutral state*, to an *alert* state. We would like such a system to detect the presence of a given number of failures

*Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France.

†Université de Lorraine, CNRS, Inria, IECL, F-54000 Nancy, France

‡Université de Lorraine, CNRS, Inria, IECL, F-54000 Nancy, France

in a totally decentralised way. More precisely, when the density of failures is below a given threshold, some components should still operate in the neutral state, while when this density crosses this threshold, all the components should be in the alert state, and remain so.

We call this problem the *decentralised diagnosis problem* and we investigate it in the cellular automata framework, that is, when the components, which are called the *cells*, are arranged according to a regular spatial structure (a grid) and all obey the same uniform transition rule.

As we aim to work with a minimal model, we will consider only ternary cellular automata with the three states mentioned above (neutral, alert, defect). The main idea of our proposition to use the phase transitions phenomenon, observed in some probabilistic cellular automata where a qualitative change of behaviour may be triggered by a small variation in the local transition probabilities. For example, it is well known that the probabilistic version of the Greenberg-Hastings model exhibits such a phase transition between an active regime where cells change their state regularly and an extinction regime where all activity disappears [1]. Interestingly enough, it was noticed that the presence of defects in the grid modifies the position of the critical threshold [1, 5, 4] and the idea to take advantage of this drift of threshold to build a diagnosis process was formulated.

Some preliminary investigations on the decentralised diagnosis problem were carried out by Gauville et al., who examined different models and compared their advantages and drawbacks [3]. The current paper aims to deepen these investigations by examining two simple ternary models. The first one, named `QuorumD`, is inspired from the Ising model and uses the von Neumann neighbourhood ; the second model is even simpler and uses Toom's size-3 neighbourhood: this neighbourhood breaks the isotropy and installs a direction in which information travels. Indeed, while this may seem a drawback for concrete applications, it has the advantage to allow us to derive a more formal analysis of the behaviour of the model.

2 The isotropic model

The cells are arranged on the two-dimensional lattice \mathcal{L} , which can be infinite ($\mathcal{L} = \mathbb{Z}^2$), or can take the form of a finite square lattice of length L with periodic boundary conditions in which case we have $\mathcal{L} = (\mathbb{Z}/L\mathbb{Z}) \times (\mathbb{Z}/L\mathbb{Z})$.

The set of states is $Q = \{0, 1, D\}$ where 0 stands for the neutral state, 1 for the alert state, and D represents a defective cell. The set of configurations is $Q^{\mathcal{L}}$.

Let us now introduce the `QuorumDrule`. It is inspired by Ising model of ferromagnetic interactions of spins. We assimilate one spin to the neutral state 0 and the other spin to the states 1 and D, grouped together. Indeed, we want the defect state to generate alert states, and we want these alert states to propagate from neighbour to neighbour and to invade the whole grid only if their number is sufficiently high. We thus propose to consider the following

rule [3]. For a cell in state q , its new state is given by $f(q, n, a, d)$ where n, a, d are respectively the number of states 0, 1, D found in the neighbourhood of this cell. In this section, we will restrict our scope to the von Neumann neighbourhood: $\mathcal{N} = \{(0, 0), (0, 1), (1, 0), (0, -1), (-1, 0)\}$.

We define the local transitions by steps. First, we express the fact that the defect state is inert: $f(D, ., ., .) = D$. Second, for $q \neq D$, we express the quiescence of the neutral and alert state by: $n = 0$ implies that $f(q, n, a, d) = 1$ and by $a = d = 0$ implies that $f(q, n, a, d) = 0$. Last, we give the general case: for $q \neq D, n > 0$ and $a + d > 0$, we have:

$$f(q, n, a, d) = \begin{cases} 1 & \text{with prob. } \pi(n, a, d) \\ 0 & \text{with prob. } 1 - \pi(n, a, d), \end{cases}$$

where

$$\pi(n, a, d) = \frac{\exp \lambda(a + d)}{\exp \lambda n + \exp \lambda(a + d)}.$$

The parameter λ , which could be called the resistivity of the environment, controls the ease with which the alert state propagates. It can also be seen as the inverse of a temperature: when λ is high, the cells tend to follow the majority state in their neighbourhood, which makes it difficult for the alert state to diffuse; reciprocally, when λ is low, more noise is introduced and a group of alerts can travel at a greater distance from the defect cell which created it.

To define the global dynamics of our cellular automaton, let us introduce the functions $C_q(\eta, x, y) = \text{card}\{(i, j) \in \mathcal{N} : \eta(x + i, y + j) = q\}$, which, for a configuration η , count the number of cells in state $q \in \{0, 1, D\}$ in the neighbourhood of a given cell $(x, y) \in \mathcal{L}$. Then, starting from an initial configuration η_0 , the system evolves according to the sequence of configurations $(\eta_t)_{t \in \mathbb{N}}$ defined by:

$$\forall (x, y) \in \mathcal{L}, \eta_{t+1}(x, y) = f(\eta_t(x, y), C_0(\eta_t, x, y), C_1(\eta_t, x, y), C_D(\eta_t, x, y))$$

where the probabilities for f are drawn independently for each cell and at each time step.

2.1 Numerical Experiments

The dynamics of the QuorumD rule can be observed on Fig. 1. In this example, we took a small grid ($L = 32$) and arbitrarily fixed the value of the resistivity to $\lambda = 1$. Each cell was initially set as a defect with probability κ and to neutral state with probability $1 - \kappa$. It can be observed that for a defect density $\kappa = 0.05$ the alert states invades the grid in a few hundred steps, while for a smaller defect density $\kappa = 0.02$, the alert state remains contained in some regions, even after a several thousands of steps.

The distinction between these two qualitative states can be measured by the density of neutral states $d_0(t)$, that is, the ratio of neutral states over the total number of cells of the grid. We define the *coexistence state* as the behaviour of the system where the diffusion of the alert states remains bounded, which




Figure 1: Evolution of a random configuration with $L = 32$, and $\lambda = 1$. The states n, a, d are respectively in white, blue and purple (colour online). The upper line shows an evolution with a density of defects $\kappa = 0.1$: the alert state diffuses in the whole grid. The bottom line shows a long-term evolution for $\kappa = 0.02$: the diffusion of the alert state remains bounded.

means that $d_0(t)$ will remain stable around a non-zero value for a long time. Conversely, the *invasion state* corresponds to the behaviour where the alert state diffuses everywhere, and where $d_0(t)$ will quickly decrease.

Figure 2 shows the temporal evolution of d_0 for different values of λ and κ . We here took a large grid ($L = 100$) and an observation time of 10 000 time steps. On the top plots of this figure, it can be observed that for $\kappa = 0.03$, the transition between the coexistence and invasion state occurs for $\lambda \in [0.8, 1.6]$, while for $\kappa = 0.05$ it occurs in the interval $[1.6, 2.4]$. Reciprocally, if we fix the value of the resistivity, we see that for $\lambda = 8$, progressively increasing the defect rate by steps of 0.04 makes the system evolve from the coexistence state to the invasion state, which corresponds to the desired behaviour for a decentralised diagnosis.

These observations are rather encouraging for a practical use. We leave for future work a more detailed analysis, where the defects would dynamically appear on the grid. It is also necessary to evaluate how the grid size influences the behaviour of the system. Ideally, when the resistivity λ is fixed, one would like to have a fixed threshold for the value of κ for large grid sizes. Unfortunately, the observations we made for larger grids and larger times revealed that the separation between the coexistence and the invasion state does depend on the lattice size, although in a limited fashion.

By analysing the behaviour of the rule at a local scale, we observed that the dynamics of the alerts in some cases closely behaves as the *bootstrap percolation*

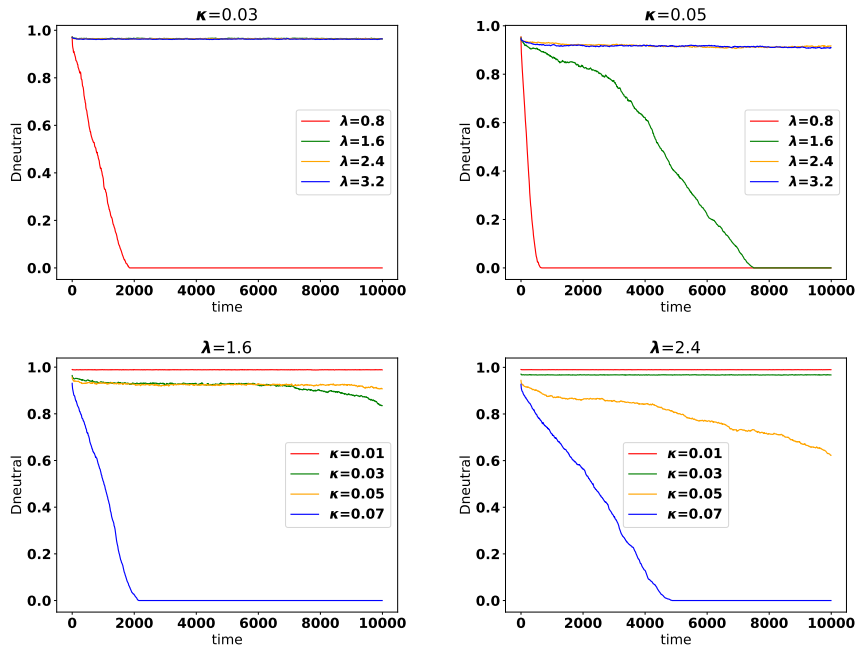


Figure 2: Temporal evolution of the density of neutral sites as function of time for different values of κ and λ , for a grid of size $L = 100$. On the two top plots, we set $\kappa = 0.03$ (top-left) or $\kappa = 0.05$ (top-right), and vary λ . On the two bottom plots, we set $\lambda = 1.6$ (bottom-left) or $\lambda = 2.4$ (bottom-right), and vary κ .

process [2]. In this model, cells have a binary state, healthy or infected, and a healthy cell gets infected if two or more neighbours are infected, after which it remains infected for ever. Although the `QuorumD` rule is stochastic and the bootstrap percolation rule is deterministic, the two dynamics have many common characteristics. In particular, in both models, rectangular zones of cells which are occupied (in the bootstrap percolation) or in the alert state (in `QuorumD`) are stable. We leave as an open question to determine more precisely to which extent the two rules have comparable dynamics. This is why we now present a simpler rule where the use of a non-isotropic neighbourhood induces a direction for the flow of information and makes the model more tractable.

3 A directional model: the biased-quiescent rule

3.1 Definition of the rule

We now consider a simplification of the previous model. The new model still has three states 0, 1 and D, but, on the non-defective cells, the local rule is

applied on a neighbourhood of size two, which consists of the North and East neighbours of a cell. The transitions are simplified: the two states 0 and 1 are still quiescent, and when a cell sees the two states, it becomes a 1 (alert) with probability p , whatever its initial state. We call this rule the **biased-quiescent rule**, or the BQ rule for short, as the rule has the quiescence property and also has a bias towards 1 with probability p .

Formally, for $p \in (0, 1)$, let ϕ be the probabilistic function defined by:

- (a) the neutral state is quiescent: $\phi(0, 0) = 0$;
- (b) the alert state is quiescent: if $q, q' \in \{1, \text{D}\}$ then $\phi(q, q') = 1$; and,
- (c) biased breaking-ties: otherwise, $\phi(q, q') = \begin{cases} 1 & \text{with prob. } p \\ 0 & \text{with prob. } 1 - p. \end{cases}$

Starting from a (random or deterministic) initial configuration $\eta_0 \in \{0, 1, \text{D}\}^{\mathbb{Z}^2}$, we denote by $(\eta_t)_{t \in \mathbb{N}}$ the sequence of random configurations obtained as follows:

$$\forall t \in \mathbb{N}, \forall (x, y) \in \mathbb{Z}^2, \eta_{t+1}(x, y) = \begin{cases} \text{D} & \text{if } \eta_t(x, y) = \text{D} \\ \phi(\eta_t(x+1, y), \eta_t(x, y+1)) & \text{otherwise,} \end{cases}$$

all the updates being independent.

In the sequel, we work on the infinite grid \mathbb{Z}^2 and we say that a configuration η has a *neutral background* if the number its non-neutral cells is finite. The main interest of this simplified and directional model is that it is analytically tractable. We now study how its asymptotic behaviour depends on the initial configuration η_0 ; we proceed in three steps:

- first, we study the evolution of an initial configuration with one alert in a neutral background,
- second, when it contains a single defect,
- third, when the defects are present with probability κ .

Our aim is to find how qualitative changes of behaviour occur when p and κ are varied, and to relate these changes to the invasion and coexistence regimes, and thus the ability to perform a decentralised diagnosis.

3.2 Change of representation and partial order on configurations

In order to simplify the notations, we introduce the following change of representation, which amounts to rotating the grid of an angle $5\pi/4$. The model now takes place on the lattice (see Fig. 3):

$$\mathcal{L} = \{(x, n) \in (1/2)\mathbb{Z} \times \mathbb{Z} : x \in \mathbb{Z} \text{ if } n \text{ is even, and } x \in (1/2) + \mathbb{Z} \text{ if } n \text{ is odd}\},$$

and the sequence of configurations $(\eta_t)_{t \in \mathbb{N}} \in \{0, 1, \text{D}\}^{\mathcal{L}}$ is defined by:

- if $\eta_t(x, n) = \text{D}$ then $\eta_{t+1}(x, n) = \text{D}$,

- otherwise, $\eta_{t+1}(x, n) = \phi(\eta_t(x - 1/2, n - 1), \eta_t(x + 1/2, n - 1))$,

all the updates being independent.

On $\{0, 1, D\}$, we consider the order $0 \leq 1 \leq D$, which naturally induces a partial order on $\{0, 1, D\}^{\mathcal{L}}$ in the following manner: we say that

$$\eta \preceq \eta' \text{ if for any } (x, n) \in \mathcal{L}, \eta(x, n) \leq \eta'(x, n).$$

By using the same Bernoulli random variables to update the processes starting from two comparable (deterministic or random) initial configurations, it is easy to see that:

Fact 1. *If $\eta_0 \preceq \eta'_0$, then we can couple the evolutions $(\eta_t)_{t \in \mathbb{N}}$ and $(\eta'_t)_{t \in \mathbb{N}}$ in such a way that: $\forall t \in \mathbb{N}, \eta_t \preceq \eta'_t$.*

3.3 Initial configuration with only one alert and no defect

In a first step, we consider the case where one alert is set in a neutral background. How will the alerts evolve as function of p ? To answer this question, we prove the following theorem.




Figure 3: The configuration η_6 with, in grey, the alert states forming a horizontal segment at time $t = 6$, with its two extremal cells in darker grey. Their possible evolutions at time $t = 7$ are shown in colour.

Theorem 1. *For the BQ rule, starting from a configuration η_0 with only one cell in state 1 and no cell in state D in a neutral background, the extinction time of the alert cells $T = \inf\{t \in \mathbb{N} : \eta_t = 0_{\mathbb{Z}^2}\}$ verifies:*

- For $p < 1/2$, then T is almost surely finite, and it is integrable ($\mathbb{E}(T) < \infty$);
- for $p = 1/2$, then T is almost surely finite, and $\mathbb{E}(T) = \infty$;
- for $p > 1/2$, T can be infinite with a positive probability ($\mathbb{P}(T = \infty) > 0$).

Proof. We assume that initially, the origin $(0, 0)$ is the only cell in state 1, all the other cells are in state 0. By induction, it is clear that at any time $t \in \mathbb{N}$, the configuration η_t is either the empty configuration $0^{\mathcal{L}}$, or a *horizontal segment*,

that is, a configuration where all the cells in state 1 have their second coordinate equal to t and have consecutive first coordinates. Let us denote by $L(t)$ the set of first coordinates of alert cells at time t :

$$\begin{aligned} L(t) &= \{x : \eta_t(x, t) = 1\}, \\ d(t) &= \max L(t) - \min L(t) + 1, \text{ with the convention } d(t) = 0 \text{ if } L(t) = \emptyset, \\ T &= \inf\{t \in \mathbb{N} : L(t) = \emptyset\} = \inf\{t \in \mathbb{N} : d(t) = 0\}. \end{aligned}$$

Note that $d(t)$ is also the cardinal of $L(t)$. Moreover, as long as $L(t) \neq \emptyset$, $(\max L(t))_{t \in \mathbb{N}}$ evolves as a random walk with steps $1/2$ with probability p and $-1/2$ with probability $1 - p$. Symmetrically, $(\min L(t))_{t \in \mathbb{N}}$ evolves as a random walk with steps $-1/2$ with probability $1 - p$ and $1/2$ with probability p . Moreover, as long as $d(t) > 1$, these two random walks are independent.

Consequently, one can check that $(d(t))_{t \geq 0}$ is a random walk on \mathbb{N} , starting from 1, absorbed in 0, with the following transition probabilities: for $i \geq 1$,

$$\begin{aligned} p_{i,i-1} &= (1-p)^2, \\ p_{i,i} &= 2p(1-p), \\ p_{i,i+1} &= p^2. \end{aligned} \tag{1}$$

and T is its absorption time. We thus have that:

- for $p < 1/2$, then $(1-p)^2 < p^2$, the random walk $(d(t))_{t \geq 0}$ has a negative bias, thus T is almost surely finite and has a finite expectation;
- for $p = 1/2$, then the random walk $(d(t))_{t \geq 0}$ is symmetric so T is almost surely finite and has an infinite expectation;
- for $p > 1/2$, then $(1-p)^2 > p^2$, the random walk $(d(t))_{t \geq 0}$ has a positive bias and thus, starting from $d(0) = 1$, it has a strictly positive probability of never hitting 0. So $\mathbb{P}(T = +\infty) > 0$.

□

Theorem 1 can directly be extended to all initial configurations containing only neutral cells and a finite number of alert cells.

3.4 Initial configuration with only one defect

In a second step, we consider the simple case where only one defect exists on the grid. Our aim is to analyse how the density of alerts evolves asymptotically. Will it grow for ever, in which case the system is said to be *transient* or will the system reach a stable regime, in which case it will visit infinitely often the same configurations and thus will be said to be *recurrent*? The following theorem answers this question.

Theorem 2. *For the BQ rule, starting from the configuration η_0 with a single defect and no alert cells in a neutral background, the first return time $T = \inf\{t \geq 1 : \eta_t = \eta_0\}$ satisfies:*

- For $p \in (0, 1/2)$, T is almost surely finite, $\mathbb{E}(T) < +\infty$ and thus the process $(\eta_t)_{t \in \mathbb{N}}$ is positive recurrent.
- For $p \in [1/2, 1)$, $\mathbb{P}(T = +\infty) > 0$ and the process $(\eta_t)_{t \in \mathbb{N}}$ is transient.

Proof. In the previous case of a single alert cell at time 0, at any time the process was either empty or one *horizontal segment*. Now, the single alert cell is replaced by a single defect cell, which generates at each time step a new horizontal segment, and these horizontal segments then evolve independently.

Let us define the *lines* of the process and their life times: for every $i \geq 1$, we define the i -th line L_i as the horizontal segment of alert states which appears at time i on the set $\{(-1/2, 1), (1/2, 1)\}$, and which then travels up by one cell at each time step until it disappears. More formally, for $t \geq 0$,

$$\begin{aligned} L_i(t) &= \{x : \eta_{i+t}(x, t+1) = 1\}, \\ d_i(t) &= \max L_i(t) - \min L_i(t) + 1, \text{ with the convention } d_i(t) = 0 \text{ if } L_i(t) = \emptyset, \\ \tau_i &= \inf\{t \in \mathbb{N} : d_i(t) = 0\}. \end{aligned} \quad (2)$$

Note that at time i , the i -th line can either be empty (and remain so) or a horizontal segment with one or two cells, in which case it will remain a horizontal segment until it disappears. The quantity $d_i(t)$ is the length of the i -th line after t time steps and τ_i is its life time (see Fig. 4).




Figure 4: The single defect is at the origin and shown in black. On the left, $t = 6$: the first line is born at time 1, and is now $L_1(5)$ in blue, while the sixth line $L_6(0)$ is just born in pink. On the right, $t = 7$, the first line has evolved in $L_1(6)$, it is displayed in blue, and the sixth line $L_6(1)$ is shown in pink.

Moreover, for $i \geq 1$, the processes $(d_i(t))_{t \geq 0}$ are independent and identically distributed random walks; their transition probabilities are given by Eq. (1).

To study the recurrence of our process $(\eta_t)_{t \in \mathbb{N}}$, we now introduce a queue $(R_t)_{t \geq 0}$. We start from an empty queue and we say that at each time $t \geq 1$, the line L_t enters the queue, and it leaves the queue at time $t + \tau_t$. For $t \geq 0$, we denote by R_t the maximal remaining living time of the lines that are in the queue at time t : it is a Markov process with values in \mathbb{N} , which satisfies the recursion relation $R_0 = 0$ and $\forall t \geq 0, R_{t+1} = \max(R_t - 1, \tau_{t+1})$ where $(\tau_i)_{i \geq 1}$ are the independent and identically distributed random variables defined in (2). The first return time T of the process $(\eta_t)_{t \in \mathbb{N}}$ to the state η_0 corresponds to the first return time of the Markov chain $(R_t)_{t \geq 0}$ to the state 0.

Lemma 1. *Assume that τ_1 is integrable, and that $\mathbb{P}(\tau_1 = 0) > 0$, then $(R_t)_{t \geq 0}$ is positive recurrent, that is, the expected amount of time to return to $R_t = 0$ has a finite first moment.*

The proof of this lemma is given in the Appendix.

- If $p < 1/2$, as noticed before, the random walk $(d_1(t))_{t \geq 0}$ has a negative bias and its absorption time τ_1 has a finite expectation. The previous lemma ensures thus that $(R_t)_{t \geq 0}$ is positive recurrent, and so is $(\eta_t)_{t \geq 0}$.
- If $p > 1/2$, the transience of $(\eta_t)_{t \geq 0}$ follows from Theorem 1 and from Fact 1.
- If $p = 1/2$, $\mathbb{P}(\tau_1 = 0) = 1/4$ and it is a classical result for the symmetric random walk that (in the following lines, C denotes some strictly positive constant whose precise value can change from line to line)

$$\mathbb{P}(\tau_1 = k | \tau_1 > 0) \sim \frac{C}{k^{3/2}} \text{ and } G(k) \stackrel{\text{def}}{=} \mathbb{P}(\tau_1 \geq k) \sim \frac{C}{k^{1/2}}.$$

Now we have, for $t > 0$,

$$\mathbb{P}(R_t = 0) = \prod_{i=1}^t \mathbb{P}(\tau_i \leq t - i) = \prod_{i=0}^{t-1} \mathbb{P}(\tau_1 \leq i) = \exp\left(\sum_{i=1}^t \ln(1 - G(i))\right),$$

$$\sum_{i=1}^t \ln(1 - G(i)) \sim -C\sqrt{t}, \text{ and thus } \mathbb{P}(R_t = 0) = \mathcal{O}\left(\exp(-C\sqrt{t})\right).$$

The Borel-Cantelli lemma ensures then that the chain is transient. \square

The results of the theorem can be extended to any initial configuration with a finite positive number of D's and a finite number of 1's. To return to the diagnosis problem, remark that for $p > 1/2$, the transient regime corresponds to a situation where the density of alerts always grows, which implies that we will not be able to discriminate between a small and a high density of failures. In other words, only the case $p < 1/2$ is interesting.

3.5 Initial configuration with a positive density of defects

For some $\kappa \in (0, 1)$, let us denote by π_κ the probability distribution on $\{0, 1, D\}^{\mathbb{Z}^2}$ for which cells are in state D with probability κ , and in state 0 with probability $1 - \kappa$, independently for different cells. In this section, we study the evolution of the BQ rule from an initial configuration η_0 distributed according to π_κ , and show (Theorem 3) that there exists a threshold value $\kappa^* > 0$ such that if $\kappa > \kappa^*$, then state 1 invades the grid (invasion regime), while it is not the case if $\kappa < \kappa^*$ (coexistence regime).

The proof of this theorem strongly relies on a comparison of the BQ rule with the *bootstrap process* $(\eta'_t)_{t \in \mathbb{N}}$, defined on $\{0, 1\}^{\mathcal{L}}$ as follows:

$$\eta'_{t+1}(x, n) = 1 \text{ if and only if } \eta'_t(x, n) = 1$$

$$\text{or } \eta'_t(x - 1/2, n - 1) = \eta'_t(x + 1/2, n - 1) = 1.$$

It is known that there exists a threshold $\beta^* \in (0, 1)$ such that from an initial configuration η'_0 where each cell, independently, has a probability β to be in state 1:

- for $\beta > \beta^*$, state 1 invades the grid, meaning that for all $(x, n) \in \mathcal{L}$, $\lim_{n \rightarrow \infty} \mathbb{P}(\eta'_t(x, n) = 0) = 0$.
- for $\beta < \beta^*$, state 1 does not invade the grid, meaning that for all $(x, n) \in \mathcal{L}$, $\lim_{n \rightarrow \infty} \mathbb{P}(\eta'_t(x, n) = 0) > 0$.

The bootstrap rule is related to directed percolation, and it can be proven that $\beta^* = 1 - p_c$, where p_c is the probability threshold for directed site percolation on the square lattice. The exact value of β^* is not known analytically, and numerical simulations give $\beta^* \approx 0.29$.

We are now able to state the following result.

Theorem 3. *For $p \in (0, \beta^*)$, there exists $\kappa^*(p) > 0$ such that from an initial configuration η_0 distributed according to π_κ ,*

- for $\kappa > \kappa^*(p)$, state 1 invades the grid, meaning that

$$\forall (x, n) \in \mathcal{L}, \quad \lim_{n \rightarrow \infty} \mathbb{P}(\eta_t(x, n) = 0) = 0,$$

- for $\kappa < \kappa^*(p)$, state 1 does not invade the grid.

Proof. For the above bootstrap rule with initial density κ , the state $\eta'_n(\cdot, n)$ of line n at time n has the same distribution as $\omega'_n(\cdot, n)$, where the sequence of configurations $(\omega'_n)_{n \in \mathbb{N}}$ is defined by

$$\omega'_n(x, n) = \psi'(\omega'_{n-1}(x - 1/2, n - 1), \omega'_{n-1}(x + 1/2, n - 1)), \text{ with:}$$

- if $q = q' = 1$ then $\psi'(q, q') = 1$,
- otherwise, $\psi'(q, q') = \begin{cases} 1 & \text{with prob. } \kappa, \\ 0 & \text{with prob. } 1 - \kappa. \end{cases}$

and with an initial configuration ω'_0 where cells are in state 1 with probability κ on line 0. This follows from the observation that for the bootstrap rule, the state of line $n + 1$ at time $n + 1$ only depends on the state of line n at time n and of the positions of the cells initially in state 1 on line $n + 1$. For each cell independently, this event has a probability κ .

The process $(\omega'_n(\cdot, n))_{n \in \mathbb{N}}$ corresponds to the evolution of a one-dimensional probabilistic cellular automaton, known under the name of *Stavskaya's rule* [6].

We are interested in the evolution of the BQ rule from an initial configuration $\eta_0 \sim \pi_d$. In the same way as for the bootstrap rule, the state $\eta_n(\cdot, n)$ of line n at time n has the same distribution as $\omega_n(\cdot, n)$, where the sequence of configurations $(\omega_n)_{n \in \mathbb{N}}$ is defined by

$$\omega_n(x, n) = \psi(\omega_{n-1}(x - 1/2, n - 1), \omega_{n-1}(x + 1/2, n - 1)), \text{ with:}$$

- if $q, q' \in \{1, \text{D}\}$ then $\psi(q, q') = \begin{cases} \text{D} & \text{with prob. } \kappa, \\ 1 & \text{with prob. } 1 - \kappa, \end{cases}$
- if $q = q' = 0$ then $\psi(q, q') = \begin{cases} \text{D} & \text{with prob. } \kappa, \\ 0 & \text{with prob. } 1 - \kappa, \end{cases}$
- otherwise, $\psi(q, q') = \begin{cases} \text{D} & \text{with prob. } \kappa, \\ 1 & \text{with prob. } (1 - \kappa)p, \\ 0 & \text{with prob. } (1 - \kappa)(1 - p) \end{cases},$

and with an initial configuration ω_0 where cells are in state D with probability κ on line 0. Furthermore, by stationarity, for any $n \in \mathbb{N}$, the distribution of η_n on line n is the same as the distribution of η_0 on line 0.

We can identify states 1 and D into a joined state denoted by $\underline{1}$. The local rule is then rewritten as follows:

$$\psi(\underline{1}, \underline{1}) = \underline{1}, \quad \psi(0, 0) = \begin{cases} \underline{1} & \text{with prob. } \kappa, \\ 0 & \text{with prob. } 1 - \kappa, \end{cases}$$

$$\psi(0, \underline{1}) = \psi(\underline{1}, 0) = \begin{cases} \underline{1} & \text{with prob. } \kappa + (1 - \kappa)p, \\ 0 & \text{with prob. } (1 - \kappa)(1 - p). \end{cases}$$

This rule dominates the bootstrap rule ψ' in the following sense: one can couple the evolutions of the two processes in such a way that there are more $\underline{1}$'s in ω than 1's in ω' . Consequently, if $\kappa > \beta^*$, this implies that the $\underline{1}$'s invade the grid.

Conversely, if $\kappa + (1 - \kappa)p < \beta^*$, the comparison with the bootstrap rule of parameter $\kappa + (1 - \kappa)p$ implies that the $\underline{1}$'s do not invade the grid.

Furthermore, by Fact 1, the BQ rule is stochastically increasing with κ . It follows that for a fixed $p \in (0, \beta^*)$, when we increase κ , we will go from the coexistence regime, where 1's do not invade the grid (valid at least as long as $\kappa + (1 - d)\kappa < \beta^*$), to the invasion regime, where 1's invade the grid (valid at least when $\kappa > \beta^*$), which ends the proof. \square

We conjecture that Theorem 3 still holds for $p \in (0, 1/2)$. More precisely, we claim that $\lim_{p \rightarrow 1/2} \kappa^*(p) = 0$, and $\lim_{p \rightarrow 0} \kappa^*(p) = \beta^*$. In other words, this would mean that the BQ rule allows us to detect any threshold of defects between 0 and β^* .

To sum up, we presented two different rules to answer the decentralised diagnosis problem. The first one, the QuorumD rule, has the advantage to be isotropic, which makes us believe that it could be applied to various other networks, in particular to irregular graphs where the number of neighbours varies from cell to cell. On the other hand, its formal analysis is quite difficult and it is an open problem to understand precisely how the density of alerts evolves according to the different parameters. By contrast, the BQ rule introduces an

artificial direction for the flow of information but this anisotropy makes the model tractable and we were able to show analytically that a qualitative change of behaviour occurs as the density of defects increases. In both cases, the study needs to be deepened and other local rules need to be studied to understand more clearly how to implement distributed consensus algorithms that would realise a form of self-diagnosis.

References

- [1] Hugues Berry and Nazim Fatès. Robustness of the critical behaviour in the stochastic Greenberg-Hastings cellular automaton model. *International Journal of Unconventional Computing*, 7(1-2):65–85, 2011.
- [2] J. Chalupa, P. L. Leath, and G. R. Reich. Bootstrap percolation on a Bethe lattice. *Journal of Physics C: Solid State Physics*, 12(1):L31, 1979.
- [3] Nicolas Gauville, Nazim Fatès, and Irène Marcovici. Diagnostic décentralisé à l'aide d'automates cellulaires. In Olivier Simonin and Stéphanie Combettes, editors, *Systèmes Multi-Agents et simulation - Vingt-septièmes journées francophones sur les systèmes multi-agents, JFSMA 2019*, pages 107–116. Cépaduès, 2019.
- [4] Leonardo Reyes and David Laroze. Cellular automata for excitable media on a complex network: The effect of network disorder in the collective dynamics. *Physica A: Statistical Mechanics and its Applications*, 588:126552, 2022.
- [5] Leonardo I. Reyes. Greenberg-Hastings dynamics on a small-world network: the collective extinct-active transition. arXiv 1505.00182, 2017.
- [6] Lorenzo Taggi. Critical probabilities and convergence time of percolation probabilistic cellular automata. *Journal of Statistical Physics*, pages 853–892, 2015.

Appendix

We recall Lemma 1.

Assume that τ_1 is integrable (i.e. $\mathbb{E}(\tau_1) < +\infty$), and that $\mathbb{P}(\tau_1 = 0) > 0$, then $(R_t)_{t \geq 0}$ is positive recurrent, that is, the expected amount of time to return to $R_t = 0$ has a finite first moment.

Proof. It is not difficult to check that $(R_t)_{t \geq 0}$ is irreducible. Let us denote by \mathbb{E}_a the expectation of this Markov chain conditionally to starting in state a . To prove the positive recurrence, we apply Foster's theorem, that we recall now in this context: if we can find a function $f : \mathbb{N} \rightarrow \mathbb{R}_+$ such that

- $\forall a \in E, \mathbb{E}_a(f(R_1)) < +\infty$,
- there exists $\varepsilon > 0$ and a *finite* subset F of \mathbb{N} such that $\forall a \in \mathbb{N} \setminus F, \mathbb{E}_a(f(R_1)) \leq f(a) - \varepsilon$,

then $(R_t)_{t \geq 0}$ is positive recurrent. In our case, we simply take f as the identity on \mathbb{N} . Since τ_1 is integrable, we have:

$$\mathbb{E}_a(R_1) = \mathbb{E}(\max(a - 1, \tau_1)) < +\infty,$$

and

$$\begin{aligned} \mathbb{E}_a(R_1) - a &= \mathbb{E}(\max(a - 1, \tau_1)) - a \\ &= \mathbb{E}(\tau_1 \mathbf{1}_{\tau_1 \geq a}) + (a - 1)\mathbb{P}(\tau_1 < a - 1) - a \\ &\leq \mathbb{E}(\tau_1 \mathbf{1}_{\tau_1 \geq a}) - \mathbb{P}(\tau_1 \leq a). \end{aligned}$$

This last upper bound goes to -1 when a tends to $+\infty$. Taking $\varepsilon = 1/2$, we fix $v > 0$ such that for any $a \geq v$, $\mathbb{E}_a(R_1) - a \leq -1/2$, and set $F = \{0, \dots, v\}$: it is a finite set and Foster's theorem applies. \square