A decentralised diagnosis method with probabilistic cellular automata

Irène Marcovici

Joint work with Nazim Fatès and Régine Marchand

Laboratoire de Mathématiques Raphaël Salem, Université de Rouen Normandie

GrHyDy 2024
October 23, 2024

Outline of the presentation

1. The decentralised diagnosis problem
2. First rule (isotropic)
3. A little digression: bootstrap percolation
4. Second rule (directional)

The decentralised diagnosis problem

- Distributed network, made of interconnected components that interact with their neighbours.

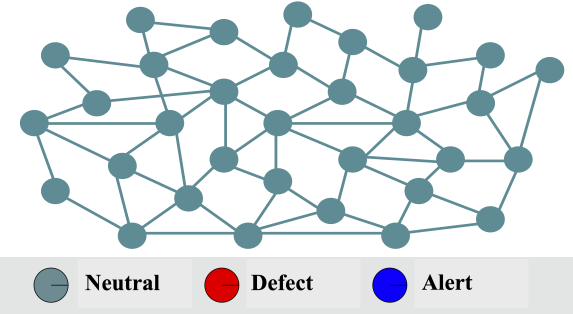

The decentralised diagnosis problem

- Distributed network, made of interconnected components that interact with their neighbours.
- The components can be subject to failures.

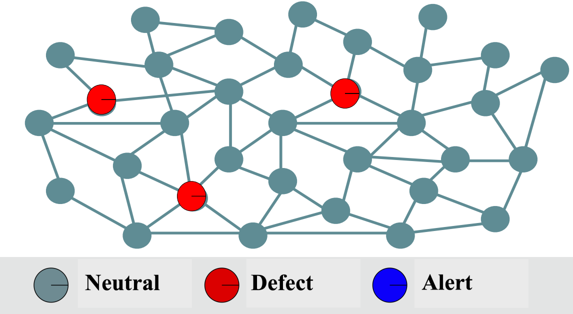

The decentralised diagnosis problem

- When the density of failures is below a given threshold, some components should still operate in the neutral state.

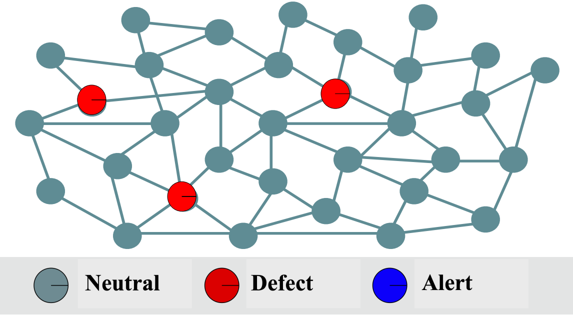

The decentralised diagnosis problem

- When the density of failures is below a given threshold, some components should still operate in the neutral state.

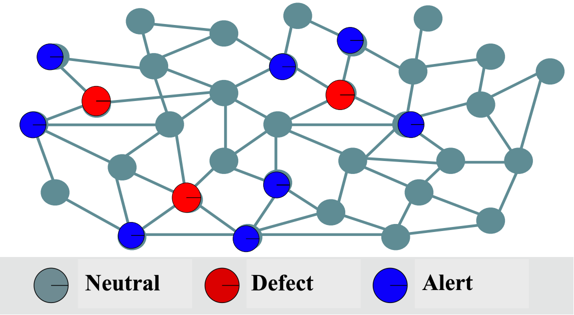

The decentralised diagnosis problem

- When the density of failures is below a given threshold, some components should still operate in the neutral state.
- When the density is beyond this threshold, all the components should be in the alert state.

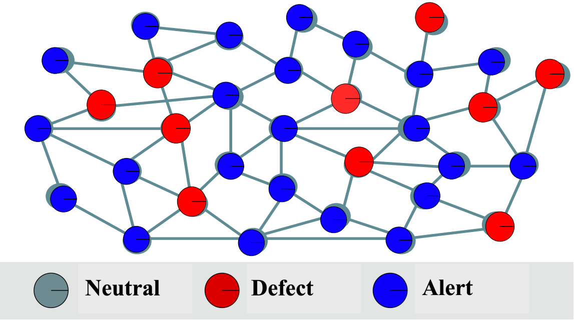

The decentralised diagnosis problem

Aim: give the entities a set of instructions to achieve this.

The goal is to gather a global information by exchanging only local information (consensus-building / self-organization).

The decentralised diagnosis problem

Aim: give the entities a set of instructions to achieve this.

The goal is to gather a global information by exchanging only local information (consensus-building / self-organization).

Difficulties:

- it is not possible to centralize information (entities all play the same role),
- conventional counting techniques cannot be used (entities have a limited memory).

The decentralised diagnosis problem

Lattice $\mathbb{Z}^2$, on which each cell can be in 3 possible states:

- N : neutral,
- D : defect (it is a fixed state),
- A : alert.

Objective

Find the simplest possible local rule such that, starting from an initial config. with N and D (indep.), state A invades the grid iff the density of D exceeds a certain threshold (which we would like to be able to choose).

Cellular automata

Let S be a finite set, and $d \geq 1$.

A function $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ is a **cellular automaton** if there exists a finite neighbourhood $\mathcal{N} \subset \mathbb{Z}^d$ a local function $f : S^{\mathcal{N}} \to S$ such that :

$$\forall x \in S^{\mathbb{Z}^d},\ \forall k \in \mathbb{Z}^d, \quad F(x)_k = f((x_{k+i})_{i \in \mathcal{N}}).$$

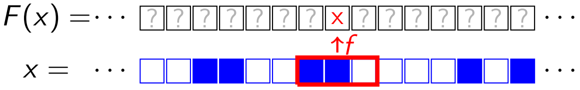

$S = \{\square, \blacksquare\}$, $\mathcal{N} = \{-1, 0, 1\}$

Cellular automata

Let S be a finite set, and $d \geq 1$.

A function $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ is a **cellular automaton** if there exists a finite neighbourhood $\mathcal{N} \subset \mathbb{Z}^d$ a local function $f : S^{\mathcal{N}} \to S$ such that :

$$\forall x \in S^{\mathbb{Z}^d},\ \forall k \in \mathbb{Z}^d, \quad F(x)_k = f((x_{k+i})_{i \in \mathcal{N}}).$$

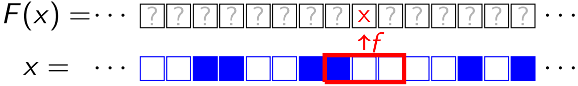

$S = \{□, ■\}$, $\mathcal{N} = \{-1, 0, 1\}$

Let S be a finite set, and $d \geq 1$.

A function $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ is a **cellular automaton** if there exists a finite neighbourhood $\mathcal{N} \subset \mathbb{Z}^d$ a local function $f : S^{\mathcal{N}} \to S$ such that :

$$\forall x \in S^{\mathbb{Z}^d},\ \forall k \in \mathbb{Z}^d, \quad F(x)_k = f((x_{k+i})_{i \in \mathcal{N}}).$$

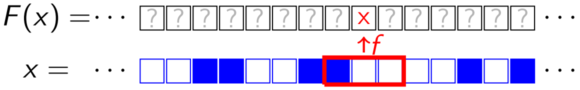

$S = \{□, ■\}$, $\mathcal{N} = \{-1, 0, 1\}$

For **probabilistic** CA, the local rule gives, for each element of $S^{\mathcal{N}}$, the probability of each state.

$$f : S^{\mathcal{N}} \to \mathcal{M}(S)$$

Let S be a finite set, and $d \geq 1$.

A function $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ is a **cellular automaton** if there exists a finite neighbourhood $\mathcal{N} \subset \mathbb{Z}^d$ a local function $f : S^{\mathcal{N}} \to S$ such that :

$$\forall x \in S^{\mathbb{Z}^d},\ \forall k \in \mathbb{Z}^d, \quad F(x)_k = f((x_{k+i})_{i \in \mathcal{N}}).$$

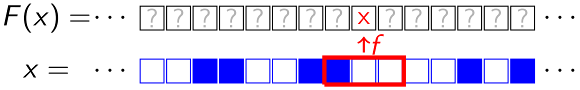

$S = \{\square, \blacksquare\}$, $\mathcal{N} = \{-1, 0, 1\}$

For **probabilistic** CA, the local rule gives, for each element of $S^{\mathcal{N}}$, the probability of each state.

$$f : S^{\mathcal{N}} \to \mathcal{M}(S) \qquad F : \mathcal{M}(S^{\mathbb{Z}^d}) \to \mathcal{M}(S^{\mathbb{Z}^d})$$

Outline of the presentation

1. The decentralised diagnosis problem
2. First rule (isotropic)
3. A little digression: bootstrap percolation
4. Second rule (directional)

First rule (isotropic)

- If all neighbours are N, the new state is N.
- If all neighbours are A or D, the new state is A.
- Otherwise, the new state of the cell is:

$$
\begin{aligned}
N \text{ with proba. } & \frac{\exp(\lambda n_N)}{\exp(\lambda n_N) + \exp(\lambda(n_A + n_D))} \\
A \text{ with proba. } & \frac{\exp(\lambda(n_A + n_D))}{\exp(\lambda n_N) + \exp(\lambda(n_A + n_D))},
\end{aligned}
$$

where n_N, n_A, n_D are resp. the nb. of neighbours N, A, D.

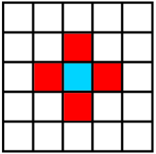

von Neumann neighbourhood

Evolution of a random configuration with $L = 32$, and $\lambda = 1$.

Density of defects $d_D = 0.1$: the alert state invades the whole grid.

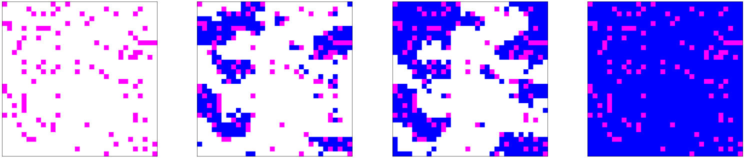

$t = 0$ $t = 20$ $t = 40$ $t = 260$

First rule (isotropic)

Evolution of a random configuration with $L = 32$, and $\lambda = 1$.

Density of defects $d_D = 0.1$: the alert state invades the whole grid.

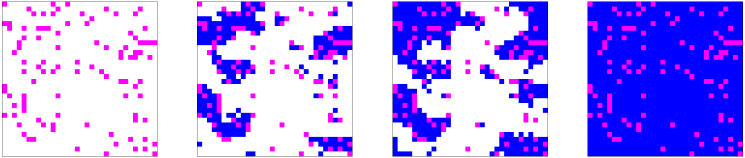

$t = 0$ $t = 20$ $t = 40$ $t = 260$

Density of defects $d_D = 0.02$: the diffusion of the alert state remains bounded.

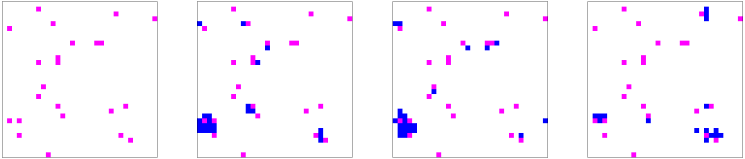

$t = 0$ $t = 1000$ $t = 10000$ $t = 40000$

Experimental observations:

- ⊕ suitable for practical use on finite grids (the parameter λ allows to adjust the threshold of defects),

Experimental observations:

- ⊕ suitable for practical use on finite grids (the parameter λ allows to adjust the threshold of defects),
- ⊖ the experimental threshold depends on the size of the grid.

First rule (isotropic)

Experimental observations:

- ⊕ suitable for practical use on finite grids (the parameter λ allows to adjust the threshold of defects),
- ⊖ the experimental threshold depends on the size of the grid.

$\rightsquigarrow$ It certainly does not work on the entire $\mathbb{Z}^2$ lattice.

Outline of the presentation

1. The decentralised diagnosis problem
2. First rule (isotropic)
3. A little digression: bootstrap percolation
4. Second rule (directional)

A little digression: bootstrap percolation

- Cells in state 1 always remain in state 1.
- Cells in state 0 with ≥ 2 neighb. in state 1 become in state 1.

Simulation with $d_1 = 0.05$.

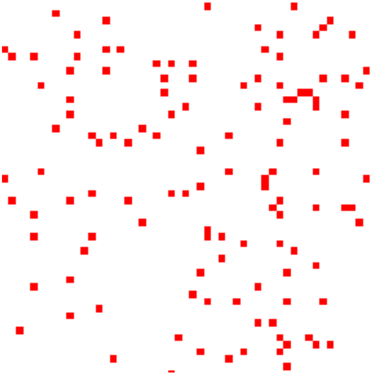

Step 0

A little digression: bootstrap percolation

- Cells in state 1 always remain in state 1.
- Cells in state 0 with ≥ 2 neighb. in state 1 become in state 1.

Simulation with $d_1 = 0.05$.

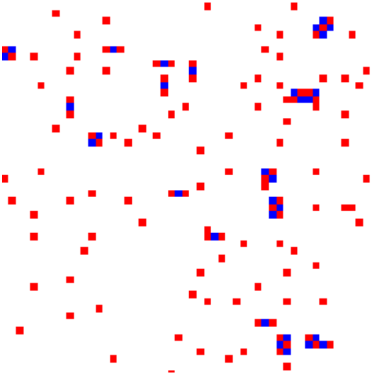

Step 1

A little digression: bootstrap percolation

- Cells in state 1 always remain in state 1.
- Cells in state 0 with ≥ 2 neighb. in state 1 become in state 1.

Simulation with $d_1 = 0.05$.

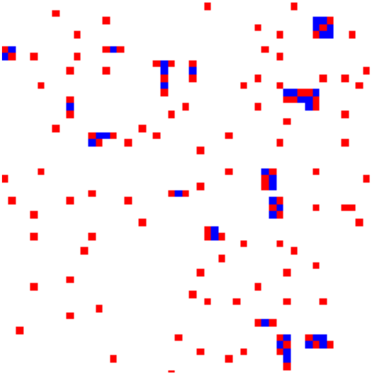

Step 2

- Cells in state 1 always remain in state 1.
- Cells in state 0 with ≥ 2 neighb. in state 1 become in state 1.

Simulation with $d_1 = 0.05$.

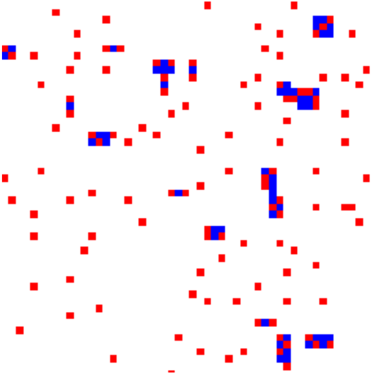

Step 3

- Cells in state 1 always remain in state 1.
- Cells in state 0 with ≥ 2 neighb. in state 1 become in state 1.

Simulation with $d_1 = 0.05$.

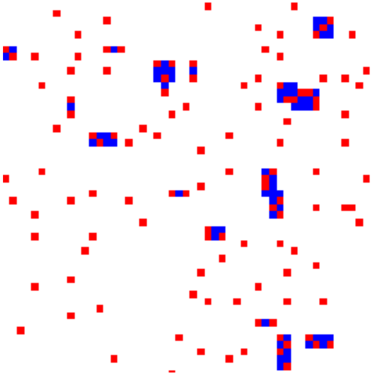

Step 4

A little digression: bootstrap percolation

- Cells in state 1 always remain in state 1.
- Cells in state 0 with ≥ 2 neighb. in state 1 become in state 1.

Simulation with $d_1 = 0.05$.

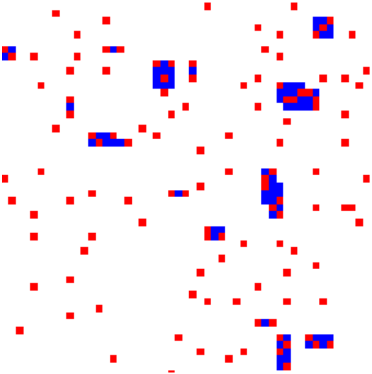

Step 5

A little digression: bootstrap percolation

- Cells in state 1 always remain in state 1.
- Cells in state 0 with ≥ 2 neighb. in state 1 become in state 1.

Simulation with $d_1 = 0.05$.

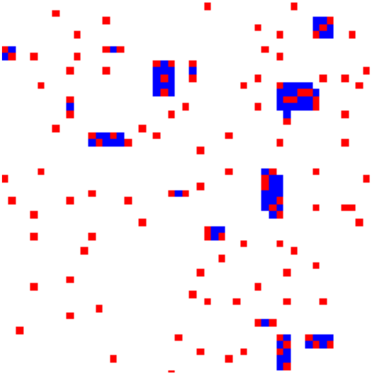

Step 6

- Cells in state 1 always remain in state 1.
- Cells in state 0 with ≥ 2 neighb. in state 1 become in state 1.

Simulation with $d_1 = 0.05$.

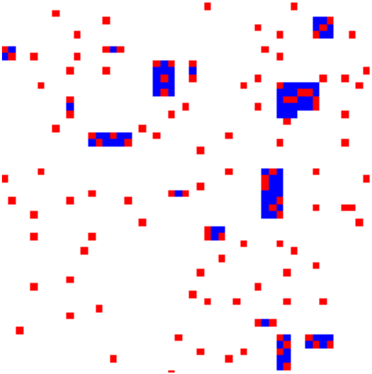

Step 7

- Cells in state 1 always remain in state 1.
- Cells in state 0 with ≥ 2 neighb. in state 1 become in state 1.

Simulation with $d_1 = 0.05$.

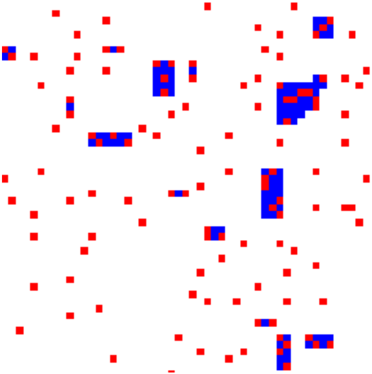

Step 8

A little digression: bootstrap percolation

- Cells in state 1 always remain in state 1.
- Cells in state 0 with ≥ 2 neighb. in state 1 become in state 1.

Simulation with $d_1 = 0.05$.

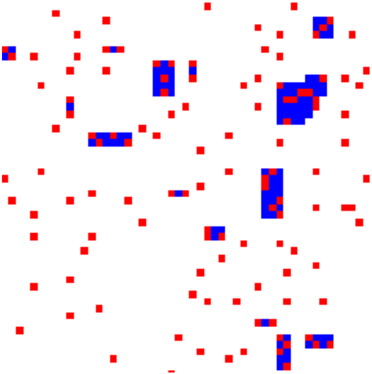

Step 9

- Cells in state 1 always remain in state 1.
- Cells in state 0 with ≥ 2 neighb. in state 1 become in state 1.

Simulation with $d_1 = 0.05$.

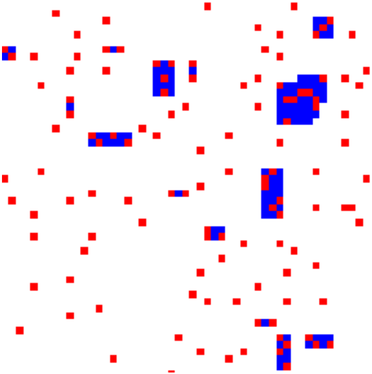

Step 10

A little digression: bootstrap percolation

- Cells in state 1 always remain in state 1.
- Cells in state 0 with ≥ 2 neighb. in state 1 become in state 1.

Simulation with $d_1 = 0.05$.

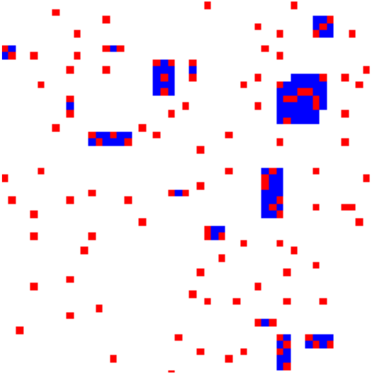

Step 11

- Cells in state 1 always remain in state 1.
- Cells in state 0 with ≥ 2 neighb. in state 1 become in state 1.

Simulation with $d_1 = 0.05$.

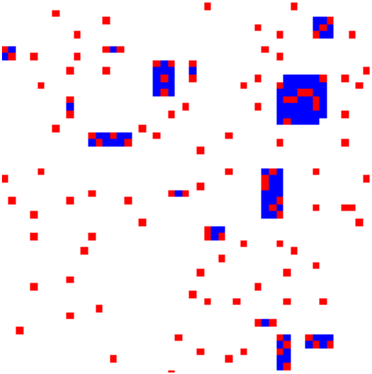

Step 12

- Cells in state 1 always remain in state 1.
- Cells in state 0 with ≥ 2 neighb. in state 1 become in state 1.

Simulation with $d_1 = 0.05$.

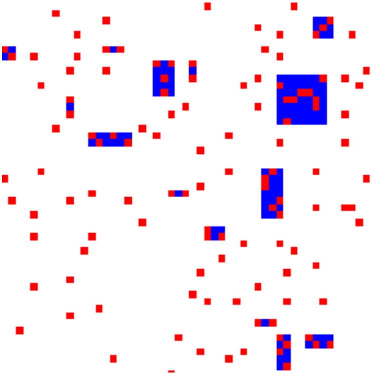

Step 13

- Cells in state 1 always remain in state 1.
- Cells in state 0 with ≥ 2 neighb. in state 1 become in state 1.

Simulation with $d_1 = 0.05$.

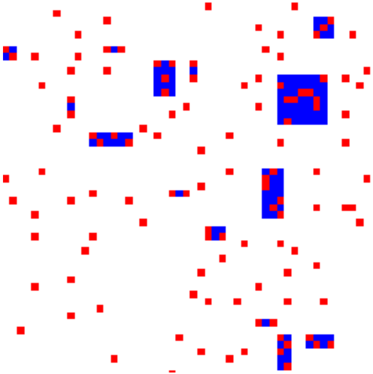

Step 14

- Cells in state 1 always remain in state 1.
- Cells in state 0 with ≥ 2 neighb. in state 1 become in state 1.

Simulation with $d_1 = 0.05$.

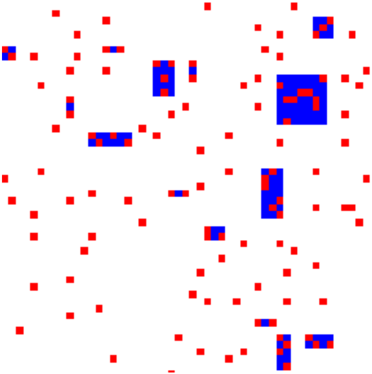

Step 15

A little digression: bootstrap percolation

Theorem [van Enter 1987]

For any $d_1 > 0$, starting from $\mathcal{B}(d_1)^{\otimes \mathbb{Z}^2}$, the bootstrap CA converges to the "all 1" configuration.

Theorem [van Enter 1987]

For any $d_1 > 0$, starting from $\mathcal{B}(d_1)^{\otimes \mathbb{Z}^2}$, the bootstrap CA converges to the "all 1" configuration.

Idea: prove that there is somewhere in the initial configuration an "all 1" square from which the whole configuration will be invaded.

This will happen iff this square is not surrounded by an "all 0" rectangle.

Theorem [van Enter 1987]

For any $d_1 > 0$, starting from $\mathcal{B}(d_1)^{\otimes \mathbb{Z}^2}$, the bootstrap CA converges to the "all 1" configuration.

Idea: prove that there is somewhere in the initial configuration an "all 1" square from which the whole configuration will be invaded.

This will happen iff this square is not surrounded by an "all 0" rectangle.

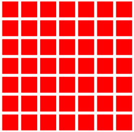

A little digression: bootstrap percolation

Theorem [van Enter 1987]

For any $d_1 > 0$, starting from $\mathcal{B}(d_1)^{\otimes \mathbb{Z}^2}$, the bootstrap CA converges to the "all 1" configuration.

Idea: prove that there is somewhere in the initial configuration an "all 1" square from which the whole configuration will be invaded.

This will happen iff this square is not surrounded by an "all 0" rectangle.

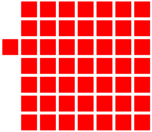

A little digression: bootstrap percolation

Theorem [van Enter 1987]

For any $d_1 > 0$, starting from $\mathcal{B}(d_1)^{\otimes \mathbb{Z}^2}$, the bootstrap CA converges to the "all 1" configuration.

Idea: prove that there is somewhere in the initial configuration an "all 1" square from which the whole configuration will be invaded.

This will happen iff this square is not surrounded by an "all 0" rectangle.

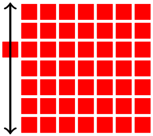

A little digression: bootstrap percolation

Theorem [van Enter 1987]

For any $d_1 > 0$, starting from $\mathcal{B}(d_1)^{\otimes \mathbb{Z}^2}$, the bootstrap CA converges to the "all 1" configuration.

Idea: prove that there is somewhere in the initial configuration an "all 1" square from which the whole configuration will be invaded.

This will happen iff this square is not surrounded by an "all 0" rectangle.

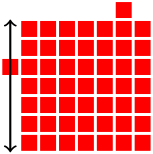

Theorem [van Enter 1987]

For any $d_1 > 0$, starting from $\mathcal{B}(d_1)^{\otimes \mathbb{Z}^2}$, the bootstrap CA converges to the "all 1" configuration.

Idea: prove that there is somewhere in the initial configuration an "all 1" square from which the whole configuration will be invaded.

This will happen iff this square is not surrounded by an "all 0" rectangle.

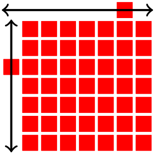

Theorem [van Enter 1987]

For any $d_1 > 0$, starting from $\mathcal{B}(d_1)^{\otimes \mathbb{Z}^2}$, the bootstrap CA converges to the "all 1" configuration.

Idea: prove that there is somewhere in the initial configuration an "all 1" square from which the whole configuration will be invaded.

This will happen iff this square is not surrounded by an "all 0" rectangle.

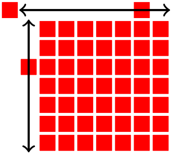

Theorem [van Enter 1987]

For any $d_1 > 0$, starting from $\mathcal{B}(d_1)^{\otimes \mathbb{Z}^2}$, the bootstrap CA converges to the "all 1" configuration.

Idea: prove that there is somewhere in the initial configuration an "all 1" square from which the whole configuration will be invaded.

This will happen iff this square is not surrounded by an "all 0" rectangle.

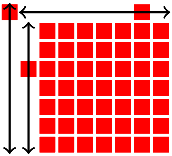

A little digression: bootstrap percolation

Theorem [van Enter 1987]

For any $d_1 > 0$, starting from $\mathcal{B}(d_1)^{\otimes \mathbb{Z}^2}$, the bootstrap CA converges to the "all 1" configuration.

Idea: prove that there is somewhere in the initial configuration an "all 1" square from which the whole configuration will be invaded.

This will happen iff this square is not surrounded by an "all 0" rectangle.

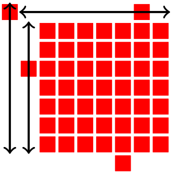

A little digression: bootstrap percolation

Theorem [van Enter 1987]

For any $d_1 > 0$, starting from $\mathcal{B}(d_1)^{\otimes \mathbb{Z}^2}$, the bootstrap CA converges to the "all 1" configuration.

Idea: prove that there is somewhere in the initial configuration an "all 1" square from which the whole configuration will be invaded.

This will happen iff this square is not surrounded by an "all 0" rectangle.

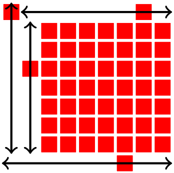

Remark: on a large $N \times N$ grid, Holroyd has proven in 2003 that

- $d_1 > \frac{\pi^2}{18 \log N} \implies$ conv. to total occupancy with high prob,
- $d_1 < \frac{\pi^2}{18 \log N} \implies$ no convergence to total occupancy.

A little digression: bootstrap percolation

NE-bootstrap percolation

- Cells in state 1 always remain in state 1.
- Cells in state 0 with North and East neighb. in state 1 become in state 1.

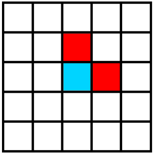

NE-bootstrap percolation

- Cells in state 1 always remain in state 1.
- Cells in state 0 with North and East neighb. in state 1 become in state 1.

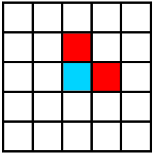

Proposition

There exists $\beta_c \in (0,1)$ such that:

- for $d_1 < \beta_c$, state 1 does not invade the whole grid,
- for $d_1 > \beta_c$, state 1 invades the grid.

$\beta_c = 1 - p_c \approx 0.29$, where $p_c =$ threshold for directed site percol.

A little digression: bootstrap percolation

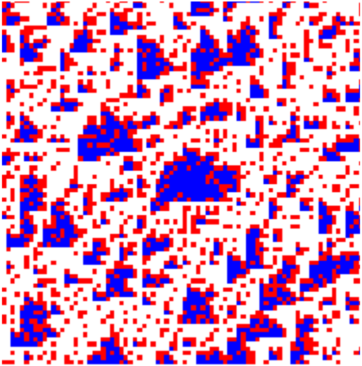

$d_1 = 0.25$

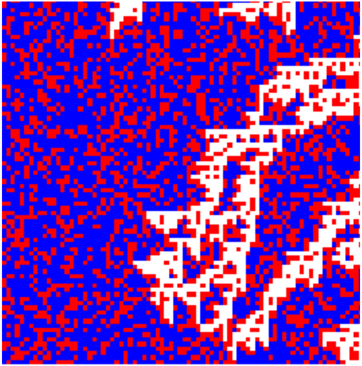

$d_1 = 0.35$

(150 iterations)

Outline of the presentation

1. The decentralised diagnosis problem
2. First rule (isotropic)
3. A little digression: bootstrap percolation
4. Second rule (directional)

Second rule (directional)

Previous rule on NE-neighb., with additional simplifications...

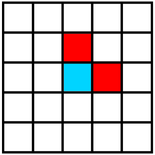

- If the two (red) neighb. are N, the new state is N.
- If the two (red) neighb. are A or D, the new state is A.
- Otherwise, state is A with prob. p and N with prob. $1 - p$.

Second rule (directional)

Parameter $p = 0.4$ (200 iterations)

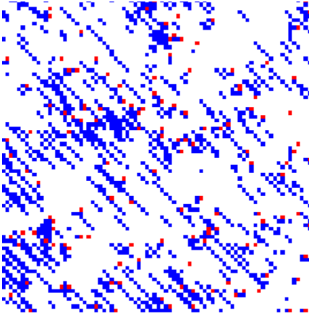

$d_D = 0.02$

Second rule (directional)

Parameter $p = 0.4$ (200 iterations)

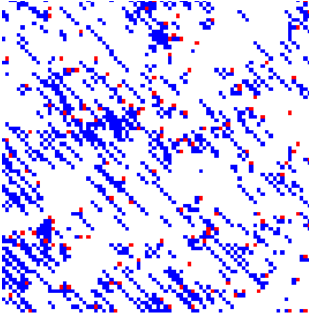

$d_D = 0.02$

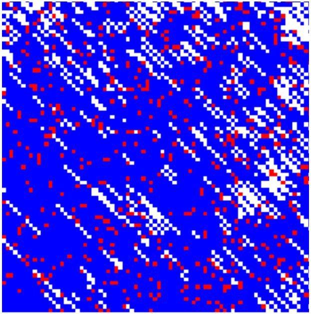

$d_D = 0.08$

Second rule (directional)

Proposition [Fatès-Marchand-M. 2023]

Initial config. with a single cell in state A and no D

T = extinction time

- For $p < 1/2$, T is a.s. finite and $\mathbb{E}(T) < +\infty$.
- For $p = 1/2$, T is a.s. finite and $\mathbb{E}(T) = +\infty$.
- For $p > 1/2$, $\mathbb{P}(T = +\infty) > 0$.

Second rule (directional)

Proposition [Fatès-Marchand-M. 2023]

Initial config. with a single cell in state A and no D

T = extinction time

- For $p < 1/2$, T is a.s. finite and $\mathbb{E}(T) < +\infty$.
- For $p = 1/2$, T is a.s. finite and $\mathbb{E}(T) = +\infty$.
- For $p > 1/2$, $\mathbb{P}(T = +\infty) > 0$.

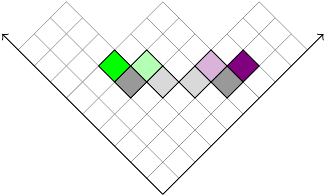

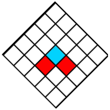

Proof: at any time, the config. is either empty or an horiz. A segment. The right boundary evolves as a random walk with steps $1/2$ with prob. p and $-1/2$ with prob. $1-p$.

Second rule (directional)

Proposition [Fatès-Marchand-M. 2023]

Initial config. with a single cell in state D and no A

T = first return time

- For $p < 1/2$, T is a.s. finite and $\mathbb{E}(T) < +\infty$.
- For $p \geq 1/2$, $\mathbb{P}(T = +\infty) > 0$.

Proof: previous proposition + Foster's theorem.

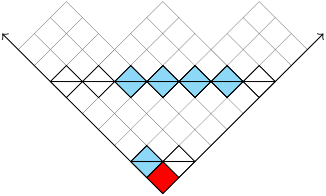

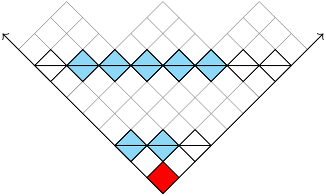

Second rule (directional)

Theorem [Fatès-Marchand-M. 2023]

For $p \in (0, \beta_c)$, there exists $d_D^c(p)$ such that:

- for $d_D > d_D^c(p)$, state A invades the grid,
- for $d_D < d_D^c(p)$, state A does not invade the whole grid.

Second rule (directional)

Theorem [Fatès-Marchand-M. 2023]

For $p \in (0, \beta_c)$, there exists $d_D^c(p)$ such that:

- for $d_D > d_D^c(p)$, state A invades the grid,
- for $d_D < d_D^c(p)$, state A does not invade the whole grid.

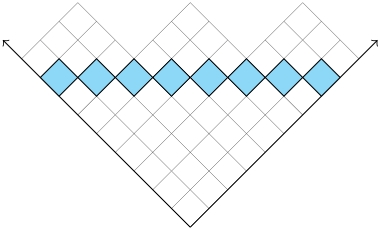

NE-bootstrap

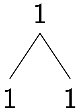

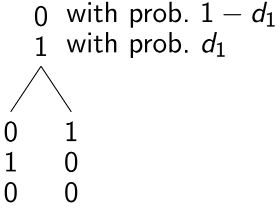

Second rule (directional)

NE-bootstrap

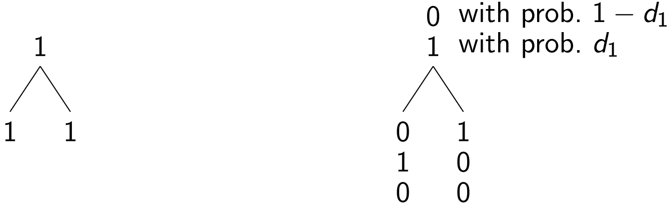

Our directional rule, with $1 = \{D, A\}, 0 = N$

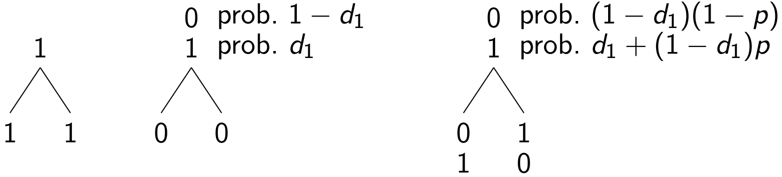

Second rule (directional)

$$\text{NE-bootstrap}(d_1) \preceq PCA(p, d_1) \preceq \text{NE-bootstrap}(d_1 + (1 - d_1)p)$$

Second rule (directional)

NE-bootstrap$(d_1) \preceq PCA(p, d_1) \preceq$ NE-bootstrap$(d_1 + (1 - d_1)p)$

- $d_1 > \beta_c \implies$ state A invades the grid
- $d_1 + (1 - d_1)p < \beta_c \implies$ state A does not invade the whole grid

The rule is stoch. increasing with d_1, so for a fixed $p \in (0, \beta_c)$, increasing d_1 takes us from the coexistence to the invasion regime.

Second rule (directional)

NE-bootstrap(d_1) $\preceq PCA(p, d_1) \preceq$ NE-bootstrap($d_1 + (1 - d_1)p$)

- $d_1 > \beta_c \implies$ state A invades the grid
- $d_1 + (1 - d_1)p < \beta_c \implies$ state A does not invade the whole grid

The rule is stoch. increasing with d_1, so for a fixed $p \in (0, \beta_c)$, increasing d_1 takes us from the coexistence to the invasion regime.

Corollary: $\frac{\beta_c - p}{1-p} \leq d_D^c(p) \leq \beta_c$, so $\lim_{p \to 0} d_D^c(p) = \beta_c$.

Conjecture: the prop. still holds for $p \in (0, 1/2)$ and $\lim_{p \to 1/2} d_D^c(p) = 0$, so one can detect any threshold of defects between 0 and β_c.

Conclusion and perspectives

- **First rule:**
  - ⊖ difficult to analyse, no phase transition on $\mathbb{Z}^2$,
  - ⊕ isotropic, can be extended to irregular networks.
- **Second rule:**
  - ⊖ artificial direction for the flow of information,
  - ⊕ tractable model, phase transition on $\mathbb{Z}^2$.
- **Perspectives:**
  - extend the analytical study of the two rules,
  - defects that appear dynamically,
  - other lattices, random graphs.

A decentralised diagnosis method with probabilistic cellular automata.
Nazim Fatès, Régine Marchand, I.M. - **AUTOMATA 2023**