



HAL
open science

Stochastic Minority on Graphs

Damien Regnault, Jean-Baptiste Rouquier, Eric Thierry

► **To cite this version:**

Damien Regnault, Jean-Baptiste Rouquier, Eric Thierry. Stochastic Minority on Graphs. 2008. hal-00272271v2

HAL Id: hal-00272271

<https://hal.science/hal-00272271v2>

Preprint submitted on 15 Apr 2008

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

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

Stochastic Minority on Graphs

Damien Regnault^{1,2}, Jean-Baptiste Rouquier^{1,2}, and Éric Thierry^{1,2,3}

¹ Université de Lyon, École Normale Supérieure de Lyon, LIP, France

² Institut des Systèmes Complexes Rhône-Alpes (IXXI)

³ CNRS, LIAFA, Université Paris 7, France

Abstract. Cellular automata have been mainly studied on very regular graphs carrying the cells (like lines or grids) and under synchronous dynamics (all cells update simultaneously). In this paper, we study how the asynchronism and the topology of cells act upon the dynamics of the classical Minority rule. Minority has been well-studied for synchronous updates. Yet, beyond its apparent simplicity, this rule yields complex behaviors when asynchronism is introduced. We investigate the transitory part as well as the asymptotic behavior of the dynamics under full asynchronism (only one random cell updates at each time step) for several types of graphs. Such a comparative study is a first step in understanding how the asynchronous dynamics is linked to the topology of the cells.

1 Introduction

Cellular automata (CA) can be seen both as a model of computation with massive parallelism and as a model for complex systems in nature. They have been studied with various fields of applications like parallel/distributed computing, physics, biology or social sciences. Most of the work concerning CA assumes that their dynamics is deterministic and synchronous (all cells update simultaneously) and that the topology of the cells is very regular (the cells are the vertices of a graph which is usually a line or a 2D or 3D grid). Such assumptions can be questioned with regard to the applications and the real life constraints. Dynamics where those assumptions are perturbed have been far less studied and their analysis is very challenging. Here is a non-exhaustive list of related works about CA in literature, the models are *stochastic CA* since the perturbations are usually introduced as stochastic processes:

- Perturbation of the updating rule: resilience to random errors [1,2], mean field analysis of general Markovian rules [3].
- Perturbation of synchronism: empirical studies about resilience to asynchronism [4,5,6,7], mathematical analysis of some 1D CA under full asynchronism (only one random cell updates at each time step) or under α -asynchronism (each cell updates independently with probability α) [8,9,10].
- Perturbation of the topology of cells: empirical studies [11].

Many rules studied in literature cannot be extended to any topology of cells: rules for lines or grids of cells may refer to specific orientations like left/right or up/down in their definition. Nevertheless a large class of CA have a definition which is independent of the graph: the *outer-totalistic CA*. Their update rule only considers the state of the updated cell

and the number of adjacent cells in each possible state (e.g. Game of Life). The *threshold CA* with states $\{0, 1\}$ are an interesting sub-class where the new state of a cell only depends on its former state and on the comparison between some threshold and the number of 1 in the neighborhood. For example, consider a simplified frequency allocation problem with an interference graph. Each vertex is a station which has to choose between two possible frequencies, say 0 or 1, when emitting. Each edge represents a possible interference between two stations if they emit with the same frequency. To avoid interferences, suppose that stations frequently update their emission frequency by choosing the one which minimizes the number of potential interferences with its neighbors. This is the *Minority* rule and the evolution of the allocated frequencies corresponds to the dynamics of this threshold CA. The Minority rule may also arise in models from biology. Consider a set of cells sharing a simple gene regulatory networks where a gene exerts a feedback inhibiting its expression. The state of a cell is whether it expresses this gene or not. Assuming that each cell starts expressing the gene when less than half of its neighbors (including itself) express it, and that otherwise it stops expressing it, leads to the Minority rule [12]. Both those models are of course extreme simplifications of any real device or phenomena, but understanding such simple rules is an indispensable step towards the study of more involved CA.

In this paper, we choose to investigate how the topology of the cells acts upon the dynamics under asynchronous updates. We focus on *Stochastic Minority* where the Minority rule applies to two possible states $\{0, 1\}$ and under full asynchronism (at each time step, only one random cell is updated with the uniform distribution). This simple rule already exhibits a surprisingly rich behavior as observed in [13,14] where it is studied for cells assembled into a torus. Minority is interesting because it is a typical CA with negative feedback. Some related stochastic models like Ising models or Hopfield nets have been studied under asynchronous dynamics (e.g. our model of asynchronism corresponds to the limit when temperature goes to 0 in the Ising model). These models are acknowledged to be harder to analyze when it comes to arbitrary graphs [15,16] or negative feedbacks [17]. Minority has been already well studied for arbitrary graphs under the synchronous regime [18] where it is shown that for any graph the dynamics eventually converges to a fixed point or a cycle of period two. As any other Markovian process on a finite number of configurations, any trajectory of Stochastic Minority ends in an attractor, i.e. a final strongly connected component of the transition graph. We study the structure of attractors and the average time to hit them, for several types of graphs: cycles (Section 3.1), cliques (Section 3.2), trees (Section 4). This hitting time can be seen as the end of the transitory part of the dynamics. The table below sums up these average hitting times: *polynomial (Poly)* or *exponential (Exp)* in the number of vertices, for the worst initial configuration. It can be compared to the worst hitting time under the synchronous dynamics.

	Fully Asynchronous	Synchronous
Path or Cycle	Poly	Exp [18]
Tree, max degree ≤ 3	Poly	Exp [18]
Tree, max degree ≥ 4	Exp	Exp [18]
Torus, von Neuman neighborhood	Poly ? [13]	Exp [18]
Torus, Moore neighborhood	Poly ? [14]	Exp [18]
Clique	Poly	Poly [18]

What is the graph parameter which enables to predict whether the average hitting time will be polynomial or exponential ? Does adding edges usually speed up or slow down the convergence ? The questions remain open, but our analyzes already disprove a simple link with the maximum degree or the mean degree, and they provide tools for further investigations about the influence of the topology on such CA dynamics.

2 The model

2.1 Topology

We consider in this part Stochastic Minority on arbitrary undirected graphs. Section 3.1, 3.2 and 4 will be respectively devoted to cycles, cliques and trees.

Definition 1 (Configuration). Let $\mathbb{G} = (\mathbb{V}, \mathbb{E})$ be a finite undirected graph with vertices \mathbb{V} and edges \mathbb{E} . $\mathbb{Q} = \{0, 1\}$ is the set of states (0 stands for white and 1 stands for black). The vertices are also called cells and $N := |\mathbb{V}|$ denotes their number. The neighborhood \mathcal{N}_i of a cell i is the set of its adjacent cells (including itself). A configuration is a function $c : \mathbb{V} \rightarrow \mathbb{Q}$ (c_i denotes the state of cell i in configuration c).

Definition 2 (Stochastic Minority). We consider the following dynamics δ that associates with each configuration c a random configuration c' obtained as follows: a cell $i \in \mathbb{V}$ is selected uniformly at random (we say that cell i is fired) and its state is updated to the minority state among its neighborhood (no change in case of equality), while all the other cells remain in their current state:

$$c'_i = \begin{cases} 1 & \text{if } \sum_{j \in \mathcal{N}_i} c_j < \frac{|\mathcal{N}_i|}{2} \text{ or } \sum_{j \in \mathcal{N}_i} c_j = \frac{|\mathcal{N}_i|}{2} \text{ and } c_i = 1 \\ 0 & \text{if } \sum_{j \in \mathcal{N}_i} c_j > \frac{|\mathcal{N}_i|}{2} \text{ or } \sum_{j \in \mathcal{N}_i} c_j = \frac{|\mathcal{N}_i|}{2} \text{ and } c_i = 0 \end{cases}$$

and $c'_k = c_k$ for all $k \neq i$. In a configuration, a cell is said active if its state changes in case the cell is fired. We denote by c^t the random variable for the configuration obtained from an initial configuration c , after t steps of the dynamics: $c^0 = c$ and $c^t = \delta^t(c)$ for all $t \geq 1$.

Definition 3 (Attractors). For the dynamics induced by δ , a set of configurations A is an attractor if for all $c, c' \in A$, the time to reach c' starting from c is finite almost surely. It is a strongly connected component with no arc leaving the component, in the transition graph where vertices are all the possible configurations and arcs (c, c') satisfy $P(\delta(c) = c') > 0$. The union of all attractors is denoted \mathcal{A} and called the limit set.

Definition 4 (Convergence). We say that the dynamics δ converges from an initial configuration c^0 to an attractor A (resp. the limit set \mathcal{A}) if the random variable $T = \min\{t : c^t \in A\}$ (resp. $T = \min\{t : c^t \in \mathcal{A}\}$) is almost surely finite.

Since we only consider finite graphs, the dynamics δ converges almost surely from any initial configuration to \mathcal{A} . The variable T is a hitting time.

2.2 Energy and Particles

As in the Ising model [16] or in Hopfield networks [17], we define a natural global parameter that one can consider to be the energy of the system since it counts the number of interactions between neighboring cells in the same state. This parameter will provide key insights on the evolution of the system.

Definition 5 (Potential). *The potential v_i of cell i is the number of its neighbors (including itself) in the same state as itself. If $v_i \leq \frac{|\mathcal{N}_i|}{2}$ then the cell is in the minority state and is thus inactive; whereas, if $v_i > \frac{|\mathcal{N}_i|}{2}$ then the cell is active. A configuration c is stable if and only if for all cell $i \in \mathbb{V}$, $v_i \leq \frac{|\mathcal{N}_i|}{2}$.*

Definition 6 (Energy). *The energy of configuration c is $E(c) = \sum_{i \in \mathbb{V}} (v_i - 1)$.*

The energy of a configuration is always non-negative. More precisely, we have the following bounds which are direct consequences of the definitions.

Proposition 7 (Energy bounds). *The energy E satisfies $2|\mathbb{E}| - 2C_{\max} \leq E \leq 2|\mathbb{E}|$, where C_{\max} is the maximum number of edges in a cut of \mathbb{G} .*

As a consequence, computing the minimum energy for arbitrary graphs is NP-hard. There exists configurations of energy 0 if and only if \mathbb{G} is bipartite. Those stable configurations are the 2-colorings of \mathbb{G} .

Fact 8 (Energy is non-increasing) *The energy is a non-increasing function of time and decreases each time a cell i with potential $> \frac{|\mathcal{N}_i|}{2}$ fires.*

Proof. When an active cell i of potential v_i fires, its potential becomes $|\mathcal{N}_i| - v_i + 1$, and the energy of the configuration becomes $E + 2|\mathcal{N}_i| - 4v_i + 2$. \square

Corollary 9. *A configuration c belongs to the limit set if and only if there is no sequence of updates that would lead the energy to decrease, i.e. if and only if $\forall t, P(E(c^t) > E(\delta(c^t)) | c^0 = c) = 0$.*

Proof. If the energy decreases when updating c to c' , then c will never be reached again (because energy is non-increasing). Reciprocally, any update that keeps the energy constant is reversible: the fired cell can be fired again to get back to the previous configuration. \square

Remark 10. Since firing a cell of odd degree makes the energy decrease, such cells are inactive in the limit set.

Definition 11. *Let c be a configuration on $\mathbb{G} = (\mathbb{V}, \mathbb{E})$, an edge $\{i, j\}$ holds a particle if $c_i = c_j$. A configuration is fully characterized (up to the black/white symmetry) by its set of particles located at $\mathbb{P} \subseteq \mathbb{E}$. Note that the converse proposition “any subset $\mathbb{P} \subseteq \mathbb{E}$ corresponds to a configuration” is true only on trees.*

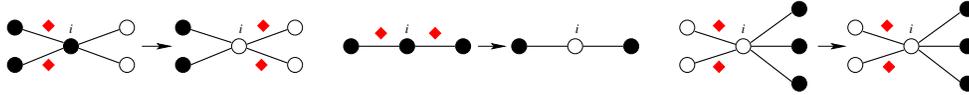


Fig. 1. Transfers of particles (red diamonds) when firing cell i .

The energy of a configuration is clearly equal to twice its number of particles. With the particle point of view, when firing a cell i of degree $\text{deg}(i) = |\mathcal{N}_i| - 1$, if the number of incident edges holding a particle is $\geq \frac{\text{deg}(i)}{2}$, these particles disappear but new particles appear on the incident edges (if any) which had no particle (as illustrated on Fig. 1). Otherwise the particles do not move.

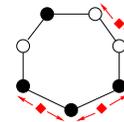
Switching between the coloring and the particle points of view may simplify the description of the configurations and the dynamics, e.g. when the energy is low and the dynamics comes to random walks of a few particles.

3 Regular graphs

In [13,14], Stochastic Minority has been studied for the torus with two different kinds of neighborhoods: the von Neumann neighborhood (4 neighbors/cell) and the Moore neighborhood (8 neighbors/cell). In both cases, the average hitting time of the limit set is conjectured to be polynomial. We investigate here two extreme cases of regular graphs: cycles and cliques.

3.1 Cycles

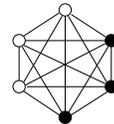
On cycles, the particle point of view is convenient and one can prove that Stochastic Minority behaves as *random walks of annihilating particles* on a discrete ring (see [8,19]). On the right, the particles are the red diamonds.



Theorem 12. *Stochastic Minority on cycles hits the limit set after $O(N^3)$ steps on expectation. If N is even, there are two attractors of energy 0: the 2-colorings of the cycle. If N is odd, there is a single attractor: a cycle in the transition graph composed of all the configurations with only one particle (energy 2).*

3.2 Cliques

We prove in this part that Stochastic Minority on cliques behaves as a *coupon collector* (see [20]). In a clique, all cells have the same neighborhood $\mathcal{N}_i = \mathbb{V}$ and thus all black (resp. white) cells have the same potential.



Theorem 13. *Stochastic Minority on cliques hits the limit set after $O(N \log N)$ steps on expectation. If N is even, there are $\binom{N}{\frac{N}{2}}$ attractors, each one is a stable configuration of energy $\frac{N^2}{2}$ which is half black and half white. If N is odd there is only one attractor which is made of the $2\binom{N}{\frac{N-1}{2}}$ configurations of energy $\frac{N^2+1}{2}$ with a difference -1 or $+1$ between black and white cells.*

4 Trees

4.1 First observations

Observations from simulations Stochastic Minority has a complex behavior on trees. There are similarities with the evolution on the torus [13,14]. The dynamics of Stochastic Minority on a random initial configuration is characterized by a fast initial energy drop. Afterwards, different regions colored by specific patterns emerge. Each region is a subtree with a 2-coloration (corresponding to minimal energy 0 inside the subtree). The energy is located on the boundaries between the regions. These regions either merge or compete with each other. Finally either one region manages to extend to the whole configuration or the boundaries between the different regions stabilize: the dynamics has reach the limit set.

Configurations of minimum energy As any bipartite graph, a tree admits exactly two configurations of energy 0 which are the two possible 2-colorings of the graph. Note that some trees admit stable configurations with non-null energy (consider an arbitrary tree all colored in black and add $|\mathcal{N}_i|$ new leaves in white to each cell i).

4.2 The Dual Rule

We now introduce dual configurations and their dual rule to ease the study of the dynamics on trees. In this dual dynamics equivalent to Stochastic Minority, the stable configurations of minimum energy are the two configurations all black and all white, and the regions which compete are all white vs. all black subtrees.

Definition 14 (Dual configurations). Consider a tree \mathbb{T} and fix a root r in \mathbb{T} . For any configuration c on \mathbb{T} , its dual configuration \hat{c} is defined as $\hat{c}_i = c_i$ if h_i is even and $\hat{c}_i = 1 - c_i$ if h_i is odd, where h_i is the distance from r to i (see Fig. 2). The mapping $c \mapsto \hat{c}$ is a bijection on the set of all the configurations: more precisely $\hat{\hat{c}} = c$.

A equivalent definition consists in making a XOR with the 2-coloring of \mathbb{T} such that r is in white (such a definition can be used for any bipartite graph, see e.g. [13]).

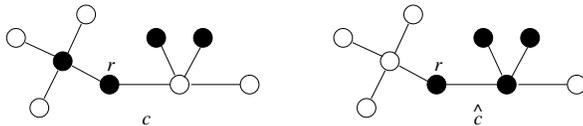


Fig. 2. A configuration c and its dual configuration \hat{c} (with regard to root r).

Proposition 15 (Dual dynamics). Consider a sequence (c^t) for the Stochastic Minority dynamics δ and the sequence (\hat{c}^t) of the dual configurations, and define the dual dynamics $\hat{\delta}$ as $\hat{\delta}(\hat{c}) = \widehat{\delta(c)}$ so that $\hat{c}^{t+1} = \hat{\delta}(\hat{c}^t)$. Then the dynamics $\hat{\delta}$ is also a stochastic CA. It

associates with each configuration \hat{c} a random configuration \hat{c}' by updating one random cell i uniformly with the rule which selects the majority state in the neighborhood of i excluding itself (in case of equality its state changes):

$$\hat{c}'_i = \begin{cases} 1 & \text{if } \sum_{j \in \mathcal{N}_i \setminus \{i\}} \hat{c}_j > \frac{|\mathcal{N}_i| - 1}{2} \text{ or } (\sum_{j \in \mathcal{N}_i \setminus \{i\}} \hat{c}_j = \frac{|\mathcal{N}_i| - 1}{2} \text{ and } \hat{c}_i = 0) \\ 0 & \text{if } \sum_{j \in \mathcal{N}_i \setminus \{i\}} \hat{c}_j < \frac{|\mathcal{N}_i| - 1}{2} \text{ or } (\sum_{j \in \mathcal{N}_i \setminus \{i\}} \hat{c}_j = \frac{|\mathcal{N}_i| - 1}{2} \text{ and } \hat{c}_i = 1) \end{cases}$$

By construction, the dual sequences (c^t) and (\hat{c}^t) as well as their corresponding dynamics δ and $\hat{\delta}$ are stochastically coupled (see [21]) by firing the same random cell at each time step.

Definition 16 (Dual potential & Energy). The dual potential \hat{v}_i of cell i is the number of its neighbors (excluding itself) in a different state than itself. If $\hat{v}_i < \frac{|\mathcal{N}_i| - 1}{2}$ then the cell is in the majority state and is thus inactive; whereas, if $\hat{v}_i \geq \frac{|\mathcal{N}_i| - 1}{2}$ then the cell is active. The dual energy \hat{E} is the sum of the dual potentials over all the cells.

Given a configuration c and its dual \hat{c} , the potential of any cell i in c is equal to the dual potential of cell i in \hat{c} plus 1. Thus the dual energy of \hat{c} is exactly the energy of c .

From now on, we consider the dual dynamics.

5 Structure of the limit set on trees

5.1 An Algorithm for the limit set

In this section we describe an algorithm (algorithm 1) that checks if a configuration belongs to the limit set, and if not, gives a sequence of updates that leads to the limit set.

Fact 17 An attractor A decomposes the graph into three sets of vertices:

1. the vertices that are in the state 0 for every configuration of A ;
2. the vertices that are in the state 1 for every configuration of A ;
3. the vertices that can be either in the state 0 or 1, depending on the configuration in A .

Algorithm 1: Membership to the limit set: check that $\hat{E}(c') = \hat{E}(c)$.

Input: A configuration c .

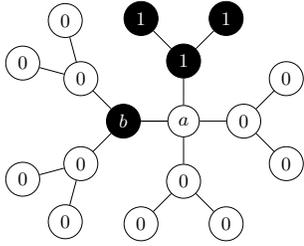
- 1 **while** There is an active cell i in state 0 **do** Fire i
- 2 **while** There is an active cell i in state 1 **do** Fire i
- 3 **while** There is an active cell i in state 0 **do** Fire i

Output: The current configuration c' .

Proposition 18. The configuration c' returned by algorithm 1 is in the limit set.

Corollary 19. The input configuration c is in the limit set if and only if the energy has not decreased during execution of the algorithm.

Three phases are necessary, as shown on Fig. 3.



Starting from this configuration, phase 1 of the algorithm makes the cell a become black. Then, phase 2 makes its left neighbor b become white, as well as a .

b is now definitely white (case 1 of Fact 17), but was not identified as such at the end of phase 1. Thus, a third phase is necessary with this algorithm to identify b .

No other cell will be active during the execution of the algorithm.

Fig. 3. Three phases are necessary.

5.2 Structure

Section 5.1 has shown that in the limit set, some vertices are fixed while other are oscillating. Let us now study more precisely the structure of the limit set.

We first pick an arbitrary node of degree 1 and set it as root r in the tree (it introduces “father” and “sons” relations). We assign a label to each cell to count the number of attractors and the size of the limit set. The number of attractors is the number of acceptable labelings. The set of labels we use is $\{(\square, 0), (\square, 1), (\triangleright, 0), (\triangleright, 1)\}$. “ \triangleright ” intuitively means “oscillating” while “ \square ” means “fixed” (like the recorder symbols play/stop). The second component is called the “preferred” state of the cell.

Definition 20 (acceptable labeling). *A labeling is acceptable if and only if, for each cell v , if the cell has label*

1. (\square, α) then it has strictly more than $\deg(v)/2$ neighbors with label (\square, α) ;
2. (\triangleright, α) then
 - (a) if the father has a label of the form (\square, β) , then $\alpha = 0$ and v has one more son labeled $(\cdot, 1 - \beta)$ than sons labeled (\cdot, β) ;
 - (b) otherwise, i.e. the father has a label of the form (\triangleright, \cdot) , v has one more son labeled (\cdot, α) than sons labeled $(\cdot, 1 - \alpha)$.

Note that only vertices of even degree can have a label of the form (\triangleright, \cdot) .

The apparent asymmetry in case 2a (imposing $\alpha = 0$) is there only to avoid double counts in Theorem 22, one could as well have defined acceptable labelings with $\alpha = 1$.

Theorem 22 shows that a labeling corresponds to an attractor, and Theorem 23 details the meaning of a labeling, thus the structure of an attractor.

Definition 21. *The configuration $\text{snd}(L)$ corresponding to a labeling L is the projection of the second component: $L(v) = (\cdot, \alpha) \Rightarrow \text{snd}(L)(v) = \alpha$.*

Theorem 22. *Given a tree, there is a bijection between attractors and acceptable labelings.*

Theorem 23 (Structure of an attractor). *Let L be an acceptable labeling. Then for every configuration c reachable by a sequence of updates from $\text{snd}(L)$, for every cell v :*

1. If $L(v) = (\square, \alpha)$ then $c(v) = \alpha$ (this is why “ \square ” intuitively means “fixed”).

2. If $L(v) = (\triangleright, \alpha)$ (in this case $\deg(v)$ is even and thus v cannot be the root, which has degree 1) then
- (a) if the father has a label of the form (\square, \cdot) , v is in the state appearing in majority among its neighbors (and can be in any state in case of equality);
 - (b) otherwise, if v is in the state α , then its sons are also in their preferred states. Moreover, if v is not in the state α , then all its sons that are not in their preferred state are in the same state.

6 Hitting time for trees

6.1 Trees of maximum degree ≤ 3

On trees where the degrees are at most 3, the dynamics ends by fixing the states of the cells of degree 1 and 3 (see Remark 10) and it may remain some isolated particles which oscillate on disjoint paths.

Theorem 24. *Stochastic Minority on trees with degrees at most 3 hits the limit set in $O(N^4)$ steps on expectation. The attractors of a tree \mathbb{T} are in bijection with the matchings of the reduced tree \mathbb{T}' where each path of \mathbb{T} has been replaced by an edge.*

6.2 Trees of maximum degree ≥ 4

In this part, we introduce biased trees (Definition 27 and Figure 5) such that the dynamics \hat{c} converges in exponential time on this topology (Theorem 34). Vertices of biased trees have a degree at most 4. In fact, biased trees simulate biased random walks (Definition 25) which converge in exponential time. Biased trees are created from small trees called widgets (Definition 26 and Figure 4) arranged on a line. Except from the ends, this line of widgets is made of gates. According to the configuration, these gates are either locked, unlocked or stable (Definition 28). On a correct configuration (Definition 29), the line of gates is split into two regions: all gates on the left side are stable and all gates on the right side are unstable (locked or unlocked). In a correct configuration, three different events may be triggered with the same probability $1/N$ (Fact 30 and Corollary 33):

- the rightmost stable gate becomes an unlocked gate.
- the leftmost unstable gate becomes stable if it is unlocked.
- the leftmost unstable gate is switched from locked to unlocked or the contrary.

Thus stable gates tend to disappear. This dynamics will ultimately converge to the stable configuration \hat{c}_f (Definition 31). To reach this configuration all gates must be stable. Thus it takes exponential times to dynamics \hat{c} to converge on biased tree with an initial correct configuration.

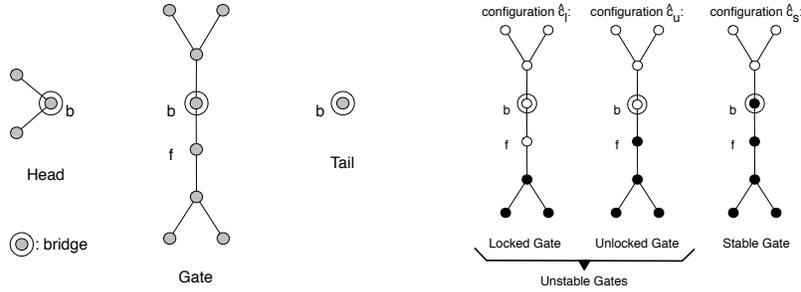
Definition 25 (Biased random walks). *Consider a sequence of random variables $(X_i)_{i \geq 0}$ defined on $\{0, \dots, n\}$ such that $X_0 = 0$ and for all $i \geq 0$:*

- $P(X_{i+1} = 1 \mid X_i = 0) = 1$ (reflecting barrier at 0).

- $P(X_{i+1} = n \mid X_i = n) = 1$ (absorbing barrier at n).
- If $0 < x < n$, $P(X_{i+1} = x - 1 \mid X_i = x) + P(X_{i+1} = x + 1 \mid X_i = x) = 1$ and $0 < a < P(X_{i+1} = x + 1 \mid X_i = x) < b < 1/2$ for some $a, b \in \mathbb{R}_+$.

This sequence is called a biased random walk. Let $T = \min\{i \mid X_i = n\}$, then $\theta(b) \leq \mathbb{E}[T] \leq \theta(a)$ where $\theta(p) = \frac{2p(1-p)}{(1-2p)^2} \left(\left(\frac{1-p}{p} \right)^n - 1 \right) - \frac{n}{1-2p}$.

Definition 26 (Widgets). A Widget W is a tree $\mathbb{T} = (\mathbb{V}, \mathbb{E}, b)$ where $b \in \mathbb{V}$ is called the bridge. We consider the three widgets described in figure 4.a: head, gate and tail and the three configurations $\hat{c}_l, \hat{c}_u, \hat{c}_s$ for gates.



4.a – The 3 widgets used in the construction of a biased tree. Gray denotes the fact that cell state is not represented

4.b – The three configurations \hat{c}_l, \hat{c}_u and \hat{c}_s .

Fig. 4. Widgets used in the construction of a biased tree.

Definition 27 (Biased trees). Let $(W_i)_{0 \leq i \leq n+1}$ be a finite sequence of widgets where $W_i = (\mathbb{V}_i, \mathbb{E}_i, b_i)$. From this sequence, we define the tree $\mathbb{T} = (\mathbb{V}, \mathbb{E})$ where $\mathbb{V} = \cup_{i=0}^{n+1} \mathbb{V}_i$ and $\mathbb{E} = (\cup_{i=0}^{n+1} \mathbb{E}_i) \cup (\cup_{i=0}^n b_i b_{i+1})$. Abusively we also denote by $(W_i)_{0 \leq i \leq n+1}$ the tree generated by this sequence. A biased tree of size n is a finite sequence of widgets $(W_i)_{0 \leq i \leq n+1}$ where W_0 is a head, for $1 \leq i \leq n$, W_i is a gate and W_{n+1} is a tail.

Definition 28 (Stable and unstable gates). Consider a biased tree $(W_i)_{0 \leq i \leq n+1}$ and a configuration \hat{c} . We denote by \hat{c}_{W_i} , the restriction of \hat{c} to widget W_i . We say that gate i is locked if $\hat{c}_{W_i} = \hat{c}_l$, unlocked if $\hat{c}_{W_i} = \hat{c}_u$ and stable if $\hat{c}_{W_i} = \hat{c}_s$. An unstable gate is a gate which is locked or unlocked.

Definition 29 (Correct configurations). Configuration \hat{c} is correct if cells of the head are black, the tail is white, and there exist j such that for all $1 \leq i \leq j$ gate i is stable and for all $j < k \leq n$ gate k is unstable. We say that configuration \hat{c} is on position j . We denote by $Pos(\hat{c})$ the position of configuration \hat{c} . The position is unlocked if $j = n$ or gate $j + 1$ is unlocked, the position is locked otherwise.

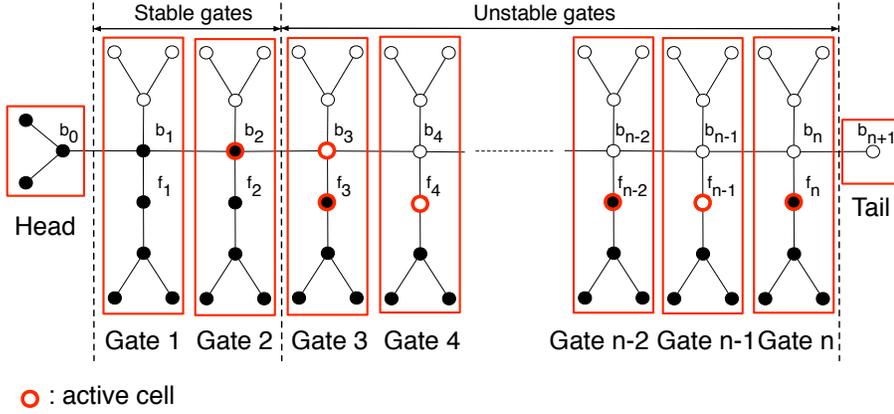


Fig. 5. A biased tree and a correct configuration on unlocked position 2.

Fact 30 (Active cells) Consider a correct configuration \hat{c} on position j . The active cells of \hat{c} are:

- Cell b_j if $j \neq 0$.
- Cell b_{j+1} if $j \neq n$ and gate W_{j+1} is unlocked.
- Cell f_i if $j < i \leq n$.
- Cell b_{n+1} if $j = n$.

Proof. Consider a correct configuration \hat{c} on position j . The only cells which may be active in \hat{c} are cells b_i and f_i for $1 \leq i \leq n$ and cell b_{n+1} . Cell b_{n+1} is active if and only if $\hat{c}(b_n) = 1$ that is to say gate W_n is stable. For all $1 \leq i \leq n$, cell f_i is active if and only if the gate W_i is unstable that is to say $j < i \leq n$. For all $1 \leq i \leq n$, cell b_i is inactive if $\hat{c}(b_{i-1}) = \hat{c}(b_i) = \hat{c}(b_{i+1})$. Thus among cells $(b_i)_{1 \leq i \leq n}$, only cells b_j and b_{j+1} may be active: cell b_j is active and cell b_{j+1} is active if gate W_{j+1} is unlocked. \square

Definition 31 (Final configuration). The final configuration \hat{c}_f is the configurations where cells of the head are black, tail is black and every gate is stable. We say that \hat{c}_f is on position $n + 1$, $Pos(\hat{c}_f) = n + 1$.

Lemma 32. Configuration \hat{c}_f is stable.

Proof. Consider the correct configuration \hat{c} on position n . According to fact 30, only cells b_n and b_{n+1} are active. If cell b_{n+1} fires, these two cells become inactive and no other cell becomes active. Firing cell b_{n+1} leads to configuration \hat{c}_f . Thus \hat{c}_f is stable. \square

Corollary 33. Consider a correct configuration \hat{c} , then configuration $\hat{c}' = \hat{\delta}(\hat{c})$ is either correct or \hat{c}_f . Moreover $|Pos(\hat{c}') - Pos(\hat{c})| \leq 1$.

Proof. Consider a correct configuration \hat{c} on position j and the configuration $\hat{c}' = \hat{\delta}(\hat{c})$. If an inactive cell fires than $\hat{c}' = \hat{c}$. Now consider that an active cell fires (see fact 30):

- if $j \neq 0$ and cell b_j fires: then gate W_j becomes unlocked and \hat{c}' is a correct configuration on unlocked position $j - 1$.
- if $j \neq n$ and cell b_{j+1} fires: then gate W_{j+1} becomes stable and \hat{c}' is a correct configuration on position $j + 1$.
- if cell f_i fires with $j < i \leq n$: then gate W_i becomes unlocked (resp. locked) in \hat{c}' if it is locked (resp. unlocked) in \hat{c} . Configuration \hat{c}' stays correct and on position j .
- if $j = n$ and cell b_{n+1} fires: then $\hat{c}' = \hat{c}_f$. □

Theorem 34. *On biased trees of size n (i.e. $N = 8n + 4$ vertices), starting from an correct configuration, Stochastic Minority converges almost surely to c^f . Moreover the hitting time T of the limit set satisfies $\Theta(1.5^n) \leq \mathbb{E}[T] \leq \Theta(n4^n)$.*

References

1. Gács, P.: Reliable cellular automata with self-organization. *Journal of Statistical Physics* **103**(1/2) (2001) 45–267
2. Toom, A., Vasilyev, N., Stavskaya, O., Mityushin, L., Kurdyumov, G., Pirogov, S.: Discrete local Markov systems. In: *Stochastic Cellular Systems : Ergodicity, Memory, Morphogenesis*. Manchester University Press (1990) 1–182
3. Balister, P., Bollobás, B., Kozma, R.: Large deviations for mean fields models of probabilistic cellular automata. *Random Structures & Algorithms* **29** (2006) 399–415
4. Bersini, H., Detours, V.: Asynchrony induces stability in cellular automata based models. In: *Proceedings of Artificial Life IV*, Cambridge, MIT Press (1994) 382–387
5. Buvel, R., Ingerson, T.: Structure in asynchronous cellular automata. *Physica D* **1** (1984) 59–68
6. Kanada, Y.: Asynchronous 1d cellular automata and the effects of fluctuation and randomness. In: *Proceedings of the Fourth Conference on Artificial Life (A-Life IV)*, MIT Press (1994)
7. Schönfisch, B., de Roos, A.: Synchronous and asynchronous updating in cellular automata. *BioSystems* **51** (1999) 123–143
8. Fatès, N., Morvan, M., Schabanel, N., Thierry, E.: Fully asynchronous behavior of double-quiescent elementary cellular automata. *Theoretical Computer Science* **362** (2006) 1–16
9. Fatès, N., Regnault, D., Schabanel, N., Thierry, E.: Asynchronous behaviour of double-quiescent elementary cellular automata. In: *Proceedings of LATIN'2006*. Volume 3887 of LNCS., Springer (2006)
10. Fukás, H.: Probabilistic cellular automata with conserved quantities. *Nonlinearity* **17**(1) (2004) 159–173
11. Fatès, N., Morvan, M.: Perturbing the topology of the game of life increases its robustness to asynchrony. In: *LNCS Proc. of 6th Int. Conf. on Cellular Automata for Research and Industry (ACRI 2004)*. Volume 3305. (Oct. 2004) 111–120
12. Demongeot, J., Aracena, J., Thuderoz, F., Baum, T.P., Cohen, O.: Genetic regulation networks : circuits, regulons and attractors. *C.R. Biologies* **326** (2003) 171–188
13. Regnault, D., Schabanel, N., Thierry, E.: Progresses in the analysis of stochastic 2d cellular automata : A study of asynchronous 2d minority. In: *Proceedings of MFCS'2007*. Volume LNCS 4708., Springer (2007) 320–332
14. Regnault, D., Schabanel, N., Thierry, E.: On the analysis of "simple" 2d stochastic cellular automata. In: *Proceedings of LATA'2008*. (2008) To appear.
15. Istrail, S.: Statistical mechanics, three-dimensionality and np-completeness : I. universality of intractability for the partition function of the ising model across non-planar lattices. In: *Proceedings of STOC'00*, ACM Press (2000) 87–96

16. McCoy, B., Wu, T.T.: The Two-Dimensional Ising Model. Harvard University Press (1974)
17. Rojas, R.: Neural Networks : A Systematic Introduction. Springer (1996) Chap. 13 - The Hopfield Model.
18. Goles, E., Martinez, S.: Neural and automata networks, dynamical behavior and applications. Volume 58 of Maths and Applications. Kluwer Academic Publishers (1990)
19. Grigoriev, S., Priezzhev, V.: Random walk of annihilating particles on the ring. Theoretical and Mathematical Physics **146**(3) (2006) 411–420
20. Grimmett, G., Stirzaker, D.: Probability and Random Process. 3rd edn. Oxford University Press (2001)
21. Lindvall, T.: Lectures on the coupling method. Wiley (1992)
22. Levin, D.A., Peres, Y., Wilmer, E.L.: Markov Chains and Mixing Times. <http://www.oberlin.edu/markov/> (2006)

A Omitted proofs

Proof (proof of Theorem 12). The best way to analyze cycles is to study the movement of particles. Firing a cell incident to a particle either attract the particle on the next edge if it is free, or annihilate both particles on each side of the cell. Consequently, the dynamics comes to the analysis of random walks of annihilating particles on a discrete ring. If the number N of cells is even (resp. odd), any configuration has necessarily an even (resp. odd) number of particles. Due to the rules of particle movements, from any initial configuration, one can always destroy pairs of particles until it remains only one particle (if N is odd) or no particle (if N is even). Thus attractors are included in this set of configurations. As a matter of fact, all those configurations belongs to an attractor since it is not possible to decrease the energy from them. The exact structure of the attractors is clear if N is even (there are only two configurations of energy 0) and when N is odd, one can easily check that configurations with only one particle are organized as a cycle in the transition graph (defined by possible transitions over all the configurations).

To bound the expected hitting time of the limit set, associate with each configuration c^t a weight X_t which is the maximum distance between two consecutive particles (for a fixed orientation of the cycle, say for instance “clockwise”) if there is at least two particles, or N if there is only one particle, or $N + 1$ if there is no particle. For all t , $X_t \in \{1, \dots, N + 1\}$ and c_t belongs to the limit set if and only if $X_t = N$ or $X_t = N + 1$. Let $\Delta X_{t+1} = X_{t+1} - X_t$. One can check that $\mathbb{E}(\Delta X_{t+1} | c_t = c) \geq 0$ for any configuration c . Moreover $\mathbb{E}((\Delta X_{t+1})^2 | c_t = c) \geq 3/N$ for any configuration c . Using a classical Stopping Time theorem from martingales [20], one gets the bound $O(N^3)$ on $\mathbb{E}[T]$ where $T = \min\{t \geq 0 \mid X_t \in \{N, N + 1\}\}$, and thus on the expected hitting time of the limit set.

A careful look at the proof shows that this bound on the hitting time of the limit set also applies to the graphs which are paths. \square

Proof (proof of Theorem 13). Let n_b be the number of black cells of a configuration. Since the neighborhood of a cell is \mathbb{V} , the potential of a black cell is n_b . The potential of a white cell is $N - n_b$. If $n_b > \frac{N+1}{2}$ (resp. $n_b < \frac{N-1}{2}$) then firing a black (resp. white) cell decreases the energy and the configuration is not in \mathcal{A} . If a configuration is in \mathcal{A} then $\frac{N-1}{2} \leq n_b \leq \frac{N+1}{2}$, consider such a configuration:

- If N is even then all cells have potential $\frac{N}{2}$ and these $C_N^{\frac{N}{2}}$ configurations are stable configurations.
- If N is odd we call C_b (resp. C_w) the set of configurations where $n_b = \frac{N+1}{2}$ (resp. $n_b = \frac{N-1}{2}$). White (resp. black) cells of a configuration in C_b (resp. C_w) are inactive and black (resp. white) cells are active, firing one of them leads to a configuration of C_w (resp. C_b). Thus from any configuration of $C_w \cup C_b$, there is no sequence of updates that causes a drop of energy and $C_w \cup C_b = \mathcal{A}$. Now, we prove that \mathcal{A} is made of only one attractor. Consider two configurations c and c' , we define the distance from c to c' as $d(c, c') = |\{i | c_i \neq c'_i\}|$. Consider two configurations c and c' such that $c, c' \in C_b$ and $c \neq c'$. Since c and c' have the same number of black cells then there exist i and j such that $c_i = 1, c'_i = 0, c_j = 0$ and $c'_j = 1$. Firing cells c_i and afterward c_j in c leads to a configuration c'' and $d(c'', c') = d(c, c') - 2$. Thus from any configuration c of $C_b \cup C_w$ there is a sequence of updates which leads to any configuration c' of $C_b \cup C_w$. The set $C_b \cup C_w$ is an attractor of size $2C_N^{\frac{N-1}{2}}$.

Now consider a configuration where $n_b > \frac{N+1}{2}$. As long as the configuration does not become a configuration of \mathcal{A} , white cells are inactive. When $\lceil n_b - \frac{N+1}{2} \rceil$ black cells have fired the configuration is in \mathcal{A} . At each time step there is a probability $\frac{n_b}{N}$ to fire a black cell. This kind of dynamics is known as coupon collector (see [20]) and $T = O(N \log N)$. \square

Proof (Proof of Proposition 18).

We first prove that the cells in state 0 at the end of phase 1 of the algorithm cannot switch to state 1, whatever the sequence of updates, i.e. there are in case 1 of Fact 17. Indeed, assume instead that there exists a cell i and sequence of configurations c_1, c_2, \dots, c_k such that

- c_1 is the configuration at the end of phase 1;
- each configuration is the result of firing one cell in the previous configuration;
- $c_1(i) = 0$ and $c_k(i) = 1$.

Let c_ℓ be the configuration just before the first update of this sequence that fires an active cell j in state 0 (there exists one since at least i will be fired in this sequence). But j must have been already active in c_1 , since it had at least as many neighbors in state 1 as in c_ℓ (this is a monotonicity argument). The algorithm thus would not have exited the first “while” loop, which is a contradiction.

Same holds for the cells in state 0 at the end of phase 3. Also, by symmetry, the cells in state 1 at the end of phase 2 are in case 2 of Fact 17.

The last thing to prove is that the remaining cells are in case 3. Indeed, those cells

- were white at the end of phase 2, which means they can be made white by a sequence of updates starting from the configuration c_1 (the configuration at the end of phase 1). By monotonicity, they can be made white by a sequence of updates starting from any configuration in the attractor.
- were black at the end of phase 3, which means by a symmetric argument that they can be made black. \square

Proof (Proof of theorem 22). Let $\text{attr}(c)$ be the attractor in which the configuration c lies. The function attr is naturally defined only on the limit set.

We define a function f that maps an attractor A to an acceptable labeling L . We then conclude the proof by showing that for every attractor A , $\text{attr} \circ \text{snd}(f(A)) = A$ and for every acceptable labeling L , $f(\text{attr} \circ \text{snd}(L)) = L$.

To define f , let A be an attractor, we construct $L = f(A)$. For all vertices that have the same state α in any configuration of A , define $L(v) := (\square, \alpha)$. All leaves are now labeled. We define the labeling of the remaining vertices inductively (in a bottom-up order). Each remaining node has an even degree and an odd number of sons. Considering a cell v having all its sons labeled:

- if the father is already labeled (thus having a label of the form (\square, \cdot)), define $L(v) := (\triangleright, 0)$;
- otherwise, the father is not already labeled, let α be the majority state among the preferred states of the sons and define $L(v) := (\triangleright, \alpha)$. Note that v can change its state in A , so $\deg(v)$ is even, so v has an odd number of sons and the majority state is well defined.

Let us show that the labeling we have just defined is acceptable. Let $\text{paint}(L, \alpha)$ be the following configuration (intuitively “set to state α as many cells as possible”):

$$\text{paint}(L, \alpha)(v) := \begin{cases} \beta & \text{if } L(v) = (\square, \beta) \\ \alpha & \text{if } L(v) = (\triangleright, \cdot) \end{cases}$$

Section 5.1 has shown that, given the way we construct L , this configuration is in A .

1. Consider a cell v labeled $(\square, 0)$ and the configuration $\text{paint}(L, 1)$. All the neighbors of v not labeled $(\square, 0)$ are in the state 1. So, there are necessarily more than $\deg(v)/2$ neighbors labeled $(\square, 0)$ (and thus in the state 0), otherwise, updating v would make it change its state, contradicting the definition of $L(v) = (\square, \cdot)$. By symmetry, point 1 of the definition of acceptable labelings is fulfilled.
2. Now consider a cell labeled (\triangleright, α) .

(a) If the father has a label of the form (\square, \cdot) then $L(v)$ has been defined as $(\triangleright, 0)$. So, point 2a is fulfilled.

(b) We consider the remaining nodes (those with a father labeled (\triangleright, \cdot)) in a bottom-up order and prove inductively that for each one, point 2b is fulfilled.

Assume for now that there is a configuration c of A in which all the nodes of the subtree of v are in their preferred state, except v which is in state $1 - \alpha$. Then if v had more than one more son labeled (\cdot, α) , updating v would make the energy decrease, contradicting the fact that c is in the limit set. So, point 2b is fulfilled.

Such a configuration c can be constructed inductively: start from the configuration $\text{paint}(L, 1 - \alpha)$, consider each node in the subtree of v in a bottom-up order, update it if it is not in its preferred state. Updating a node makes it go to its preferred state, because it has an acceptable label.

Let us now show that for any acceptable labeling L , $\text{snd}(L)$ is in the limit set. As is the last paragraph, start from the configuration $\text{paint}(L, 0)$, consider each node in a bottom-up order, update it if it is not in its preferred state. Updating a node makes it go to its preferred state, because it has an acceptable label. Then consider each node in a top-down order and update it if it is not in the state 1. Again, thanks to the definition of an acceptable labeling, updating a node makes it change its state and be in the state 1.

Symmetrically, there is a sequence of updates leading from $\text{paint}(L, 1)$ back to $\text{paint}(L, 0)$. This is a cycle, so the energy cannot decrease.

Moreover, there is no sequence of updates leading to change the state of a cell labeled (\square, α) : the first change of state of such a cell v would contradict the fact that it has more than $\deg(v)/2$ neighbors with label (\square, α) and thus in the state α .

Thus, the energy cannot decrease and the whole cycle, including $\text{snd}(L)$, is in the limit set (see Proposition 9).

We have even shown that if we let A be $\text{attr} \circ \text{snd}(L)$, then

- all vertices labeled (\square, α) are in the state α in every configuration in A
- there are two configurations c, c' in A such that, for every cell v labeled (\triangleright, \cdot) , $c(v) = 0$ and $c'(v) = 1$.

In the definition of f , the value of α for vertices labeled (\triangleright, α) is entirely determined by the labeling of vertices labeled (\square, \cdot) . Since L and $f(\text{attr} \circ \text{snd}(L)) = L$ have the same set of vertices labeled (\square, α) for $\alpha \in \{0, 1\}$, they are equal.

Finally, Section 5.1 has shown that for an attractor A , the configuration c in which all oscillating vertices are in the state 1 belongs to A . We have also shown that $c' := \text{paint}(f(A), 1)$ is in the limit set. But c' assigns state 1 to vertices labeled (\triangleright, \cdot) , i.e. oscillating vertices, and assigns the state they have in all configuration of A to the remaining vertices. Thus, $c = c'$, which means that the attractors $\text{attr} \circ \text{snd}(f(A))$ and A have a common element and are thus equal.

This concludes the proof by implying that f is a bijection. □

Proof (Proof of Theorem 23). This is clearly true for $\text{snd}(L)$. So it is sufficient to show recursively that if this is true for a configuration c , then it is true for the resulting configuration c' after firing one cell.

Any cell v labeled (\square, α) has strictly more than $\deg(v)/2$ in state α , so v is not active in c . This proves point 1.

Proof of Theorem 22 has shown that $\text{snd}(L)$ is in the limit set. Thus, any cell v has always at least $\deg(v)/2$ neighbors in the same state as v (otherwise, firing it would decrease the energy). This proves point 2a.

Now, consider a cell v in case 2b. There are only two cases to consider:

- v is in its preferred state and a son v' of v is fired. But then v' must also have been in its preferred state α , and same for all its sons. Which means v' has one more son labeled (\cdot, α) than sons labeled $(\cdot, 1 - \alpha)$, and is thus inactive.
- v not in its preferred state and is fired. Then v has as many neighbors in each state (because we are in the limit set), all its sons that are not in their preferred state are in

the same state (recursion hypothesis), and v has one more son labeled (\cdot, α) than sons labeled $(\cdot, 1 - \alpha)$ (definition of an acceptable labeling). Which implies that all the sons of v are actually in their preferred state. \square

Proof (proof of Theorem 24). We study here the movements of particles in the initial tree \mathbb{T} . One can divide \mathbb{T} into its induced subgraphs which are paths, denoted $\mathbb{P}_1, \dots, \mathbb{P}_\ell$. Those paths link the cells of odd degree (1 or 3). The reduced tree \mathbb{T}' is obtained by replacing each path \mathbb{P}_j by an edge. Consider a configuration c on \mathbb{T} which belongs to an attractor. There can not be two particles on the same path \mathbb{P}_j , otherwise a sequence of updates could lead to the collision of these two particles and thus to an energy decrease. In the same way, there can not be two particles on two paths \mathbb{P}_j and \mathbb{P}_k which share a common extremity. This extremity would necessarily be a cell of degree 3, then a sequence of updates could position the two particles on the edges incident to this cell. Firing this cell at that time would decrease the number of particles by at least 1 and lead to an energy decrease. Consequently, in configurations of the limit set, there is at most one particle by path \mathbb{P}_j and there can not be two particles on two incident paths. Reciprocally, one can easily check that such configurations always belong to the limit set. Moreover in such configurations, once a particle is stuck on a path \mathbb{P}_j , it can not leave it any longer. One can deduce from all those remarks that the attractors are in bijection with the matchings of \mathbb{T}' (the matching indicates where the isolated cells are located in \mathbb{T}).

To prove the bound on the expected hitting time of the limit set, we find a bound on the time until at least one particle disappears. Consider a configuration where there exist two particles on a same path \mathbb{P}_j of length n . One can suppose that these two particles follow a random walk on this path with reflecting barriers at each extremity, unless they collide with another particle (leading to the loss of two particles in the tree) or unless one of the two particles leaves the path (leading to the loss of at least one particle in the tree since leaving necessarily involves a cell of degree 3 fired with two incident particles). Thus under the condition that they have not disappeared before, a bound on the expected time elapsed until they collide can be derived from classical studies of random walks with reflecting barriers [20]: this expected time is bounded by $O(n^3)$, and consequently it is also a bound on the time until at least one particle disappear.

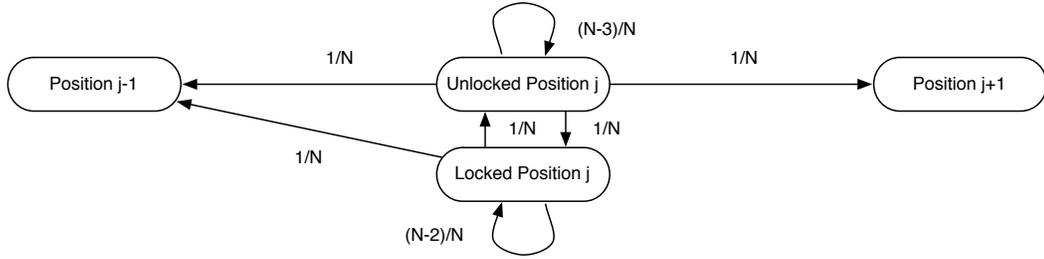
Consider a configuration where there exist two particles on two paths \mathbb{P}_j and \mathbb{P}_k (of respective length n and m) sharing a common extremity: the cell i . With the same reasoning, assuming that the two particles have not led to the removal of another particle means that they follow random walks on their respective paths with reflecting barriers at the extremities. Then they can only disappear by being both incident to cell i when cell i is fired. By analyzing the two-dimensional random walk corresponding to the evolution of the respective distances to cell i , this event occurs after at most $O(\max(n, m)^3)$ steps on expectation, as proved in Section C using standard tools for multi-dimensional random walks.

Finally, from any configuration which does not belong to the limit set, at least one particle disappear within $O(N^3)$ steps on expectation. Since the number of particles in any initial configuration on a tree is bounded by N , the expected time to hit the limit set is bounded by $O(N^4)$.

Proof (proof of Theorem 34). Consider a biased tree of size n , an initial correct configuration \hat{c}^0 on position 0 and the sequence $(\hat{c}^t)_{t \geq 0}$. Dynamics $\hat{\delta}$ converges almost surely from initial

configuration \hat{c}^0 and $c^T = c_f$. We define the sequence of random variable $(t_i)_{i \geq 0}$ as $t_0 = 0$ and $t_{i+1} = \min\{t > t_i | \text{Pos}(\hat{c}^{t_{i+1}}) \neq \text{Pos}(\hat{c}^{t_i}) \text{ or } \text{Pos}(\hat{c}^{t_{i+1}}) = n+1\}$. Consider the sequence of random variable $(X_i)_{i \geq 0}$ such that $X_i = \text{Pos}(\hat{c}^{t_i})$. According to corollary 33, $|X_{i-1} - X_i| = 1$.

Consider a configuration \hat{c}^t on locked position $n > j > 0$ then firing cell b_j leads to a configuration on position $j - 1$ and firing cell f_j leads to a configuration on unlocked position j . Firing other cells does not affect the position of the configuration. Consider a configuration \hat{c}^t on unlocked position $n > j > 0$ then firing cell b_j leads to a configuration on position $j - 1$, firing cell f_j leads to a configuration on locked position j and firing cell b_{j+1} leads to a configuration on position $j + 1$. Firing other cells does not affect the position of the configuration. A cell has a probability $1/N$ to fire where $N = 4 + 8n$. Thus, the evolution of a configuration on position $0 < j < n$ can be summarized as:



A basic analysis yields that:

- if $1 \leq x \leq n$ then $P(X_{i+1} = x + 1 | X_i = x) = 1 - P(X_{i+1} = x - 1 | X_i = x)$ and $1/5 \leq P(x_{i+1} = x + 1 | X_i = x) \leq 2/5$.
- $P(X_{i+1} = 1 | X_i = 0) = 1$.
- $P(X_{i+1} = n + 1 | X_i = n + 1) = 1$.

Thus the behavior of $(X_i)_{i \geq 0}$ is as described in definition 25. We define the random variable $T' = \min\{i | X_i = n\}$ which corresponds to the first time when all gates are stable, then $\Theta(\frac{3}{2}^n) \leq \mathbb{E}[T'] \leq \Theta(4^n)$ on expectation (see Def. 25). We call c^{f-1} the correct configuration on position n (*i.e.* all gates are stable). Then $c^T = c^f$, $c^{T-1} = c^{f-1}$ and $P(c^{t+1} = c^f | c^t = c^{f-1}) = 1/2$. Thus, $\mathbb{E}[T] = \Theta(\mathbb{E}[t_{T'}])$. By definition, $t_{T'} = \sum_{i=1}^{T'} (t_i - t_{i-1}) = \sum_{i=1}^{\infty} [(t_i - t_{i-1}) 1_{t_i < T'}]$. Since there are at most 2 cells which may modify the position of a correct configuration, we have $1 \leq \mathbb{E}[t_{i+1} - t_{i-1}] \leq \Theta(n)$. Thus $\sum_{i=1}^{\infty} (1_{t_i < T'}) \leq \mathbb{E}[t_{T'}] \leq \Theta(\sum_{i=1}^{\infty} (n 1_{t_i < T'}))$. We conclude that $\Theta((\frac{3}{2})^n) \leq \mathbb{E}[T] \leq \Theta(n 4^n)$. \square

B Biased random walk

Consider a sequence of random variable $(X_i)_{i \geq 0}$ defined on $\{0, \dots, n\}$ such that $X_0 = 0$ and for all $i \geq 0$:

- $P(X_{i+1} = 1 | X_i = 0) = 1$ (reflecting barrier at 0).
- $P(X_{i+1} = n | X_i = n) = 1$ (absorbing barrier at n).
- If $0 < x < n$, $P(X_{i+1} = x - 1 | X_i = x) + P(X_{i+1} = x + 1 | X_i = x) = 1$ and $0 < a < P(X_{i+1} = x + 1 | X_i = x) < b < 1/2$ for some $a, b \in \mathbb{R}_+$.

Theorem 35. For all $0 \leq k \leq n$, let $T_k := \min\{i | X_i = k\}$ be the hitting time of position k . Then

$$\theta_k(b) \leq \mathbb{E}[T_k] \leq \theta_k(a)$$

$$\text{where } \theta_k(p) = \frac{2p(1-p)}{(1-2p)^2} \left(\left(\frac{1-p}{p} \right)^n - \left(\frac{1-p}{p} \right)^k \right) - \frac{n-k}{1-2p}.$$

This theorem is a direct consequence of classical analyzes of simple random walks on $\{0, \dots, n\}$ where for all $0 < x < n$, $P(X_{i+1} = x+1 | X_i = x) = p$ and $P(X_{i+1} = x-1 | X_i = x) = q$, with $p+q=1$. If $P(X_{i+1} = 1 | X_i = 0) = 1$ and $P(X_{i+1} = n | X_i = n) = 1$, then for all $0 \leq k \leq n$, the expectation of T_k is exactly $\mathbb{E}[T_k] = \frac{2p(1-p)}{(1-2p)^2} \left(\left(\frac{1-p}{p} \right)^n - \left(\frac{1-p}{p} \right)^k \right) - \frac{n-k}{1-2p}$. This formula is obtained by solving the following system of equations:

$$\begin{cases} \mathbb{E}[T_k] = p(1 + \mathbb{E}[T_{k+1}]) + (1-p)(1 + \mathbb{E}[T_{k-1}]) & \text{for all } 0 < k < n \\ \mathbb{E}[T_n] = 0 \\ \mathbb{E}[T_0] = 1 + \mathbb{E}[T_1] \end{cases}$$

C Bound on the hitting time of a 2D finite Markov Chain

C.1 Background on Markov chain theory

We recall only the necessary background on Markov chains to get a bound on the hitting time of a 2D finite Markov chain. For a gentle introduction and proofs, we refer for instance to Chapter 7 of [22].

Let $(X_t)_{t \in \mathbb{N}}$ be a Markov chain. We note τ_b the **hitting time** of b , i.e. the first time the Markov chain is in state b :

$$\tau_b := \min\{t \mid X_t = b\}$$

If (X, P) is a Markov chain reversible with respect to the probability π , the **conductance** of an unoriented edge (x, y) is

$$c(x, y) := \pi(x)P(x, y) = \pi(y)P(y, x)$$

Let a and b be two distinguished vertices, representing source and sink. We use the following potential, or **voltage**, of a vertex:

$$V(x) := \mathbb{P}(\tau_a < \tau_b \mid X_0 = x)$$

Clearly $V(a) = 1$ and $V(b) = 0$. Now, define the **current flow** on oriented edges as

$$I(x, y) := c(x, y)(V(x) - V(y)) \quad \text{and} \quad \|I\| := \sum_x I(a, x)$$

The **effective resistance** between a and b is

$$\mathcal{R}(a, b) := \frac{V(a) - V(b)}{\|I\|}$$

Theorem 36 (Commute time identity).

$$\mathbb{E}(\tau_b | X_0 = a) + \mathbb{E}(\tau_a | X_0 = b) = c\mathcal{R}(a, b)$$

In our case, $c = 1$, so $\mathcal{R}(a, b)$ is an upper bound for the average hitting time $\mathbb{E}(\tau_b | X_0 = a)$. Here is how one can bound $\mathcal{R}(a, b)$. A **flow** from a to b is a function on oriented edges which is antisymmetric: $\theta(x, y) = -\theta(y, x)$ and which obeys Kirchhoff's node law:

$$\forall v \notin \{a, b\} \quad \sum_x \theta(x, v) = 0$$

This is just the requirement “flow in equals flow out”. The **strength** of a flow is

$$\|\theta\| := \sum_x \theta(a, x)$$

Theorem 37 (Thomson's Principle). *For any finite connected graph,*

$$\mathcal{R}(a, b) = \inf \{ \mathcal{E}(\theta) \mid \theta \text{ a unit flow from } a \text{ to } b \}$$

where $\mathcal{E}(\theta) := \sum_{x,y} \frac{(\theta(x, y))^2}{c(x, y)}$.

C.2 Application to our case

If a neighbor does not exist (because the node is on the border), the edge points to the node itself. Each edge has the same probability $\frac{1}{4}$.

An invariant probability is the uniform probability $\pi : x \mapsto \frac{1}{nm}$. Our Markov chain is invariant with respect to this probability. So we can use the definitions of section C.1:

$$\forall x, y \quad c(x, y) = \frac{1}{4nm}$$

Thanks to Thomson's principle, it is sufficient to construct a flow from a to b to get an upper bound on $\mathcal{R}(a, b)$. If $a = (i, j)$ and $b = (n, m)$, we consider the trivial (and far from optimal) flow

$$\begin{cases} \theta((k, j), (k+1, j)) := 1 & \text{if } i \leq k < n \\ \theta((n, k), (n, k+1)) := 1 & \text{if } j \leq k < m \\ \theta(x, y) := 0 & \text{elsewhere.} \end{cases}$$

That is, a flow on a single path from a to b .

$\mathcal{E}(\theta) \leq (n+m)4nm$. We conclude with the commute time identity that the average hitting time is $O(n^3)$, assuming wlog that $n \geq m$.