



HAL
open science

Is there a trade-off between simplicity and robustness? Illustration on a lattice-gas model of swarming

Nazim Fatès, Vincent Chevrier, Olivier Bouré

► To cite this version:

Nazim Fatès, Vincent Chevrier, Olivier Bouré. Is there a trade-off between simplicity and robustness? Illustration on a lattice-gas model of swarming. Pierre-Yves Louis; Francesca R. Nardi. Probabilistic Cellular Automata, Springer, 2018, Emergence, Complexity and Computation, 978-3-319-65556-7. 10.1007/978-3-319-65558-1_16 . hal-01230145

HAL Id: hal-01230145

<https://inria.hal.science/hal-01230145v1>

Submitted on 29 Feb 2016

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.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

Is there a trade-off between simplicity and robustness? Illustration on a lattice-gas model of swarming

Nazim Fatès^{1,3} Vincent Chevrier^{2,3}
Olivier Bouré

(1) Inria Nancy – Grand Est (2) Université de Lorraine (3) Loria laboratory - UMR 7503

November 17, 2015

“Everything should be made as simple as possible, but not simpler”. The aim of this contribution is to examine a simple discrete model of collective motion in the light of this semi-apocryphal word of Einstein¹. We re-visit in detail a cellular automaton rule that was proposed some twenty years ago to explain the formation of swarms of particles: under some conditions, particles that are initially randomly scattered on a lattice may form coherent groups and move in a common direction, decided by a form of consensus.

Our study focuses on the model introduced by Deutsch [Deu96, BDG97]: in this particular kind of cellular automata, called lattice-gas cellular automata, the cells contain particles which can move along the four cardinal directions. The interactions between particles are local and stochastic: the particles tend to align with the neighbouring particles according to a single parameter, called the alignment sensitivity. Observations show the presence of a “minimal cohesion force”: the particles self-organise into patterns only when the alignment sensitivity is set higher than a threshold which depends on the number of particles in the grid. If this sensitivity is too low, no pattern emerges. When it becomes sufficiently high, it was observed that a spontaneous symmetry breaking occurs: the particles form a stripe pattern that moves diagonally on the lattice; the direction it takes results from a distributed consensus between particles.

The organisation in stripes is only one of the various patterns that can be observed [BFC13b]. Depending on the number of particles and on the value of the sensitivity, particles may also organise into clusters or, more surprisingly, they may “anti-align” and form what we call “checkerboard patterns”. The purpose of this book chapter is to explore the various behaviours of the model and see how they depend on the sensitivity and on the density. In a second time, we ask if this behaviour resist various modifications of the model or if its simple and discrete nature make it sensitive to such modifications.

¹See <http://quoteinvestigator.com/2011/05/13/einstein-simple/> for more details (consulted Jan. 2015)




Figure 1: Stipe formation: evolution of the system from a Bernoulli random initial configuration. The grid size is $L = 32$, the sensitivity is $\sigma = 2.0$ and the initial density is equal to 0.2. (See Sec. 2 for the formal definition of the system, the meaning of the colors and the parameters.)

1 Formal description of the system

1.1 First approach

Before we start describing formally the model, we recommend to the readers to observe visually the phenomenon of self-organisation as shown on Figure 1. We should take this figure as a metaphor: in the same way that, sitting on the sea shore, we can observe hundreds of birds group and form a flock, the particles that are randomly spread on the lattice progressively unite their movements and make a group that travels diagonally.

This simple experiment illustrates how randomness is a means to attain a consensus in a decentralised system. Indeed, the rule that governs the particles is stochastic, and despite this randomness, or more exactly *thanks to* the randomness, an agreement can be reached. To understand this phenomenon, many questions need to be raised. First, the role of the local rule has to be clarified. In particular, what level of “noise” must be introduced to allow the particles to form patterns? Beyond the local rule, how does the self-organisation phenomenon depend on the form of the grid? How does the system react to perturbations of its updating? etc.

In the lines that follow, we will first define formally our system and then examine with numerical simulations the behaviour of the model.




Figure 2: Color code associated to each local flux in $\{-1, 0, 1\} \times \{-1, 0, 1\}$ (color online).

1.2 The LGCA swarming model

The main elements of the model are the *cells*. The cells are arranged as a regular lattice, denoted by \mathcal{L} . Here, since our approach mainly relies on simulations, we will use a finite grid with periodic boundary conditions; we thus define the grid as $\mathcal{L} = (\mathbb{Z}/X.\mathbb{Z}) \times (\mathbb{Z}/Y.\mathbb{Z})$. For the sake of simplicity, when the grid has square dimensions we use $L = X = Y$. The generalisation to the infinite grid \mathbb{Z}^2 can be obtained without any modification.

Each cell of the lattice consists of four *channels*, denoted by $\mathbf{n}, \mathbf{e}, \mathbf{s}, \mathbf{w}$, and respectively associated to North, East, South and West directions. A channel can be *occupied* or *empty*, depending on whether it contains or not a *particle*. We thus represent the state of a cell as a quadruple $(q_{\mathbf{n}}, q_{\mathbf{e}}, q_{\mathbf{s}}, q_{\mathbf{w}}) \in Q = \{0, 1\}^4$, called a *tile*, where the 0 or 1 value respectively represents the absence or presence of a particle in a given channel.

The state of the system at a given time is called a *configuration* and we here choose to represent a configurations a *mapping* from \mathcal{L} to Q . The space of configurations is denoted by $Q^{\mathcal{L}}$.

The evolution of the system is described by a discrete dynamical system: $F : Q^{\mathcal{L}} \rightarrow Q^{\mathcal{L}}$. The *global function* F is decomposed into two steps: (a) the *interaction* step, where particles of a cell are “rearranged” in this cell, that is each particle is affected to a channel, and (b) the *propagation* step, where particles are moved from a cell to its neighbours. This movement is simple: each particle moves in the direction of its channel (there is thus no possibility of collisions).

Formally, we write that: $x^{t+1} = F(x^t) = F_{\text{prop}} \circ F_{\text{int}}(x^t)$. Let us now describe these two steps. We begin with the *propagation step*. Noting $y = F_{\text{prop}}(x)$ and: $\forall c \in \mathcal{L}, y(c) = (y_{\mathbf{n}}(c), y_{\mathbf{e}}(c), y_{\mathbf{s}}(c), y_{\mathbf{w}}(c))$, we simply have:

$$\begin{aligned} y_{\mathbf{n}}(c) &= x_{\mathbf{n}}(c - v_{\mathbf{n}}) & \text{with} & & v_{\mathbf{n}} &= (0, 1) \\ y_{\mathbf{e}}(c) &= x_{\mathbf{e}}(c - v_{\mathbf{e}}) & & & v_{\mathbf{e}} &= (1, 0) \\ y_{\mathbf{s}}(c) &= x_{\mathbf{s}}(c - v_{\mathbf{s}}) & & & v_{\mathbf{s}} &= (0, -1) \\ y_{\mathbf{w}}(c) &= x_{\mathbf{w}}(c - v_{\mathbf{w}}) & & & v_{\mathbf{w}} &= (-1, 0) \end{aligned}$$

where v_d represents the movement by one cell in direction d . In words, the North channel of a cell c *after* propagation is equal to the value of the North channel of the cell that is at the *South* of c *before* propagation (see Fig. 3).

We now define the interaction step. Let $q = (q_{\mathbf{n}}, q_{\mathbf{e}}, q_{\mathbf{s}}, q_{\mathbf{w}})$ be a *tile* (the state of a given cell). The interaction step consists of permuting the positions




Figure 3: Example of an evolution of five cells on one time step. Cells are represented with their four channels (triangles); each channel can be empty (white) or filled with a particle (red or blue, color online). The time step is decomposed into two steps: the interaction step (I) and the propagation step (P).

of the particles; this amounts to replacing this tile by a tile which contains the same number of particles. To define the choice of the new tile, let us denote by $N(q) = q_n + q_e + q_s + q_w$ the number of particles of a given tile q and by $\pi(q)$ the set of its “permuted tiles”: $\pi(q) = \{q' \in \{0, 1\}^4, N(q') = N(q)\}$. Now, since the interaction step is stochastic, we assign to each permuted tile a *weight* and select this tile with a probability that is proportional to this weight.

Let us now explain how we calculate the weights; we need to introduce intermediary definitions. For a configuration x and a given cell c , two quantities will be used:

- The *local flux* $\varphi : Q \rightarrow \{-1, 0, 1\}^2$ is the resulting “momentum” of the particles contained in this cell. Noting $x(c) = q = (q_n, q_e, q_s, q_w)$, we have:

$$\varphi(q) = q_n \cdot v_n + q_e \cdot v_e + q_s \cdot v_s + q_w \cdot v_w.$$

- The *director field* $D : Q^{\mathcal{L}} \times \mathcal{L} \rightarrow \mathbb{Z}^2$ is the sum of the local fluxes of the four immediate neighbours of c :

$$D(x, c) = \varphi(x(c + v_n)) + \varphi(x(c + v_e)) + \varphi(x(c + v_s)) + \varphi(x(c + v_w)).$$

For a configuration x , the probability to update a cell c with a tile $q = x(c)$ into a tile $q_I \in \pi(q)$ is then given by:

$$P(q \rightarrow q_I) = \frac{\exp[\sigma \cdot \langle D(x, c), q_I \rangle]}{Z(x, c)}$$

where:

- σ is a positive constant called the *alignment sensitivity*,
- $\langle \cdot, \cdot \rangle$ denotes the scalar product of two vectors of \mathbb{Z}^2 , and
- $Z(x, c) = \sum_{q \in \pi(x(c))} \exp[\sigma \cdot \langle D(x, c), \varphi(q) \rangle]$ is the renormalisation factor.

In words, given that σ is a positive factor, the tiles which are “in agreement” with the local field have the greatest weight while the tiles that are “in opposition” (negative scalar product) are assigned a little weight.

By noting $x_I = F_{\text{int}}(x)$ the result of the interaction step, our system can now be defined with the probability $P(x \rightarrow x_I)$ to update x into x_I :

$$P(x \rightarrow x_I) = \prod_{c \in \mathcal{L}} P(x(c) \rightarrow x_I(c)).$$

2 Elementary properties of the system

Before presenting numerical simulations of the model, let us examine two of its properties: the existence of a “hidden” invariant and its recurrence, that is, the fact that it almost surely comes back to its initial condition (possibly in a very long time).

Checkerboard invariant

By construction, the evolution of the system conserves the number of particles. But the LGCA construction also introduces another conserved quantity. Let us take a grid with even dimensions, and name “even” and “odd” particles, the particles which are respectively located on an even or odd cell, that is, a white or black cell of the corresponding checkerboard. Then, the total number of even and odd particles is exchanged at each time step. Formally, let (x^t) represent a given evolution of the system, let us denote by $\mathcal{L}_0 = \{c = (c_x, c_y) \in \mathcal{L}, c_x + c_y \in 2\mathbb{Z}\}$ and by $\mathcal{L}_1 = \mathcal{L} \setminus \mathcal{L}_0$ the set of even and odd cells, respectively. We also denote by: $N_0(t) = \sum_{c \in \mathcal{L}_0} N(x^t(c))$ and by $N_1(t) = \sum_{c \in \mathcal{L}_1} N(x^t(c))$ the number of even and odd particles, respectively. Remark that a cell that belongs to \mathcal{L}_0 has its four neighbours in \mathcal{L}_1 , and reciprocally. Also remark that the interaction step does not modify N_0 and N_1 . Then, by the composition of the interaction and propagation step: $(N_0, N_1)(t+1) = (N_1, N_0)(t)$.

The presence of this “checkerboard invariant” is often seen as a drawback of the model: indeed, it was not included in it deliberately. This is a well-known problem of LGCA on a square grid and this is why some authors prefer to use hexagonal grids, for which this invariant disappears. Note that using grids with odd dimensions is a dangerous expedient: indeed, even if the invariant does not exist at the global scale, it is still locally conserved on a short-time evolution. We refer to the work of Barberousse and Imbert for an in-depth discussion on how the existence of invariant quantities may affect the interpretation of the numerical simulations obtained on LGCA [BI13].

About patterns

In what follows, the term “pattern” will be used informally to designate a given subset of the configuration space $Q^{\mathcal{L}}$. These configurations have the same visual appearance and their evolution shows a form of visual stability during long




Figure 4: Example of a recurrent behaviour. The first time step (arrows 1 and 2) shows an arbitrary move of the particles. The three following time steps shows a sequence that reverses the system to the initial state.

periods (typically a few thousand steps). Note that the difficulty is that the patterns that we visually observe are not perfect; their formal definition would thus need to “tolerate” a given degree of noise...

Ideally, we would like to study the system by calculating the stationary distributions of the Markov chain which defines the dynamics of the system. Unfortunately, even though this Markov chain depends only on two parameters, coming up with a closed formula that would help us predict the properties of the system is out of reach. (This is already difficult for one-dimensional binary systems, see e.g. Ref. [MM14, Tag15].) Note that even if we use numerical analysis, the state space on which we need to compute the stationary distribution grows exponentially with the size of the grid and the number of particles.

To our knowledge, only few analytic results exist on this model. A mean-field approximation was used to predict the phase transition between non-ordered and ordered states (see Ref. [BDG97] and Ref. [DD05, Chap. 8]). However, this predicts the existence of a phase transition but does not reflect the various types of stable patterns observed in the system. Partial results have also been given by Marcovici and concern the study of invariant measures of an infinite-size system for some specific cases [Mar13, Chap. 2.6].

Recurrence of the Markov chain

The second important property to remark is the *recurrence* of the Markov chain that represents the evolution of the system. Various proofs of this property can be given ; we will here give an informal argument by showing that if a configuration y is reachable from x in one step, then x is reachable from y in three steps. As an example, consider the case illustrated in Figure 4. We now consider a particular cell c and a given particle, e.g. the which points to the

East on step (1) of the figure (in blue or dark grey). After the first time step, composed of an interaction (1') and a propagation (2), the particle p , is now in the cell at the North of c .

Let us now show that there is a sequence of updates that reverses this move.

- The interaction step makes p reverse its direction (2'); the propagation step (3) will thus make p come back to c , but with a direction not necessarily equal to the original direction.
- The interaction orients p in the inverse direction of its initial direction. In our example, it now points to West (3'), and it is propagated to the Western neighbour cell (4).
- In this cell, the direction of p is again inverted (4'). After propagation, p is again in c ; it points East, its initial direction (4'=1).

It can be observed that the same argument also applies to the other cell of Figure 4. Since there is no possibility of collision between particles in a lattice-gas cellular automaton, this argument can be generalised to any cell with any number of particles. Consequently, if a configuration y can be reached from x , then there is a non-zero probability that x can be reached from x , on other words, the system is recurrent.

This recurrence property has an important consequence: when the system “moves” from a disordered state to an ordered state (a pattern) that seems stable, there is always a non-zero probability to destroy the order and to return to the disordered initial condition. As a consequence, for *finite* grids, the patterns that we will observe are all subject to *metastability*: they may be stable for an average time that is very long, typically exponential in the grid size, and therefore unreachable by simulation, but they are always bound to be destroyed.

Patterns can thus be characterised by their *attractivity*, which quantifies the probability to reach a given pattern when starting from a random initial condition and their *stability*, which quantifies the average time that the system will remain on a given pattern before leaving it. In the following experiments, we will be mainly interested in attractivity: taking a random initial condition, we want to know which are the attractors that will “absorb” the greatest number of trajectories. However, we will also implicitly evaluate the stability of the patterns: we will require that their properties remain stable in large intervals of simulation time (typically a few thousand steps). It is a crucial problem to determine to which extent the stability of the patterns may be estimated by the means of mathematical analysis.

3 Exploration of the stable patterns

In the following experiments, we use Bernoulli random initial conditions, *i.e.*, each channel is set to be occupied independently of the other channels with a probability d , called the *initial density*, or simple the *density*, still this is a conserved quantity. Note that we do not consider densities higher than $1/2$ since




Figure 5: Stripe formation: Evolution of the mean alignment for the sampled trajectory observed in Fig. 1: $L = 32$, $d = 0.2$, $\sigma = 2.0$.

the local rule behaves symmetrically by exchanging particles and “holes”, that is, empty channels.

The experiments that we will present are in part a re-examination of the experiments presented in previous articles [BFC13a, BFC13b]. They are completed by new experiments which aim at showing the limits and difficulties in the study of this model.

3.1 Quantification of the order

We denote by $\gamma(x)$ the *mean alignment* of a configuration x : this parameter evaluates the tendency of a particle to be aligned with the particles of its neighbouring cells. It is defined with:

$$\gamma(x) = \frac{1}{4N(x)} \sum_{c \in \mathcal{L}} \langle \varphi(x, c), D(x, c) \rangle$$

Note that since each particle counts 4 times in the calculus of $D(x)$ (one per direction), we need to divide by 4 in order to assign to γ a value in $[-1, 1]$. A value close to 1 indicates that in general particles are aligned with their neighbours, a value close to 0 indicates an absence of alignment and a value close to -1, the fact that particles have a direction which is opposed to that of their neighbours. In an abuse of language we will call this phenomenon *anti-alignment*.




Figure 6: Stripe formation: distribution of the time to reach $\gamma = 0.8$ for $L = 32$ and $L = 64$.

3.2 Temporal evolution of the system

Figure 5 shows the evolution of the mean alignment parameter for the same trajectory as the one observed in Fig. 1. The evolution can be roughly divided into three phases: first, the formation of the stripe occurs in a relatively short time. In the second step, the stripe strengthens progressively and its shape gets more and more regular. After a thousand steps, the system has reached a metastable state and its characteristics do not seem to change.

3.3 Statistical distributions

Since we are working on a stochastic system, it can be asked whether this scenario can be repeated or if there is a large variability from one experiment to the other. Our informal observations of the time to form a stripe showed that the process may greatly vary from one run to the other. For example, the trajectory shown in Fig. 5 shows that the stripe appears in approximately 200 time steps, but for other simulations with $L = 32$, it was possible to observe the absence of a stripe during more than 1200 steps.

In order to assess the statistical variation of the time needed to form a stripe, we repeated the experiment, taking the same density but other initial conditions. For each trajectory, we measured the time needed to attain a state with a high level organisation. This level of organisation was arbitrarily decided to be quantified by the threshold of $\gamma = 0.8$. The value of this threshold corresponds to a point where the stripe has almost finished its formation and, in practice, never disappears. The order that is already present is in some sense “too high” to be destroyed by the random fluctuations of the system. However, remark

that since the system is recurrent, the stripe is bound to disappear.

Figure 6 shows the distribution of the time needed to reach the threshold $\gamma = 0.8$ for $L = 32$. The statistical sample consists of $N = 10\,000$ experiments obtained with $d = 0.2$. It can be asked whether the distribution of this time obeys a Poisson law ; we observed informally that the “tail” of the distribution is heavier than expected but larger statistical samples are needed to confirm the observations.

When we double the size of the lattice to $L = 64$, we notice a paradoxical effect : the average time to reach the threshold increases slightly but the “tail” of the distribution becomes *less* developed. One possible explanation of this paradox is that for small lattices there are locally stable patterns that form and destroy, whereas for large lattices, such patterns are less present. We refer to our previous work [BFC13b] for more details on this issue. It is an open problem to infer the behaviour of the system for lattice sizes that tend to infinity.

3.4 The checkerboard phase

To observe our second pattern, we keep the alignment sensitivity constant $\sigma = 0.2$, but we double the density taking $d = 0.4$.

Observing a trajectory. Figure 7 shows one evolution with these settings: contrarily to what was observed before, the particles now tend to *anti-align* with their neighbours! We observe that rapidly, homogeneous and stable zones appear; in these zones, the colours, which represent local fluxes, are alternating as in a checkerboard.

The stability of this pattern can be explained by looking at the effects of the interaction step and the propagation step. Indeed, in the interaction step, a cell tends to re-organise its particles in order to make them align with the neighbours. However, as the same movement happens in the neighbouring cells, all the particles reverse their direction. After the propagation step, the situation is unchanged: we are back to the initial condition.

The argument above shows that a checkerboard patterns are a kind of “equilibrium points” : if they are present, they will tend to stay, but we still do not know how the pattern forms. Empirically, we can only see that checkerboards merge when they touch each other. There exists a random movement of the frontiers which make the checkerboard grow or shrink. After a possibly very long time, they merge to form a single checkerboard that spans through the lattice².

Time evolution. The process of checkerboard formation can be (partially) quantified by monitoring the evolution of the alignment. Figure 8 shows how the alignment evolves for the same trajectory as observed on Fig. 7. We see that

²These dynamics is qualitatively similar to what is observed for the appearance of a checkerboard with the two-dimensional minority rule under a fully asynchronous dynamics [RST09] or with an Ising spin system with ferromagnetic alignment sensitivity.




Figure 7: Checkerboard formation: evolution of the system from a random initial configuration. The grid size is $L = 32$, the sensitivity is $\sigma = 2.0$ and the initial density is equal to $d = 0.4$.




Figure 8: Checkerboard formation: evolution of the mean alignment for the sampled trajectory observed in Fig. 7: $L = 32$, $d = 0.4$, $\sigma = 2.0$.




Figure 9: Checkerboard formation: distribution of the time to reach $\gamma = -0.5$ for $L = 32$ and $L = 64$.

it decreases until it reaches a metastable state characterised by $\gamma = -0.8$ (approximately). This value corresponds to the existence of a single checkerboard “zone” that covers a large part of the lattice while in some parts of the lattice, we find no particle or one particle per cell.

Statistical distributions. As for the stripe, a rough statistical estimation of the time needed to form a checkerboard was performed. For 1000 initial conditions, we measured the time needed to attain the arbitrary threshold of $\gamma = -0.5$. The distribution of this time is shown on Fig. 9. We observe that for $L = 32$, some trajectories take more than 2000 time steps to stabilise. However, for $L = 64$, such outliers are no longer observed and the average time to attain the threshold is 140 steps.

3.5 The cluster phase

We now observe the third pattern : we take a density of $d = 0.2$ and a “high” alignment sensitivity: $\sigma = 4$. (This value corresponds to the double of the one used in the first experiment.)

Single trajectory observation. The evolution of the system from a Bernoulli random initial condition is shown on Fig. 10. We see that the system stabilises on a stationary state formed of groups of particles, the *clusters*, which move in the same directions. These group of particles are coherent and give the impression to traverse each other without interfering much. However, the shape of the groups evolve slowly by loosing or gaining particles.

Temporal evolution. Figure 11 shows the temporal evolution of the mean alignment γ as a function of time for the same trajectory as the one observed on Fig. 10. For these initial conditions, the phenomenon of clusters formation is in general very rapid.

Statistical distribution. We measured the distribution of the time needed to attain $\gamma = 0.8$. We were surprised to observe that for $L < 128$, it is relatively frequent to observe trajectories that do *not* form clusters. (See in particular Ref. [BFC13a].) It is thus necessary to increase the size to values much higher than $L = 128$ in order to obtain stable statistical results that do not suffer from finite-size effects (see Fig. 12).

To conclude this part, we identified three main types of organisations of the system, or to use the vocabulary of statistical physics, three phases. In fact only clusters correspond to what we first expected: coherent groups of particles that move in the same direction. Stripes are surprising because they imply a symmetry breaking (one direction is chosen) and checkerboards are even more surprising because their particles are anti-aligned.




Figure 10: Clusters formation: evolution of the system from a random initial configuration. The grid size is $L = 32$, the sensitivity is $\sigma = 4.0$ and the initial density is equal to $d = 0.2$.




Figure 11: Cluster formation: evolution of the mean alignment for the sampled trajectory observed in Fig. 10: $L = 32$, $d = 0.2$, $\sigma = 4.0$.




Figure 12: Cluster formation: distribution of the time to reach $\gamma = 0.8$ for $L = 64$ and $L = 128$ (1000 samples).




Figure 13: Stripe formation for rectangular grid dimensions: $(X, Y) = (150, 30)$ and $(d, \sigma) = (0.2, 2)$; (top): $t = 2000$, the process is still under evolution (bottom): $t = 4000$, the metastable state is reached (a multi-loop stripe).

4 Robustness of the patterns

We now examine if the patterns are due to the simple and discrete nature of cellular automata. Do they come from the regularity of the lattice? Are they a “perverse effect” of the perfect synchronous updating of the lattice?

In the following, we will examine what happens to the three main patterns when they are submitted to the following perturbations:

- change in the lattice size,
- change in the type of boundary conditions in the lattice,
- change in the type of updating.

We will give a succinct view on our previous observations and indicate in which directions the research could be continued.

4.1 The size ratio of the lattice is a key parameter

The change in the ratio of the lattice reveals interesting facts on the stripe formation. We refer to our previous studies for more information on this phenomenon [BFC13b]. The main observation is that if the lattice ratio is an integer value k (for a small k), the stripe “loops” in order to form a closed pattern.

It is interesting to note that for some particular values of the grid ratio, a new “bifurcation” appears. In some cases, the stripe does not totally succeed in forming a “closed” pattern and in some other cases, the stripe forms a multi-loop closed pattern. For instance, for $(X, Y) = (100, 75)$, it was observed that the system may stabilise on *three* different patterns depending on the value of σ . As shown on Figure 14, for a sensitivity $\sigma \in [1, 1.5]$, an irregular stripe appears; for $\sigma \in [1.5, 1.8]$, a regular multi-loop stripe appears; and for higher values of σ , we recover the clusters pattern. This situation contrasts with what




Figure 14: Patterns observed for rectangular grid dimensions $(X, Y) = (100, 75)$ at $t = 5000$. Depending on the value of σ , we observe: (left) a non-perfect stripe, (middle) a regular stripe with multiple loops on the grid, (right) clusters.

can be observed for a square grid of size $(100, 100)$ in which only simple loops are observed with the same settings of σ and d .

These observations made us describe the stripe as a “resonance effect”: in order for a stripe to appear, particles need to cross the grid several times and to interact regularly with the particles that come in some perpendicular directions. In case where a good “harmony” in these interactions is found, an amplification effect is created, which stabilises greatly the pattern. It is an open question to obtain a more precise explanation of this phenomenon.

4.2 Reflecting borders suppress the single stripe

In order to get a better understanding of the mechanisms involved in the stripe formation, we propose to observe a new experiment where we use *reflecting borders*. For the sake of brevity, we put the technical descriptions of these particular boundary conditions in the Appendix (see p. 20).

Figure 15 presents a comparison of two trajectories of the system for the simulation conditions that normally lead to the formation of a stripe: $d = 0.2$ and $\sigma = 2.0$. It is surprising to observe that in both cases we no longer observe the formation of a single stripe: (a) For $L = 64$, we observe the formation of “big” clusters that travel horizontally and vertically and periodically “bounce” on the corners. (b) For $L = 128$, we observe the formation of *four* diagonal stripes that coexist and with particles that periodically bounce on the “walls” and change their directions. We observed that these two different qualitative behaviours co-exist for intermediary values of L between 64 and 128: the higher L is, the higher the probability to observe the coexistence of four diagonal stripes.

These observations corroborate the assertion that the observation of a single stripe is a kind of “resonance effect”. With reflecting borders, it is no longer possible for a single stripe to “absorb” the other stripes and the four stripes may coexist. It is an open problem to understand why there seems to exist a “minimal” grid size for the observation of the four coexisting stripes.

Contrarily to the stripe pattern, the checkerboard pattern is robust to the




Figure 15: Patterns observed with reflecting borders with two square lattices: (top) $L = 64$ (bottom) $L = 128$. Settings are: $d = 0.2$ and $\sigma = 2.0$.

change of boundary conditions: for $d = 0.4$ and $\sigma = 2.0$, reflecting borders do not perturb the formation of a checkerboard.

The situation is more complex for the clusters pattern. Once again, we observe a strong dependence on the system's size. It is common to observe the co-existence of stable horizontal and vertical patterns, diagonal stripes and, near the borders, checkerboards. (This latter phenomenon can be explained by the existence of higher concentration of particles near the borders.) However, after a long simulation time (a few thousand steps), it is observed that the diagonal stripes “win” and that the system stabilises on a four-axis-stripe.

The observations for such borders are only preliminary and they need a more extensive study. It is important to note that they often reveal a strong and unexpected dependence on the lattice size. These emergent patterns can thus not be interpreted as resulting *only* from the local rule.

4.3 Asynchronism creates a new pattern

Asynchronism in cellular automata is a rich source of questions. It has given rise to numerous studies, with various motivations and approaches. We refer to a recent survey paper for readers who would like to have an overview on the topic [Fat14].

The first version of an asynchronous lattice gas was proposed in Ref. [BFC13a]. Note that applying an asynchronous updating in this model is less straightfor-




Figure 16: Application of an asynchronous interaction at time $t = 1000$ after a checkerboard was formed. Simulation conditions: $L = 32$, $d = 0.4$, $\sigma = 2.0$.




Figure 17: Two evolutions of the Tartan observed with an synchronous interaction of $\alpha_I = 0.5$ (top) A Tartan pattern that transforms into a tripe : $d = 0.4$ (bottom) A Tartan pattern that remains stable for at least 50 000 steps. Simulation conditions: $L = 64$, $\sigma = 2.0$.

ward than for “classical” cellular automata. The updating of the lattice gas is done in two steps and we may apply the asynchrony to one step or the other, or to both steps. We will here focus on the asynchronism of *interaction*: we apply the interaction rule with probability α_I and leave the cells’ state unchanged with probability $1 - \alpha_I$. This provides a way to progressively “deviate” from the regular synchronous case. The question we ask is: What happens if we apply a small amount of asynchronism?

We empirically observed that if we start from a stripe or from a “clustered” state, no qualitative change is visible. However, if we start from a checkerboard, then we have a quick destruction of the pattern. This dissolution is shown on Fig. 16.

In some cases a new pattern can be formed: horizontal and vertical lines appear, which span throughout the lattice and “loop”. This pattern was named

tartan, with a reference to the Scottish textile. This pattern is not always stable: in most cases, the tartan is destroyed after a few thousand steps and is replaced by a diagonal stripe (see Figure 17).

The dissolution of the tartan pattern greatly depends on how close the density is to the value $1/2$. For instance, trajectory shown on Fig. 17-(bottom) shows a tartan pattern that remains stable for more than $5 \cdot 10^5$ steps. For this simulation, a fusion of two zones occurs at $t \sim 10700$. It is another open problem to understand the dynamics of the tartan pattern. We believe that, contrarily to the other “major” patterns whose stability time would be exponential in the grid size, the stability of the tartan is “only” quadratic.

5 Discussion

Let us review the main elements of our study, pattern by pattern. In particular, we come back to our initial question: can we use a simple LGCA to model the swarming phenomenon or does this make our model too simplistic?

About the diagonal stripe pattern. This pattern is an example of spontaneous symmetry breaking: from a state that is initially disordered, the system “cools down” and “chooses” one of the four possibilities of diagonal orientations for the stripe. Experiments tend to show the presence of a “resonance effect”: the pattern is created by an amplification of the waves that spontaneously emerge and travel across the lattice. Waves enter into collision and merge until they form one single stripe. The size ratio of the lattice plays a key role: if the grid is a square, a stripe forms easily, otherwise, the system may stabilise on various shapes depending on the value of the sensitivity and the size ratio. If we replace periodic boundary conditions by reflexive boundary conditions, then the symmetry breaking is no longer observed and the system shows four coexisting stripes that periodically “bounce” on the borders.

About the checkerboard pattern. The stability of the checkerboard patterns is linked to the checkerboard invariant (see Sec. 2), but it is not a direct consequence of it. On the contrary, the very existence of this pattern is surprising because the local evolution rule is *stochastic*. Recall that the transition from diagonal stripes to the checkerboards is observed by keeping the sensitivity constant and by only increasing the number of particles on the lattice. All happens as if this increase would remove some degrees of freedom and force the system to oscillate. In other words, with this pattern, although the cellular automaton is stochastic, the “noise” in the local rule is not sufficient to smooth out the discrete nature of the system.

We also observed that adding a small amount of asynchrony is sufficient to destroy the pattern. It is not yet clear whether there exists a minimal amount of asynchrony to introduce or if any quantity makes the checkerboard patterns unstable.

About the clusters pattern. Clusters are groups of particles that move coherently. They have the strange property of traversing each other without interacting too much. This property results both from the local rule *and* from the fact that the LGCA model allows particles with different directions to be simultaneously present on a single cell. It is an open problem to measure precisely the transition between the diagonal stripe pattern and the clusters pattern. Note that a simple quantitative parameter to discriminate between the two behaviours would be the average “occupancy” per non-empty cell: in the case of diagonal stripe, this parameter should be close to 2 while it should be observed close to 1 for the clusters pattern.

Finite vs. infinite lattices As mentioned earlier, the central problem to understand the behaviour of the model is its *metastability*. The patterns we encountered may remain an exponential time in the lattice size, but not an *infinite* time. We thus aim at studying *infinite lattices*, but, we have to remember that the behaviour of an infinite lattices may not always reflect the limit behaviour for finite lattices whose size tend to infinity. For example, the mechanism that produces diagonal stripes is probably different on an infinite lattice and will require a different approach than the one we had here.

To sum up, this model has an extraordinary richness of behaviour although it involves only a minimal set of definitions. Obviously, the propagation step plays a key role: the particles move and locally change their density, which allows them to produce new phenomena that would not be seen if they were static. The question is to understand why the three main patterns react so differently to small changes in the simulation conditions (for instance to an asynchronous updating or to the use of a non-square lattice). Would it be that the existence of non-robust patterns is a signature that the model is too simple? But at the same time, if we are unable to correctly analyse it mathematically, does it mean that the model is already too complex?

References

- [BDG97] Harmen J. Bussemaker, Andreas Deutsch, and Edith Geigant. Mean-field analysis of a dynamical phase transition in a cellular automaton model for collective motion. *Physical Review Letters*, 78(26):5018–5021, 1997.
- [BFC13a] Olivier Bouré, Nazim Fatès, and Vincent Chevrier. First steps on asynchronous lattice-gas models with an application to a swarming rule. *Natural Computing*, 12(4):551–560, 2013.
- [BFC13b] Olivier Bouré, Nazim Fatès, and Vincent Chevrier. A robustness approach to study metastable behaviours in a lattice-gas model of swarming. In Jarkko Kari, Martin Kutrib, and Andreas Malcher, editors, *Proceedings of Automata’13*, volume 8155 of *Lecture Notes*

- in *Computer Science*, pages 84–97. Springer, 2013. Extended version available as a tech. report at <https://hal.inria.fr/hal-00768831>.
- [BI13] Anouk Barberousse and Cyrille Imbert. New mathematics for old physics: The case of lattice fluids. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 44(3):231–241, 2013.
- [DD05] Andreas Deutsch and Sabine Dormann. *Cellular Automaton Modeling of Biological Pattern Formation – Characterization, Applications, and Analysis*. Modeling and Simulation in Science, Engineering and Technology. Birkhäuser, 2005.
- [Deu96] Andreas Deutsch. Orientation-induced pattern formation: swarm dynamics in a lattice-gas automaton model. *International Journal of Bifurcation and Chaos*, 06(09):1735–1752, 1996.
- [Fat14] Nazim Fatès. A guided tour of asynchronous cellular automata. *Journal of Cellular Automata*, 9(5-6):387–416, 2014.
- [Mar13] Irène Marcovici. *Automates cellulaires probabilistes et mesures spécifiques sur des espaces symbolique*. PhD thesis, Université Paris 7, 2013. text in English.
- [MM14] Jean Mairesse and Irène Marcovici. Around probabilistic cellular automata. *Theoretical Computer Science*, 559:42–72, 2014.
- [RST09] Damien Regnault, Nicolas Schabanel, and Eric Thierry. Progresses in the analysis of stochastic 2D cellular automata: A study of asynchronous 2D minority. *Theoretical Computer Science*, 410(47-49):4844–4855, 2009.
- [Tag15] Lorenzo Taggi. Critical probabilities and convergence time of percolation probabilistic cellular automata. *Journal of Statistical Physics*, 159(4):853–892, 2015.

Acknowledgments

The remarks of Théo Rubenach and Irène Marcovici were a precious help for the improvement of this text.

Appendix : simulation of reflecting borders in a LGCA

The simulation of the *reflecting borders* boundary conditions is applied as follows.

We initialise the system by letting each channel contain a particle with probability d . There are exceptions: (a) The border cells are all empty. (b) For the

cells situated immediately next to the northern, eastern, southern and western borders cells, we respectively empty the North, East, South and West channels. Formally, we use $\mathcal{L} = \{0, \dots, X\} \times \{0, \dots, Y\}$ and:

- $\mathcal{B} = \{(i, j) \in \mathcal{L}, i = 0 \text{ or } i = X - 1 \text{ or } j = 0 \text{ or } j = Y - 1\}$,
- $\mathcal{B}_N = \{(i, j) \in \mathcal{L}, j = Y - 1\}$,
- $\mathcal{B}_E = \{(i, j) \in \mathcal{L}, i = X - 1\}$,
- $\mathcal{B}_S = \{(i, j) \in \mathcal{L}, j = 0\}$,
- $\mathcal{B}_W = \{(i, j) \in \mathcal{L}, i = 0\}$.

By noting the initial condition $x \in Q^{\mathcal{L}}$ for each $(i, j) \in \mathcal{L}$, taking $x(i, j) = (q_n, q_e, q_s, q_w)$, we have $q_n = 0$ if $(i, j) \in \mathcal{B} \cup \mathcal{B}_N$, $q_e = 0$ if $(i, j) \in \mathcal{B} \cup \mathcal{B}_E$, etc. We call this last set of 4 conditions, the *integrity condition*.

The integrity condition guarantees that no particle will travel to a border cell (in \mathcal{B}). It is easy to see that the propagation step preserves the integrity condition, but not the interaction step. Our method thus consists in checking if the north channel of a cell of \mathcal{B}_N is occupied. In this case, the particle is re-affected among the free channels of the cell, with a uniform probability. All happens as if the particle has “bounced” on a northern wall. Clearly, such a rearrangement is always possible as this cell can not contain four particles. The same procedures is applied for the three other directions.