

Emergence of Macro Spatial Structures in Dissipative Cellular Automata

Andrea Roli¹ and Franco Zambonelli²

¹ DEIS

Università degli Studi di Bologna

`aroli@deis.unibo.it`

² DISMI

Università di Modena e Reggio Emilia

`franco.zambonelli@unimo.it`

Abstract. This paper describes the peculiar behavior observed in a class of cellular automata that we have defined as *dissipative*, i.e., cellular automata that are *open* and makes it possible for the environment to influence their evolution. Peculiar in the dynamic evolution of this class of cellular automata is that stable macro-level spatial structures emerge from local interactions among cells, a behavior that does not emerge when the cellular automaton is *closed*, i.e., when the state of a cell is not influenced by the external world. Moreover, we observed that Dissipative Cellular Automata (DCA) exhibit a behavior very similar to that of dissipative structures, as macro-level spatial structures emerge as soon as the external perturbation exceeds a threshold value and it stays below the “turbulence” limit. Finally, we discuss possible relations of the performed experiments with the area of open distributed computing, and in particular of agent-based distributed computing.

1 Introduction

In this paper, we present and discuss a set of experiments that we have performed on a new class of cellular automata that we have defined as *Dissipative Cellular Automata* (DCA). DCA differ from “traditional” cellular automata in two characteristics: while “traditional” cellular automata are composed of cells that interact with each other in a synchronous way and that are influenced in their evolution only by the internal state of the automata themselves, dissipative ones are *asynchronous* and *open*. On the one hand, cells update their status independently of each other, in an “autonomous” way. On the other hand, the automata live dipped in an environment that can directly influence the internal behavior of the automata, as in open systems.

The reported experiments show that DCA exhibit peculiar interesting behaviors. In particular, during the evolution of the DCA, and despite the out-of-equilibrium situation induced by the external environment, stable macro-level spatial structures emerge from local interactions among cells, a behavior that does not emerge when the cellular automaton is synchronous and closed (i.e.,

when the state of a cell is not influenced by the environment). Furthermore, ordered patterns emerge, like in dissipative systems [11], when the external perturbation is higher than a critical value and they are present for a specific perturbation strength range.

On this basis, the paper argues that similar sort of macro-level behaviors are likely to emerge as soon as multiagent systems (or likes) will start populating the Internet and our physical spaces, both characterized by intrinsic and unpredictable dynamics. Such behaviors are likely to dramatically influence the overall behavior of our networks at a very large scale. This may require new models, methodologies, and tools, explicitly taking into account the environmental dynamics, and exploiting it during software design and development either defensively, to control its effects on the system, or constructively, as an additional design dimension.

This paper is organized as follows. Sect. 2 defines DCA as CA characterized by asynchronous dynamics and openness. In Sect. 3 we describe experiments and we discuss the results obtained. In Sect. 4 the relation between DCA and dissipative systems is further investigated, by showing the typical system behavior as a function of the external perturbation. We conclude with Sect. 5 outlining potential applications and future work.

2 Dissipative Cellular Automata

In this section we first briefly recall the definition of Cellular Automata (CA) and introduce the terminology that will be used in the following. Then, we define Dissipative Cellular Automata (DCA) as CA characterized by *asynchronous* dynamics and *openness*.

A CA is defined by a quadruple $\mathcal{A} = (S, d, V, f)$, where S is the finite set of possible states a cell can assume, d is the dimension of the automaton, V is the neighborhood structure, and f is the local transition rule. In this work we assume what follows:

- The automaton structure is a 2-dimensional discrete grid closed to a 2-dimensional torus (namely, $N \times N$ square grids with wraparound borders).
- The neighborhood structure is regular and isotropic, i.e., V has the same definition for every cell.
- f is the same for each cell (uniform CA).

The quadruple \mathcal{A} specifies the static characteristics of an automaton. The complete description of a CA requires the definition of its dynamics, i.e., of the dynamics ruling the update of the state of CA cells. In general, the dynamics of a CA assumes a discrete time. The usual definition of CA is with synchronous dynamics: cells update their state in parallel at each time step.

Synchronous CA of this kind have been deeply studied [19,1] and have also an interesting biological/systemic interpretation: cells can be interpreted as alive/dead, or system elements active/inactive depending on their state.




Fig. 1. A synchronous CA having reached a cyclic attractor.




Fig. 2. A fixed point reached by an asynchronous CA. The initial state is the same of the synchronous one.

2.1 Asynchronous Dynamics

Accordingly to the most accepted terminology [6,10,13], a CA is *asynchronous* if cells can update their state independently from each other, rather than all together in parallel, according to a dynamics that can be either *step-driven* or *time-driven*.

In *step-driven* dynamics, a kind of global daemon is introduced, whose job is to choose at each time step one (and only one) cell to update. In *time-driven* dynamics, each cell is assumed to have an “internal clock” which wakes up the cell and makes it update. Also, time-driven dynamics provides for a more continuous notion of time. The updating signal for a cell can be either deterministic (e.g., every time steps) or probabilistic (e.g., the probability that the cell update its state is uniformly distributed), and the next state of a cell is selected on the basis of the current state of neighboring cells.

In the experiments presented in this paper, CA have an asynchronous time-driven dynamics: at each time step, a cell has a probability λ_a to wake up and update its state. The update of a cell has been implemented as atomic and mutually exclusive among neighbors, without preventing non-neighbor cells to update their state concurrently.

In general, it has been observed that the asynchronous CA exhibits behaviors which are very different from the ones of their synchronous counterparts, both in terms of transient and final attractor. Both the dynamics have the same fixed points [13], i.e., attractors that are fixed points under synchronous dynamics are fixed points also under asynchronous dynamics and vice versa. Nevertheless, trajectories in the state space and basins of attraction can be very different and some of the final attractors reached under asynchronous dynamics may be reached with lower probability under synchronous one.

As an example, Fig. 1 and Fig. 2 show the steady states reached by a synchronous and an asynchronous CA, starting from the same initial (random) state. These are characterized by a Moore neighborhood structure (the neighbors of a cell are the 8 one defining a 3×3 square around the cell itself) and the following transition rule:

$$f = \{ \text{a dead cell gets alive iff it has 2 neighbors alive; a living cells lives iff it has 1 or 2 neighbors alive} \}.$$

Under asynchronous regime, CA usually reaches a fixed point that its synchronous counterpart has never been observed to be able to reach in all the experiments we performed.

2.2 Openness

Most of CA studied so far are closed systems, as they do not take into account the interaction between the CA and an environment. Instead, the class of CA that we have studied is, in addition to asynchronous, *open*: the dynamic behavior of the CA can be influenced by the external environment.

From an operative point of view, the openness of the CA implies that some cell can be forced from the external to change its state, independently of the cell having evaluated its state and independently of the transition function (see Fig. 3).

From a thermodynamic perspective, one can consider this manifestation of the external environment in terms of energy flows: forcing a cell to change its state can be considered as a manifestation of energy flowing into the system and influencing it [11]. This similarity, together with the fact that the activities of the cells are intrinsically asynchronous and that the externally forced changes in the state of cells perturb the CA in an irreversible way, made us call this kind of CA as *Dissipative Cellular Automata* (DCA).

From a more formal point of view, a DCA can be defined as follows:

- $\mathcal{A} = (S, d, V, f)$,
- asynchronous time-driven dynamics (with probability λ_a),
- a perturbation action $\varphi(\alpha, \mathcal{D}, \lambda_e)$.

where \mathcal{A} is the quadruple defining the CA, the dynamics is the one already discussed in Subsect. 2.1, and the perturbation action φ is a transition function which acts concurrently with f and can change the state of any of the CA cells to a given state $\alpha \in V$ depending on some probabilistic distribution \mathcal{D} , independently of the current state of the cells and of their neighbors. Specifically, in our experiments $\alpha = 1$ (i.e., the cell if forced to be “alive”) and \mathcal{D} is a distribution such that each cell has probability λ_e to be perturbed.




Fig. 3. The basic structure of a dissipative cellular automaton: the environment influences the state of cells by injecting “energy”.

3 Emergent Behaviors

The behavior exhibited by DCA is dramatically different from both their synchronous and closed asynchronous counterparts.

In general, when the degree of perturbation (determined by λ_e) is high enough to effectively perturb the internal dynamic of the DCA (determined by the rate of cell updates λ_a) but it is still not prevailing over it so as to make the behavior of the DCA almost random (which happens when λ_e becomes comparable λ_a), peculiar patterns emerge. Since the external perturbation strength is relative to the cells update rate, i.e., the amount of perturbation is given by the ratio between external and internal update rate, the actual control parameter is the ratio λ_e/λ_a .

We have observed that the perturbation on the cells induced by the external – while keeping the system out of equilibrium and making impossible for it to reach any equilibrium situation – makes DCA develop large-scale regular spatial structures. Such structures show long-range correlation between the state of the cells, emerged despite the strictly local and asynchronous transition rules, and break the spatial symmetry of the final state. In addition, such structures are stable, despite the continuous perturbing effects of the external environment.

Our experiments involved many combinations of rules for cells to live/die and get alive and number of neighbors. We tested most of the rules which generate local patterns, excluding those leading to trivial attractors (i.e., all cells alive/dead). For each combination of rules and neighborhood structures we simulated the CA dynamics starting from 20 different random initial states. The interested reader can refer to the Web page: <http://polaris.ing.unimo.it/DCA/> to access our simulation files in the form of applets, and appreciate the dynamic evolution of these DCA by reproducing our experiments. In the following we will discuss and show some among the typical results of the general phenomenon we observed.




Fig. 4. Two Behaviors Evolved in a Dissipative Cellular Automata. Despite the out-of-equilibrium situation forced by the external environment, stable large-scale and symmetry-breaking patterns emerge.

For example, Fig. 4 shows two different patterns emerged from a DCA, both exhibiting stable macro-level spatial structures. For this DCA, the transition rules and the neighborhood structure are the one described in Subsect. 2.1. In both cases, the presence of global scale patterns – breaking the rotational symmetry of the automata – is apparent. By comparing these patterns with the ones observed in the same CA under asynchronous but close dynamics, one can see that openness has provided for making small scale patterns, emerged from local transition rules, enlarge to the whole CA size. Once this global states has emerged, they are able to restabilize autonomously, despite the fact that the perturbing effects tends to modify them.

As another example, Fig. 5 shows two typical patterns emerged for a DCA with a neighborhood structure made up of 12 neighbors (the neighbors of a cell are all cells having a maximum distance of 2 from the cell itself) and with the following transition rule:

$$f = \{ \text{a dead cell gets alive iff it has 6 neighbors alive; a living cells lives iff it has 3,4,5, or 6 neighbors alive} \}$$

Again, it is possible to see large symmetry-breaking patterns emerge, extending to a global scale the local patterns that tends to emerge under asynchronous but closed regime (Fig. 6). The patterns are stable despite the continuous perturbing effect of the environment. Moreover, the pattern shown on the left of Fig. 5 is dynamic. First, the long diagonal stripes change continuously in their microlevel shape, while maintaining the same global structure. Second, all this stripes translate horizontally at a constant speed in the DCA lattice.

DCA share common characteristics with *Stochastic Cellular Automata* [2, 14]). Stochastic Cellular Automata (SCA) are synchronous CA with a transition function characterized by an *outgoing probability distribution*, which biases the choice of the next cell's state. The main difference between DCA and SCA




Fig. 5. Two different behaviors evolved in a Dissipative Cellular Automaton, large-scale patterns emerge. The left picture shows a step of a dynamic pattern, with horizontally translating diagonal stripes.




Fig. 6. A stabilized situation in an asynchronous close cellular automaton following the same rules of the DCA in Fig. 5: no large-scale patterns emerge.

is that DCA's transition function is deterministic and the non-determinism is introduced by external perturbations. The DCA model enables us to describe systems composed of several independent interacting entities (asynchronous and concurrently acting) which can be affected by external perturbations.

Behaviors similar to the ones we observed in DCA have also been obtained for synchronous CA in [20], where long-range patterns are generated by means of peculiar transition functions which explicitly introduce symmetry-breaking rules. Therefore, even though the observed behavior is similar, the emergence of macro-level spatial structures has a different origin: in our case, no symmetry-breaking rules are introduced and the regular patterns are generated by the combination of "symmetric" transition functions, asynchrony and external perturbation.

4 Explaining DCA Dynamics

DCA behavior exhibits a strong analogy with the behavior of dissipative systems [11], e.g., Benard's cells. A fluid between two plates is in thermodynamic equilibrium if no thermal energy flows from the external to perturb the equilibrium. In presence of small differences between the temperature of the two plates, the thermal energy is still not enough to perturb the fluid, and energy flows between the two plates in the form of thermal diffusion. However, as soon as the temperature gradient reaches a critical point, thermal flow in the fluid starts occurring via convection. This motion does not occur in a disordered way: regular spatial patterns of movement emerge, with wide-range and symmetry breaking correlation among cell movements. This behavior is maintained until the temperature gradient between the two plates become too high, in which case the regular patterns disappear and the fluid motion becomes turbulent.

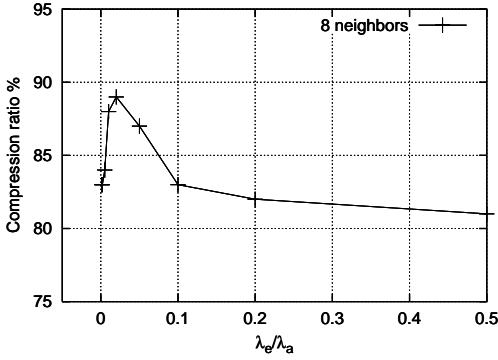
By analogy, we conjecture that the behavior of DCA might be subject to the same phenomenon, where the temperature gradient between the two plates is substituted by the ratio λ_e/λ_a . When this ratio is 0, the system is in equilibrium, and no perturbation from the external occur. For very small perturbation, the dynamic behavior of the DCA does not substantially change. As soon as the ratio becomes high enough, the DCA dynamics changes and regular spatial patterns appears. For very high ratio, spatial patterns disappear and the DCA dynamics becomes highly disordered.

A rough measure of the emergence of macro-level structures can be provided by the compression percentage achieved by compression algorithms. The higher the compression factor, the lower the randomness of the CA configuration. Although this measure does not directly evidence long-range correlations, it nevertheless provides meaningful information about the amount of structure of a CA state. Fig. 7 shows typical results obtained with DCA with different number of neighbors and transition function. The states of DCA have been measured once the equilibrium have been reached¹. As we can observe, when the ratio λ_e/λ_a approaches a critical value θ_1 the compression ratio cr abruptly increases. This corresponds to the onset of structure in the system. cr reaches a maximum approximately located at $\lambda_e/\lambda_a \approx 0.05$. Then it decreases till reaching again the initial values, indicating the disappearance of macro-level structures in DCA states. We observe that, for the first two DCA in Fig. 7, cr quickly decreases, on the opposite of the last one, for which it seems that structures are still present for higher values of the ration λ_e/λ_a . In general, for all the experiments performed, we observed that the critical value for the onset of structured patterns is approximately $\lambda_e/\lambda_a \approx 0.05$.

The above similarity suggests that the same causes that determine the behavior of Benard cells also determine the behavior of DCA.

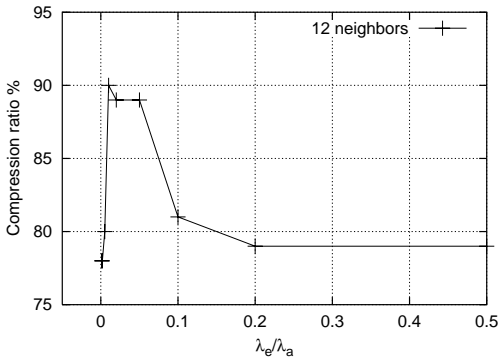
Without any perturbation, or in the presence of small one, each autonomous component (a molecule or a DCA cell), acting asynchronously accordingly to

¹ The compression algorithm used is that provided by usual compression utilities like *gzip*, at maximum compression level.



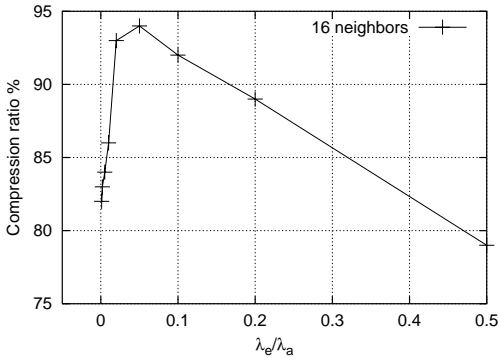
8 neighbors

A dead cell gets alive iff it has 2 neighbors alive; a living cells lives iff it has 1 or 2 neighbors alive.



12 neighbors

A dead cell gets alive iff it has 6 neighbors alive; a living cells lives iff it has 3-6 neighbors alive.



16 neighbors

A dead cell gets alive iff it has 6 neighbors alive; a living cells lives iff it has 3-6 neighbors alive.

Fig. 7. Amount of structure as a function of the ratio λ_e/λ_a . The compression ratio is evaluated by a usual compression algorithm (*gzip*). Observe that macro-level structures appear only in a specific range, as in dissipative structures.

strictly local rules, tend to reach a local equilibrium (or a strictly local dynamics), which produce a global uniform equilibrium of the whole system.

When the system is kept in a substantial out-of-equilibrium situation, the locally reached equilibrium situations are continuously perturbed, resulting in continuous attempt to locally reestablish equilibrium. This typically ends up with cell groups having found new equilibrium states more robust with regard to the perturbation (or compatible with it). Such stable local patterns start soon dominating and influencing the surrounding, in a sort of positive feedback, until a globally coordinated (i.e., with large-scale spatial patterns) and stable situation emerges.

When the degree of perturbation is high enough to avoid local stable situations to persist for enough time, they can no longer influence the whole systems, and the situation becomes turbulent.

5 Conclusion and Future Work

This paper has reported the outcomes of a set of experiments performed on a new class of cellular automata, DCA, which are open to the environment and can be perturbed by its dynamics. These experiments have shown that the perturbation makes large-scale symmetry breaking spatial structures, not observed under closed regime, emerge. By introducing a measure of the randomness of DCA states we have shown that structures emerge when the external perturbation is higher than a critical value and below the turbulence limit.

The experiments reported in this paper are indeed preliminary, and further work is in progress:

- we are currently exploring different measures for evaluating the emergence of large-scale patterns. For example, we may consider techniques analogous to the ones presented in [4,3,5], where structure is measured by evaluating the complexity of the probabilistic automaton reconstructed from the data series representing the CA evolution. Other possibilities rely on the application of techniques derived from image analysis (for example, we may use spatial correlation measures);
- we are extending our DCA simulation framework so as to study the behavior of network structures other than the regular ones of DCA, such as small-world graphs [18] and boolean networks [9], as well as networks with mobile nodes;
- we intend to perform further experiments to evaluate the behavior of DCA under different perturbation regimes and to experiment with more complex DCA, i.e., DCA with large set of states and/or with non-uniform transition functions [17,16].

The results presented in this paper promise to have several potential implications in the area of distributed computing. In fact, DCA exhibit characteristics (i.e., autonomy of components, locality in interactions, openness to the environment) that are typical of modern distributed computing environments, e.g., sensor networks and multi-agent systems.

Agents are autonomous entities [7], as their execution is not subject to a global flow of control. Indeed, the execution of an agent in a multiagent system may proceed asynchronously, and the agent's state transition occur according to local internal timings. This is actually what happens in DCA, because of the adopted time-driven dynamics. Moreover, agents are situated entities that live dipped in an environment, whether a computational one, e.g., a Web site, or a physical one, e.g., a room or a manufacturing unit to be controlled. The agent is typically influenced in its execution (i.e., in its state transitions) by what it senses in the environment. In this sense, agents and multi-agent systems are "open systems": the global evolution of a multi-agent system may be influenced by the environment in which it lives. And, in most of the cases, the environment possesses a dynamics which is not controllable or foreseeable. For instance, computational resources, data, services, as well as the other agents to be found on a given Web site cannot be predicted and they are likely to change in time. This sort of openness is the same that we can find in DCA, where the perturbation of the environment, changing the internal state of a cell, can make us consider the cell as situated in an environment whose characteristics dynamically change in an unpredictable way.

Given the above similarities, we argue that similar sort of macro-level behaviors are likely to make their appearance also in such systems, raising the need for models, methodologies, and tools, explicitly taking into account the autonomy and environmental dynamics and exploiting them either constructively, to achieve globally coordinated behaviors, or defensively, to control the behavior of the system. On the one hand, one could think at exploiting the environmental dynamics to control and influence a multi-agent system from "outside the loop" [15], that is, without intervening on the system itself. In a world of continuous computations, where decentralized software systems are always running and cannot be stopped (this is already the case for Internet services and for embedded sensors) changing, maintaining and updating systems by stopping and re-installing them is not the best solution, and it could not be always feasible. On the other hand, the reported experiments open up the possibility that a software system immersed in a dynamic environment may exhibit behaviors very different from the ones it was programmed for. Obviously, this is not desirable and may cause highly damaging effects.

Of course, we are not the first discussing the possibility of emergence of complex self-organizing behaviors in multi-agent systems. However, most of the studies (apart from a few exceptions [12]) have focused on "closed" agent systems, in which the internal dynamics of the systems totally drive its behavior. Instead, we have shown, via a very simple and "minimal" multi-agent system, as a DCA can be considered, that complex non-local behaviors can emerge due to the influence of the environmental dynamics. The impact of this observation in the modeling, engineering, and maintaining of distributed agent systems may be dramatic [8,21].

References

1. Y. Bar-Yam. *Dynamics of Complex systems*. Addison Wesley, 1997.
2. T.D. Barfoot and G.M.T. D'Eleuterio. Multiagent Coordination by Stochastic Cellular Automata. Proceeding of IJCAI 2001, Seattle, 2001.
3. J.P. Crutchfield. Discovering Coherent Structures in Nonlinear Spatial Systems. In *Nonlinear Dynamics of Ocean Waves*, 1992.
4. J.P. Crutchfield and J.E. Hanson. Turbulent Pattern Bases for Cellular Automata. *Physica D* **69**:279–301, 1993.
5. J.E. Hanson and J.P. Crutchfield. The Attractor-Basin Portrait of a Cellular Automaton. *J. Statistical Physics* **66**:1415–1462, 1992.
6. T.E. Ingerson and R.L. Buvel. Structure in asynchronous cellular automata. *Physica D*, 10:59–68, 1984.
7. N.R. Jennings. On Agent-Based Software Engineering. *Artificial Intelligence*, 117(2), 2000.
8. T. Kuhn. *The Structure of Scientific Revolutions*. University of Chicago Press, 3rd Edition, Nov. 1996.
9. S.A. Kauffman. *The origins of order*. Oxford University Press, New York, 1993.
10. E.D. Lumer and G. Nocolis. Synchronous versus asynchronous dynamics in spatially distributed systems. *Physica D*, 71:440–452, 1994.
11. G. Nocolis and I. Prigogine. *Exploring Complexity: an Introduction* W. H. Freeman (NY), 1989.
12. V. Parunak and S. Bruekner. Entropy and Self-Organization in Agent Systems. 5th International Conference on Autonomous Agents, ACM Press, May 2001.
13. B. Schönfisch and A. de Roos. Synchronous and asynchronous updating in cellular automata. *BioSystems*, 51(3):123–143, 1999.
14. B. Schönfisch and M.O. Vlad. Physical approach to the ergodic behavior of stochastic cellular automata with generalization to random processes with infinite memory. *Physica A* **229**:273–294, 1996.
15. D. Tennenhouse. Proactive Computing. *Communications of the ACM*, May 2000.
16. M. Sipper. The Emergence of Cellular Computing. *IEEE Computer*, 37(7):18–26, July 1999.
17. M. Sipper and M. Tomassini. Computation in artificially evolved, non-uniform cellular automata. *Theoretical Computer Science*, 217(1):81–98, March 1999.
18. D. Watts. *Small-Worlds*. Princeton University Press, 1999.
19. S. Wolfram. *Cellular Automata and Complexity*. Addison–Wesley, 1994.
20. S. Wolfram. *A New Kind of Science*. Wolfram Media Inc., 2002.
21. F. Zambonelli and H.V. Parunak. From Design to Intention: Signs of a Revolution. Proceedings of the 1st ACM Joint Conference on Autonomous Agents and Multi-Agent Systems, 2002.