

Multiagent Coordination by Stochastic Cellular Automata

T D Barfoot and G M T D'Eleuterio

tim.barfoot@utoronto.ca, gde@utias.utoronto.ca

University of Toronto Institute for Aerospace Studies

4925 Dufferin Street, Toronto, Ontario, Canada, M3H 5T6

Abstract

A coordination mechanism for a system of sparsely communicating agents is described. The mechanism is based on a stochastic version of cellular automata. A parameter similar to a temperature can be tuned to change the behaviour of the system. It is found that the best coordination occurs near a phase transition between order and chaos. Coordination does not rely on any particular structure of the connections between agents, thus it may be applicable to a large array of sparsely communicating mobile robots.

1 Introduction

The term *multiagent system* encompasses large bodies of work from engineering, computer science, and mathematics. Examples include networks of mobile robots [Mataric, 1992], software agents [Bonabeau *et al.*, 1994], and cellular automata [Wolfram, 1984]. A common thread in all multiagent systems is the issue of *coordination*. How are a large number of sparsely coupled agents able to produce a coherent global behaviour using simple rules? Answering this question will not only permit the construction of interesting and useful artificial systems but may allow us to understand more about the natural world. Ants and the other social insects are perfect examples of local interaction producing a coherent global behaviour. It is possible for millions of ants to act as a *super-organism* through local pheromone communication. We seek to reproduce this ability on a fundamental level in order to coordinate artificial systems.

It can be argued that *cellular automata* (CA) are the simplest example of a multiagent system. Originally studied by [von Neumann, 1966], the term CA is used to describe systems of sparsely coupled difference equations. Despite their simple mechanics, some extremely interesting behaviours have been catalogued (e.g., Conway's *Game of Life*). The word *self-organization* is used in many contexts when discussing multiagent systems which can lead to confusion. Here we use it to mean *multiagent coordination in the face of more than one alternative*. We will be describing a stochastic version of cellular automata. The goal will be to have all cells choose the same symbol from a number of possibilities using only sparse communication. We maintain that rules able to

succeed at this task are self-organizing because the cells are not told which symbol to choose, yet they must all coordinate their choices to produce a globally coherent decision. If we told the cells which symbol to choose, the task would be very easy and no communication between cells would be necessary. This can be dubbed *centralized organization* and is in stark contrast to *self- or decentralized organization*. We believe that coordination in the face of more than one alternative is at the very heart of all multiagent systems.

This paper is organized as follows. Related work is described, followed by a description of the model under consideration. Results of its performance on the multiagent coordination task are presented. Statistical analysis of the rule are provided followed by discussions and conclusions.

2 Related Work

In the following note that typically cellular automata do not operate in a stochastic but rather a deterministic manner. Unless explicitly stated (e.g., *stochastic cellular automata* (SCA)), the term cellular automata will imply determinism.

[von Neumann, 1966] originally studied cellular automata in the context of self-reproducing mechanisms. The goal was to devise local rules which would reproduce and thus spread an initial pattern over a large area of cells, in a tiled fashion. The current work can be thought of as a simple case of this where the tile size is only a single cell but there are multiple possibilities for that tile. Furthermore, we wish our rules to work starting from any random initial condition of the system.

Cellular automata were categorized by the work of [Wolfram, 1984] in which four *universality classes* were identified. All rules were shown to belong to one of class I (fixed point), class II (oscillatory), class III (chaotic), or class IV (long transient). These universality classes can also be identified in SCA and we will show that in our particular model, choosing a parameter such that the system displays long transient behaviour (e.g., class IV) results in the best performance on our multiagent coordination task.

[Langton, 1990] has argued that natural computation may be linked to the universality classes. It was shown that by tuning a parameter to produce different CA rules, a phase transition was exhibited. The relation between the phase transition and the universality classes was explored. It was found that class IV behaviour appeared in the vicinity of the phase transition. The current work is very comparable to this study in

that we also have a parameter which can be tuned to produce different CA rules. However, our parameter tunes the amount of randomness that is incorporated into the system. At one end of the spectrum, completely random behaviour ensues while at the other completely deterministic behaviour ensues. We also relate the universality classes to particular ranges of our parameter and find a correlation between performance on our multiagent coordination task and class IV behaviour. We attempt to use similar statistical measures to [Langton, 1990] to quantify our findings.

[Mitchell *et al.*, 1993], [Das *et al.*, 1995] study the same coordination task as will be examined here in the case of deterministic CA. However, their approach is to use a genetic algorithm to evolve rules successful at the task whereas here hand-coded rules are described. They found that the best solutions were able to send long range *particles* (similar to those in the *Game of Life*) [Andre *et al.*, 1997] in order to achieve coordination. These particles rely on the underlying structure of the connections between cells, specifically that each cell is connected to its neighbours in an identical manner. The current work assumes that no such underlying structure may be exploited and that the same mechanism should work for different connective architectures. The cost for this increased versatility is that the resulting rules are less efficient (in terms of time to coordinate) than their particle-based counterparts.

[Tanaka-Yamawaki *et al.*, 1996] studies the same problem to that considered here. They use *totalistic* [Wolfram, 1984] rules which do not permit exploitation of the underlying structure of the connections between cells but rather rely on the intensity of each incoming symbol. They also vary a parameter to produce different rules and find that above a certain threshold, “global consensus” occurs but below it does not. However, they consider large clusters of symbols to be a successful global consensus. We do not and thus turn to a stochastic version of their totalistic rules to destroy these clusters and complete the job of global coordination.

3 The Model

In deterministic cellular automata there is an *alphabet* of K symbols, one of which may be adopted by each cell. Incoming connections each provide a cell with one of these symbols. The combination of all incoming symbols uniquely determines which symbol the cell will display as output. Stochastic cellular automata (SCA) work in the very same way except at the output level. Instead of there being a single unique symbol which is adopted with probability 1, there can be multiple symbols adopted with probability less than 1. Based on this *outgoing probability distribution* over the K symbols, a single unique symbol is drawn to be the output of the cell. This is done for all cells simultaneously. It should be noted that deterministic CA are a special case of SCA.

We consider a specific sub-case of SCA in this paper which corresponds to the totalistic rules of CA. Assume that cells cannot tell which symbols came from which connections. In this case, it is only the intensity of each incoming symbol which becomes important. Furthermore, we desire that our rules work with any number of incoming connections thus rather than using the number of each of the incoming K sym-

bols, we use this number normalized by the number of connections which can be thought of as an *incoming probability distribution*. In summary the model we consider is as follows.

Totalistic SCA. Consider a system of N cells, each of which is connected to a number of other cells. Let \mathcal{A} represent an alphabet of K symbols. The state of Cell i at time-step t is $x_i[t] \in \mathcal{A}$. The input probability distribution, \mathbf{p}_{in} , for Cell i is given by

$$\mathbf{p}_{in}[t] = \sigma_i(x_1[t], \dots, x_N[t]) \quad (1)$$

where σ_i accounts for the connections of Cell i to the other cells. The output probability distribution \mathbf{p}_{out} is given by the map, φ ,

$$\mathbf{p}_{out}[t+1] = \varphi(\mathbf{p}_{in}[t]) \quad (2)$$

The probability distributions \mathbf{p}_{in} and \mathbf{p}_{out} are stochastic columns. The new state of Cell i at time-step $t+1$ is randomly drawn according to the distribution $\mathbf{p}_{out}[t+1]$ and is represented by $x_i[t+1]$.

It should be noted that in (1) if the connections between the cells are not changing over time then the functions, $\sigma_i(\cdot)$, will not be functions of time. However, we could allow these connections to change which would make them functions of time.

Once the connections are described through the $\sigma_i(\cdot)$ functions, the only thing that remains to be defined is the φ -map. We assume that each cell has the same φ -map but this need not be the case. The possibilities for this map are infinite and thus for the remainder of this paper we discuss a parameterized subset of these possibilities. This subset will be called *piecewise- φ* and is defined as follows.

Piecewise- φ . Let

$$\mathbf{p}_{in} = [p_{1,in} \cdots p_{K,in}]^T \quad (3)$$

The (unnormalized) output probabilities are given by

$$p_{k,out} = \begin{cases} 1, & \text{if } \frac{1}{K} + \beta(p_{k,in} - \frac{1}{K}) \geq 1 \\ 0, & \text{if } \frac{1}{K} + \beta(p_{k,in} - \frac{1}{K}) \leq 0 \\ \frac{1}{K} + \beta(p_{k,in} - \frac{1}{K}), & \text{otherwise} \end{cases} \quad (4)$$

where β is derived from the tunable parameter λ as follows:

$$\beta = \begin{cases} 2\lambda, & \text{if } 0 \leq \lambda \leq \frac{1}{2} \\ \frac{1}{2}(1-\lambda)^{-1}, & \text{if } \frac{1}{2} \leq \lambda < 1 \end{cases} \quad (5)$$

The (normalized) output probability column is

$$\mathbf{p}_{out} = \frac{1}{p_{\Sigma,out}} [p_{1,out} \cdots p_{K,out}]^T \quad (6)$$

where $p_{\Sigma,out} = \sum_{k=1}^K p_{k,out}$.

Note that in (5), the tunable parameter, λ acts in a similar manner to a temperature parameter. When $\lambda = 1$ we have a completely deterministic rule while when $\lambda = 0$ we have a completely random rule. Figure 1 shows what the rule looks like for different λ when $K = 2$.

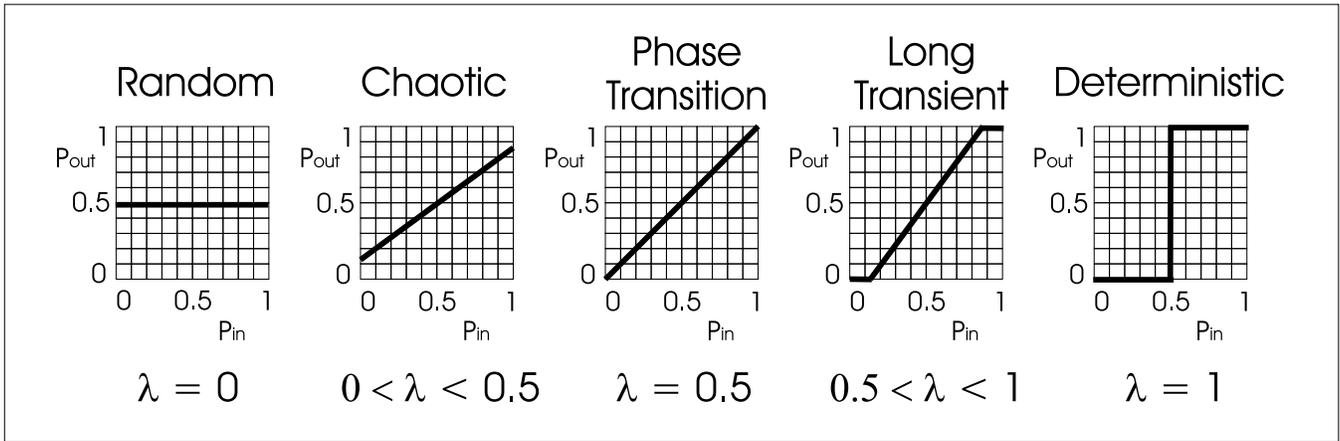


Figure 1: The *piecewise- φ* rule for different values of λ and $K = 2$.

An equilibrium point, \mathbf{p}^* , in a φ -map is one for which the following is true

$$\mathbf{p}^* = \varphi(\mathbf{p}^*) \quad (7)$$

The idea behind the *piecewise- φ* rule was to create an instability in the probability map at the uniform distribution equilibrium point

$$\mathbf{p}_{\text{uni}}^* = \left[\frac{1}{K} \quad \frac{1}{K} \quad \dots \quad \frac{1}{K} \right]^T \quad (8)$$

such that a small perturbation from this point would drive the probability towards one of the stable equilibria

$$\mathbf{p}_1^* = [1 \ 0 \ \dots \ 0]^T \quad (9)$$

$$\mathbf{p}_2^* = [0 \ 1 \ \dots \ 0]^T \quad (10)$$

$$\vdots \quad (11)$$

$$\mathbf{p}_K^* = [0 \ 0 \ \dots \ 1]^T \quad (12)$$

It turns out that when $0 \leq \lambda < \frac{1}{2}$, the equilibrium point, $\mathbf{p}_{\text{uni}}^*$, is the only stable equilibrium. However when $\frac{1}{2} < \lambda \leq 1$, $\mathbf{p}_{\text{uni}}^*$ becomes unstable and the other equilibria, $\mathbf{p}_1^*, \dots, \mathbf{p}_K^*$, become stable. This is similar to the classic pitchfork bifurcation as depicted in figure 2 for $K = 2$. However, with K symbols in the alphabet the pitchfork will have K tines.

It is important to stress that we have designed the stability of our system at a local level. The question of global stability and success on the multiagent coordination problem does not follow directly from the local stability of each cell. It might be possible to study the global stability of a system of cells with the *piecewise- φ* rule analytically. The approach in this paper has been to study it through simulation and statistical measures.

4 Simulation

We now present simulations of cells running the *piecewise- φ* rule. In order to ensure that the connections between cells are not regular, we consider each cell to exist in a Cartesian box (of size 1 by 1). The N cells are randomly positioned in this

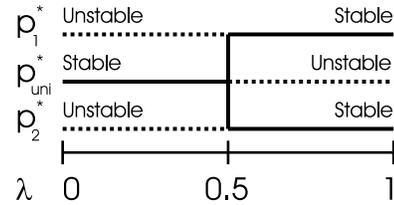


Figure 2: Pitchfork stability of the *piecewise- φ* rule for $K = 2$. λ is a parameter analogous to a temperature.

box and symmetrical connections are formed between two cells if they are closer than a threshold Euclidean distance, d , from one another. Figure 4 shows example connections between $N = 100$ cells with $d = 0.2$. Figure 3 shows example time series for different values of λ . When $\lambda < 0.5$, chaotic global behaviour arises, with $0.5 < \lambda < 1$ fairly successful behaviour results but with $\lambda = 1$ clusters form. The formation of clusters means that the global system has stable equilibria which we did not predict from the local rule. However, as λ is decreased towards 0.5, these equilibria are no longer stable and the system continues to coordinate.

It would seem that there is a good correlation between the stability on the local level and the behaviour type of the global system. As λ moves from below 0.5 to above, it appears there is a dramatic phase transition in the behaviour of the system (totally chaotic to fixed point). In the neighbourhood of 0.5 there is long transient behaviour. It turns out that the best value for λ , from the point of view of multiagent coordination, is approximately $\lambda = 0.6$.

5 Statistics

In an attempt to quantify the qualitative observations of the previous section a number of statistical measures were employed in the analysis of the SCA time series. These were used also by [Langton, 1990]. The first measure is taken from [Shannon, 1948] and will be referred to as *entropy* (H). It is defined as follows.

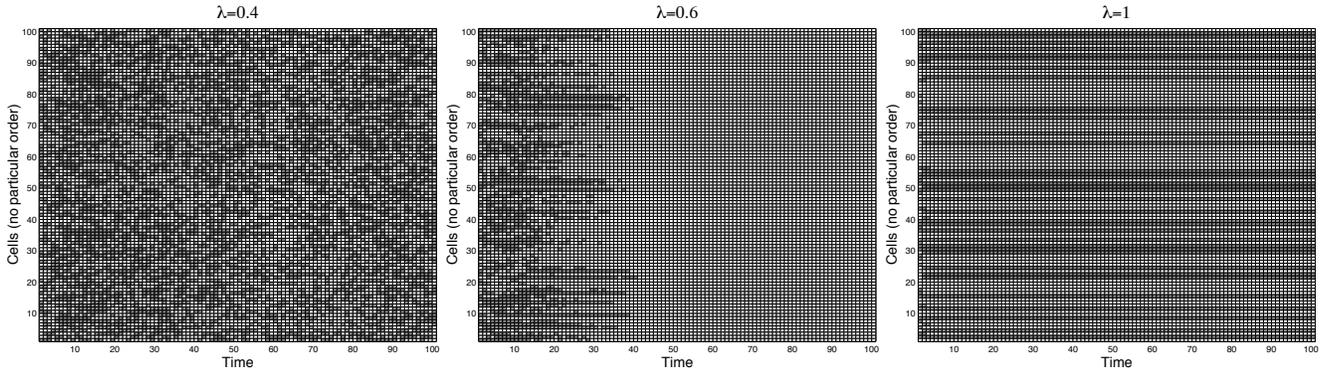


Figure 3: Example time series for different values of λ and $N = 100$, $K = 2$, $d = 0.2$. (left) chaotic behaviour, (middle) successful coordination, (right) clusters. The 2 colours represent the 2 symbols of the alphabet.

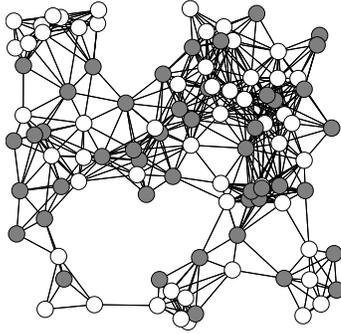


Figure 4: Example connections between $N = 100$ cells with $d = 0.2$. K colours represent the K symbols of the alphabet. Here an initial random condition is displayed for $K = 2$.

Entropy Given a sequence of M symbols

$$\mathbf{s} = [s_1 \ s_2 \ \dots \ s_M]^T \quad (13)$$

from an alphabet of size K , the entropy of the sequence may be computed as follows. First compute the frequency, n_k , of each of the K symbols $\forall k \in 1 \dots K$ which is simply the number of occurrences of symbol k in the sequence, \mathbf{s} . From the frequencies, compute the probability, p_k , of each of the K symbols $\forall k \in 1 \dots K$ as

$$p_k = \frac{n_k}{n_\Sigma} \quad (14)$$

where $n_\Sigma = \sum_{i=1}^K n_i$. Finally, the entropy of sequence, $H(\mathbf{s})$, is defined as

$$H(\mathbf{s}) = \frac{-\sum_{k=1}^K p_k \ln(p_k)}{\ln(K)} \quad (15)$$

where the $\ln(K)$ denominator is a normalization constant to make $H(\mathbf{s}) \in [0, 1]$.

This entropy function produces a value of 0 when all the symbols in \mathbf{s} are identical and a value of 1 when all K symbols are equally common. The second measure is based on the first and will be referred to as *mutual information* (I). It is defined as

Mutual Information Given two sequences of M symbols

$$\mathbf{s}_1 = [s_{1,1} \ s_{1,2} \ \dots \ s_{1,M}]^T \quad (16)$$

$$\mathbf{s}_2 = [s_{2,1} \ s_{2,2} \ \dots \ s_{2,M}]^T \quad (17)$$

from an alphabet of size K , the mutual information of the sequence, $I(\mathbf{s}_1, \mathbf{s}_2)$, may be defined as

$$I(\mathbf{s}_1, \mathbf{s}_2) = H(\mathbf{s}_1) + H(\mathbf{s}_2) - H(\mathbf{s}_1, \mathbf{s}_2) \quad (18)$$

where $H(\mathbf{s}_1, \mathbf{s}_2)$ is the entropy of the two sequences considered as a joint process (i.e., with an alphabet of size $K \times K$).

These two measures may be computed on any sequence of symbols. We tested them on *spatial* sequences (e.g., time series *columns* from figure 3) and *temporal* sequences (e.g., time series *rows* from figure 3). The most interesting measures were *average spatial entropy* (average of entropies computed from all columns in a time series) and *average temporal mutual information* (average of all I s computed from all rows in a time series. I was computed between a row and itself shifted by one time-step).

Figure 5 show various measures for 1000 values of λ . At each value of λ , 100 simulations were done on different random connections between cells and initial conditions. Thus, all displayed measures are actually averaged over 100 simulations. Each simulation was run for 300 time-steps with $N = 100$, $K = 2$, and $d = 0.2$.

Figure 5 (left) shows the average number of clusters¹ at the final time-step for different values of λ . Clearly there is an optimal value of λ near 0.6. Figure 5 (middle) shows average spatial entropy for different values of λ . This measure has a good correlation with average number of clusters. Again, there is a minimum occurring at approximately $\lambda = 0.6$ which corresponds to the best performance at multiagent coordination.

Figure 5 (right) displays average temporal mutual information for different values of λ . This is a very interesting plot.

¹Number of clusters was computed by considering the SCA as a Markov chain with connections deleted between cells displaying different symbols. The number of clusters is then the number of eigenvalues equal to 1 from the Markov transition matrix.

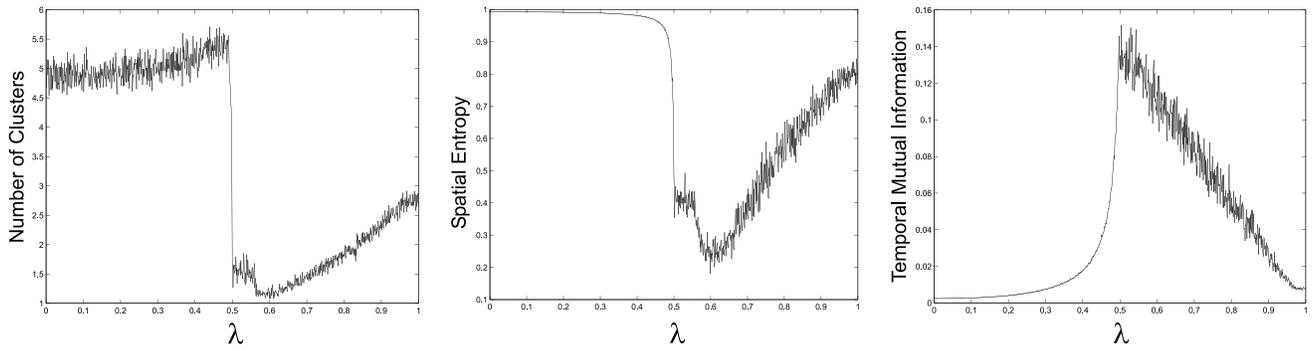


Figure 5: (left) Average number of clusters at final time-step for 1000 values of λ . (middle) Average spatial entropy for 1000 values of λ . (right) Average temporal mutual information for 1000 values of λ . All plots show average of 100 simulations at each value of λ .

Temporal mutual information seems to capture the length of the global transient behaviour of the system. As discussed in [Langton, 1990], the random pattern in the chaotic region is not considered transient but rather the steady state behaviour. The peak in temporal mutual information occurs at $\lambda = 0.5$, the phase transition, and drops away on either side. [Langton, 1990] has a similar plot. Figure 6 shows how the number of clusters at the final time-step (300 time-steps) changes as the problem scales up to more cells; it appears to be a linear relationship. Figure 7 shows how the number of clusters at the final time-step changes for different message sizes.

6 Discussion

The strong correlation between the local stability of the *piecewise- φ* rule and the type of global behaviour is quite interesting. It appears that $\lambda \gg 0.5$ corresponds to fixed point behaviour (class I), $\lambda \ll 0.5$ corresponds to chaotic behaviour (class III), and λ near 0.5 corresponds to long transient behaviour (class IV). The correlation most likely has something to do with the way in which the incoming probability distribution is computed in (1). This step delivers information averaged from all connected cells. This averaging serves to smooth out differences between connected cells. However, if this smoothing occurs too quickly (i.e., $\lambda = 1$) the system does not have time to smooth globally resulting in the formation of clusters. The addition of noise in the particular form of the *piecewise- φ* rule aids in slowing the smoothing process thus destroying the clusters. This has been called *critical slowing down* [Haken, 1983] in other systems. As we approach the critical point ($\lambda = 0.5$ or $\beta = 1$) from above, the strength of the instability decreases which slows down the decision-making process. It is a balance of these two effects which seems to be the most effective at multiagent coordination. The optimal operating value of λ is not right at the phase transition but a little bit towards the deterministic end of the λ spectrum (approximately $\lambda = 0.6$).

Note that we did not find any oscillatory behaviour (class II) which is likely because the connections between the cells are symmetrical. However, if the *piecewise- φ* rule in figure 1 is reflected (left-right) then the system ‘blinks’ and global coordination corresponds to all cells blinking in phase with one another.

In this model of multiagent coordination, the boundaries between clusters have purposely been made unstable. This forces them to move randomly until they contact one another and annihilate, leaving a single cluster. The results presented here used $N = 100$ cells and required on average 150 time-steps to get to a single cluster with $d = 0.2$, $K = 2$ and $\lambda = 0.6$. Clearly the time required to form a single cluster will increase with the number of cells in the system. Figure 6 confirms this by showing that at the end of 300 time-steps, increasing the number of cells, N , results in more clusters. The linear relationship suggests that scaling-up may be possible but more in depth studies are required. Figure 7 shows how the system scales to different message sizes, K . Here the relationship between number of clusters (after 300 time-steps) to message size is a bit surprising, first dropping then increasing as K increases. It levels off again as the number of symbols exceeds the number of cells (since at most N symbols can be represented in the random initial condition). Again, the nature of this scaling should be studied more closely.

The *piecewise- φ* rule is not the only map that can be used to achieve multiagent coordination in SCA. Replacing it with other monotonically increasing functions (i.e., in figure 1) with the same equilibria also works. We had comparable success to the *piecewise- φ* map using

$$p_{k,\text{out}} = (p_{k,\text{in}})^\beta \quad \forall k \in 1 \dots K \quad (19)$$

with the outgoing probability column normalized as in (6).

The model considered here does not require knowledge of the underlying structure of the connections between cells. This was a design requirement as it was originally motivated by a network of communicating mobile robots whose connections might be changing over time and thus difficult to exploit. It is thus natural to question whether the model still works as the connections are varied over time. To this end, a small amount of Gaussian noise was added to the positions of the cells in the Cartesian box of figure 4 at each time-step. As the cells moved, the connections between them changed (since they are limited by the range, d). The SCA model was still able to form single clusters. This was possible even when $\lambda = 1$ which does make sense since there is still some noise being added. However, the nature of the noise is at the connection level rather than the signal level. This aspect is currently under further investigation.

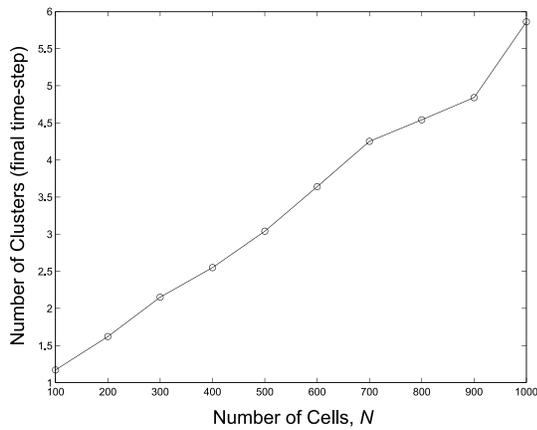


Figure 6: Number of clusters at final time-step as the number of cells is varied from $N = 100$ to $N = 1000$. Parameters for each run were 300 time-steps, $\lambda = 0.6$, $K = 0.2$ and $d = \frac{0.2}{\sqrt{N}}$ to keep the density of connections constant. Plot shows average of 100 simulations at each value of N .

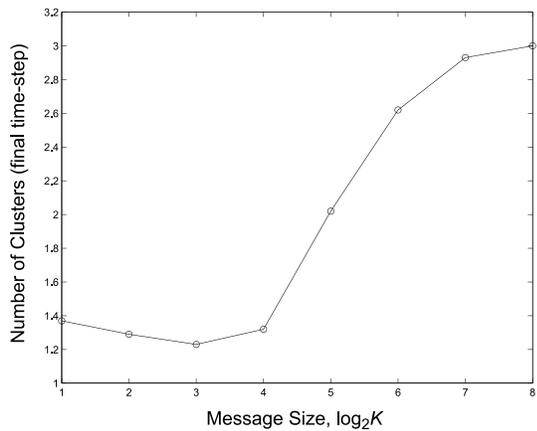


Figure 7: Number of clusters at the final time-step as the message size is varied from $K = 2$ (1 bit) to $K = 256$ (8 bits). Parameters for each run were 300 time-steps, $\lambda = 0.6$, $N = 100$ and $d = 0.2$. Plot shows average of 100 simulations at each value of K .

7 Conclusion

A mechanism for multiagent coordination has been presented based on stochastic cellular automata. We consider this to be an example of self-organizing behaviour in that global coordination occurs in the face of more than one alternative. It was shown that by using stochastic rules, sparsely communicating agents could come to a global consensus. A parameter in the coordination mechanism was tuned and it was found that coordination occurred best when the system was near a phase transition between chaotic and ordered behaviour (the optimum was a little bit towards the ordered side).

It is hoped that this model will shed light on self-organization as a general concept while at the same time providing a simple algorithm to be used in practice.

Acknowledgements

We would like to acknowledge support of this work by NSERC (Natural Science and Engineering Research Council) and CSA (Canadian Space Agency).

References

- [Andre *et al.*, 1997] David Andre, Forrest H. Bennett, and John R. Koza. Evolution of intricate long-distance communication signals in cellular automata using genetic programming. In Chris G. Langton and Katsunori Shimohara, editors, *Artificial Life V, Proceedings of the Fifth International Workshop on the Synthesis and Simulation of Living Systems*. MIT Press, 1997.
- [Bonabeau *et al.*, 1994] Eric Bonabeau, Guy Theraulaz, Eric Arpin, and Emmanuel Sardet. The building behaviour of lattice swarms. In Rodney A. Brooks and Pattie Maes, editors, *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*. MIT Press, 1994.
- [Das *et al.*, 1995] Rajarshi Das, James P. Crutchfield, Melanie Mitchell, and James E. Hanson. Evolving globally synchronized cellular automata. In L.J. Eshelman, editor, *Proceedings of the Sixth International Conference on Genetic Algorithms*, pages 336–343, San Francisco, CA, April 1995.
- [Haken, 1983] H. Haken. *Synergetics, An Introduction, 3rd Edition*. Springer-Verlag, Berlin, 1983.
- [Langton, 1990] Chris G. Langton. Computations at the edge of chaos: Phase transitions and emergent computation. *Physica D*, 42:12–37, 1990.
- [Mataric, 1992] Maja J. Mataric. Designing emergent behaviours: From local interactions to collective intelligence. In J. A. Meyer, H.L. Roitblat, and S. W. Wilson, editors, *Simulation of Adaptive Behaviour: from Animals to Animals 2*. MIT Press, 1992.
- [Mitchell *et al.*, 1993] Melanie Mitchell, Peter T. Hraber, and James P. Crutchfield. Revisiting the edge of chaos: Evolving cellular automata to perform computations. *Complex Systems*, 7:89–130, 1993. SFI Working Paper 93-03-014.
- [Shannon, 1948] C E Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423, July 1948.
- [Tanaka-Yamawaki *et al.*, 1996] Mieko Tanaka-Yamawaki, Sachiko Kitamikado, and Toshio Fukuda. Consensus formation and the cellular automata. *Robotics and Autonomous Systems*, 19:15–22, 1996.
- [von Neumann, 1966] Jon von Neumann. *Theory of Self-Reproducing Automata*. University of Illinois Press, Urbana and London, 1966.
- [Wolfram, 1984] Stephen Wolfram. Universality and complexity in cellular automata. *Physica D*, 10:1–35, 1984.