Minority Games and Distributed Coordination in Non-Stationary Environments

Aram Galstyan and Kristina Lerman
Information Sciences Institute
University of Southern California
4676 Admiralty Way
Marina del Rey, CA 90292-6695

Abstract - In this paper we examine emergent coordination in a network of competing boolean agents. The agents play so called Generalized Minority Game where the capacity level is allowed to vary externally. We study the properties of such a system for different values of the mean connectivity K of the network, and show that the system with $K = 2$ shows a high degree of coordination for relatively large variations of the capacity level. We also show that for $K > 2$ coordination can be achieved by tuning the homogeneity parameter of the agents' boolean strategies.

I. Introduction

Autonomous multi-agent systems (MAS) have become one of the more prominent and active research areas in the AI community. The interest is easy to understand: in the future, many of the military and industrial tasks, from intelligence gathering and reconnaissance, to toxic waste clean-up to running a large organization, will be assigned to and carried out by autonomous intelligent agents, whether they are information agents, mobile robots, or nodes in a sensor or a communications network. It is, therefore, crucial to understand the issues of MAS design, robustness, control, scalability, dynamics of collective behavior, and adaptability in uncertain hostile environments.

Multi-agent systems in real environments will have to have the following characteristics: they have to be (1) scalable: composed of thousands or more agents; (2) robust: robust to individual agent failure; (3) flexible: agents can be dynamically added or removed, they can be given the ability to reallocate and redistribute themselves in a self-organized way; (4) adaptable: will need to adapt to a changing environment; (5) low complexity: will need to operate with minimal communications overhead. Distributed market-based control satisfies the above requirements, and it is, therefore, preferable over architectures based on central control.

Previous research has shown that market-based (or game–dynamical) control strategies, in which agents make "economically" motivated decisions with the goal of maximizing individual payoff but which result in the optimization of some global system property, is a feasible distributed control mechanism for multi-agent systems. Hence, understanding the nature of emergent coordination in these type of systems is of great importance.

In this paper we present a model of distributed coordination in non–stationary environments. We consider a network of interconnected boolean agents that compete for a resource with a limited capacity. At each time step the agents face a binary choice of whether to use the resource or not, and those who used the resource are rewarded (punished) if their number is less (greater) than the resource capacity. The non–stationarity is modelled by time dependent capacity level. We show that under certain conditions our model shows globally adaptive and coordinated behavior, resulting in efficient resource allocation.

II. Minority Games

The Minority Game [1] (MG) is a simple model of competing MAS yet it has a very rich and complex dynamics. It was introduced by Challet and Zhang as a simplification of Arthur's El Farol Bar attendance problem [2]. The MG consists of N agents with bounded rationality that repeatedly choose between two alternatives labelled 0 and 1 (*e.g.*, staying at home or going to the bar). At each time step, agents who made the minority decision win. In the Generalized Minority Game [3], the wining group is 1 (0) if the fraction of the agents who chose "1" is smaller (greater) than the capacity level η, $0 < \eta < 1$. For $\eta = 0.5$, the game reduces to the the traditional MG. Each agent uses a set of S strategies to decide its next move and reinforces strategies that would have predicted the winning group. A strategy is simply a lookup table that prescribes a binary output for all possible inputs. In the original version of the game, the input is a binary string containing the last m outcomes of the game, so the agents interact by sharing the same global signal. If the agents choose either action with probability $1/2$ (the random choice game), then, on average, the number of agents choosing "1" (henceforth referred to as attendance) is $(N-1)/2$ with standard deviation $\sigma = \sqrt{N}/2$ in the limit of large N. The most interesting phenomenon of the minority model is the emergence of a coordinated phase, where the standard deviation of attendance, the volatility, becomes smaller than in the random choice game. The coordination is achieved for memory sizes for which the dimension of the reduced strategy space is comparable to the number of agents in the system, $2^m \sim N$ [4], [5]. It was later pointed out that the dynamics of the game remains mostly unchanged if one replaces the string with the actual histories with a random one [6], provided that all the agents act on the same signal. Analytical studies based on this simplification has revealed many interesting properties of the minority model[7], [8].

In addition to the original MG, different versions of the game where the agents interact using local information only (cellular automata [9], evolving random boolean networks [10], personal histories [11]), have been studied. In particular, it was established that coordination still arises out of local interactions, and the system as a whole achieves "better than random" performance in terms of the utilization of resources.

In all previous studies the capacity level has been fixed as an external parameter, so the environment in which the agents compete is stationary. As we mentioned above, however, we are interested in a situation where the environment is changing. It is interesting to see if a coordinated behavior still emerges, and to what degree agents can adapt to the changing environment. Namely, we study a system of boolean agents playing a generalized minority game, assuming that the capacity level is not fixed but varies with time, $\eta(t) = \eta_0 + \eta_1(t)$, where $\eta_1(t)$ is a time dependent perturbation. The framework of the interactions is based on Kauffman NK random boolean nets [12], where each agent gets its input from K other randomly chosen agents, and maps the input to a new state according to a boolean function of K variables, which is also randomly chosen and quenched throughout the dynamics of the system. The generalization we make is that agents are allowed to adapt by having more than one boolean function, or strategy, and the use of a particular strategy is determined by an agent based on how often it predicted the winning group throughout the game. Note that this approach is very different from adaptation through evolution studied previously in the context of the minority model [10].

III. The Model

We consider a set of N boolean agents described by "spin" variables $s_i = \{0, 1\}, i = 1, \ldots, N$. Each agent gets its input from K other randomly chosen agents, and maps the input to a new state:

$$s_i(t+1) = F_i^j(s_{k_1}(t), s_{k_2}(t), ..., s_{k_K}(t)) \tag{1}$$

where $s_{k_i},\ i = 1, \ldots, K$ are the set of neighbors, and $F_i^j, j = 1, \ldots, S$ are randomly chosen boolean functions (called strategies hereafter) used by the i-th agent. Both the set of neighbors and strategies are chosen randomly and quenched throughout the game. For each strategy F_i^j, the agent keeps a score $U_i^j(t)$ that monitors the performance of that strategy, adding (subtracting) a point if the strategy predicted the winning (loosing) side. Let the "attendance" $A(t)$ be the cumulative output of the system at time t, $A(t) = \sum_{i=1}^N s_i(t)$. Then the winning choice is "$1''$ if $A(t) \leq N\eta(t)$, and "$0''$ otherwise. Those in the wining group are awarded a point while the others loose one. Agents play the strategies that have predicted the winning side most often, and the ties are broken randomly.

As a global measure of efficiency we introduce $\delta(t) = A(t) - N\eta(t)$, that describes the deviation from the optimal resource utilization. In the following we will be primarily interested in the cumulative "waste" over a certain time window:

$$\sigma^2 = \frac{1}{T_0} \sum_{t=t_0}^{t_0+T_0} \delta(t)^2 \tag{2}$$

For $\eta_1(t) = 0$ this quantity is simply the squared standard deviation as defined in the traditional minority game.

We can compare the performance of our system to a default random choice game, defined as follows: assume that the agents are told what is the capacity $\eta(t)$ at a given time step, and they choose to go to the bar with probability $\eta(t)$. In this case the main attendance is close to $\eta(t)N$ at each time step, and the fluctuations around the mean are given by the standard deviation

$$\sigma_0^2 = N \frac{1}{T} \int_{T_0}^{T_0+T} dt' \eta(t')[1 - \eta(t')] \tag{3}$$

IV. Results

We performed intensive numerical simulations of the system described above, with the number of agents ranging from 100 to 5000, and for different choices of $\eta(t)$ and network connectivity K. For each K, a set of strategies was chosen for each agent randomly and independently from a pool of 2^{2^K} possible boolean functions, and was quenched throughout the game. In all simulations we used $S = 2$ strategies per agent. Starting from a random initial configuration, the system evolved according to the specified rules. The duration of the simulation T_0 was determined by the particular choice of $\eta(t)$. Although in our simulations we used different forms for the perturbation $\eta_1(t)$, in this letter we consider periodic perturbations only.

Our main observation is that networks with $K = 2$ show a tendency towards self organization into a phase characterized by small fluctuations, hence, an effective utilization of the resource, even for relatively large variations in the capacity level $\eta(t)$. Note, that in the Kauffman nets with $K > 2$ the dynamics of the system is chaotic with an exponentially increasing

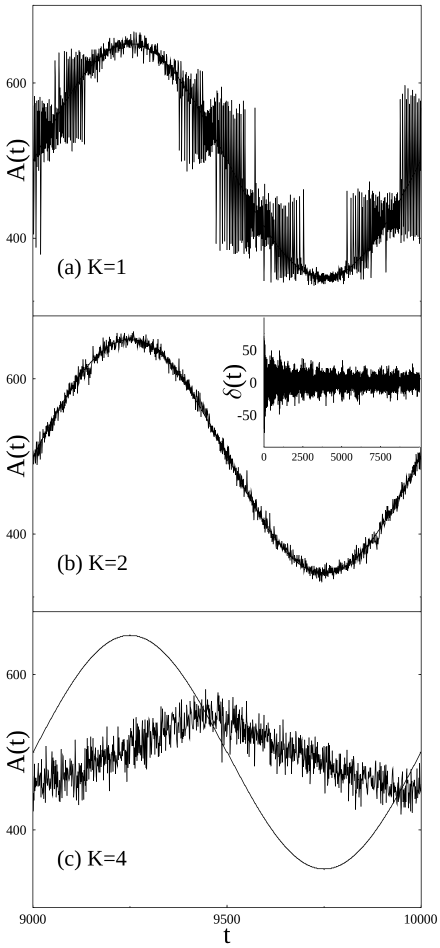

Fig. 1. A segment of the attendance time series for $\eta(t) = 0.5 + 0.15 sin(2\pi t/T)$,$T = 1000$ and different network connectivity K.

length of attractors as the system size grows, while for $K < 2$ the network reaches a frozen configuration. The case $K = 2$ corresponds to a phase transition in the dynamical properties of the network and is often referred as the "edge of the chaos". We would like to reiterate, however, that our system is different from a Kauffman network since the agents have an internal degree of freedom, characterized by their strategies. Specifically, our system does not necessarily have periodic attractors, while in Kauffman nets periodic attractors are guaranteed to exist due to the finite phase space and quenched rules of updating.

Fig. 1 shows a typical segment of the time series of the attendance $A(t)$ for a system of size $N = 1000$, a sinusoidal perturbation $\eta_1(t)$, and different network connectivities. For $K = 1$ the agents react to the changes in the capacity level, however there are strong fluctuations in the attendance series.

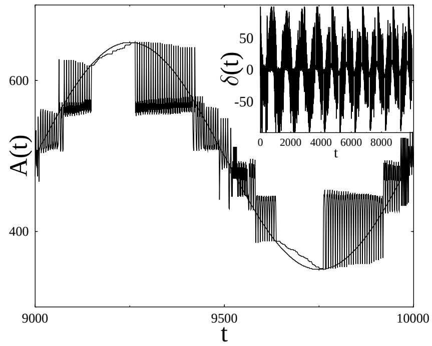

Fig. 2. Same as in Fig 1 for the traditional MG (global histories) with memory size m=5

For $K = 4$, on the other hand, response of the system to the environmental dynamics is very week, and, as our results indicate, becomes even weaker for larger K. Remarkably, the system with $K = 2$ adapts very efficiently to changes in the capacity level. The inset of Fig. 1 b) shows the time series of the deviation $\delta(t)$ for $K = 2$. Initially there are strong fluctuations, hence poor utilization of the resource, but after some transient time the system as a whole adapts and the strength of the fluctuations decreases. In fact, the standard deviation of the fluctuations is considerably smaller than in the random choice game as defined by Eq. (3). Note also, that the agents have information only about the winning choice, but not the capacity level. This suggests that the particular form of the perturbation may not be important as long as it meets some general criteria for smoothness.

For comparison, we also studied the effect of the changing capacity level in the traditional (generalized) minority model whit publicly available information about the last m outcomes of the game. We plot the attendance and deviation time series for a system with a memory size $m = 6$ (corresponding to the minimum of σ) in Fig. 2. One can see that in this case also the system reacts to the external change. However, the structure of adaptation is very different from the previous case. Indeed, an analysis of Fig. 2 shows that even though the total "wealth", i.e., the total points accumulated by the agents in the system, increases with time, the overall performance in terms of resource allocation as described by σ is much poorer compared to the previous case. Another important difference is that in the traditional system the distribution of wealth among the players is much wider than in the system with local information exchange, i.e., the later mechanism of adaptation is socially more "fair".

In Fig. 3 we plot the attendance and deviation time series for $K = 2$ and a different choice of $\eta(t)$, that contains both slow and fast components in the perturbation. Even in the presence of two vastly different time scales, the system utilizes the resource very efficiently. Another interesting observation is that if we run the simulations long enough, the response of the system to the changing capacity level gets "out of phase" with the perturbation, and gradually becomes chaotic, leading to a deterioration in the performance of the system. We found that the time during which the efficient phase is stable depends strongly on the rate of change in the capacity level. Although we are still examining the reasons for this behavior, our preliminary results indicate that it is due to the "freezing" of a certain fraction of agents because of an increased gap in the strategy scores. In other words, it seems that the system as a whole becomes less flexible.

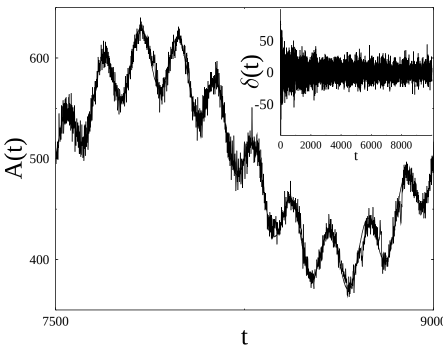

Fig. 3. A segment of the attendance time series for $K = 2$, $\eta(t) = 0.5 + 0.1sin(2\pi t/T_1) + 0.03\sin(2\pi t/T_2)$, $T_1 = 1500$, $T_2 = 150$.

In Fig. 4 we plot the variance per agent versus network connectivity K, for system sizes $N = 100, 200, 500, 1000$. For each K we performed 32 runs and averaged results. Our simulations suggest that the details of this dependence are not very sensitive to the particular form of the perturbation $\eta_1(t)$, and the general picture is the same for a wide range of functions, provided that they are smooth enough. As we already mentioned, the variance attains its minimum for $K = 2$ independent of the number of agents in the system. Note that this is different from the traditional minority game, where the position of the minimum scales logarithmically with N. For bigger K it saturates at a value that depends on the amplitude of the perturbation and on the number of agents in the system. We found that for large K the time series of the attendance closely resembles the time series in the absence of perturbation. This implies that for large K the agents do not "feel" the change in the capacity level. Consequently, the standard deviation increases linearly with the number of agents in the system, $\sigma \propto N$. For $K = 2$, on the other hand, the scaling has the form $\sigma \propto N^{1/2}$.

V. Discussion

Though the results presented here look very interesting, we currently do not have an analytical theory for the observed emergent coordination. In contrast to the traditional minority game, where global interactions and the Markovian approximation allow one to construct a mean field description, our model seems to be analytically intractable due to the explicit emphasis on local information processing. We strongly believe, however, that the adaptability of the networks with $K = 2$ is related to the peculiar properties of the corresponding Kauffman nets, and particularly, to the phase transition between the chaotic and frozen phases. It is known[13] that the phase transition in the Kauffman networks can be achieved by tuning the homogeneity parameter P which is the fraction of 1's or 0's in the output of the boolean functions (whichever is the majority), with the critical value given by $P_c = 1/2 + 1/2\sqrt{1 - 2/K}$. To test our hypothesis, we studied the properties of networks with $K = 3$ for a range of homogeneity parameter P. In Fig. 5 we plot σ^2/N ver-

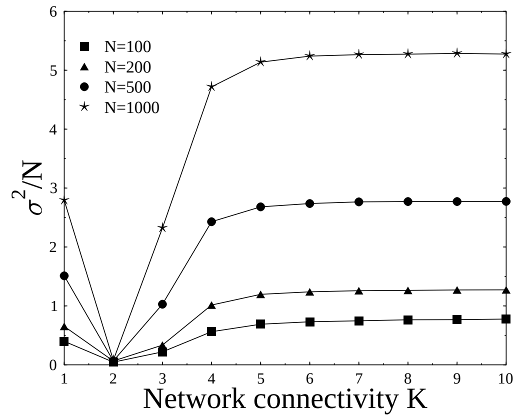

Fig. 4. σ^2/N vs the network connectivity for different system sizes

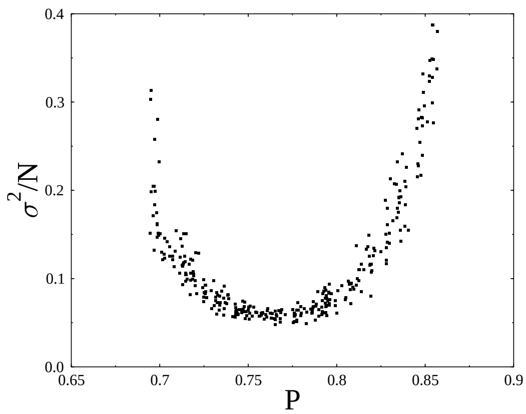

Fig. 5. Standard deviation per agent vs homogeneity coefficient P for K=3 networks: N=1000, $\eta(t) = 0.5 + 0.15\sin(2\pi t/T)$, $T = 1000$

sus the homogeneity parameter P. One can see that the optimal resource allocation is indeed achieved in the vicinity of the $P_c \sim 0.78$, indicating that the properties of Kauffman networks at the phase transition might be responsible in this emergent coordination. However, more studies are needed for a definite answer.

VI. Conclusion

In conclusion, we studied a network of adaptive boolean agents competing in a dynamic environment. We established that networks with connectivity $K = 2$ can be extremely adaptable and robust with respect to capacity level changes. For $K > 2$ the coordination can be achieved by tuning the homogeneity parameter to its critical value. Remarkably, the adaptation happens without the agents knowing the capacity level. Interestingly, the system that uses local information is much more efficient in a dynamic environment than a system that uses global information. This suggests that our model may serve as a feasible mechanism for distributed resource allocation in a multi-agent system.

VII. Acknowledgements

This research was supported by the Defense Advanced Research Projects Agency (DARPA) under contract number F30602-00-2-057.

References

[1] D. Challet and Y.-C. Zhang, "Emergence of Cooperation and Organization in an Evolutionary Game", Physica A **246**, 407 (1997)

[2] W. Arthur, "Inductive Reasoning and Bounded Rationality", Am. Econ. Assoc. Papers Proc. **84**, 406 (1994)

[3] N.F. Johnson, P.M. Hui, Dafang Zheng and C.W. Tai , "Minority Game with Arbitrary Cuttoffs", Physica A **269**, 493 (1999).

[4] D. Challet and Y.-C. Zhang, "On the Minority Game: Analytical and Numerical Studies", Physica A **256**, 514 (1998)

[5] R. Savit, R. Manuca, and R. Riolo , "Adaptive Competition, Market Efficiency, Phase Transition", Phys. Rev. Lett., 82(10), 2203 (1999)

[6] A. Cavagna, "Irrelevance of Memory in the Minority Game", Phys. Rev. E **59**, R3783 (1999); A. Cavagna, "Comment on "Adaptive Competition, Market Efficiency, and Phase Transitions"", Phys. Rev. Lett. **84** 5, 1058 (2000); R. Savit, "Reply to Comment on "Adaptive Competition, Market Efficiency, and Phase Transitions"", Phys. Rev. Lett. **84**, 1059 (2000)

[7] D. Challet, M. Marsili, "Phase Transition and Symmetry Breaking in the Minority Game", Phys. Rev. E **60**, R6271 (1999)

[8] See http://www.unifr.ch/econophysics/minority/ for an extensive collection of articles and refernecs

[9] T. Kalinowski, H.-J. Schulz and M. Briese, "Cooperation in the Minority Game with local information", Physica A **277**, 502 (2000)

[10] M. Paczuski, K. E. Bassler, "Self-organized Networks of Competing Boolean Agents", Phys. Rev. Lett. **84**, 3185 (2000)

[11] M.A.R. de Cara, O. Pla, F. Guinea, "Learning, Competition and Cooperation in Simple Games", Eur. Phys. J. **B** 13, 413 (2000)

[12] S.A. Kauffman, *The Origins of Order*, (Oxford University Press, New York, 1993)

[13] B. Derrida and Y. Pomeau, "Random Networks of Automata: A Simple Annealed Approximation" Europhys. Lett. **1**, 45 (1986)