Engineering Self-Adaptation in Agent Societies: An Infection-Based Mechanism Overview*

Norman Salazar, Juan A. Rodriguez-Aguilar, and Josep Ll. Arcos

IIIA, Artificial Intelligence Research Institute
CSIC, Spanish National Research Council
{norman,jar,arcos}iiia.csic.es

Abstract. We propose a computational *self-adapting* mechanism that facilitates agents to distributedly *evolve* their social behavior to reach the best social conventions. Our approach borrows from the social contagion phenomenon: social conventions are akin to infectious diseases that spread themselves through members of the society. Furthermore, we experimentally show that our mechanism helps a MAS to regulate itself by searching and establishing (better) social conventions on a wide range of interaction topologies and dynamic environments.

1 Introduction

Distributed mechanisms that regulate the behavior of autonomous agents in multi-agent systems (MAS) have become necessary because centralized approaches relying on global knowledge are not viable in open MAS. Furthermore, it is difficult for centralized approaches to cope with dynamic environments. We observe that in actual-world societies, social behavior is self-regulated through social conventions. These conventions emerge in a decentralized manner to balance personal interests with respect to the society's, so that each member can pursue its individual goals without preventing other members to pursue theirs.

From a sociological point of view, conventions result when members of a population adhere to some behavior, which is neither dictated nor enforced by a central authority. They can be regarded as rules followed by most members of the society, which are created and self-perpetuated by such members. Thus, the emergence of social conventions can be regarded as a self-organizing process.

One of the trends of thought in social studies is that conventions emerge by propagation or contagion, where social facilitation and imitation are key factors [4, 3]. From the MAS point of view, the studies in [13, 12] show that convention emergence is possible. However, these works limit to analyze propagation, leaving out innovation (the discovery of rules), which is a very important factor for the evolution of societies. When the aim is to help a MAS reach conventions in dynamic environments, propagation may not be enough since this assumes that at least some agent in the society knows the appropriate behavior, and this is

* An extended version of this paper has been published in IEEE's SASO 08 [11]

not always the case. Additionally, the problem can become even more difficult when the aim is to also reach the best conventions.

In this work we attempt at going beyond finding conventions. We propose an evolutionary computational mechanism that facilitates agents in a MAS to self-organize and self-adapt in such a manner that the best conventions dynamically emerge for a wide range of interaction topologies.

Notice that further evolutionary approaches appear in the literature. Indeed, convention emergence has been described as an evolutionary process [2], and so evolutionary algorithms (EA) have been employed to find conventions in agent societies. Nevertheless, they are usually applied either: (i) as a centralized process [6]; or (ii) as an individual self-contained process for each agent [7]. Both approaches can be potentially slow and tend to be off-line processes. This makes them unsuitable for dynamically adapting conventions, which is our purpose.

2 An Evolutionary Infection-Based Mechanism

We propose a computational mechanism that helps agents in a MAS reach *conventions* that maximize the social welfare. At this aim, we assume that we can accomplish our goal by maximizing agents' *individual welfares*. Thus, we stay in page with the distributed nature of the problem.

Social contagion [3] is a phenomenon that relates the spreading of behaviors between individuals, in a society, to an infectious disease. Hence, we chose to model the social contagion phenomenon into a MAS framework. However, we target beneficial conventions that if possible tend to maximize the social welfare. Considering the social welfare as a composition of individual welfares, it makes sense to let the individual behaviors that impact positively on it, here named *good behaviors*, be more infectious. Nevertheless, positive infection at most achieves a total replication of the best-known behavior among agents. Therefore, behavior innovation is also required.

By this means we expect that a MAS can reach conventions that are dominant in the society so that no better ones can be found and no worst ones can upstage them. However, if some unaccounted factor(s) alter(s) the MAS in such a manner that the current conventions become obsolete, the infectious process will reconfigure the conventions toward a better social welfare.

In our infection-based mechanism, each agent has a set of genes that encodes its behavior. Agents can infect other agents with their genes following the *survival of the fittest* concept: the fittest the agent (the highest its individual welfare), the more infectious. Furthermore, our algorithm realizes innovation (exploration) by letting agents mutate their genes. Importantly, this process runs distributedly: each agent decides whether to infect or mutate based on local knowledge.

Thus, each agent is endowed with: i) an **evaluation function** to assess its individual welfare; ii) a **selection process** to select an infecting peer, out of its local neighborhood, based on its fitness; iii) an **infection operator** to inject some of its genes into the selected agent ; and iv) an **innovation operator** to create new behaviors. See [10] for a detailed description.

3 Case Study: Coordination Game

Agents in the coordination game interact with each other by engaging in iterative games. Each game has multiple rounds. During a round each agent randomly selects a neighbor agent to play with. A play consists in both agents doing an action, either A or B. The actions are constrained in each agent by its current social rules. Plays are rewarded with a payoff, which is accumulated after each game round. The payoff for a round can be: -1,1 or α based on the agent's current action and the action of the playing neighbor (different, both B and both A respectively). This payoff captures pure coordination games [13][12] ($\alpha = 1$) and coordination games with equilibrium differing in social efficiency [9] ($\alpha > 1$).

Each agent has two parameterized rules: one to help it decide what action to take based on the last opponent's past action; and another one to decide the action to take when no past action is known. Thus, our mechanism has the task of finding for each agent the appropriated actions for these rules.

We know beforehand that four cooperative-only conventions exist (conventions that always try to cooperate), and also that they are the strongest attractors. Two of them always make agents do A (*A-conventions*) and the other two always do B (*B-conventions*). A-conventions give higher payoffs when $\alpha > 1$.

It is well known that the behavior of infections its affected by the type of topology on which a population interacts [14, 8]. Therefore, in order to empirically analyze such effects in our mechanism we chose the following interaction topologies: *small-world*, $W_{1000}^{10,0.1}$; *scale-free*, $S_{1000}^{10,-3}$; and *random graphs*, R_{1000}^{10}.

4 Empirical Results

Our experiments where designed to show that our infection-based mechanism can: i) self-organize the agents in the MAS to reach the best social convention(s) for a wide range of initial social rule configurations and under the most common interaction topologies; and (ii) realize self-adaptation in the presence of dynamic (changing) conditions.

At this aim, each experiment is defined by a combination of: i) an interaction topology model; ii) a payoff: $\alpha \in \{1, 1.5, 2\}$; and iii) an initial rule distribution, drawn from: a) **random** (rules are randomly set); b) **attractor-free** (rules set from the non-cooperative-only convention); c) **low sub-optimal** (25% of the agents with B-Conventions rules); d) **high sub-optimal** (75% of agents in a B-Convention); and e) **fully sub-optimal** (all agents with B-Convention rules).

We run 50 simulations of each experiment. In a simulation agents interact and infect each other, as described above, during 20000 ticks. We coutend at each tick the agents with the same rules, and the agents doing A or B. These counts where then aggregated per experiment using the inter-quartile mean.

4.1 Self-Organization Empirical Results

In the pure coordination game ($\alpha = 1$) the MAS establishes one of the best conventions regardless of the initial rule distribution and independently of the inter-

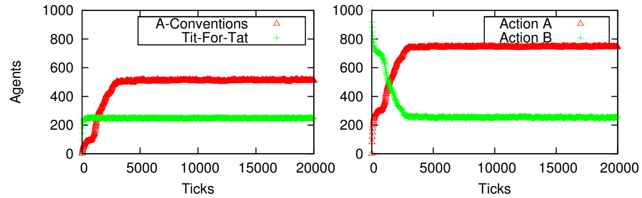

(a) Scale-free

Fig. 1. Results of experiments with full sub-optimal initialization. On the left the agents per convention; on the right the agents per action.

action topology. Whereas for different social efficiencies case ($\alpha > 1$) the results depend mainly on the initial rule distribution. When using **random**,**attractor-free** or **low sub-optimal** initial distribution, a MAS readily establishes in an A-convention for $\alpha > 1.0$ independently of the interaction topology.

Departing from a **high sub-optimal** distribution, a MAS establishes in a B-convention when $\alpha = 1.5$ for all interaction topologies. However, by setting α to 2.0, the small-world networks manage to establish an A-convention. Thus, we conclude that agents will not consider a new convention if its benefit is not significant enough. And that for the scale-free case a greater benefit is needed.

The **fully sub-optimal** distribution represents the worst case scenario. In this case, innovation becomes a key factor. When the innovation probability is low, (experiments above), the MAS is unable to converge to the best convention, because innovating agents are not able to overcome the high peer pressure. Even more, infected scale-free networks are hard to overcome [5, 8].

At this aim, we increased the innovation rate. In this manner scale-free ($\alpha = 2.0$) and small-world ($\alpha > 1$) converge to an A-Convention. This happens because a small group of agents playing *tit-for-tat* kind of rules starts to appear (see figure 1). Agents with this strategy can coexist with B-Convention agents with a small or non-negative effect to their accumulated payoffs (left-hand plots of figure 1. Therefore, when agents with an A-convention rules appear, they have a higher chance of having neighbors that will cooperate with them. However, a high mutation presents the disadvantage that a small part of the population will be constantly trying to innovate. In our case this translate in around 80% of action A convergence (Figure 1).

Overall from the experiments we conclude that *highly-clustered* agent communities (e.g. small-world) are more open to positive infections, where as the *low-clustered* ones (e.g. scale-free) are harder to infect if a stable infection is already in place. This is similar to some results shown by the scenarios studied in [9]. However, our mechanism can overcome the difficulty of re-infecting *low-clustered* networks by using a high innovation through mutation rate. Nevertheless, there is an associated cost to this high innovation: a small subgroup of agents unable to settle on a convention.

Finally, we can claim that i) a convention is always reached, and ii) under certain conditions this convention is the best one for all topologies. Moreover, when

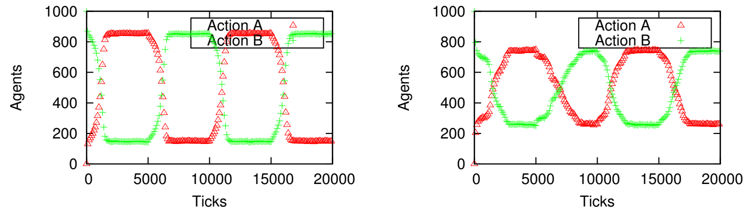

(a) Small-world (b) Scale-free

Fig. 2. Results of the dynamic payoff matrix experiments.

these conditions are not met, e.g. a suboptimal convention is fully established, our mechanism can still reach the best convention through innovation.

4.2 Self-Adaptation Empirical Results

In the previous subsection, it was shown that our infection-based mechanism endows social agents in a MAS with self-organization capabilities. Next, we shall show through experiments that it also functions as a self-adaptation mechanism that allows agents to re-organize themselves in the presence of dynamic changes.

The experiment definition is the same as the one used in section 4 with the addition of a dynamic component. This dynamic component can take the form of either run-time changes in the payoff matrix, *dynamic payoff matrix*, or an ever-changing agent population , *dynamic population and neighborhood*.

Dynamic payoff matrix. We simulate a dynamic environment by introducing changes into the payoff matrix at run time. The changes take the form of swapping the efficient action, which means that if A is the most efficient action then after the swap B will become the most efficient one (i.e the payoff values of cooperating in A are swapped with the values of cooperating in B) and vice versa. Notice that agents are not explicitly informed when the payoff matrix changes. Instead, they realize that the games they play lead to different results.

We performed experiments using: the matrices with different social efficiencies ($\alpha = 1.5$ and $\alpha = 2.0$); the scale-free ($S_{1000}^{<10>,-3}$) and small-world ($W_{1000}^{<10>,0.1}$) topologies;the full sub-optimal social rules initialization (the worst case scenario); and the matrix change occurred every 5000 ticks.

Figure 2 shows the number of agents performing each of the possible actions. We observe that after each matrix change (at ticks 5000, 10000 and 15000) the agents quickly re-organize into conventions that result in performing the most efficient action. We also observe that in the small-world, the *reaction time* is faster than the scale-free. By reaction time we mean the time elapsed between the change and the re-organization. This result was expected since the small-world is highly clustered. The results from the experiments clearly show that our infection-based mechanism endows agents with self-adapting capabilities.

Dynamic population and neighborhood. In this environment, the number of agents in the MAS changes and so their neighborhoods. In practice these environment changes are achieved by dynamically changing the network topology. Specifically we inter-wove the Barabasi-Albert (BA) network generation algorithm [1], and the MAS simulation. In other words, the MAS and the BA algorithm are executed at the same time.

The component combination used for the experiment was: a payoff matrix with $\alpha = 2.0$; scale-free topology that started at $S_{400}^{<10>,-3}$ and ended at $S_{2400}^{<10>,-3}$; and a full sub-optimal social rules initialization for both the initial agents and the new ones. The new agents were added every 50 simulation ticks.

The experiments show that, even in a MAS with a continuous influx of agents with less than optimal social rules, our mechanism is able not to only reach the best convention, but also to sway most of the incoming agents into performing the most efficient action. Therefore, i) our mechanism allows agents to reach the best convention in dynamic populations; and ii) when the best social convention has emerged, it empowers incoming agents with pre-established social rules to *adapt* its rules to the best social convention.

In summary, we claim that our infection-based mechanism presents self-adaptation properties, since it allows agents (using only local information) to dynamically change their rules in response to dynamic changes in the system.

Acknowledgments

The authors thank CONACyT for the scholarship of the first author. This work was funded by IEA (TIN2006-15662-C02-01), OK (IST-4-027253-STP), eREP(EC-FP6-CIT5-28575) and AT (CONSOLIDER CSD2007-0022, INGENIO 2010).

References

1. R. Albert and A.-L. Barabasi. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74:47, 2002.
2. R. Axelrod. An evolutionary approach to norms. *The American Political Science Review*, 80(4):1095–1111, 1986.
3. R. Burt. Social contagion and innovation: Cohesion versus structural equivalence. *American J. of Sociology*, 92:1287–1335, 1987.
4. R. Conte and M. Paolucci. Intelligent social learning. *Artificial Society and Social Simulation*, 4(1):1–23, 2001.
5. Z. Dezso and A.-L. Barabási. Halting viruses in scale-free networks. *Physical Review E (Statistical, Nonlinear, and Soft Matter Physics)*, 65(5):055103, 2002.
6. I.-K. Jeong and J.-J. Lee. Evolving multi-agents using a self-organizing genetic algorithm. *Applied Mathematics and Computation*, 88:293–303, 1997.
7. D. Moriarty, A. Schultz, and J. Grefenstette. Evolutionary algorithms for reinforcement learning. *Artificial Intelligence Research*, 11(1-1):241–276, 1999.
8. R. Pastor-Satorras and A. Vespignani. Epidemic dynamics and endemic states in complex networks. *Physical Review E*, 63:066117, 2001.
9. J. Pujol, J. Delgado, R. Sangüesa, and A. Flache. The role of clustering on the emergence of efficient social conventions. In *IJCAI 2005*, pages 965–970, 2005.
10. N. Salazar, J. A. Rodriguez-Aguilar, and J. L. Arcos. Infection-based self-configuration in agent societies. In *GECCO 08*, 1945–1952, USA, 2008. ACM.
11. N. Salazar, J. A. Rodriguez-Aguilar, and J. L. Arcos. An Infection-Based Mechanism for Self-Adaptation in Multi-Agent Complex Networks. In *SASO 08*, 161–170, Venice, Italy, 2008.
12. Y. Shoham and M. Tennenholtz. On the emergence of social conventions: Modeling, analysis, and simulations. *Artificial Intelligence*, 94(1-2):139–166, 1997.
13. A. Walker and M. Wooldridge. Understanding the emergence of conventions in multi- agent systems. In *ICMAS95*, pages 384–389, 1995.
14. D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, June 1998.