

# Robust Convention Emergence in Social Networks through Self-Reinforcing Structures Dissolution

DANIEL VILLATORO and JORDI SABATER-MIR, IIIA - CSIC  
SANDIP SEN, University of Tulsa

Convention emergence solves the problem of choosing, in a decentralized way and among all equally beneficial conventions, the same convention for the entire population in the system for their own benefit. Our previous work has shown that reaching 100% agreement is not as straightforward as assumed by previous researchers, that, in order to save computational resources fixed the convergence rate to 90% (measuring the time it takes for 90% of the population to coordinate on the same action). In this article we present the notion of social instruments as a set of mechanisms that facilitate and accelerate the emergence of norms from repeated interactions between members of a society, only accessing local and public information and thus ensuring agents' privacy and anonymity. Specifically, we focus on two social instruments: rewiring and observation. Our main goal is to provide agents with tools that allow them to leverage their social network of interactions while effectively addressing coordination and learning problems, paying special attention to dissolving meta-stable subconventions.

The first experimental results show that even with the usage of the proposed instruments, convergence is not accelerated or even obtained in irregular networks. This result leads us to perform an exhaustive analysis of irregular networks discovering what we have defined as *Self-Reinforcing Structures* (SRS). The SRS are topological configurations of nodes that promote the establishment and persistence of subconventions by producing a continuous reinforcing effect on the frontier agents. Finally, we propose a more sophisticated composed social instrument (observation + rewiring) for robust resolution of *subconventions*, which works by the dissolution of the stable frontiers caused by the Self-Reinforcing Substructures (SRS) within the social network.

Categories and Subject Descriptors: I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence

General Terms: Experimentation, Algorithms, Performance

Additional Key Words and Phrases: Norms, conventions, emergence, MAS, adaptation

## ACM Reference Format:

Villatoro, D., Sabater-Mir, J., and Sen, S. 2013. Robust convention emergence in social networks through self-reinforcing structures dissolution. *ACM Trans. Autonom. Adapt. Syst.* 8, 1, Article 2 (April 2013), 21 pages.

DOI: <http://dx.doi.org/10.1145/2451248.2451250>

---

This work was supported by the European Community under the FP6 programme (eRep project CIT5-028575), by the Spanish Education and Science Ministry (AEI project TIN2006-15662-C02-01, AT project CONSOLIDER CSD2007-0022, INGENIO 2010), by Proyecto Intramural de Frontera MacNorms (PIFCOO-08-00017), and by the Generalitat de Catalunya (2009-SGR-1434).

Authors' addresses: D. Villatoro (corresponding author) and J. Sabater-Mir, Artificial Intelligence Research Institute (IIIA), Spanish Scientific Research Council (CSIC), Bellaterra, Barcelona, Spain; email: [dvillatoro@iiia.csic.es](mailto:dvillatoro@iiia.csic.es); S. Sen, Department of Mathematical and Computer Sciences, University of Tulsa, Tulsa, OK 74104.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).

© 2013 ACM 1556-4665/2013/04-ART2 \$15.00

DOI: <http://dx.doi.org/10.1145/2451248.2451250>

## 1. INTRODUCTION

Conventions are a special type of norms, used as a mechanism for sustaining social order, increasing the predictability of behavior in the society, and specifying the details of these unwritten laws. Following Coleman's theory [Coleman 1998], conventions emerge to solve coordination problems, where there is no conflict between the individual and the collective interests, as what is desired is that everyone behaves in the same way, without any major difference on which action agents are coordinated. Therefore, the selection of the focal action in such norms is arbitrary.

In the multiagent systems community, the problem of convention emergence has been commonly approached through what has been defined as *social learning* [Mukherjee et al. 2007; Sen and Airiau 2007]. This approach has resulted in special interest because of the decentralized nature of the self-organization task. In *social learning* of norms, where each agent is learning concurrently over repeated interactions with randomly selected neighbors in the social network, a key factor influencing success of an individual is how it learns from the "appropriate" agents in their social network. Therefore, agents can develop subconventions depending on their position on the topology of interaction. As identified by several authors [Salazar-Ramirez et al. 2008; Toivonen et al. 2009; Villatoro et al. 2009], metastable subconventions interfere with the speed of the emergence of more general conventions. The problem of subconventions is a critical bottleneck that can derail emergence of conventions in agent societies, and mechanisms need to be developed that can alleviate this problem. Subconventions are conventions adopted by a subset of agents in a social network who have converged to a different convention than the majority of the population. Subconventions are facilitated by the topological configuration of the environment (isolated areas of the graph which promote endogamy) or by the agent reward function (concordance with previous history, promoting cultural maintenance). Assuming that agents cannot modify their own reward functions, the problem of subconventions has to be solved through the topological reconfiguration of the environment. In order to produce the lowest impact on the original structure of the social network, the subconventions need to be dissolved with the help of agents in the conflicting area where two different conventions confront, that we define as frontier.

Being part of a social network, agents can exercise certain control over it so as to improve one's own utility or social status. We define *social instruments* to be a set of tools available to agents to be used within a society to influence, directly or indirectly, the behavior of its members by exploiting the structure of the social network. Social instruments are used independently (an agent does not need any other agent to use a social instrument) and have an aggregated global effect (the more agents use the social instrument, the stronger the effect). In this work we study the effects of social instruments on the emergence of conventions through the dissolution of subconventions. Previous research has provided results proving that emergence of conventions based solely on private interactions is affected by factors like topology [Delgado 2002] or learning algorithms [Sen and Airiau 2007].

The main contributions of this work are the following: (1) establishment of a brief classification and analysis of the social instruments present in the existing literature, (2) the application of social instruments to facilitate the convention emergence process, (3) the identification of *Self-Reinforcing Substructures* (SRS) in social networks, that cause the metastable subconventions, and (4) the dissolution of the SRS with the help of the proposed social instruments. More specifically, we study two social instruments in this work: rewiring and observation. *Rewiring* allows agents to control the links that relate them with other agents by replacing them intelligently. This direct control of the topology of the social network allows agents to control whom they interact with,

resulting in increased reward without actually altering the reward function. On the other hand, *observation* allows agents to obtain partial information of the convention emergence process by observing other agents in the neighborhood. The access to this information allows agents to consider extra information over what they receive from direct interactions. This observation process also has an impact on the reward of the agents by speeding up their convention emergence within the society. After a first experimental phase, we discover the pertinence of the frontier in the subconventions problem. The frontier is the topological area of the network at the edge of settled groups (with the same shared convention within the group, and different outside the group).

Later we also propose a combination of the two presented social instruments to develop a robust solution to the subconventions formation problem. We present experimental results and their analysis about the efficacy of these social instruments for convention emergence by expediently dissolving the SRSs.

The structure of this article is as follows: in Section 2 we present the motivation of this work and a more detailed explanation of our vision of social weapons; we then present a small categorization of social weapons from a review of the literature of multiagent systems in Section 3, we then propose two new social weapons in Section 4. In Section 5 we present the agent interaction and reward model that we have used; experimental results are presented in Section 6. We perform an exhaustive analysis of irregular networks in Section 7, and after identifying the SRS, the combined social instrument is presented and experimented with to observe its efficacy in Section 8. Finally, conclusions from the analysis of the results and future work are presented in Section 9.

## 2. MANIPULATING THE NETWORK

Normally, in *social learning* scenarios the reward function<sup>1</sup> for an agent is determined either by the environment (e.g., if one overexploits resources, they will be exhausted) or by individual social interactions (e.g., drive on the same side of the road as the rest of the drivers for one's own and others' safety) and typically individual agents cannot modify it. However, in social learning interactions, the reward is determined based on the actions of the partners involved in the interaction. Therefore, an agent can indirectly modify the reward obtained and change its success and visibility and hence social status by controlling or modifying the social network it is located in, for example, selectively choosing which other agents it interacts with, giving different importance to the relationships with other agents, sharing information about the interactions with other agents, etc.

Social instruments are used by agents to improve their utility, without directly altering the utility function. Given a set of social instruments, agents can exploit the state of their social environment to have tangible effects on the utility obtained from their interactions with other agents in that environment.

A society where individual agents repeatedly use a social instrument has the properties of a complex system: the effect of the individual elements aggregated together exhibits, as a whole, properties that are not present in the properties of the individual parts. Moreover, the emergent aggregate behavior may change over time. The impact of the social instruments is reflected on its aggregated effect which often transforms the social environment.

---

<sup>1</sup>We use the definition of reward function typically used in the Markov decision processes framework, where it is understood as a private function that, given a state and a set of actions chosen by the players, and maybe their past history, returns to the agent a scalar reward value.

For inspiration, we can look to human societies and observe how humans have developed different social instruments to exploit and produce changes in their social networks. These human social mechanisms have direct or side-effects on the social networks. There exists some literature on classifying and analyzing the effects of these social instruments on the network and therefore on the agents' interactions [Albert et al. 2000; de Pinninck et al. 2007; Savarimuthu et al. 2007; Urbano et al. 2009; Villatoro et al. 2009; Babaoglu and Jelasity 2008; Conte and Paolucci 2002; Hales 2002; Hales and Arteconi 2006].

However, there is a critical dearth of work in this area of socially inspired tools for MultiAgent Systems (MAS). Our goal is to fill this gap by formalizing and adapting the social mechanisms used by humans for agent societies. As a social network can be represented as a type of graph, we profit from the good ground formalism provided by graph theory and we establish a parallel between network operators and social operators. However, from a societal point of view, only a subset of all the possible graph operations make sense.

### 3. SOCIAL INSTRUMENTS REVIEW

In a social network, agents are represented by nodes and relationships by edges. Assuming agents cannot create or destroy other agents, the set of graph operations that are socially relevant are the following: destroy a link, create a link, add attributes to nodes (node coloring<sup>2</sup>), add attributes to links (link coloring<sup>3</sup>), and transmit information (any node can transmit information to any other node following the right protocols using the communication network, which might be different from the social network).

The social instruments available to the agents should be able to perform these operations. To alter the state of the social network in which agents are located, researchers have developed different social mechanisms that we now categorize. We identify some social instruments found in the MAS literature and observe their effect on the social network. The mechanisms we are interested in are those that can be used by the agents without institutional support.

*Partner Selection (Edge Addition/Removal).* People continually select who they do and do not want to interact with. Giving agents the capability of removing edges attached to them in their social network gives them direct control over the network that directly affects them, improving the overall behavior of the system [Albert et al. 2000; de Pinninck et al. 2007].

*Tags (or Node Coloring).* Recognizing certain attributes and characteristics of other agents before interacting with them reduces social friction (e.g., by reducing the number of unsuccessful interactions) and improves coordination. To facilitate this process, some kind of externally visible social markings are needed, for example, tag mechanisms [Chao et al. 2008; Holland 1993].

By carrying a tag, having a role [Savarimuthu et al. 2007] or a certain *force* [Urbano et al. 2009], an agent allows all other agents to recognize a certain characteristic or quality even before any direct interaction.

*Social Position (or Link Coloring).* Humans recognize that our behaviour has to be adapted depending on who we are interacting with, for example, I behave differently with a friend from school than with my boss. Agents can extract information from the topology of a social network to gauge the influence or status of other agents

<sup>2</sup>Node coloring is a special case of node labeling in order to recognize certain characteristics of a node represented by its color.

<sup>3</sup>Link coloring is a special case of link labeling in order to recognize certain characteristics of the relationship between two nodes represented by its color.

and use this information strategically to further their own interests [Villatoro et al. 2009].

*Mimicking (or Node Imitation)*. Humans consciously or subconsciously imitate the strategies of more successful individuals as they aspire to improve their own performance. Access to the strategies of other nodes in the network can be used as a social instrument [Hales and Arteconi 2006]. Such access can be applied to compare the efficiencies of several strategies and mimic the most effective ones.

#### 4. PROPOSED SOCIAL INSTRUMENTS

In the last section, we have analyzed several social instruments that have been studied in the literature. In this section we propose two new additional social instruments: rewiring and observation. As we will see in the next sections, rewiring directly changes the structure of the network and observation exploits the information of certain parts of the network to bias agents' decisions. For the sake of clarity, we will clarify certain terms that might seem similar and are used in the rest of this article: agents choose an action at each timestep and this action, except when exploring, expresses the agents' preference over a finite set of actions. In this article, as we deal with learning algorithms, the preference represents the action which has the highest value estimated by the learning algorithm.

##### 4.1. Rewiring (or Intelligent Link Removal and Creation)

If we use human societies as inspiration, agents should have the ability of choosing whom they want to interact with. We have adapted this idea into a social instrument that allows agents to "break" the relationships from which they are not receiving any benefit and try to create new ones. This social instrument allows an agent to remove links with other agents it is connected with and intelligently substitute those links by new ones, making this last step the crucial difference with *edge removal*. Agents decide to rewire a link after the number of unsuccessful interactions<sup>4</sup> with another agent crosses a certain *tolerance* threshold. Agents also need to decide whom they want to establish the new link with<sup>5</sup>. In this research we are especially interested in developing methods that access only local and public information, ensuring agents privacy and anonymity. We have developed an intelligent rewiring scheme, (*neighbor's advice*), which will be compared with two other control methods: the zero-intelligence method (*random rewiring*) and the full-information method (*global advice*).

- (1) *Random Rewiring*. Agents rewire to a randomly selected agent from the population.
- (2) *Global Advice*. Agents rewire to an agent that is randomly selected by the system from those that have selected the same action in the last timestep. If no such agent exists, random rewiring is applied.
- (3) *Neighbor's Advice*. Agents rewire to an agent recommended by a neighbor. The rewiring agent X asks a neighbor which in the last timestep selected the same action as that of X (agent referent) for another agent which also selected the same action but is not a neighbor of X (agent referred). If these conditions do not apply, random rewiring is applied.

Despite the similarity of our approach with the work of Griffiths and Luck [2010] in that both use an evolutionary approach, there is a crucial difference. Griffiths and Luck [2010] measure the results of their techniques after the reproduction of a number

<sup>4</sup>Unsuccessful interaction in our convention emergence scenario corresponds to being uncoordinated or not sharing the same convention for that interaction.

<sup>5</sup>As in this work we only consider undirected networks, when one agent decides to break or establish one link, the other agent will also be affected by this decision.

of generations, and using a certain mutation rate. On the other hand, we use a more *online approach* where agents can modify their social network at runtime, without the necessity of evolving new generations. In addition, our rewiring methods do not access any private agent information (used only in the global advice which is used as a control case) such as their actual reward. Our rewiring methods only access the last played action of other players, respecting the privacy of agents' information. There is also a clear difference between this social weapon and the *partner selection* mechanism: rewiring produces a tangible change in the network, while partner selection on the other hand does not directly affect the network but prioritizes some nodes over the rest.

#### 4.2. Observation

In a social learning scenario, allowing agents to observe the actions selected by other agents outside their circle of interaction can provide useful information to support the convention emergence process. However, there has to be a trade-off between observing and interacting. Allowing agents to observe other agents' state might help when deciding which action should be reinforced in its learning algorithm in order to become the established convention. In order to analyze the effects of observation, we will allow agents to observe, at certain timesteps, a subset of other agents' states in the population. Therefore, agents will be assigned an *observation probability* that will determine the probability to use the observation instrument each timestep. Moreover, agents need to know the maximum number of agents they can observe (*observation limit*<sup>6</sup>) and how they want to observe (*observation method*). As we previously did for rewiring, we propose three different observation methods two intelligent methods (*local observation* and *random focal observation*) contrasted against the zero-intelligence method (*random observation*).

- (1) *Random Observation*. Agents observe random agents from the society.
- (2) *Local Observation*. Agents observe their immediate neighbors in the social network.
- (3) *Random Focal Observation*. Agents select one random agent from the society and observe that agent and its direct neighbors.

After the observation process, the agent will choose the majority action taken by the selected observed agents and will send a strong reinforcement signal to its learning algorithm for that action. We are aware that in certain configurations, choosing the observed "majority action" might not help agents achieve the general convention. Nonetheless, and because of the intrinsic nature of the convention emergence process, the observed "majority action" will lead in most of the situations to the general majority action.

Despite the similarity, this instrument and *mimicking* (presented in the previous section) behave differently: mimicking agents copy from a successful agent the strategy (which in most of the cases is private) as well as the interaction neighbors (which might also be private). On the contrary, with the observation instrument agents observe the last decision (hoping this represents the actual strategy) of a group of agents (without knowing the actual success of the agents) and update their estimates based on the majority. With observation, agents only access information that has been previously made public by the observed agent, while with mimicking, they access information that can be considered private.

<sup>6</sup>If an agent has more neighbors than specified by the observation limit, it will randomly select the additional agents to observe. If it has less neighbors than the observation limit, it will observe all its neighbors.

## 5. MODEL OF REWARD

The social learning situation for norm emergence that we are interested in is that of learning to reach a social convention. We borrow the definition of a social convention from Shoham and Tennenholtz [1997]: *A social law (a restriction on the set of actions available to agents) that restricts agents' behavior to one particular action is called a social convention.* For this reason, in our social learning scenario norms are implicit. Agents do not have any internal representation of norms, only preference for one action (the one specified by the norm) over the others.

To test the efficiency of the proposed social instruments, we test them under the most accepted convention emergence model (used by Delgado et al. [2003], Kittock [1993], Mukherjee et al. [2007], Sen and Airiau [2007], Shoham and Tennenholtz [1997], and Walker and Wooldridge [1995]): agents converge to a convention through repeated bilateral interactions with other agents from their social neighborhood. Any interaction between two agents is represented as a 2-person  $n$ -action game. At each timestep, each agent is paired with another agent and they independently decide about their actions. This decision is made without observing the paired agent's identity or strategy, preserving all agents' anonymity and privacy.

In our approach a social convention will be reached the exact timestep *all* the agents are in the same state or consistently choose the same action (the actual state reached or action chosen is immaterial). Other researchers have used a convergence threshold of 90% [Delgado et al. 2003; Kittock 1993]. However, we have observed that with certain reward functions and on certain topologies, after 90% of the society has converged to a convention, the convergence dynamics can be very different than those produced until the 90% convergence rate. Hence, a threshold of 90% can produce misleading or inaccurate results.

An agent state is determined by the action the agent is choosing: when two agents are randomly paired to interact, each one will choose an action (or state) for this interaction, and will receive a reward from the environment according to the game they are playing (in this case, the coordination game used in Sen and Airiau [2007]). A convention is reached if all agents present in the population choose the same action (or state) in any interaction.

As in several other research in convention emergence (see Delgado et al. [2003] and Kittock [1993]), the interactions between agents in our framework are constrained by one of two different underlying structures: (i) a *one-dimensional s-lattice* with connections between all  $s$  neighboring vertex pairs and (ii) a *scale-free network*<sup>7</sup>, whose node degree distribution asymptotically follows a power law. Each agent is represented by a node in the network and the links represent the possibility of interaction between nodes (or agents). All the links in the network are undirected and unweighted.

As in Kittock [1993], Shoham and Tennenholtz [1997], and Villatoro et al. [2009], agents are endowed with a limited memory of past interactions (same memory size for all agents). Agents save in their memory *when* an interaction occurred, the *action* chosen, and the *reward* obtained. As we will see shortly, the information stored in memory is used differently depending on the type of decision-making procedure adopted.

Agents cannot observe the other agent's memory, current decision, or immediate reward, and hence cannot calculate the payoff for any action before actually interacting with the opponent.

Once the model of interaction is fixed, the decision-making procedure needs to be defined. We test our social instruments using three well-known decision-making functions.

<sup>7</sup>The method used in this work for the generation of the scale-free networks is the classic *Barabási Albert preferential attachment* [Barabási and Albert 1999] (with  $m_0 = 2$ ).

		Player 2	
		Action A	Action B
Player 1	Action A	+1	-1
	Action B	-1	+1

Fig. 1. Coordination game.

- (1) *Best Response Rule (BRR)* [Mukherjee et al. 2007; Sen and Airiau 2007]. Each agent chooses the action with which it has obtained the highest payoff in the last iteration. A positive reward is given to agents if they are coordinated. This function gives preference to the myopic utility-maximizing action.
- (2) *Highest Cumulative Reward Rule (HCRR)* [Shoham and Tennenholtz 1997; Kittock 1993]. “[A]n agent switches to a new action iff the total payoff obtained from that action in the latest  $m$  iterations is greater than the payoff obtained from the currently chosen action in the same time period.” A reward is generated based on a coordination game. This function gives preference to the action that has obtained the highest accumulated payoff in the last  $m$  interactions (rather than only in the last one).
- (3) *Memory-Based Rule (MBR)* [Villatoro et al. 2009]. A positive reward is given to agents if they are coordinated, and it is proportional to the actions they have chosen in the past. This function gives preference to the action that has provided the largest payoff while taking into consideration in the reward function also the previous actions (promoting concordance with previous history).

The system’s payoff that each agent receives with the first two functions is shown in Figure 1, which shows a classical *coordination game*. Then, the internal payoffs that each agent uses is different depending on the function: for BRR the payoffs received are the same used in the internal reasoning of the agents; for the HCRR, an accumulative value of the rewards obtained in the past with the coordinated action is used; and finally, for the MBR a proportional value of the past actions equal to the coordinated action taken by the agent. The HCRR also specifies the action that each agent has to take in each interaction. On the other hand, for BRR and MBR, agents use a learning algorithm to estimate the worth of each action. Agents choose their action in each interaction in a semideterministic fashion: a certain percentage of the decisions will be chosen randomly, representing the exploration of the agent, and for the rest of the decisions, agents deterministically choose the action estimated to be of the highest utility. In all the experiments presented in this article, and in order to follow the methodology established in previous work of social learning [Mukherjee et al. 2007; Sen and Airiau 2007; Villatoro et al. 2009], the exploration rate has been fixed at 25%, that is, one-fourth of the actions are chosen randomly.

The learning algorithm used here is a simplified version of the Q-learning algorithm [Watkins and Dayan 1992]. The Q-update function for estimating the utility of an action  $a$  is

$$Q^t(a) \leftarrow (1 - \alpha) \times Q^{t-1}(a) + \alpha \times \text{reward}, \quad (1)$$

where  $\text{reward}$  is the payoff for the current interaction and  $Q^t(a)$  is the utility estimate of action  $a$  after selecting it  $t$  times. When agents decide not to explore, they will choose the action with the highest Q value.

The simulation process for repeated interactions in the agent society is presented in Algorithm 1.



**ALGORITHM 1:** Simulation Process

---

```

for timesteps do
  forall agents do
    Select another partner agent from population;
    Each selected agent chooses an action;
    The joint action from the selected agents and their history determines rewards;
    Selected agent(s) use received reward to update action estimates;
    Apply Social Instrument;
  end
end

```

---

**6. EXPERIMENTS**

We use the simulation model presented in the previous section to evaluate how the social instruments that we propose influence the process of emergence of social conventions.

Some of the simulation parameters are common to all the experiments:

- population size = 100;
- agents located in a social network with different topologies: s-lattices (varying the value of  $s$  attending to the neighborhood size  $NS = 2 \times s$ ) and scale-free network;
- memory size 5 (for HCRR and MBR).<sup>8</sup>

We experiment with different sets of available actions (2, 5, and 10 actions) for agents to choose from. Based on the underlying assumptions of the convention emergence problem, all actions (or states) are potentially of the same value, as long as all agents converge to the same action (i.e., given the set of actions =  $\{A, B\}$ , converging to action  $A$  produces the same utility to agents as converging to action  $B$ ). Agents are initialized with no preference between the available actions and randomly choose actions with equal probabilities. One specific metric that will be used for the valuation of the experiments is the *convergence time* that is defined as the number of timesteps that it takes for the whole population to agree on a common convention. All the results presented are the average over 25 runs of the simulation.

**6.1. Effects of Rewiring**

We have experimented with the three rewiring methods introduced in Section 4.1 on three topologies: a low-clustered<sup>9</sup> one-dimensional lattice (lattice with neighborhood size = 10), a high-clustered one-dimensional lattice (lattice with neighborhood size = 30), and a scale-free network. We have explored the search space of the tolerance levels, observing how they affect the convergence time and the number of components (maximally connected subgraphs) created when convergence is reached with the different decision-making functions.

*6.1.1. Influence of Rewiring Methods.* From the experimental data we have observed how, in general, the *Global Advice (GA)* rewiring method produces the best convergence time due to its centralized nature and access to global information. Nonetheless the decentralized methods, especially the *Neighbor's Advice (NA)* method, also demonstrate good performance. The *NA* method improves the *Random Rewiring (RR)* method as it more expediently resolves the subconventions that appear in the one-dimensional lattices during the convention emergence process. These metastable subconventions

<sup>8</sup>These values have chosen in order to compare our results with those from the literature [Shoham and Tennenholtz 1997; Villatoro et al. 2009] where they conventionally use a memory size of 5.

<sup>9</sup>Clustering coefficient is a measure of degree to which nodes in a graph tend to cluster together.




Fig. 2. Component sizes histogram for scale free.

have been already identified in Salazar-Ramirez et al. [2008], Toivonen et al. [2009], and Villatoro et al. [2009] to be the preferred action of a group of nodes who have converged to a different convention than the rest. These subconventions, as noted in Villatoro et al. [2009], do persist as long as the frontier region remains metastable<sup>10</sup>. The frontier region is the group of nodes in the subconvention that directly interacts with other nodes that have adopted a different convention.

The explanation of why, when using the *neighbor's advice* method, these subconventions are resolved more expediently is the following: Agents in the frontier use the rewiring instrument as they cross the tolerance level faster than those not in the frontier. For this reason, the *RR* method will relink an agent with a more suitable agent with a probability of  $\frac{1}{\text{Number Of Actions}}$  (as this is the probability that a random agent plays the same action). In contrast, the *NA* method will relink the agent with another one playing the same action if it is accessible. In case there is no other agent which fulfills the necessary conditions to connect with, random rewiring will be applied, obtaining in the worst-case scenario the same results.

These results are reaffirmed for the scale-free networks, although an interesting phenomenon concerning the final number of components is observed in this case. Figure 2 shows the size of the components<sup>11</sup> when a convention was reached. The horizontal axis represents the tolerance levels and the vertical axis presents ranges (of size 10) of component sizes, and the histogram darkness corresponds to the amount of components. When using the *RR* (Figure 2(a)) and the *GA* (Figure 2(c)) methods, the number of components obtained is always above one, resulting in a fragmented society. The number of components is slightly higher for the *NA* method than when using the other methods. This phenomenon is observed because of the structure of the scale-free networks and the dynamics of the rewiring methods. Nonhub nodes reach their tolerance levels faster than hub nodes (as nonhub nodes interact with a lower number of agents). Once the tolerance level is reached, the rewiring process kicks in. We do not observe any distinctive behavior separating the *RR* or the *GA* methods as they both incorporate a random component. However, when using the *NA* scheme, the number of components is variable, depending on the tolerance level: the number of components is higher with low tolerance, although we have observed that most of these components are small in size (as shown in Figure 2(b)).

Nonhubs (when using the *NA* method) will rewire to any of their leaf nodes (nodes with degree 1) separating completely from the initial complete component (e.g., in Figure 3

<sup>10</sup>Metastable conventions are those that should preferably be broken, but remain stable because of a combination of factors like interaction dynamics, topology, and decision-making functions.

<sup>11</sup>A component of an undirected graph is defined as a subgraph in which any two vertices are connected to each other by some path, and no node in this subgraph is connected to additional nodes.




Fig. 3. Example of components evolution in scale-free NA rewiring method.

we can observe how one agent has applied rewiring, breaking the link that linked itself with the main component, and rewiring to another agent forming a new component of the network), creating a greater number of components, but of smaller sizes.

*6.1.2. Influence of Topology.* When observing the effects of the topology, we find that the convergence time is increased under the effects of rewiring when the neighborhood size is increased in regular networks. This effect is due to the average degree of the nodes in the network (intuitively proportional to the clustering coefficient)<sup>12</sup>. The one-dimensional lattices with higher neighborhood sizes are more fragmented than those with more restricted neighborhoods. Therefore, when increasing the neighborhood size, the number of links between agents also increase (thereby increasing the clustering coefficient). High-degree networks in highly clustered societies are more resistant to rewiring, as the node that wants to use the rewiring would have to apply it to a higher number of nodes, and then be rewired to the same amount of nodes with the appropriate strategy, induced by the last played action. This scenario highlights the difference between our social instrument and that presented in Hales and Arteconi [2006] and Griffiths and Luck [2010]: our instrument needs to substitute the selected links one by one; on the other hand, the instrument in Hales and Arteconi [2006] and Griffiths and Luck [2010] accesses the private information of another agent and copies their entire set of neighbors in one shot. Our social instrument sacrifices convergence time to preserve privacy.

Related to these results, we have also observed that under the effects of rewiring there exists a trade-off between the diameter of the regular network and the number of components: with higher diameters (smaller neighborhood size) the number of components remains constant and is greater than one (see Figure 4(a)). However, we can observe in Figure 4(b) that when the diameter is smaller (with higher neighborhood size), the number of components is reduced to one for certain tolerance values. The explanation of this phenomenon is again related to the clustering coefficient of the network combined with the dynamics of the rewiring process. When the clustering level of the network is low, agents have less links to be rewired. Initially, and before the tolerance level is reached, agents also change their preferences without rewiring. Eventually, and after the tolerance level is reached, agents start rewiring their links (and reducing their number of preference changes). As explained previously, low-clustered societies have higher tendency to become disconnected and form different components.

Experimental results (shown in Figure 5) also show interesting properties with scale-free networks (see Section 6.1.1). Similar to what was observed by Villatoro et al. [2009] we found that subconventions have a stronger effect on scale-free networks when using MBR, producing, in general, a larger number of components (see Figure 4(c)). However,

<sup>12</sup>The clustering coefficient quantifies how well connected are the neighbors of a vertex in a graph. We borrow the definition introduced in Watts and Strogatz [1998], where “the local clustering coefficient of a node is given by the proportion of links between the vertices within its neighborhood divided by the number of links that could possibly exist between them.”

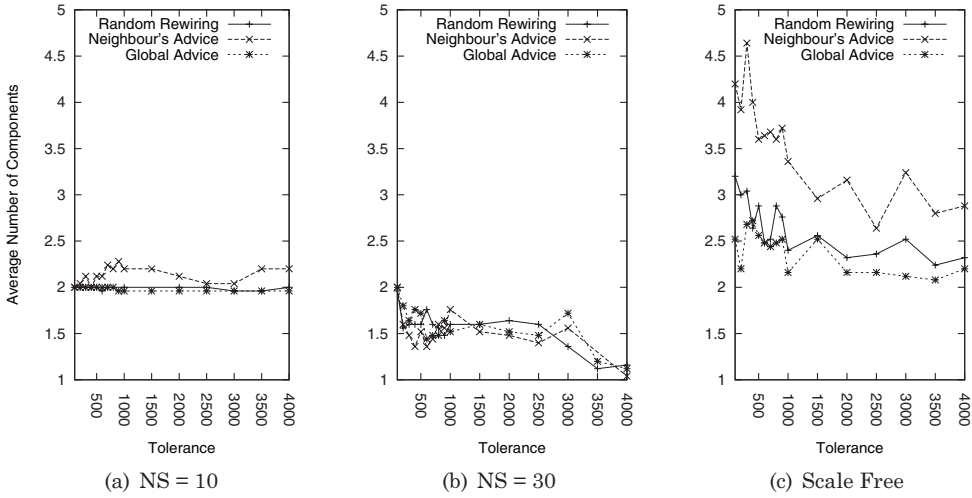


Fig. 4. Rewiring methods comparison for number of components in a one-dimensional lattice (a) and (b) and scale free (c).

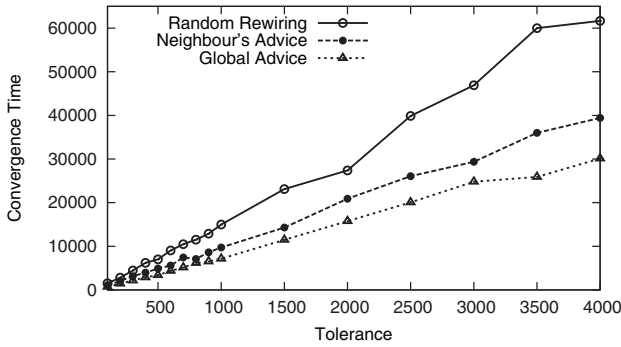


Fig. 5. Convergence times with different rewiring methods in a scale-free network.

global convergence is obtained in this type of networks when using the rewiring social instrument.

We can conclude that rewiring performs better in low-clustered societies, producing a stratified population which results in a significant reduction in convergence time. In more clustered networks, the tolerance level has to be chosen carefully (depending on the other experimental parameters) to produce an effective technique for norm emergence.

## 6.2. Effects of Observation

In this section we analyze the effects of the *observation* social instrument when used by agents. We test and compare the three different methods proposed, exploring the search space with a representative range of observation probability values. As we did in the previous section, we experiment with three topologies: a low-clustered lattice, a high-clustered lattice, and a scale-free network. To determine the effects of the different developed methods, we fix the observation limit to 10 for the experiments.

**6.2.1. Influence of Observation Methods.** Comparing the results from the three observation methods we observe in Figure 6 that the Random (*RO*) and the Random Focal




Fig. 6. Comparison on effects of convergence time on one-dimensional lattice.

Observation (*RFO*) methods are the most effective ones (with different topological effects), and have very similar results when compared with the Local Observation (*LO*) method. The reason for this phenomenon can be attributed to the frontier effect. When agents use the *LO* method, they observe their direct neighbors. If the observing agent is in the frontier area, then this observation is pointless. However, observing different areas gives a better understanding of the state of the world, and hence the *RO* and the *RFO* methods perform better.

**6.2.2. Influence of Topology.** For the BRR and MBR decision-making functions, we have observed that the different observation methods produce a more pronounced effect in societies with higher diameters, as we can see in Figure 6 (only results for MBR are plotted, although similar results are observed for BRR). We notice that a small percentage of observation drastically reduces convergence times. The reason for this effect can again be found in the frontier and the subconvention effects previously discussed. Subconventions emerge more readily when the social network has a small diameter and the frontier region represents the unsettled area. These subconventions are more easily resolved at these frontiers by observation rather than by learning through interactions.

**6.2.3. Influence of Decision-Making Techniques.** The main experimental result obtained for BRR and MBR is that greater observation leads to faster convergence. However, in most cases, even smaller observation percentages help significantly reduce the convergence time. This result is important if observation has a concomitant cost. However, observation has a poor effect on HCRR: we find that the convergence time is proportional to the observation probability (the higher the observation, the worse it performs). The reason for this effect is found on the HCRR's lack of exploration policy. When observing, an action is reinforced, and this action might be different from the one that the agent has converged to within its environment. This new reinforcement destabilizes the convergence that was already obtained. This effect is not produced with the other reward rules as they have an exploration rate that quickly corrects this destabilization.

### 6.3. Discussion on the Frontier Effect

The experiments performed up to now have shown us how the emergence of social conventions can be facilitated with the usage of our proposed social instruments, achieving good results. However, some practical information has to be kept for the appropriate usage of them: On the one hand, tolerance levels have to be carefully chosen when using rewiring. On topologies with large diameters, the social network might suffer important changes affecting the number of components. Moreover, we have observed an interesting phenomenon on scale-free networks: without the usage of the social

instruments, full convergence is hard to obtain. With the use of rewiring instrument convergence is reached and a larger amount of components are produced when using the neighbor's advice method.

On the other hand, a small percentage of observation is very helpful in the achievement of full convergence. Interestingly, when using the local observance method, this improvement is not as prominent as with the other methods.

With these pieces of information we can hypothesize that there exists a special type of subconventions that cannot be resolved at a local level with the methods proposed, especially in scale-free networks. As our local information methods do not produce improvement in the process of convention emergence, agents in the frontier have to be provided additional mechanisms to dissolve the subconventions with local information.

## 7. UNDERSTANDING SUBCONVENTIONS

The results from experiments presented earlier, together with the observations narrated by other authors [Epstein 2000; Toivonen et al. 2009; Villatoro et al. 2009], convinced us that subconventions are problematic obstacles to the emergence of global conventions. These subconventions thrive (amongst other reasons) because of the topological structure of the network where they emerge.

Because of the inherent structure of the scale-free networks, interesting properties of the network are highlighted in the dynamics of the convention emergence process. To have a more detailed idea of how the topology affects the convergence emergence process, we have performed a detailed study of its dynamics.

### 7.1. Which is the Strongest Node?

In previous works in the literature [Sen and Airiau 2007; Mukherjee et al. 2008], some authors fixed the behavior of a number of agents in the population to observe how their immutable behavior would affect the emergence of a convention. These experiments were performed in regular networks and the position of the fixed learners on the network was not considered. Inspired by those works, we have decided to apply such technique in irregular networks, locating the fixed learners in selected positions of the network.

Scale-free networks are defined as irregular networks where very few nodes have a large number of connections (hubs), and a large number of nodes have a very few connections (leaves). To understand which nodes have a stronger effect in the emergence of conventions, we experiment by fixing a certain percentage of the population to a specific convention. Three experimental variants are tested.

*Fixing hub agents' behavior.* The behavior of the nodes with higher degree of connectivity are fixed.

*Fixing leaf agents' behavior.* The behavior of the nodes with lower degree of connectivity are fixed.

*Fixing random agents' behavior.* The behavior of randomly selected nodes are fixed.

We hypothesize that by fixing the behavior of hub agents, the rest of the population will converge faster than in any other possible situation. Our intuitions are founded on the dynamics of the convention emergence: given that a hub node will interact with more agents, its fixed behavior will remain unaffected by the others' decisions. However, as others are learners, they will learn from the decisions of the hubs.

Experimental results presented in Figure 7 partially confirm our hypotheses. We have performed extensive evaluation of the search space, exploring different percentages of fixed agents in different locations of the network and analyzing their

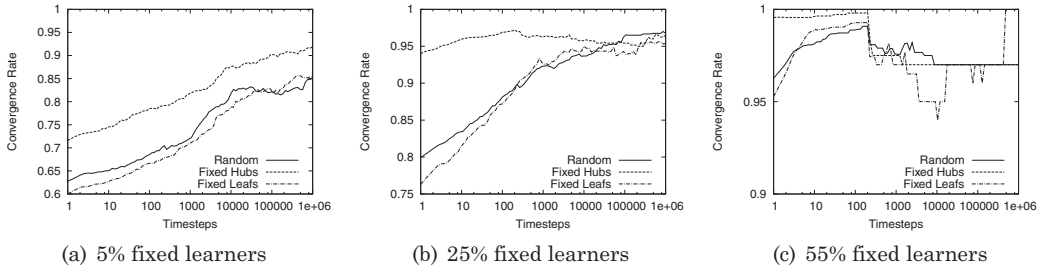


Fig. 7. Fixing agents' behavior.

convergence rate<sup>13</sup>. For small numbers of fixed agents, fixing the hubs produces better convergence rates than when fixing the same amount of leaves or random agents (see Figures 7(a) and 7(b)). On the other hand, when the number of fixed agents is larger, fixing the leaves results in faster convergence (see Figure 7(c)).

This empirical result leads us to pay attention to the leaves. Experimental data confirms that when the majority of the leaves have reached a common convention (artificially by fixing its behavior, or naturally through the convention process) reaching full convergence can be easily achieved. Consequently, to identify why learning leaves (by fixing the hubs' behavior) produces such delay in the convention emergence, we analyzed carefully the networks when the metastability was reached. By taking snapshots of the state of the convention emergence process (the structure and state of the network: which node is in which state and connected with which other nodes) at the timestep where no improvement in the convergence rate has been made, we identified *Self-Reinforcing Substructures (SRS)*. These substructures are a group of nodes that, given the appropriate configuration of agents' preferences and network topology, do maintain subconventions, as the agents in the frontier remain in a metastable non-convergent state because of the reinforcement received from the interactions with the agents in the SRS.

We have grouped the topology of these structures into two general categories.

- The *Claw* SRS is formed by connecting a node with a number of *hangers*<sup>14</sup> connected to it smaller than the number of links with the main component of the network. When the hangers coordinate to the same convention among themselves and with the connecting node, we have a self-reinforcing structure. For example, in Figure 8(a), A is the central node, having one connection with the rest of the network and 3 hangers: B (that it is another claw), C (plain hanger) and D (chain's connecting node).
- The *Caterpillar* SRS is a structure formed by a central path, whose extreme nodes are connected to the main network component and from its members other SRSs (such as claws, chains, or plain hangers) can hang. For example, in Figure 8(b), A, B, C, and D are members of the central path and the other nodes reinforce them.

These two abstract (examples in Figure 8) structures can be found as subnetworks of scale-free and random networks. As we have observed, the existence of these SRSs (74% of the generated networks with the methods described in Delgado et al. [2003] contain SRSs) are the main reason why convergence to a 90% level (as observed by Delgado

<sup>13</sup>Convergence rate is the proportion of the population sharing the majority action with respect to the whole population.

<sup>14</sup>A *hanger* is formed by nodes that are connected to a member of a cyclic component, but which do not themselves lie on a cycle [Scott 2000], and a *chain* is a walk in which all vertices and edges are distinct.




Fig. 8. Self-reinforcing structures.

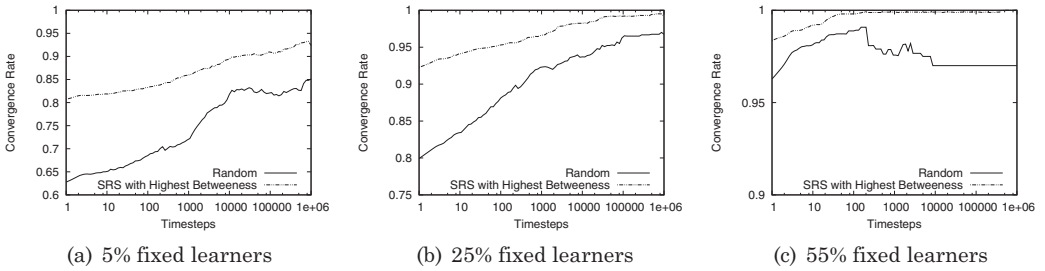


Fig. 9. Comparison with fixed SRS nodes.

et al. [2003]) is achieved relatively quickly, but overcoming the last 10% (containing the SRS) is significantly harder to achieve.

To check the validity of our hypotheses, we have extended the previous experiment, adding another variant where we fix the behavior of the nodes in the SRS with higher betweenness<sup>15</sup>.

Experimental results in Figure 9 show a comparison of the results of convergence rates of the best performing variants of those presented previously and the one that fixes the agents' behavior in the SRS. We can observe that fixing the SRS nodes outperforms any of the other strategies.

This result is very interesting as it helps us understand the specific substructures that generate subconventions within social networks. We are now aware that only by dissolving the subconventions in the frontier of a SRS can full convergence be ensured in scale-free networks. To more robustly tackle and eliminate the subconvention's effect, we propose the combination of both analyzed instruments in the next section.

## 8. COMBINING INSTRUMENTS: SOLVING THE FRONTIER EFFECT

After experimenting with both social instruments in Section 6, we observed that the subconventions need to be resolved in what we consider to be the "frontier" region of self-reinforcing structures. Otherwise, the agents within the SRS will continue reinforcing a frontier node with one specific convention and which will be at conflict with the convention of the rest of the network.

A subconvention in a regular network is not metastable (i.e. the subconvention is continuously affected, tending to disappear), but unfortunately, slows down the process

<sup>15</sup>Betweenness measures the extent to which a particular node lies "between" the various other nodes in the graph [Freeman 1979].



of emergence. On the other hand, in other network types, such as random or scale-free subconventions, they seem to reach metastable states<sup>16</sup> because of the existence of the self-reinforcing structures identified in Section 7.1. Therefore, by giving agents the tools to dissolve these frontiers, we hypothesize that convention emergence will be achieved faster and full convergence rates will be obtained.

However, because of the way social instruments have been designed, the instruments can be activated in situations where it is not necessary (observation follows a probabilistic approach, and rewiring is activated using a rewiring tolerance). Therefore, and by combining both social instruments developed in this work, we have designed a composed instrument for resolving subconventions in the frontier in an effective and robust manner. This composed instrument allows agents to “observe” when they are in a frontier, and then apply rewiring, with the intention of breaking subconventions. To effectively use this combined approach, agents must first recognize when they are located on a frontier.

We define a frontier as a group of nodes in the subconvention that are neighbors to other nodes with a different convention and at the same time are not in the frontier with any other group. To provide a more precise definition we first need to define our system. We have

$$\text{System}^t = \{A, \text{Rel}^t, S^t\}, \quad (2)$$

where  $A$  is a set of agents,  $\text{Rel}^t$  is a neighborhood function at time  $t$ , and  $S^t$  is the actual state of agents at time  $t$ .

We can therefore formalize the notion of subconvention  $\text{Sub}$ .

$$\text{Sub}^t \subset A \text{ where } \exists a \in \text{Sub}^t, \nexists b \in \text{Sub}^t \mid \text{Rel}^t(a, b) \wedge S_a^t \neq S_b^t$$

And now we can define when an agent is located in a frontier.

$$\begin{aligned} \text{Frontier}^t(a) = \{ & a, c \in \text{Sub}^t \wedge \exists b \notin \text{Sub}^t \wedge \text{Rel}^t(a, b) \wedge S_a^t \neq S_b^t \wedge \text{Rel}^t(a, c) \\ & \wedge (\forall d \in A \mid \text{Rel}^t(c, d) \rightarrow S_c^t == S_d^t) \} \end{aligned} \quad (3)$$

This formula basically means that an agent  $a$  is in the frontier when it is connected to another agent  $b$  not sharing the same convention, and at the same time, is being reinforced by other agents  $c$  that are not in the frontier.

As the cases are different for regular and irregular networks, two types of frontiers need to be defined:

- weak frontiers* as the ones that are not metastable in regular networks; and
- strong frontiers* as the ones generated by the SRSs in irregular networks.

The most important characteristic that defines a frontier is the existence of a confrontation. Confrontation occurs when two agents in an interaction do not share the same convention<sup>17</sup>.

Before proceeding further, we will define three characteristics of agents with respect to their convention and the topological position in the network. An agent is *in equilibrium* if it has the same number of neighbors in its own convention as in the other convention. An agent is a *weak* node if the number of neighbors in its own convention is lower than those in the other, and an agent is a *strong* node otherwise (if the number of neighbors in its own convention is greater than those in the other). In

<sup>16</sup>By experimentation, we have observed that around 99% of the generated scale-free networks do not achieve full convergence before one million timesteps with any of the decision-making functions used in this work and without any social instrument.

<sup>17</sup>Not sharing the same convention, choosing a different action, or choosing a different state to be, are considered equivalent expressions for our purpose.

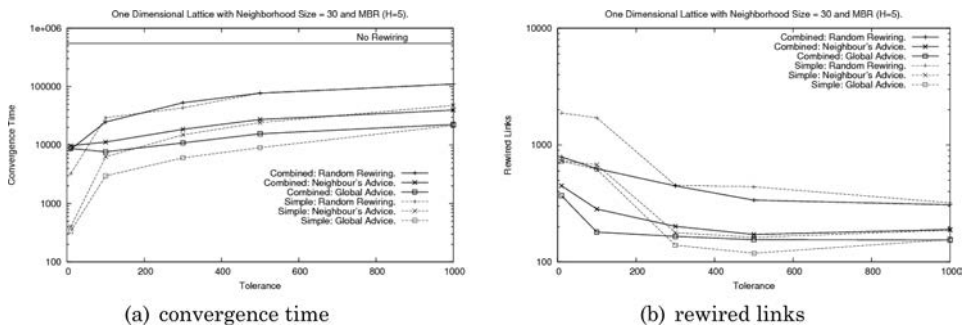


Fig. 10. Comparison with simple and combined social instruments on regular network using MBR.

regular networks, two confronted agents are in a frontier region iff: (1) At least one of the confronted agents is in an equilibrium position, and (2) all the neighbors of an inequilibrium confronted agent are strong nodes. In irregular networks, two confronted agents are in a frontier region iff both agents are strong nodes.

With these formalizations and identification of what the frontiers are in each of the topologies, agents can be equipped with this knowledge and the tools necessary to recognize when they are located in a frontier and repair that situation. To achieve this task we will use a combination the simple instruments presented in Section 4: first, agents will use observation (with the local observance method) to identify if they are located in a frontier (and therefore part of a self-reinforcing structure), and secondly, they will use rewiring to solve the frontier problem, by disconnecting from another agent in the SRS and reconnecting to another random agent.

### 8.1. Results

We have conducted exhaustive experimentation with the composed instrument on the three topologies and using the different decision-making functions described in the previous section. The use of the composed instrument on regular networks does not produce an improvement on convergence time with respect to simple rewiring (one example of topology and decision-making technique can be observed in Figure 10(a)). However, an important improvement is observed in the number of rewired links (one example of this improvement can be seen in Figure 10(b)). In general, this improvement is observed for lower tolerances. The reason for this effect is because higher-tolerance rewiring works in the same way as the composed social instrument, but without observing. For those smaller values the effect is significant, reducing the number of rewiring links down to half of the original value.

On the other hand, we observe an important improvement for convergence times when using the composed instrument (with the recognition of SRS) on irregular networks. The results presented in Figure 11 represent the average results from 25 different scale-free networks with and without using the combined social instrument. By comparing Figures 11(a) and 11(b) we notice the trade-off between the improvement in convergence times and the amount of rewiring to be done. The reason for this phenomena is because the composed social instrument decomposes the SRS differently than the simple rewiring, which only rewires the node in the actual frontier.

## 9. CONCLUSIONS AND FUTURE WORK

We have introduced the use of *social instruments* as tools that facilitate norm emergence. We have identified the characteristics and opportunities for effectively utilizing these social instruments for facilitating norm emergence through social learning. Social

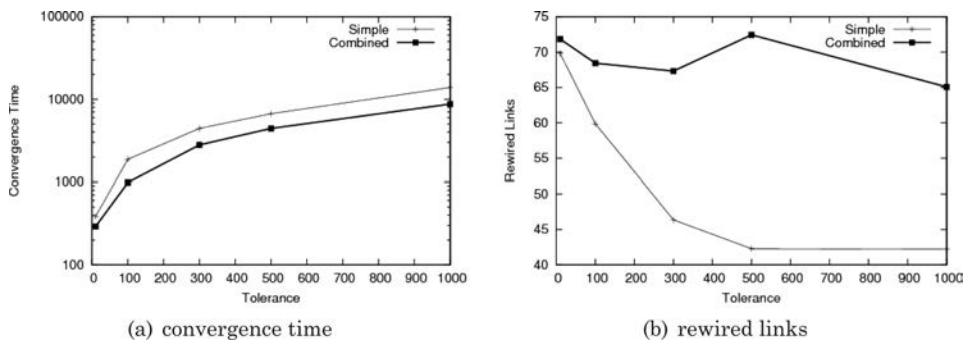


Fig. 11. Comparison with simple and combined social instruments on scale-free network using BRR.

instruments are attractive since they do not require centralized monitoring or enforcement mechanisms, are generally easy to use, have very low computational costs, ensure the privacy and anonymity of agents within the network (as our instruments only access public local information), and are scalable to large systems. We present an initial classification of related instruments in multiagent systems and networks.

Experimental results with the identified social instruments have shown that the emergence of transitory subconventions are the cause of the delay of the emergence of global conventions. From results presented in this article for the two simple social instruments studied, we observe that the most effective social instruments are those that more expediently solve this subconvention formation problem in the frontier regions. In the case of the rewiring social instrument, we have identified three methods, one that uses global information and two others using local information, that improve the performance of the system. We observed how the rewiring methods perform better in societies with a low clustering coefficient. In societies with a relatively large clustering coefficient, a large number of rewirings have to be performed, resulting in longer convergence times. A future line of research should carefully analyze the relationship between the clustering coefficient of the networks and the performance of the rewiring instrument, which at the moment remain described vaguely.

For the *observation* social instrument, we also analyzed three methods, and these produce improved convergence times. We emphasize that even a very small percentage of observation does drastically reduce convergence times. The first experimental results led us to perform a more exhaustive study to identify the self-reinforcing structures that are present in social networks with other topologies, for example, scale-free networks. Finally, we have presented a composed social instrument as a robust solution that overcomes persistence of subconventions generated by the identified SRS, improving the convergence times obtained with simple rewiring.

In a world where almost 950 million users belong to an online social networking platform (where virtual agents could also exist) [Radwanick 2010], it is important to understand what mechanisms this virtual entities should be equipped with to facilitate the emergence of common conventions (for the sake of the whole group) as quickly as possible. Moreover, as a system manager, the results from this work highlight the harmful potential of self-reinforcing structures within the network for delaying the emergence process, and provide effective solutions for such critical problems.

In the near future, our primary aim is to include and compare other social instruments presented in the literature review in our platform and observe their effects on the convention emergence process. The integration of other social instruments in our platform would allow us to improve their relative effectiveness by analyzing novel combinations of social instruments. The integration of these concepts in social systems and

the analytical study of them would represent another step forward in the integration of social mechanisms into multiagent systems.

## REFERENCES

- ALBERT, R., JEONG, H., AND BARABASI, A. L. 2000. Error and attack tolerance of complex networks. *Nature* 406, 6794, 378–382.
- BABAOGU, O. AND JELASITY, M. 2008. Self-\* properties through gossiping. *Philos. Trans. Series A Math. Physical, Engin. Sci.* 366, 1881, 3747–3757.
- BARABASI, A.-L. AND ALBERT, R. 1999. Emergence of scaling in random networks. *Sci.* 286, 509–512.
- CHAO, I., ARDAIZ, O., AND SANGUESA, R. 2008. Tag mechanisms evaluated for coordination in open multiagent systems. In *Engineering Societies in the Agents World VIII*, A. Artikis, G. OHare, K. Stathis, and G. Vouros, Eds., Lecture Notes in Computer Science, vol. 4995, Springer, 254–269.
- COLEMAN, J. 1998. *Foundations of Social Theory*. Belknap Press of Harvard University Press.
- CONTE, R. AND PAOLUCCI, M. 2002. *Reputation in Artificial Societies: Social Beliefs for Social Order*. Vol. 7, Kluwer.
- DE PINNINCK, A. P., SIERRA, C., AND SCHORLEMMER, M. 2007. Friends no more: Norm enforcement in multi-agent systems. In *Proceedings of the 6<sup>th</sup> International Joint Conference on Autonomous Agents and Multiagent Systems*. 640–642.
- DELGADO, J. 2002. Emergence of social conventions in complex networks. *Artif. Intell.* 141, 1–2, 171–185.
- DELGADO, J., PUJOL, J. M., AND SANGUESA, R. 2003. Emergence of coordination in scale-free networks. *Web Intell. Agent Syst.* 1, 2, 131–138.
- EPSTEIN, J. M. 2000. Learning to be thoughtless: Social norms and individual computation. Working papers 00-03-022, Santa Fe Institute.
- FREEMAN, L. C. 1979. Centrality in social networks: Conceptual clarification. *Social Netw.* 1, 3, 215–239.
- GRIFFITHS, N. AND LUCK, M. 2010. Changing neighbours: Improving tag-based cooperation. In *Proceedings of the 9<sup>th</sup> International Conference on Autonomous Agents and Multiagent Systems*. 249–256.
- HALES, D. 2002. Group reputation supports beneficent norms. *J. Artif. Soc. Social Simul.* 5, 4.
- HALES, D. AND ARTECONI, S. 2006. Slacer: A self-organizing protocol for coordination in p2p networks. *IEEE Intell. Syst.* 21, 29, 35.
- HOLLAND, J. 1993. The effects of labels (tags) on social interactions. Working papers 93-10-064, Santa Fe Institute.
- KITTOCK, J. E. 1993. Emergent conventions and the structure of multi-agent systems. In *Proceedings of the Complex Systems Summer School*. Santa Fe Institute Studies in the Sciences of Complexity Lecture Vol. VI. Lectures in Complex Systems, Santa Fe Institute. Addison-Wesley, 507–521.
- MUKHERJEE, P., SEN, S., AND AIRIAU, S. 2007. Norm emergence in spatially constrained interactions. In *Proceedings of the Working Notes of the Adaptive and Learning Agents Workshop at the 6<sup>th</sup> International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'07)*.
- MUKHERJEE, P., SEN, S., AND AIRIAU, S. 2008. Norm emergence under constrained interactions in diverse societies. In *Proceedings of the 7<sup>th</sup> International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'08)*. 779–786.
- RADWANICK, S. 2010. Facebook captures top spot among social networking sites in india. [http://www.comscore.com/Insights/Press\\_Releases/2010/8/Facebook\\_Captures\\_Top\\_Spot\\_among\\_Social\\_Networking\\_Sites\\_in\\_India](http://www.comscore.com/Insights/Press_Releases/2010/8/Facebook_Captures_Top_Spot_among_Social_Networking_Sites_in_India)
- SALAZAR-RAMIREZ, N., RODRIGUEZ-AGUILAR, J. A., AND ARCOS, J. L. 2008. An infection-based mechanism for self-adaptation in multi-agent complex networks. In *Proceedings of the 2<sup>nd</sup> IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO'08)*. P. R. S. Brueckner and U. Bellur, Eds., 161–170.
- SAVARIMUTHU, B. T. R., CRANFIELD, S., PURVIS, M., AND PURVIS, M. 2007. Role model based mechanism for norm emergence in artificial agent societies. In *Proceedings of the International Conference on Coordination, Organizations, Institutions, and Norms in Agent Systems III (COIN'07)*. 203–217.
- SCOTT, J. 2000. *Social Network Analysis: A Handbook* 2<sup>nd</sup> Ed. Sage Publications.
- SEN, S. AND AIRIAU, S. 2007. Emergence of norms through social learning. In *Proceedings of the 20<sup>th</sup> International Joint Conference on Artificial Intelligence (IJCAI'07)*. 1507–1512.
- SHOHAM, Y. AND TENNENHOLTZ, M. 1997. On the emergence of social conventions: Modeling, analysis, and simulations. *Artif. Intell.* 94, 139–166.
- TOIVONEN, R., CASTELLO, X., EGUILUZ, V. M., SARAMAKI, J., KASKI, K., AND SAN MIGUEL, M. 2009. Broad lifetime distributions for ordering dynamics in complex networks. *Phys. Rev. E Statist. Nonlin. Soft Matter Phys.* 79, 1–8.

- URBANO, P., BALSÀ, J., ANTUNES, L., AND MONIZ, L. 2009. Force versus majority: A comparison in convention emergence efficiency. In *Proceedings of the International Workshop on Coordination, Organizations, Institutions, and Norms in Agent Systems IV (COIN@AAMAS'09)*. Lecture Notes in Computer Science, vol. 5428, Springer, 48–63.
- VILLATORO, D., SEN, S., AND SABATER-MIR, J. 2009. Topology and memory effect on convention emergence. In *Proceedings of the International Conference of Intelligent Agent Technology (IAT)*. IEEE Press.
- WALKER, A. AND WOOLDRIDGE, M. 1995. Understanding the emergence of conventions in multi-agent systems. In *Proceedings of the 1<sup>st</sup> International Conference on Multi-Agent Systems*. V. Lesser, Ed., MIT Press, 384–389.
- WATKINS, C. AND DAYAN, P. 1992. Q-Learning. *Mach. Learn.* 8, 3-4, 279–292.
- WATTS, D. J. AND STROGATZ, S. H. 1998. Collective dynamics of “small-world” networks. *Nature* 393, 6684, 440–442.

Received March 2011; revised March 2012; accepted July 2012