Force versus Majority: A Comparison in Convention Emergence Efficiency

Paulo Urbano, João Balsa, Luis Antunes and Luis Moniz

Departamento de Informática da Faculdade de Ciências de Lisboa
Edifício C6 – Campo Grande
1749-016 Lisboa - Portugal
pub@di.fc.ul.pt

Abstract

In open societies such as multi-agent systems, it is important that coordination among the several actors is achieved efficiently. One economical way of capturing that aspiration is consensus: social conventions and lexicons are good examples of coordinating systems, where uniformity promotes shared expectations of behavior and shared meanings. We are particularly interested in consensus that is achieved without any central control or ruling, through decentralized mechanisms that prove to be effective, efficient, and robust. The nature of interactions and also the nature of society configurations may promote or inhibit consensual emergence. Traditionally, preference to adopt the most seen choices (the majority option) has dominated the emergence convention research in multi-agents, being analyzed along different social topologies.

Recently, we have introduced a different type of interaction, based on force, where force is not defined a priori but evolves dynamically. We compare the Majority class of choice update against Force based interactions, along three dimensions: types of encounters, rules of interaction and network topologies. The experiments we have made show that interactions based on Force are significantly more efficient for group decision making.

Introduction

Shoham and Tennenholtz (1992) have defended that in multi-agent systems agents have to agree on common rules to decrease the number of conflicts and promote cooperative behaviour. These rules take the form of conventions that the agents share to favour coordination. Research goals have been centred around the idea of emergent collective choice in a decentralised way. This has been the theme of such works such as (Delgado 2002; Delgado, Pujol and Sangüesa 2992, Sen and Airiau 2007).

Essentially, the problem can be stated as follows. A group of homogeneous agents has to decide to adopt a behavioural strategy out of a given set. There are no good reasons to prefer one strategy over the other, the individual decision, and even the collective selection are arbitrary. What is important is that everyone adopted the same strategy. We can think of examples such as the direction in which traffic flows: it is quite unimportant whether it is right or left, as long as everyone uses the same lane.

In dynamical systems this problem proves intractable to solve beforehand or in a centralized fashion (Shoham and Tennenholtz 1997), so efforts have been concentrated in developing emergent co-learning processes that allow consensus to be achieved among all the agents (actually, usually only 90% of consensus defines convergence, to allow for isolated agents to exist without any mind-changing interactions).

Agents have a chance to change their convention when they interact with other agents, in a decentralised and locally confined manner. Shoham and Tennenholtz (1994, 1997) have studied the efficienty of convergence in these conditions, and selected the highest cumulative reward (HCR) as the most efficient individual update rule.

In the initial research on convention emergence (Shoham and Tennenholtz 1992; Shoham and Tennenholtz 1994; Shoham and Tennenholtz 1997], there were no restrictions on interactions; any agent could interact by chance with any other individual. Kittock (Kittock 1995) introduced interaction graphs in order to specify restrictions on interactions and made experiences with the HCR update rule along different interaction graphs. Based on his experiments with regular and fully connected graphs, he conjectured that efficiency depends on the diameter of the graph. In what concerns the number of interactions needed to accomplish consensus, Kittock observed a variation with the number of agents of O(N3) for regular graphs and O(NlogN) for fully connected ones.

However, regular graphs are not very realistic. "If we pay attention on real networks, we find out that most of them have a very particular topology, they are complex networks. (…) Complex networks are well characterized by some special properties, such as the connectivity distribution (either exponential or power-law) or the small-world property" research (Delgado, 2002). Delgado et al (Delgado, 2002; Delgado, Pujol and Sangüesa 2002) have made experiments with HCR update rule for fully connected, regular, scale-free and small-world graphs and their results were consistent with Kittock's, confirming the relation between efficiency and graph diameter. Delgado observed also that scale-free and small-world networks were as efficient as fully connected ones, but small-world networks were slower to converge to a unique choice.

In the latter strategy update rules, agents only interact along a succession of pair wise encounters, i.e, they apply their strategy update rule after meeting one agent, randomly chosen inside the group of its "neighbors". Walker and Wooldridge (Walker and Wooldridge 1995) introduced the deterministic Simple Majority (SM) strategy update rule, adapted to simultaneous encounters, where each agent has access to the strategies of all its neighbors before strategy update. Delgado (Delgado 2002) introduced a stochastic variation to SM—he named it the Generalized Simple Majority (GSM).

In (Urbano and Coelho 2005) a new rule for strategy update was introduced, named Recruitment based on Force with Reinforcement (RFR). This rule showed faster convergence than HCR in the case of fully connected networks. In RFR, agents have different power to influence others, but their force is not defined a priori in a hierarchy network as in (Kittock 2004), rather it evolves dynamically along the interactions. Agents submit to stronger agents, copying their strategies, but also inheriting their force, in a double mimetic process.

In (Urbano et al. 2008) we compare Recruitment based on Force against HCR for different topology graphs and RFR proved to be more efficient for Fully connected, Small-world, Scale-free, Regular and Random networks.

The co-existence of several concomitant social networks, and mechanisms permitting permeability among contexts in different networks, has allowed more frequent and quicker convergence that were not possible in some of the hardest cases (Antunes et al. 2007; Antunes et al. 2008).

In the present paper, we advance research on the emergent collective adoption of a common strategy. We expand and summarize our results concerning RFR versus HCR/EM, and introduce a new behavior called Recruitment of the Strongest with Reinforcement, adapted to simultaneous encounters.

This paper is organized as follows: in the next section we describe the different network topologies used in the experiments. Then we introduce two strategy update rules (one based on force, and the other on majority) for pair wise encounters and compare them along the different social graphs. Then we continue by describing strategy update rules (one based on force, and the other on majority) for simultaneous interactions, which are compared. Finally we conclude pointing some future directions. Since this works extends previous work presented in (Urbano et al 2008) it is unavoidable to describe the results of that research here.

Interaction Graph Topologies

The interaction graph topology is a general way of modeling restrictions on inter-actions. Restrictions could be due to spatial barriers, communicating links, different castes, social groups, etc. We have experimented with five network topologies: fully connected, regular, scale-free, small-world and random. The average path length is calculated by finding the shortest path between all pairs of nodes, adding them up, and then dividing by the total number of pairs. It indicates us, on average, the number of steps it takes to get from one member of the network to another. The diameter of a graph is the longest shortest-path between nodes. The clustering coefficient is a measure of "all-my-friends-know-each-other" property. When it is high, we may say: "the friends of my friends are my friends." The clustering coefficient of a node is the ratio of existing links connecting a node's neighbors to each other to the maximum possible number of such links. The clustering coefficient for the entire network is the average of the clustering coefficients of all the nodes.

Regular Graphs

By definition, a graph is considered regular when every node has the same number of neighbors. We are going to use a special kind of regular graph, explored in (Kittock 1995) and named Contract Net with Communication Radius K in (Tennenholtz 1996). $C_{N,K}$ is the graph (regular ring lattice) on N nodes such that node i is adjacent to nodes (i+j) mod N and (i−j) mod N for $1 \leq j \leq K$. In a $C_{N,K}$ graph, every node has connectivity 2*K. These are highly clustered graphs but have very long path lengths (average path length and diameter grow linearly with the number of nodes).

Fully Connected Graphs

In this type of graph topology, named K_N, there are no restrictions on the pattern of interactions: each agent is connected to every other agent in the society. This means that an agent can potentially interact with any other agent. K_N is a special case of a regular graph where each agent has N-1 neighbors, in a group of N agents.

Random Graphs

$R_{N,K}$ are random graphs with N nodes and average connectivity of K. Every node, has on average, K neighbors chosen randomly. The clustering coefficient of $R_{N,K}$ tends to 0 and the average path length is small and grows logarithmically with N.

Scale-free Graphs

This network type, $S_{N,\gamma}$ has a large number of nodes connected only to a few nodes and a small number of well-connected nodes called hubs. The power law distribution highly influences the network topology. It turns out that major hubs are closely followed by smaller ones. These ones, in turn, are followed by other nodes with an even smaller degree, and so on. As the network changes in size, the ratio of hubs to the number of nodes in the rest of network remains constant—this is why it is named scale-free. The connectivity of a scale-free network follows a power law $P(k) \sim k^{-\gamma}$. Such networks can be found in a

surprisingly large range of real world situations, ranging from the connections between websites to the collaborations between actors.
To generate the scale-free graphs we have used the Albert and Barabási extended model (Albert and Barabási 2000), since Delgado argues that it allows some control over the exponent (γ) of the graph (Delgado 2002). The inspiration of this algorithm is that of "preferential attachment," meaning that the most "popular" nodes get most of the links. The construction algorithm relies on four parameters: m_0 (initial number of nodes), m (number of links added and/or rewired at every step), p (probability of adding links), and q (probability of edge rewiring). The algorithm starts with m_0 isolated nodes and at each step performs one of these three actions until the desired number N of nodes is obtained:

(1) with probability p, add m ($\le$ m0) new links. Pick two nodes randomly. The starting point of the link is chosen uniformly and the end point of the link is chosen according to the probability distribution:

$$\Pi_i = (k_i + 1) \Big/ \sum_j (k_j + 1) \ ,$$

where Π_i is the probability of selecting the ith node and k_i is the number of edges of node i. This process is repeated m times.

(2) with probability q, m edges are rewired. That is, repeat m times: choose uniformly at random one node i and one link l_{ij}. Delete this link and choose a different node k with probability $\{\Pi_l\}_{l=1,\ldots,N}$ and add the new link l_{ik}.

(3) with probability 1-p-q add a new node with m links. These new links will connect the new node to m other nodes chosen according to $\{\Pi_l\}_{l=1,\ldots,N}$.

Using this algorithm, the parameter γ is a function of m and p:

$$\gamma = (2m(1-p)+1)/m+1$$

Small-world Graphs

The Small World graphs are highly clustered graphs (like regular graphs) with small average path lengths (like random graphs, described above). To generate small world graphs we use the Watts-Strogatz model (Watts 1999; Watts and Strogatz 1998). It depends on two parameters, connectivity (K) and randomness (P), given the size of the graph (N).
This model starts with a $C_{N,K}$ graph and then every link is rewired at random with probability P, that is, for every link l_{ij} we decide whether we change the "destination" node with probability P; if this is the case, we choose a new node k uniformly at random (no self-links allowed) and add the link l_{ik} while erasing link l_{ij}. In fact, for P = 0 we have $W_N = C_{N,K}$ and for P = 1 we have a completely random graph (but not scale-free). For intermediate values of P there is the "small-world" region, where the graph is highly clustered (which means it is not random) but with a small characteristic path length (a property shared with random graphs).
Albert-Barabási model graphs have not the small-world property and reciprocally the Watts-Strogatz model does not generate scale-free graphs (it generates an exponential connectivity distribution, not a power law).

Strategy Update Rules in Pair wise Interactions

Agent societies consist of N agents on a graph, where every agent is located on a node of the graph. Its adjacent nodes are its neighbors. In order to make experiments and simulations we have adopted a simple agent model where they have at their disposal a finite repertoire of strategies. We only deal with the two strategies case. We use here the concept strategy in a very abstract way: it can be a social norm, like driving on the left or on the right lane, the meaning of a word, an orientation for flocking, etc. In order to focus on the essential features of agent interactions, the agent environment consists solely of other agents, which in turn depend on the network topology. So, each agent has to adopt one of the strategies from the repertoire and through mutual interactions they can change their adopted strategies along time. A consensus, or collective choice, exists when all the agents are using the same particular strategy.
From the point of view of each agent, there is an interaction scenario of a sequence of pair wise asymmetric encounters, where it meets randomly one of its neighbors. After an encounter, each agent updates its strategy, i.e., it selects the strategy it will use in the next interaction—the result need not necessarily be a change in strategy adoption. Therefore, agents need strategy update rules (behaviors). We assume that each agent updates its strategy at each encounter. Shoham and Tennenholtz (Shoham and Tennenholtz 1992; Shoham and Tennenholtz 1997) have studied the effects of updating less frequently on the efficiency of global choice emergence. We only consider asymmetric encounters, where only one of the agents applies its strategy update rule, based on the strategies used by all the individuals involved in the interaction. Thus, interactions are always considered from the point of view of some particular agent. We now describe the two strategy update rules whose performance we subsequently compare. In the first scenario the agent and its selected partner strategies are crucial for the update, and in the second case, it is the simultaneous strategies of its neighbors and its own strategy that influence the update rule.

External Majority/ Highest Cumulative Reward/ Feedback Positive with Score

The External Majority strategy update rule (EM) was introduced by Shoham and Tennenholtz (Shoham and Tennenholtz 1992) and is the following: adopt the strategy that was observed more often in other agents in the last m interactions, and remain with your current strategy otherwise—in case of a draw do not change. In EM,

memory is used to register the strategies observed during the last interactions. An agent updates its memory after observing its partner strategy and then decides to change to a new strategy only in case it was more frequently observed than the current one.
In the context of lexical emergence, Kaplan (Kaplan 2000) introduced a strategy update rule called Positive Feedback with Score, which is pretty much the same as EM. The only difference is that, in case of equality, the agent does not necessarily remain with its current strategy but chooses randomly one of the previously most seen strategies. Kaplan considered the full history of encounters for strategy update.
The most referred strategy update rule is the Highest Cumulative Reward update rule (HCR), which was developed in the context of game theory by Shoham and Tennenholtz (Shoham and Tennenholtz 1994). Intuitively, a game involves a number of players each of which has available to it a number of strategies. Depending on the strategies selected by each agent, they each receive a certain payoff. The payoffs are captured in a payoff matrix. Thus, returning to the context of this paper, when two agents meet they play a pure coordination game, which is an instance of the class of coordination games introduced by Lewis (Lewis 1969). The pure coordination game is defined by the following symmetric payoff matrix:

	A	B
A	+1,+1	-1,-1
B	-1,-1	+1,+1

Suppose that every player has two available strategies, say A and B. If both players play A, both players receive a payoff of 1. If they play B they receive a payoff of 1. When the players do not agree, for example, player 1 plays A and player 2 plays B, they will both receive a payoff a -1; the remaining situation is symmetric. The condition on the entries of the payoff matrix makes it clear that the best action consists in playing the same strategy, i.e., coordinating.
According to the HCR update rule, an agent switches to a new strategy if and only if the total payoff obtained from that strategy in the latest m interactions is greater than the payoff obtained from the current strategy in the same last m interactions. The m parameter may not have limit, implying that the full history of pair wise meetings will play a role in the strategy selection process, or we can implement a forgetting mechanism by limiting m. The agents' memories register the payoffs that each strategy has received during the last m encounters. When an agent receives new feedback it discards its old memory to maintain the memory at a fixed size.
Shoham and Tenneholtz (Shoham and Tennenholtz 1997) showed that EM and HCR are equivalent strategy update rules in the case there is a repertoire of two strategies to select.

Recruitment based on Force with Reinforcement

In this strategy update rule there is a new attribute, besides the strategy, called force. Thus, agents are characterized by two attributes: strategy and force. These attributes can be observed during encounters. During a dialogue (asymmetric), involving two agents, one is the observing agent and the other is the observed one. The observing agent "fights" metaphorically with its partner, comparing its force with the partner's force. If the observing agent is the stronger one, or if they have identical force, it will loose the fight; otherwise it will be the winner. The winner's behavior is: (1) if they have the same current strategy its force is reinforced by 1 unit, otherwise (2) it does nothing. The loser's behavior is: (1) it imitates both strategy and force in case they have different current strategies, otherwise imitates the force of the winner agent and increments its force by 1 unit. In sum, stronger agents recruit weaker agents for their parties, enlarging the influence of their options. As the recruited agents will be at least as strong as the winners, they will be better recruiters.

At the beginning of an experiment, every agent has the same value of force (0) and their forces evolve along with interactions. Therefore, there is no a priori (off line) power hierarchy. This shows a clear contrast with the work of Kittock on authority (Kittock 2004), where agents have fixed different influences on one another's behavior, modeled by the probability of receiving feedback during encounters: the more influential agents receive feedback with some probability, while the less influential agents always receive feedback.

The force attribute can be interpreted not as the strength of an individual because, being imitated, it is diffused along agents, and does not belong to any agent, but as the force of the strategy the player is adopting. The more the strategy is diffused the more it will have stronger representatives. So when a player observes a stronger agent it is recruited, inheriting his force, i.e., updating the information about the strategy it is now adopting. There is a positive reinforcement when an agent faces another with the same strategy during a meeting, which is the amplification mechanism for strategy diffusion. RFR does not try to simulate any natural behavior and it was introduced in (Urbano and Coelho 2005) as the best outcome of experimenting with several strategies involving the idea of emergence of hierarchies in consensus emergence. We think it is simpler than the EM as agents do not need to maintain the recent history of encounters in spite of using force as an extra attribute. It maintains the essential properties of EM, which is the capability to adapt, locality (an agent relies only is the information gathered in interactions), and no more cognitive skills than the capability to imitate.

Experiments and Results for the Pair wise Strategy Update Rules

Experiments were conducted using the most recent version of the Netlogo platform (Wilensky 2003), version 4.0.2, released December 2007.
The system starts with half of the agents adopting randomly one of the strategies (50% possibilities for each). In each step, every agent, in an asynchronous way, is selected and chosen for asymmetric strategy updating. The order of selected agents is completely random and changes in each iteration. We use the same measure of performance as in (Kittock 2005; Delgado 2002): average number of interactions to a fixed convergence, where convergence means the fraction of agents using the majority strategy. We made 100 runs for each parameter setting and in each run we have measured the number of encounters until 90% convergence and calculated the average performance of the different runs.
Our main goal was to compare the performance of the two strategy update rules: External Majority (EM) and Recruitment based on Force with Reinforcement (RFR). In order to choose the size of memory of EM update rule we have made a lot of experiments with different types of networks and the size = 3 achieved the best performance. Both Kittock (Kittock 2005; Kittock 2004) and Delgado (Delgado 2002; Delgado, Pujol and Sangüesa 2002) chose size = 2 in their convention experiments, but size 3 EMs out-performed size 2 EMs in our own tests.

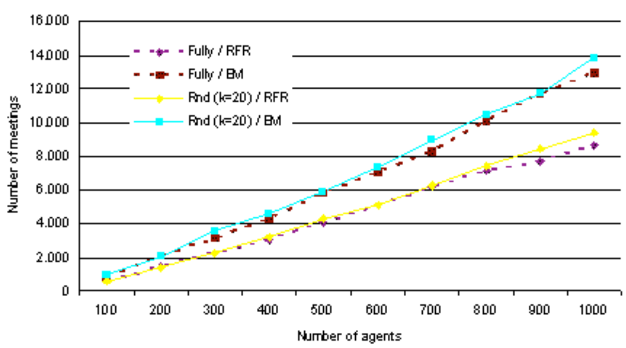

Figure 1 – RFR vs EM in and random networks (with 40 neighbors per agent on average - $R_{N,20}$) and fully connected networks (K_N).

So we will only present the comparison between RFR and the best EM (EM-3). The comparison between these two behaviors was made along the different kinds of networks described before, using different parameter settings. Again, we only present here the most representative experiments.
For all settings we made the number of agents range from 100 to 1000, using a step of 100.
In figure 1, we can see a comparison between the average number of meetings needed for a 90% convergence using the two behaviors in fully connected (dashed lines) and random networks (K=20, solid lines), and with RFR (diamonds) and EM (squares) as update rules. In both cases RFR clearly outperforms EM, since consensuses are reached in a much lower number of meetings. Besides, the difference appears to increase with N. In fact, in the above experiments, improvements vary between 26% and 37%.

Using other types of networks, the difference between the two behaviors is even clearer. In figure 2 we compare the performance of the two behaviors in small-world networks (with P = 0.1, dashed lines), regular networks (with K = 20, each agent with 40 neighbors, solid lines) and scale-free networks (with γ=2.5, dotted lines). In these cases we need to use a logarithmic scale in the y axis. The average number of meetings needed to reach a consensus is again much lower when RFR (diamonds) is used. In Scale-free networks, the performance is improved between 25% and 37%; in Small-world networks improvements reach 80%. Using regular networks the difference with N greater than 500 is so huge that it is not represented.

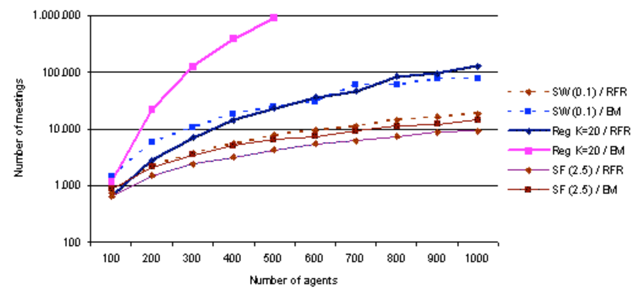

Figure 2 - RFR vs EM in small-world, W_N(K=12; P=0,1), regular, $C_{N,20}$, and scale-free networks, $S_{N,2.15}$.

Strategy Update Rules in Simultaneous Interactions

In contrast with previous experiences, now, in each encounter an agent interacts not only with a randomly chosen neighbor, but with all his neighbors at the same time. Interaction is asymmetric, seen from the point of view of an agent—the simultaneous state of neighbors is used to update agents' strategy, and naturally will have to be taken into account in the strategy update rules. Again, each agent updates its strategy at each encounter.

Generalized Simple Majority

The natural strategy update rule to use in simultaneous interactions with all neighbors, equivalent to EM, for the pair wise encounter, would be the Simple Majority created by Walker and Wooldridge (Walker and Wooldridge 1995). Following this rule, an agent only changes strategy when more than half of the neighboring agents adopt a different strategy than its current one. This deterministic rule does not guarantee convergence to a consensual situation for some networks, specially the regular ones. In some cases, even in small groups, agents get stuck in a deadlock, never reaching a situation of full convergence or even 90% convergence. Delgado (Delgado 2002)

developed a stochastic version of Simple Majority, which they named the Generalized Simple Majority (GSM). We are going to describe GSM, not exactly the sane way as was originally defined by Delgado.
Suppose we have N agents in a graph with a well-defined neighborhood for every agent. If agent j has K neighbors it will adopt state S with a probability that depends on the number of neighbors adopting S (K_S):

$$f_\beta(K_S) = \frac{1}{1 + e^{2\beta(2(1-K_S/K)-1)}}$$

The formula above is a corrected version of the one introduced by (Delgado 2002). This rule generalizes Simple Majority since, for $\beta \to \infty$, an agent adopts state S only when at least half of the neighbors are in that state. Note that when a tie occurs, the probability of changing is 50%. In our experiences, we used $\beta = 10$ – it was also used in (Delgado 2002).

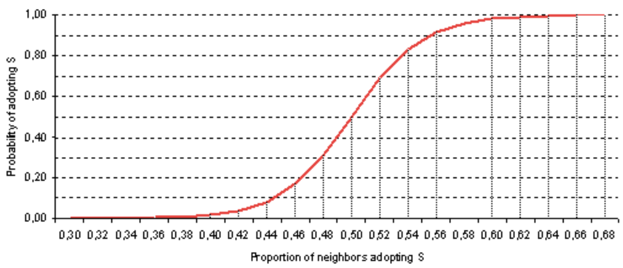

Figure 3 – Probability of changing to a new state S, for β=10, given the percentage of neighbors that are adopting S.

As illustrated in figure 3, the probability of adopting S is positive even when the neighbors adopting S are not in majority. This is true for values greater than 38%. Conversely, it is possible to adopt a state that is adopted by the minority of the neighbors, which can happen if the majority is no greater than a percentage of 62%. Note that when a tie occurs, the probability of changing is 50%. Delgado provided some analytical evidence for the convergence of GSM, but no theorem exists that guarantees it. In our experiments, GSM has always converged to a consensual situation of 90%, in all types of networks, varying the number of agents until 1000.

Recruitment by the Strongest with Reinforcement

In order to adapt the RFR update rule to simultaneous interactions, partners in encounters will not be chosen randomly as before. Now an agent can access the strategies of all its neighbors—it will choose the strongest to compare forces and eventually imitate it and for reinforcement. The behavior is pure Recruitment based on Force with Reinforcement (RFR), but it will be with the strongest of its neighbors, not with a random chosen neighbor. We will name this strategy update rule Recruitment by the Strongest with Reinforcement (RSR). There is simultaneous access but real interaction (applying RFR) is only with the strongest.

Experiments and Results for the Simultaneous Strategy Update Rules

The comparison between RSR and GSM was made along the different kinds of networks described in section 2, using different parameter settings. Again, we only present here the most representative experiments and for all settings we made the number of agents range from 100 to 1000, using a step of 100. We counted the number of encounters necessary for 90% consensus along 100 runs for each parameter setting and values were averaged.
In figure 4, we can see a comparison between the average number of meetings needed for a 90% convergence using the two behaviors in fully connected (dashed lines), random networks (K=20, dotted lines) and scale-free networks (with γ=2.15, solid lines) with RSR (squares) and GSM (diamonds) as update rules.

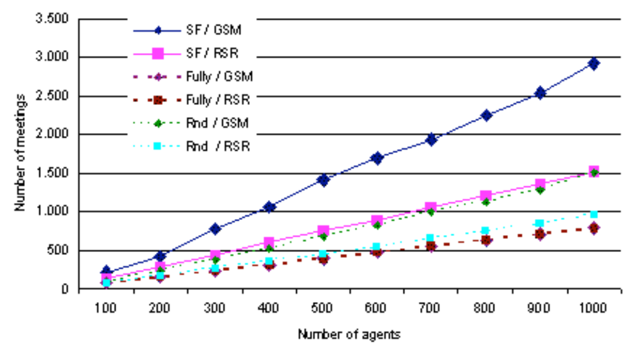

Figure 4 – RSR vs GSM in and random networks (with 40 neighbors per agent on average - $R_{N,20}$), fully connected networks (K_N), and scale-free networks, $S_{N,2.15}$.

With the exception of fully connected networks, RSR clearly outperforms GSM, since consensuses are reached in a much lower number of meetings. In the fully-connected case, GSM is slightly better, around 1%, which is marginal and not perceptible in the chart. In the Random networks case, we have an improvement for RSR around 16% for N=100 and gradually the difference increases, and stabilizes around 35% for N greater than 500. In the Scale-free case, RSR, again, exceeds GSM, with an improvement which varies between 33% and 48% (N > 500).

Using other types of networks, the difference between the two behaviors is even clearer. In figure 5, we compare the performance of the two behaviors in small-world networks (with K=12, P = 0.1, dashed lines) and regular networks (with K = 20, solid lines). In these cases we need to use a logarithmic scale in the y axis. The average number of meetings needed to reach a consensus is now much lower when RSR (squares) is used. In Regular networks, the performance is improved in 99%, reaching 99,9% for N above 600; in Small-world networks improvements reach 93% for N=100, reaching 98% for N bigger than 300.

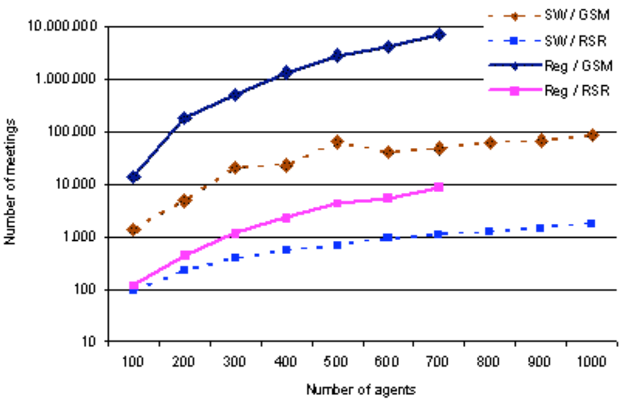

Figure 5 – RSR vs GSM in small-world, W_N(K=12; P=0,1), and regular networks, $C_{N,20}$.

Conclusions and Future Work

In what respects our primary goal of comparing strategy update rules based on majority against ones based on force, the main conclusion is that those based on force almost always perform better than the ones based on majorty. This conclusion is valid for all types of networks with different parameter settings, except the case of fully connected graphs in simultaneous encounters. Even considering that we did not exhaust all type of networks and its parameters, this is an impressive result. According to our experiments strategy update rules based on force, represent at least a 25% improvement over the majority one, but is much higher in many settings, most extremely in regular networks.

An interesting point is that our results for the RFR strategy update rule show that the network diameter strongly influences the performance, as noted in other settings by authors such as Kittock (Kittock 2005) and Delgado (Delgado 2002). The smaller the network diameter is, the better. Also, the performance of fully-connected networks and scale-free ones seems to be quite similar (as can be observed in figure 1, comparing lines with similar marks, that correspond to the same behavior). This is a very important result since scale-free networks are much less expensive than fully-connected ones. Nevertheless, we must perform experiments with greater values of N in order to obtain definite conclusions on this matter.

Besides performing experiments with larger values of N, three other aspects are scheduled for short-term future work. One is to consider that agents can choose between more than two strategies. The other is to explore the performance of these behaviors in networks with dynamical structure, and finally we want to see force strategy update rules applied to situations where agents can choose between strategies valued differently.

References

Albert, R., Barabasi, A.-L. 2000. Topology of evolving networks: Local events and universality, Phys. *Rev. Lett. 85, 5234-5237.*

Antunes, L., Balsa, J., Urbano, P., Coelho, H. 2007. *The Challenge of Context Permeability in Social Simulation*. Proceedings of the The Fourth European Social Simulation Association Conference.

Antunes, L., Balsa, J., Urbano, P., Coelho, H. 2008. *Exploring Context Permeability in Multiple Social Networks*. (submitted)

Barabasi, A.-L. 2002. *Linked: The New Science of Networks*, Perseus Publishing, Cambridge, Massachusetts.

Delgado, J. 2002. Emergence of Social Conventions in Complex Networks. *Artificial Intelligence*, 141, 171-185.

Delgado, J., Pujol, J.M., Sangüesa, R. 2002. Emergence of coordination in scale-free networks. *Web Intelligence and Agent Systems* 1(2) , 131-138.

Kaplan, F. 2000. *L'Emergence d'un Lexique dans une Population d'Agents Autonomes*. These de doctorat de l'Université de Paris VI.

Kittock, J.E. 1995. Emergent conventions and the structure of multi-agent systems, in: L. Nadel, D. Stein (Eds.), *1993 Lectures in Complex Systems*, in: SFI Studies in the Sciences of Complexity, Addison-Wesley, Reading, MA.

Kittock, J.E. 1994. The impact of locality and authority on emergent conventions: Initial observations, in *Proc. AAAI-94*, Seattle, WA, AAAI Press, pp. 420–425.

Lewis, D.K., 1969. *Convention: A Philosophical Study*, Harvard Univ. Press, Cambridge, MA.

Sen, S., Airiau, S. 2007. Emergence of norms through social learning. In *Proceedings of the twentieth International Joint Conference on Artificial Intelligence (IJCAI'07).*

Shoham, Y, Tennenholtz, M. 1992. Emergent conventions in multi-agents systems: initial experiments results and observations. In *Proceedings of the 3rd International Conference on Principles of Knowledge and Reasoning*, (1992) 225-231.

Shoham, Y, Tennenholtz, M. 1994. Co-learning and the evolution of social activity. *Technical Report STAN-CS-TR-94-1511*, Department of Computer Science, Stanford University, Stanford, CA, USA.

Shoham, Y, Tennenholtz, M. 1997. On the Emergence of Social Conventions: modeling, analysis and simulations. *Artificial Intelligence* 94, 139-166.

Tennenholtz, M. 1996 Convention evolution in organizations and markets. *Computational and Mathematical Organization Theory* 2, 261-283.

Urbano P., Coelho, H. 2005. From Consensus to Consensus: random evolution of collective choices, in *Proceedings of EPIA 2005 MASTA Workshop*, Carlos Bento, Amilcar Cardoso and Gaël Dias (eds.), Covilhã (Portugal), December 5-7, IEEE, 2005..

Urbano, P., Balsa., J., Antunes, L. and Moniz, L. 2008. Efficiency of the Emergence of Consensus in Complex Networks – assessing force influence. (Submitted).

Watts, D.J. 1999. Small Words, the Dynamics of Netwroks between Order and Randomness, Princeton University Press, Princeton, NJ.

Watts. D.J. and Strogatz, S.H. 1998. Collective dynamics of 'small-world' networks, *Nature*, 393, 440-442.

Wilensky, U. 1999. NetLogo. http://ccl.northwestern.edu/netlogo/. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.