Exploring Indirect Reciprocity in Complex Networks using Coalitions and Rewiring

Ana Peleteiro and Juan C. Burguillo
Dept. of Telematics Engineering
University of Vigo, Spain
{apeleteiro,J.C.Burguillo}@gti.uvigo.es

Siang Yew Chong
School of Computer Science
University of Nottingham, Malaysia Campus
Siang-Yew.Chong@nottingham.edu.my

ABSTRACT

It is generally known that cooperation can be achieved in complex real-world interactions that are not limited to direct interactions only. In particular, cooperation can consider prior interactions with other players, i.e., indirect reciprocity. Moreover, coalition based mechanisms have shown to facilitate cooperation among self-interested agents. Also, research on games over dynamic topologies has found empirical evidence showing that partner switching leads to cooperative behavior. In this paper we present a new mechanism to improve cooperation among self-interested agents placed in a complex network. Our mechanism is based on three main pillars: indirect reciprocity, coalitions and rewiring. Thus agents play against each other an indirect reciprocity game where they can create coalitions to share information about agents' reputation or change their personal network (social contacts). Altogether, we explore the conditions to enhance cooperation in complex networks. Finally, in our experiments we determine how, by using our mechanism, cooperation is improved in our reputation-based game, and how topology highly influences cooperation in our scenario.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence-Multiagent Systems

Keywords

donation game, indirect reciprocity, coalitions, rewiring

1. INTRODUCTION

Games are useful mathematical constructs to abstract and model real-world problems involving strategic decision-making in a variety of contexts ranging from politics, economics, to biology among others [2]. In particular, they can capture the intrinsic properties of these problems through the specification of rules that constrain strategies to certain behaviors (legal moves they can make as responses to stimulus, such as historical plays), goals for strategies to meet (to win the game), and rewards under finite resources (payoffs) [6]. There are many settings with which games have been proposed as models for study. For example, the prisoner's dilemma game embodies the *dilemma* of two interacting individuals (players) who are better off mutually cooperating than mutually defecting, being vulnerable to exploitation by one who defects [4]. In multi-player games, there is an interesting situation known as *the Tragedy of the Commons*. Such framework reflects the *tragedy* of all players pursuing self-maximizing behaviors that lead to the worst outcome for the community, rather than collectively cooperating for a better result [11, 15].

It is generally known that cooperation can be achieved in complex real-world interactions that are not limited to direct interactions only [16]. In particular, cooperation can consider prior interactions with other players, i.e., indirect reciprocity. In a highly simplified example, the donation game is used to show how the mechanism of indirect reciprocity operates using players' reputation to promote cooperation [17]. Unlike the case of direct reciprocity, whereby any altruistic act of helping to another player is returned, in indirect reciprocity the altruistic act of helping others is perceived by the community as *helpful*, providing good reputation, and receiving help in return by other players. Indirect reciprocity is also associated with interactions having short encounters (e.g., one-shot interactions) whereby the effects of direct reciprocity on the interaction outcome are minimized.

Indirect reciprocity has been used by players that compare reputation of potential recipients and cooperate only when the recipient has the same or higher reputation than the donor's strategy. It can be shown that a population of such players can evolve cooperative plays through discriminators that can distinguish players with high reputation (those that have cooperated with other players in past interactions) and cooperate only with such players [17]. Other studies have applied the mechanism of indirect reciprocity in complex interactions where cooperative plays are difficult to evolve. Chong et al. [7] have shown that the mechanism of indirect reciprocity through repeated interactions is less effective in promoting cooperation for interactions with higher number of alternative choices and shorter encounter (e.g., lower number of rounds in a repeated game). However, strategies can evolve to use reputation as a mechanism to estimate behaviors of future partners and to elicit cooperation right from the start of interactions. Cooperation occurs when strategies evolve to maintain high reputation scores.

On the one hand, the notion of coalitions has been studied by the game theory and multi-agent community for decades. In fact, coalition formation [21, 22] is one of the fundamental approaches for establishing collaborations among self-interested agents. For instance, coalition-based mechanisms as [3, 12, 20] confirm that coalitions indeed facilitate cooperation among self-interested agents. However, those works consider that there is a leader imposing the behavior of the whole coalition, and that it gets taxes for it. Also, in those approaches, the decision to join a coalition is based on the benefits that the coalition provides to its members. Not only that, but most approaches in coalitions formation consider that agents within a coalition must cooperate and/or agree to act in a prede-

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

fined way. However, coalitions may also be seen as flat groups where agents join to share some information since they may get benefits from it, even if part of such information is private. Not only that, but in this case an agent may use the coalition reputation as a group to decide whether to join or not, instead of its benefits.

On the other hand, research on games over dynamic topologies has found empirical evidence showing that partner switching leads to cooperative behavior. Along this line, Fu et al. [9] propose a coevolutionary model of the prisoner's dilemma that allows agents to either adjust their strategies or switch their defective partners. Thus, they show that partner switching may help to stabilize cooperation. Also in [8], Fu et al. focus on the effect of reputation on an individual's partner switching problem in a network, showing that using their mechanism, cooperation can prevail. Although in a different framework (the investigation of tag-based coordination), Griffiths et al. [10] show that partner switching can help to increase coordination resilience in the face of malicious behavior. However, previous work in partner switching considers that agents can only interact with their direct neighbors only, while it may be the case that, even having their social networks, an agent may interact with any other agent from the population. Moreover, previous approaches consider only the individual reputation of the neighbors, while using the group reputation may be also beneficial.

In this paper we present a mechanism to improve cooperation among self-interested agents placed in a complex network, where agents play the donation game with any other members of the population. Our mechanism is based on three main pillars: indirect reciprocity, coalitions and rewiring. Our coalition formation mechanism differs from previous approaches since agents in a coalition do not agree to control or behave in a certain way neither with agents inside the coalition, nor with agents outside of it. Instead, coalitions are groups of agents that share information regarding reputation that might result beneficial for them. This is why we propose to use a coalition reputation measurement to decide to which coalition to join. Concerning the dynamics of agent behavior, and as our agents are placed in a social network, they may imitate their neighbors' strategies, if they seem successful in terms of payoff. Finally, to improve cooperation even further, we include a rewiring mechanism that uses the reputation of the neighbors to change their social links (i.e., rewire). In our experiments we determine that cooperation is improved when we include our coalition and rewiring mechanism. Moreover, we analyze how topology influences cooperation in this scenario.

The rest of the paper is organized as follows. First, in Sect. 2 we consider the basic donation game model that we consider in our framework. Then, in Sect. 3 we extend such model using coalitions and partner switching (rewiring). In Sect. 4 we describe the simulation results obtained from our framework. Finally, Sect. 5 presents the conclusions, and points out some promising directions for future work.

2. DONATION GAME RULES

Our donation game is based in the classic donation game published by Nowak and Sigmund [17] involving image scoring strategies which are a measure of reputation. As described in their paper, the game is composed of several rounds where N agents play the donation game. In each round, a small set of m donor-recipient pairs are chosen. Therefore, the chance that a given player meets the same player again is negligibly small. Thus, direct reciprocity cannot work here.

From each each pair of agents, one is selected as the donor, and the other one as the recipient. Every agent i has a strategy represented by the integer $k_i \in [-5, 6]$ and an image score (reputation) given $s_i \in [-5, 5]$ that depends on its behavior in the past. The donor i has to decide, depending on its strategy (k_i), and its the opponent j score (s_j), if it cooperates (donates) with the other agent. If $k_i \leq s_j$, then agent i donates a benefit b to agent j at a cost c to itself, and increases its image score (s_i) by 1. Otherwise (i.e., $k_i > s_j$), no donation or cost are involved (both obtain zero payoff) but the image score of the donor (s_i) is decremented by 1. Note that the image score of the recipient does not change in any case.

Hence, strategies with $k \leq 0$ are termed cooperative, because individuals with these strategies cooperate with individuals that have not had an interaction. Then, we can observe two extreme game-playing strategies, i.e., the strategy with $k_i = -5$ represents cooperation regardless other agent's score, while the strategy $k_i = 6$ represents defection in all cases. Other strategies represent various degree of discriminating play, e.g., $k_i \in [-4, 0]$ are discriminators that lean towards cooperation [17].

In our case, after finishing a round, agents imitate the best strategies in their neighborhood, while in [17] agents reproduce themselves, to produce a new population, depending on their obtained payoff. Note that that in both cases, depending on the value of m and the random selection, it may happen that there are differences in the amount of times that different agents have played the donation game in a round. However, what is relevant is the evolution of the whole game, and not what happens to a particular agent.

3. MODEL

We consider a population of N agents where any agent can interact with any other agent (i.e., panmictic interaction) to play the donation game (see Sect. 2). However, agents are connected in a complex network, having each of them a set of peers that constitute their neighborhood. We want to model real world interactions over social networks, so agents' neighbors are their close related contacts from which agents obtain information. However, in real world, apart from having a set of direct contacts, people usually belong to several clubs, associations, organizations, or groups in general. We model this second set of contacts with the notion of coalitions, as a way that agents may share some information about the environment where they play. Thus if an agent agrees to become a member of a coalition, it also agrees to share information with the rest of the coalition members. This information sharing helps agents while interacting with the whole population in the panmictic game.

In Algorithm 1 we present the basic game behavior, that will be explained detail in the following sections. As a short description, we can see how pairs of agents (line 3) play the donation game during a round (set of encounters), and that any agent has to decide:

- Its action (to donate or not) depending on its own strategy, and the other's image score (line 4). This influences its payoff and image score (line 5).
- To keep independent or join a coalition, and if joining, to which one (line 8).
- Deciding its new strategy for the next round (line 9).
- Changing their neighbors, depending on the image score of the neighborhood (line 10).
- Finally, the payoff and image score are reset for the next round (line 11).

Algorithm 1 Game Behavior

1: **function** $PlayRound(m \leq N)$
2:   **for** $m\ times$ **do**
3:     $(a_i, a_j) = FindPlayers(i, j \in [1, N]; i \neq j)$
4:     $a_i.DecideAction(s_i, k_j)$
5:     $a_i.ChangePayoffScore()$
6:   **end for**
7:   **for** $all\ a_i$ **do**
8:     $a_i.ChangeCoalition()$
9:     $a_i.ChangeStrategy()$
10:     $a_i.Neighborhood = Rewire()$
11:     $a_i.ResetPayoffScore()$
12:   **end for**
13: **end function**

3.1 Reputation Sharing

In order to decide their strategy, and to maximize their payoff, agents need to know their opponents' image score. This is a challenging issue, since each of the agents can play with any other in the population. In Nowak's model [17], they use two approaches to solve this problem. First, they consider that image score is public, and that all agents know the image score of any other agent in the population. Second, the authors consider that there exist a small percentage of agents (neighbors) that can observe a particular interaction; and only those agents, plus the recipient, update the other agent's image score. The first scenario is an idealistic one, while in the second scenario, each agent has a different perception about the image score of the others.

In this paper, we model reputation sharing in a different manner. Each agent has a set of neighbors, and this neighborhood represents the direct contacts (friends or mates) that an individual has. We assume that each agent knows the image score of its neighbors. At the same time, we assume that agents may belong to coalitions, that models groups of interest, or organizations, that shares reputation information among its members. Therefore it models a global exchange of information biased by the different coalitions.

Thus, differing to [17] and as in [8], in our model agents are connected to others in a complex network, where each of the agents has a neighborhood. However, as in [17] and differing to [8], each agent may interact with any other agent of the population. We do not consider agents playing only in their neighborhood, since then agents could have a direct reputation from its neighbors. Therefore, as each player may interact with any other in the population, direct reciprocity does not work, since the chances of one player interacting again with the same player are negligibly small [17].

3.2 Action Selection

In previous sections we have presented the donation game and how reputation information flows among the agents. Now, in Algorithm 2 we proceed to explain how a donor acts in our model when it encounters with a recipient (line 4, Algorithm 1).

Once a random pair of agents a_i and a_j has been randomly selected to interact, and their roles are defined, the donor (a_i) checks if the recipient (a_j) belongs either to its neighbors or to its coalition mates (line 2, 2). If it belongs to any of those groups, then we assume that the donor knows the score of the recipient. In the contrary case, as it has no information, it assumes that the image score of a_j is 0 (following [17]). After this, the donor has to decide, depending on its strategy (k_i), if it donates to the recipient, providing a benefit b with a cost c to itself (line 7). This action increases its image score (line 8). On the contrary, if a_i does not donate, both individuals receive zero pay-off, but the image score of the donor is decreased by one (lines 10 and 11).

Algorithm 2 Behavior of a donor a_i

1: **function** $ChangePayoffScore(a_i, a_j)$
2:   **if** $a_i.InCoaOrNeighbor(a_j)$ **then**
3:     $s_{ij} = a_j.GetScore()$
4:   **else**
5:     $s_{ij} = 0$
6:   **if** $k_i \leq s_{ij}$ **then**
7:     $a_i.Donate(a_j, b, c)$
8:     $a_i.ChangeScore(+1)$
9:   **else**
10:     $a_i.Donate(a_j, 0, 0)$
11:     $a_i.ChangeScore(-1)$
12: **end function**

3.3 Coalition Formation

In our approach, we allow agents to form coalitions in order to share reputation information and therefore to improve cooperation. We consider that when an agent joins a coalition, it agrees to share its image score with the rest of the coalition members, but also obtains the image score of the other members of that coalition. We impose that any agent can only belong to a unique coalition at a time, since we consider that coalitions somehow compete in the game. Moreover, it is a simplified assumption in the model to avoid excessive complexity. Finally, agents belonging to a coalition are not necessary neighbors.

Each coalition has an image score that depends on the average image score of its members. Let Coa_j represent the coalition j with coalition member agents indexed by i. The size of the coalition $|Coa_j|$ gives the number of agents in the coalition. The coalition score, CS_j, is specified as follows:

$$CS_j = ln|Coa_j| \cdot \frac{\sum_{i \in Coa_j} s_i}{|Coa_j|} \tag{1}$$

where s_i is the image score of member agent i. We include a scaling factor $ln|Coa_j|$ to model that larger-sized coalitions have more influence to attract agents to join those coalitions, as the amount of information they may share is bigger. At the same time, if a cooperative agent joins a bigger coalition it has more chances to be identified as a cooperator (if behaving as a donor) in successive encounters, and obtaining more donations in average than if it was isolated or in a smaller coalition. Thus, the benefits for the agent are double: a higher probability to obtain donations, and better options to rewire.

In Algorithm 3 we present the rules for coalition dynamics, which we adapt from other approach that considers direct interactions [12]. The decision to join a coalition is based on simple rules, as in [3], which precludes modeling sophisticated agents that can learn about the rules to form coalitions.

Firstly, if an agent that belongs to a coalition is isolated from its coalition mates, i.e., none of its neighbors belongs to its coalition, then it becomes independent (line 4). We do this since we consider that each agent in a coalition must have at least one connection to another coalition member to transmit/receive information. Otherwise, it checks the payoff of its neighbors to see if its payoff P_i has been the worst in its neighborhood (line 5). If this is the case, it searches among its neighbors the agent or the coalition with the best reputation (s_j and CS_j, lines 6 and 7) to join them depending on the value of CS_j with respect to s_j (lines 9 and 11, respectively).

Algorithm 3 Rules for coalition formation and independence

```
function ChangeCoalition
    Coa_i = GetCoalition(a_i)
    if ((Coa_i ≠ ∅) & a_i.Isolated(Neighbors, Coa_i)) then
        a_i.GetIndependence()
    else if WorstPayoff(Neighbors, P_i) then
        (s_j, a_j) = BestIndScore(IndepNeighbors)
        (CS_j, Coa_j) = BestCoaScore(CoaNeighbors)
        if CS_j ≥ s_j then
            JoinCoalition(Coa_j)
        else
            CreateCoalition(a_i, a_j)
end function
```

3.4 Changing the Strategy

At the beginning of the game, each agent is randomly assigned a strategy. However, depending on the payoffs it is obtaining, it may change it in order to increase its benefits. When agents are neighbors (directly connected in the network), we consider that they know each other's image score, as well as their payoff and their strategy in the previous game. Thus we assume that agents have access to local information about reputation, payoff and strategies from their neighbors, since they can directly observe them. With that information, an agent changes its strategy to copy the one with the highest payoff in its neighborhood, if higher than its own payoff.

3.5 Network Topologies

We place agents in a complex network since they provide a realistic model of the topological features found in many nature, social and technological networks [19, 18, 23] (e.g., computer and social networks). For our experiments, we focus on small-world and scale-free network topologies, since they model the most common networks appearing in our human societies and in nature.

- Small-world networks: They model real world complex systems, as neural networks, food webs, scientific collaboration networks, and computer networks [13]. These networks present the small-world phenomenon, in which nodes have small neighborhoods, but yet it is possible to reach any other node in a small number of hops. This type of networks are highly-clustered (i.e., have a high clustering coefficient). Formally, we note them as $W_V^{k;p}$, where V is the number of nodes, k the average connectivity, i.e., the average size of the node's neighborhood, and p the rewiring probability. We used the Watts & Strogatz model [23] to generate these networks.

- Scale-free networks: They model real-world networks, as the world-wide web [1], the Internet, and some biological networks [14]. These networks are characterized by having a few nodes acting as highly-connected hubs, while the rest of them have a low connectivity degree. Scale-free networks are low-clustered networks. Formally we denote them as $S_V^{k;-\gamma}$, where V is the number of nodes, and its degree distribution is given by $P(k) \sim k^{-\gamma}$, i.e., the probability $P(k)$ that a node in the network connects with k other nodes is roughly proportional to $k^{-\gamma}$. We used the Barabasi-Albert algorithm [19] to generate these networks.

3.6 Rewiring

In most real-world network interactions, relationships are not static, i.e., agents can change the individuals that they are linked to. We denote this change in the network topology as rewiring. By using rewiring agents can modify their neighborhood if they are not satisfied with their neighbors.

As a difference with [8], where one agent is randomly chosen to change its neighbors, in our model, we specify a neighborhood measure of satisfaction to decide if an agent wishes to change it or not. In Eq. 2 we define the probability of rewiring for an agent i, which depends on the aggregate image score of all the neighbors, i.e., it depends on the average neighborhood reputation.

$$p_i^{rew} = \frac{(10 - \frac{\sum_{j=1}^{F}(s_j+5)}{F})}{10} \tag{2}$$

where s_j is the image score of each of the neighbors of a_i and F is the number of neighbors (friends) that the agent a_i has. Observe that $s_j \in [-5, 5]$, thus the maximum difference between scores is 10. Once agent a_i computes this probability, then it samples a Bernoulli distribution to decide if rewiring or not (Algorithm 1, line 10). If agent a_i decides to rewire, then it leaves its neighbor with the lowest image score, and joins the one with the highest one in its coalition. The reason for this is that, as we stated above, we consider that coalitions are communities that share reputation information, so agents can benefit from it to change their neighbors. We point out that this rewiring procedure only happens if the agent with the lowest image score does not become isolated, i.e., we do not allow disconnected nodes in our network.

4. EXPERIMENTS

In this section, we present the performance of our mechanism, using the final strategy selected by the agents, after the simulation has converged, as a measure of the cooperation level achieved by the population. Firstly, in Sect. 4.1 we present the empirical setting for our experiments. Secondly, in Sect. 4.2 we analyze how our mechanism of coalitions and rewiring allows for the emergence of cooperation. Finally, in Sect. 4.3, we analyze the differences on results depending on the initial topology.

4.1 Experimental Settings

We have performed a in-depth experimental study. However, due to space constraints, we only show the most relevant ones. Thus in these experiments we perform simulations in which the number of agents N is set to 400. Each run is composed of a set of iterations in which agents repeatedly play the donation game. The number of iterations varies in each particular run depending on the simulation convergence and stability. We consider convergence when there are no changes in the strategy of the agents during ten consecutive iterations. Finally, the parameters used for building the networks are $W_{400}^{5;0.1}$ and $S_{400}^{5;-2}$.

4.2 Emergence of Cooperation

In this section we analyze the effects in cooperation using coalitions and the rewiring mechanism in the networked donation game.

Firstly, in Fig. 1 we see the results of a typical simulation when we do not use coalitions nor rewiring. In the histogram we represent the percentage of agents with a certain strategy when the simulation has converged. We see that all the agents end up playing $k \geq 0$. This means that agents lean toward playing defective (remember that $k = 6$ means that an agent defects independently to the other agent image score).

Secondly, we allow agents to use only rewiring to change their neighborhood. We have observed that both for scale-free and small-world networks, the results are similar to the case where we do not use coalitions nor rewiring, since $k > 0$ for all agents (we do not

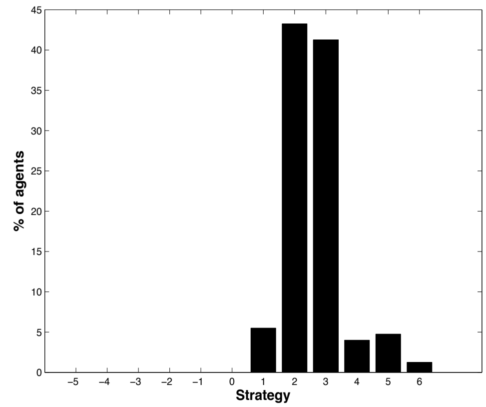

Figure 1: Percentage of agents' strategies with no coalitions, no rewiring, in a scale-free network.

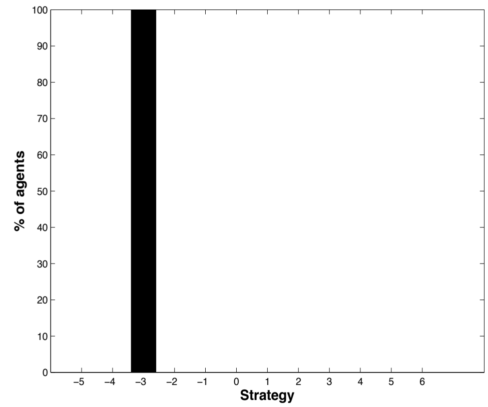

(a) Percentage of agents' strategies. All agents end with $k = -3$.

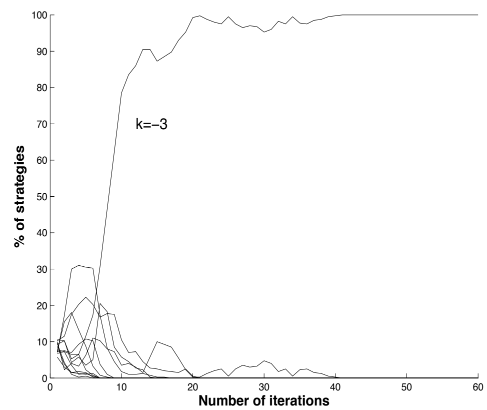

(b) Evolution of strategies on time. All agents end with $k = -3$.

Figure 2: Strategies in scale-free without rewiring. Convergence to cooperative strategy k=-3.

depict it since it is similar to Fig. 1). This differs to the results obtained by Fu et al. [8], where they successfully use rewiring to improve cooperation among agents. However, here we propose a different environment, where even if agents are connected to others, they can play with any agent in the population. In fact, as they have no information about other agents' reputation, since there are no coalitions for information sharing, the rewiring is done randomly, and it might even worsen their neighborhood.

Thirdly, we endow agents only with our coalition formation mechanism, but without allowing them to rewire. We find that in scale-free networks, allowing them to join coalitions is enough to achieve cooperation. In Fig. 2 we present the final percentage of agents per strategy when convergence is reached, as well as the evolution of strategies in time. We see that in this case all agents converge to strategy $k = -3$. Moreover, we have observed that in different simulations the results vary from one strategy to other, but being $k \leq 0$ in all cases. However, when using small-world networks, we have observed that agents converge to a single strategy, which is not cooperative ($k > 0$). In Fig. 3a we show an example where all agents converge to $k = 2$. Moreover, in Fig. 3b we see the evolution of strategies on time, where to ease the display, we only name the two strategies that survive longer. We see that the strategies $k = -3$ and $k = 2$ compete to dominate the divided population. However, in the end the non-cooperative strategy prevails. This pattern is repeated in different simulations, but with different strategies arising.

Now, we study if cooperation is improved when we add rewiring to the coalition formation mechanism. In Fig. 4 we present the percentage of strategies after one simulation when using our mechanism, both starting with an initial scale-free and small-world topology (Fig. 4a and Fig. 4b, respectively). We see that in both cases, all agents end up a using a cooperative strategy ($k \leq 0$). Moreover, not only agents converge to a cooperative strategy, but we have observed that in every simulation, all agents converge to the same cooperative strategy (but different in sucessive runs).

In order to see that our combined mechanism allows only cooperative strategies to arise ($k \leq 0$), in Fig. 5 we present the results for ten different simulations. We represent the percentage simulations in which all agents end with each strategy, keeping in mind that in each simulation, all agents converge to the same strategy. We only present it for a scale-free initial topology, since results for small-world are similar. We see that in 30% of the simulations, all agents converge to $k = -4$, while other three cooperative strategies ($k = -5, -3, -2$) appear in 20% of the simulations each. Only 10% of agents use $k = 0$, which is the most discriminating among cooperative strategies.

Thus, we found that by using coalitions and rewiring cooperation emerges. This happens mainly by two reasons: firstly, because one single super coalition is formed (see Fig. 6, where we see the evolution of the number of coalitions along a simulation). As an agent has information not only about its neighbors, but also about its coalition mates, this results in agents having more information about the image score of the whole population as the simulation evolves. Secondly, as an agent can change its neighborhood, it can discover and join other agents with higher image score. This allows an agent to donate with higher probability, also increasing its image score, and therefore its chances for obtaining a donation next time it becomes a recipient.

We have further investigated the effects of adding rewiring to the coalition formation mechanism. For this, we used Pajek [5], which is a tool for analysis and visualization of large networks. We have observed that in scale-free networks, hubs (agents with higher number of connections) have a strong influence over the rest of agents, and also more information than them. This eases the process of convergence to one single coalition, where all the agents use the same cooperative strategy. This happens even when we only use our coalition formation mechanism. But, when we introduce rewiring the process of convergence is even faster. This happens

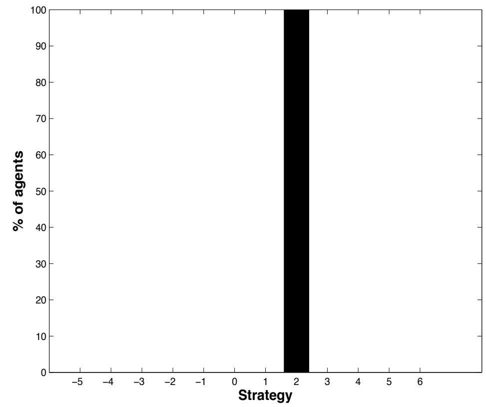

(a) Percentage of agents' strategies.

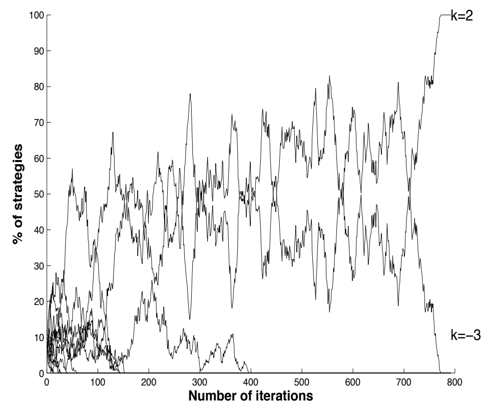

(b) Evolution of strategies along time.

Figure 3: Strategies in small-world, with coalitions but without rewiring. All agents end with $k = 2$.

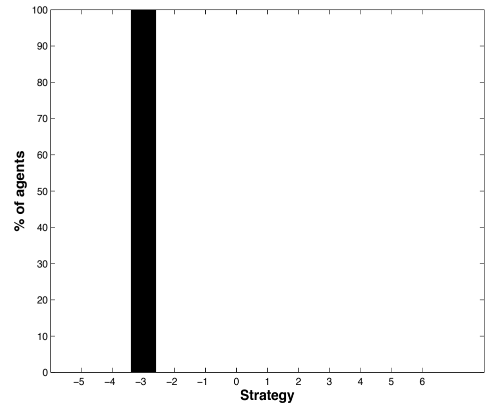

(a) Percentage of agents' strategies in a simulation with a scale-free topology. All agents end with $k = -3$.

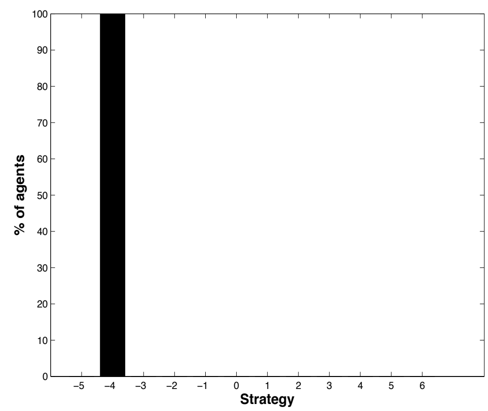

(b) Percentage of agents' strategies in a simulation with a small-world topology. All agents end with $k = -4$.

Figure 4: Strategies obtained after two simulations with scale-free and small-world initial topologies, using coalitions and rewiring.

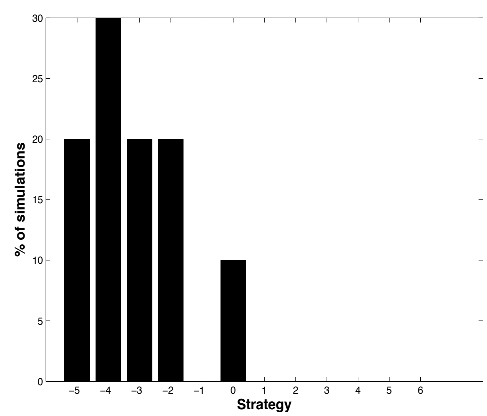

Figure 5: Average percentage of strategies after ten simulations.

because, as we allow agents to choose their neighbors, hubs are the most successful ones, making their own influence even higher, and also the influence of the coalition they belong to.

In the case of small-world networks, all agents have more or less the same number of connections, meaning that all agents have a similar level of information at the beginning. However, when we add rewiring, agents start to create influence groups composed by some agents which have higher connections than the others. In Fig. 7 we depict an example of a final configuration when we start with a small-world network topology (here we used only 25 agents to ease its display). We see that agents self-reorganize in a structure, where some of them have much more links than the others. Thus, as in the scale-free case, bigger and more influential coalitions (regarding their image score and size) are formed.

Finally, we compare our results with the ones obtained in [17], which is the basis for our work. The comparison is not easy as that paper presents a panmictic scenario, and genetics are used to evolve the most popular strategies in the population. In the scenario with public image score, the obtained strategy was $k = 0$. But in a second scenario where agents have a limited view of others' image score, agents tend to be defective ($k > 0$). In our case, our coalitions and rewiring mechanism allows to achieve cooperation even in this second scenario with limited information.

4.3 Topology Influence

In previous section, we have presented how regardless of the initial network configuration, all agents converge to the same cooperative strategy, with $k \leq 0$ (Fig. 4 and Fig. 5), and one single super coalition emerges (Fig. 6). However, we have noticed differences between scale-free and small-world about how they reach convergence. Salazar et al. [20] also addressed this issue, although in a

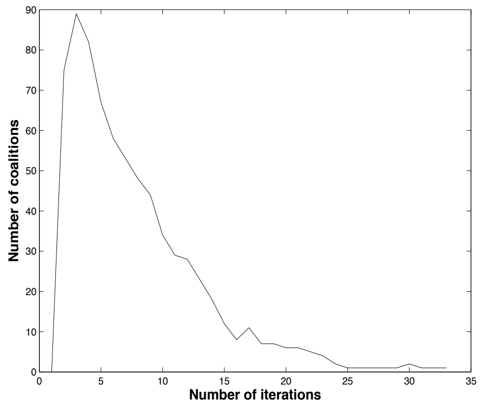

(a) Scale-free.

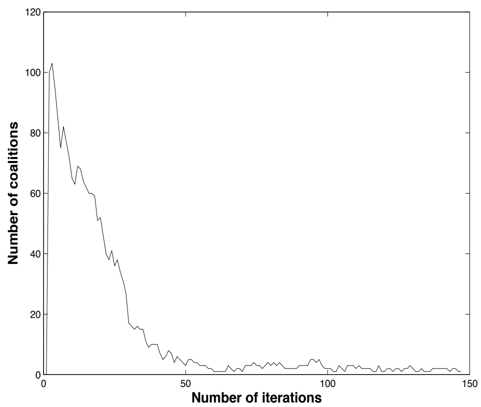

(b) Small-world.

Figure 6: Evolution of the number of coalitions along the iterations.

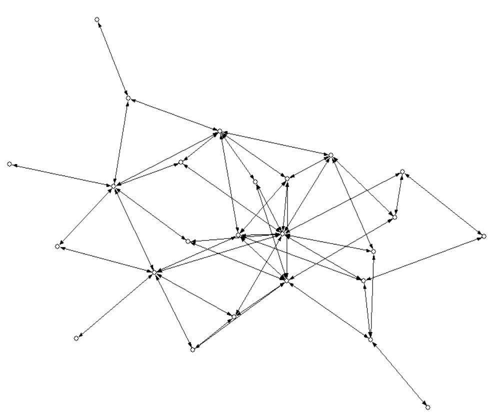

Figure 7: Final topology, after starting with a small-world network with 25 agents, using coalitions and rewiring.

different problem, and with a different focus. Now we investigate the reasons for those differences in our scenario.

We have noticed that the time required for the convergence varies depending on the topology. In Fig. 8 we see the evolution of strategies along iterations. It is noticeable that to reach cooperative convergence, starting with a scale-free topology the convergence is much faster.

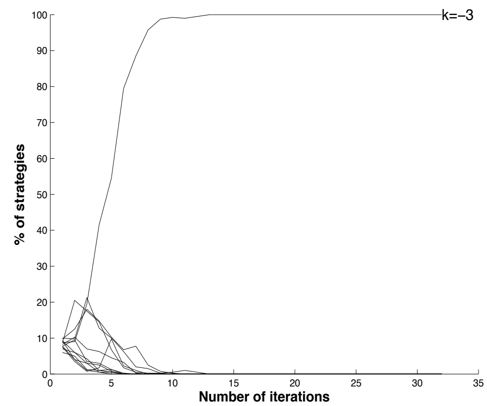

(a) Scale-free.

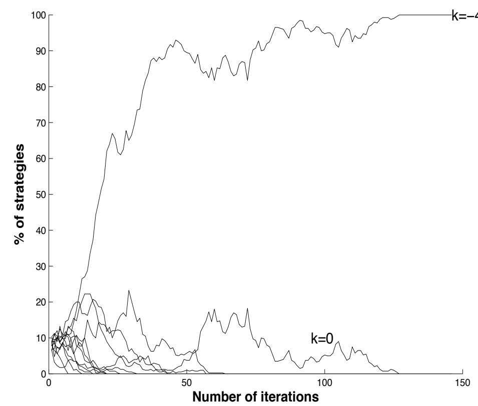

(b) Small-world.

Figure 8: Evolution of agents' strategy along the iterations.

On the one hand, the faster convergence with scale-free is due to the strong influence that hub agents have over the rest of the population. As a hub has a considerable higher number of neighbors than the rest of the agents, it has more information to play initially (as agents know the image score of their neighbors), which increases its chances to get higher benefits. This puts them in an excellent influence position, since as they are the ones with highest benefits, other agents copy their strategy. Moreover, as there are several agents with only one link to the hubs, they promptly join the hub to form a coalition, thus less and bigger coalitions are formed faster. Besides, as they belong to bigger coalitions, and have more neighbors, those hubs are very popular to rewire to them. This causes that they increase even more their individual and coalition influences.

On the other hand, in small-world networks each agent has a similar number of neighbors, so all the agents have more or less the same level of influence. Hence, this explains why multiple coalitions coexist longer (Fig. 6b). Therefore the path to form one single coalition, and converge to the same strategy, is slower. However, with the use of rewiring, agents with highest score start having more neighbors than the others, which results in more influence (see Fig. 7). Afterwards, more agents imitate them, and the coalitions they belong start to grow faster, allowing to finally reach a single cooperative strategy.

5. CONCLUSIONS

In this paper, we have presented a new scenario where agents are connected in a network, but where any agent may interact with any other in the population. Agents' neighbors are their close related contacts from which agents obtain information, as strategy, payoff and reputation. Thus the neighborhood of an agent models a friendly exchange of information. However, in real world, apart from having a set of direct contacts, people usually belong to several clubs, associations, organizations, or groups in general. We model this second set of contacts with coalitions, as a way that agents may share information about the environment where they play. Contrary to most approaches, we do not consider coalitions as a way for agents to coordinate and act together, but as a way for them to share some information. Moreover, we also introduce a measure of the coalition image score, in order to decide which one an agent should join. Finally, we include a rewiring mechanism using the neighbors' reputation as a criteria to change their worst social links, and coalitions information as a way to rewire to the best coalition members.

We have confirmed that the use of coalitions and rewiring indeed improves cooperation when we play the donation game in our social scenario. Moreover, we have analyzed the differences between the results obtained when we use a scale-free or a small-world topology. In our experiments, firstly we determined that only using rewiring does not allow cooperation to emerge. This is because rewiring is done mainly randomly to any other agent, which can even worsen the neighborhood. Secondly, we determined that using our coalition formation mechanism only, cooperation emerges only in the case of scale-free networks. However, in small-world networks, we observed that the use of coalitions is not enough to achieve convergence to a cooperative strategy. The reason is that, in scale-free networks, hubs have a strong influence; allowing to create bigger coalitions in less time, and speeding up the appearance of cooperation. Finally, when using rewiring together with coalitions, both in scale-free and small-world networks; all agents in the population converge to a cooperative strategy. In the case of scale-free, the convergence to a cooperative strategy is faster, since again hubs speed up the convergence process. Thus, we have seen the positive effects that grouping and social networking have over cooperation in complex networks with indirect reciprocity.

As future work, we plan to study how much improvement can be obtained avoiding non-cooperative agents. We also plan to study mechanisms to cope with malicious agents, that may not share their real image score with the rest of the coalition. We will also build new simulations using more real-world topologies, together with in-depth research about the changes in the networks along the simulations. Finally, we will also study the influence of allowing agents to belong to multiple coalitions at a time.

Acknowledgments

The first author thanks the grant Formación de Profesorado Universitario (FPU), reference AP2010-1742. Ana Peleteiro and Juan C. Burguillo have been supported by the European Regional Development Fund (ERDF) and the Galician Regional Government under grant CN 2012/260 (Consolidation of Research Units: AtlantTIC) and by Mineco, Spain, under grant CALM (TEC2010-21405-C02-01).

6. REFERENCES

[1] Adamic, L.A. and Huberman, B.A. Power-Law distribution of the world wide web. *Science*, 287(5461):2115, Mar. 2000.

[2] R. Axelrod. *The Evolution of Cooperation*. Basic Books, New York, 1984.

[3] R. Axelrod. *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*. 1st printing edition, Aug. 1997.

[4] R. Axelrod and W. D. Hamilton. The evolution of cooperation. *Science*, 211:1390–1396, 1981.

[5] V. Batagelj and A. Mrvar. Pajek - analysis and visualization of large networks. In *Graph Drawing Software*, volume 2265, pages 77–103, 2003.

[6] K. Chellapilla and D. B. Fogel. Evolution, neural networks, games, and intelligence. *Proceedings of the IEEE*, 87(9):1471–1496, Sept. 1999.

[7] S. Y. Chong and X. Yao. More choices and reputation in multi-agent interactions. *IEEE Transactions on Evolutionary Computation*, 11(6):689–711, 2007.

[8] F. Fu, C. Hauert, M. A. Nowak, and L. Wang. Reputation-based partner choice promotes cooperation in social networks. *Physical Review E*, 78(2):026117+, Aug. 2008.

[9] F. Fu, T. Wu, and L. Wang. Partner switching stabilizes cooperation in coevolutionary prisoner's dilemma. *Phys. Rev. E*, 79(3):036101, 2009.

[10] N. Griffiths and M. Luck. Changing neighbours: improving tag-based cooperation. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1*, AAMAS '10, pages 249–256, 2010.

[11] G. Hardin. The tragedy of the commons. *Science*, 162:1243–1248, 1968.

[12] J. C. Burguillo . A memetic framework for describing and simulating spatial prisoner's dilemma with coalition formation. In *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems*, pages 441–448, 2009.

[13] S. Kniesburges, A. Koutsopoulos, and C. Scheideler. A self-stabilization process for small-world networks. In *IPDPS*, pages 1261–1271. IEEE Computer Society, 2012.

[14] M. E. J. Newman. The structure and function of complex networks. *SIAM Review*, 45(2):167–256, 2003.

[15] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, Cambridge, UK, 2007.

[16] M. A. Nowak and K. Sigmund. The dynamics of indirect reciprocity. *Journal of Theoretical Biology*, 194:561–574, 1998.

[17] M. A. Nowak and K. Sigmund. Evolution of indirect reciprocity by image scoring. *Nature*, 393:573–577, 1998.

[18] R. Pastor-Satorras and A. Vespignani. Epidemic dynamics and endemic states in complex networks. *PHYS.REV.E*, 63:066117, 2001.

[19] A. Reka and A. Barabási. Statistical mechanics of complex networks. *Rev. Mod. Phys.*, 74:47–97, June 2002.

[20] N. Salazar, J. A. Rodriguez-Aguilar, J. L. Arcos, A. Peleteiro, and J. C. Burguillo-Rial. Emerging cooperation on complex networks. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, AAMAS '11, pages 669–676, 2011.

[21] T. Sandholm, K. Larson, M. Andersson, O. Shehory, and F. Tohmé. Coalition structure generation with worst case guarantees. *Artif. Intell.*, 111(1-2):209–238, July 1999.

[22] O. Shehory and S. Kraus. Task allocation via coalition formation among autonomous agents. In *Proceedings of the 14th international joint conference on Artificial intelligence - Volume 1*, IJCAI'95, pages 655–661, 1995.

[23] D. J. Watts and S. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, June 1998.