Mind Soc
DOI 10.1007/s11299-009-0061-6

Hybrid models for achieving and maintaining cooperative symbiotic groups

İlker Yıldırım · Pınar Yolum

Received: 10 November 2008 / Accepted: 15 May 2009

Abstract Societies are composed of groups that interact. Symbiotic groups are those in which agents complement each other in resources that they have in excess. Symbiotic groups are useful especially when the resources in an environment are distributed unevenly, because they enable agents to trade resources easily. However, for trading to happen successfully, agents in symbiotic groups need to cooperate, i.e., they should be willing to donate resources when appropriate. Similarly, if some agents in a symbiotic group are defectors, they should be identified by others and eliminated from the group for the well-being of the remaining agents. Accordingly, we first study Edmonds' tag-based model of symbiotic groups to understand the lifespan of symbiotic groups (e.g., why some groups live shorter than others). Then, we enhance Edmonds' model by adding the capability of reciprocal interactions to agents, thus achieving a hybrid model. We capture reciprocity in three different models and study their effects on the elimination of defectors in symbiotic groups. Our experimental results show that the groups that are built with the proposed hybrid model can eliminate more defectors and earlier than tag-based models. Further, the hybrid approach can generate symbiotic groups more effectively and efficiently.

Keywords Cooperation · Reciprocity · Tags · Emergent behavior

Electronic supplementary material The online version of this article (doi:10.1007/s11299-009-0061-6) contains supplementary material, which is available to authorized users.

İ. Yıldırım (✉) · P. Yolum
Department of Computer Engineering, Boğaziçi University,
Bebek, 34342 Istanbul, Turkey
e-mail: ilker.yildirim@boun.edu.tr

P. Yolum
e-mail: pinar.yolum@boun.edu.tr

Published online: 16 June 2009

1 Introduction

In many real life settings, resources are distributed unevenly among individuals. Since individuals may need different amounts of different resources, it is beneficial for them to form groups with those who can contribute to their resource pool. Symbiotic groups are groups in which individuals complement each other in resources that they have in excess. By trading excess resources, they can help each other maintain a balance of resources. However, not all individuals in such groups are equally cooperative. That is, some individuals may prefer not to share their resources, but instead only to receive donations from others. While such defective behavior is useful for the individual itself, it may obstruct others to collect necessary resources to continue their living. Hence, it is crucial for individuals to exist in a *cooperative* symbiotic group that help each other. This paper studies how such cooperative symbiotic groups can be formed.

Our study follows multiagent-based modeling and simulation (MABS) approach. This approach models each individual in the system as an agent and simulates their actions and interactions with each other. Multiagent-based simulation offers strong insights about complex evolutionary processes as well as cooperation. By careful and comprehensive application of MABS, one can study more realistic settings in an experimental setup (Hales et al. 2003; Norling et al. 2008). Simulation of such settings serves as a natural tool to observe emergent phenomena in a complex adaptive system (Tang et al. 2006).

In our model, individuals are represented as software agents that can perceive their environment, reason about their actions, and act accordingly. Further, the agents are autonomous and heterogeneous (Huhns and Singh 1998). Because the agents are autonomous, there is no guarantee that they will help others. Because they are heterogeneous, they may have different resources, or capabilities. These properties match human societies well and thus supports our choice for multiagent-based modeling. To understand cooperation, we study two leading techniques: one based on *tags*, the other based on *reciprocity*. Put simply, tag-based cooperation approaches assume that each agent has a public tag and that agents with similar tags are inclined to cooperate with each other. Reciprocity-based cooperation approaches assume that agents keep track of who is helping them and try to cooperate with those that are helpful.

Ideally, one strives for a cooperation model that will allow symbiotic groups to be formed effectively (that allow trading of resources) and efficiently (that is formed with few interactions). After the groups are formed then the groups should live long. That is, the groups should be formed once and exploited for the duration of agents' life time. In order to provide long-lived communities, the defectors that are threat to the life-span of the groups should be eliminated fast.

1.1 Contribution

To provide these properties, we propose to marry a tag-based model Edmonds (2006) with a reciprocity model Sen (1996) and Sen and Dutta (2002). Our proposed hybrid model starts with a tag-based model but also applies reciprocity. Hence, it benefits from the strengths of both tag-based and reciprocity-based approaches. For our hybrid approach, we study three reciprocity models. The reciprocity models are targeted to help agents find other agents who have been helpful before. We analyze our hybrid models in terms of how well and how fast they eliminate defectors. Our analysis shows that the hybrid models can eliminate defectors better than a tag-based or a reciprocity-based approach. Further, we investigate how these reciprocity models affect the formation of symbiotic groups in terms of effectiveness and efficiency and show that our proposed hybrid models perform better than a tag-based or a reciprocity-based model.

1.2 Organization

Section 2 gives the necessary background including Edmonds' model on symbiotic groups. Section 3 develops different models of reciprocity for building communities and shows how these models can be applied in conjunction with a tag-based model. Section 4 studies capabilities of the developed model in terms of how well and how fast it eliminates defectors. Sections 5 and 6 compare the proposed hybrid model with tag-based and reciprocity-based models in terms of effectiveness and efficiency, respectively. Finally, Sect. 7 discusses the literature with comparisons to our work.

2 Edmonds' symbiotic group model

We start with a promising tag-based model of symbiotic groups that was developed by Edmonds (Edmonds 2006). Tags are observable properties of agents that can be used as a heuristic to recognize types of individuals (Holland 1993). A tag of an agent does not necessarily say anything specific about the agent's abilities or its behavior. However, in real life tags are often used by people to build connections. For example, having a green beard may be interpreted as a tag (Keller and Ross 1998). People with green beards may feel part of a group because of this, even though nothing specific about these people's capabilities or behavior is known. Edmonds' model is based on this intuition. The model is used to show how cooperation can arise in groups.

Edmonds models the environment as consisting of agents and different types of resources. Each agent represents an individual with a certain skill, a tag value, a tolerance value, and particular amount of each kind of resource. The *skill* of an agent determines the type of food it can harvest. The *tolerance value* determines how defector an individual is. The *tag value* is the only value that can be observed by other agents. Tag and tolerance values are real values between 0 and 1. The *maximum tolerance value* denotes the largest tolerance value of the agent in a given population. The number of skills that a population will have in the simulation is a parameter that can be set at the beginning of each run. It can be a value between 2 and 6. For example if the number of skills parameter is set to 3, then there are 3 different types of food in the environment and there are agents of 3 different types of skill which can harvest the corresponding type of food.

Each agent needs to possess a minimum amount of each type of resource in order to survive, but one agent can harvest only one type of food. Therefore, agents need to cooperate in order to survive. At every time step, each individual is randomly paired with a number of other individuals. These pairs are used to share resources such that at each time step each individual either makes donations or receives donations. Say agents x and y have been paired. Agent x makes a donation, if two conditions hold: (1) The difference in the tag values of x and y is strictly less than x's tolerance value and (2) x has an excess in any of its resource stores. When the first condition holds, x and y are either similar to each other (the difference between their tag values is small) or x is very generous (its tolerance value is high) such that it shares resources even with those that have significantly different tag values. The second condition is necessary since x cannot share its resources if it does not have them in excess. When these two conditions hold, x gives a share of its resource to y.

At each step, agents that have enough resources reproduce to create offspring. If there is no noise added to the tag and tolerance values, the corresponding values of the offspring are exactly copied from its parent. If there is reproduction noise, then tag and tolerance values of the offspring are set as tag and tolerance values of its parent plus Gaussian noise. Furthermore, the skill of the offspring is either exactly as that of its parents, or it is set randomly regardless of its parent's skill.

Finally an agent can die for one of two reasons: it starves (i.e., doesn't have enough food) or it reaches a maximum age that is defined in the simulation. An agent that cannot obtain resources dies from starvation after a few steps of the simulation.

The evolution of the agent society can be described with the following steps:

- A small "seed" collection of cooperative individuals with similar tags arises initially.
- The cooperative agents grow in number due to their interaction. However, there are some agents who cannot interact with others successfully, either because they do not have enough resources, or their tag values are much different than others, and so on. These agents starve and die.
- Eventually defectors arise in the group. The reason that defectors arise is that after their creation they both harvest on their own and take donations from non-defector agents. Hence, their resources are large so they do not starve and create offspring who are again defectors.
- These defectors do even better than the others in that group since they always receive donations but never donate themselves. Hence they reproduce until they come to dominate that group.
- When the number of defectors is high, the group does not do so well compared to other cooperative groups since there is no one left to supply food. That is, there is little sharing or no sharing and so the group eventually dies.

In Edmonds' Model we observe that the agents form groups but the groups do not live long, since the agents cannot live long. In other words, the groups cannot be maintained since they are infected by defector members of the group, especially by those that enter the system. This means that even if the cooperating agents in the population identify the existing defectors, they cannot immediately identify the newly entering defectors.

Groups cannot live long enough to detect defectors. That means defectors can find peers to receive donations from. For example, assume agent x is a defector. If the size of the group is large, then agent x can be paired with agent y who has excess resource and has never met agent x before. If the difference between their tags is small enough, the donation takes place. Hence, agent x survives even though it is a defector.

If the agents can cooperate to eliminate defectors early on, they can form persistent groups that live long and that are not wiped out by existence of a few defectors. In order to achieve maintained groups for long lived societies, we add reciprocal behavior to agents so that the defectors can be eliminated in a growing group.

3 Modeling reciprocity

The intuition behind reciprocity is that the agents should help those that help them in return (Axelrod 1990). Such a broad definition can be modeled in different ways. We study three variations of reciprocity that differ in the way agents represent each other. In order to model reciprocity, we add memory into each individual. Each agent is associated with a unique identifier so that the agents can be identified uniquely. The memory of each individual contains information about other agents. More specifically, for a particular agent we store the number of donations it has made to others as well as the number of donations it has received from each individual. In addition, for each type of the reciprocity models, reciprocity is activated only after a small, predefined number of cycles elapse.

3.1 Strict individual reciprocity model

The first model (*Strict Individual Model*) treats reciprocity strictly and claims that an agent will only help other agents if they have necessarily been helpful before. Let T_{ij} denote the number of donations agent i made to agent j. We assume that there are no message failures so that the number of donations received by agent i from agent j is equal to T_{ji}. Agent i stores the id of agent j, as well as T_{ij} and T_{ji} in its memory.

If agent i donates some food to agent j, both agent i and agent j update their corresponding values to reflect the donation. To maintain consistent memories for each individual, we update each individual's memory after every death and every birth. The values in agents' memories are used before donations occur. More specifically, each individual checks if $T_{ij} - T_{ji} < \alpha$, where α is a reciprocity threshold and may be varied from simulation to simulation. Donation occur only if $T_{ij} - T_{ji} < \alpha$. If that requirement is not satisfied, then agent i considers j as a defector and does not donate to agent j.

3.2 Relaxed individual reciprocity model

In the *Relaxed Individual Model*, each agent maintains a general opinion about each agent but does not model each agent precisely as in the *Strict Individual Model*. More specifically, each individual marks others as defectors, donating or unknown. At the beginning, each individual marks all others as unknown. After reciprocity is activated, individuals start to mark each other as defector or donating. When a new agent is added to the society, it is marked as unknown by the agents in the population and the new member marks all agents in the population as unknown. The meanings of marking an individual as defector, donating or unknown are as follows:

- *Donating*: Agent x marks agent y as donating if x and y are paired and y donates to x.
- *Defector*: Agent x marks agent y as defector, when y does not donate to x when there is a pairing between them. If at a later interaction, y donates to x, then x relabels y as donating.
- *Unknown*: Agent x marks agent y as unknown if x has no prior information about agent y.

As in Edmonds' Model, x generates a random pair pool. x can remove any agent y from its pair pool if the difference between the tag values of y and x are larger than the tolerance value of x. If y was in pair pool of x but was removed by x from pair pool, then y marks x as defector. But, if x does not remove y, and donates from its excess resources to y, than y marks x as donating.

3.3 Group reciprocity model

Contrary to the previous two models, now we assume that agents in the same group share information among the group. Specifically, agents share information about the number of donations each makes. The agents share this information by keeping a common table that records the donations. Whenever an agent receives or makes a donation, it updates this table to reflect the changes in the donations.

Assume that agent x and agent y are in the same pair and agent x donates to agent y. When this donation occurs, we increment number of donations made by x by one. Simultaneously, we increment number of donations received by agent y by one. Agent x checks the number of donations received by and made by agent y. Let R_x denote the number of donations received by agent x and M_x denote the number of donations made by agent x. If at the next time point, agent y is paired with agent x, agent y requires that $R_x - C < M_x$ where C is a constant. That is, agent y expects x to have received at most C more donations than the number of donations it has made. In our simulations we selected C as 5.

We study these models in static populations, where new agents do not enter the population. We apply a procedure of *isolation* to generate these static populations. That is, we stop accepting new agents into the population when two conditions hold: (1) The size of population reaches Φ and (2) the difference between sizes of the largest and the smallest skill groups is smaller than Δ. The first condition is necessary to be able to evaluate the groups sufficiently. If the population is too small, then the results may not be representative. The second condition is necessary to reach a balanced population. For instance, if the number of agents with skill 0 is the largest, and the size of the agents with skill 1 is the smallest, we expect the difference between these two sizes to be smaller than Δ so that the population can still continue to exchange resources and be able to survive. Φ and Δ are tunable parameters in the model and we vary them in our simulations to show their effects on the results.

4 Eliminating defectors

We study how well different models generate cooperative symbiotic groups. Our first criterion is how successfully each model eliminates defectors. We measure this through the tolerance value of starving agents. The tolerance value of an agent shows its willingness to cooperate with agents that are substantially different than itself. That is, an agent with a small tolerance value will only cooperate with agents that are strictly similar to itself. Intuitively, high tolerance values correspond to cooperative agents and low tolerance values correspond to defectors. Thus, if the agents that starve has low tolerance values, then they are actually defectors. If their average tolerance value is high, then cooperative agents starve (Sect. 4.1). Our second criterion is how *fast* defectors are eliminated. We measure this by inspecting the simulations at different time intervals (Sect. 4.2).

4.1 Success in eliminating defectors

In the isolated models, we observe the tolerance values of the agents that die from starvation in a cooperative group. To retrieve these values, we first let the population form symbiotic groups (for four time steps). In these time steps, first agents who cannot cooperate to begin with (because they do not have any resources) die. The remaining agents become part of a group. Then, in the remaining time steps, we analyze the tolerance values of agents that manage to become part of a cooperative group but later die because of starvation. If the tolerance values of such agents are actually low, we can conclude that the defectors are eliminated.

Table 1 shows the average tolerance values of agents that starve in 20 runs with 5,000 steps when the population size is above 150 and the number of food types is 2, 3 or 4. The *MaxTol* parameter shows the maximum tolerance value for the simulation. The values in the parentheses are the number of runs in which isolation

Table 1 Average tolerance values of agents that starve under isolation of 20 runs with 5,000 steps

MaxTol = 0.05, Φ = 150, Δ = 60	2	3	4
Edmonds' model	0.025 (0)	0.029 (0)	0.030 (8)
Strict individual model	0.013 (0)	0.011 (0)	0.012 (9)
Relaxed individual model	0.019 (0)	0.015 (0)	0.026 (11)
Group model	0.018 (0)	0.021 (0)	0.031 (11)

does not start, meaning cooperation does not occur. Note that in all of our simulations we use isolated populations including Edmonds' Model.

In Edmonds' Model where there is no reciprocity but only a tag mechanism, we observe that the average tolerance values of the starving agents is 0.025 when the food type is set to 2. This number increases as the number of food types increase. This results imply that among starving agents, defectors are rarely seen, and mostly non-defectors die. For the remaining runs (where agents starve), the tolerance values of the starving agents are 0.029 and 0.030, respectively.

For the *Strict Individual Model*, the values for average tolerance are much smaller (around 0.013). This shows that the agents that starve are mostly the defectors. Hence the *Strict Individual Model* is successful in eliminating the defectors agents from the groups. This confirms our result that agents can detect and eliminate the defectors.

Next, we repeat the same experiments for *Relaxed Individual Model*. We again study the average tolerance values of the individuals that die from starvation. Again, looking at Table 1 we see that the average tolerance values of the agents die from starvation after isolation started are always smaller than the results of Edmonds' Model. However, compared to the *Strict Individual Model*, the number of populations in which no starvation takes place is higher. This shows that *Relaxed Individual Model* eliminates defectors more than Edmonds' Model but less than the *Strict Individual Model*.

Lastly, we repeat the same experiments for the *Group Model* and study the average of tolerance values of the individuals die from starvation. As Table 1 indicates, for food types 2 and 3, the *Group Model* performs better than the *Edmonds' Model* but worse than the *Relaxed Individual Model*. However, when food type is set 4, this model performs slightly worse than Edmonds' Model and cannot eliminate defectors as well as the other two reciprocity models.

The results presented here can be intuitively explained by the inherent characteristics of the four models. The Edmonds' Model is only tag-based. It encourages cooperation among agents with similar tags but does not create incentives to be cooperative. Hence, defectors can survive in the symbiotic groups. The addition of reciprocity on top of the tag-based model establishes an incentive layer for the agents to help cooperative agents that have been helpful before, but not help uncooperative agents. Thus, defectors are eliminated. Among the three reciprocity models, the *Strict Individual Model* models agents more accurately than the *Relaxed Individual Model*, which models agents more accurately than the *Group Model*. Our results are compatible with this. That is, the more accurate a reciprocity model, the better it eliminates defectors.

4.2 Time-line of eliminating defectors

Next, we need to show that these eliminations of defectors take place quickly and that the groups stabilize in a short time. This is important because after a short elimination phase, the groups should live consistently for a long time. If the defectors are not eliminated fast, then during the elimination phase defectors may benefit unnecessarily from other agents.

Table 2 Percentage of remaining defectors after isolation started of 20 runs with 5,000 steps

MaxTol = 0.05, $\Phi = 150$, $\Delta = 60$, no. of foodtypes = 2 (%)	0%	1%	2%	3%	10%	50%	90%
Edmonds' model	0.16	0.16	0.16	0.16	0.16	0.16	0.15
Strict individual model	0.15	0.14	0.14	0.14	0.07	0.07	0.07
Relaxed individual model	0.1	0.09	0.06	0.06	0.06	0.06	0.06
Group model	0.11	0.11	0.11	0.11	0.11	0.11	0.11

We compare the populations in terms of when defectors die. Again, we made 20 runs with 5,000 steps. For each run the time after isolation started is divided into 100. And we give the percentage of the defectors after just isolation started: after 1%, after 2%, after 3%, after 10%, after 50% and after 90% time passed from the step isolation started. This way we can observe when the defectors are removed from the groups and when groups stabilize. We report the averages of the 20 runs.

Table 2 shows the percentage of defectors for different models. To show the percentages, we take agents with tolerance values smaller than the half of the average of the entire population as defector. Since the tolerance values follow a normal distribution, this helps to set the number of defectors in all the populations as approximately equal. We see that there is no significant decrease in the percentage of the defectors in Edmonds' Model or in the *Group Model*, whereas percentage of the defectors decreases fast in the *Strict Individual Model* and *Relaxed Individual Model*. For these two reciprocity models, after a few steps (at most 10% of the simulations), many defectors are eliminated.

In profiles in Figs. 1, 2, and 3, each agent is shown by a line. The center of the line corresponds to the tag of that agent, and length of this line is the tolerance of the same agent. And the darkness of the line is the skill of the agent and all agents move one unit upwards as they age by each simulation step.

In the first profile in Fig. 1, we see that two possible groupings are likely to emerge, one is around tag 0.5 and the other around tag 0.8. But in the following steps, the group around tag value 0.5 dominates in a winner-take-all kind of dynamics. We observe the group as isolated in the profile in Fig. 2. Although groups emerged in Edmonds' model are prone to selfish agents, and therefore they get wiped out as they emerge, our agents in our hybrid models achieve to maintain the groups very long, i.e. till the end of the simulation, as seen in Fig. 3.

5 Effectiveness of symbiotic groups

Effectiveness of a population denotes how well the agents in the population cooperate. The intuition is that the more agents cooperate, the more effective they can work together. In the present setup, the cooperation denotes how much agents donate to each other. Thus, we measure the effectiveness of a population by its average donation rate. If agents have high average donation rates, then they are more willing to share their resources with each other; thereby leading to high cooperativeness in the population.

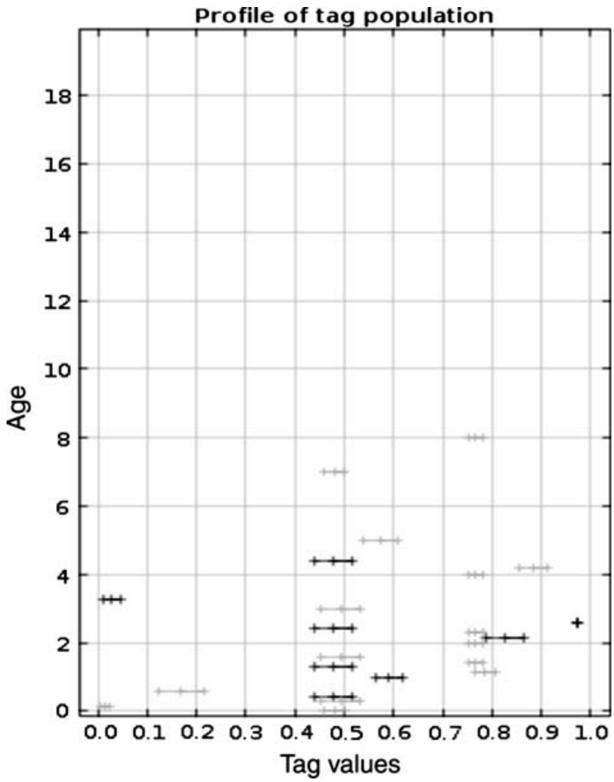

Fig. 1 The initial state of the agent population

We made 20 runs of 5,000 steps, the minimum size of population to start isolation, $\Phi = 150$, maximum difference between minimum and maximum sized sub-populations to start isolation, $\Delta = 60$, and the number of food types is 2.

Table 3 gives the the average donations for the Edmonds' Model, tagged reciprocity models and the non-tag reciprocity models. What we mean by non-tag reciprocity model is the model in which we do not apply tag mechanism but we apply the reciprocity model. By that way groups can emerge between every agent regardless of their tag values. We observe that Edmonds' Model achieves high effectiveness rates. The only two models that achieve better effectiveness are the *Tag* and *Non-Tag Relaxed Individual Models*. This shows that *Relaxed Individual Model* by itself is capable of generating effective groups. Interestingly, the other two reciprocity models do not contribute much to the tag-based model.

When the reciprocity model is the *Group Model*, agents are more tolerant to each other in terms of cooperation. That is, agent *A* may still help agent *B* even if agent *B* has not been particularly helpful to *A* before (but has been helpful to some other agents). As a result, this time we are relaxing our requirement of cooperativeness in the opposite direction. Hence, populations of the *Group Model* cannot achieve as high cooperation values as the *Relaxed Individual Model*.

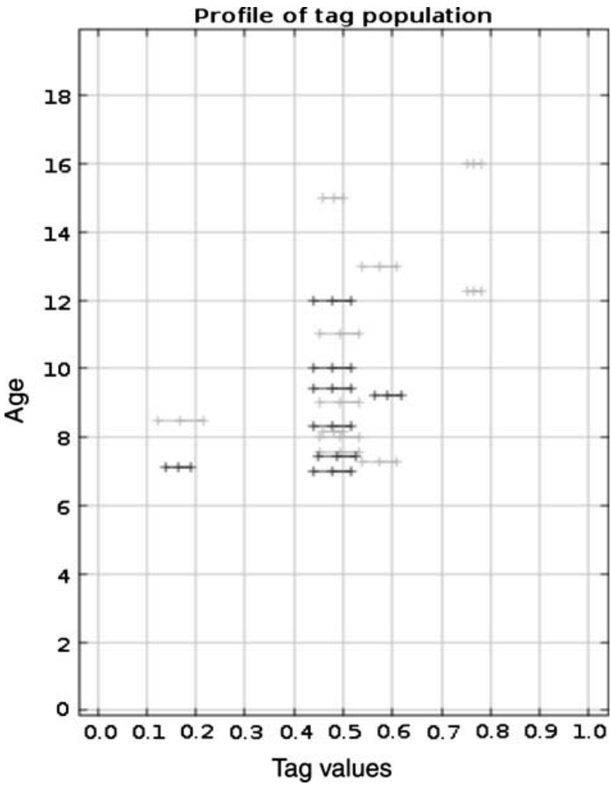

Fig. 2 A group is formed and it dominates others

On the other hand, *Strict Individual Model* is too strict about who will help whom. If an agent is defector (even if it has a high donation rate), it is eliminated from the population. Such agents help very few agents but when they do they donate more. Since such individuals are also eliminated from the population, *Strict Individual Model* is not as effective as the *Relaxed Individual Model*.

6 Efficiency of symbiotic groups

Efficiency of a model denotes how easily it allows generation of groups. For example, if agents need to interact with too many other agents before settling of a group, then the efficiency of the underlying mechanism is low. Ideally, agents should be able to form cooperative groups as quickly as possible (with few interactions).

To measure the efficiency of a model, we count the average number of agents that are contacted at each round. More specifically, at each interaction, an agent considers a set of agents as possible *donation candidates*. For Edmonds' Model, these are the agents whose tag values match. For the non-tag reciprocity models, these agents are those that are not defectors. For the hybrid model, these agents are the ones who have matching tags and are not defectors.

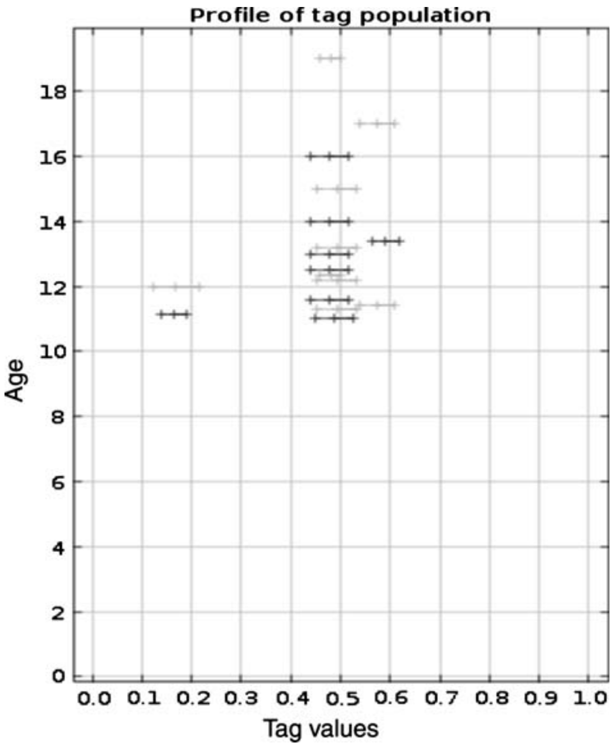

Fig. 3 The group is maintained, and not defected by the defectors

Table 3 Average donations over 20 run with 5,000 steps

$\Phi = 150$, $\Delta = 60$	2	3	4
Edmonds' model	0.881 (0)	0.877 (0)	0.942 (6)
Non-tag strict individual model	0.863 (0)	0.863 (0)	0.863 (0)
Non-tag relaxed individual model	0.999 (0)	0.999 (0)	0.999 (0)
Non-tag group model	0.877 (0)	0.876 (0)	0.877 (0)
Strict individual model	0.740 (0)	0.866 (0)	0.900 (11)
Relaxed individual model	0.890 (0)	0.980 (0)	0.993 (7)
Group model	0.723 (0)	0.802 (0)	0.867 (6)

To compare different models in terms of efficiency, we count the number donation candidates until isolation starts. A more efficient model is expected to achieve groups with a smaller number of average donation candidates.

Table 4 shows the efficiencies for different models. The non-tag reciprocity models achieve the highest rates and thus are clearly inefficient. The main reason for this is that non-tag reciprocity models consider too many agents as possible partners. Edmonds' Model performs better than non-tag reciprocity models. These results are compatible with Riolo's findings (Riolo 1997) about tag and non-tag

Table 4 Average number of donation candidates over 20 runs with 5,000 steps

$\Phi = 150$, $\Delta = 60$	2	3	4
Edmonds' model	2.020 (0)	1.595 (0)	0.453 (5)
Non-tag strict individual model	5.063 (0)	5.407 (0)	5.565 (0)
Non-tag relaxed individual model	5.065 (0)	5.414 (0)	5.623 (0)
Non-tag group model	4.829 (0)	4.989 (0)	5.110 (0)
Strict individual model	1.179 (0)	1.303 (0)	0.565 (5)
Relaxed individual model	1.789 (0)	1.325 (0)	0.464 (0)
Group model	1.651 (0)	1.290 (0)	0.434 (7)

based cooperation models. The hybrid models achieve the highest efficiency rates. All three reciprocity models improve efficiency. Among the three modes, *Group Model* achieves the highest efficiency. This is intuitive because instead of trying all agents individually, this model allows sharing of information among agents. Hence, even with few interactions, agents know who the defectors are. Since *Relaxed Individual Model* requires more interactions, its performance is slightly worse than the *Group Model*. *Strict Individual Model* tries to model each individual accurately and hence requires the most interactions among the three reciprocity models. Accordingly, its performance is lower than the other two hybrid approaches. However, even then its efficiency is comparable to that of Edmonds' model.

7 Discussion

This paper proposes a hybrid cooperation model by combining Edmonds' model on symbiotic groups with reciprocity. Edmonds' model shows how such groups can emerge but does not study the quality of these groups. Our study of Edmonds' model is that groups collapse and are reformed because of the defectors. If defectors can be eliminated from the groups, then the groups can be stable. To achieve this, we have developed different models of reciprocity for donations and applied these models in static populations.

Cooperation exists at all levels in the nature, from social insects to human society. As discussed in Nowak (2006), this observation implies specific mechanisms are necessary for the evolution of cooperation. Among others, two proposed mechanisms are direct and indirect reciprocity. The principle under direct reciprocity is that an individual helps others who have helped her before. For the indirect reciprocity, the principle is that an individual helps others who have been helpful to others. Our *Relaxed and Strict Individual Models* capture the idea of direct reciprocity, because agents always keep track of other agents only by their direct interactions. But, *Group Model* illustrates indirect reciprocity, in which reputation of agents, which is the group level table of interactions, is available to the entire population, and the agent act according to these reputation values.

In this paper, we showed that our reciprocity models can eliminate defectors and can eliminate them fast. Among the three reciprocity models, the *Strict Individual*

Model performed the best for this activity. Also, we showed that hybrid models of tag and reciprocity generate more effective groups in the most efficient way. Our simulation results show that when agents follow the *Relaxed Individual Model*, they are able to achieve significantly high levels of effectiveness when compared to the Edmonds' Model. We observe that when agents follow the *Group Model*, they are able to achieve better efficiency rates than Edmonds' Model.

Our results suggest that a rather demanding model does not directly imply better cooperation, and in fact that the reverse is true. Indeed, although *Strict Individual Model* is very demanding, it performs poorly when compared to *Relaxed Individual Model* which is rather simple. Our findings are in agreement with previous evidence about the computational capacity and efficiency of a multiagent system which relies on fast and frugal decision rules (Gigerenzer et al. 1999; Posada et al. 2006; Posada and Lopez-Paredes 2008).

Our experiments show that agents of low cognitive capacity—by exploiting physical and social environment—can achieve cooperation, not be invaded by defectors, and maintain efficient personal exchange. Various finding in the literature also point in the direction of simplicity. Aktipis shows that a rather simple strategy, namely 'Walk Away' (move after partner defects) can survive against defectors and outperform other relatively complex strategies in the environment (Aktipis 2004). Joyce et al. report similar findings where agents with a relatively simple strategy called MOTH (unconditionally cooperate, but disassociate when partner defects), achieve better than TIT FOR TAT under most conditions (Joyce et al. 2006).

Importantly, we observed that non-tag populations achieve highly effective groups, but with the cost of low efficiency. However tag-based populations not only maintain effective groups, but they do it highly efficient when compared to others. A similar observation can be found in Riolo (1997). Riolo investigates the role of tag mechanism in leading a population to mutual cooperation (reciprocity). He particularly studies the effects of varying mutation pressure in a population, the extent and cost of tag searching (the cost of trials to find someone similar enough), the number of other agents played per generation, and the number of rounds played per game. Similar to our findings, he reported that an increase in exploration enables tag populations to achieve cooperation faster and more often (non-tag like dynamics). But the greater amount of exploration increases instability in the populations. And the instability is far greater for non-tag populations. Furthermore, Riolo compares populations that use tag-mechanism with those that do not. And he also concluded that using tag-mechanism results in shorter time to high cooperation (more efficient) when compared to non-tag mechanism.

Our hybrid models are tag based, in which agents are randomly paired but it is necessary that two agents have close enough tags to cooperate. A different way of modeling whom to cooperate is to assume a network over agents. Using such networks, Paolucci et al. (2006) study an interesting example of non-kin altruism from nature: blood sharing among vampire bats. They investigate the role of grooming networks (analogous to a kind of friendship network) in the formation and maintenance of roosts, which are a social structure within the group. They assume that agents have a fixed ordering among their grooming network. Under this setting, they show that the whole population builds larger and maintained groups, when agents are allowed to switch around subgroups, and not all newcomer requests to subgroups are rejected.

Hales also demonstrates tag processes that are sufficient to produce altruistic behavior, which is not kin related (Hales 2003). Groups of individuals cooperate and evolve to increase group level fitness based on tag similarity. The model is set such that individuals cannot help the same-kin agents directly. Each agent has a skill value, a tag value and a tolerance value. In each step each agent is awarded some number of resources. The type of the resource is assigned such that agents cannot receive the same type of food of its skill. So it must search for the agent with required skill and tag values. Donation requires a recipient to be found with the required skill type and similar enough tag. In his experiments, Hales study dumb and smart search methods for making donations. In the dumb search, an agent selects a random agent, donates if it has similar enough tag value and matching skill. In the smart search, the agent searches for an agent that has similar enough tag value and matching skill, and then selects it. The results show that the smart search populations achieve better average donation values.

Although our hybrid models achieve to maintain groups effectively and efficiently they can be realized in some other way as well. For instance one can alternate Strict Individual Model such that the agents renormalize their number of interactions to decide whether to donate. That change would be interesting, because not all agents are able interact as much as others. And renormalization will avoid effects of variations in terms of number of interactions that the agents are being involved in. One other point is about the Group Model, where agents can access all the information among the population and that information is correct. But agents may not always be truthful, or there may be privacy concerns of agents. That is an interesting future work for us. In our future work, we also plan to study if and how reciprocity can be learned by agents in a group. Further, we plan to incorporate models in which agents have resource constraints, especially in terms of memory.

Acknowledgments We thank Rosaria Conte and the anonymous reviewers for their helpful comments. This work has been partially supported by The Scientific and Technological Research Council of Turkey (TUBITAK) under grant 107E021.

References

Aktipis CA (2004) Know when to walk away: contingent movement and the evolution of cooperation. J Theor Biol 231(2):249–260

Axelrod RM (1990) The evolution of cooperation. Penguin Books, London

Edmonds B (2006) The emergence of symbiotic groups resulting from skill-differentiation and tags. JASSS 9(1):10. ISSN 1460–7425. URL: http://jasss.soc.surrey.ac.uk/9/1/10.html

Gigerenzer G, Todd PM, ABC Research Group et al (1999) Simple heuristics that make us smart. Oxford University Press, New York

Hales D (2003) Evolving specialisation, altruism and group-level optimisation using tags. In: Sichman JS, Bousquet F, Davidsson P (eds) Multi-agent-based simulation II. Lecture notes in artificial intelligence, 2581, pp 26–35

Hales D, Rouchier J, Edmonds B et al. (2003) Model-to-model analysis. JASSS 6(4). URL http://jasss.soc.surrey.ac.uk/6/4/5.html

Holland J (1993) The effect of labels (tags) on social interactions. Technical report, Santa Fe Institute, NM

Huhns MN, Singh MP (1998) Agents and multiagent systems: themes, approaches, and challenges. In: *Readings in Agents*, pp 1–23. Morgan Kaufmann

Joyce D, Kennison J, Densmore O, Guerin S, Barr S, Charles E, Thompson NS et al (2006) My way or the highway: a more naturalistic model of altruism tested in an iterative prisoners' dilemma. JASSS, 9(2):4. ISSN 1460-7425. URL http://jasss.soc.surrey.ac.uk/9/2/4.html

Keller L, Ross KG (1998) Selfish genes: a green beard in the red fire ant. Nature 394(6693):573–575

Norling E, Powell CR, Edmonds B et al (2008) Cross-disciplinary views on modelling complex systems. In: Proceedings of the 9th workshop on multiagent-based simulation

Nowak M (2006) Five rules for the evolution of cooperation. Science, 314, December 2006

Paolucci M, Conte R, Tosto GD et al (2006) A model of social organization and the evolution of food sharing in vampire bats. Adapt Behav 14(3):223–238. ISSN 1059-7123. http://dx.doi.org/10.1177/105971230601400305

Posada M, Lopez-Paredes A (2008) How to choose the bidding strategy in continuous double auctions: imitation versus take-the-best heuristics. JASSS 11(1):6. ISSN 1460-7425. URL http://jasss.soc.surrey.ac.uk/11/1/6.html

Posada M, Hernandez C, Lopez-Paredes A et al (2006) Strategic behaviour in continuous double auction. Lect Notes Econ Math Syst 584:31–43

Riolo RL (1997) The effects of tag-mediated selection of partners in evolving populations playing the interated prisoner's dilemma. Santa Fe Institute Working Paper

Sen S (1996) Reciprocity: a foundational principle for promoting cooperative behavior among self-interested agents. In: Victor L (ed) Proceedings of the first international conference on multiagent systems. MIT Press, pp 322–329

Sen S, Dutta PS (2002) The evolution and stability of cooperative traits. In: *AAMAS '02:* proceedings of the first international joint conference on Autonomous agents and multiagent systems. New York, USA, pp 1114–1120

Tang Y, Parsons S, Sklar E et al (2006) Modeling human education data: from equation-based modeling to agent-based modeling. In: Antunes L, Takadama K (eds) Proceedings of the multiagent-based simulation, LNAI 4442, pp 41–56