

Changing Neighbours: Improving Tag-Based Cooperation

Nathan Griffiths
Department of Computer Science
University of Warwick
Coventry, CV4 7AL, UK
nathan@dcs.warwick.ac.uk

Michael Luck
Department of Computer Science
King's College London
London, WC2R 2LS, UK
michael.luck@kcl.ac.uk

ABSTRACT

In systems of autonomous self-interested agents, in which agents' neighbourhoods are defined by their *connections* to others, cooperation can arise through observation of the behaviour of neighbours to determine values of trust and reputation. While there are many techniques for encouraging cooperative behaviour within such systems, they often require a centralised authority or rely on reciprocity that is not always available. In response, this paper presents a decentralised mechanism to supporting cooperation without requiring reciprocity. The mechanism is based on tag-based cooperation, supplemented by assessing neighbourhood context and using simple rewiring to cope with cheaters. In particular, the paper makes two key contributions. First, it provides a technique for increasing resilience in the face of malicious behaviour by enabling individuals to *rewire* their connections to others and so modify their neighbourhoods. Second, it provides an empirical analysis of several *strategies* for rewiring, evaluating them through simulations.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence — *Multiagent systems*

General Terms

Experimentation, Algorithms, Reliability

Keywords

Cooperation, Tags, Rewiring, Reputation, P2P

1. INTRODUCTION

Many existing approaches for supporting cooperation in decentralised environments are based on *reciprocity*, namely the notion that repeated encounters imply that altruistic or selfish acts performed by an agent may eventually be returned. Reciprocity can be *direct* or *indirect*. Direct reciprocity occurs in situations where two agents have repeat interactions, and over time each provides assistance to the other. Indirect reciprocity exists in situations where agents are likely to interact with others whose behaviour with third

Cite as: Changing Neighbours: Improving Tag-Based Cooperation, Nathan Griffiths and Michael Luck, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 249–256.

Copyright © 2010, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

parties they have previously observed. Trust and reputation are arguably the most widespread mechanisms for supporting cooperation in decentralised systems. However, for an agent to have sufficient information to make effective decisions using trust (based on its own interactions) the environment must contain a high degree of direct reciprocity. Similarly, for reputation (based on the recommendations of others) to be effective, the environment must contain sufficient indirect reciprocity to provide the information needed to determine an individual's reputation.

As systems get larger, more open, and more dynamic, the degree of direct and indirect reciprocity typically reduces, since the likelihood of repeat interactions between any two individuals reduces, as does the likelihood of agents having observed the interactions of others with third parties. P2P systems and the Internet exhibit these characteristics which, in a network context, correspond to moderate (or large) distances between nodes, and low clustering coefficients [15]. Consequently, it becomes more difficult to support cooperation effectively in such environments. In addition, in the presence of individuals seeking to accept the benefits of cooperation without reciprocating, the problem is made even harder. In environments with a low degree of reciprocity there is the temptation for self-interested agents to *cheat* by taking advantage of cooperation from others, but not cooperating in return. The lack of reciprocity means that trust and reputation are not appropriate solutions. In response, in this paper we describe an approach that supports cooperation without requiring reciprocity, and present a technique to improve resilience to malicious behaviour by enabling individuals to change their connections to others, i.e. *rewire* their neighbourhood. We provide alternative strategies for rewiring and evaluate their effectiveness using simulations. Our approach builds on the tag-based approach proposed by Riolo, Cohen and Axelrod (RCA) [18], and Griffiths' modification to RCA's approach [5], in which an assessment of neighbourhood context is used to improve cooperation. Existing work has given little consideration to coping with malicious behaviour, which we address in this paper.

This paper does not focus on a specific application domain. Instead, we consider an abstract environment in which agents are required to interact, but in which there is little direct or indirect reciprocity. We assume that each agent has a neighbourhood defined by its *connections* to others, such that where there is a connection there is the possibility for interaction. These connections are assumed to be bidirectional. We also assume that agents are able to observe the general nature of their neighbours' behaviour (namely,

whether they are acting cooperatively or not), although not the specific details (in terms of with whom they interact and for what task). This environment reflects the form of many real-world settings. For example, in *ad hoc* communication networks any two nodes are unlikely to have many repeat interactions and, although nodes can observe the behaviour of others (e.g. in terms of packet forwarding), it is unlikely that there will be sufficient observations of a given node of interest to determine a meaningful assessment of reputation. Similarly, in a P2P content sharing system, repeat interactions between a given pair of nodes are relatively few, and a node is unlikely to have many observations of the behaviour of a specific individual. However, nodes do have the ability to observe the general nature of their neighbours' behaviour (e.g. in terms of whether they are providing content).

2. BACKGROUND

Mechanisms for enabling cooperation without reciprocity have been of interest for many years to biologists and social scientists investigating how cooperative societies of selfish individuals might evolve [3, 4]. One approach that has received particular interest is the idea that cooperation might be based on the recognition of cultural artefacts or traits [1, 4, 7, 8]. In particular, there have been promising results for using simple observable traits, or *tags* [9], as cultural artefacts to engender cooperation where no reciprocity exists [2, 18, 20]. Tag-based cooperation has recently been shown as successfully able support cooperation in a multi-agent system, in the form of a simulated P2P network [6]. Existing work on tags, however, has given little consideration to the possibility that some members of the population may deviate from the rules of the system, by not cooperating when the rules of the system dictate that they should, and become "cheaters" — the issue that we address in this paper. Before we introduce our mechanism for coping with cheating agents, however, we give an overview of tag-based cooperation.

RCA's tag-based approach to cooperation has been the starting point for much of the work in this area, and it is the base of the mechanism presented in this paper. In RCA's approach an individual's decision to cooperate is based on whether an arbitrary tag associated with it, representing some observable trait, is sufficiently similar to that associated with a potential recipient [18]. The approach is illustrated using a simple *donation scenario* in which each agent acts as a potential donor with a number of randomly selected neighbours. Should an agent opt to donate, it incurs a cost c , and the recipient gains a benefit b (it is assumed that $b > c$), otherwise both receive nothing. Each agent i is initially randomly assigned a tag τ_i and a tolerance threshold T_i with a uniform distribution from $[0, 1]$. An agent A will donate to a potential recipient B if B 's tag is within A 's tolerance threshold T_A , namely $|\tau_A - \tau_B| \leq T_A$. Agents are selected to act as potential donors in P interaction pairings, after which the population is reproduced proportionally to their relative scores, such that more successful agents produce more offspring. Each offspring is subject to mutation, so that with a small probability a new (random) tag is received or Gaussian noise (with mean 0 and a small standard deviation) is added to the tolerance.

RCA have shown that a high cooperation rate can be achieved with this simple approach, without requiring reciprocity. Their results show cycles in which a cooperative population is established, which is then invaded by a mu-




Figure 1: HE's rewiring approach showing an original and rewired neighbourhood.

tant whose tag is similar (and so receives donations) but has a low tolerance (and so does not donate). Such mutants initially do well, leading to them taking over the population subsequently lowering the overall rate of cooperation, but eventually the mutant tag and tolerance become the most common and cooperation again becomes the norm [18].

Hales and Edmonds (HE) apply RCA's approach in a P2P setting, with two main changes [6]. First, RCA's *learning interpretation* of the reproduction phase is adopted, meaning that for reproduction each agent compares itself to another at random and adopts the other's tag and tolerance if the other's score is higher (again subject to potential mutations) [18]. Second, HE interpret a tag as being an agent's set of neighbours in the P2P network, i.e. an agent's connections to others in the population. In RCA's work each agent is connected to each other agent, and there is no corresponding notion of neighbourhood. In HE's model, the process of an agent adopting another's tag is equivalent to it dropping all of its own connections, and copying the connections of the other agent (also adding a connection to the other agent itself) [6]. This process is illustrated in Figure 1 which shows agent A dropping its own connections and adopting those of agent B . Again, there is also a small probability of mutation, interpreted as replacing a randomly selected neighbour with another node in the network.

Using simulations, HE have shown this approach to be promising in situations where agents are given free reign to rewired the network, and replace all of their connections at each reproduction phase. This rewiring is an *all-or-nothing* operation, in that although an agent can adopt a completely new set of neighbours (*replacing* its existing neighbourhood), it cannot *modify* its existing neighbourhood. Our view is that such extreme rewiring, where the neighbourhood topology might completely change with each new generation, is not practicable in all scenarios. For example, in a communication network this would imply that all existing routes

become outdated and need to be re-established, while in a content sharing system an agent would lose all information about the content available in its neighbourhood. In this paper we consider a less extreme situation, in which agents are able to rewire a proportion of their neighbourhood.

Both RCA and HE assume that agents do not deviate from the rules of the system, i.e. they assume that there are no cheaters. We define a *cheater* to be an agent that accepts donations, but will not donate to others, even if the “rules” of the system dictate that it should. We assume that if a cheater reproduces, then its offspring will also cheat. In this paper we assume that cheaters do not falsify their tags or additionally manipulate their tolerance thresholds. In standard tag-based cooperation, introducing even a small proportion of cheaters into the population causes cooperation to collapse [5]. Our aim in this paper is to provide a mechanism that copes with the presence of cheaters but, unlike trust and reputation, does not rely on reciprocity. In common with RCA we view tags as simple arbitrary observable traits, and motivated by HE’s work we allow agents to partially modify their neighbourhood connections. Building on Griffiths’ context awareness modification to RCA’s approach [5], we show that by allowing agents to modify their neighbour connections we are able to significantly improve cooperation in the presence of cheaters.

3. ENHANCING COOPERATION

We consider a population of agents, each of which has (bi-directional) connections to n neighbours, such that agents are only able to interact with their neighbours (although for reproduction we still consider the population as a whole). Unlike RCA and HE, we assume that a proportion of the population are cheaters and will not cooperate with others even when their tags are within the tolerance threshold. In this paper, we adopt the *donation scenario* used by RCA, but we note that this could be extended to more realistic settings, for example in the manner of HE to a P2P scenario [6]. We use RCA’s parameter values for benefit and cost, of $b = 1$ and $c = 0.1$ [18] (the addition of a cost of 0.1 avoids negative payoffs [16]). Each agent i is initially assigned an arbitrary tag τ_i and tolerance T_i with uniform distribution from $[0, 1]$ ¹. Our approach uses two techniques for improving cooperation in the presence of cheaters: augmenting RCA’s approach with neighbourhood context assessment (first introduced in [5]) and the use of neighbourhood rewiring, which is the main contribution of this paper.

3.1 Context Awareness

Our first technique uses Griffiths’ modification of RCA’s model, in which agents assess their neighbourhood context in terms of how cooperative they perceive their neighbours to be [5]. The donation decision is modified so that an agent’s assessment of its neighbourhood context becomes a factor in the decision to donate. This approach relies on the assumption that agents can observe their neighbours’ donation behaviour, and is realistic in many real-world settings. For example, in a communication network nodes can detect whether packets have been forwarded, and in a file sharing

¹More strictly we apply a lower bound of -10^{-6} to *tolerance* to address Roberts and Sherratt’s concerns regarding agents with identical tags being forced to cooperate [19]. This also allows the population to contain non-cooperative agents of the form considered by Masuda and Ohtsuki [12].

system nodes can observe whether others’ downloads have completed. These observations allow an agent to assess its neighbourhood context, i.e. how cooperative its neighbours are. Agents are given a fixed length FIFO memory to record the last l donation behaviours observed for each neighbour. When the neighbour donates, an observation value of +1 is recorded, and when it does not -1 is recorded. This memory is fairly sparse, since the number of interactions is small compared to the number of agents, and so the overhead incurred is relatively small (2 bits per observation for $n \times P$ observations, where n is the number of neighbours and P the number of pairings).

In order to assess its neighbourhood context, an agent considers each of its n neighbours in turn, and determines the contribution to the context assessment c_i of neighbour i , which is simply the proportion of observed interactions in which the neighbour donated, given by:

$$c_i = \begin{cases} \frac{\sum_{j=1}^{l_i} o_i^j}{l_i}, & \text{if } l_i > 0 \\ 0, & \text{otherwise} \end{cases}, \quad \text{if } l_i > 0 \quad (1)$$

where o_i^j represents the j ’th observation of neighbour i , and l_i is the number of observations recorded of i ’s donation behaviour ($l_i < l$). By considering each of its n neighbours, agent A ’s assessment of its current neighbourhood context C_A is given by:

$$C_A = \frac{\sum_{i=1}^n c_i}{n} \quad (2)$$

This context assessment can be used to influence the donation decision. The intuition is that that agents “expect” that by donating they are more likely to receive a future donation from some other (observing) agent. However, because the number of interactions is small compared to the number of agents, this is a *weak* notion of indirect reciprocity, which is insufficient to support a “traditional” notion of reputation. One agent’s donation to another is unlikely to be directly repaid or directly observed by a third party, so that there is little direct or indirect reciprocity. Instead, context assessment gives a general impression of the donation behaviour in a neighbourhood, indicating the likelihood of receiving future donations, and hence can be viewed as a weak notion of indirect reciprocity. An agent’s assessment of its neighbourhood context is incorporated into the model by adapting the decision to donate, such that both tolerance and neighbourhood context are considered (provided some minimum number of observations have been made). Thus, an agent A will donate to B if:

$$|\tau_A - \tau_B| \leq (1 - \gamma) \cdot T_A + \gamma \cdot C_A \quad (3)$$

The parameter γ , which we call the *context influence*, allows us to tune the technique. If $\gamma = 0$, the technique is identical to RCA’s method, while if $\gamma = 1$ then the donation decision is determined solely by an agent’s assessment of its neighbourhood context.

We adopt RCA’s learning interpretation of reproduction (as do HE), such that after a fixed number of interaction pairings P an agent compares itself to another selected at random from the population. If the other agent is more successful, then its tag and tolerance are copied (subject to a small probability of mutation), meaning that the other agent reproduces. Otherwise the tag and tolerance are unchanged, meaning that the first agent itself reproduces.




Figure 2: The effect of context on donation.

The effect of including context assessment in the decision to donate is illustrated in Figure 2, which shows the average donation rate across the population after 100 generations (by which time the donation rate has stabilised) in a population containing 10% cheaters against the context influence γ . The results in Figure 2 were obtained using a population of 100 agents, a neighbourhood size of $n = 10$, and 10 interactions per agent per generation ($P = 10$). The leftmost point, with $\gamma = 0$, is equivalent to RCA’s approach with no consideration of context, and the rightmost point, with $\gamma = 1$, uses only context assessment in the donation decision. Without context assessment, RCA’s approach gives a donation rate of below 40%, while using a context influence of $\gamma > 0.3$ we see a significant rise in donation rate, achieving over 60% for $\gamma \geq 0.5$. It is clear that considering context in the donation decision can provide an increase in the donation rate. In this paper, our focus is on the rewiring technique and strategies, and so we do not include further results on this part of the mechanism (a discussion of the relevant factors can be found in [5]). In the remainder of the paper, therefore, we assume $\gamma = 0.5$ unless stated otherwise, since in general we have found that this provides a significant improvement in donation rate, and that there are relatively small gains in using larger values (although we briefly discuss the effect of context influence on *rewiring* when considering the results presented in Figure 8).

3.2 Simple Rewiring

Our second technique enables agents to *rewire* their network neighbourhoods, such that after reproduction an agent is able to remove a proportion λ (the *rewire proportion*) of connections to neighbours, and replace them with connections to new neighbours. This approach is motivated by HE’s promising results, but unlike HE we do not assume that an agent can replace *all* of its existing connections since, as discussed above, this is likely to be impractical in real-world settings. Our hypothesis is that (i) cooperation can be improved by removing connections to agents that are not cooperative, and (ii) adding new connections based on the experiences of others can improve cooperation. We investigate this hypothesis by considering the following strategies, which are in increasing order of sophistication:

(i) *random*: λ neighbours are removed at random, and λ new neighbours are added at random. (ii) *randomReplaceWorst*: the λ worst neighbours are removed, and λ new neighbours are added at random. (iii) *individualReplaceWorst*: the λ worst neighbours are removed, and then the best λ neighbours of the agent’s best neighbour are added. Additional randomly selected neighbours are added if necessary to prevent the neighbourhood shrinking due to duplication (an agent has at most one connection to another individual, with duplicate connections having no meaning). (iv) *groupReplaceWorst*: the λ worst neighbours are removed, and then the best (non-duplicate) neighbour from each of the agent’s λ best neighbours are added. The neighbours are considered in descending rank order and, for each, the best non-duplicate neighbour is added. Again, additional randomly selected neighbours are added if necessary, to ensure that each agent remains connected to n others.

For *randomReplaceWorst*, *individualReplaceWorst* and *groupReplaceWorst* we determine which connections to remove, by using the contribution to the context assessment c_i for each neighbour i (as defined in Equation 1) as the metric for ranking agents. Thus, an agent will remove connections to the λ agents that have the lowest c_i values. Similarly, for *individualReplaceWorst* and *groupReplaceWorst* we also use the contribution to the context assessment to determine which connections to add. If the c_i values of two or more agents are equal then one is selected arbitrarily. The parameter λ , which we call the *rewire proportion*, determines the extent to which the network is rewired each generation. Note that the latter two strategies can be thought of as simplistic reputation mechanisms, in that agents update their connections based on the experiences of others. However, unlike typical reputation mechanisms, the assessment is based on relatively little information, which is not predicated on a notion of (direct or indirect) reciprocity [11, 17].

The *individualReplaceWorst* and *groupReplaceWorst* strategies are illustrated in Figure 3. If A ’s neighbours in order of preference are B, C, D, E, F , and 2 neighbours are to be replaced ($\lambda = 2$), then the connections to E and F will be dropped. If B ’s neighbours, in preference order, are H, D, I, G, A then for *individualReplaceWorst* A will add H and I to its neighbourhood, as shown in (b). D is not added since it is already in A ’s neighbourhood and so the next preferred neighbour is used. For *groupReplaceWorst*, shown in (c), A will add a connection to H (from B ’s neighbours) and to C ’s most preferred neighbour, which we suppose is J .

4. STRATEGY EVALUATION

Using the PeerSim P2P simulator², we have built a simulation of our model and have experimented with several alternative configurations. In this section we give an overview of the main findings and explore the parameters in turn. Unless stated otherwise, the results presented here represent an average of 10 runs using a population of 100 agents, a neighbourhood size of $n = 10$ (using a regular random network as the initial topology), 10 pairings per agent per generation ($P = 10$), and a context influence of $\gamma = 0.5$. After reproduction, the probability of mutating the tag by selecting a new random value is 0.001, and there is a 0.001 probability of adding Gaussian noise to the tolerance (with mean 0 and standard deviation of 0.01).

²<http://peersim.sourceforge.net/>




Figure 3: A's original neighbourhood (a) and the result of rewiring using *individualReplaceWorst* (b), and *groupReplaceWorst* (c).




Figure 4: Rewiring strategies (10% cheaters).

Figure 4 compares the donation rate for a population containing 10% cheaters using the four alternative rewiring strategies with different rewire proportions. RCA's standard approach and Griffiths' context assessment modification to RCA's approach (without rewiring) are included for comparison. The lower of the two dashed horizontal lines represents RCA's standard approach, and the upper dashed horizontal line represents RCA's approach with context assessment. These correspond to the leftmost ($\gamma = 0.0$) and middle points ($\gamma = 0.5$) of the results shown in Figure 2. The first conclusion that can be drawn from Figure 4 is that although the *random* strategy improves upon RCA's standard approach, it performs worse than context assessment alone. Second, it is clear that the more sophisticated strategies that replace the *worst* λ neighbours perform better than using context assessment alone, except for very high replacement proportions. For $\lambda \leq 0.7$ there is little difference in the performance of the three strategies. (For clarity we have omitted the standard deviations from Figure 4, but for $\lambda \leq 0.7$ the strategies are within one standard deviation of each other.) For $\lambda > 0.7$ the two most sophisticated strategies of *individualReplaceWorst* and *groupReplaceWorst* give the best performance, and there is not a significant difference between them in this experimental configuration, although the latter gives a marginally better donation rate. As discussed above, in a real-world system such high degrees of




Figure 5: Rewiring strategies (20% cheaters).

rewiring may not be possible, and in many situations the rightmost side of the graph may not be reachable due to an upper limit on λ of less than 1. In all of our simulations, the *random* strategy is consistently significantly worse than the other rewiring strategies, and often worse than context assessment alone; we therefore omit it from the remaining results presented in this paper.

In order to investigate further the performance of the strategies, we varied the proportion of cheaters present. Figures 5 and 6 give the results for 20% and 30% cheaters respectively, with RCA's standard approach and context assessment shown for comparison. RCA's standard approach again performs the worst, followed by context assessment. As the proportion of cheaters is increased, the donation rate reduces for each of the strategies, as we would expect. Again, there is relatively little difference in performance between *individualReplaceWorst* and *groupReplaceWorst*, although on average the latter gives a marginal improvement. With 30% cheaters, *randomReplaceWorst* is significantly worse than the more sophisticated strategies for $\lambda > 0.3$. Therefore, adding new connections based on the experiences of others, rather than at random, is beneficial. The highest donation rate for both 20% and 30% cheaters is achieved with *groupReplaceWorst* and $\lambda = 0.6$. Statistical significance test results are omitted due to space constraints, but the benefits of *individualReplaceWorst* and *groupReplaceWorst* over other strategies are significant for all λ . For moderately




Figure 6: Rewiring strategies (30% cheaters).




Figure 7: Standard deviation in donation rate (30% cheaters).

high rewire proportions ($\approx 0.6..0.9$) with $\leq 20\%$ cheaters *groupReplaceWorst* is significantly better, otherwise the difference is marginal.

To further illustrate the performance of the strategies we have plotted the standard deviation in donation rate for *randomReplaceWorst*, *individualReplaceWorst*, and *groupReplaceWorst* in Figure 7 (for a population containing 30% cheaters). From these results it is clear that the standard deviation of the *randomReplaceWorst* strategy is the highest, and that for a rewire proportion of $\lambda \geq 0.4$ *groupReplaceWorst* has a slightly lower standard deviation than *individualReplaceWorst*. Taking into account both the donation rates obtained for the various cheater proportions and the standard deviations, the best donation rate is likely to be achieved with *groupReplaceWorst* using a rewire proportion of $\lambda \geq 0.4$ (although on average the improvement over *individualReplaceWorst* is $< 3\%$).

In relation to our hypothesis that (i) cooperation can be improved by removing connections to agents that are not cooperative, and (ii) adding new connections based on the experiences of others can improve cooperation, we see these




Figure 8: The effect of context influence.

results as confirmation that this is indeed the case. Certainly, each of the strategies that removes connections to the worst λ neighbours perform better than those strategies that do not utilise such rewiring. The evidence for the latter part of the hypothesis is less strong, although we see a clear improvement in donation rate for *individualReplaceWorst* and *groupReplaceWorst* over *randomReplaceWorst* where there are higher proportions of cheaters, and in particular where there are larger rewire proportions ($\lambda > 0.4$). We also see a lower standard deviation for *groupReplaceWorst* and *individualReplaceWorst*. However, for lower proportions of cheaters, or low rewire proportions, the advantage in using the experiences of others is marginal.

4.1 Context influence

We have performed simulations to compare the use of rewiring with context assessment against the use of rewiring alone, and with different weightings given to *context influence*, i.e. exploring different values for γ . This enables us to evaluate whether our results arise from the combination of rewiring and context assessment or purely from rewiring alone. The results are shown in Figure 8, in which the *groupReplaceWorst* strategy is used (we have obtained similar results for the other strategies). It can be seen that the worst donation rate is achieved with a context influence $\gamma = 0.0$, which is equivalent to using RCA's standard approach with the addition of rewiring. As more influence is given to the context, the donation rate improves, but there is no clear best value for γ , with all values used providing an improvement over not using neighbourhood context assessment. This supports our use of $\gamma = 0.5$ as a reasonable parameter value in our other simulations. Our conclusion from these results is that both the context assessment and rewiring techniques improve the donation rate, but that there are no specific values for γ and λ that perform best in all situations. In general, we find that values in the middle of the range perform well, i.e. $\gamma \approx 0.5$ and $0.4 \leq \lambda \leq 0.8$ but we suggest that tuning through empirical experimentation is needed for any given configuration.

4.2 Interaction pairings

The number of interaction pairings per generation, P ,




Figure 9: Varying pairings per generation.




Figure 10: Varying population size ($n = 10$).

also affects the donation rate, as can be seen in Figure 9 (we omit *randomReplaceWorst* from the remaining results since it gives lower performance). The donation rate with *groupReplaceWorst* is fairly stable for more than 10 pairings (i.e. $P \geq 10$) with an average donation rate of 66%, while for $P < 10$ the average donation rate drops to 61%. With *individualReplaceWorst* we see a fairly stable donation rate for $P \leq 25$ of 64%, which reduces for $P > 25$ to 54%. Overall, *groupReplaceWorst* performs better than *individualReplaceWorst*, by 3.6% on average, although the difference is most pronounced with larger number of pairings ($P > 25$). The *groupReplaceWorst* strategy also has a slightly lower standard deviation than *individualReplaceWorst*, of around 8% compared to 11%. Thus, *groupReplaceWorst* is again preferred to *individualReplaceWorst* (although the improvement is marginal for $P \leq 25$), and on average the highest donation rate is obtained for $10 \leq P \leq 25$.

4.3 Population size

Population size has relatively little influence on the donation rate obtained, as shown in Figure 10, in which we fix the neighbourhood size to $n = 10$ and vary the population

Strategy	n			
	5	10	25	50
<i>groupReplaceWorst</i>	0.68	0.69	0.68	0.57
<i>individualReplaceWorst</i>	0.60	0.68	0.63	0.52

Table 1: Donation rate for various neighbourhood sizes, with 100 agents and 30% cheaters.

size between 100 and 500. It can be seen that, although there are fluctuations, there is not a significant change in donation rate as the population size is increased. Again, *groupReplaceWorst* gives a marginally higher average donation rate (67%) than *individualReplaceWorst* (65%), with a lower standard deviation (of 4% compared to 7%). Based on these results our view is that population size itself is not a major factor in the donation rate achieved, but that *groupReplaceWorst* is (marginally) the preferred strategy.

4.4 Neighbourhood size

We have also explored the effect of neighbourhood size on the donation rate. Table 1 shows the donation rate using *individualReplaceWorst* and *groupReplaceWorst* with neighbourhood sizes of $n = 5, 10, 25$ and 50 for a population size of 100 with 30% cheaters. For *groupReplaceWorst* there is relatively little difference in donation rate with $n \leq 25$, but $n = 50$ gives a reduction in donation rate of around 10%. For *individualReplaceWorst*, however, there is more variation in donation rate depending on the neighbourhood size, with the highest donation rate of 68% being achieved for $n = 10$, but lower rates achieved for other values, and again the donation rate is significantly reduced for $n = 50$. From these results we see that *groupReplaceWorst* gives better results than *individualReplaceWorst* for a wider range of neighbourhood sizes. In both cases the best donation rate is achieved for $n = 10$, although for *groupReplaceWorst* the advantage over a neighbourhood size of 5 or 25 is marginal.

4.5 Summary of evaluation

We have found that our results support the hypothesis that (i) cooperation can be improved by removing connections to agents that are not cooperative, and (ii) adding new connections based on the experiences of others can improve cooperation. Furthermore, we see that *groupReplaceWorst* gives the best donation rate on average, closely followed by *individualReplaceWorst*. In situations with low proportions of cheaters, or where only low degrees of rewiring are possible then *randomReplaceWorst* also improves donation rate, but in general it does not perform as well as the strategies that consider the experiences of others when adding connections. In the presence of cheaters the *random* strategy improves the donation rate over RCA's standard approach, but does not perform as well as context assessment alone. In general, moderate settings for the parameters perform the best, with a context influence of $\gamma \approx 0.5$ and a rewiring proportion of $0.4 \leq \lambda \leq 0.8$. Similarly, the best results are achieved for a moderate number of interaction pairings per generation of $10 \leq P \leq 25$. We find that a neighbourhood size of $n = 10$ gives the highest donation rate, but that population size has little influence on the donation rate.

5. RELATED WORK AND CONCLUSIONS

In this paper we have presented and demonstrated a tag-

based mechanism for supporting cooperation in the presence of cheaters. Other existing work on tag-based cooperation focuses on other aspects and does not consider cheaters as we define them. For example Jansen and van Baalen [10] investigate the effect of tag space size (i.e. the set of possible values for tags), and find that a relatively large tag space is required to allow cooperation to persist over generations in weakly structured populations. In this paper we also use a large tag space, by allowing tags to take any value in the interval $[0, 1]$. Masuda and Ohtsuki [12] describe a mechanism for improving cooperation in populations containing agents that are permitted to have a negative tolerance (i.e. agents that will not cooperate with others that have an identical tag). Their approach allows agents to observe the tolerance of others, and not cooperate with those that have a negative tolerance, thereby causing such agents to be unsuccessful and not reproduce. In this paper, we do not assume that agents can observe the tolerance of others, but our mechanism is designed to cope with non-cooperative agents in the form of cheaters. Matlock and Sen [13, 14] generalise the tag matching mechanism to enable cooperation between different groups of (non-cheating) individuals with different tags. They use *tag matching patterns* (where tags are strings of bits against which pattern matching is performed), payoff sharing for agents to share payoffs with their “opponent”, and modified reproduction and mutation mechanisms that preserve tag matching patterns. In this paper we focus on the issue of cheaters, and do not consider cooperation between different social groups, and although we allow cooperation between individuals with different tags, as per RCA we view these agents to be part of the same group (i.e. sharing the same tag within some tolerance threshold).

The mechanism presented in this paper supports cooperation in populations of autonomous self-interested agents, some of whom may be cheaters, without requiring reciprocity. Our approach adds a novel technique for neighbourhood rewiring to Griffiths’ context assessment modification [5] of RCA’s approach. We also presented and evaluated a number of rewiring strategies. Our results show that cooperation can be improved by enabling agents to change their neighbour connections, in particular by removing connections to the worst neighbours and replacing them with connections to new neighbours. The best results are obtained by considering the experiences of others, by connecting to neighbours with whom others have had positive experiences. We also found that it is better to consider the experiences of several other agents rather than those of a single agent. Our results show that context assessment and rewiring techniques both improve cooperation, and that moderate values of context influence and rewiring proportion give the best results.

There are several areas of ongoing work. We intend to continue our experimental evaluation to further understand the relationships between the parameters that define an agent’s environment and behaviour. Previous tag-based approaches give evidence that mutation rate influences whether cooperative behaviour can be established [20], and that tag space size is important [10]. We aim to investigate how these factors, among others, influence cooperation in our approach. We will also investigate different initial topologies, and how these evolve, with a view to evaluating our approach in a mobile ad-hoc communication environment in which agents can influence, but not completely control, their neighbourhoods by moving around the environment.

6. REFERENCES

- [1] P. D. Allison. The cultural evolution of beneficent norms. *Social Forces*, 71(2):279–301, 1992.
- [2] R. Axelrod, R. A. Hammond, and A. Grafen. Altruism via kin-selection strategies that rely on arbitrary tags with which they coevolve. *Evolution*, 58(8):1833–1838, 2004.
- [3] R. Boyd and P. J. Richerson. The evolution of indirect reciprocity. *Social Networks*, 11:213–236, 1989.
- [4] R. Dawkins. *The Selfish Gene*. Oxford University Press, 1976.
- [5] N. Griffiths. Tags and image scoring for robust cooperation. In *Proc. 7th Int. Conf. on Autonomous Agents and Multiagent Systems*, pages 575–582, 2008.
- [6] D. Hales and B. Edmonds. Applying a socially inspired technique (tags) to improve cooperation in P2P networks. *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, 35(3):385–395, 2005.
- [7] W. D. Hamilton. The genetical evolution of social behaviour I. *Journal of Theoretical Biology*, 7(1):1–16, 1964.
- [8] W. D. Hamilton. The genetical evolution of social behaviour II. *Journal of Theoretical Biology*, 7(1):17–52, 1964.
- [9] J. H. Holland. *Hidden order: How adaptation builds complexity*. Addison-Wesley, 1995.
- [10] V. A. A. Jansen and M. van Baalen. Altruism through beard chromodynamics. *Nature*, 440:663–666, 2006.
- [11] A. Jøsang, R. Ismail, and C. Boyd. A survey of trust and reputation systems for online service provision. *Decision Support Systems*, 43(2):618–644, 2007.
- [12] N. Masuda and H. Ohtsuki. Tag-based indirect reciprocity by incomplete social information. *Proc. Royal Society B*, 274(1610):689–695, 2007.
- [13] M. Matlock and S. Sen. Effective tag mechanisms for evolving coordination. In *Proc. 6th Int. Conf. on Autonomous Agents and Multiagent Systems*, pages 1345–1352, 2007.
- [14] M. Matlock and S. Sen. Effective tag mechanisms for evolving coordination. In *Proc. 8th Int. Conf. on Autonomous Agents and Multiagent Systems*, pages 489–496, 2009.
- [15] M. E. J. Newman. The structure and function of complex networks. *SIAM Review*, 45:167–256, 2003.
- [16] M. Nowak and K. Sigmund. Evolution of indirect reciprocity by image scoring. *Nature*, 393:573–577, 1998.
- [17] P. Resnick, R. Zeckhauser, E. Friedman, and K. Kuwabara. Reputation systems. *Communications of the ACM*, 43(12):45–48, 2000.
- [18] R. Riolo, M. Cohen, and R. Axelrod. Evolution of cooperation without reciprocity. *Nature*, 414:441–443, 2001.
- [19] G. Roberts and T. N. Sherratt. Does similarity breed cooperation? *Nature*, 418:499–500, 2002.
- [20] A. Traulsen. Mechanisms for similarity based cooperation. *The European Physical Journal B - Condensed Matter and Complex Systems*, 63(3):363–371, 2008.