

Cooperation through Tags and Context Awareness

Nathan Griffiths¹

Abstract. In recent years a range of techniques such as trust, reputation and social norms have been used to support cooperation. Attention has tended to focus on situations where a degree of reciprocity, either direct or indirect, exists between agents, and existing techniques typically rely on such reciprocity to engender cooperative behaviour. Increasingly, environments are emerging where large numbers of agents interact without ongoing repeat interactions, in which there is little or no reciprocity. In this paper, we propose a mechanism to support cooperation without requiring reciprocity. Our approach supplements tag-based cooperation with an assessment of neighbourhood context to cope with cheaters. Using a simple peer-to-peer scenario we show how cooperative behaviour is favoured, and the effect of cheating agents is reduced.

1 Introduction

A range of techniques including trust, reputation and social norms have been used to establish and maintain cooperation in multi-agent systems. Many successful approaches have been developed for a number of environments. However, the increased use of large distributed systems such as peer-to-peer (P2P) networks, and the emergence of ubiquitous computing environments, mean that enabling and maintaining cooperation remains an important question. Such environments typically contain a large number of agents that must cooperate, but the environment characteristics are such that repeat interactions between agents may be rare. Many of the common approaches for supporting cooperation in multi-agent systems, although helpful, are not a complete solution, since little may be known about potential interaction partners and there is a relatively low likelihood of any subsequent interactions with the same partner. In this paper we propose a mechanism, that combines ideas from biology and the social sciences, to support cooperation in such environments. Our approach is an extension of the tag-based mechanism proposed by Riolo, Cohen and Axelrod (RCA) [14]. The approach we propose in this paper is related to that we describe in [6], in which an alternative extension to RCA's approach is considered.

Most existing approaches to cooperation are based on reciprocity, namely the notion that repeated encounters imply that any altruistic or selfish act performed by an agent may eventually be returned by the recipient. Direct reciprocity is the simplest, and most common, approach where two agents have repeat interactions in which there is the opportunity to “cooperate” or “defect”. The iterated “prisoner’s dilemma” is a quintessential example of such a setting. In large scale systems, such as P2P networks, interactions between a given pair of agents are infrequent and often single-shot, and so there is minimal direct reciprocity present. An alternative is indirect reciprocity, where a third party is involved in repeat interactions. Agents are unlikely to have direct repeat interactions, but are likely to interact with

others whose behaviour with third parties they have previously observed. Nowak and Sigmund characterise direct reciprocity through the principle of “You scratch my back, and I’ll scratch yours”. Similarly, indirect reciprocity is characterised as “You scratch my back, and I’ll scratch someone else’s” or “I scratch your back and someone else will scratch mine” [11]. In some circumstances, however, even indirect reciprocity might be limited, and we may need to enable cooperation without reliance on reciprocity of any form, for example if there is no memory of past encounters [14].

Trust and reputation are the most common approaches to supporting cooperation in multi-agent systems [9, 12, 13]. However, such techniques are based on the notion of reciprocity and so are of limited use in situations where reciprocity is lacking. In this paper we extend RCA’s tag-based mechanism [14], to provide a model for establishing and maintaining cooperation that does not assume reciprocity, and is suitable for situations where repeat interactions are rare. In the following section we introduce the theoretical approaches upon which our model is based, along with the promising initial results obtained by others. The model itself is introduced in Section 3. Our experimental setting, in the form of a simple P2P system, and selected experimental results are discussed in Section 4, and Section 5 concludes the paper.

2 Related Work

Indirect reciprocity is not an novel idea: biologists and social scientists have long considered cooperation in environments where the individuals concerned may not directly meet again, but where cooperative strategies are favoured [1, 3, 10]. Furthermore, theoretical models of cooperation exist that do not require any reciprocity, but instead are based on the recognition of cultural artifacts, such as the “green beard effect” and “kin” recognition [2, 4]. Promising results have recently been obtained using “tags” [8] as cultural artifacts to enable cooperation without reciprocity [14], which in has in turn led to a technique to improve cooperation in P2P networks [7]. Existing work on tags, however, has given only limited consideration to the existence of “cheaters” in the population, and it is this issue that we address in this paper.

Riolo, Cohen and Axelrod describe a tag-based approach to cooperation in which an agent’s decision to cooperate is based on whether an arbitrary “tag” associated with it is sufficiently similar to that associated with the potential recipient [14]. In RCA’s model each agent i is initially randomly assigned a tag τ_i and a tolerance level 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 $|\tau_A - \tau_B| \leq T_A$. Thus, agents with a high tolerance will donate to others with a wide range of tags, while those with a low tolerance only donate to others with very similar tags [14]. When an agent donates it pays a cost c , and the recipient receives a benefit

¹ University of Warwick, UK, email: nathan@dcs.warwick.ac.uk

b (it is assumed that $b > c$). RCA have performed simulations in which each agent acts as a potential donor in P interaction pairings, after which the population of agents is reproduced in proportion to their relative scores. Each offspring’s tag and tolerance is subject to a potential mutation, such that with some small probability a new (randomly selected) tag is received or the tolerance is mutated by the addition of Gaussian noise (with mean 0 and a small standard deviation). RCA found that a high cooperation rate can be achieved with this simple model, in which no reciprocity is required. Their results show oscillations in which a cooperative population is established, only to be invaded by a mutant whose tag is similar (and so receives donations) but with low tolerance (and so does not donate). Such mutants initially do well and take over the population, lowering the overall rate of cooperation. Eventually, the mutant tag becomes the most common and cooperation again becomes the norm [14].

RCA’s approach is an effective mechanism for achieving cooperation without relying on reciprocity, but their model relies on an assumption that no cheaters are present in the population. A cheating agent is one that accepts donations, but will not donate to others, even if the “rules” of the system dictate that it should. Thus, a cheater in RCA’s scenario would accept donations, but never donate to others regardless of tag similarity. We assume that cheaters follow the usual rules of reproduction in terms of offspring characteristics (e.g. tag and tolerance), but that their offspring will also be cheaters.

Hales and Edmonds (HE) apply RCA’s approach in the context of a P2P network, with two important changes [7]. The first change is to adopt RCA’s “learning interpretation” of the reproduction phase, such that each agent compares itself to another and adopts the other’s tag and tolerance if the other’s score is higher (again subject to potential mutations) [14]. The second change is that HE interpret a tag as an agent’s set of neighbours in the P2P network. Thus, adopting another agent’s tag is equivalent to re-wiring the P2P network such that the other agent’s connections are adopted [7]. Again, there is a small probability of mutation, which is interpreted as replacing a randomly selected neighbour with another node in the network. Simulations performed by HE have shown this approach to be very promising in situations where agents are able to re-wire the network, and in which there are no cheaters. In this paper, motivated by HE’s promising results, we focus on achieving cooperation in the presence of cheaters, without permitting agents to re-wire their network neighbourhoods. Our approach is based on RCA’s model, and HE’s application of it (minus re-wiring), supplemented by a mechanism to cope with cheaters.

3 Extending Tags through Context Assessment

In this paper we use a P2P network as an illustrative scenario, and although we intend our approach to be fairly generic, our discussion will focus on a P2P setting. We consider a network of nodes, or agents, in which each agent has a fixed number n of connections to neighbours. The network topology is assumed to be fixed, and we do not permit agents to re-wire their network connections. Furthermore, unlike RCA and HE we assume that a proportion of the population will be cheaters, meaning that they will take all the benefits offered to them but will always refuse to act cooperatively towards others. For simplicity, we adopt the “donation scenario” used by RCA in which each agent is chosen to act as a potential donor with a number of neighbours. If the agent donates it incurs a cost c and the recipient receives a benefit b , otherwise both agents receive nothing. We use RCA’s parameter values of $b = 1$ and $c = 0.1$. (These values are in turn adopted from Nowak and Sigmund, and the addition of a cost of

0.1 is to avoid negative payoffs [10].) It should be noted that although this is an artificial scenario, it could be extended in the manner of HE to more realistic P2P applications such as file sharing [7].

Our approach is founded upon RCA’s tag-based technique, but we incorporate a simple mechanism to combat cheaters in which agents assess their current context, in terms of their neighbours’ donation behaviour, as part of the decision to donate. Each agent i is initially assigned an arbitrary tag τ_i and tolerance T_i with uniform distribution from $[0, 1]^2$. As in RCA’s model, an agent A will donate to a potential recipient B if B ’s tag is within a certain threshold of its own. To demonstrate the impact of cheaters, initially suppose that this threshold corresponds to A ’s tolerance (as per RCA’s model), meaning that if $|\tau_A - \tau_B| \leq T_A$ then A will donate to B . Later, we will expand this interpretation to include A ’s assessment of its current context. RCA’s learning interpretation of reproduction is adopted (i.e. that used by HE) such that after a fixed number P of interaction pairings an agent compares itself to another selected at random. If the other agent is more successful than itself then the other’s details (its tag and tolerance) are copied, meaning that the other agent reproduces, otherwise no change is made. If the parent agent is a cheater, then its offspring will also be a cheater, regardless of its other characteristics. After reproduction there is a potential mutation of the offspring’s tag and tolerance, with probabilities m_τ and m_T respectively.

In common with RCA we find that a relatively stable donation rate (i.e. cooperation) over a large number of generations is established for appropriate parameters, provided that cheaters are not introduced into the population. Figure 1 shows the dynamics of the donation rate for a configuration that mirrors RCA’s setting. Specifically, we use the parameter values $m_\tau = m_T = 0.01$ and $P = 3$. Note that in our P2P setting an agent has a restricted set of neighbours (in this case $n = 49$ where the network size $N = 100$) whereas in RCA’s approach an agent has all others in the population as “neighbours” in this sense. Our values differ from those used by RCA in that the probability of tag mutation and tolerance mutation are lower (RCA use $m_\tau = m_T = 0.1$). Using these parameters the form of our results matches those obtained by RCA in [14]. If we use RCA’s parameter values for m_τ and m_T we get a significantly lower donation rate than in their simulations. The reasons for this are unclear, and require future investigation. However, Edmonds and Hales notice similar differences from RCA’s results, and suggest that bias in reproducing agents with equal scores and automatic donation to “tag clones” in RCA’s simulation are potential contributory causes [5].

Figure 1 shows that RCA’s model (parameter values aside) allows cooperation to be established in the absence of cheaters. Unfortunately, when cheaters are introduced cooperation soon disappears. Figure 2 shows the effect of creating a population where a proportion of agents act as cheaters, who accept donations from others but never donate (regardless of tag similarity). Where there are no cheaters (the upper dotted line) cooperation is established as before. Introducing 5% of the population as cheaters reduces the donation rate (the solid line) and allowing 10% of agents to be cheaters (the dashed line) leads to minimal cooperation (with under 10% of interactions being cooperative). Without modification, therefore, RCA’s approach soon fails to provide cooperation in the presence of cheaters, with even relatively small proportions of cheaters significantly reducing the average donate rate.

² More strictly we allow tolerance to have a lower bound of -10^{-6} to address Roberts and Sherratt’s concerns that RCA’s approach forces agents with identical tags to always cooperate [15]. The results discussed in this paper permit this small negative tolerance.




Figure 1. Donation rate with no cheaters using RCA's approach.

To cope with the presence of cheaters we extend RCA's approach such that the decision to donate is related to the context in which an agent is situated, in addition to its tolerance. Each agent has a fixed set of n connections to its neighbours, and we assume that these neighbours are able to observe the agent's donation behaviour. This observation assumption is realistic in many real-world settings. For example, in a file sharing system nodes can observe whether other nodes' downloads have completed, or in a communication network nodes can detect whether packets have been forwarded. Using the observations of its neighbours' donation behaviour, an agent is able to assess the context in which it is situated, with respect to how cooperative its neighbours are (i.e. how often they donate). Agents have a fixed length memory, in which they records the last l donation interactions observed for each of their neighbours. Where the neighbour donated to another agent a value of +1 is recorded, and where it refused to donate a value of 0 is recorded. The memory operates as a FIFO queue, such that new entries are appended until the maximum capacity of l is reached, at which point the oldest entry is removed from the head of the queue to allow the new entry to be appended. Based on the set of observations across all neighbours an agent can estimate the current context. Note that this memory is fairly sparse, since the number of interactions is relatively small compared to the number of agents, and so the overhead incurred is fairly small.

In order to assess its current context an agent considers each of its neighbours in turn, and taking them together builds an assessment of how cooperative its context is (in terms of donations). The contribution to the context c_n of neighbour n is simply the proportion of observed interactions in which the neighbour donated, given by:

$$c_n = \begin{cases} \frac{\sum_{j=1}^{l_n} o_n^j}{l_n} & \text{if } l_n > 0 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where o_n^j represents the j 'th observation of n , and l_n is the number of observations of n 's donation behaviour ($l_n < l$). By considering the donation behaviour of each of its n neighbours, an agent can assess its current context C_A as follows:

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

An agent can now consider its context when deciding whether or not to donate. Our assumption is that an agent is more likely to do-




Figure 2. The effect of cheaters on donation rate with RCA's approach.

nate if in a cooperative context. This is related to the notion of indirect reciprocity in that agents "expect" that by donating they are likely to receive a donation from some other (observing) agent in the future. However, because the number of interactions is small compared to the number of agents, this is a *weak* notion of indirect reciprocity. Specifically, we do not assume that a donor will have directly observed a recipient's past behaviour, but only that donors are able to make a general assessment of their current context. The notion of context is incorporated into the model by adapting the decision to donate, such that both tolerance and context are considered. To ensure that an agent has sufficient observations on which to base its assessment we introduce a minimum observations threshold σ . If the total number of observations exceeds σ then context is incorporated into the donation decision, otherwise RCA's standard approach is used. Thus, if $\sum_{i=1}^n l_n \geq \sigma$ then context is incorporated into the donation decision, while if $\sum_{i=1}^n l_n < \sigma$ tolerance alone is used as per RCA's approach. Assuming that there are sufficient observations, then an agent A will donate to B if:

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

where T_A is A 's tolerance and C_A its assessment of the current context. The parameter γ allows us to tune the model. A value of $\gamma = 0$ means that the model is identical to RCA's approach, while a value of $\gamma = 1$ implies that the decision to donate is based solely on the agent's assessment of its context, with tolerance having no bearing. Values between 0 and 1 allow both tolerance and context to influence the donation decision.

Our approach differs from typical approaches to achieving cooperation through trust and reputation, since there is less reliance on the existence of specific observations. Trust and reputation mechanisms typically assume that, taken together, a group of agents will have sufficient information about an individual's past behaviour to estimate its reputation [9, 13]. Such information is not guaranteed in a P2P setting, and so we use a general assessment of an agent's context, rather than attempting to assess an individual's cooperative nature. Our experimental results show that using this approach we can still achieve a significant improvement in cooperation. (Although certainly, if sufficient information was available to use a more standard reputation mechanism, then it would be likely to perform better.)

The second area in which we consider an alternative to RCA's ap-




Figure 3. Donation rate using “standard” reproduction.

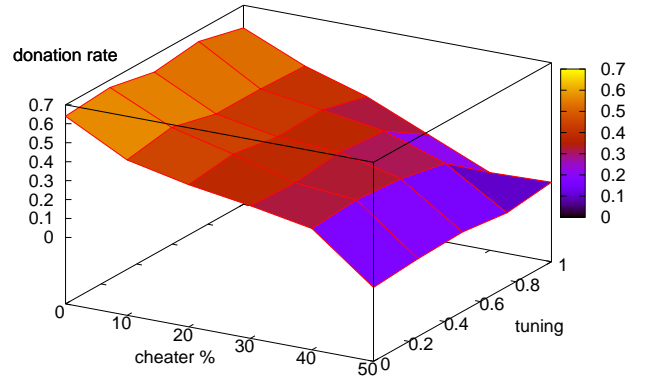


Figure 4. Donation rate using “context-based” reproduction.

proach is with respect to reproduction. In RCA’s model, after a certain number of interactions an agent will compare itself to another at random. If the other agent is more successful then its tag and tolerance values are copied (subject to minor mutations), i.e. the successful agent reproduces. In addition to replacing tolerance in the decision to donate by a combination of tolerance and context, we also consider using context for reproduction. If on comparison with another agent the other is more successful, its tag is copied, as is its assessment of its context. For the resulting offspring, the decision to donate becomes a consideration of a combination of its current context and its parent’s context. Thus, offspring A will donate to B if:

$$|\tau_A - \tau_B| \leq (1 - \vartheta) \cdot C_{parent(A)} + \vartheta \cdot C_A \quad (4)$$

where $C_{parent(A)}$ refers to A ’s parent’s assessment of its context (at time of reproduction), and ϑ is a tuning parameter that allows us to determine the influence of the current and parent’s context assessments. If ϑ is 1 then only the current context assessment is considered, while a value of 0 means that only the parent’s context is considered.

4 Results and Discussion

We have performed a number of simulations to investigate the effectiveness of our model, the influence of the tuning parameters (γ and ϑ), and the alternative reproduction mechanisms. Our simulations are built using the PeerSim P2P simulator³. We have experimented with various networks sizes, neighbourhood sizes and parameter values. Our simulations typically ran for 1500 generations, although longer simulations have been undertaken to check for long term stability. In this section we discuss the main findings based on our results. In the previous section, Figures 1 and 2 show how the donation rate evolves and oscillates over generations in a single simulation run. In this section, however, we average the donation rate over the whole simulation (1500 generations), and further take an average over 10 runs of the simulation. This allows us to compare donation rates for different configurations, without having to consider the inherent oscillations that the tag-based approach produces.

The initial tag and tolerance assigned to an agent in the simulation results presented here are randomly selected uniformly from

$[0, 1]$. We have also followed RCA and explored high initial tolerance ($T = 0.5$) and low initial tolerance ($T = 0.005$) settings. Our results mirror those found by RCA in that other than for short transients the end results are not substantially different from using a random initial tolerance [14].

The main characteristics that determine the donation rate in our model are the tuning parameters γ and ϑ for the ‘standard’ reproduction and ‘context-based’ reproduction approaches respectively. Figures 3 and 4 show how the donation rate is affected by the proportion of cheaters for standard reproduction and context-based reproduction, for varying values of the tuning parameters. The results are based on a network size $N = 100$ with each agent having $n = 49$ neighbours, a minimum observation threshold of $\sigma = 3$, and a history window size of $l = 5$. The standard deviation of the donation rate in all settings shown in these results is fairly low (below 0.05), showing that the donation rate achieved is fairly consistent.

As expected, higher proportions of cheaters significantly reduce the donation rate achieved using both reproduction approaches. Figure 3 allows us to compare the effectiveness of using context in the donation decision in comparison to RCA’s approach. Where the tuning parameter γ is set to 0 our model is identical to RCA’s, and as the proportion of cheaters rises from 0% to 50% the donation rate drops from around 0.35 to 0.15. It is clear from Figure 3 that increasing the influence of context in the donation decision, by increasing the tuning parameter γ , improves the donation rate. This improvement is most significant at low cheater proportions, but remains even at high cheater rates. Comparing Figure 4 and Figure 3 we note that using context-based reproduction also gives an improvement over RCA’s unmodified approach. Again, the improvement is most pronounced at low to medium cheater proportions (below 40%). However, for high proportions of cheaters (50%) using context-based reproduction actually gives a worse performance than RCA’s mechanism regardless of the tuning parameter ϑ . This effect can be better observed if we consider the donation rate for standard reproduction versus context-based reproduction for the extremes of the tuning parameters, as shown in Figure 5.

It is clear from Figure 5 that we can improve on the donation rate achieved by RCA’s unmodified approach (standard reproduction with $\gamma = 0$), shown by the solid line, for all settings of cheater proportion. The best results are obtained using our modification for consid-

³ <http://peersim.sourceforge.net/>

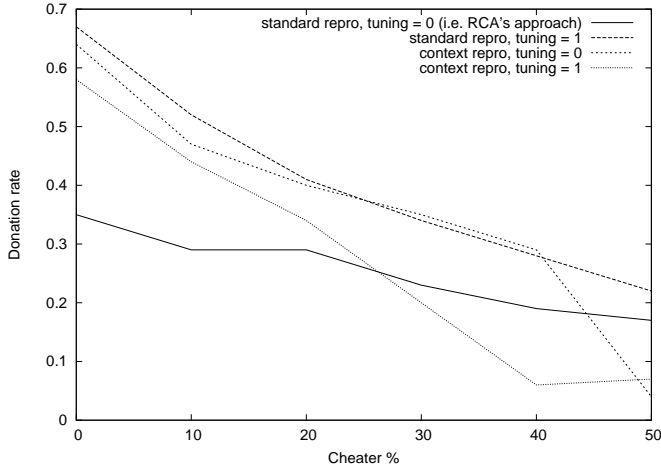


Figure 5. Donation rate for the tuning parameter extremes of both reproduction methods.

ering context in the donation decision rather than tolerance (tuning parameter $\gamma = 1$), but using RCA's standard reproduction method, shown by the upper dashed line. Context-based reproduction focusing on the parent's context ($\vartheta = 0$), the short-dashed line, instead of standard reproduction gives a slight reduction in performance for less than 20% cheaters, a very small increase for 20–40%, and a very significant decrease (much worse than RCA) for above 40% cheaters. Context-based reproduction where the current context is considered rather than the parent's context, shown as the dotted line in Figure 5, performs better than RCA, but worse than our other configurations, for cheater rates of around 0–25%. For rates above 25% it performs worse than all the other approaches.

From the results presented so far we can conclude that our modification to include context in the donation decision does give a significant improvement over RCA's approach. Using context-based reproduction focusing on the parent's context gives little advantage in low–medium cheater proportions (for $\vartheta = 0$) and performs worse than standard reproduction for high cheater proportions (or where the current context is emphasised with $\vartheta = 1$).

The effect of different history window sizes (l) on the donation rate is given in Figure 6. Note that for these results the number of pairings between reproduction cycles P was increased accordingly. It can be seen that effect of window size on donation rate is minimal for both reproduction methods (both tuning parameters are set to 1). There is a very small increase in donation rate when larger histories are considered, but the improvement is negligible. Furthermore, the memory overhead of maintaining longer observation histories for each neighbour is likely to outweigh the small improvement in donations for most settings.

We also consider the effect of an agent's neighbourhood size on the donation rate. Figure 7 shows the donation rate for both reproduction methods (with tuning parameters $\gamma = \vartheta = 1$) in a population of 10% cheaters. Neighbourhood size is shown as the percentage of the total nodes in the network N that are in an agent's neighbourhood, i.e. $n/N \times 100$. In this case the network size was restricted to 100 for efficiency of simulation, but we have obtained selected corresponding results for large networks of up to 2500 nodes (the current practical limit of our simulator's capabilities). It is clear that regardless of network size, using standard reproduction with context in the donation decision again outperforms the use of context in both the donation decision and reproduction. Higher donation rates

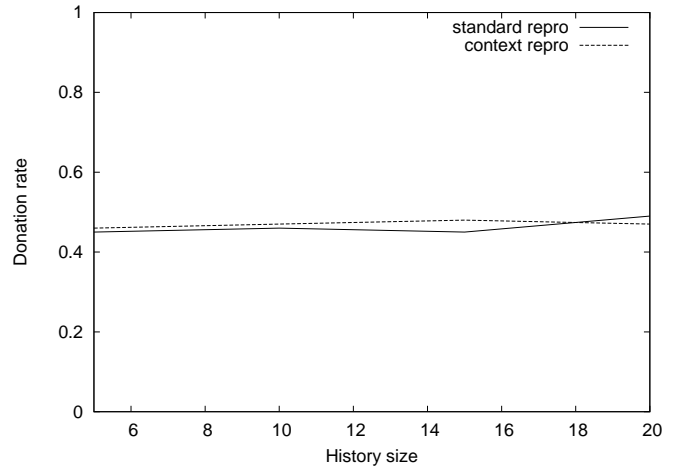


Figure 6. Donation rate for varying history window sizes.

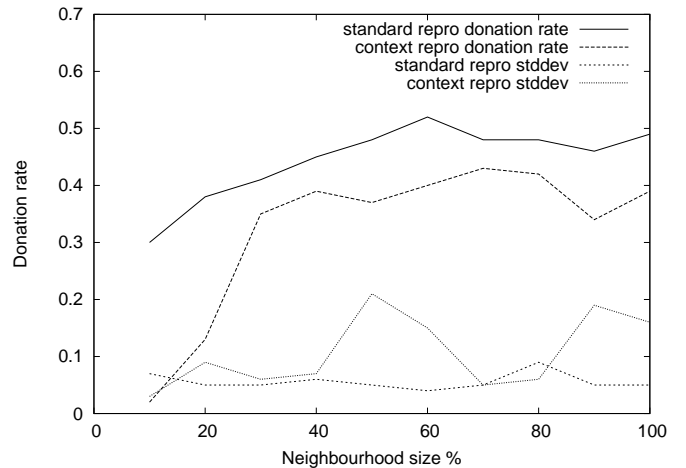


Figure 7. Donation rate for varying neighbourhood sizes.

are generally achieved for larger neighbourhood sizes. For the standard reproduction approach donation rate improves with neighbourhood size up to 60% (i.e. an agent has 60% of the network as neighbours), after which there is a slight decline in performance. It should be noted, however, that large neighbourhoods (e.g. above 40%) are likely to be impractical in most real-world systems, due to the large numbers of agents involved, and so the results for below 40% are the most relevant to real-world applications. For the context-based reproduction approach we also see a significant increase in donation rate as the neighbourhood is initially expanded. This increase again reduced for medium to large neighbourhoods, resulting in a slight decline for very large sizes. Figure 7 also shows the standard deviation of the runs used to obtain the donation rates. For the standard reproduction approach the standard deviation is fairly low and consistent (around 0.05). However, using context-based reproduction gives an inconsistent standard deviation (in the range of 0.05–0.2), illustrating the instability of this approach in comparison to using standard reproduction.

5 Conclusions

In this paper we have described a mechanism for establishing cooperation amongst agents without a reliance on reciprocity. Building on RCA's tag-based approach we have shown how incorporating an assessment of an agent's current context into the donation decision improves the donation rate. Context assessment is dependent on the extent of the interaction history recorded and on the neighbourhood size. Our results show that the history window size has minimal impact on donation rate, while increasing neighbourhood size does increase donation rate (at least for practical neighbourhood sizes below approximately 40%). We also considered an alternative to RCA's reproduction method, in which a parent's context was inherited by its offspring and subsequently used in donation decisions. Our results demonstrate that this alternative reproduction method was not effective. Overall, our simulations show that augmenting RCA's approach with context assessment for the donation decision is successful, and gives a significant increase in cooperation (of over 30% in some settings), but that RCA's standard method for reproduction is the most effective.

There are several areas of ongoing work. Primarily, we aim to explore a more sophisticated mechanism for assessing context, and to consider alternative methods for enabling offspring to use their parent's context assessment in the donation decision. Our aim is to investigate whether the donation rate can be further improved, without relying on reciprocity. Further in the future we will explore incorporating a simple trust model to exploit the limited reciprocity that exists, even in the kind of large scale environment we consider. Finally, we aim to simulate our approach in a more realistic P2P setting, such as the file-sharing example used by HE [7].

REFERENCES

- [1] R. D. Alexander, *The Biology of Moral Systems*, Aldine de Gruyter, 1987.
- [2] P. D. Allison, 'The cultural evolution of beneficent norms', *Social Forces*, **71**(2), 279–301, (1992).
- [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] B. Edmonds and D. Hales, 'Replication, replication and replication — some hard lessons from model alignment', *Journal of Artificial Societies and Social Simulation*, **6**(4), (2003).
- [6] N. Griffiths, 'Tags and image scoring for robust cooperation', in *Proceedings of the Seventh International Conference on Autonomous Agents and Multiagent Systems (AAMAS-08)*, (2008, to appear).
- [7] 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).
- [8] J. H. Holland, *Hidden order: How adaptation builds complexity*, Addison-Wesley, 1995.
- [9] 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).
- [10] M.A. Nowak and K. Sigmund, 'Evolution of indirect reciprocity by image scoring', *Nature*, **393**, 573–577, (1998).
- [11] M.A. Nowak and K. Sigmund, 'Evolution of indirect reciprocity', *Nature*, **437**, 1291–1298, (2005).
- [12] S. D. Ramchurn, D. Huynh, and N. R. Jennings, 'Trust in multi-agent systems', *Knowledge Engineering Review*, **19**(1), 1–25, (2004).
- [13] P. Resnick, R. Zeckhauser, E. Friedman, and K. Kuwabara, 'Reputation systems', *Communications of the ACM*, **43**(12), 45–48, (2000).
- [14] R. Riolo, M. Cohen, and R. Axelrod, 'Evolution of cooperation without reciprocity', *Nature*, **414**, 441–443, (2001).
- [15] G. Roberts and T. N. Sherratt, 'Does similarity breed cooperation?', *Nature*, **418**, 499–500, (2002).