Tags and Image Scoring for Robust Cooperation

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

ABSTRACT

Establishing and maintaining cooperation is an enduring problem in multi-agent systems and, although several solutions exist, the increased use of online trading systems, peer-to-peer networks, and ubiquitous computing environments mean that it remains an important question. Environments are emerging in which large numbers of agents are required to cooperate, but where repeat interactions between agents may be rare or non-existent. Most existing approaches to cooperation rely on reciprocity to establish notions of trust and reputation. However, where repeat interactions are rare such approaches are not always effective. In this paper we use ideas from biology and the social sciences to provide a mechanism that supports cooperation in such environments. Our mechanism combines a tag-based method to enable cooperation given a lack of reciprocity, with an adaptation of a simple image scoring reputation model to cope with cheating agents. Using a simple peer-to-peer scenario we show how cooperative behaviour is favoured, and how the influence of cheating agents can be reduced using only minimal information about an agent's neighbours.

Categories and Subject Descriptors

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

General Terms

Experimentation, Algorithms, Reliability

Keywords

Cooperation, Tags, Reputation, Image Scoring, Evolution, P2P

1. INTRODUCTION

The question of how cooperation is established and maintained is an enduring problem for understanding biological systems and social networks. A similar issue exists in multi-agent systems where mechanisms are needed to support cooperation between autonomous agents. Although several successful approaches exist for some environments, the increasing use of online trading systems, peer-to-peer (P2P) networks, and ubiquitous computing environments mean that enabling and maintaining cooperation remains an important question. Such environments typically have a large number of agents that are required to cooperate, but repeat interactions between agents may be rare. The majority of existing approaches for establishing cooperation do not completely fit, since little may be known about the 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.

Reciprocity is the basis for the majority of existing approaches to cooperation — 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 historically the most common, approach where two agents have repeat interactions in which there is the opportunity to "cooperate" or "defect". In recent years economists, social scientists, and computer scientists have become increasingly interested in indirect reciprocity, where a third party is involved in repeat interactions — agents are not likely to directly have repeat interactions, but are likely to interact with others whose behaviour with third parties they have previously observed. Nowak and Sigmund [14] 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". This increased interest in indirect reciprocity is due in part to a growth in global markets and large scale systems in which one-shot interactions between (relatively anonymous) partners become the norm. Indeed, it may also be necessary to engender cooperation where no reciprocation of any form exists, for example if there is no memory of past encounters [17]. Where transactions take place online between strangers, the success and robustness of a marketplace is traditionally based to a large extent on the ability to utilise whatever reciprocity exists in order to establish some form of reputation and trust [10, 16].

Many models exist for supporting cooperation between agents using trust and reputation [10, 11, 15, 16]. However, the majority target scenarios in which agents have ongoing repeat interactions, and so they use direct reciprocity to ground cooperation. In this paper we propose a model for establishing and maintaining cooperation in situations where

Cite as: Tags and Image Scoring for Robust Cooperation, Nathan Griffiths, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. 575-582.

repeat interactions are rare and direct reciprocity is not appropriate. We use P2P networks as an illustrative domain, but our approach is generic and applicable in a range of cooperative applications in which direct reciprocity cannot be used. In the following section we describe the theoretical approaches upon which our model is based, along with the promising initial results obtained by others. In Section 3 we introduce the P2P context for our model, and the model itself is introduced in Section 4. Our experimental setting and selected experimental results are described in Sections 5 and 6. Finally, Section 7 concludes the paper.

2. RELATED WORK

Indirect reciprocity is not an novel idea: biologists and social scientists have long considered a kind of cooperation that does not require individuals to directly meet again, but where cooperative strategies are favoured [1, 3, 13]. Furthermore, theoretical models of cooperation exist that do not require any reciprocity, but instead are based on the recognition of cultural artifacts or "kin" recognition [2, 4]. Previously, the agent community has tended to focus on scenarios where direct or indirect reciprocity exists, and numerous successful models of cooperation based on trust and reputation have been developed [10, 16]. In large open environments, however, very little reciprocity may exist and existing models of trust and reputation become less appropriate. In a P2P system, for example, agents that interact may never directly interact again, and may only have a very small number of interactions with, or observed by, the same third party (as required for indirect reciprocity). Therefore, a mechanism for cooperation is needed that does not rely on reciprocal interactions.

In recent years, there have been promising results for an approach to cooperation that uses "tags" [9] as cultural artifacts to engender cooperation where no reciprocity exists [17], which has in turn led to a technique to improve cooperation in P2P networks [8]. Note that tags in this context are simple observable traits that are ascribed to each agent, and the notion is unrelated to the folksonomy use of the term "tags" for collaborative and social content classification. 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. However, before we discuss how to combat the problem of cheaters, we introduce the tag-based approach to cooperation.

Riolo, Cohen and Axelrod (RCA) 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 [17]. Their approach can be viewed as related to Nowak and Sigmund's (NS) model of "image scoring" in which a simple reputation-like mechanism allows agents to only cooperate with others that are known to be sufficiently generous [13]. Image scoring gives rise to indirect reciprocity, since an agent's actions are observed and directly influence whether others in the future will be cooperative towards it. Tag-based cooperation, however, is independent from past or future interactions within the current generation[1] — image depends on past actions, while tags do not.

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, 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 [17]. When an agent donates it pays a cost c, and the recipient receives a benefit b (it is assumed that $b > c$). RCA have performed simulations in which each agent acts as a potential donor in P interaction parings, after which the population of agents is reproduced in proportion to their relative score. 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 cycles 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 so take over the population, lowering the overall rate of cooperation. Eventually, the mutant tag becomes the most common and cooperation again becomes the norm [17]. RCA's approach is an effective mechanism for achieving cooperation without relying on reciprocity, but their model relies on the 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 [8]. 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) [17]. The second change is that HE interpret a tag as an agent's set of neighbours in a P2P network. Thus, adopting another agent's tag is equivalent to re-wiring the P2P network such the other agent's connections are adopted [8]. 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 we are interested in networks in which the population may include cheaters, and where agents cannot re-wire the network. The assumption that agents cannot re-wire the network (i.e. the topology is fixed) is for reasons of simplicity. Our approach, therefore, will be based on RCA's model, and HE's application of it (minus the re-wiring), supplemented by a mechanism to cope with cheaters.

A number of other tag models have been proposed [6], using alternative formulations of tag representation, and matching and mutation processes. Hales and Edmonds [7]

[1]Where agents reproduce, an offspring's tag will typically depend on the success of the parent's actions, however, tags are independent from action within the current generation.

represent tags as fixed length bit strings, and consider how they can be used to achieve cooperation where a recipient must find a potential donor with an appropriate resource, rather than finding an arbitrary donor. Their results are promising, and led to the P2P-based model described above [8]. However, as with other approaches the presence of cheaters is not considered. Matlock and Sen (MS) [12] add the flexibility for agents to cooperate with other tag groups, by generalising the tag matching mechanism to enable cooperation between individuals with different tags. This is achieved through the use of "tag matching patterns" (where tags are a string of bits against which pattern matching is performed), payoff sharing such that agents can share payoffs with their "opponent", and a more sophisticated reproduction and mutation mechanism that preserves tag matching using patterns. In this paper we also allow for cooperation between individuals with different tags, as in RCA's model, but as per RCA we take these agents to be part of the same group (i.e. sharing the same tag within some tolerance threshold). We are not concerned in this paper with enabling cooperation between different social groups. Our focus is on addressing the issue of cheaters, a problem not considered by MS. We see our work as complementary to MS, since we address a different problem. Indeed, future work may investigate whether our approach to cope with cheaters is effective in setting described by MS.

Many sophisticated models of reputation have been proposed by agent researchers, which have proved successful in a variety of domains [10, 16]. However, they tend to assume that even where there is no direct reciprocity, there will be many interactions with a common third party, meaning that a high level of indirect reciprocity exists. Such models also tend to require agents to maintain a relatively large amount of information. In the context of mathematical biology, Nowak and Sigmund (NS) [13] have developed a simple mechanism for reputation using indirect reciprocity, that avoids the complexity inherent in typical agent-based approaches. In their model each agent has an image score s that is known to all agents. If a donor helps a recipient its image score is incremented, and its score is decremented if it declines to help. Each agent has a strategy k such that it will help others whose image score is above k, i.e. agent i will cooperate with j if $s_j \geq k_i$. After a certain number of interactions the agents are reproduced, with the number of offspring proportional to an agent's success. In simulations, NS find that for certain k values, cooperation will become prevalent.

In this paper we describe how RCA's tag-based approach can be used in a P2P setting, where cheaters exist and rewiring is not possible. We use a reputation mechanism based on NS's model to cope with cheaters, but we extend their method to reduce the requirement for indirect reciprocity. In particular, an agent uses an estimate of the combined image score of its neighbours to reduce the impact of cheaters.

3. THE P2P CONTEXT

How best to achieve cooperation in a given domain depends in part on the characteristics of that domain. In this paper we use P2P networks as an illustrative scenario, and although our approach is fairly generic, our discussion will focus on a P2P setting. We assume that there is a large number of agents in comparison to the number of interactions, meaning that repeat interactions are rare. Thus, direct reciprocity is negligible, and although indirect reciprocity exists it is not always guaranteed (meaning that there may not exist a common third party with whom any two agents have interacted or that has observed their actions).

We are concerned with establishing cooperation in 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 agents are not able to re-wire the network, or make use of network overlays (as is HE's approach). Furthermore, unlike RCA, HE, and NS we assume that a proportion of the population will be cheaters, meaning that they will receive 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 and NS 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 NS's parameter values of $b = 1$ and $c = 0.1$ (the addition of a cost of 0.1 is to avoid negative payoffs) [13]. 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 [8].

4. COMBINING TAGS AND IMAGE SCORING

Our mechanism is founded upon RCA's tag-based approach, but we incorporate a simple mechanism to combat cheaters that is inspired by NS's image scoring method. 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's tolerance threshold, namely $|\tau_A - \tau_B| \leq T_A$. We use RCA's learning interpretation of reproduction (i.e. that used by HE) where after a certain number of interactions an agent compares itself to another selected at random. If the other agent is more successful then its details are copied (meaning that the other agent reproduces), otherwise no change is made. There are two ways of determining when to reproduce: *interval-based* or *rate-based*. For interval-based reproduction an agent reproduces after a fixed number P of interactions (this is the approach taken by RCA, HE and NS). In rate-based reproduction we specify a reproduction probability, which represents the probability that an agent will reproduce after each interaction. In this paper, for simplicity, we concentrate on interval-based reproduction (however, we have obtained similar experimental results using a rate-based approach). After reproduction there is a potential mutation of the offspring's tag and tolerance, with probabilities m_τ and m_T respectively. Specifically, with probability m_τ a new (randomly selected) tag is received, and with probability m_T the tolerance is mutated by the addition of Gaussian noise with mean 0 and a small standard deviation (a value of 0.01 was used in these results).

In common with RCA we find that relatively stable donation rate (meaning cooperation) is established over a large

[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 [18]. The results presented here permit this small negative tolerance.

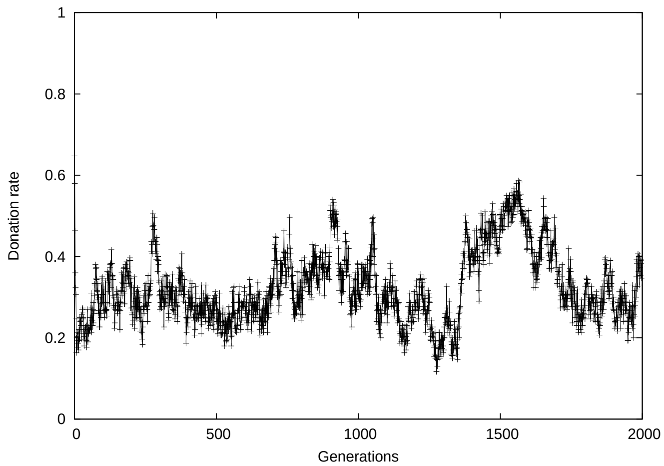

Figure 1: Donation rate with no cheaters using RCA's approach.

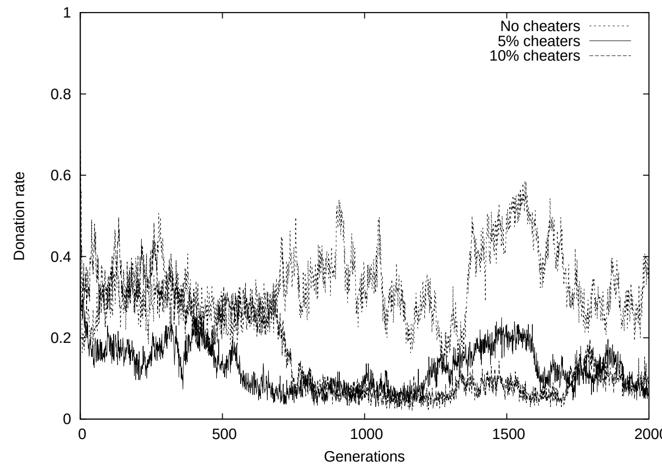

Figure 2: The effect of cheaters on donation rate with RCA's approach.

number of generations for appropriate parameters, provided that cheaters are not introduced into the population. As per RCA's results, the donation rate fluctuates with each generation, but across generations the average behaviour is cooperative. Figure 1 shows the dynamics of the cooperation 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 probably 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 [17]. 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].

Our method thus far is identical to RCA's (parameter values aside) and 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) soon results in a fairly stable state of 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, and for relatively small proportions of cheaters the average donate rate is significantly reduced.

RCA's method makes no use of any reciprocity that may exist in the environment. In our setting of P2P systems, there is potential for a (limited) degree of indirect reciprocity in that an agent has a restricted neighbourhood with whom to interact, and there may be repeat interactions via a third party. As noted by Alexander [1], indirect reciprocity can be thought of as originating from direct reciprocity in the presence of an interested audience. In our P2P setting we assume that the number of interactions relative to the number of neighbours is small, meaning that there is little direct reciprocity between a given pair of agents. However, for simplicity we also assume that agents cannot re-wire the network, and so have a fixed neighbourhood of potential interaction partners. Thus, for any interaction, the neighbourhood within which it takes place represents an interested audience, since an agent in that audience may in the future be required to interact with one of the agents involved in the interaction. This notion of an interested audience is the basis of NS's image scoring mechanism to establish indirect reciprocity. The premise of their approach is that an agent should only donate to others that have been observed to be sufficiently generous. In NS's method all agents effectively observe all interactions, via each agent (honestly) maintaining a public image score. An agent's strategy then determines the minimum value of image score to which it will donate. We cannot directly use their method for two reasons: first, agents have a restricted neighbourhood and so cannot observe all interactions, and second because we allow for the presence of cheaters (who one might assume would be dishonest in publishing their image score). We therefore take the essence of NS's approach — using the observations of interested parties to establish indirect reciprocity — and re-cast it in our P2P environment.

In our environment an agent has a fixed set of n connections to its neighbours, and these neighbours may be able to observe its actions. Assume that any given neighbour might observe an interaction with some probability p_o. If the probabilities for all neighbours are equal to 1 then an interaction is observed by all n neighbours, and conversely probabilities of 0 imply that no neighbours observe. Note that this assumption of the potential for observation 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 others have forwarded packets. In reality the probability of observation is likely to change over time for each individual agent, according to the activities that it is involved in. For simplicity, however, we assume that p_o is fixed, and is the same for all agents.

Based on the observations of its neighbours' activities thus far, an agent is able to build a (limited) picture of the cooperative behaviour of its neighbours. Each agent has a simple memory recording (up to) the last l donation interactions that it has observed its neighbours take part in. For each interaction a value of $+1$ is recorded to represent a donation and -1 to represent refusal to donate. 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. Using this information, an agent can estimate the general donation behaviour of its neighbours. It is important to note that this memory is fairly sparse, since the number of interactions is relatively small compared to the number of agents. This differs from other approaches to reputation in which it is typically assumed that, taken together, a group of agents will have sufficient information about an individual's past behaviour to estimate its reputation [10, 16]. The overhead of maintaining such a memory is fairly small, with an upper bound of just $n \times l$ values, which is smaller than typical reputation mechanisms.

The decision to donate in RCA's approach is determined by an agent's tolerance, i.e. A will donate to B if $|\tau_A - \tau_B| \leq T_A$. In order to make use of its current observations, we allow an agent to modify its tolerance at run time, for each interaction. The general principle used by the majority of existing reputation approaches is that if a potential recipient is thought to be cooperative then one is more likely to be cooperative towards it, i.e. tolerance is increased. Such approaches typically use information obtained from third parties about the potential recipient B's behaviour. However, in our setting we are assuming that the number of interactions is small in relation to the number of agents. Therefore, we cannot rely on specific information from third parties about the potential recipient, because such information will often not be available. Instead, we relax the general principle such that if an agent is situated in a neighbourhood that is thought to be cooperative then tolerance is increased. In other words, we change the emphasis from trying to make assessments about specific individuals, to assessing the general context in which an agent is situated. Our aim is to provide a mechanism in which if an agent would cooperate using RCA's approach then it will also do so in our mechanism, but in a cooperative setting it may cooperate in our approach where in RCA's it would decline.

In order to determine the tolerance level to use, an agent considers the observations it has made so far of each agent in its neighbourhood. From these we calculate a change component for each neighbour, by summing the observations recorded scaled by the extent of observations made. Recall that the observations recorded are $+1$ for a donation and -1 for a refusal to donate. If we simply sum the observations, then the change component for an agent would be negative if it has declined to donate more often than it has donated. When the change components from all neighbours are combined together (as defined below in Equation 2) this could result in a reduction of tolerance. This is inconsistent with the general principle outlined above of increasing tolerance in neighbourhoods that are perceived to be cooperative — we do not wish to reduce tolerance, only to increase it where appropriate. Thus, we avoid negative change components by defining the change component δ_n for neighbour n as:

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

where o_n^j represents the j'th observation of n, and l_n denotes the number of observations made of n's interactions as a donor (thus $l_n < l$). We can tune the model using the parameter λ, which determines the extent to which we scale an agent's tolerance. A value of $\lambda = 0$ means that the model is identical to RCA's approach, while a value of $\lambda = 1$ implies that the new tolerance could be increased up to double the original value. (Values of λ above 1 are not excluded, but we find that they have little additional effect.)

Once the change components for each neighbour have been calculated they can be combined to give an overall change in tolerance. For agent A the updated tolerance T'_A is given by:

$$T'_A = T_A + \left(T_A * \frac{\sum_{i=1}^{n} \delta_i}{n}\right) \quad (2)$$

The decision to donate is now based on this new tolerance value. Thus, A will donate to B if $|\tau_A - \tau_B| \leq T'_A$.

This model does not use indirect reciprocity in the usual manner, since there is less reliance on the existence of specific observations. Typical reputation approaches to combating cheaters tend to assume that sufficient information exists to estimate the cooperativeness of a particular agent. Such information is not guaranteed in a P2P setting, and so our model relaxes the information needed, by using an general assessment of an agent's neighbourhood, rather than assessing an individual. Our experimental results show that using this approach can still provide 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.)

Our model provides a number of parameters that determine the rate of donation that is achieved in a given setting. These are, the maximum length of the observation memory l, the probability with which an interaction is observed p_o, and the scaling factor λ, that determines the extent to which tolerance is adjusted. Several characteristics of the environment also determine the rate of donation, and in a practical application these are typically outside our control. Specifically, the network size N in relation to the number of neighbours n has a significant influence, along with the proportion of cheaters in the environment. Also of significance is the number of interactions an agent enters into before reproduction, i.e. the interval size P or the reproduction rate used. This number of interactions strongly influences the extent to which observations can be made. More interactions between generations implies that more observations can be made, meaning that there is more likelihood of having observed a given agent, and (towards the end of a generation) it is likely that an agent's memory will be populated (i.e. $l_n = l$ for neighbour n). We have performed a number of simulations to explore the effectiveness of our model, and the influence of the various parameters. In the remainder of this paper we summarise our findings.

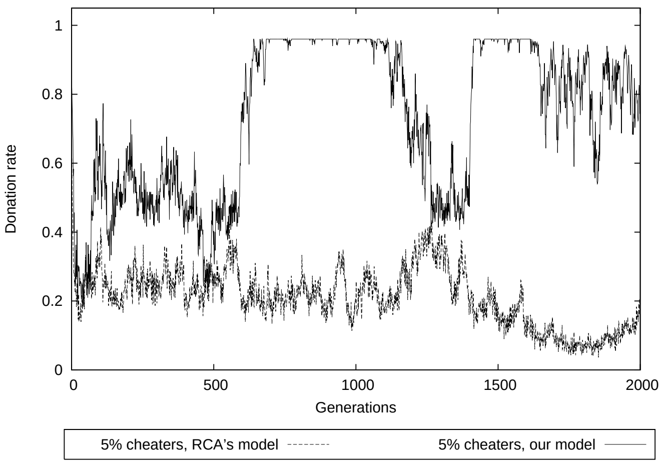

Figure 3: Donation rate for our model and RCA's approach with 5% cheaters.

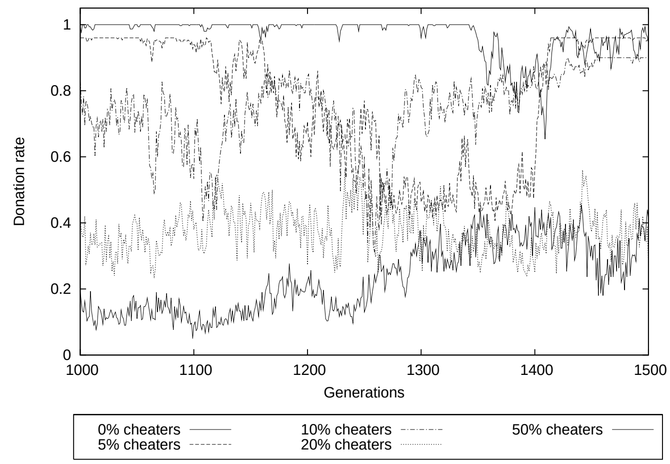

Figure 4: Donation rate as the proportion of cheaters in the population is varied.

5. SIMULATION ENVIRONMENT

We have undertaken simulations using several configurations of network to investigate our model. Our simulations have been built using the PeerSim P2P simulator[3]. We have experimented with networks of ranging from 100 to 5000 nodes, where each node has a (randomly assigned) neighbourhood varying between 5% and 100% of the population. Our simulations ran for between 500 and 5000 generations, using both interval reproduction and rate-based reproduction. The results show little difference between the two reproduction approaches, and for simplicity of presentation we focus on interval reproduction in this paper.

The initial tag and tolerance assigned to an agent are randomly selected uniformly from $[0, 1]$, although we have also used RCA's method of exploring high initial tolerance ($T = 0.5$) and low initial tolerance ($T = 0.005$) settings. In this respect our results mirror those found by RCA, in that other than for short transients the end result is not substantially different from using a random initial tolerance [17].

In the following section we present illustrative results that show the effectiveness of our model. The results shown are representative runs obtained from many runs of the simulator (typically 10 per configuration). Due to the random initial allocation of tolerances and tags the results obtained vary from run to run, but with the exception of a very small number of outliers the results mirror those included here. The figures presented in this section show how the donation rate evolves over individual simulation runs. We also give the average donate rate across all generations, averaged over 10 runs of the simulation.

6. RESULTS AND DISCUSSION

Before we consider the effect of the various environment and model parameters, we begin showing the effectiveness of our approach in the 5% cheater setting shown earlier (in Figure 2). Figure 3 shows the donation rate for both our model and RCA's approach in the presence of 5% of the population being cheaters. The environment parameters are identical to those used to generate the results in Figure 2. We use a network size of $N = 100$, a neighbourhood of $n = 49$, and set the tuning parameter as $\lambda = 1$. It can be seen that after around 250 generations the population becomes significantly more cooperative than where RCA's unmodified approach is used (the dotted line). The donation rate still undergoes significant fluctuations, reducing as mutants with lower tolerances take advantage, until the population re-stabilises with a new dominant tag cluster. This effect was also observed, and is discussed in detail, by RCA [17].

Figure 4 shows an extract of 500 generations from runs using our model with various proportions of cheaters in the environment. We show an extract of only a small number of generations for clarity of presentation (i.e. we show generations 1000–1500 rather than 0–2000 in Figure 3). A base case of 0% cheaters is shown (plotted as the upper solid line), which compares very favourably to the donation rate achieved by RCA's approach given 0% cheaters (as shown in Figure 1). The lowest line on the graph (shown as a solid line) is for an environment in which 50% of the agents are cheaters. In this case our approach achieves a donation rate of approximately 27% over all generations (once the simulation has stabilised, after around 250 generations). This compares with approximately 11% achieved using RCA's model without modification (which is not plotted in Figure 4 for clarity reasons). We also include the donation rate for 5% cheaters (the upper dashed line), 10% cheaters (the dotted and dashed line) and 20% cheaters (the dotted line). When averaged across generations for several runs of the simulation we get the following donation rates.

	Cheater Proportion				
	0%	5%	10%	20%	50%
Our model	0.95	0.69	0.61	0.41	0.27
RCA's model	0.62	0.18	0.16	0.17	0.11

The results show that as more cheaters are introduced the donation rate declines, as is expected. However, the relationship between the number of cheaters and donation rate is not linear. As a small number of cheaters are introduced into the population they have a relatively large effect on donation (e.g. a population with 5% cheaters has much lower donation rate than 0% cheaters). Conversely, for

[3]http://peersim.sourceforge.net/

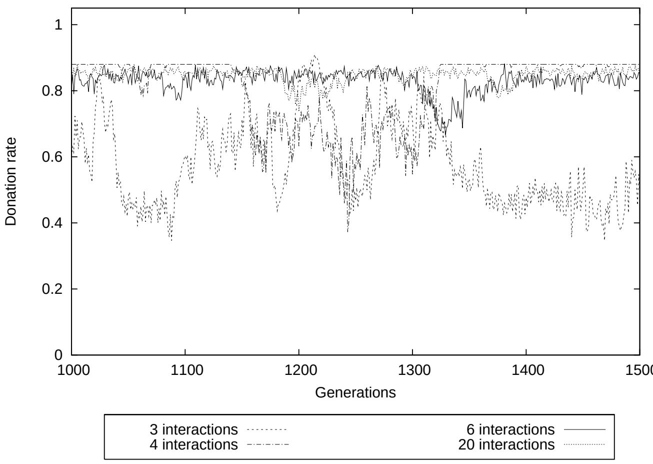

Figure 5: The effect on donation rate of the number of interactions (and memory length) before reproduction.

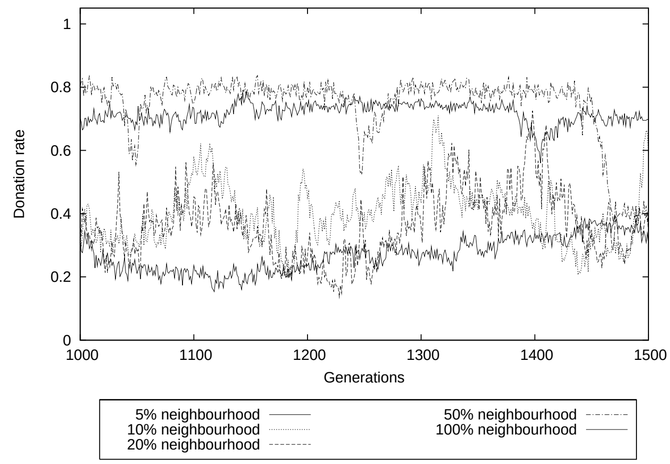

Figure 6: The effect of neighbourhood size on donation rate.

higher proportions of cheaters a small increase in the number of cheaters has a relatively small effect (e.g. the rate for 25% is only slightly less than as shown here for 20%).

The number of observations that are used to calculate an agent's tolerance has a fairly large effect on the donation rate. Figure 5 illustrates this by considering different values for the number of interaction pairings P between each reproduction cycle, and setting the memory length to the maximum possible value, i.e. $l = P$, and allowing an agent to observe all interactions in its neighbourhood ($p_o = 1$). Note that if $l < P$ then an agent would only record the last l interactions regardless of the number of pairings. Other parameters are as for previous results. In this setting we have allowed 10% of the population to act as cheaters. The lowest donation rate of around 60% is achieved where there are only 3 interactions per generation (the lower dashed line), while 20 interactions gives the highest rate (the dotted line) of 82%. Using 4 interactions (the dotted and dashed line) gives an improvement to approximately 65% on average. However, this improvement undergoes major oscillations — reducing to similar rates as when less interactions are used, but with periods of over 85%. As we increase the number of interactions to 6 (the solid line) we again see an improvement, to 77%, which is much more stable. Again, the donation rate oscillates, but to a much lesser extent. The average donation rate across several simulation runs are as follows.

Interactions	3	4	6	20
Donation Rate	0.60	0.65	0.77	0.82

We have performed other simulations in which we vary the memory length l independently from the number of parings P, and consider alternative values for the observation probability p_o. Overall, our results show that as the number of observations recorded increases, so does the donation rate. However, the relationship between these parameters is not straightforward, and our results require more analysis before conclusions can be drawn.

We have also considered the effect of neighbourhood size (as a proportion of network size) on the donation rate. We used a cheater population of 10% and $P = 3$ interaction pairings per generation to obtain the results shown in Figure 6. Agents were permitted to observe all of their neighbours' interactions, i.e. $p_o = 1$. The lowest donation rate of 38% was achieved when an agent's neighbourhood size was 5% of the network size, i.e. $n/N = 0.05$. This compares with an donation rate of approximately 15% using RCA's unmodified approach. As the neighbourhood size is increased, we initially get a small corresponding increase in donation rate. Neighbourhoods of 10% and 20% result in donation rates of 39% and 40% respectively (shown in Figure 6 by the dotted line and dashed line). A neighbourhood of 50% gives a much larger increase (as shown by the dotted and dashed line) to 74%. The maximum donation rate of 85% was achieved where an agent could observe all other agents in the network (shown by the solid line). This is analogous to the setting used by NS in which each individual (honestly) maintains a public image score. Note that configurations with over 40–50% neighbours are not realistic for a practical P2P setting since the number of agents, and so the number of potential neighbours, is large. In a real-world P2P application of the model the neighbourhood size is likely to be fairly small, but these results demonstrate that improvements can be made where observations are limited (e.g. the 5% case), but that there is benefit in trying maximise the neighbourhood size that is observed. The average donation rate across several simulation runs are as follows.

Neighbourhood	5%	10%	20%	50%	100%
Donation Rate	0.38	0.39	0.40	0.74	0.85

7. CONCLUSIONS

In this paper we have proposed a novel mechanism that supports cooperation amongst agents where there is minimal indirect reciprocity. Building on RCA's tag-based approach we have shown how agents can use simple observations of their neighbours' behaviour to assess the extent to which their environment is cooperative. As the number of observations is increased, agents are able to improve the donation rate in the population. Our mechanism has less reliance on indirect reciprocity than existing agent reputation systems, and is more realistic in scenarios where agents might cheat than NS's image scoring approach on which it is based. We have illustrated the effectiveness of the method though simulation, and presented the main results in this paper.

There are several areas of ongoing work. Of highest priority is to continue our experimental evaluation to understand more about the relationship between the parameters that determine an agent's behaviour. Specifically, although we have shown that the donation rate increases as the number of observations recorded increases, the relationship between memory length l, the number of parings P, and the observation probability p_o are not yet fully understood. Our secondary concern is to explore the effect of alternative reproduction biases and "tag clone" donation behaviours (as suggested by [5]), along with different scheduling methods in our simulation to further analyse the reasons behind the need for lower mutation rates in our implementation compared to RCA's. (PeerSim supports both cycle and event driven simulations. Our current implementation is cycle driven, and we intend to re-implement the simulation in an event driven manner.) We also aim to explore the effect of introducing re-wiring into the network, such that agents have a different set of neighbours for each generation, and of creating larger networks with smaller neighbourhood sizes. Finally, we aim to re-cast our simulation environment in a more realistic P2P setting, for example the file-sharing example used by HE. Future work will also consider alternative definitions for the tolerance update function, including reducing tolerance in certain contexts as well as increasing it.

8. ACKNOWLEDGEMENTS

The author would like to thank the developers of the PeerSim P2P simulator, which facilitated the development of our simulations, and the anonymous reviewers for their useful comments which informed revision of the paper.

9. 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] D. Hales. Understanding tag systems by comparing tag models. In B. Edmonds, C. H. Iglesias, and K. G. Troitzsch, editors, *Social Simulation: Technologies, Advances and New Discoveries*. Idea Group, 2007.

[7] D. Hales and B. Edmonds. Evolving social rationality for MAS using tags. In *Proceedings of the Second International Conference on Autonomous Agents and Multi-agent Systems (AAMAS-03)*, pages 497–503, 2003.

[8] 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.

[9] J. H. Holland. *Hidden order: How adaptation builds complexity*. Addison-Wesley, 1995.

[10] 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.

[11] R. M. Kramer. Trust and distrust in organizations: Emerging perspectives, enduring questions. *Annual Review of Psychology*, 50:659–598, 1999.

[12] M. Matlock and S. Sen. Effective tag mechanisms for evolving coordination. In *Proceedings of the Sixth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-07)*, pages 1345–1352, 2007.

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

[14] M. Nowak and K. Sigmund. Evolution of indirect reciprocity. *Nature*, 437:1291–1298, 2005.

[15] S. D. Ranching, D. Huynh, and N. R. Jennings. Trust in multi-agent systems. *Knowledge Engineering Review*, 19(1):1–25, 2004.

[16] P. Resnick, R. Zeckhauser, E. Friedman, and K. Kuwabara. Reputation systems. *Communications of the ACM*, 43(12):45–48, 2000.

[17] R. Riolo, M. Cohen, and R. Axelrod. Evolution of cooperation without reciprocity. *Nature*, 414:441–443, 2001.

[18] G. Roberts and T. N. Sherratt. Does similarity breed cooperation? *Nature*, 418:499–500, 2002.