Understanding Tag Systems by Comparing Tag models[1]

David Hales
Department of Computer Science
University of Bologna, Italy.
dave@davidhales.com

Abstract

Since Holland (1993) introduced the concept of tags a number of tag models with intriguing and potentially very useful, properties have been advanced. However there is currently little understanding as to the exact mechanisms that produce these results. Specifically it is not know what (if any) are the necessary conditions for tag systems to produce high levels of cooperation in social dilemmas. In this paper by comparing existing tag models to formulate a hypothesis and then using simulation we identify what appears be a necessary condition for high cooperation. Previous tag models implicitly contained the condition but authors did not identify the significance of it.

1. Introduction

Tags are markings or social cues that are attached to individuals (agents) and are observable by others (Holland 1993). They evolve like any other trait in a given evolutionary model. The key point is that the tags have no direct behavioral implication for the individuals that carry them. Through indirect effects, however, they can evolve from initially random values into complex ever changing patterns that serve to structure interactions between individuals.

Riolo (1997) showed how tags could boost cooperation in a scenario involving agents playing the iterated prisoners dilemma (IPD). Agents bias their game playing towards individuals with similar tags (the indirect effect). In these studies tags were represented by a single real number attached to each agent.

Subsequently Hales (2000) advanced a model, using binary tag strings that demonstrated the evolution of cooperative interactions in the single round Prisoners

[1] This work partially supported by the EU within the 6th Framework Programme under contract 001907 (DELIS).

Dilemma (PD). Further work (Riolo et al, 2001) showed the emergence of altruistic giving behavior and the evolution of cooperation and specialization (Hales 2002)[2].

These latter models are important because they advance a novel mechanism for evolving coordinated and cooperative interactions between unrelated agents that have no knowledge of each other and have never met previously (i.e. strangers). This obviates the need for repeated interactions (Tivers 1971), "genetic" relatedness (Hamilton 1964), "image scoring" (Nowak and Sigmund 1998) or strict spatial relationships (Nowak and Sigmund 1992) in the production of cooperation. Tag mechanisms therefore have potential engineering applications where these other methods are not applicable (Hales and Edmonds 2003).

Although the general mechanism by which tags produce these results appears to be the result of a dynamic group formation and dissolution process (Hales 2000, Riolo et al 2001, Sigmund and Nowak 2001) with selection at the group-level, there has been little analytical or empirical exploration of this hypothesis. Indeed it is not even currently understood what the necessary and / or sufficient conditions might be to produce tag systems what give rise to these properties of interest (other than the specific existence proofs of the simulation results presented). In this paper we begin this process.

2. Paper Outline

In this paper we identify what appears to be a necessary condition that all previous models implicitly contained. In each case the authors had not identified this property as significant, yet without it the phenomena of interest disappears. We report the results of computational simulations that demonstrate the necessity of the condition and begin to sketch out a way towards analytically capturing the condition.

The necessary condition is that the mutation rate of the tag must be much higher than the mutation rate applied to any behavioral traits. In this way cooperative "groups" (agents sharing the same or similar tags and interacting cooperatively with each other) can be "cloned" before being invaded by exploitative mutants that "kill" or "dissolve" the group. We demonstrate this by varying a parameter (the tag / action trait mutation ratio) over many of runs of a simulation model and measure cooperation. The result is a (non-linear) sigmoid-like relationship, indicating a transition threshold for the relative mutation rate in a given system.

Since recent work (Hales and Edmonds 2003 and in press) has indicated how tag mechanisms might be applied to the solution of complex engineering problems a deeper understanding of the necessary and sufficient conditions of application would be timely. Such mechanisms have application in self-organizing adaptive Peer-2-Peer networks (Hales and Edmonds in press) and distributed and spontaneously self-

---

[2] It should be noted that the conclusions of these further studies have been questioned (Roberts and Sherrat 2002, Edmonds and Hales 2003). Essentially the scenarios do not bear too close a comparison to a PD because there is no dilemma.

organizing mobile agent based applications (where issues of trust and cooperation are paramount but can not easily be dealt with using traditional techniques).

3. Some Previous Tag Models

There have been a number of tag (simulation) models implemented each demonstrates how higher-than-expected levels of cooperation and altruism are produced when tags are employed. In all cases the models implement evolutionary systems with assumptions along the lines of the replicator dynamics (i.e. reproduction into the next generation proportional to utility in the current generation and no "genetic-style" cross-over operations but low probability mutations on tags and strategies).

Riolo (1997) gave results of studies applying tags in a scenario where agents played dyadic (pair wise) Iterated Prisoner's Dilemma games (IPD). Tags (represented as a single real number) allowed agents to bias their partner selection to those with similar tags (probabilistically). He found that even small biases stimulated high levels of cooperation when there were enough iterations of the game with each pairing.

In Hales (2000) a tag model was applied to a single round PD, one where pairings resulted in a single game of PD. Tags were represented as binary strings. Pairing was strongly biased by tag identity (rather than probabilistic similarity). In this model very high levels of cooperation were produced between strangers in the singe round game if the binary tag strings were long enough.

In Riolo et al (2001) a tag model was applied to a resource-sharing scenario in which altruistic giving was shown to emerge. Agents were randomly paired (some number of times) and decided if to give resources or not. The decision to give was based on tag similarity mediated by a "tolerance gene" as well as the "tag gene" (both represented by real numbers). The utility to the receiving agent of any given resource was greater than to that of the giving agent. It was shown that if each agent was paired enough times in each generation and the cost/benefit ratio was low enough then high levels of cooperation were found.

In Hales and Edmonds (2003) tags were applied to a simulated robot coordination scenario, originally given by Kalenka and Jennings (1999), producing high levels of cooperative help giving.

3.1 Mutation in the Models

We will now describe in a little detail how mutation was applied to the agents in each of the above models. We will not discuss the specific details of the reproduction process since we do not consider this relevant to the focus of this paper - in all cases it is safe to assume that variants of "roulette wheel" selection and "tournament selection" were used. These produce probabilistic selection into the next generation following the replicator dynamics assumptions stated earlier. Neither will we focus on the interactions or specific payoffs applied in each model, suffice to say all models capture some kind of collective coordination / cooperation problem in which cheating or free riding is possible under certain conditions.

In order to examine and compare mutation schemes we make a distinction between the mutation rate applied to the tag and that applied to the strategy. In all cases agents are represented in the models using sets of artificial "genes" (some set of data types) that are mutated when copied into the next generation.

The published descriptions of the models all explicitly state that the mutation rate applied to the tag and the strategy is the same (some probability). We label this rate *m*. However, models vary in the mutation *operation* applied with probability *m* and in the way they represent tags and strategies. Here it is claimed that this variation of mutation operation and tag / strategy representation can *hide* what is best understood as a variation in *mutation rate*.

<Tag type> [range]	**Strategy**	**Examples**	**Reference**
<R> [0..1]	<R R R>	<0.05><1,1,1>	Riolo 1997
<B B B B> [0 | 1]	< B >	<0 1 1 0><0>	Hales 2000
< I > [1..500]	< B B >	<324><0 1>	Hales & Edmonds 2003
< R > [0..1]	< R >	<0.6><0.5>	Riolo et al 2001

Table 1. Examples of the representations of tags and strategies in various tag models. For details see the text.

In Hales (2000) tags are represented as fixed length bit strings (experiments were conducted using various lengths of strings – in table 1 we show the four bit case only) and strategies as a single bit (either to cooperate or to defect). The mutation rate was $m = 0.001$ and the population size was $p = 100$. Since each agent is completely represented by a binary string the mutation operation is simply to flip each bit with probability *m* (both tag and strategy bits). It would superficially appear that strategy and tag are therefore mutated at the same rate and in the same way. However the results given in Hales (2000) show that high cooperation only occurred when the number of tag bits L was large ($L \geq 32$). In these cases the tag is more prone to mutation than the strategy. Any mutation in the tag creates a new distinct tag because pairing in the model is based on tag *identity* not *similarity*. The effective mutation rate on the tag as a whole is $1-(1-m)^L$ so for $L = 32$ bits the mutation rate on the tag is 32 times that on the strategy.

In Riolo et al (2001) each agent is composed of two *real* numbers (see table 1) - one representing its tag and one representing a so-called "tolerance". The tolerance is a kind of "proxy strategy". Essentially (simplifying) a smaller tolerance value means a less cooperative agent. Mutation is applied to both the tag and tolerance with probability $m = 0.1$. Again, superficially, it appears that both are being mutated at the same rate. However, the mutation *operation* applied to the tag and tolerance is not the same. When mutation is applied to the tag it is replaced with a new random value drawn informally from the range but when the tolerance is mutated it has Gausian noise (of mean 0 and standard deviation 0.01) added to it. So tags, when mutated, get *new* values but tolerances get *modified* by small values. We would expect the absolute average tag change amount[3] to be one third when mutation is applied. Since $m = 0.1$ we might characterize the average overall tag change amount to be ≈ 0.0333. In the case of tolerance the absolute average change would be almost two orders of magnitude lower (≈ 0.0008).

In both Riolo (1997) and Hales and Edmonds (2003) our analysis becomes slightly less straightforward. In both cases strategies are composed of multiple "genes" which do not simply relate to unconditional cooperation or selfish behavior. This is in part due to the scenarios. In Riolo (1997) agents play the IPD with agents having similar tags for a number of rounds. The level of cooperation produced is not high and constant but fluctuates with periods of high and low cooperation. Tags are represented by single real values [0..1], strategies by triples of real values $<i, p, q>$ (see table 1) capturing a probabilistic IPD strategy space (i is the probability of cooperation for the first round, p the probability of cooperation if in the previous round the other agent cooperated, q the probability of cooperation if the other agent defected on the previous round). So a space comprising tit-for-tat as well as pure defection and pure cooperation is formed (along with probabilistic variants). The mutation rate $m = 0.1$ is the same for each trait as is the operation (adding Gaussian noise with mean 0 and standard deviation 0.5). Here we have an interesting counter-point to the previous model (Hales 2000) where we stated (above) that because the tag was split in to several parts the effective mutation rate was higher than the strategy that was specified as a single "gene". Here, we have the reverse, so surely this suggests that the mutation rate applied to the tag is lower than that applied to the strategy? In one sense this is true. However, what is important is not the mutation of the representation as such, the stored value, but *how that value relates to behavior*. Since the strategy is a triple, in which pure cooperation is represented as all values being 1 and pure defection all values being 0, the relationship between mutation and the resultant change in strategy is not simple. However we can note that the probability of going from a triple of zeros to a triple of ones (from pure defection to pure cooperation) in a single (or even multiple) mutation event is approaching zero. However, since we are talking about IPD not single round PD the situation is more complex and we leave detailed treatment to a future work[4].

In Hales and Edmonds (2003) simulated robots work in teams to unload trucks in a warehouse. Here again we have a strategy composed of multiple parts. In the model tags are represented as single cardinal values [1..500] and strategies as pairs of binary values. Again the way the strategy effects behavior is complex and moderated by the scenario. A strategy represented by bit values "11" represents full cooperation whereas a value of "00" represents completely selfish behavior. Mutation is applied

---

[3] Here (and in the following examples) we make a few simplifying assumptions (namely that tags are treated as random variables).

[4] It is worth noting that the cooperation found in the paper (Riolo 1997) was not of the "strong" single interaction kind given in Hales (2000) and Riolo et al (2001). Indeed one of the findings of the paper was that tags did not produce cooperation in the single round game.

to the triple of traits with rate m = 0.1. The mutation operation is to replace the existing value with another value chosen uniformly randomly over the space. Again simplifying things a little we can say that the probability of a strategy changing from 11 to 00 (or vice versa) is the probability that two bits are replaced with their compliment $0.25(m^2) = 0.0025$. The probability of a completely new tag (tags are distinct integers, matching is based on identity) is 0.998(m) = 0.0998.

So, our analysis of these existing models shows that tags mutate more quickly than strategies under algorithms that presents a uniform *mutation rate*. Of importance (as stated before) is the *representation* of tags and strategies and *mutation operators* taken together with the mutation rate. Only by considering all these factors can an underlying average relative rate of change be determined between the two entities (tag and strategy). In each case when we do this we find that the tag changes much more quickly than the strategy. Now we advance a hypothesis based on this.

4. Hypothesis and Theory

From our analysis of the mutation schemes in the previous tag models we now advance a qualitative hypothesis concerning a necessary condition for tag models to produce high cooperation in one-time interactions: f*or tag based systems to support high levels of cooperation tags must mutate faster than strategies.* We can also state a qualitative "mini-theory" to explain this: *Cooperative tag groups need to spread (by mutation of tags) before free-riders (by mutation on strategies) invade the group.*

We don't have a quantitative complement to these two statements. It would appear that in order to determine the specific numbers in a specific scenario (model) we would need to consider the nature of the tag space, the nature of the strategy space and the way agents specifically interacted (the game). This is an aspect of on-going work.

5. Testing the Hypothesis

In order to test (at least partially) our hypothesis we implemented a new (minimal) tag model in which agents play single rounds of PD. We consider the result of high cooperation in the single round PD the most significant result so far advance for tags. Additionally the scenario is well understood and there are many existing models that allow for comparison. The singe-round PD captures, in a minimal way, many of the essential features of the problems of cooperation in collective interactions. In our tag model (described below) we varied the relative mutation rate between the tag and strategy to examine if this had an effect on the amount of cooperation produced. Firstly we describe the PD.

5.1 The Prisoner's Dilemma

The Prisoner's Dilemma (PD) game captures a scenario in which there is a contradiction between collective and self-interest. Two players interact by selecting one of two choices: Either to "cooperate" (C) or "defect" (D). For the four possible outcomes of the game players receive specified payoffs. Both players receive a reward payoff (R) and a punishment payoff (P) for mutual cooperation and mutual defection respectively. However, when individuals select different moves, differential payoffs of temptation (T) and sucker (S) are awarded to the defector and the cooperator respectively. Assuming that neither player can know in advance which move the other will make and wishes the maximize her own payoff, the dilemma is evident in the ranking of payoffs: $T > R > P > S$ and the constraint that $2R > T + S$. Although both players would prefer T, only one can attain it. No player wants S. No matter what the other player does, by selecting a D move a player ensures she gets either a better or equal payoff to her partner. In this sense a D move can't be bettered since playing D ensures that the defector cannot be suckered. This is the so-called "Nash" equilibrium for the single round game. It is also an evolutionary stable strategy for a population of randomly paired individuals playing the game where reproduction fitness is based on payoff. So the dilemma is that if both individuals selected a cooperative move they would both be better off but both evolutionary pressure and game theoretical "rationality" selected defection.

5.2 The TagWorld model

The TagWorld model presented here is similar to Hales (2000). What is new is that we explicitly vary the mutation rate applied to the tag while keeping the rate constant for the strategy.

Agents are represented by a single binary digit (the strategy bit) and a single real number in the range [0..1] (the tag). The strategy bit represents a pure strategy: either unconditional cooperation or unconditional defection. Initially the population have their strategy and tag values set to randomly with uniform probability over the space of all possible values. The following evolutionary algorithm is then applied.

In each generation each agent (a) is selected from the population in turn. A game partner is then selected. Partner selection entails the random selection of another agent (b) from the population such that (a) ≠ (b) but the tags of (a) and (b) are identical. If no agent exists with identical tags to (a) then (b) is selected at random from the entire population regardless of tag value. Consequently (a) will always find a partner even if its tag does not match any other agent in the population. During game interaction (a) and (b) invoke their strategies and receive the appropriate PD payoff (T, R, P or S)). After all agents have been selected in turn and played a game a new population is asexually reproduced. Reproductive success (fitness) is proportional to average payoff (i.e. the total score divided by the number of games played). The entire population of agents is replaced using a "roulette wheel" selection method (Davis 1991)[5]

5.2.1 Parameters used in the model

For our initial experiments (presented below) we used the following parameter value: the population size was N = 100 and the number of generations for each run of the model was 1000. The PD payoffs were T = 1.1, R = 1, P = S = 0.0001. These values were selected to give a very high incentive to cheat (T is high and P and S are low). P and S were selected as a small value but greater than zero (indicating a very small chance for agents, with Sucker or Punishment payoffs, of reproduction). If a small value is added to P (enforcing T > R > P > S) results are not significantly changed.

For the strategy bit the mutation rate was fixed constant at m = 0.001 (a low value). But for the tag a mutation factor f was applied to m changing the mutation rate. We varied f from [0..10] in increments of 2. Mutation of the strategy involved flipping the bit value. Mutation of the tag involved replacing the tag value with another uniformly randomly selected tag from the range [0..1]. To summarize, when an agent is selected for reproduction into the next generation, mutation is applied to the strategy bit (resulting in the bit being flipped with probability m) and to the tag (resulting in it being replaced with a new randomly selected tag with probability mf).

5.2.2 Results

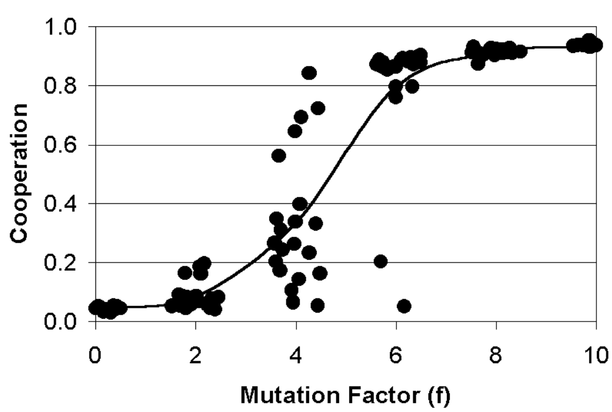

Figure 1. Results from simulations plotting mutation factor (f) against cooperation.

[5] Using this method the probability that an agent will be reproduced into the next generation is probabilistically proportional to average payoff.

The results are given in figure 1. Cooperation increases as the mutation factor is increased. For each value of the mutation factor (*f*) given on the x-axis are plotted 20 points from 20 individual runs (to 1000 generations). Cooperation given on the y-axis represents the proportion of all game interactions in a run that were mutually cooperative. Since we have 100 agents, with one game each per generation and 1000 generations per run, each point represents a proportion of mutual cooperation over 10^5 games. Each run had the same parameters but was initialized with different pseudo-random number seeds. The (smoothed) line joins the plotted average of the 20 points. The average is therefore over 2 x 10^6 individual games. To improve readability noise has been added to the x-coordinate of each point (+/-0.5).

There are a number of interesting characteristics presented in figure 1. Firstly, we do indeed see an increase (on average) of cooperation when we increase the relative mutation rate of the tag with respect to the strategy. Given this we have a little more confidence that our hypothesis may be correct since it allowed us to predict this property. The increase is non-linear, the average curve approximating a sigmoid shape with a threshold that would appear to be around f = 5. Notice that above f = 6 we see no results below 0.8 cooperation and below f = 4 we see no results above 0.2 cooperation[6]. Around the threshold we get high variance of results – indicating both high and low cooperation outcomes. So it would appear that at the threshold things become unpredictable and chaotic (i.e. the initial random variations of the runs send the model into different cooperation regimes) but that either side of the threshold the outcome is predictable.

5.2.3 Further Results

Our initial results are encouraging and appear to indicate that applying a high mutation rate to the tag relative to the strategy produces high levels of cooperation (at least in the PD game). However, does this explain the cooperation demonstrated in those other models described previously? Also, we have only tested our hypothesis with one kind of tag (a real number) and one kind of task (playing the PD). In order to begin to address these issues we re-implemented the Hales (2000) model[7]. In order to do this we extended the TagWorld model to incorporate tags represented as strings of binary digits with mutation being applied to *each bit in the string* with the same probability (m = 0.001) as the strategy bit.

Firstly we docked the model with the results previously reported. Figure 2 shows the original results from Hales (2000) and figure 3 shows the new results. Due to computation and time limitations we only executed each run to 10,000 generations. In the original results runs were to 100,000. Since variance can be quite high in the mid range of L (as shown in figure 2) we did not expect exact matching – what we did expect was the same overall pattern (high cooperation where L >= 32 bits over

[6] Points that appear to violate this are a result of the added noise.

[7] The code for the original model has long since gone to the big hard disk in the sky.

all of the T payoffs examined). Each bar in the chart is an average of 5 independent runs with the same parameters but different random number seeds. Figure 3 shows the absolute difference between figure 2 and 3 for comparison purposes.

Next we changed the model such that mutation was applied to the tag *as a whole* with the same probability as the strategy (m = 0.001). This was achieved by replacing the entire binary tag string with a new binary tag string randomly selected from the set of possible tag strings (i.e. replacing each bit with a randomly selected one). Those results are shown in Figure 5. As can be seen, cooperation completely disappears over all of the parameter range tested – this indicates that the previous results in Hales (2000) were related on the specific kind of mutation used rather than simply the structure of the tag (a binary string).

Next we increased the mutation factor (f) applied to the tag by powers of 2 (f = 1, 2, 4, 8..64) while keeping the strategy mutation rate at m = 0.001. This means that we apply the same mutation rate to the tag as when each bit in the string was been mutated with probability m but we simply replace the tag completely with a new binary string rather than mutating each bit separately. The results can be seen in figure 6 - cooperation reappears. This result indicate that it is not the specific kind of mutation (independently applied to each bit of the tag with probability m) applied that is necessary for high cooperation but rather the probability of mutation applied to the tag.

Finally we kept the last scheme but replaced the binary string with a single real value. Figure 7 shows the results from this. From our findings we would expect that replacing the binary string with a real value should make no difference to the level of cooperation so long as the same mutation values are used and indeed this is evidenced. We also got the same results when we used an integer for the tag (with values between 0..30,000).

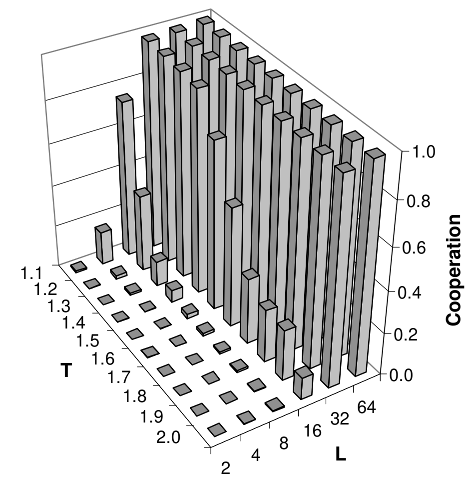

Figure 2. Results given in Hales (2000). T = PD temptation payoff and L is the tag length in bits. Runs were to 100,000 generations.

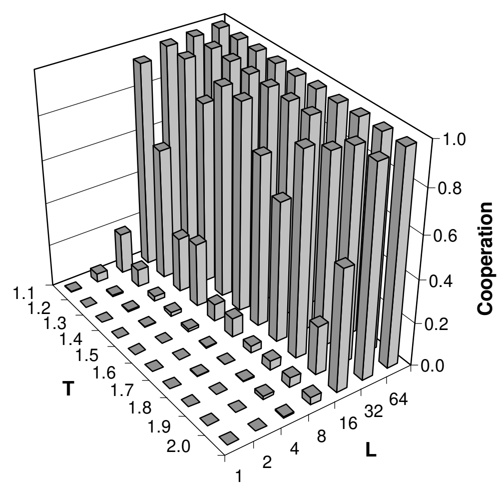

Figure 3. Reproduction of results using same parameters as in figure 2 (except that here runs were to only 10,000 generations)..The main differences are where L = 8 and 16.

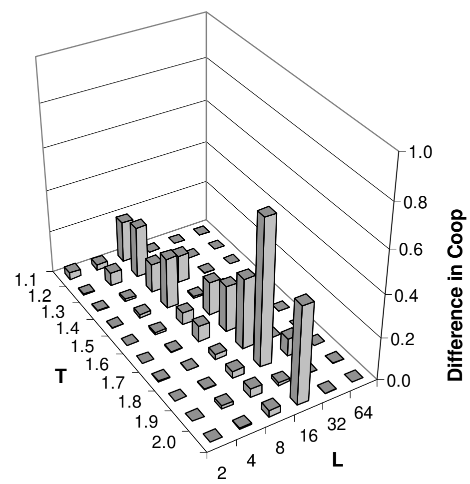

Figure 4. Shows the absolute differences between figure 2 and figure 3 (i.e. the docking errors). As stated previously, we consider this to be due to the difference in generations – but we need to test this.

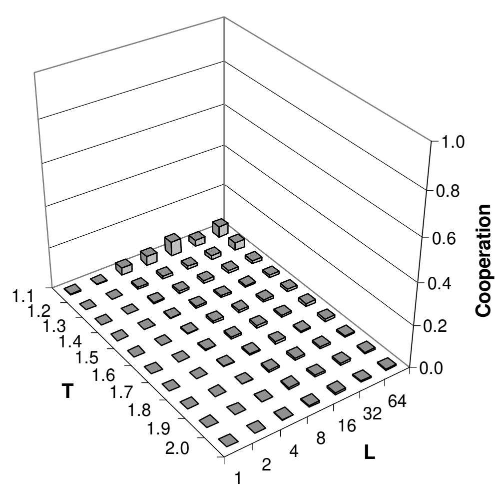

Figure 5. Results when the same mutation rate is applied equally to the tag and strategy. T is the PD temptation payoff and L is the tag length in bits.

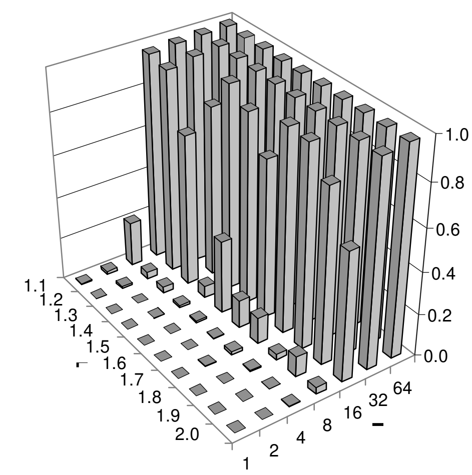

Figure 6. Results when mutation is increased by the tag length L such that the mutation factor f = L in all cases. Cooperation is restored.

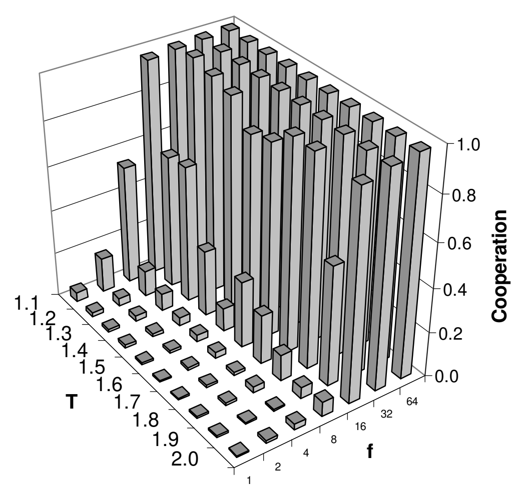

Figure 7. The tag is replaced with a single real number but the mutation factor applied to tags (f) is increased by the same values as previously the binary tag length was increased. The tag mutation factor (f) appears to be the necessary condition to produce cooperation in these scenarios.

6. Conclusions

From a detailed analysis of existing tag models we identified an implicit assumption – the mutation rate of the tags was higher than that applied to the strategies. We tested this hypothesis in a new tag model by varying the mutation rate of the tag while keeping the rate applied to strategies constant. We found that there was a non-linear relationship between amount of cooperation and the ratio of tag to strategy mutation rate. High cooperation was only produced when tag mutation was much higher than the strategy mutation rate. However, more work needs to be done in order to predict, for given scenarios, what the tag / strategy mutation ratio threshold value would be[8].

We then took this result and tested it over a larger parameter space by re-implementing a previous model, docking with that model and then obtaining the same results with changed tag structures by increasing mutation rates. I would appear that the hypothesis holds allowing for a degree of prediction and control.

Although our initial motivation for exploring tags was to understand aspects of human social phenomena our current motivation for this work is to understand how to program artificial systems such as truly decentralized (i.e. severless, where there is no central server but rather a collection of cooperating peer nodes) P2P self-organizing networks (Jelasity et al in press). In these systems each node or "peer" needs to offer bandwidth and processing capacity to other nodes without necessary getting a payback from those nodes. Additionally each node has partial and often changing views of node members meaning that storage of reputation information becomes unwieldy and non-scalable. In on-going work we are attempting to import ideas such as tags developed in the complex system and social simulation communities into the engineering realm.

We believe that the single-round PD captures *one kind* of P2P engineering problem. If we can get nodes to cooperate in the PD then we believe we can engineer them to share bandwidth and processing time, altruistically, in real systems. We have already demonstrated that the lessons learned here can be used in P2P file sharing scenarios (using simulation) Hales (2004a, 2004b). We have therefore practically shown how results from PD type simulations can be applied to engineering problems. However we still have *many* issues to address such as how systems can be engineered in which agents (nodes) can not "whitewash" a system (that is, simply defect all the time never adapting – i.e. not acting in a boundly rational way) or how to stop agents from presenting different tags to different agents. Our next step is to apply these techniques to more realistic P2P simulations.

[8] This will depend on a number of factors and a discussion is beyond the scope of, and space allowed for, this paper. See Hales (2000, 2001) for a little more on this.

References

Davis, L.(1991) Handbook of Genetic Algorithms. Van Nostrand Reinhold, NY, USA.

Edmonds, B. and Hales, D. (2003) Replication, Replication and Replication - Some Hard Lessons from Model Alignment. Journal of Artificial Societies and Social Simulation 6(4).

Hales, D. (2000), Cooperation without Space or Memory: Tags, Groups and the Prisoner's Dilemma. In Moss, S., Davidsson, P. (Eds.) Multi-Agent-Based Simulation. Lecture Notes in Artificial Intelligence, 1979:157-166. Berlin: Springer-Verlag.

Hales, D. (2001) Tag Based Cooperation in Artificial Societies. PhD Thesis (Dept. Of Computer Science, University of Essex, U.K. 2001).

Hales, D. (2002) Evolving Specialisation, Altruism and Group-Level Optimisation Using Tags. In Sichman, J. S., Bousquet, F. Davidsson, P. (Eds.) Multi-Agent-Based Simulation II. Lecture Notes in Artificial Intelligence 2581:26-35 Berlin: Springer Verlag.

Hales, D. and Edmonds, B. (2003) Evolving Social Rationality for MAS using "Tags", In Rosenschein, J. S., et al. (eds.) Proc of the 2nd International Conference on Autonomous Agents and Multiagent Systems, Melbourne, July 2003 (AAMAS03), ACM Press, 497-503

Hales, D. and Edmonds, B. (2004) Can Tags Build Working Systems? - From MABS to ESOA. Presented at the ESOA workshop at the AAMAS 2003 Conference (15th July 2003). Published as Lecture Notes in Artificial Intelligence 2977, Springer.

Hales, D. (2004a) Self-Organising, Open and Cooperative P2P Societies - From Tags to Networks. Presented at the 2nd Workshop on Engineering Self-Organsing Applications (ESOA 2004) located with the AAMAS 2004 conference, NY, July 2004. To be published by Springer.

Hales, D. (2004b) From Selfish Nodes to Cooperative Networks - Emergent Link-based Incentives in Peer-to-Peer Networks. To be presented at The Fourth IEEE International Conference on Peer-to-Peer Computing (P2P2004), 25-27 August 2004, Zurich, Switzerland. To be published by IEEE press.

Hamilton, W. D. (1964) The genetical evolution of social behaviours, J. Theor.Biol. 7, 1-52.

Holland, J. (1993) The Effect of Lables (Tags) on Social Interactions. Santa Fe Institute Working Paper 93-10-064. Santa Fe, NM.

Jelasity, M., Montresor,A., and Babaoglu, O. (2004) A modular paradigm for building self-organizing peer-to-peer applications. Proceedings of the 1st International Workshop on Engineering Self-Organising Applications (ESOA 2003), Springer.

Kalenka, S., and Jennings, N.R. (1999) Socially Responsible Decision Making by Autonomous Agents. Cognition, Agency and Rationality (eds. Korta, K. et al) Kluwer 135-149.

Nowak, M. & May, R. (1992) Evolutionary Games and Spatial Chaos. Nature, 359, 532-554.

Nowak, M. & Sigmund, K..(1998) Evolution of indirect reciprocity by image scoring. Nature, 393, 573-557.

Riolo, R. (1997) The Effects of Tag-Mediated Selection of Partners in Evolving Populations Playing the Iterated Prisoner's Dilemma. SFI Working Paper 97-02-016. Santa Fe, NM.

Riolo, R. L., Cohen, M. D. & Axelrod, R. (2001) Evolution of cooperation without reciprocity. Nature 414, 441-443

Roberts, G. & Sherratt, T. N. (2002) Nature 418, 449-500

Sigmund, K. and Nowak, A, M. (2001) Tides of Tolerance. Nature 414, 403-405.

Trivers, R. (1971) The evolution of reciprocal altruism. Q. Rev. Biol. 46, 35-57.