Change Your Tags Fast! – A necessary condition for cooperation?[1]

David Hales

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

Abstract. Since Holland (1993) introduced the concept of tags as a possible cooperation forming mechanism in evolving system (among other things) a number of tag models with intriguing, and potentially very useful, properties have been advanced. However there is currently little detailed understanding of the underlying processes 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. We identify what appears to be a necessary condition that previous tag models *implicitly* contained. We formulate our hypothesis by detailed analysis of the previous models and then explore and test it (in simulation) with a new model that demonstrates for the first time the importance of high mutation applied to tags. In order to import MABS techniques into engineering of MAS these kinds of analysis will be required. We need to understand the *necessary conditions* of desirable emergent properties not just *existence proofs* of them.

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 agents 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.

In the computational models discussed here tags are modeled using some number (either a binary bit string, a real number or an integer). When agents interact they preferentially interact with agents possessing the same (or similar) tag value. One way to visualize this is to consider a population of agents partitioned between different colors.

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

Each agent carries a single color. In a system with only 3 different possible tag values we could think of this as each agent carrying a flag of red, green or blue. Agents then preferentially interact with agents carrying the same color (forming "interaction groups"). When agents evolve (using some form of evolutionary algorithm) they may mutate their tag (color). This equates to moving between interaction groups.

In the models presented here, tags take on many possible unique values (by say using a real number, there are many possible unique tags rather than just 3 colors) however, the basic process is the same – agents with the same tags preferentially interact and tags evolve like any other genotypic trait.

Another way to think of tags is that some portion of the genotype of an agent is visible directly in the phenotype but the other agents.

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

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 appearing to occur at the group-level, there has been little analytical or empirical exploration of this hypothesis.

2. Some Previous Tag Models

There have been a number of tag models implemented previously. All generally show 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, no "genetic-style" cross-over operations but low probability mutations on tags and strategies during reproduction).

---

[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.

Riolo (1997) gave results of expansive and detailed 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. Again interaction was dyadic. 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 one shot game.

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 as 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.

2.1 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 (variants of "roulette wheel" selection and "tournament selection" were used, and these produced 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.

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 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. It is this variation of mutation operation and tag / strategy representation that can hide what is best understood as a variation in mutation rate.

In Hales (2000) tags are represented as fixed length bit strings and strategies as a single bit (either to cooperation in the single-round PD or to defect). The mutation rate is m = 0.001 and the population size is 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 of the paper show that high cooperation only occurred when the number of tag bits L was large (L = 32 or more). In these cases the tag is more prone to mutation than the strategy because it contains more bits. Any change in the tag effectively creates a new distinct tag because pairing in the model is based on tag identity not similarity. So the effective mutation rate on the tag as a whole is $1-(1-m)^L \approx 0.0315$ (more than 30 times that on the strategy).

In Riolo et al (2001) each agent is composed of two real numbers - 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 bother the tag and tolerance with probability m = 0.1. Again it appears that both are being mutated with the same rate. However, the mutation operation applied to the tag is to replace it with a random value drawn informally from the range but the tolerance has Gausian noise (of mean 0 and standard deviation 0.01) added to it. So tags, when mutated, get new values chosen randomly from the range but tolerances get modified by small values. Simplifying the analysis somewhat, we could expect the absolute average tag change amount to be ≈ 0.333 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 we can see that 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 relate to simple strategies of 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 into periods of high and low cooperation. Tags are represented by single real values [0..1], strategies by triples of real values <i, p, q> each a probability 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 several parts the effective mutation rate was higher than the strategy. 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 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 mutation event is approaching zero. However, since we are talking about IPD not just a single round interaction the situation is more complex and we leave detailed treatment to a future paper[3].

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. However, to simplify, a strategy represented by bit values "11" represen ts full cooperation whereas a value of '00" represents completely selfish behavior. Mutation is applied to the triple of traits with rate m = 0.1. The mutation operation is to replace the existing value with another value chosen uniformally 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 (again tags are distinct, matching on identity) is $0.998(m) = 0.0998$.

So in all these cases it appears tags change more quickly than strategies under an algorithm 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 estimated 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. Next we advance a hypothesis based on this.

3. 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: In general for tag based systems to support high levels of cooperation tags must mutate faster than strategies. We can also state a

[3] The cooperation found here (Riolo 1997) was not for the single interaction kind given in Hales (2000) and Riolo et al (2001). Indeed one of the findings of the paper was that the given model did not produce cooperation in the single round game.

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[4].

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.

3.1 Testing the Hypothesis

In order to test 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 to be 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 new model we varied the relative mutation rate between the tag and strategy to examine if this had an effect on the amount of cooperation produced. The model and results are described below but firstly we briefly outline the single-round PD.

3.2 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

---

[4] Unfortunately space here does not allow an illustration of the tag-group process. We refer interested readers to Hales (2000).

selected a cooperative move they would both be better off but both evolutionary pressure and game theoretical "rationality" selected defection.

3.3 The TagWorld model

Our TagWorld model is a variation on Hales (2000). We use a single real number to represent tags (as in Riolo et al 2001) rather than a binary string. 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 (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 payoff. After all agents have been selected in turn and played a game a new population is asexually reproduced. Reproductive success is proportional to average payoff. The entire population of agents is replaced using a "roulette wheel" select ion method (Davis 1991)[5].

3.3.1 Parameters used in the model

For the results presented here we used similar parameters to Hales (2000), though here we did not execute a scan over the parameter space. The population size was N = 100 and the number of generations for each run of the model was 1000. The PD payoffs were T = 1.9, 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

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

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 real 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*).

3.3.3 Results

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.

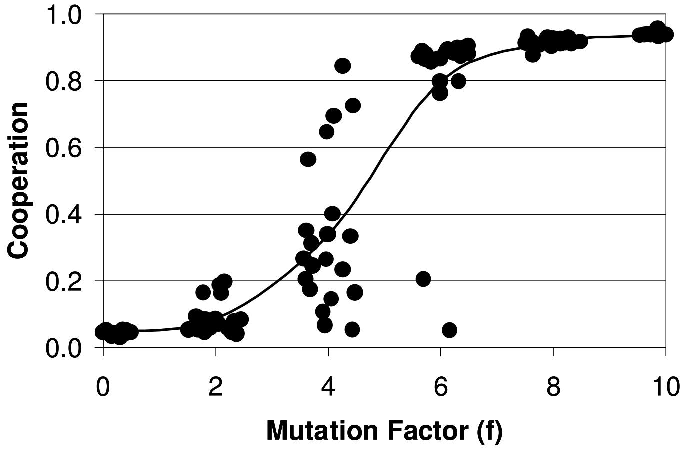

Fig 1. Results from each simulation run plotting mutation factor (*f*) against cooperation.

The increase is non-linear, the average curve approximating a sigmoid shape with three zones: A first zone with convergence to low value, a zone where it is unpredictable and a zone with convergence to high value. Where $f < 4$ we find convergence to low cooperation (no results above 0.2 cooperation[6]). For $f > 6$, cooperation converges to a high value (no results below 0.8 – note points that appear to violate this statement are a result of the added noise as mentioned above. In the "unpredictable area" $4 \leq f \leq 6$ we get high variance of results – indicating both high and low cooperation outcomes. Here, it would seem, results become unpredictable and chaotic (i.e. influenced by random variations due to the different pseudo-random number seed used in each run).

4. 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[7].

We believe that the single-round PD potentially captures many kinds of engineering problem. *One kind* we are currently exploring is a P2P engineering problem. If we can get nodes to cooperate in the PD then we believe we can use a similar technique to get them to share bandwidth and processing time, altruistically, in real systems. But we still have *many* issues to address. On-going work with network-like P2P simulation scenarios (Hales 2004b) has shown that high mutation rates on tags (or network links in these cases) is important in maintaining high cooperation and scalability. However our medium term aim is to produce a deployable service following a modular approach (Jelasity et al 2004).

References

Davis, L. (1991) Handbook of Genetic Algorithms. Van Nostrand Reinhold, New York.

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

[7] 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.

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.) Proceedings 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. In Di Marzo Serugendo, G.; Karageorgos, A.; Rana, O.F.; Zambonelli (eds.) Engineering Self-Organising Systems - Nature-Inspired Approaches to Software Engineering. Lecture Notes in Artificial Intelligence 2977, Springer, Berlin.

Hales, D. (2004b) Self-Organizing, Open and Cooperative P2P Societies – From Tags to Networks. To be presented at the 2nd Workshop on Engineering Self-Organizing Applications at AAMAS 2004, July 2004, New York.

Hamilton, W. D. (1964) The genetical evolution of social behaviours, I and II. 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. Procceedings of the 1st International Workshop on Engineering Self-Organising Appplications (ESOA 2003), Springer.

Kalenka, S., and Jennings, N.R. (1999) Socially Responsible Decision Making by Autonomous Agents. Cognition, Agency and Rationality (eds. Korta, K., Sosa, E., Arrazola, X.) 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. Santa Fe Institute 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.