Altruism "For Free" using Tags

David Hales
David Hales, Dept. of Computer Science, The University of Bologna
dave@davidhales.com

Keywords: Altruism, Evolution of cooperation, Tags, Prisoner's Dilemma.

A number of recent models have shown how "tags" (Holland 1993) - arbitrary observable phenotypic markers - can produce altruistic behaviour between initially selfish individuals in an evolving system (Riolo et al 2001). These models use tag and tolerance values associated with each individual such that *those with identical tags are forced to be altruistic* (Roberts et al 2002, Edmonds et al 2003). Here we outline a simpler approach in which individuals play the single round Prisoner's Dilemma (PD) game. Contrary to the previous tag models altruism is demonstrated in the single round game *without forced altruism between those with identical tags* or knowledge of previous interactions. The system is reverse-scalable (the more individuals in the population the quicker altruism emerges) and robust to noise. This *"altruism for free"* property has already been adapted and applied into robotic scenarios (Hales et al 2003), and peer-to-peer networks (Hales 2004) producing light-weight robust algorithms that allow selfish agents to work as cooperative teams in low trust environments without the need for complex trust or market mechanisms.

We demonstrate that tags produce cooperation in the singe-round PD in an agent-based simulation. In the PD agents are paired and play a game by selecting one of two strategies (either cooperate or defect). Depending on what they select they each receive a payoff value. If both cooperate they both get R, if both defect they get P, otherwise the cooperator gets S and the defector gets T. The dilemma arises from the payoff constraints: T>R>P>S and 2R > T+S. Agents are represented as fixed length bit strings (of length L+1) comprising a tag of length L bits and a single strategy bit. The strategy bit represents a pure strategy, either unconditional cooperation or unconditional defection. Initially the population of agents is set to random bit strings (with each bit decided by a fair coin toss). The following evolutionary algorithm is then applied:

```
LOOP some number of generations
  LOOP for each agent (a) in the population
    Select a game partner agent (b) with matching tag
    Agent (a) and (b) play single round PD
  END LOOP
  Reproduce agents in proportion to their average payoff
END LOOP.
```

In each generation, each agent (a) is selected from the population (of size N) 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) and tags match identically. If no agent in the population has matching tag to then (a) interacts with a randomly chosen agent. 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 N agents is replaced using a "roulette wheel" selection method. Mutation is applied to each bit of each reproduced player with probability M = 0.001, consequently *tags and strategies* are mutated in reproduced agents. The PD payoffs are parameterised over T such that T > 1. The reward R for cooperation is 1. The punishment P for mutual defection and the sucker payoff S for cooperation with a defector are both some small value.

A set of runs to 100,000 generations with a population of size N = 100, was executed for various values of T and L. We found high levels of cooperation emerged quickly when L was sufficiently large (L>16). The results obtained indicate that very high levels of cooperation can be sustained between selfish greedy optimisering agents in the single round PD via simple tag biasing. There is no requirement for knowledge of past performance or recognition of individual agents

The tag space can be visualised as an L-dimensional hyper-cube with corners representing unique tag values. Agents sharing a tag, share a corner. Mutation produces movement between corners. Game interaction is therefore taking place in an abstract "tag space". Cooperative groups sharing matching tags will form in corners of the hyper-cube. These groups will outperform non-cooperative groups and hence tend to increase in size over generations. However, if mutation introduces defecting agents into a cooperative group they will tend to outperform the cooperators within the group (by suckering them). From this the seeds of the destruction of the group are planted, since as the number of defectors increases within a group the overall fitness of agents within the group decreases. Other more cooperative groups (if they exist) will tend to expand. While this process is occurring, mutation of tag bits will produce a slow migration of agents between corners of the hyper-cube, possibly founding new groups in previously empty corners.

Figure 1 is a visualisation of the process over time taken from a single run. Each line on the vertical axis represents a corner of the hyper-cube (i.e. unique tag value). The horizontal axis represents time in generations. If no agents have a particular tag value in a given generation then the line is left blank (white). Alternatively, if a corner contains all cooperative agents then the line is light grey. For a mixed group in which there are both cooperators and defectors the line is dark grey. For an entirely defective group the line is black. Examination of figure 2 shows the time evolution of groups in tag space. Initially cooperative groups (light grey lines) become invaded by defectors producing mixed groups (darker grey) that very swiftly become entirely defective (black) and then quickly go extinct (white). See Hales (2000) for more results from this model.

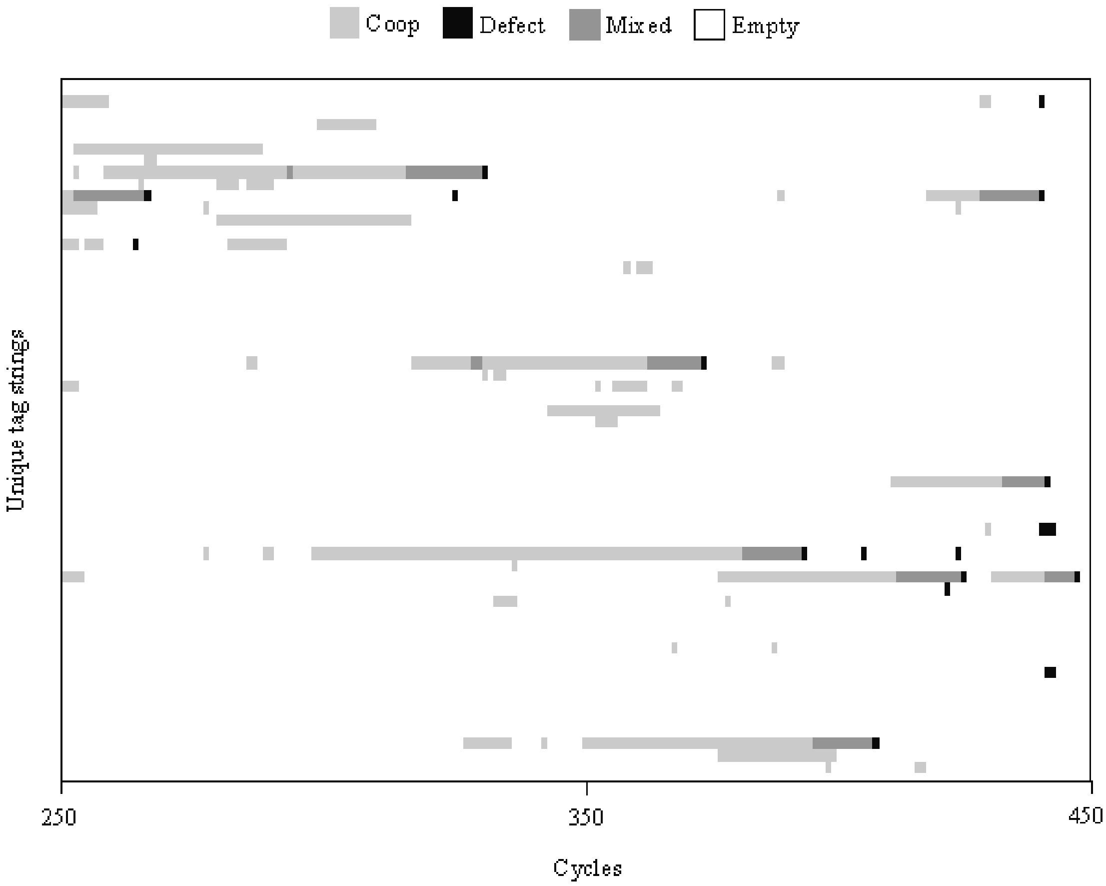

Fig. 1. Visualisation of 200 generations from a simulation run showing cooperative groups coming into and going out of existence. Each line on the vertical axis represents a unique tag value (of which only a subset is shown). If all agents sharing a tag value are cooperative then the line is light grey. If all agents are defectors then the line is coloured black. A mixed group is shown as dark grey.

References

D. Hales, “From selfish nodes to cooperative networks -- emergent link based incentives in peer-to-peer networks”, Proc. of the 4th IEEE International Conference on Peer-to-Peer Computing (P2P2004). IEEE Computer Soc. Press, (2004).

D. Hales, 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. (2000).

D. Hales. & B. Edmonds. Evolving Social Rationality for MAS using "Tags", In Rosenschein et al. (eds.) Proceedings of the 2nd International Conference on Autonomous Agents and Multi-agent Systems, (AAMAS03), ACM Press, 497-503. (2003).

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

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

R. Riolo, M. D. Cohen, R. Axelrod, “Cooperation without Reciprocity”, Nature 414, 441-443, (2001).

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