Agent Cooperation using Simple Fixed Bias Tags And Multiple Tags

Enda Howley and Colm O'Riordan

Department Of Information Technology
National University Of Ireland, Galway.
`enda.howley@nuigalway.ie`
`colmor@it.nuigalway.ie`

Abstract. Much research in multi-agent systems has focussed on the emergence of cooperation in societies where individually optimal behaviour for agents leads to low levels of cooperation in the society. This conflict between individual and collective rationality can be modelled through the use of social dilemmas such as the prisoner's dilemma. Tagging schemes have been shown to increase levels of cooperation through biasing interactions in a manner comparable to that of kin selection and trust mechanisms. We outline some simulations involving a simple tagging system and outline the main factors which are vital to increasing cooperation. This paper also outlines the effects of multiple tags.

1 Introduction

Agent interactions are often heavily biased by certain group structures. These structures may be based upon certain models of trust or kin selection. These groups reflect an approximation of which peers are most likely to remain altruistic. Some schemes are based solely on past behaviour and therefore reflect aspects of trust and relationships. Other models use physical proximity as a guide to bias interactions.

Significant research has been conducted involving agents whose interactions are determined by spatial proximity. Nowak and May(1993) describe the significance of group structuring techniques with special attention to spatially determined interactions[6]. This structuring can also be represented through the use of tagging which is an abstract method of biasing agent interactions based on membership of certain tag groups. Holland(1993) describes tags as markings or social cues that are attached to individuals (agents) and are observable by others[5]. When determining which agents should interact, we calculate the proximity of their tag values, as opposed to their physical proximity. Tagging is an abstraction which allows us reflect the relatedness between agents based on any possible grouping analogy, not simply spatial or genetic proximity. Tagging can be used as a general case to represent all these possible grouping structures without the specific complexities which they entail.

Hales(2004)[4] states that tags can evolve from initially random values into complex ever-changing patterns that serve to structure interactions between individuals. Tagging schemes are highly accurate mechanisms for biasing agent interactions based on their relatedness with each other.

This paper will analyse a simple tagging scheme and review some of the factors which contribute to its success. This paper will describe a series of experiments which we have designed involving a simple tagging scheme. For example, we will review the levels of cooperation achieved using different amounts of tags. We also investigate multiple tags which allow us to model agent societies, where agents can exist in a number of distinct groups.

Our primary motivation throughout this paper involves studying the effects of biasing agent interactions using various parameter space values. We extend previous research on tagging schemes to allow agents participate in multiple groups. This will be outlined through a series of simulations involving multiple tags. Throughout this paper all simulations involve populations of agents competing through the iterated Prisoner's Dilemma.

2 Related Research

The Prisoner's Dilemma (PD) is a simple two-player game where each player must make a decision to either cooperate (C) or defect (D). Both players decide simultaneously and therefore have no prior knowledge of what the other has decided. If both players cooperate they receive a specific payoff. If both defect they receive a lower payoff. If one cooperates and the other defects then the defector receives the maximum payoff and the cooperator receives the minimum. The payoff matrix outlined in Table 1 demonstrates the potential payoffs for each player.

Table 1. Payoff Matrix

Players Choice	Cooperate	Defect
Cooperate	$(\lambda1, \lambda1)$	$(\lambda2, \lambda3)$
Defect	$(\lambda3, \lambda2)$	$(\lambda4, \lambda4)$

The dilemma is a non-zero-sum, non-cooperative and simultaneous game. For the dilemma to hold in all cases, certain constraints must be adhered to: $\lambda2 < \lambda4 < \lambda1 < \lambda3$. These conditions result in $\lambda2$ being the sucker's payoff, $\lambda1$ is the reward for mutual cooperation, $\lambda4$ is the punishment for mutual defection, and $\lambda3$ provides the incentive or temptation to defect. The dilemma also states $2\lambda1 > \lambda2 + \lambda3$. This constraint prevents players taking alternating turns receiving the sucker's payoff ($\lambda2$) and the temptation to defect ($\lambda3$), therefore maximising their score. The following values were used throughout this research: $\lambda1 = 3, \lambda2 = 0, \lambda3 = 5, \lambda4 = 1$.

In the non-iterated game, the rational choice is to defect, while in the finitely repeated game, it is rational to defect on the last move and by induction to defect all the time. However, if there exists a non-zero probability the two players will play again, then cooperation may emerge. Within a society of social groups, repeated meetings are common and as with tag group members who meet repeatedly, significant levels of cooperation can emerge.

2.1 Previous Tagging Models

Holland[5] initially outlined the concept of tags and since then significant numbers of tag models have emerged[8][4]. Holland describes tags as markings or social cues that are attached to individuals (agents) and are observable by others. Riolo has described a number of tagging approaches throughout a series of papers[8][9], each focusing on the effects of tagging on levels of cooperation among players in the Iterated Prisoner's Dilemma. These papers outline basic forms of tagging: fixed-bias tagging, variable-bias tagging and evolved-bias tagging.

These models comprise two-player Iterated Prisoner Dilemma (IPD) tournaments. Players are not paired randomly as in traditional tournaments such as those researched by Axelrod[2]. Instead, each player is given a numeric tag value and this is used to probabilistically bias interactions towards players of similar tag values. In these tagging models, significant increases in cooperation can be observed. Simple replicator dynamics are used to reflect the fitness of agents over successive generations. The proportional fitness of a genome is used to determine representation of that genome in subsequent generations.

We hope to gain a greater understanding of how tagging mechanisms successfully increase cooperation among agents playing the IPD. We address the following questions:

- What factors allow tagging schemes boost cooperation among agents?
- How does multiple tagging effect levels of cooperation?

In later sections, we outline a number of experiments and the results obtained. In the following section we will discuss the design of our simulator and tagging model.

3 Experimental Setup

3.1 Strategies

The design of any IPD simulator requires the simple creation of an initial set of player strategies. The research by Nowak et. al.[7] provides a basis for our method of strategy definition. The following is a strategy genome with 3 genes representing 3 possible behaviour values when P_i *(Probability of cooperation in the initial move of a game)*, P_c *(Probability of cooperation after opponent has cooperated)*, P_d *(Probability of cooperation after opponent has defected).*

$$Genome = P_i, P_c, P_d, \tag{1}$$

Our initial population of agent genomes is randomly generated with random gene values and therefore the initial population has an initial average fitness of about 2.25[1][8]. A fourth gene representing a tag value is also given to each genome, this tag value is randomly assigned to all strategies in the population and is a simple integer in the range [1...Maximum number of tags]. Multiple tagging involves extending this model to include a number of tag genes representing each of the tag groups an agent is a member.

3.2 Tag Model

Our simulator consists of a population of 100 players; each is randomly generated from a normally distributed set of possible strategies. These strategies are each given random tag values. Games lasting 20 iterations between all members of the same tag value are conducted. Agents of different tag values do not play each other. Agents can play themselves. If an agent is the solitary holder of a tag value it may still play itself, even though there are no other members of the population it can interact with.

The representation of a genome in successive generations is based on its fitness in the current generation. Therefore this dictates representation in the next generation and so on, through successive generations. This simple replicator dynamic is similar to that used by Riolo[8][9] in his simulations. In our research and the research of Riolo, all reproduction is asexual. As a result, no crossover between agents occurs.

This simulator extends previous tag models to allow some parameters vary over certain experiments. For example, this simulator can conduct experiments across a number of possible tag parameters.

4 Results

In this section we outline the results of a number of experiments. Our primary goal is to address the questions which we outlined earlier in Section 2.1 of this paper.

4.1 Performance of Tagging Scheme

The first set of experimental results represent a fitness comparison between two populations. One implements a simple tagging scheme while the other uses no tagging technique. Across 50 simulations, 100 agents are allowed to compete for survival while playing each other. The tagged model randomly distributes up

---

[1] If we use the following IPD payoff matrix values $\lambda 1 = 3, \lambda 2 = 0, \lambda 3 = 5, \lambda 4 = 1$, a initial population of random agents will have an average of these fitness values $(3 + 0 + 5 + 1)/4 = 2.25$

to 50 different tag values to the 100 agents. Throughout each simulation, these players will play IPDs of 20 iterations with peers of the same tag value. 50 generations are simulated and these tag groups are represented through successive generations based on their fitness through the use of a replicator dynamic. Our non-tagged model operates in exactly the same manner but all players play IPD's of 20 iterations with every other member of the population.

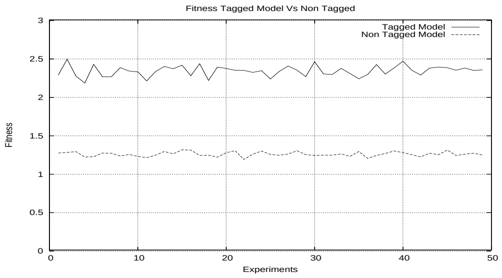

Fig. 1. Tagged Model Vs Non-Tagged Model

In Figure 1, we can clearly see the improved fitness achieved through the use of tagging. This data is comparable with previous research in the domain of tagging[8]. These results are consistent across all experiments while our simulation without tagging displays much lower levels of average fitness. The improved performance of our tagged model is as a result of tag groups insulating themselves against invasion from defecting strategies. In a non-tagged model, the effects of the defecting strategies propagates throughout the population and they dominate very quickly. This is less likely in a tagged model which partitions the population into tag groups through limiting their interactions. As a result groups containing defectors are impeded and loose representation to fitter competitive groups.

4.2 The Significance of Tag Group Size

To fully understand the factors determining the success of a tag model, we investigate the effect of varying the number of tags. The following experiment illustrates how the number of tags can influence levels of cooperation among agent populations. Four simulations represent populations using different numbers of permissible tags in their initial composition (1, 5, 20, 50) are shown in Figure 2.

In the experiment represented in Figure 2 we observe that as the number of tags increases, the partitioning effect increasingly limits the effects of defection. The model performs best when the number of tags is high. Therefore, in any one tag group, there exists a higher probability that all the members will be cooperative and as a result the group will be fitter.

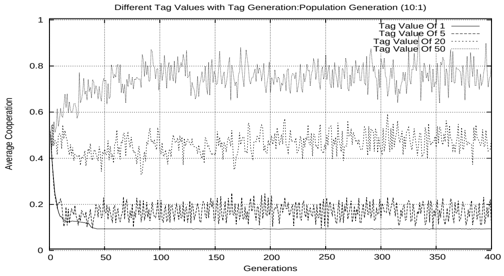

Fig. 2. Different Tag Ranges

4.3 The Evolution Of Tags

We have observed the importance of tag group size over time. The following experiment examines how the group dynamics are effected by the evolutionary process. In this simulation we record the percentage of possible tags in use at each generation. We also plot the average cooperation at each generation.

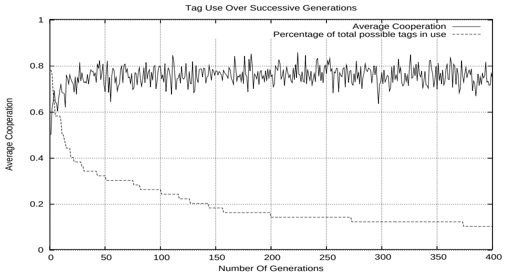

Fig. 3. Numbers Of Tags In Use

From this experiment we see that the number of total tag values falls significantly after the initial generations. The levels of cooperation remain very high. This is explained by the ability of tags to marginalise non-cooperative behaviour over the initial generations. The tagging system results in non-cooperative agents becoming extinct in the earlier generations while the most cooperative tag groups increase rapidly in size. These tag groups with the most altruistic members takeover the total population and lead to the extinction of tag groups which contain any non-cooperative agents. The resulting small number of tag groups all contain highly cooperative agents and are about equal in size. Usually this experiment resulted in about 5 of a possible 50 tag groups surviving, each with an average membership of 20.

4.4 The Effects of Multi-Tagging

In the following series of experiments, we have extended the simple tagging model through allowing each agent participate in multiple tag groups. This reflects real life interactions where individuals often participate in many social groups. Each agent may hold membership of up to five tag groups simultaneously. This is modelled through displaying each of these tag values. An agent cannot interact with peers which do not contain at least one common tag value. Therefore, they must each hold membership of at least one common tag group.

In the following experiment we outline the effects of varying the number of tags a population of agents may simultaneously hold. The actual tag values are specified at random in the range of 1 to 50 among a population of 100 agents. We ran 5 simulations allowing agent membership of different numbers of groups ranging from 1 to 5.

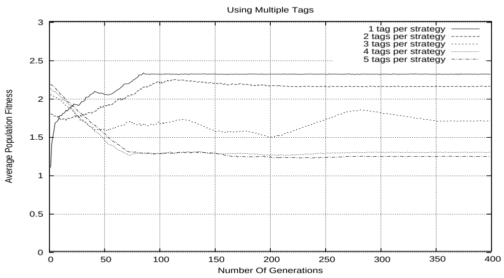

Fig. 4. Number Of Tags Per Strategy

From this experiment we observe the decreased levels of fitness among populations which permit agent membership of multiple tag groups. This is explained through the increased interactions which multiple group membership facilitates. As explained earlier altruism benefits from partitioning small groups of cooperative strategies away from defective peers. Multiple tags undermine this system and result in lower levels of fitness throughout the population.

A less contrived experiment specifies a population with a mixed degree of membership among many tag groups. For example each agent will have a high probability of holding one random tag membership, a smaller probability of holding two and so on. We model this as follows. The probability Y is calculated with respect to the number of tag groups X and some negative power N.

$$Y = x^{-N} \tag{2}$$

The following experiment shows the effects of varying the possible number of tags in use throughout the population while applying the above function with a N value of 2.2. It has been shown in many studies of various social structures and social networks, that connections and interrelationships between individuals follow this type of distribution[3]. Albert et. al.[1] show for a range of examples including web topology, citation patterns and food webs, the degree of connection can be described by a power law. This value significantly biases agents towards containing less tags as can be seen from the following probabilities.

$$P_1, P_2, P_3, P_4, P_5 = 1.0, 0.21, 0.08, 0.04, 0.02 \tag{3}$$

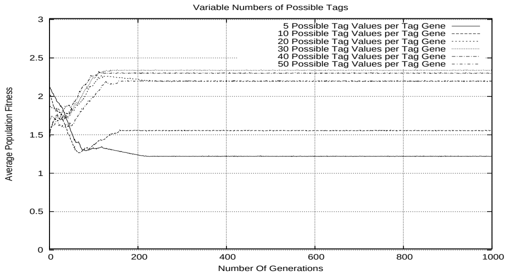

Fig. 5. Varying the number of possible tags in the population

The function specified limits the number of multiple tag groups agents may participate within. Subsequent fitness depends on the number of tag values permitted within the population. But these levels of fitness are dependant upon reducing the number of tag groups agents participate in. While this is reflected through results outlined earlier in this paper it is augmented through examining our final test case experiment which follows. We can observe the opposite effect of greater use of multiple tags through changing the values used in our formula. In the following case we can observe the direct effects of increased multiple tag use on the population.

In the following experiment we replicate its predecessor through simulating a new value of $N = 0.1$. The resulting function represents a greater probability agents will hold membership of multiple tag groups. Here we examine the negligible effect of varying the number of tags within a population which is biased towards holding membership of multiple tag groups. We observe this bias through the following probabilities:

$$P_1, P_2, P_3, P_4, P_5 = 1.0, 0.93, 0.89, 0.87, 0.85 \tag{4}$$

This experiment confirms that multiple tags undermine cooperation and as a result diminish the fitness of the agent population. Multiple tags counteract the benefits of using many tag values to boost cooperation in a population. The conflicting nature of the two parameters is confirmed through reviewing the dominant strategies which proceed to win our final simulations. In simulations depicting the emergence of cooperation, dominant strategies were predominantly single tag holders. Alternatively in non-cooperative populations multiple tag holders prevailed and displayed dominance.

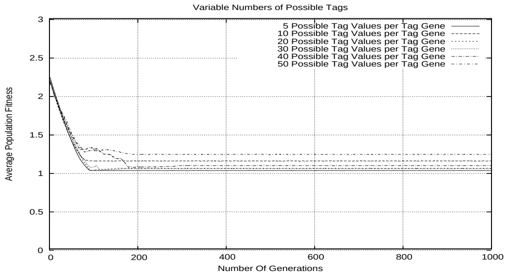

Fig. 6. Varying the number of possible tags in the population

5 Conclusions

Our experiments have highlighted the primary factors leading to the success of certain tagging models. We have observed the importance of the number of tags used in a population. This has a direct effect on the overall levels of cooperation. The success of tagging is based upon preventing invasion from greedy agents. This partitioning effect which is synonymous with tagging and spatial models is fundamentally important to their success. A tag group of only one agent never encounters the “invasion” difficulties which usually jeopardise cooperation among larger groups of agents. In successful tagging models we have observed a small set of tag groups emerging to dominate a population after a number of generations. This small number of tag groups composed of highly cooperative strategies experience rapid growth because of their fitness. Tag groups containing non-cooperative agents experience high attrition rates.

The effects of multiple tagging resulted in undermining the overall fitness of our population. This feature reflected that agents which spanned multiple tag groups were at a disadvantage and more susceptible to exploitation. Reinforcing our previous evidence that less interactions among agents improved cooperation, multiple tags increased such interactions and as a result decreased levels of cooperation.

In answer to the two questions posed earlier (Section 2.1) of this paper we conclude that the primary factors contributing to the success of tagging schemes all involve limiting the number of agent interactions to a minimum. This finding is based on results across all our experiments and is further reinforced by previous research in the domain.

Future work involves more elaborate tagging models. Various aspects of learning, evolution and communication of tags are also possible through extensions of current tag models.

Acknowledgment

The primary author would like to acknowledge the Irish Research Council for Science, Engineering and Technology(IRCSET) for their assistance through the Embark initiative.

References

1. R. Albert, A. Barabsi, and H. Jeong. Scale-free characteristics of random networks: The topology of the world wide web, 2000.
2. R. Axelrod. *The Evolution of Cooperation.* Basic Books, New York, 1984.
3. G. Csyani and B. Szendroi. Stability by mutation in evolutionary games. *Physical Review*, E 69, 2004.
4. D. Hales. Understanding tag systems by comparing tag models. In *Presented at the 2nd Model-to-Model Workshop M2M2 co-located with ESSA04*, 2004.
5. J. Holland. The effects of labels (tags) on social interactions. *Working Paper Santa Fe Institute 93-10-064*, 1993.
6. M. Nowak and R. May. The spatial dilemmas of evolution. *Int Journal of Bifurcation and Chaos*, 3:35–78, 1993.
7. M. Nowak and K. Sigmund. The evolution of stochastic strategies in the prisoner's dilemma. *Acta Applicandae Mathematicae*, 20:247–265, 1990.
8. R. Riolo. The effects and evolution of tag-mediated selection of partners in populations playing the iterated prisoner's dilemma. In *ICGA*, pages 378–385, 1997.
9. R. Riolo. The effects of tag-mediated selection of partners in evolving populations playing the iterated prisoners dilemma. *Nature 414*, pages 441–443, 2000.