

Effective Tag Mechanisms for Evolving Cooperation

Matt Matlock
Department of Computer Science
University of Tulsa
matt-matlock@utulsa.edu

Sandip Sen
Department of Computer Science
University of Tulsa
sandip-sen@utulsa.edu

ABSTRACT

Certain observable features (tags), shared by a group of similar agents, can be used to signal intentions and can be effectively used to infer unobservable properties. Such inference will enable the formulation of appropriate behaviors for interaction with those agents. Tags have been previously shown to be successful in social dilemma situations such as the prisoner's dilemma, and more recently have been shown to be applicable to other games by augmenting the standard tag mechanisms. We examine these more general tag mechanisms, and explain previously reported results by more thoroughly examining their fundamental designs. We show that these new tag mechanisms, along with some adjustments and augmentations, can be effective in enabling stable, socially optimal, and fair cooperative outcomes to emerge in general sum games. We focus, in particular, on general-sum conflicted games, where socially optimal outcomes do not necessarily yield the best results for individual agents. We argue that the improvements and understanding of these mechanisms expands the usability of tag mechanisms for facilitating coordination in multiagent systems. We argue that they allow agents to effectively reuse knowledge learned from interactions with one agent when interacting with other agents sharing the same features.

Categories and Subject Descriptors

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

General Terms

Algorithms, Performance, Economics

Keywords

tags, cooperation, evolution, learning, games

1. INTRODUCTION

The world is a phenomenally complicated place in which there exist a large number of entities that we can classify as agents (humans, animals, computer programs, etc). Unfortunately, each of these agents has only bounded cognitive

capabilities. In order to effectively respond to their environments, they rely on experience and limited data to make predictions about their best plan of action. These generalizations do not produce ideal outcomes, since they are not even guaranteed to accurately model the state of the agent's environment. However, they form an important component of human and animal reasoning. Without these generalizations an agent would likely be unable to process enough information to make robust, timely decisions [6]. For example, our experience with pressing large red buttons may have often been negative (false fire alarms, unnecessary emergency stops, accidental nuclear war). Therefore, we shall cease to press such buttons except in extreme circumstances. Fortunately, we can conclude that such generalizations are actually effective in real life (notice that we are all still alive).

In multi-agent systems research, we are often concerned with artificial agents which learn to play effectively with one another through the observations of other agents' behaviors [10, 13]. The focus of multi-agents research on tag mechanisms is complementary to this standard method. We are interested in investigating how strategies which have been shown to be successful against certain agents can be reused with other agents to develop compact behaviors which are effective and can be discerned by an agent with bounded cognitive capabilities. The particular method we are interested in is creating a clustering in the agent space which groups agents who possess similar observable characteristics. We believe that doing so will allow a strategy which was effective with a single agent in this classification group to be used effectively with other agents in that group. This clustering in agent space significantly reduces the number of interactions that must take place before an agent can learn to act effectively with many different agents.

The natural grouping of genetic schema in evolutionary computation is a convenient way of modeling this clustering [1, 2]. John Holland [5, 11] proposed to add to such models a primitive means of communication, called a "tag", to aid agents in identifying the groups they belong to. This tag corresponds to the observable features that humans and animals use in real life. They do not necessarily correspond to the behavior of the agents. Thus, an agent in this space consists of a strategy and a tag, which change via the evolutionary process as the interactions proceed. However, because this is an evolutionary process, it is likely that agents who possess similar tags also behave similarly due to common genetic ancestry. This clustering based on tags allows us to conclude that strategies which are useful against one agent will also be effective against other agents with similar

Cite as: Effective Tag Mechanisms for Evolving Cooperation, Matt Matlock, Sandip Sen, *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, Decker, Sichman, Sierra and Castelfranchi (eds.), May, 10–15, 2009, Budapest, Hungary, pp. 489–496
Copyright © 2009, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org), All rights reserved.

tags. Thus, limited interactions can produce robust strategies to effectively handle large and complicated agent spaces. Recent research into tag mechanisms applied to populations of interacting agents playing single shot interactions represented as stage games [9] have shown them to promote cooperation in variations of the Prisoner’s Dilemma [11, 12]. Most of these papers offer high-level explanations of how tags promote cooperation. However, a detailed analysis that clearly explains the fundamental subtleties and characteristics of these interactions is lacking in a large portion of the literature. More careful study of the choice of various parameters in tag simulations is also desirable.

We intend to show that the set of tag based mechanisms proposed by Matlock and Sen [7] succeed in producing high levels of Pareto and Social optimal outcomes on the general set of conflicted games, as well as characterize their individual behavior on the Prisoner’s Dilemma and the Coverage Game ¹ (see Figure 2). Secondly, we give a more in-depth analysis of the effect of the payoff sharing mechanism on the rate of socially optimal and Pareto optimal outcomes in the set of conflicted games. We believe that for cooperation to be sustained, fairness of outcomes should also be a key consideration. The primary result of this paper is an extension of the analysis of the paired reproduction mechanism so as to show that it performs as well as Payoff Sharing mechanisms, thus providing a more robust solution. In particular, we will introduce both mathematical and empirical rationale for the social optimality and fairness of the paired reproduction mechanism under all game scenarios, with a focus on the success of this mechanism on conflicted game scenarios.

2. PREVIOUS RESEARCH

McDonald and Sen [8] show that prior tag mechanisms can only achieve cooperation in situations where cooperation is equivalent to behavior imitation. They tested prior tag algorithms on anti-coordination games (see Figure 2) where complementary actions achieved the highest payoff. The previous tag algorithms were found to perform no better than random algorithms in this situation.

Matlock and Sen [7] introduced a number of augmentations to the basic tag scheme to remove this limitation of tag based algorithms. These augmentations initially centered around a new and improved partner selection function, but eventually required some further augmentations to achieve desirable performance levels. Their analysis of the problems with the original tag algorithm (expanded on from the work by McDonald and Sen [8]) indicated that the problem lay in the requirement that agents play only with agents who possess similar tags. This similarity implies a common genetic ancestry and thus a common behavior. Hence, only games which are solved using imitation will achieve desirable performance. They introduced new schemes by which agents could select partners arbitrarily from among the agent population using a simple function mapping from the tag space onto a boolean space. This boolean output indicated whether or not play with this individual was acceptable. They showed that this mechanism did improve upon the tag algorithm’s ability to handle situations requiring complementary strategies for effective per-

¹In the coverage game, two agents with different capabilities have to cover two areas with different coverage requirements. If both agents cover the same area their payoffs are less.

formance. Their mechanisms, however, failed on complex situations such as the prisoner’s dilemma. It appeared that cooperative groups under the prisoner’s dilemma, once invaded via mutation by a defector, quickly died out. Ultimately, the void in the tag space could not be quickly filled with more cooperative groups and the rather sub-par Nash Equilibrium was the consistent result of runs on the PD. Returning to the McDonald and Sen’s [8] analysis of standard tag algorithms, it became apparent that the reason older tag algorithms achieved such a high performance lay in the fact that simple mutation allowed for a great deal of interaction diversity. That is, under the original algorithm there were enough small groups that the death of one group was not a significant enough impact on the population as a whole.

To counteract this effect, Matlock and Sen [8] proposed two other mechanisms that further augmented their matching schemes. These were payoff sharing and paired reproduction, both of which (when combined with unilateral and mutual matching respectively) enabled their tag schemes to perform well on both Prisoner’s Dilemma and Anti-Coordination type games. However, it is felt that the investigation of the fundamental reasons for which these mechanisms succeeded in effectively solving the Prisoner’s Dilemma, and furthermore, the apparent lack of further investigation into paired reproduction, despite the fact that the latter is applicable in a much larger number of cases, was in sore need of remedy. In particular, payoff sharing performed much better than paired reproduction, but is limited because it cannot be used in situations where side-payments are not allowed. Thus, a more thorough analysis of these mechanisms is desired, hopefully one which will indicate the reasons why paired reproduction suffers in terms of performance and amends the identified maladies.

Let O be the set of outcomes available to agents playing the game G , N be the set of agents playing G , S the set of socially optimal outcomes, P the set of Pareto optimal outcomes and F the set of fair outcomes. Let $x_o(n)$ be the payoff to agent n in outcome o . We would like to see to what extent these tag mechanisms can fulfill the following three requirements:

Pareto Optimality, P:

$$o \in P \iff \forall o' \in O, \exists n \in N x_o(n) \geq x_{o'}(n)$$

For any other outcome preferred by some agent, another agent will prefer the current outcome.

Social Optimality, S:

$$o \in S \iff \forall o' \in O, \sum_{n \in N} x_o(n) \geq \sum_{n \in N} x_{o'}(n)$$

The sum of payoffs is maximized under this outcome.

Fairness, F:

$$o \in F \iff \forall o' \in O$$

$$\max_{n_1, n_2 \in N} |x_o(n_1) - x_o(n_2)| \leq$$

$$\max_{n_1, n_2 \in N} |x_{o'}(n_1) - x_{o'}(n_2)|$$

The maximum difference of payoffs under this outcome is minimized.

	C	D
C	3 3	1 4
D	4 1	2 2

Figure 1: Prisoner’s Dilemma

We believe that a good empirical test of whether or not these mechanisms fulfill the favorable criteria is to run the mechanisms on every game in the class of conflicted games. A conflicted game is defined as a game in which the following condition holds $\forall o \in O, \exists n \in N, o' \in O : x_o(n) < x_{o'}(n)$ that is, for any outcome there exists an agent who prefers some other outcome to the current one. Worded in another way, there exists no outcome which every agent prefers (no dominating outcome). The reason why this is a useful classification is that it draws a sharp difference in the convergence of algorithms depending on whether or not they encourage agents to be selfish (that is to seek the best outcome they can achieve) or to optimize the social welfare, which often results in a somewhat lower payoff for the agent. The prisoner’s dilemma (see Figure 1) is an excellent example of this sort of interaction. The social welfare optimizing outcome (C,C) results in a lower payoff than one of the (D,C) or (C,D) outcomes, thus algorithms which emphasize selfishness will tend to converge to (D,D) giving us a division between socially optimal and selfish algorithms. Of course, this criteria is not necessarily sufficient to guarantee that the outcome is also fair. Fairness, as we have defined above, essentially states that the agents achieve payoffs which are as closely balanced as possible, ensuring that one socially optimal agent is not getting, so to speak, the bad end of the stick. Thus, examining the strategy convergence of conflicted games under our augmented tag algorithm can tell us about the effectiveness of these mechanisms at coercing desirable social behavior out of self-interested agents.

3. TAG BASED MECHANISMS

Matlock and Sen [8] proposed to augment the partner selection process by creating a function mapping $f(T)$ which mapped tag values onto a boolean space which indicated whether or not play with a partner possessing tag T would be acceptable. Since the tag space they planned on using was the space of n -bit binary numbers, a ternary string of the form $M = \{m_i\}_{i=1}^n$ $m_i \in \{0, 1, *\}$ was used for $f(T)$ where $*$ is a don’t care symbol. The string M is said to match T iff $\forall i \in \{1, \dots, n\}, m_i \in M, t_i \in T \Rightarrow m_i = t_i \vee m_i = *$. $f(T)$ which returns true on a match and false otherwise. The advantage of using such a model lay in the fact that it lent itself readily to standard genetic algorithm mutation, reproduction and random generation and allowed a reasonable balance between general and specialized selection from within the tag space. They then defined the following matching schemes for partner selection.

3.1 Generalized Unilateral Matching

In this scheme an agent A can select a partner B to interact with iff the function $f_A(T_B) = true$. That is, if the matching string of agent A matches the tag of partner B. If this is fulfilled, A and B invoke their strategies and A receives the corresponding payoff from the interaction. B has no preference or benefit from this interaction. This

matching scheme performed well on AC type games but failed on PD type games. Insufficient performance in this matching scheme was attributed to complexities of interaction that arise from the fact that unilateral matching allows arbitrarily sized cyclic groups to form, according to the following mathematical condition: $f_1(T_2) = true, f_2(T_3) = true, \dots, f_{n-1}(T_n) = true, f_n(T_1) = true$. This complex cycle, once broken, would cause the death of all of the n agents in the cycle. This was only a contributor to the overall problem, but a requirement that would simplify these interactions was desired.

3.2 Generalized Mutual Matching

In this scheme and agent A can select a partner B to interact with iff both functions $f_A(T_B) = true$ and $f_B(T_A) = true$. Thus, both partners must consent to play with one another. Under this scheme B has preference but once again has no benefit from interaction since only A receives the payoff dictated by their strategy profile. This scheme unfortunately still exhibited poor performance on PD type games, but succeeded in solving AC games.

3.3 Payoff Sharing Mechanisms

In coalitional games it is often convenient to allow agents to give side payments to other agents in order to maintain a balance of fairness between the received payoffs in a game. In our situation, such a mechanism seemed particularly relevant in the generalized unilateral matching scheme. Returning to our condition of cyclic interaction $f_1(T_2) = true, f_2(T_3) = true, \dots, f_{n-1}(T_n) = true, f_n(T_1) = true$ we may well notice that any agent in this cycle whose payoff is low is unlikely to survive the fitness proportionate selections. Thus, in this cycle, some amount of payoff sharing (denoted by α) between an agent i and its partner $i + 1$ may serve to maintain this interaction (where $i + 1$ is the agent with low payoff). However, this fundamentally changes the payoff matrix of most conflicted games, and often works because the payoff matrix reflected no longer exhibits a lack of strictly dominating outcomes. For example, in Matlock and Sen [8], payoff sharing was noted to solve the prisoner’s dilemma game with the coefficient of sharing set to 0.4 (that is, 40 percent of an agent’s payoff was given to the partner at that iteration). This situation made the socially optimal outcome of the PD dominate under reciprocal interactions (that is, interactions where $f_1(T_2)$ and $f_2(T_1)$ are both true).

3.4 Paired Reproduction

In light of the shortcomings of payoff sharing and its limited domain of use (especially with respect to conflicted game scenarios), it was necessary to investigate other methods. Under both of the matching schemes discussed above, a tag mutation does not produce new cooperative interactions, but instead preserves the previous interaction via the matching string, or alters the current interaction, though not necessarily preserving cooperative strategy. Therefore, it was recognized that some mechanism to introduce a population with diverse interactions would be beneficial for convergence to socially optimal solutions. The particular mechanism was applied to the population with probability p_r (with the constraint that $p_r + p_c = 1$ where p_c is the probability of normal cloning reproduction) and selected an agent proportional to fitness, and subsequently selected some agent with whom a reciprocal interaction was possible (that is $f_1(T_2) =$

	M	C
M	1 1	4 3
C	3 4	2 2

Figure 2: Coverage Game.




Figure 3: Percentage of Cooperation in the Coverage Game Under Unilateral Matching

$f_2(T_1) = true$) under a fitness proportionate scheme. These agents were then reproduced, underwent tag mutations and then their matching strings were coerced so that the interaction condition $f_1(T_2) = f_2(T_1) = true$ was preserved. Though this mechanism showed promise as a robust solution to the two games examined, a more in-depth analysis yields far more useful results.

4. RESULTS

As has been discussed, we will be examining the effectiveness of the tag algorithm on the general class of conflicted games, which as of yet have not been extensively shown to be solvable via tag methods. We will run several tests on the set of all possible 2x2 conflicted games with ordinal payoffs, and also examine the results on a couple of games of particular interest. Namely, a coverage game (see Figure 2) and the Prisoner’s dilemma (see Figure 1). Hales and Edmonds showed that the original tag algorithm performed extremely well on games similar to the prisoner’s dilemma. In such games, effective coordination required both agents playing the same strategy. However, McDonald and Sen’s [8] analysis of the Tag Algorithm showed that it was unable to sustain cooperation levels better than random algorithms in such situations as anti-coordination games. It is thus of interest to us to examine closely both types of games. The coverage game gives us a conflicted interest anti-coordination game to serve as a suitable test system. In this game, the best outcomes are achieved when each agent plays an opposing strategy. We believe that a close examination of all of these conflicted games will demonstrate that the tag augmentations proposed here frequently produces socially optimal, fair outcomes. We will also present a mathematical discussion of the properties of the proposed paired reproduction operator.

4.1 Unilateral and Mutual Matching

The results for the effectiveness of Unilateral and Mu-

	Average Pareto Optimal Outcomes	Standard Deviation	Average Socially Optimal Outcomes	Standard Deviation
1	0.86	0.23	0.74	0.35
2	0.72	0.32	0.61	0.38
3	0.93	0.09	0.80	0.28
4	0.94	0.06	0.84	0.22
5	0.93	0.15	0.80	0.31
6	0.96	0.02	0.89	0.20
7	0.79	0.02	0.77	0.07
8	0.90	0.01	0.87	0.05

- 1 Unilateral Matching
- 2 Mutual Matching
- 3 Fitness Sharing 20% with Unilateral Matching
- 4 Fitness Sharing 20% with Unilateral Matching
- 5 Fitness Sharing 40% with Mutual Matching
- 6 Fitness Sharing 40% with Mutual Matching
- 7 Paired Reproduction 10% Frequency
- 8 Paired Reproduction 95% Frequency

Table 1: Tag Algorithm Performance Averaged Over All Conflicted Games at 500 iterations per run, 5 runs per game.




Figure 4: Percentage of Cooperation in the Prisoner’s Dilemma Under Unilateral Matching

tual matching schemes on the general class of 2x2 conflicted games with ordinal payoff ² is given in Table 1. This format, however, does not convey a number of interesting empirical properties of these two mechanisms. Thus, we will examine the games of interest under each particular matching scheme. Figures 3 and 4 show that the level of Pareto and socially optimal outcomes in each of the two games that were chosen for detailed testing under a unilateral matching scheme.

For the Coverage Game (2), we observe a high level of cooperative behavior under our matching scheme (see Figure 3), but an inadequate level under the old self matching system used by Hales and Edmonds [4] (see Figure 5). This is because self-matching matching tag schemes promote the play of identical strategies and hence the original tag mechanisms cannot perform any better than a random algorithm (approximately fifty percent optimal outcomes). On the other hand, the new tag algorithm allows agents to

²In a 2x2 ordinal payoff game each player rank orders the four possible outcomes. A game is conflicted if no outcome is most preferred by both players. There are 54 distinct 2x2 conflicted games with ordinal payoffs [3].




Figure 5: Percentage of Cooperation in the Coverage Game Under Hales' and Edmonds Tag System




Figure 6: Percentage of Cooperation in the Coverage Game Under Mutual Matching

select arbitrary partners from the tag space which enables anti-coordination strategies to propagate themselves. This is a notable improvement. It should be noted, that where Hales and Edmonds succeed (the Prisoner's Dilemma) our mechanism fails. The reason for the problems in this sort of interaction lies in the natural complexity arising in single matching schemes. The following situation is highly probable: let G_1, \dots, G_n be a sequence of tag groups such that the agents in each group have a matching function that is approximately identical to f_1, \dots, f_n where $f_i(G_{i+1}) = true$, and $f_n(G_1) = true$. Thus, the agents match each other's groups in a cyclic fashion. Suppose this entire cyclic chain is interacting in an optimal manner (that is, they are all playing the socially optimal outcome C C in the prisoner's dilemma game). Now, if a single defector is introduced into group G_i , the agents in group G_{i-1} will begin to perform poorly, because they have a non-zero probability of playing our defective agent in G_i , while the agent in G_i will take advantage of G_{i+1} to gain a higher level of fitness. Under this scheme, group G_i will fill with defectors, while G_{i-1} will gradually die off. Ultimately, group G_i will kill off group G_{i-1} which will cause G_{i-2} to be unable to find interaction partners. This will give them a low fitness and continue a chain of events which leads to the eventual extinction of the cyclic interaction chain G_1, \dots, G_n . This instability would encourage the simulation to converge to stable equilibria (such as D, D) in dilemma driven games such as the Prisoner's Dilemma. Furthermore, this instability makes it difficult for the simulation to maintain an equilibrium, as evidenced by the generally wild oscillations in each of the figures 3 and 4.

In order to remedy the problems with complex cyclic interactions, it was thought that introducing the requirement that an interaction can occur only if both agents match one another (that is $f_i(A_j) = f_j(A_i) = true$). This would alleviate all problems with cyclic interactions. Under mutual matching schemes, the comparison to self matching (at least empirically) is similar, with the expected exception that the graphs exhibit somewhat greater stability in the percentage of optimal outcomes in the Coverage Game (see figure 6). Unfortunately, as evidenced by figure 7, the prisoner's dilemma still exhibits inadequate levels of cooperation when compared to traditional tag mechanisms.

4.2 Payoff Sharing




Figure 7: Percentage of Cooperation in the Prisoner's Dilemma Under Mutual Matching

Given the unstable nature of the cyclic interactions discussed in the section on the unilateral matching scheme, it is desirable to come up with mechanisms by which this instability may be remedied. One such possibility is payoff sharing. This mechanism allows an agent to share a percentage, α , of his payoff with the agent with whom it is interacting. The goal is that this extra side payment will increase the fitness of a suboptimal agent so that it can be reproduced in further generations. For example, in the cyclic case, an agent n at some point in the chain may be interacting with another agent $n + 1$ who is played a strategy which yields a suboptimal payoff for agent n . Thus, the agent $n - 1$ who is interacting with n and may be receiving a better payoff, may share some percentage of its payment with n , hopefully sustaining this cycle.

In practice, the mechanism performs extremely well, both on the general set of conflicted games (see Table 1) and on the Prisoner's Dilemma game (see Figure 8). We also see an extremely low standard deviation on the conflicted games set, meaning that many of the games are solved reasonably well. Additionally, individual games are extremely stable, as evidenced by the representative graph on the PD game (see Figure 8). This is notably the only proposed mechanism that has succeeded in achieving any level of cooperative behavior in the prisoner's dilemma game. In order to more




Figure 8: Percentage of Cooperation in the Prisoner's Dilemma under Payoff Sharing with $\alpha = 0.4$

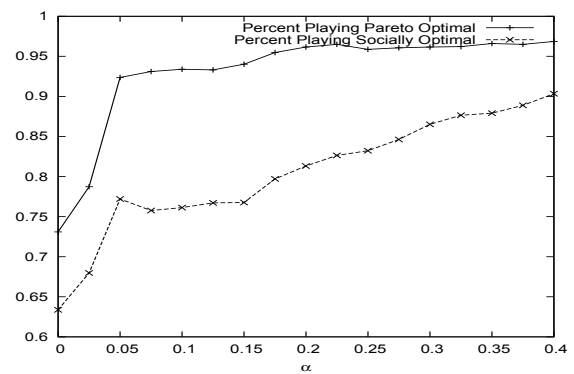


Figure 10: Effect of Varying α under Mutual Matching with Payoff Sharing.




Figure 9: Effect of Varying α under Unilateral Matching with Payoff Sharing.

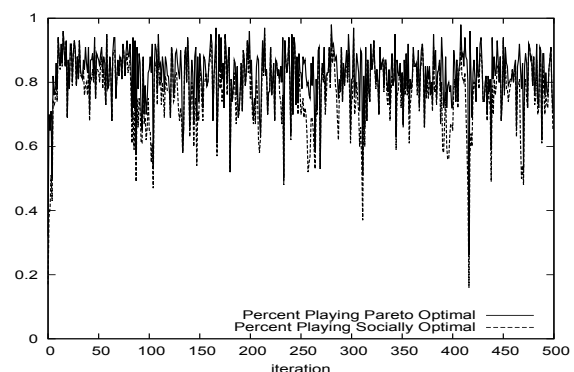


Figure 11: Percentage of Cooperation in the Prisoner's Dilemma under Paired Reproduction with Rate 0.1.

effectively grasp the effect of varying levels of fitness sharing α on the percentage of cooperative interactions over all conflicted 2x2 games, we have provided graphs for values of α between 0 and 40 percent (see Figure 9). These results are promising, however more stability is always positive, and we have tested these payoff sharing schemes under the mutual matching system as well. Given that mutual matching provides more stability in the general case, we see that, indeed, it also provides greater cooperation levels and game stability in both the conflicted games, and the 3 chosen games (see Table 1). Additionally, we have given a graph of the effect of values of alpha from 0 to 40 percent for the mutual matching payoff sharing scheme (see Figure 10). The unfortunate downside to the use of payoff sharing is that it is only applicable in environments in which the formation of coalitions (and thus side payments) are allowed. However, in the general sense, not every game can be treated as coalitional, thus we must find some mechanism which is general enough to be used in games where such side payments are not an available strategy. Thus, we introduce the paired reproduction operator.

4.3 Paired Reproduction

Originally, the paired reproduction operator was conceived of as an infrequently applied operator with the simple purpose of introducing diversity into the interactions that exist in the tag space (not strategic diversity, but rather, a

larger tag domain, with more refined matchings). Testing the paired reproduction operator with a low level of frequency (10 percent) on the Prisoner's Dilemma game, we see a marked increase in the stability and level of social optimality of outcomes in the simulation (see Figure 11), particularly compared to unilateral and mutual matching (figures 4 and 7). The level of Pareto and socially optimal outcomes under this technique, when compared with the effectiveness of payoff sharing, is unfortunately lacking. However, it seems clear that the mechanism is a necessary and useful alternative in non-coalitional environments.

Further analysis of the paired reproduction operator that we have introduced now seems necessary. For low values of p_r (consequently, high values of p_c the normal cloning reproduction probability) we have a marked increase in stability and cooperation in the broad set of conflicted games. However, we need to better characterize this mechanism's behavior by varying simulation parameters. The first experiment is to examine the percentage of socially optimal and Pareto optimal outcomes produced as the percentage of paired reproduction is varied from zero to one hundred percent. As can be seen in Figure 12, this results in a convex graph, indicating that values between 0 and 100 percent suffer from some limiting factor resulting from the introduction of paired reproduction. One possible explanation for the observed dip lies in the percentage of agents in any generation




Figure 12: Effect of Varying Mutual Reproduction Rate on Percentage of Pareto, Social, and Non-matching Outcomes.




Figure 13: Effect of Varying Paired Reproduction Probability on Percentage of Fair Outcomes in the General Set of Conflicted Games

for whom the evaluation of their matching function produced no viable partners. That is $f_i(j) = f_j(i) = true$ is not satisfied for any j in N (under the mutual matching scheme). This results from the allowance of a certain percentage of cloning and mutation, so that an agent who is cloned, but whose partner is not cloned, often belongs to this classification, along with those agents whose mutations are significant enough to affect their viability as partners. Examining Figure 12, we see that this convexity does indeed coincide with the concavity of the graph of the percentage of the agents in the simulation who exhibit the non-matching criteria as the percentage of paired reproduction is varied. Thus, it seems we have an all or nothing scenario with regard to our selection of reproduction rates under this scheme. However, an examination of the effect of higher, i.e., greater than 90 percent, paired reproduction rates on the percentage of fair outcomes in the simulation yields some interesting results.

Examining the effect of varying the reproduction rates on the percentage of fair outcomes in the set of conflicted games reveals the following curiosity: fair outcomes are preferred given the correct choice of p_r . For the general set of conflicted games (see Figure 13) we see a slight rise in the level of fair outcomes as we move away from 100 percent paired reproduction. This empirical evidence suggests a preference for fair outcomes under this reproduction scheme. Thus,

the mathematical hypothesis is that, when an outcome is socially optimal, then it is preferred to every other outcome under paired reproduction, if the probability of cloning is sufficiently small. Furthermore, when two outcomes are socially optimal, then the fairer of the two outcomes is preferred by the reproduction choice mechanism. Therefore, both socially optimal and fair outcomes are preferred by the mechanism.

We can characterize the probability of reproduction of a pair of agents n_1, n_2 by the following equation:

$$p(n_1, n_2) = p_c \left(\frac{x(n_1)}{x(N)} * \frac{x(n_2)}{x(N)} \right) + p_r \left(\frac{x(n_1)}{x(N)} P_{n_1}(n_2) + \frac{x(n_2)}{x(N)} P_{n_2}(n_1) \right)$$

where N is the set of all agents $x(n_i)$ the payoff to agent i in the last round, $x(N)$ is the sum of all payoffs in the last round, p_c is the probability of cloning and p_r is the probability of paired reproduction and $p_c + p_r = 1$. $P_{n_i}(n_j)$ is the probability choice function which returns the probability that agent n_j will be selected from agent n_i 's pool of viable partners. We now prove that paired reproduction favors socially optimal, fair outcomes.

PROPERTY 1. *When an outcome is socially optimal, then it is preferred to every other outcome under paired reproduction.*

Define $x_o(n_i)$ to be the payoff to agent n_i under outcome o . And let $x_o(n_1) + x_o(n_2) \geq x_{o'}(n_1) + x_{o'}(n_2)$ for outcomes o and o' . Then, clearly

$$p_r \left(\frac{x_o(n_1)}{x(N)} P_{n_1}(n_2) + \frac{x_o(n_2)}{x(N)} P_{n_2}(n_1) \right) \geq p_r \left(\frac{x_{o'}(n_1)}{x(N)} P_{n_1}(n_2) + \frac{x_{o'}(n_2)}{x(N)} P_{n_2}(n_1) \right)$$

we now must simply choose p_c so that:

$$p_c \left| \left(\frac{x_o(n_1)}{x(N)} * \frac{x_o(n_2)}{x(N)} \right) - \left(\frac{x_{o'}(n_1)}{x(N)} * \frac{x_{o'}(n_2)}{x(N)} \right) \right| < p_r \left| \left(\frac{x_o(n_1)}{x(N)} P_{n_1}(n_2) + \frac{x_o(n_2)}{x(N)} P_{n_2}(n_1) \right) - \left(\frac{x_{o'}(n_1)}{x(N)} P_{n_1}(n_2) + \frac{x_{o'}(n_2)}{x(N)} P_{n_2}(n_1) \right) \right|$$

Then we will have

$$p_c \left(\frac{x_o(n_1)}{x(N)} * \frac{x_o(n_2)}{x(N)} \right) + p_r \left(\frac{x_o(n_1)}{x(N)} P_{n_1}(n_2) + \frac{x_o(n_2)}{x(N)} P_{n_2}(n_1) \right) \geq p_c \left(\frac{x_{o'}(n_1)}{x(N)} * \frac{x_{o'}(n_2)}{x(N)} \right) + p_r \left(\frac{x_{o'}(n_1)}{x(N)} P_{n_1}(n_2) + \frac{x_{o'}(n_2)}{x(N)} P_{n_2}(n_1) \right)$$

Therefore, the mechanism prefers socially optimal outcomes given correct choice of p_c .

PROPERTY 2. *When two outcomes are socially optimal, the mechanism will prefer whichever of them satisfies the criteria of fairness.*

Suppose that we have two outcomes o and o' such that

$$x_o(n_1) + x_o(n_2) = x_{o'}(n_1) + x_{o'}(n_2)$$

and

$$|x_o(n_1) - x_o(n_2)| \leq |x_{o'}(n_1) - x_{o'}(n_2)|$$

and suppose without loss of generality that

$$x_o(n_1) \geq x_o(n_2)$$

and

$$x_{o'}(n_1) \geq x_{o'}(n_2)$$

Then, let

$$k = \frac{(x_{o'}(n_1) - x_{o'}(n_2)) - (x_o(n_1) - x_o(n_2))}{2}$$

so that

$$\begin{aligned} x_{o'}(n_1)x_{o'}(n_2) &= (x_o(n_1) + k)(x_o(n_2) - k) \\ &= x_o(n_1)x_o(n_2) + k(x_o(n_2) - x_o(n_1) - k) \end{aligned}$$

now, $k \geq 0$ and $x_o(n_1) \geq x_o(n_2)$ so that

$$k(x_o(n_2) - x_o(n_1) - k)$$

is a negative quantity, therefore

$$x_o(n_1)x_o(n_2) + k(x_o(n_2) - x_o(n_1) - k) \leq x_o(n_1)x_o(n_2)$$

so that we now have

$$\left(\frac{x_o(n_1)}{x(N)} \frac{x_o(n_2)}{x(N)} \right) \geq \left(\frac{x_{o'}(n_1)}{x(N)} \frac{x_{o'}(n_2)}{x(N)} \right)$$

and since both outcomes have equivalent payoff sums, we have established that fair outcomes are preferred.

Therefore, by choosing p_r to be adequately large (but not so large that $p_c = 0$), socially optimal fair outcomes will be preferred by this choice mechanism.

The additional benefit of using an argument based upon the smoothing principle in order to demonstrate this mathematical justification is that, in the event that an n -agent paired reproduction paradigm is implemented for playing n -agent games, we can apply this argument pairwise to agents in an n -group, showing that, in the limit (that is, as the smoothing argument is applied iteratively to a group of n agents) the result is that the preferred outcome converges on a socially optimal and fair situation.

5. CONCLUSIONS

We have argued for the need to develop tag mechanisms for the effective reuse of knowledge and strategies in interactions with diverse agents. In general, identifying similar agents can be extremely difficult, however, when agent's external features are inherited genetically, along with their internal strategies, the correlation between these two datasets enables agents to identify and reuse strategies which were found effective against similar agents.

In particular, using simple representations of external features, and simple models of interaction, recently developed mechanisms have been shown to be effective both in situations where the formation of coalitions and the exchange of side payments is allowed (payoff sharing in coalitional games) in general sum games.

Furthermore, we have demonstrated that these mechanisms are effective, not only for a few selected games, but on the broad category of conflicted games. This result is important, as it means that these mechanisms will encourage agents with differing interests and strategies to converge to socially optimal outcomes in most cases.

We believe that this is just the tip of the iceberg with regards to the richness of tag mechanisms, both in representations and complexity of interactions. It is clearly too early to attempt to represent nuanced visual representations, however the careful study and mathematical abstraction of external feature based communication seems to promise a yield of highly robust mechanisms for fostering cooperation among diverse agent groups.

Acknowledgment: A DOD-ARO Grant #W911NF-05-1-0285 partially supported this work.

6. REFERENCES

- [1] R. Axelrod. *The complexity of cooperation: Agent-based models of conflict and cooperation*. Princeton University Press, Princeton, NJ, 1997.
- [2] J. Bendor and P. Swistak. The evolutionary stability of cooperation. *American Political Science Review*, 91(2):290–307, June 1997.
- [3] S. J. Brams. *Theory of Moves*. Cambridge University Press, Cambridge: UK, 1994.
- [4] D. Hales and B. Edmonds. Evolving social rationality for mas using "tags". In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 497–503, Melbourne, Australia, July 2003. ACM Press.
- [5] J. H. Holland, K. Holyoak, R. Nisbett, and P. Thagard. *Induction: Processes of Inferences, Learning, and Discovery*. MIT Press, Cambridge, MA, 1986.
- [6] C. M. Judd and B. Park. Definition and assessment of accuracy in social stereotypes. *Psychological Review*, 100(1):109–128, Jan 1993.
- [7] M. Matlock and S. Sen. Effective tag mechanisms for evolving coordination. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 1340–1347, 2007.
- [8] A. McDonald and S. Sen. The success and failure of tag-mediated evolution of cooperation. In K. Tuyls, P. J. Hoen, K. Verbeeck, and S. Sen, editors, *LAMAS*, volume 3898 of *Lecture Notes in Computer Science*, pages 155–164. Springer, 2005.
- [9] R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1991.
- [10] L. Panait and S. Luke. Cooperative multi-agent learning: The state of the art. *Autonomous Agents and Multi-Agent Systems*, 11(3):387–434, 2005.
- [11] R. Riolo. The effects and evolution of tag-mediated selection of partners in populations playing the Iterated Prisoner's Dilemma. In *Proceedings of the Seventh International Conference on Genetic Algorithms*, pages 378–385. Morgan Kaufmann Publishers, Inc., 1997.
- [12] R. Riolo, M. D. Cohen, and R. Axelrod. Cooperation without reciprocity. *Nature*, 414:441–443, 2001.
- [13] P. Stone and M. Veloso. Multiagent systems: A survey from a machine learning perspective. *Autonomous Robots*, 8(3), July 2000.