The Evolution of Agent Strategies and Sociability in a Commons Dilemma

Enda Howley
Department of Information Technology
National University of Ireland
Galway
enda.howley@nuigalway.ie

Jim Duggan
Department of Information Technology
National University of Ireland
Galway
jim.duggan@nuigalway.ie

ABSTRACT

This paper explores the evolution of strategies in a n-player dilemma game. These n-player dilemmas provide a formal representation of many real world social dilemmas. These include issues as widespread as littering, voting and sharing common resources such as sharing computer processing time. This paper explores the evolution of altruism using an n-player dilemma. Our results show the importance of sociability in these games. Using a tag-mediated interaction model we will demonstrate the significance of this social characteristic.

Keywords

Evolution, Learning, Cooperation, Agent Interactions, Tragedy of the Commons, Tag-Mediated Interaction Models

1. INTRODUCTION

When a common resource is being shared among a number of individuals, each individual benefits most by using as much of the resource as possible. While this is the individually rational choice, it results in collective irrationality and a non Pareto-optimal result for all participants. These n-player dilemmas are common throughout many real world scenarios.

The computer science community is particularly concerned with how finite resources can be used most efficiently where conflicting and potentially selfish demands on those resources are common. Those resources may range from access to processor time or bandwidth.

One example commonly used throughout existing research is the *Tragedy of the Commons* [5]. This outlines a scenario thereby villagers are allowed to graze their cows on the village green. This common resource will be over grazed and lost to everyone if the villagers allow all their cows to graze, yet if everyone limits their use of the village green, it will continue to be useful to all villagers. Another example is the *Diners Dilemma* where a group of people in a restaurant agree to equally split their bill. Each has the choice to exploit the situation and order the most expensive items on the menu. If all members of the group apply this strategy, then all participants will end up paying more [4].

These games are all classified as n-player dilemmas, as they involve multiple participants interacting as a group. These games are similar to a number of well known two player games such as the *Prisoner's Dilemma* [1] and the *Battle of the Sexes* [2]. These games involve only two players interacting through simple pairwise interactions. N-player dilemmas have been shown to result in widespread defection unless agent interactions are structured. This is most commonly achieved through using spatial constraints which limit agent interactions through specified neighbourhoods on a spatial grid. Limiting group size has been shown to benefit cooperation in these n-player dilemmas [17].

In this paper we will examine an n-player dilemma, and study the evolution of strategies when individuals can bias their interactions through a tag mediated environment. Furthermore, we will show how certain strategies evolve with respect to their sociability towards their peers. The simulations presented in this paper use the n-player Iterated Prisoner's Dilemma (NIPD). The purpose of this paper is to examine the evolution of cooperation and sociability throughout the agent population in the NIPD. The research presented in this paper will deal with a number of specific research questions:

1. Can a tag-mediated interaction model be used to determine group interactions in a game such as the NIPD? This approach has been used to great effect in the two player Prisoner's Dilemma.
2. If agents have an evolvable trait which determines their sociability, then will this trait prove significant to the emergence of cooperation in the agent society?

The following section of his paper will provide an introduction to the NIPD and a number of well known agent interaction models. In the Experimental Setup Section we will discuss our simulator design and our experimental parameters. Our Results Section will provide a series of game theoretic simulations. Finally we will outline our conclusions and future work.

2. BACKGROUND RESEARCH

In this section we will introduce the NIPD game while also discussing some existing background research relevant to this paper.

Cite as: The Evolution of Agent Strategies and Sociability in a Commons Dilemma, Enda Howley and Jim Duggan, *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. XXX-XXX.

2.1 The N-Player Prisoner's Dilemma

The domain of modern game theory can be traced back to the seminal work of von Neumann and Morgenstern [16], and the subsequent surge of interest in the Prisoner's Dilemma [15]. In this paper we are most interested in n-player games. The n-player Prisoner's Dilemma is also known as the Tragedy of the Commons [5] and the payoff structure of this game is shown in Figure 1.

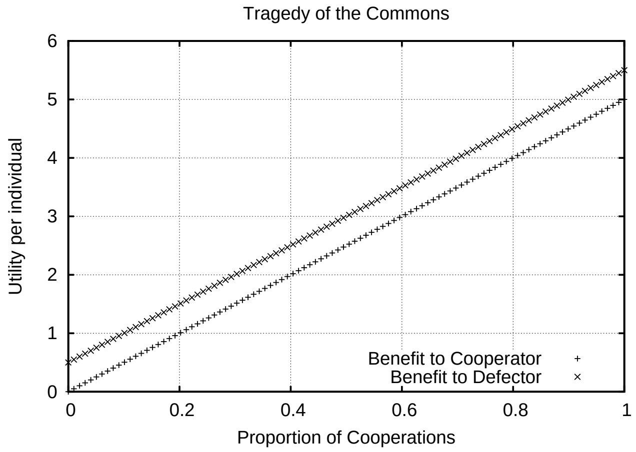

Figure 1: The N-Player Prisoner's Dilemma

On the horizontal axis is the fraction of cooperators in the group of n players in a particular game. On the vertical axis is the payoff for an individual participating in a game. There is a linear relationship between the fraction of cooperators and the utility received by a game participant. Importantly, the payoff received for a defection is higher than for a cooperation. The utility for defection dominates the payoff for cooperation in all cases. Therefore, an individual that defects will always receive a higher payoff than if they had chosen to cooperate. The result of this payoff structure should result in an advantage to defectors in the agent population. Despite this, a cooperator in a group of cooperators will do much better than a defector in a group of defectors.

This game is considered a valid dilemma due to the fact that individual rationality favors defection despite this resulting in state which is less beneficial to all participants. In our case where all individuals defect they all receive 0.5. This state is a non-pareto, sub-optimal, and collectively irrational outcome for the agent population. For all values of x this can be expressed as follows:

$$U_d(x) > U_c(x) \tag{1}$$

x is the fraction of cooperators while U_d and U_c are utility functions based on the fraction of cooperators in the group.

In this game a group of cooperators will receive a much higher payoff than a group of defectors. Since all moves are made simultaneously players cannot anticipate their peers actions. In this paper we will examine the iterated form of the n-player Prisoner's Dilemma (NIPD). Therefore, each game interaction is played over a number of iterations.

2.2 Agent Interaction Models

A number of alternative agent interaction models have been proposed and examined, such as spatial constraints and tag mediated interactions. Existing research involving the Prisoner's Dilemma has shown the significant impact of spatial grids on levels of cooperation [10, 9, 8]. Commonly used spatial structures involve Moore and vonNeumann neighbourhoods. These spatial constraints determine which individuals each agent interacts with.

Increasingly complex aspects of agent interactions have been examined. In relation to the NIPD, a number of authors have examined the effects of community structure on the evolution of cooperation [12, 11, 3]. These have shown that neighbourhood structure benefits cooperation in the NIPD game.

Tags are visual markings or social cues which can help bias social interactions [6]. They are a commonly used agent interaction model and can be considered akin to football supporters identifying each other through wearing their preferred team colours or jersey. Similarly individuals can identify each other in conversations through a common language, dialect, or regional accent. Tag-mediated interaction models are often considered as more abstract interaction models, and thereby useful to represent agent interactions more abstractly without the specific characteristics of a specific topology or implementation. The research presented by Riolo has demonstrated how tags can lead to the emergence of cooperation in the Prisoner's Dilemma [13, 14].

In this paper we will augment existing research to show the effects of using a tag-mediated interaction model to determine group interactions in the NIPD. The following section will provide a detailed specification of our simulator and the overall design of our experiments.

3. EXPERIMENTAL SETUP

In this section we will outline our agent structure, our agent interaction model and our evolutionary algorithm.

3.1 Agent Genome

In our model each agent is represented through an agent genome. This genome holds a number of genes which represents how that particular agent behaves.

$$Genome = G_C, G_T, G_S, \tag{2}$$

The G_C gene represents the probability of an agent cooperating in a particular move. Each agent has G_C gene which never changes throughout their lifetime. The G_T gene represents the agent tag. This is represented in the range $[0 \ldots 1]$ and is used to determine which games each agent participates. Finally, the G_S gene represents the sociability of each agent. This gene is also a number in the range $[0 \ldots 1]$ which acts as a degree of sociability for that individual agent. Initially these agent genes are generated using a uniform distribution for the first generation. Over subsequent generations new agent genomes are generated using our genetic algorithm.

3.1.1 Tag Mediated Interactions

In our simulations each agent interacts through a simple tag mediated interaction model. We adopt a similar tag implementation as that outlined by Riolo [13]. In our model each agent has a G_T gene which is used as their tag value. Each agent A is given the opportunity to make game offers to all other agents in the population. The intention is that this agent A will host a game and the probability other agents will participate is determined as follows.

$$d_{A,C} = |A_{GT} - C_{GT}| \tag{3}$$

This equation represents the absolute value between the tag values of two agents A and C. This value is used to generate two roulette wheels R_a and R_c for A and C. These two roulette wheels will then be used to determine agent A's attitude to C and agent C's attitude to A.

$$R_a = \text{ A's acceptance of C} \tag{4}$$

$$R_c = \text{ C's acceptance of A} \tag{5}$$

C will only participate in the game when both roulette wheels have indicated acceptance. The distribution of these roulette wheels are also influenced by each agents sociability gene. This gene acts like a scalar value which is used to reflect that some agents are more sociable than others and will therefore be more willing to play with their peers. This is shown in the following equation, where R_a represents the roulette wheel probability of entering a game.

$$R_a = d_{A,C} \times A_{GS} \tag{6}$$

Each agent in the population makes a game offer to all other agents, and the set of agreed players then participate in the NIPD game.

3.1.2 Genetic Algorithm

In our simulator we have implemented a simple genetic algorithm. In each generation individuals participate in varying numbers of games. Therefore, fitness is determined by summing all their payoffs received and getting an average payoff per game. In each generation, the top 10% of agents are copied directly into the following generation. The other 90% of the agent population in generation $G+1$ is generated through evolving new strategies based on agent fitness in G.

Individuals are selected through roulette wheel selection based on their fitness from generation G. Parent pairs are selected and then these are used to generate a single new agent offspring for generation $G+1$. Crossover occurs through averaging the two parents strategy genes $G_C, G_T, G_S,$. These averaged strategy genes are then used for the new agent. A 5% chance of mutation on each of these strategy genes is also used, and once this occurs a gaussian distribution is used to determine the degree of change.

4. EXPERIMENTAL RESULTS

In this section we will present a series of simulations showing the results of our experiments. Firstly, we will examine a set of graphs depicting the results from a single run over 1000 Generations. The aim of this single run is to show the inherent links between certain agent gene values and the overall cooperation throughout the agent population. Later in this section we will present simulations showing results from a number of experimental runs. These will demonstrate the overall stability of our results over multiple runs. All our simulations were conducted using an agent population of 100 agents.

4.1 A Single Run

The following simulations show data collected from a single run over 1000 generations.

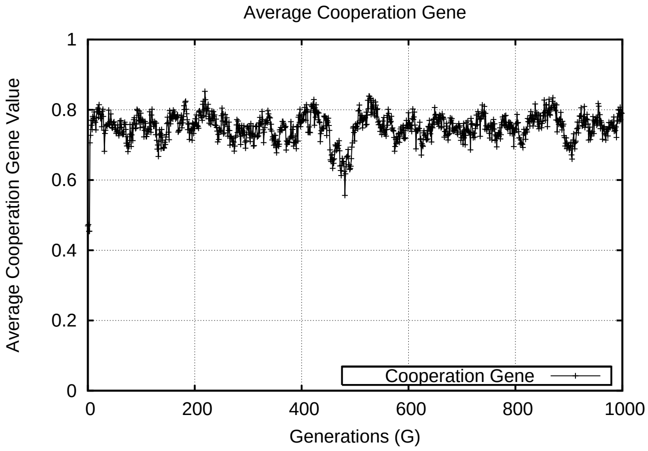

Figure 2: Average Cooperation Gene (1 Run)

In Figure 2, we can identify the rapid emergence of cooperation throughout the agent population. This graph depicts the average G_C gene throughout the agent population in each generation. The average at Generation 0 is around 0.5 as the initial population has an even distribution throughout all possible G_C gene values.

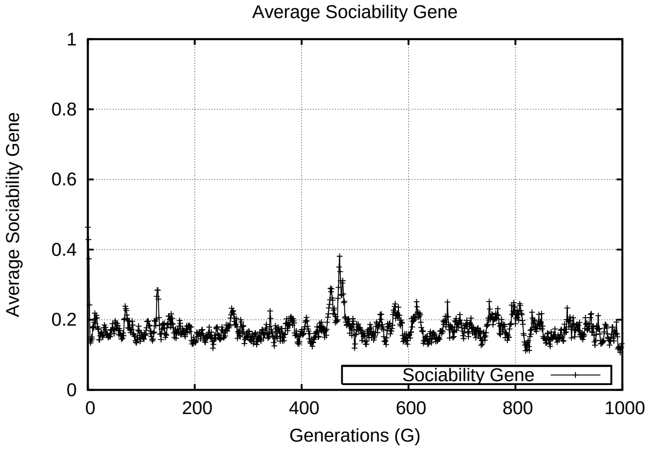

Figure 3: Average Sociability Gene (1 Run)

The results shown in Figure 3, show the average G_S gene throughout the agent population. These results show a rapid drop in the average G_S gene which reflects the tendency of the agent population to interact with fewer peers. The data shown here augments the results identified in Figure 2. The increased levels of cooperation throughout the population are closely linked with the tendency of individuals to act less sociably. It is clear from the results that the heightened cooperative gene is linked directly with the lower sociability gene.

In Figure 4, the average G_T gene is shown. As with most tag models, this functions as a random marker which serves to bias interactions and in itself holds no actual meaning.

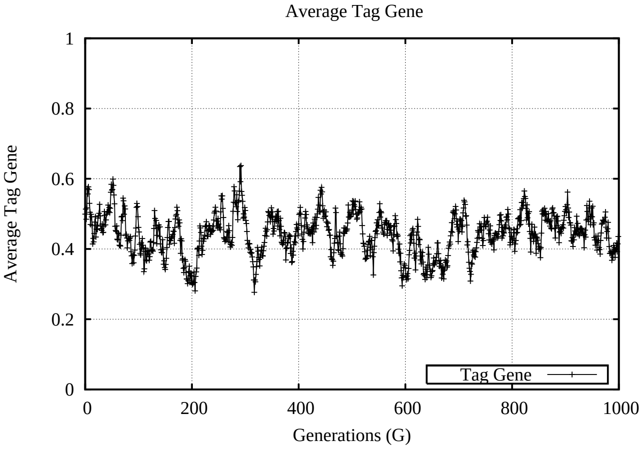

Figure 4: Average Tag Gene (1 Run)

As can be seen here when its value is averaged throughout the agent population it in influenced by the population evolutionary dynamics and occurrences of mutation that occur throughout the simulation.

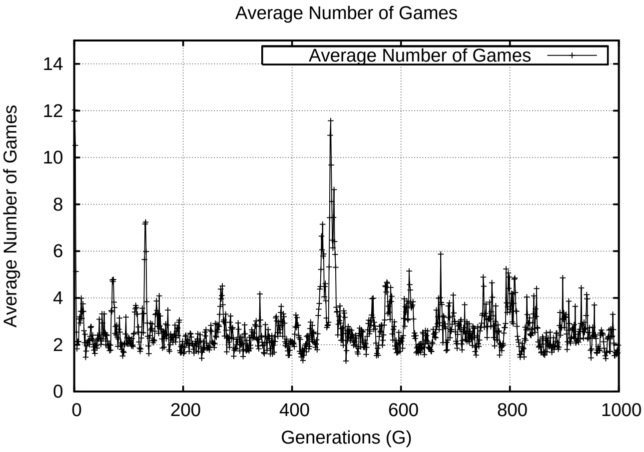

Figure 5: Average Number of Games (1 Run)

The results in Figure 5, depict the average number of games each agent participates in throughout successive generations. These results show the underlying dynamics that resulted in the heightened average cooperation shown in Figure 3. Once agents begin to participate in multiple n-player dilemmas they are exposed to exploitation and they are then heavily penalised. It is clear that cooperation is achieved through agents participating in as few games as possible. This serves to limit their exposure to potential exploitative peers.

The simulations shown are from a single run over 1000 generations. These simulations show the close relationship between the various agent gene values, and the collective behaviour of the agent population. For example around Generation 440 we can identify a period of increased sociability and a corresponding drop in cooperativeness throughout the population. This feature is clearly identifiable through examining the average cooperative gene results (Figure 2), the sociability gene results (Figure 3) and also the average game participation results (Figure 5).

4.2 Multiple Runs

In this section we will examine results recorded and average across many experimental runs of our simulator. The following graphs are averaged over 25 experimental runs. The purpose of these experiments is to demonstrate whether the overall trends identified previously are repeated over many runs.

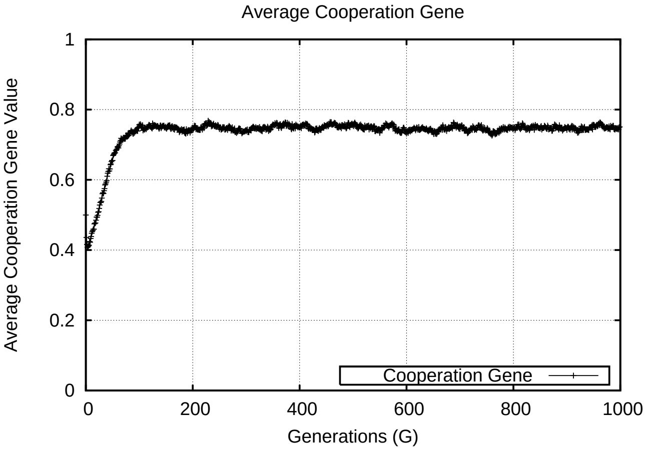

Figure 6: Average Cooperation Gene (25 Runs)

The data shown in Figure 6 show the average cooperation gene averaged over many experiments. It is clear from the results shown that the agent population consistently converges on cooperation throughout multiple experiments. This behaviour is a function of the most successful strategies being selected and used to evolve new strategies. Once the most successful strategies are cooperative, then this process is relatively clear. The reason for these strategies being the most successful is clarified in the following graph.

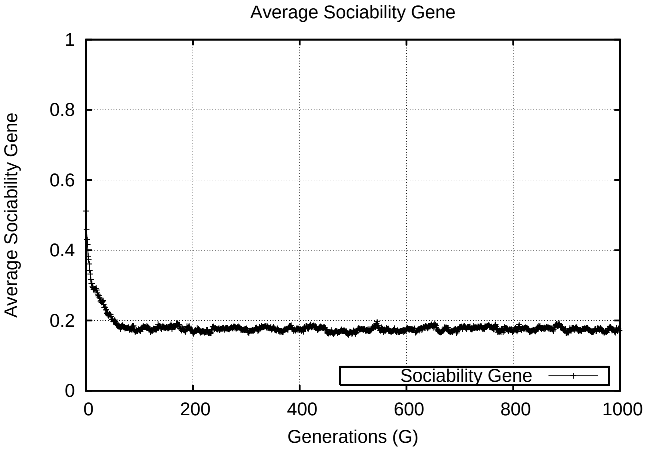

Figure 7: Average Sociability Gene (25 Runs)

In Figure 7 we can identify the prevalence of low G_S genes throughout the agent population. Through limiting game participating to a tiny number of games, each agent minimises the opportunity of less cooperative individuals to exploit them. Once cooperative strategies benefit heavily by limiting their interactions they receive heightened payoffs and then this feature is propagated throughout new agents in the population.

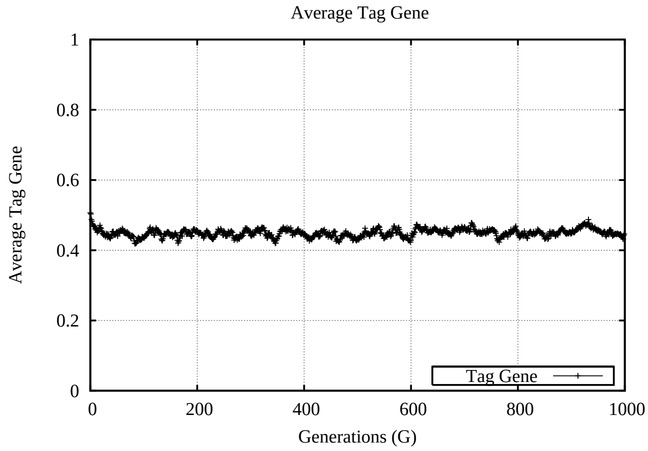

Figure 8: Average Tag Gene (25 Runs)

The data shown in Figure 8 shows the average tag gene values over successive generation. These tags are simply arbitrary markings and hold no great significance to the agent population. Due to this we observe no significant convergence of these tag values over multiple experimental runs.

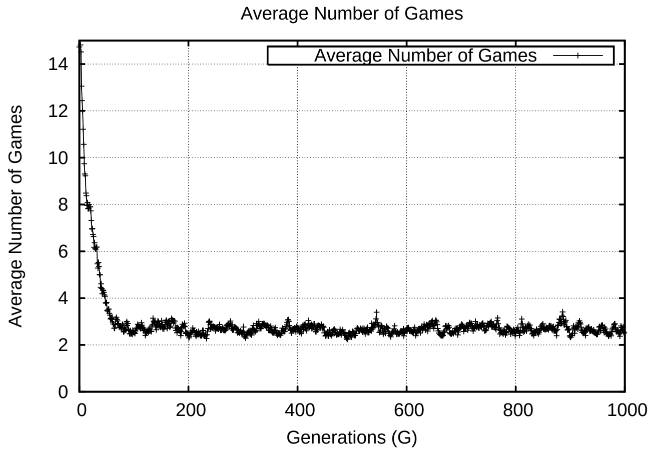

Figure 9: Average Number of Games (25 Runs)

The tendency to interact in as few games as possible is clearly shown in Figure 9, which depicts the average number of games each agent participates in. Once the most prevalent behaviour throughout the agent population is tending towards being less sociable, then the potential for exploiters to gain a foothold is very difficult. Exploiters require cooperators in order to exploit and thereby achieve high payoffs. If most individuals in the population are playing very few games then this task of exploiting is much more difficult and thereby these individuals are at a disadvantage.

The results in this section have focussed on the overall levels of cooperation and sociability of individuals playing the NIPD game. The advantage to an individual agent who chooses to defect is clear in the initial game definition, yet the experimental results presented in this section show an alternative reality. The results indicate a clear benefit to individuals who are less sociable and thereby choose to be far more discerning regarding game participation. This facilitates the emergence of cooperation and helps to maintain it over successive generations.

5. CONCLUSIONS

In this paper we have presented a series of game theoretic simulations examining the NIPD game. This typical Tragedy of the Commons game offers a formal interpretation of many real world scenarios involving an individuals choice between what is individually most beneficial or collectively most beneficial to the group. Examples of these include choosing to pay ones income tax, or deciding to recycle ones litter.

The results outlined in this paper demonstrate that despite there being a clear incentive to defect, cooperation can still emerge. This stems from the ability of individuals in our agent population to determine their degree of sociability towards their peers. This reinforces some existing literature involving the traditional Prisoner's Dilemma [7]. Howley and O'Riordan have shown the benefit of explicitly limiting limiting peer interactions in a fixed bias tag environment. In the case of the Prisoner's Dilemma this has the effect of limiting the potential of less-cooperative individuals exploiting cooperative individuals. Similarly, Yao and Darwin [17] have demonstrated the effects of limiting group size in the NPD. By explicitly limiting the number of participants in each n-player dilemma the benefit to cooperation is demonstrated. Our models reinforces these observations through an alternative approach. In our case we have not explicitly determined the social behaviour of our agent population. Instead we have allowed the agent population to evolve with respect to their cooperative and sociability genes.

Our results have demonstrated the importance of sociability in games such as the NIPD. Furthermore, we have also demonstrated the advantage to cooperative individuals who act less sociably towards their peers. Limiting game participation provides a very effective defence against exploiters. Earlier in our introduction we posed two specific research questions.

1. Our first research question queried the suitability of tags as an interaction model for an n-player game. Our tag-mediated interaction model successfully biased interactions in the agent population. Without our tag mediated interaction model, cooperation would have been less likely to emerge. As has been stated, tags have been used previously as a means of determining pairwise interactions in games such as the Prisoner's Dilemma. In this paper we have extended existing research to provide a tag model which can allow a population of agents to determine n-player game participation.

2. The second research question involved the sociability gene. The significance of the sociability gene in our simulations is clear from the direct link between the cooperativeness and sociability genes in all our simulations. While reinforcing observations from existing research, we have also demonstrated the ability of individuals to evolve cooperation through limiting interactions in an evolutionary context.

This paper has presented an evolutionary model capable of modeling sociability within the agent strategy genome. We have also shown how tags can be used to determine n-player games. Finally, our results have shown through an evolutionary model that there is a clear benefit to agent strategies who are cooperative in tandem with being less sociable through limiting their exposure to exploitation.

6. SUMMARY AND FUTURE WORK

This paper has examined a number of issues involving how agents bias their interactions in n-player games. We have shown that tags can be successfully adapted to bias agent interactions in a n-player game such as the NIPD. Furthermore, we have demonstrated how an agent population can engender and maintain cooperation through an evolvable sociability trait. In future work we hope to examine how cooperation can be engendered without limiting game participation so dramatically. It is possible that game participation can be increased once cooperation emerges, yet this will expose game participants to exploitation. We hope to examine a number of mechanisms which would allow this to occur.

7. ACKNOWLEDGMENTS

The authors would like to acknowledge the continued support of Science Foundation Ireland (SFI).

8. REFERENCES

[1] R. Axelrod. *The Evolution of Cooperation.* Basic Books, New York, 1984.

[2] J. S. Banks and R. L. Calvert. A battle-of-the-sexes game with incomplete information. *Games and Economic Behavior*, 4(3):347–372, July 1992.

[3] R. Chiong, S. Dhakal, and L. Jankovic. Effects of neighbourhood structure on evolution of cooperation in n-player iterated prisoner's dilemma. In *IDEAL*, pages 950–959, 2007.

[4] N. S. Glance and B. A. Huberman. The dynamics of social dilemmas. *Scientific American*, 270(3):76–81, 1994.

[5] G. Hardin. The tragedy of the commons. *Science*, 162(3859):1243–1248, December 1968.

[6] J. Holland. The effects of labels (tags) on social interactions. *Working Paper Santa Fe Institute 93-10-064*, 1993.

[7] E. Howley and C. O'Riordan. The emergence of cooperation among agents using simple fixed bias tagging. In *Proceedings of the 2005 Congress on Evolutionary Computation (IEEE CEC'05)*, volume 2, pages 1011–1016. IEEE Press, 2005.

[8] M. Nowak and R. May. Evolutionary games and spatial chaos. *Nature*, 359:826–829, Oct. 1992.

[9] M. Nowak and R. May. The spatial dilemmas of evolution. *Int Journal of Bifurcation and Chaos*, 3:35–78, 1993.

[10] M. Oliphant. Evolving cooperation in the non-iterated prisoner's dilemma: the importance of spatial organisation. In *Proceedings of Artificial Life IV*, 1994.

[11] C. O'Riordan and H. Sorensen. Stable cooperation in changing environments. In *AAMAS (3)*, pages 1529–1532, 2008.

[12] C. O'Riordan and H. Sorensen. *Stable Cooperation in the N-Player PrisonerŠs Dilemma: The Importance of Community Structure*, volume 4865 of *Lecture Notes in Computer Science (lncs) 4865*, pages 157–168. Springer-Verlag Berlin, 2008.

[13] 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.

[14] R. Riolo, M. Cohen, and R. Axelrod. Evolution of cooperation without reciprocity. *Nature*, 414:441–443, 2001.

[15] A. Tucker. A two person dilemma. In W. Poundstone, editor, *Prisoner's Dilemma.* Anchor Books, New York, 2 edition, 1950. Lecture at Stanford University.

[16] J. vonNeumann and O. Morgenstern. *Theory of Games and Economic Behavior.* Princeton, N.J.: Princeton University Press, 1980, c1972, 1944.

[17] X. Yao and P. J. Darwen. An experimental study of n-person iterated prisoner's dilemma games. *Informatica*, 18:435–450, 1994.