



Centrum voor Wiskunde en Informatica

REPORTRAPPORT

Stabilization of Tag-Mediated Interaction by Sexual Reproduction
in an Evolutionary Agent System

F. Alkemade, D.D.B. van Bragt, J.A. La Poutré

Software Engineering (SEN)

SEN-R0042 December 31, 2000

Report SEN-R0042
ISSN 1386-369X

CWI
P.O. Box 94079
1090 GB Amsterdam
The Netherlands

CWI is the National Research Institute for Mathematics and Computer Science. CWI is part of the Stichting Mathematisch Centrum (SMC), the Dutch foundation for promotion of mathematics and computer science and their applications.

SMC is sponsored by the Netherlands Organization for Scientific Research (NWO). CWI is a member of ERCIM, the European Research Consortium for Informatics and Mathematics.

Copyright © Stichting Mathematisch Centrum
P.O. Box 94079, 1090 GB Amsterdam (NL)
Kruislaan 413, 1098 SJ Amsterdam (NL)
Telephone +31 20 592 9333
Telefax +31 20 592 4199

Stabilization of Tag-Mediated Interaction by Sexual Reproduction in an Evolutionary Agent System

F. Alkemade D.D.B. van Bragt J.A. La Poutré
CWI

P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

ABSTRACT

The evolution of cooperation in a system of agents playing the iterated prisoner's dilemma (IPD) is investigated. We present results for the standard two-person IPD as well as the more general N-person IPD (NIPD) game. In our computational model, agents can recognize each other and decide whether to interact or not, based upon "tags" (labels). We consider the evolutionary stability of the evolving populations. Previous work is extended by introducing sexual reproduction (recombination) of agents and by analyzing its influence on the evolving populations. We observed the occasional formation of very stable cooperative societies, as opposed to previous results without sexual reproduction. These cooperative societies are able to resist invasions of "mimics" (defecting agents with the tag of a cooperating agent).

2000 Mathematics Subject Classification: 68T05, 91A05, 91A06, 91A28

1998 ACM Computing Classification System: I.2.6, I.2.11, J4

Keywords and Phrases: prisoner's dilemma game, evolutionary algorithms, tagging

Note: Work carried out under theme SEN4 "Evolutionary Systems and Applied Algorithmics". An extended abstract of this paper has been published in the proceedings of the First International Workshop on *Computational Intelligence in Economics and Finance* (CIEF'2000) in Atlantic City, U.S.A, February 27-March 3, 2000, page 945-949. This research was part of the project "Evolutionary Exploration Systems for Electronic Markets", funded by the Netherlands Organization for Scientific Research (NWO), project 612-60-008.

1. INTRODUCTION

Recently, interest in the evolution of strategic choice and the simulation of adaptive agents has surged among economists and game theorists [4]. An important topic in this field is the "spontaneous" emergence of cooperation in multi-agent systems [1]. It has been suggested [6, 8] that under *difficult* circumstances (i.e., when agents are tempted to display uncooperative behavior) "tags" might be useful in promoting the evolution of cooperation. Tags (labels or exterior characteristics) enable agents to recognize other players and bias their behavior accordingly.

Empirical evidence for the potential of tag-mediated systems has recently been provided by Riolo [8]. Riolo studied a multi-agent system in which agents play a short iterated prisoner's dilemma (IPD) game (of only four rounds) against each other. Under these conditions, it is very difficult for cooperation to emerge. When Riolo added a simple tagging mechanism, however, population dynamics changed dramatically and the agents were able to reach mutual cooperation earlier and over extended periods of time. However, results in [8] indicate that the evolving tag-using populations are still relatively unstable.

Riolo's pioneering work is extended in this paper by investigating the stability of the evolving populations, together with the role of the evolutionary algorithm (EA) that updates the agents' strategies over time. EAs are stochastic search methods based on the principles of natural genetic systems [7]. These algorithms deal with a population of individuals (referred to as *agents* here). Riolo used a very simple EA in his experiments. In particular, the reproduction of the agents was modeled as an *asexual* process (i.e., each parent produces one offspring with mutation as the only genetic operator). This model does not allow for the exchange of parts of strategies between agents

(recombination). This recombination process may play an important role in the evolutionary search process. Earlier experiments by Axelrod [2] have demonstrated, for instance, that cooperative societies form more frequently if recombination of the agents’ strategies occurs.

Recombination of strategies is also called *sexual* reproduction. During this process well-performing agents (the “parents”) exchange parts of their strategies to produce offspring. We will show that the sexual reproduction mechanism has a remarkably stabilizing effect on the evolving populations, as compared to asexual reproduction. This monotonicity has as a side effect that the cooperative populations, on the average, emerge after a longer period of time than in the asexual case. To obtain a speed-up of the emergence of cooperation, we therefore also consider a natural extension of EAs, incorporating a model for tag-directed parent (mate) selection.

Tag-directed mating replaces the standard (random) mating of parents with a more sophisticated matching algorithm, in which parents can select their own co-parent (based upon exterior mating characteristics). Preliminary results with this extended model already indicate an increasing number of robust cooperative populations.

This paper is organized as follows. First, we give a brief outline of our computational experiments in Section 2. Our extension of Riolo’s work is then presented in Section 3. The influence of sexual reproduction is investigated first in Section 3.1. Section 3.2 then discusses the influence of tag-directed mate selection. In Section 3 the two-person IPD is considered. The N-person IPD (NIPD), which is a more suitable model for the so-called “social dilemmas” [5], is evaluated in Section 4. Conclusions are drawn in Section 5.

2. EXPERIMENTAL SETUP

All computational experiments have been performed using Swarm,¹ a multi-agent software platform for the simulation of complex adaptive systems (see Fig. 1). The central component of Swarm is the “ModelSwarm”, which contains the population of players and the (evolutionary) algorithm to construct a new population (“newList”) from the previous one (“popList”). The “IPD” component contains the implementation of the IPD game (payoff values are the same as in [8], where the punishment payoff $P=1$, the temptation payoff $T=5$, the suckers payoff $S=0$, and the mutual cooperative outcome leads to a reward payoff $R=3$ for both players). Finally, the “ObserverSwarm” is a component for data extraction and visualization.

The EA that we implemented in the “ModelSwarm” consists of a canonical fitness-proportional selection scheme. Three different recombination operators were implemented: single-point crossover, two-point crossover and uniform crossover (see [7]). Because we are mainly interested in the *relative* performance of the agents, the raw fitness f_i (the average payoff over all played rounds) is normalized by taking $\hat{f}_i = (f_i - \mu)/\sigma + 1$, where μ is the mean population fitness. This implies that a player performing one standard deviation above the mean will (on average) get two offspring. Negative fitness values ($\hat{f}_i < 0$) were reset to 0.1 (see [7]) so that individuals with a very low fitness still have some small chance of reproducing.

The representation of the agents’ strategies differs from the representation Riolo used in his experiments. Here, we use pure (i.e., deterministic) strategies that are encoded as binary-valued chromosomes, whereas Riolo used mixed (i.e., probabilistic) strategies that were encoded as real-valued chromosomes. We adopt the genetic representation of IPD players as proposed by Axelrod [2]. Each agent has a memory capacity of three previous moves (one move of his own, and two of the opponent) in our model. The agent’s strategy (which specifies the agent’s next move) is then encoded in $2^3 = 8$ strategy bits $[s_0, \dots, s_7]$, located on the agent’s chromosome (see Fig. 2).

Three additional bits $[m_0, m_1, m_2]$ are present on this chromosome to determine the agent’s first move. Because the IPD game is very short in our simulations (only four rounds, as in [8]), the agents only have a memory capacity of three previous moves (instead of a memory of six moves in Axelrod’s 151-round simulations).

¹<http://www.santafe.edu/projects/swarm>.

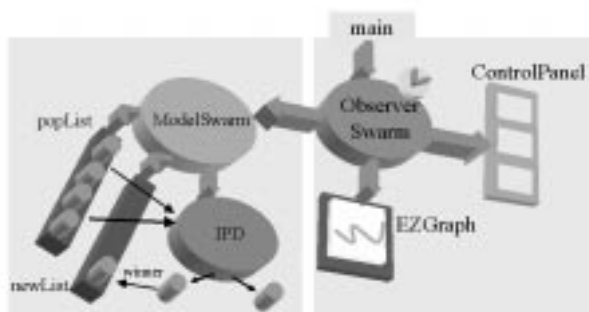


Figure 1: The Swarm model. On the left, the main components “ModelSwarm” and “IPD” are depicted. On the right, we see the “ObserverSwarm” component that extracts and displays data from the components on the left.

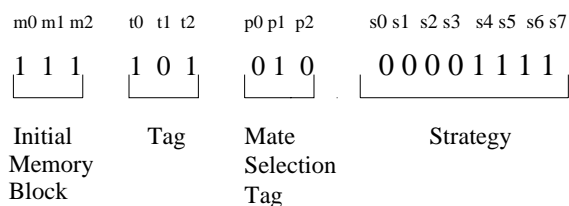


Figure 2: The structure of an agent’s chromosome when tag-based opponent selection and mate selection are both used.

An agent’s tag consists of three tag bits $[t_0, t_1, t_2]$ (see Fig. 2). We adopted the same algorithm for tag-based opponent selection as Riolo [8]. In this algorithm (see Fig. 3) each player can inspect the tag of a limited number of opponents. An opponent search cost (of 0.02, see Table 1) is associated with each tag trial. When the allowed number of tag search trials is exceeded the player is matched with a random opponent. After an opponent has been selected, the IPD game is played and the opponent search costs are subtracted from the player’s average payoff. In the experiments described here, the tag bias, which specifies the maximum allowed Hamming distance between two tags, is set to 0. In this case the agent in Fig. 2 will only accept opponents with a similar $[1,0,1]$ tag.

The agent’s chromosome, depicted in Fig. 2, also has a mate selection tag. The mate selection mechanism is implemented analogously to the opponent selection algorithm. After calculating the fitness of all agents in the population, a pool of parents is generated (using fitness proportional selection). Each parent then compares his mate selection tag with the mate selection tags of other parents in order to find a matching co-parent. When the maximum number of mate tag trials is exceeded, the parent will mate with a randomly selected co-parent. Note that the mate search costs are set equal to zero (see Table 1).

An overview of the model settings is given in Table 1. Whenever possible similar parameter settings as in [8] were chosen. Increasing or decreasing the mutation rate only had a small influence on the results when the mutation range was varied between 0.001 and 0.1 (a range of commonly recommended

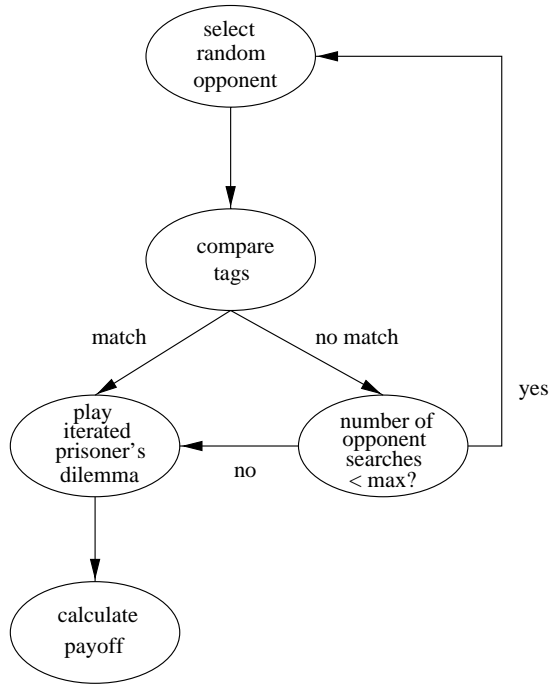


Figure 3: The algorithm for tag-based opponent selection. This algorithm will be repeated until ten opponents have been selected from the population of 400 individuals.

values [3]).

<i>Parameter</i>	<i>Value</i>
Population size	400
Number of tag trials	5
Tag size (in bits)	3
Tag bias	0
Opponent search costs	0.02
Mate search costs	0.00
Mutation rate (per bit)	0.025
Crossover probability	0.9
Number of moves per game	4
Number of games to play	10

Table 1: Model settings. We adopted the same terminology as in [8].

3. RESULTS AND DISCUSSION

3.1 Asexual vs. sexual reproduction

Figure 4 shows a typical run of a population of agents playing the IPD. The oscillatory pattern of the mean population fitness indicates that the population alternates between a state of mutual cooperation (when the fitness is close to the reward level $R = 3$) and a society of defectors (when the fitness is close to the punishment level $P = 1$). Societies of cooperators are frequently undermined by “mimics”:

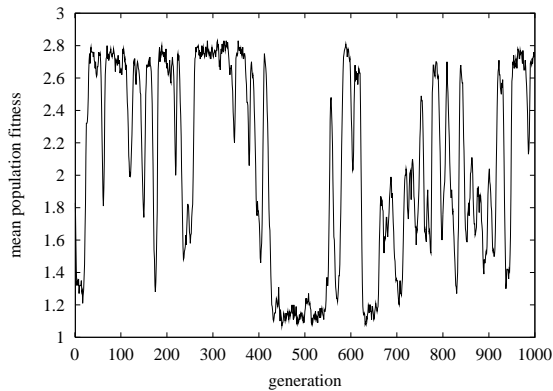


Figure 4: Mean fitness of a population of *asexual* agents playing the IPD. The fluctuations in fitness show that stable societies of cooperating agents do not form.

defecting agents with a tag associated with a group of cooperators. These mimics are not recognized as being defectors and can therefore successfully exploit the cooperative agents [8].

After introducing sexual instead of asexual reproduction, we obtain a significant change in population dynamics, see Fig. 5. The oscillatory behavior visible in Fig. 4 disappears and the individual

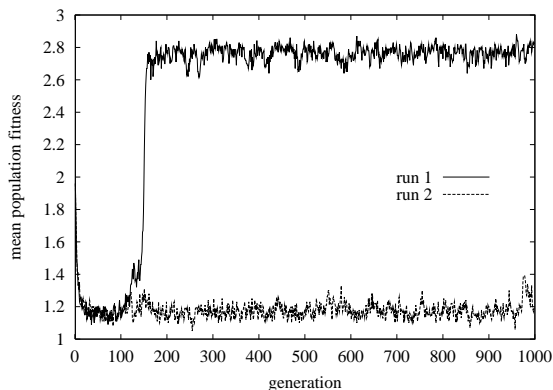


Figure 5: Mean fitness of a population of *sexual* agents playing the IPD: two typical runs. Notice that the oscillatory behavior observed in Fig. 4 disappears.

runs can now be classified as (1) runs in which a high mean fitness level is achieved and sustained, and (2) runs in which a society of (mainly) defectors forms. Examples of both cases are shown in Fig. 5. An important aspect is the monotonicity of the observed behavior: once cooperation emerges it persists over long periods of time. We, for instance, extended some runs for as long as 10,000 generations without the mean fitness dropping a single time below the 2.3 level after a cooperative period occurred.²

Average results (over 30 runs) are presented in Table 2.

In case of sexual reproduction, the mean fitness over the entire evolution history (MHF) remains rather low. This is due to the fact that it takes longer for cooperative societies to emerge when

²A society with cooperative periods is defined here as a society in which the mean population fitness remains above 2.3 for 20 successive generations at least once during the entire run.

	<i>Asexual</i>	<i>Sexual</i>	
		<i>(two-point)</i>	<i>(single-point)</i>
<i>Runs with cooperative periods (RCP)</i>	0/30	8/30	24/30
<i>Sustained stable cooperative runs (RSSC)</i>	0/30	8/30	24/30
$G_{over2.3}$	0.59 (0.09)	1.0 (0.0)	1.0 (0.0)
$G_{under1.7}$	0.21 (0.06)	0.0 (0.0)	0.0 (0.0)
$G_{to2.3}$	22 (21)	$7.7 \cdot 10^2$ ($7.6 \cdot 10^2$)	$2.5 \cdot 10^2$ ($2.1 \cdot 10^2$)
<i>Mean historical fitness (MHF)</i>	2.24 (0.1)	1.54 (0.5)	1.95 (0.44)

Table 2: Average fitness performance for asexual experiments and sexual experiments. (Numbers are calculated for 30 runs of 1,000 generations; standard deviations in brackets.)

recombination is used. Notice for instance in Table 2 that, when two-point crossover is used, the average number of generations it takes before the fitness first exceeds a value of 2.3, $G_{to2.3}$, increases to $7.7 \cdot 10^2$ in the experiments with sexual reproduction. But once a population has exceeded this fitness level, the population fitness never drops below 1.7 ($G_{under1.7}=0$) and in fact always remains above 2.3 ($G_{over2.3}=1$). $G_{over2.3}$ is defined as the fraction of generations in which the mean fitness is above 2.3, counting only generations after $G_{to2.3}$. Analogously, $G_{under1.7}$ counts the fraction of generations in which the mean fitness has dropped below 1.7. When one-point crossover or uniform crossover was used instead of two-point crossover, the same stabilizing effect was observed. However, cooperation was achieved more often using the less disruptive single-point crossover operator. Populations using single-point crossover evolved to cooperative societies in 24 out of 30 runs compared with only 8 out of 30 runs for two-point crossover. The mean fitness increased to 1.95 (0.44) compared with 1.54 (0.5) in case of two-point crossover (see Table 2). In case of uniform crossover, the mean fitness was only 1.34 (0.33) and cooperative societies emerged in 5 out of 30 runs.

The striking stability of the evolving cooperative societies could be explained by the influence of sexual reproduction on the number of mimics. Our experiments show that in cooperative societies distinct tagging groups form, i.e., for each tag one agent type becomes dominant. A mimic is then defined as an agent with the same tag as the dominating agent type, but with at least 5 different strategy bits. Defined this way, the proportion of mimics is 20-30 % in the experiments with asexual reproduction. This large proportion of exploiting agents contributes to the large fluctuations in mean fitness observed in Fig. 4. The proportion of mimics is much smaller (below 10 %) in the experiments with sexual reproduction.

3.2 Tag-directed mate selection

Results obtained using the relatively disruptive two-point crossover operator slightly improved when selective mating was introduced. Mate selection yields an increase in mean fitness (from 1.54 (0.5) to 1.68 (0.5)), mainly because, on average, mutual cooperation is discovered earlier ($G_{to2.3}$ decreases from $7.7 \cdot 10^2$ to $7.2 \cdot 10^2$). Again we found that, once was cooperation established, average fitness stayed above the 2.3 level. We also performed experiments with an evolving tag bias for mate selection (located on the chromosome). In this setup, a mate was only accepted if the Hamming distance between the two mating tags was equal to the tag bias. We found, in general, that agents have a strong preference for partners with a similar mating tag (the average mate tag bias converged to a small value).

4. TAGGING IN THE N-PERSON IPD GAME

The two-person IPD game can be used to model many social processes where cooperation is desirable but not easily obtained or sustained. There is an important class of cooperation problems, however, which can not be modeled adequately by the two-person IPD game. These problems are the so-called social dilemmas [5]. Social dilemmas can, however, be modeled by the N-person IPD (NIPD) game. Therefore, we extend our research on tagging to the NIPD game. Previous computer simulations of the NIPD with evolutionary algorithms [9] have shown that it becomes substantially more difficult to evolve cooperative societies if the number of players increases (i.e., when $N > 2$). To investigate whether the tagging mechanism also fosters stable cooperation in the NIPD, we performed a series of experiments for $N > 2$.

The NIPD game is described in detail in [9]. Like in the 2-person IPD, an agent can only choose between cooperation or defection in each round. The payoff (per round) is then determined as follows. If an agent cooperates, his payoff is equal to $2n_c - 2$, where n_c is the total number of cooperating agents. If he defects, he earns a payoff of $2n_c + 1$. In the 4-person IPD this payoff scheme would lead to a payoff of 1 in a society of defectors and a payoff of 6 in a cooperative society. The payoff matrix for an agent in the 4-person IPD is shown in Table 3.

Number of cooperators among the other $N - 1 = 3$ players		0	1	2	3
agent A	<i>Cooperates</i>	0	2	4	6
	<i>Defects</i>	1	3	5	7

Table 3: Payoff matrix for an agent in the 4-person IPD [9].

In our computational model, the strategy of an agent is now depending on (1) the agent's previous moves and (2) the number of cooperating agents in these rounds. If we consider the 4-person IPD, the length of the strategy block (see Fig. 2) is therefore equal to $2^9 = 512$ bits (3 bits to denote the previous 3 moves of the player, and 3 times 2 bits to denote the number of cooperators in the previous 3 rounds). The initial memory block then also has a size of 9 bits. As in the previous experiments with the 2-person IPD, the length of the tag is equal to the size of the initial memory block (i.e., 9 bits). Without a mating tag, the total chromosome length for agents in the 4-person IPD is therefore equal to 530 bits. An agent in the 4-person IPD repeatedly applies the algorithm for tag-based opponent selection as is shown in Fig. 3. The agent is thus allowed 5 tag trails per opponent.

In the evolutionary algorithm, the mutation probability (per bit) is reduced to 0.002 (from 0.025, see Table 1) to avoid an excessive increase of the number of mutations due to the much longer chromosome length for $N=4$. (With this mutation rate, on average one bit per chromosome is mutated for $N=4$ agents.) Values for the remaining parameters were kept the same as in the 2-person IPD (see Table 1).

Computational results for $N=4$ are reported in Table 4. The degree and stability of the emerging cooperation is measured by monitoring the mean historical fitness (MHF), the number of runs with cooperative periods (RCP), and the number of runs with sustained stable cooperation (RSSC). The horizontal lines in Fig. 6 indicate the mean population fitness for different values of the number of cooperators n_c (accounting for tag search costs). Notice that it is very difficult to achieve cooperative societies (the MHF remains low). Although the tagging mechanism increases average fitness levels, population-wide cooperation does not emerge. Remember that the average population fitness would still be equal to 2.25 if there is only one cooperator in each round (i.e., $n_c = 1$, corrected for the tag search costs). Note, however, that the number of runs with cooperative periods (RCP) and the number of runs with sustained stable cooperation (RSSC) significantly increase in case of tag-using agents.

To gain more insight in the nature of the cooperation that occurs, we examined the number of cooperators per tag group and in each round of the game. As in the 2-person IPD, distinct tag groups

emerge after approximately 100 generations. Most of these tag groups exhibit defective behaviour. Sometimes, however, a tag group discovers cooperative strategies. In most runs, this “cooperating” group was of a substantial size, periodically increasing average fitness levels to 3 or even higher. The maximum fitness measured during the experiments was approximately equal to 5, which indicates the emergence of a large group of cooperators (also given the fact given that the agents have to pay tag search costs).

The MHF data presented in Table 4 suggests that sexual cooperation does not help the emergence of cooperation. Notice for instance that the MHF decreases from ≈ 2.08 in the asexual experiments to ≈ 1.14 in the experiments with single-point crossover (when tag use is allowed). A more careful analysis however shows that the higher mean fitness in the asexual experiments is due to the fact that in these experiments incidentally a very high fitness is achieved. This cooperation level cannot be sustained however. The experiments further show that (due to slow convergence), fitness levels in the runs without sexual reproduction decrease very slowly, but once all cooperation is completely lost, it is very difficult for the asexual agents to reestablish it (at least not within the 10,000 generations we have examined). Figure 7 shows that this is not the case for experiments with sexual agents: after an initial period of low fitness levels a transient towards an increased level of cooperation occurs (after approximately 600 generations). This increased level of cooperation is then sustained in the remainder of the experiment.

<i>Sexual reproduction</i>	<i>Tags</i>	<i>MHF</i>	<i>RCP</i>	<i>RSSC</i>
No	No	1.08 (0.02)	0/30	0/30
<i>(single-point)</i>	No	1.11 (0.03)	0/30	0/30
<i>(two-point)</i>	No	1.11 (0.02)	0/30	0/30
No	Yes	2.08 (0.47)	25/30	3/30
<i>(single-point)</i>	Yes	1.14 (0.13)	8/30	4/30
<i>(two-point)</i>	Yes	1.62 (0.50)	20/30	14/30

Table 4: Influence of tagging and sexual reproduction in the 4-person IPD. Note that both the mean historical fitness (MHF), the number of runs with cooperative periods (RCP) and the number of runs with sustained stable cooperation (RSSC) increase if the agents can use tags. (Statistics are calculated for 30 runs of 10,000 generations; standard deviations in brackets.)

<i>Sexual reproduction</i>	<i>Tags</i>	<i>MHF</i>	<i>RCP</i>	<i>RSSC</i>
No	No	6.8 (0.2)	5/10	2/10
<i>(single-point)</i>	No	5.7 (0.3)	4/10	2/10
<i>(two-point)</i>	No	6.2 (0.2)	5/10	3/10
No	Yes	9.5 (2.4)	8/10	4/10
<i>(single-point)</i>	Yes	12.9 (1.3)	7/10	7/10
<i>(two-point)</i>	Yes	10.6 (1.5)	8/10	8/10

Table 5: Experimental results for the 8-person IPD when the number of iterations per game is increased to 10. Note that the mean historical fitness (MHF) increases when the use of tags is allowed. (Statistics are calculated for 10 runs of 10,000 generations; standard deviations in brackets.)

When the number of players was increased to 8 or 16 players, only defective societies were observed, with fitness levels always lower than 1.5.³ We found in additional experiments that this difficulty in

³In the 8-person IPD, the mutation probability is set equal to 0.0002; in the 16-person IPD this probability is set equal to 0.00003.

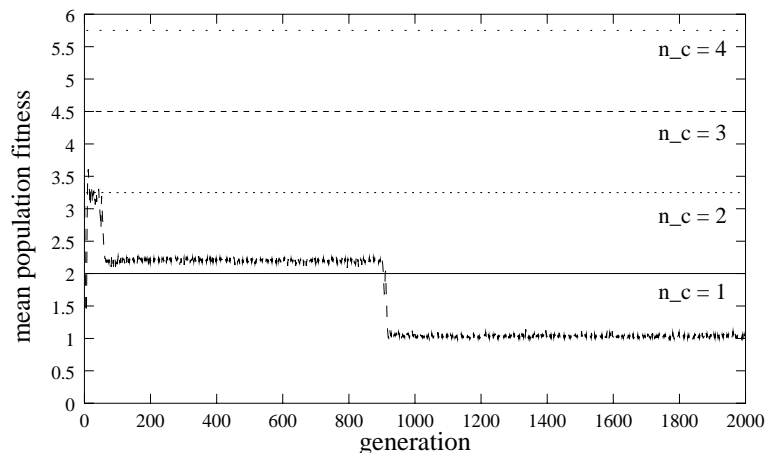


Figure 6: Influence of tagging and asexual reproduction in the 4-person IPD, one typical run with cooperation. Note the step-wise decrease of fitness as evolution proceeds. The horizontal lines indicate the fitnesses of populations with on (average) 1, 2, 3, or 4 cooperators.

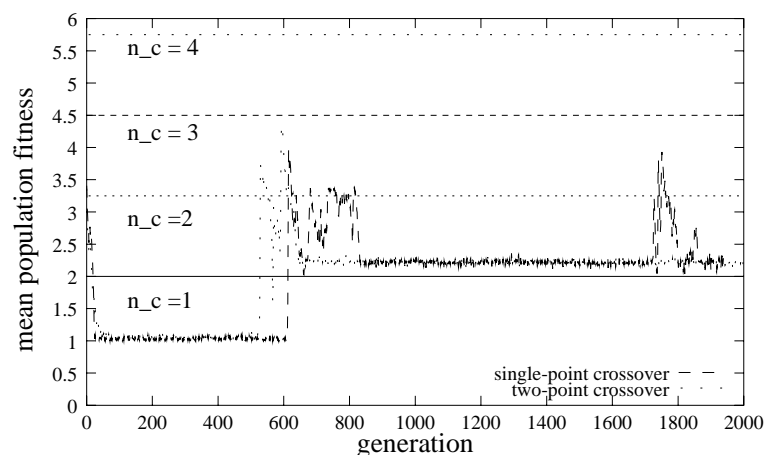


Figure 7: Influence of tagging and sexual reproduction in the 4-person IPD. Two typical runs with cooperation. Notice that, once some cooperation is achieved, the population stays out of the defective zone throughout the entire run.

achieving cooperation was caused mainly by the small number of rounds (namely 4, see Table 1) in the game. When the game length increases, average fitness levels rise, and cooperation is achieved more often. As an example, Table 5 shows the results for the 8-person IPD when the number of iterations is increased to 10. Remember that fitness values in the 8-person IPD lie between 1 (for $n_c = 0$) and 14 (for $n_c = 8$). Again we see that tags help to establish cooperation in societies of agents playing the NIPD. If we look at the population after 10,000 generations we also observe a strong convergence per tag group as was the case in the 2-person IPD.

In the 16-person IPD, increasing the number of iterations to 10 causes a small increase of the level of cooperation. In 10 runs of 10,000 generations each, cooperation (i.e., $fitness > 1$) emerged only once with tag-using agents. These results are roughly compatible with experiments from Yao and

Darwen [9], where no cooperation was found in a 16-person IPD without tags (with 100 iterations and agents with a memory of size 2). The failure to reach cooperation can be caused by the large search space (the chromosome length without tags is 32,783 in the 16-person IPD without tags).

5. CONCLUSIONS

We have studied evolutionary processes in multi-agent systems. In this paper, we have investigated the “evolution of cooperation” in a population of agents playing the tag-mediated (iterated) prisoner’s dilemma (IPD). Computational experiments have been performed using evolutionary algorithms (EAs). We have shown that the tagging mechanism and the reproduction process of the agents play a major role in the formation of stable cooperative societies.

In the 2-person IPD the population alternates between a state of mutual cooperation and a society of defectors in a model with asexual reproduction (i.e., when children are simple copies of their parents and mutation is the only genetic operator). A distinct behavior emerges if reproduction of the agents is sexual (i.e., when the parental strategies are recombined during the reproduction process). We observed, for instance, the formation of very stable societies of cooperative agents, a phenomenon not observed in the experiments with asexual reproduction. Furthermore, we found that cooperative societies emerge more frequently when the recombination operator is not too disruptive (e.g., a single-point crossover scheme). Finally, we proposed a tagging mechanism to enable biased partner selection. First results for this extended model were presented.

Results for the N-person IPD showed that (1) it becomes more difficult to evolve cooperative societies if the number of players is increased (i.e., $N > 2$), and that (2) tagging does help to achieve cooperation in the N-person IPD game. Furthermore, stable long-term cooperation emerges more frequently when sexual recombination of the agents’ strategies occurs (as in the two-person game studied in this paper).

References

1. R. Axelrod. *The Evolution of Cooperation*. Basic Books, New York, 1984.
2. R. Axelrod. The evolution of strategies in the iterated prisoner's dilemma. In Lawrence Davis, editor, *Genetic Algorithms and Simulated Annealing*, pages 32–41. Pitman, London, 1987.
3. T. Bäck, U. Hammel, and H.P. Schwefel. Evolutionary computation: Comments on the history and the current state. *IEEE Transactions on Evolutionary Computation*, 1:3–17, 1997.
4. K. Binmore and N. Vulkan. Applying game theory to automated negotiation. *Netnomics*, Vol. 1(1):1–9, 1999.
5. R.M. Dawes. Social dilemmas. *Annual Review of Psychology*, 31:169–193, 1980.
6. J.H. Holland. *Hidden order: How adaptation builds complexity*. MIT Press/Addison-Wesley, Reading (MA), 1995.
7. M. Mitchell. *An introduction to genetic algorithms*. MIT Press, Cambridge (MA), 1996.
8. R.L. Riolo. The effects and evolution of tag-mediated selection of partners in populations playing the iterated prisoner's dilemma. In Th. Bäck, editor, *Proceedings of the 7th International Conference on Genetic Algorithms*, pages 378–385. Morgan Kaufmann, 1997.
9. X. Yao and P.J. Darwen. An experimental study of N-person iterated prisoner's dilemma games. *Informatica*, 18:435–450, 1994.