Clans and Cooperation in the Iterated Prisoner's Dilemma

Bryant A. Julstrom
Computer Science and Information Technology
St. Cloud State University
St. Cloud, MN 56301 USA
julstrom@stcloudstate.edu

ABSTRACT

In evolutionary algorithms that evolve populations of strategies for the Iterated Prisoner's Dilemma, higher levels of cooperation evolve when the strategies engage in longer contests. When IPD-playing organisms are segregated into clans, with different strategies for contests with members of the same and other clans, populations evolve higher levels of cooperation for contests within clans. Cooperation between members of different clans increases with the length of their contests, but decreases when there are more clans. When organisms use the same strategy for contests with friends and with strangers, the difference between cooperation within and between clans is smaller, but persists.

Categories and Subject Descriptors

G.2.1 [**Mathematics of Computing**]: Discrete Mathematics—*Combinatorics*; I.2.8 [**Problem Solving, Control Methods, and Search**]: Heuristic Methods

Keywords

Iterated Prisoner's Dilemma, contest length, clans, cooperation

1. INTRODUCTION

In the Prisoner's Dilemma, two players decide, without communicating, whether to *cooperate* or *defect*. When both cooperate, each receives a payoff R; when both defect, each receives P. When one cooperates and one defects, the cooperator receives S and the defector receives T.

The dilemma occurs when $T > R > P > S$ and $2R > (T + S)$. The largest payoff goes to defection in the face of cooperation and the smallest to a betrayed cooperator; mutual defection receives the second lowest payoff. The players receive a higher average payoff when both cooperate. Thus, each player's best choice is defection, but when both defect, their payoffs are smaller than had both cooperated.

Repeated rounds of the Prisoner's Dilemma yield the Iterated Prisoner's Dilemma (IPD). Each player's payoff for the entire contest is the sum of its payoffs in the individual rounds. The players remember previous rounds and can adjust their play in response.

Many researchers have used the IPD to investigate competition and cooperation in social interactions. Axelrod [1, 2] hosted two tournaments of IPD strategies. The strategy called tit-for-tat (TFT) won both. TFT cooperates, then copies its opponent's last move; it rewards cooperation and punishes defection. Later, Axelrod [3] wrote a genetic algorithm that evolved strategies as effective as TFT. Crowley [4] represented strategies as sets of rules whose left hand sides were patterns of recent play and whose right hand sides were plays. He found that higher levels of cooperation evolved when the number of rules was moderate and contests were longer. Fogel [5] found that longer contests led to higher levels of cooperation, a result confirmed by Julstrom [6].

Organisms do not, however, interact equally often, and they may apply different strategies with strangers than with friends. Here, IPD-playing organisms are segregated into *clans*, and each organism consists of *two* IPD strategies. The organism uses one strategy for (longer) contests with members of its own clan and the other for (shorter) contests with members of other clans.

GECCO'12 Companion, July 7–11, 2012, Philadelphia, PA, USA
ACM 978-1-4503-1178-6/12/07.

2. A GENETIC ALGORITHM

Two-strategy IPD-playing organisms evolve in a genetic algorithm. A string of 70 C's and D's encodes each strategy, as Axelrod described [3]. The outcomes of the three previous IPD rounds determine a strategy's play in the current round. The GA initializes its population with random strategies. In each generation, it evaluates the organisms in a round-robin tournament. Payoffs are $T = 5$, $R = 3$, $P = 1$, and $S = 0$; these values satisfy the constraints that define the Prisoner's Dilemma. Each strategy's fitness is the sum of its payoffs over all its contests. Reproduction takes place within clans, via two-point crossover followed by mutation one-quarter of the time.

Following Axelrod [3], when an organism's score is more than one standard deviation above the mean of its clan, it participates in crossover twice. When its score falls within one standard deviation of the mean, it participates once. Otherwise, it does not reproduce. In the three sets of tests described below, each clan contained 50 organisms, and the GA ran through 500 generations, and each test was run independently 50 times.

3. EXTERNAL CONTEST LENGTH

In the first set of tests, the GA's population contained three clans. Contests within clans consisted of $L = 100$ rounds of the simple Prisoner's Dilemma, while contests between members of different clans consisted of $\ell = 5$, 10, 20, and 50 rounds. Figure 1 summarizes the tests, in which internal cooperation always exceeded external, and longer external contests lead to higher average cooperation.

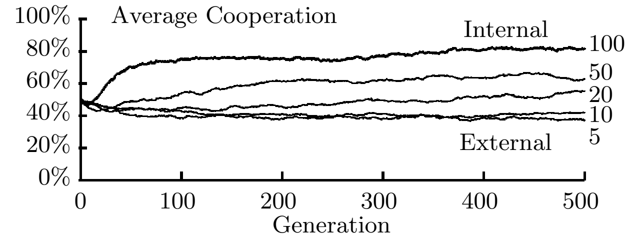

Figure 1: Average cooperation over 50 trials for external contests with lengths $\ell = 5$, 10, 20, and 50 rounds among three clans and for internal contests with length $L = 100$ rounds (and $\ell = 5$).

The average payoff per round cannot exceed 3, but in external contests, only one organism's payoffs accrue to its clan. It is possible, then, for a clan's average payoff per round to exceed 3; its members prosper by defecting when their competitors in other clans cooperate.

4. THE NUMBER OF CLANS

In the second set of tests, the GA was run with the internal and external contest lengths fixed at $L = 100$ and $\ell = 5$, respectively, but with $N = 3$, 5, 10, and 20 clans. Figure 2 summarizes these trials, in which external cooperation decreases as the number of clans increases. More clans provide more opportunities for each clan to find others of which its members can take advantage, and this outweighs the potential benefits of mutual cooperation.

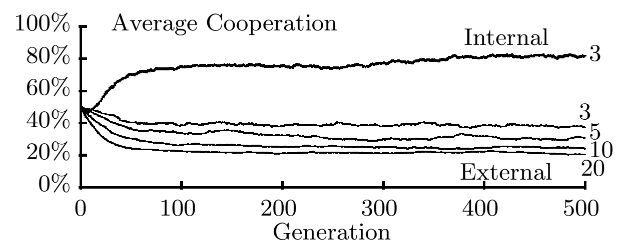

Figure 2: Average external cooperation as a function of the number N of clans. The top curve depicts the development of internal cooperation within three clans. The bottom four curves show the decline in evolved cooperation as the number of clans increases.

5. ONE STRATEGY

In the third set of trials, each organism applied the same strategy in both its internal and its external contests. The internal contest length was again $L = 100$, and the external contest length varied.

In all the sets of trials, average internal cooperation still exceeded average external cooperation, though by a smaller margin than with two-strategy organisms. Unlike the earlier tests, external cooperation tracked internal cooperation. Both increased as the generations succeeded each other, and the gap between the two averages narrowed. Figure 3 illustrates these observations for the 50 trials in which the external contest length ℓ was set to 10. Even with a single strategy, organisms respond to the fitness reward obtained by exploiting other clans.

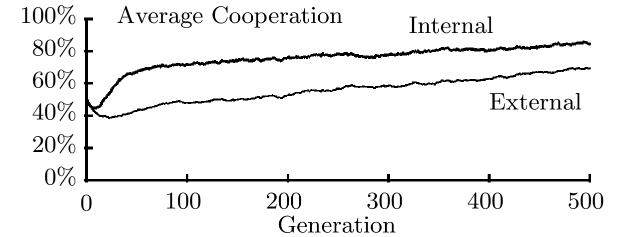

Figure 3: The evolution of average internal and external cooperation in 50 populations of one-strategy organisms, organized into three clans. The internal contest length was 100 and the external contest length was 10.

6. CONCLUSION

In a genetic algorithm, organisms compete in repeated rounds of the Iterated Prisoner's Dilemma. The population is divided into clans of equal size, and organisms reproduce only with members of their own clan. Each organism's fitness is the sum of its payoffs in IPD contests with all the other organisms in the population, amd organisms can use different strategies with members of the same clan and with strangers. Contests within clans are longer than contests between members of different clans, representing more interaction with friends than with strangers.

In repeated trials of the genetic algorithm, populations developed higher average levels of cooperation for contests within clans than for contests between members of different clans. When organisms applied different strategies in contests with friends and with strangers, average external cooperation increased as the external contest length increased, and decreased as the number of clans increased. Even when organisms applied the same strategy in all their contests, average external cooperation trailed average internal cooperation, though the gap narrowed and both were higher.

7. REFERENCES

[1] R. Axelrod. Effective choice in the prisoner's dilemma. *Journal of Conflict Resolution*, 24(1):3–25, 1980.

[2] R. Axelrod. More effective choice in the prisoner's dilemma. *Journal of Conflict Resolution*, 24(3):379–403, 1980.

[3] R. Axelrod. The evolution of strategies in the iterated prisoner's dilemma. In L. Davis, editor, *Genetic Algorithms and Simulated Annealing*, pages 32–41. Morgan Kauffmann, Los Altos, CA, 1987.

[4] P. H. Crowley. Evolving cooperation: strategies as hierarchies of rules. *BioSystems*, 37:67–80, 1996.

[5] D. B. Fogel. Evolving behaviors in the iterated prisoner's dilemma. *Evolutionary Computation*, 1(1):77–97, 1993.

[6] B. A. Julstrom. Effects of contest length and noise on reciprocal altruism, cooperation, and payoffs in the iterated prisoner's dilemma. In T. Bäck, editor, *Proceedings of the Seventh International Conference on Genetic Algorithms*, pages 386–392, San Francisco, CA, 1997. Morgan Kaufmann Publishers.