Toward Adaptive Cooperative Behavior

Elpida S. Tzafestas

Institute of Communication and Computer Systems
Electrical and Computer Engineering Department
National Technical University of Athens
Zographou Campus
Athens 15773, GREECE
brensham@softlab.ece.ntua.gr

Abstract

We analyze the Iterated Prisoner's Dilemma and the performance of GRADUAL, the best behavior found so far. This behavior has the undesirable property of permanent memory, which would be detrimental for stability. As a solution to the permanent memory problem, we propose an adaptive tit-for-tat behavior that uses a self-regulated estimate of the opponent's friendliness. On a second level, we demonstrate that an additional self-regulation loop, parallel to the first, is necessary to ensure performance comparable to GRADUAL's. A number of theoretical conclusions are drawn, the most prominent being that the actual cooperative potential of the behavior is the by-product of the double self-regulation loop and that the second regulation loop concerns the parameters that define the temporal dynamics of behavior. Our results trigger a discussion on stability as well as on the nature of the cooperation problem itself.

1. Introduction

A major issue on the intersection of artificial life and theoretical biology is cooperative behavior between selfish agents. The cooperation problem states that each agent has a strong personal incentive to defect, while the joint best behavior would be to cooperate. This problem is traditionally modeled as a special two-party game, the Iterated Prisoner's Dilemma (IPD).

At each cycle of a long interaction process, the agents play the Prisoner's Dilemma. Each of the two may either cooperate (C) or defect (D) and is assigned a payoff defined by the following table.

Agent	Opponent	Payoff
C	C	3 (= Reward)
C	D	0 (= Sucker)
D	C	5 (= Temptation)
D	D	1 (= Punishment)

Usual experiments with IPD strategies are either tournaments or ecological experiments. In tournaments, each strategy plays against all others and scores are summed in the end. In ecological experiments, populations of IPD strategies play in tournaments and successive generations retain the best strategies in proportions analogous to their score sums.

The first notable behavior for the IPD designed and studied by Axelrod (Axelrod & Hamilton 1981, Axelrod 1984) is the Tit For Tat behavior (TFT, in short) :

Start by cooperating,

From there on return the opponent's previous move.

This behavior has achieved the highest scores in early tournaments and has been found to be fairly stable in ecological settings. TFT demonstrates three important properties, shared by most high scoring behaviors in IPD experiments :

- It is good (it starts by cooperating)
- It is retaliating (it returns the opponent's defection)
- It is generous (it forgets the past if the defecting opponent cooperates again).

Further strategies include stochastic ones (Nowak & Sigmund 1992), the Pavlov strategy (Nowak & Sigmund 1993) that cooperates when it has played the same move as its opponent etc. In the literature we may also find studies in an evolutionary perspective (Fogel 1993), theoretical or applied biological studies (Axelrod & Dion 1988, Feldman & Thomas 1987, Milinski 1987) and studies of modified IPD versions (Stanley et al. 1994).

The best designed behavior found so far in the literature is GRADUAL (Beaufils et al. 1996) which manages to achieve the highest scores against virtually all other designed behaviors. This behavior starts by cooperating and then plays Tit For Tat, except that it does not defect just once to an opponent's defection. Instead, it responds by playing blindly (nxD)CC, where n is the opponent's number of past defections. That is, GRADUAL responds with DCC to the first opponent's defection, DDCC to the second, etc. The justification given for the performance of this behavior is that it punishes the opponent more and more, as necessary, and then calms him down with two successive cooperations.

The motivation for our work has been our conviction that a behavior comparable to GRADUAL could be found, that has not permanent, irreversible memory. Instead, we are after a more adaptive tit-for-tat based model that would demonstrate behavioral gradualness and possess the potential for stability in front of changing worlds (opponent replacement etc.).

Our medium term objective is to integrate such adaptive cooperative behaviors together with regular task-achieving behaviors in animats acting in simulated or artificial worlds.

Finally, another background motivation for our work has been the will to apply the same regulation principles used in (Tzafestas 1995b, chapters 4 and 5) and (Tzafestas 1998), as another step toward the development of a regulation theory.

2. Analysis of GRADUAL

Let us examine the high scores that GRADUAL obtains against other behaviors. Designed behaviors found in the literature usually fall in one of three categories :

- Behaviors that use feedback from the game, usually cooperative behaviors unless the opponent defects, in which case they use a retaliating policy (for example, "tft" retaliates once, "grim" retaliates forever, "gradual" retaliates increasingly long, etc.).
- Behaviors that are essentially cooperative and retaliating, but start suspiciously by playing a few times D in the beginning, so as to probe their opponent's behavior and decide on what they have to do next (for example, suspicious tft (STFT) and the "prober" behavior of Beaufils et al. 1996)
- Behaviors that are clearly irrational, because they don't use any feedback from the game (for example, the random behavior and all blind periodic behaviors such as CCD, DDC etc.).

A behavior will maximize its score, if it is able to converge to cooperation with all behaviors of the first two categories and converge to defection against behaviors of the third category. Steady defection against periodic behaviors is necessary in order to achieve the highest possible score. For example, in the case of CCD, a TFT agent converges quickly to responding DCC and gets an average score of (5+3+0)/3=2.66 per move. On the contrary, the same agent (CCD) and the Always-Defecting (ALLD) agent get against CCD an average score of 3+3+1=2.33 and 5+5+1=3.66 per move, respectively. The GRADUAL behavior fulfills both of the above specifications, because it responds with two consecutive C's after a series of defections, giving the chance to STFT or prober behaviors to revert to cooperation, and converges to ALLD against irrational behaviors. A solution to the permanent memory problem has to demonstrate the same property.

3. The solution : Adaptive tit-for-tat

The adaptive behavior that we are seeking should be essentially tit-for-tat, in the sense of being good, retaliating and forgiving. Moreover, it should demonstrate fewer oscillations between C and D. To this end, it should have an estimate of the opponent's behavior, whether cooperative or defecting, and react to it in a tit-for-tat manner. The estimate will be continuously updated throughout the interaction with the opponent. The above may be modeled with the aid of a continuous variable, the world's image, ranging from 0 (total defection) to 1 (total cooperation). Intermediate values will represent degrees of cooperation and defection. The adaptive tit-for-tat model can then be formulated as a simple linear model :

Adaptive tit-for-tat

If (opponent played C in the last cycle) then

world = world + r(1-world), r is the adaptation rate*

else

world = world + r(0-world)*

If (world >= 0.5) play C, else play D

The usual tit-for-tat model corresponds to the case of r=1 (immediate convergence to the opponent's current move). Clearly, the use of fairly small r's will allow more gradual behavior and will tend to be more robust to perturbations.

Now, let us simulate the behavior of the adaptive tit-for-tat agent against all three types of behaviors described earlier.

For initially cooperative behaviors with feedback and a retaliation policy, the model cooperates steadily and converges quickly to total cooperation (as is shown in the value of the world variable, see fig. 1 against grim or tft).

For suspicious or prober behaviors, the model plays exactly like tit-for-tat, while the value of the world variable oscillates around the critical value of 0.5 (see fig. 2 against suspicious tft).

For periodic behaviors, the value of the world variable converges quickly to oscillations around the characteristic value of "number_of_C's/number_of_D's" in the opponent's period (see figs. 3 and 4 against CCD and CDD, respectively).

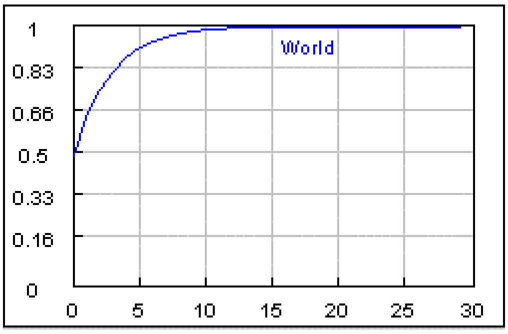

Figure 1. History of the world variable during interaction of the adaptive tit-for-tat agent with a tit-for-tat or grim behavior (r=0.2, world(0)=0.5).

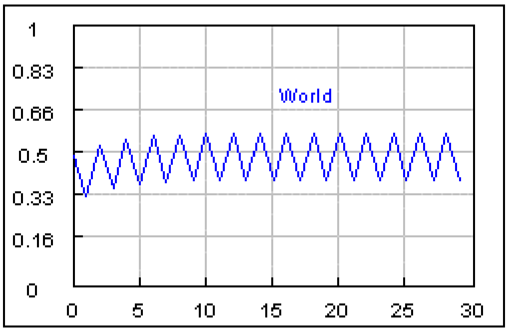

Figure 2. History of the world variable during interaction of the adaptive tit-for-tat agent with a suspicious tit-for-tat behavior (r=0.2, world(0)=0.5).

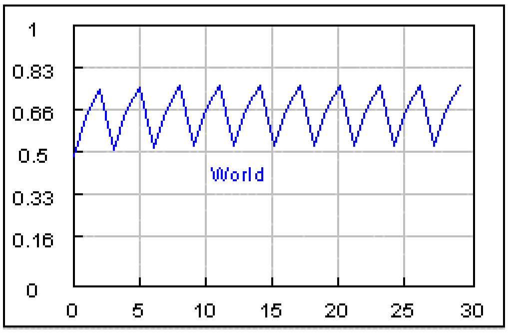

Figure 3. History of the world variable during interaction of the adaptive tit-for-tat agent with a CCD behavior (r=0.2, world(0)=0.5).

From the above, it can be seen that this first version of the model suffers from manipulation of the world variable by the opponent. This shows as stabilization of the agent to an oscillatory behavior (as is the case against stft) or a steady cooperative behavior against irrational agents (as is the case against CCD). To bypass this problem, we exploited our observation that different rates for cooperation and defection (r_c and r_d, respectively) yield different results. More specifically, we observed that the adaptive tit-for-tat agent manages to get opponents such as stft or the prober to cooperate if $r_c>r_d$, while it manages to fall to steady defection against periodic behaviors if $r_c<r_d$ (see figs. 5 and 6, against stft and CCD).

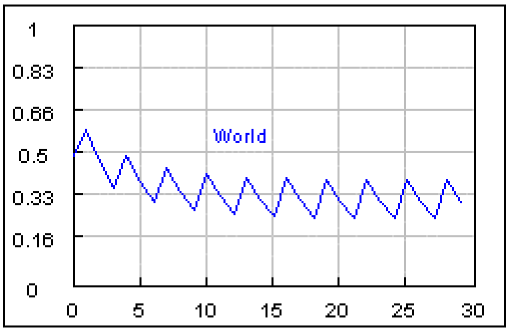

Figure 4. History of the world variable during interaction of the adaptive tit-for-tat agent with a CDD behavior (r=0.2, world(0)=0.5).

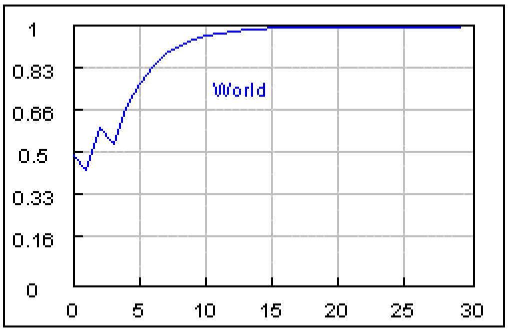

Figure 5. History of the world variable during interaction of the adaptive tit-for-tat agent with a suspicious tit-for-tat behavior (r_c=0.3, r_d=0.1, world(0)=0.5).

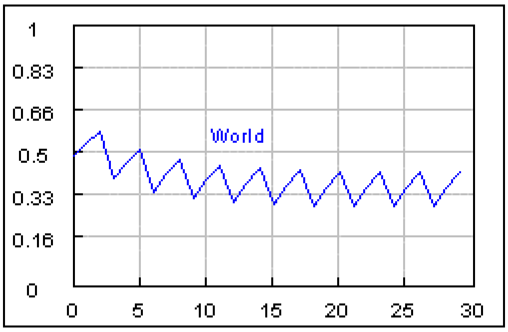

Figure 6. History of the world variable during interaction of the adaptive tit-for-tat agent with a CCD behavior (r_c=0.1, r_d=0.3, world(0)=0.5).

Thus, what we need at this point is a method for the adaptive tit-for-tat agent to discover whether the opponent uses a retaliating behavior or is just irrational and to adopt accordingly the proper rate setting. We have designed and examined several such variants for estimating the opponent's irrationality and we have finally found the following rule :

Throughout an observation window, record how many times (n) the agent's move has coincided with the opponent's move. At regular intervals (every "window" steps) adapt the rates as follows :

If (n>threshold) then

$r_c = r_{min}, r_d = r_{max}$

else $r_c = r_{max}, r_d = r_{min}$

The rule may be translated as :

If (the world is cooperative enough) then*

$r_c = r_{min}, r_d = r_{max}$

else $r_c = r_{max}, r_d = r_{min}$

(*) *recall that "my move = opponent's move" is the so-called pavlovian criterion of cooperation (*Nowak & Sigmund 1993*)*

Note that the agent drops its cooperation rate when the world is assumed cooperative, and increases it otherwise, that is, it uses negative feedback at the rate regulation level.

Another alternative interpretation for a cooperative world is a world that tries to manipulate the agent (so as to get it to respond with the same value found in the world). In this case it makes sense to drop the cooperation (potential manipulation) rate and become less adaptive to the world.

We have shown in simulations that the adaptive tit-for-tat agent with the meta-regulation mechanism converges to the proper behavior against both retaliating and irrational agents. Figures 7 and 8 give the behavior of the meta-regulated adaptive tit-for-tat agent against STFT and CCD.

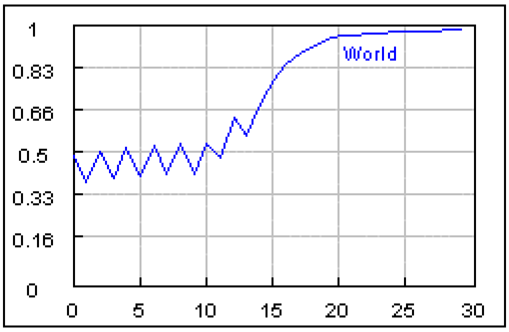

Figure 7. History of the world variable during interaction of the meta-regulated adaptive tit-for-tat agent with a suspicious tit-for-tat behavior ($r_c(0)$=0.2, $r_d(0)$=0.2, r_{max}=0.3, r_{min}=0.1, world(0)=0.5, window=10, threshold=2).

Note also how the agent manages to differentiate between a retaliating agent and an irrational one that has initially the same behavior. Figure 9 gives the behavior of the meta-regulated adaptive tit-for-tat agent against CDCD, that resembles STFT in the beginning. The agent first assumes that the opponent is retaliating and becomes increasingly cooperative, but soon finds out that the opponent is actually irrational and reverts to defection. Figure 10 gives the corresponding rates.

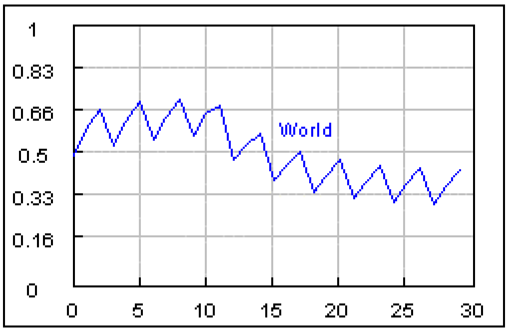

Figure 8. History of the world variable during interaction of the meta-regulated adaptive tit-for-tat agent with a CCD behavior ($r_c(0)$=0.2, $r_d(0)$=0.2, r_{max}=0.3, r_{min}=0.1, world(0)=0.5, window=10, threshold=2).

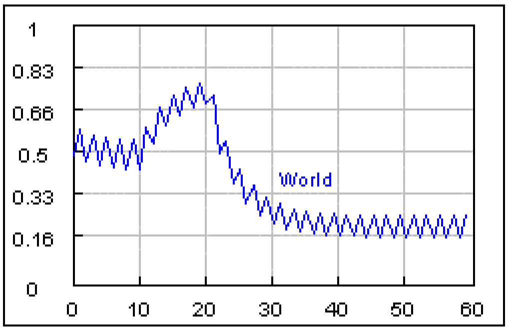

Figure 9. History of the world variable during interaction of the meta-regulated adaptive tit-for-tat agent with a CDCD behavior ($r_c(0)$=0.2, $r_d(0)$=0.2, r_{max}=0.3, r_{min}=0.1, world(0)=0.5, window=10, threshold=2).

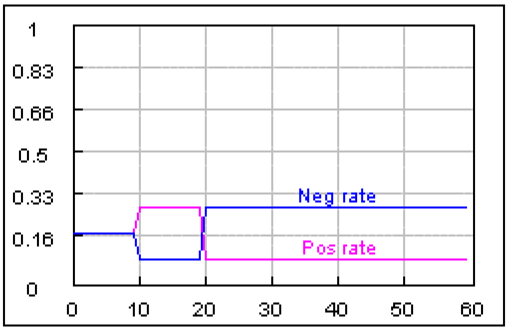

Figure 10. History of the rates for the experiment of the above figure.

The meta-regulated model is insensitive to the initial value of its world variable, provided that it is at least equal to 0.5 (remember that a tit-for-tat like behavior should start by cooperating). However, even for a defective initial value of the world variable, the adaptive agent may converge to cooperative behavior against the tit-for-tat agent. In this case, the resulting history is almost identical to the one of figure 7. The model is also insensitive to the exact values of r_{max} and r_{min}. Different values for the two rates will only result in scaling or stretching of the curves, the qualitative performance remaining intact. The same thing applies to the values of the observation window and the threshold, although they must be constrained so that the window will be sufficiently large and the threshold sufficiently small compared to the window. For example, a window of 5 with a threshold of 2 or a window of 10 with a threshold of 6 cannot get to discover periodic behaviors and defect to them. On the contrary, a window of 15 with a threshold of 3, yields the same results as a window of 10 with a threshold of 2, only stretched in the x-axis. Figure 11 gives the behavior of the meta-regulated adaptive tit-for-tat agent against CDCD for a window of 15 with a threshold of 3. The reader is invited to compare this figure to figure 9.

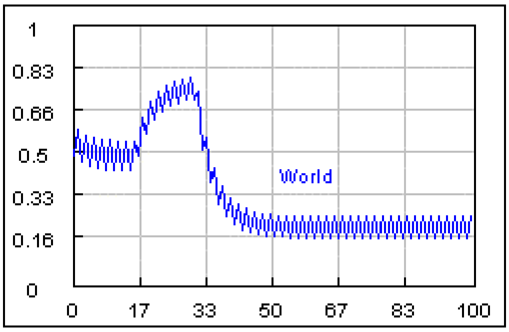

Figure 11. History of the world variable during interaction of the meta-regulated adaptive tit-for-tat agent with a CDCD behavior ($r_c(0)$=0.2, $r_d(0)$=0.2, r_{max}=0.3, r_{min}=0.1, world(0)=0.5, window=15, threshold=3).

We performed extensive tournament experiments, as well as a few initial ecological ones, with the behavior set used by (Beaufils et al. 1996) that contains all three kinds of behaviors, beside ADAPTIVE :

- **GRADUAL**
- **TFT**
- **MAJOR_C** Plays the opponent's most frequent move, cooperates in case of equality
- **GRIM** Cooperates until the opponent defects, then defects all the time
- **PROBER** Begins by playing CDD, then if the opponent has cooperated on the 2nd and 3rd move, defects all the time, otherwise it plays tit for tat
- **PAVLOV** Begins by cooperating, then cooperates if the two players made the same move (either C or D)
- **STFT**
- **ALLC** Always cooperate
- **CCD**
- **DDC**
- **ALLC** Always defect
- **LUNATIC** Cooperate or defect with equal probabilities

The tournament ranking of these behaviors is the following :

ADAPTIVE	39836
GRADUAL	39741
TFT	37395
MAJOR_C	36573
GRIM	35924
PAVLOV	34888
PROBER	32858
ALLC	31518
STFT	29881
LUNATIC	27816
CCD	27158
ALLD	26194
DDC	25995

The meta-regulated adaptive tit-for-tat agent's behavior has been verified against many other agents of the three types. In almost all cases it has been found to converge on the right behavior, either cooperation or defection, and this quickly enough so as to slightly outperform GRADUAL. Two notable exceptions are the OPPOSE behavior that returns the opposite of the opponent's last move (C on D and D on C) and sparse periodic behaviors (for example, CCCCCCCD). The case of the OPPOSE behavior is special because it involves the potential for extreme exploitation by the opponent. Actually, OPPOSE is a masochist behavior and it will be happy to play against a sadist one. Clearly, our adaptive behavior, being inherently "good", cannot demonstrate sadism, on the contrary it will try repeatedly to establish cooperation with the opponent. Since this cannot be achieved it will maintain oscillations forever. Note also that the fittest behavior against OPPOSE is ALLD, just as against ALLC! However, all good retaliating behaviors (including TFT, GRADUAL and ADAPTIVE) always cooperate against ALLC. Since there is no criterion for differentiating between ALLC and OPPOSE in the long term, a behavior has to demonstrate some degree of sadism (responding to C with D), in order to be able to exploit them. Similar things happen in the case of sparse periodic behaviors, because a behavior has to ignore large periods of C's and converge to the ALLD behavior.

These observations bring us to the issue of stability in front of perturbations, i.e., occasional defections. A stable agent should be able to demonstrate some generosity and forget all about it, provided that sufficient cooperative evidence exists or is given afterwards.

The importance of resistance to perturbations may be justified as follows : if an agent's opponent is replaced by a new one with a different behavior, the agent should be able to react to the new opponent and minimize the effects of past experience on its behavior. In everyday terms, often enough our opponent in a game changes (for example, the clerk that serves on us in the bank). In such cases, we should be able to find the proper behavior against our opponent without being tied too much to our previous idea about him, and, if necessary, to altogether forget the past. This may also be applied to a single opponent who, for some reason, changes radically its behavior. In other words, true cooperative behavior should be not only responsive but adaptive as well. This is clearly not the case of the GRADUAL behavior, because it becomes increasingly slow in forgiving, or, equivalently, increasingly nasty.

A demonstration is the following : let us suppose that a GRADUAL agent interacts with an ALLD agent (that always defects) for 15 steps before the ALLD agent is replaced by another newborn GRADUAL agent. A 1000-cycle interaction between the two results to a score of 1051 for each one of them, while the first few steps are given next :

GRADUAL CDCCD DDDCC DDDDD DDDDD CCDDD …

OPPONENT DDDDD DDDDD DDDDD CDCCD DDDCC …

This transcript shows that both agents become increasingly defecting, hence obtaining lower scores, while we would like them to be able to converge to cooperative behavior. The same thing applies to the adaptive tft behavior, i.e., it cannot always manage to converge to cooperation with a fellow adaptive tft agent if the latter starts defectively. This is so because, rather surprisingly, the defection rate is higher than the cooperation rate (cf. figure 10), so that the agent will need plenty of time before becoming cooperative against its opponent. Some opponents (including newborn adaptive ones) will not be patient enough to wait and will revert to defection, too. Thus, the adaptive tit-for-tat model is initially convergent to the proper behavior but is still unstable.

The above observations suggest a ***rationality criterion*** :

A behavior is rational if it cooperates with itself.

Of course, this criterion is fulfilled for "good" retaliating behaviors (i.e., behaviors that start by cooperating, for example, tit for tat, gradual, grim, pavlov, adaptive etc.) if the interaction starts from the beginning, without past experience. This is however generally not the case if the interaction starts from a point where one of the agents is in defecting mode. For example, when tit for tat plays against suspicious tit for tat, they alternate C and D : stft plays DCDCDC… and tft plays CDCDCD…). A soft rationality criterion may thus be formulated as :

A behavior is rational if it ends up cooperating with an initially defecting version of itself.

A hard rationality criterion would be :

A behavior is rational if it ends up cooperating with itself, even if they both start from defecting initial conditions.

Fulfillment of this criterion would also allow the behavior to resist to local perturbations, other than opponent replacement. For example, in everyday settings it is often the case that an agent's move is perturbed by environmental conditions or just misinterpreted by its opponent. A rational agent should be able to overcome this local trouble and re-converge to cooperation.

According to the above, our adaptive tit-for-tat agent may be thought of as only half-adaptive, because it has reversible memory and can apparently adapt to arbitrary new environments, but is still unstable to perturbations. A full adaptive model would be rational as well. We have no way to know whether such a full adaptive model exists, and if so, whether this model constitutes an Eldorado, i.e. a universal solution to all our cooperation problems. In any case, it appears that the true cooperation problem is essentially a problem of adaptation and stability against perturbations and may be studied "in the empty", i.e. with an agent against itself under varying initial conditions for all parameters involved. An agent should therefore play against an opponent as it would play against itself.

The difficulty of the stability problem described above is an indication that the formulation of the cooperation problem is perhaps too awkward to capture naturally some parameters involved intuitively in the process, such as tolerance. A continuous game model (like the one studied in (Tzafestas 1995a)) would be more allowable to such phenomena.

4. Theoretical discussion

We have shown above that the agent's behavior is based on a critical continuous variable (the world variable) that drives its motivation to cooperate or to defect. This variable represents for the agent an estimate of the friendliness or hostility of the world; it has therefore *cognitive value*. By regulating its own variable, the agent tries to get the world to a stable value, preferably C, otherwise D.

The good performance of the behavior is ensured through an additional self-regulation mechanism acting on the adaptation rates. This is an important observation, since it is compatible with the dynamical approach to cognition (van Gelder and Port 1995), stating that the most important factor in cognitive mechanisms is the nature of dynamics involved. Note also that there is no need to have a rate dynamics other than the "bang-bang" dynamics (high-low value), because what counts is the relation between r_c and r_d, rather than their absolute values.

The double regulation loop implies a different point of view on the problem. While it has been traditionally tackled as a score maximization problem, in our work we are proposing the inverse point of view. The agent may be regarded as trying to regulate within bounds some internal variables. The regulated variables appear to be critical for an agent's survival or operationality, so that Ashby (1960) called them *essential variables*. The buildup and reinforcement of the world estimate in stable environments is a by-product of agent self-regulation when a perturbation occurs. The driving force of the agent's behavior is thus the state of its essential variables, whereas the final generation of cooperative or defective moves constitutes the metabolic part of the overall mechanism.

It is noteworthy that exactly the same qualitative conclusions have been drawn in the cases of agents exploring an environment with more or less uniform distribution of sources (Tzafestas 1995b, chapters 4 and 5) and of ant agents building and reinforcing trails to a food source (Tzafestas 1998). In both cases, of course, different cognitive variables and different types of first-level adaptation are required.

5. Conclusions and perspectives

We have investigated the classical IPD problem and studied the behavior and performance of the best behavior found so far, GRADUAL. This behavior has permanent memory and thus it does not constitute a good basis for a truly perturbation-resistent adaptive behavior. What is necessary is a regulation mechanism ensuring that the agent takes the opponent's behavior into account from a certain distance without committing itself to it. On top of a basic regulation loop that estimates the opponent's friendliness or hostility, a second regulation loop is introduced that acts on the rates of the first one. The meta-regulation mechanism allows for the agent to achieve performances comparable to that of GRADUAL. Theoretically, the overall model relies on the definition of a *cognitive variable* for the agent that is adapted throughout the interaction with the world. The adaptation rates that define the dynamics of the system follow a bang-bang regulation dynamics between two limits and this constitutes the meta-regulation loop. Overall, the agents may be regarded as self-regulating some internal "essential" variables, with the by-product being the final cooperative or defective behavior.

As far as regulation is concerned, our approach has been already validated in the past for the exploration problem in a uniform source distribution and for the trail building problem of ant agents and the same principles have been found to apply. The next step is to formulate and solve in the same way a few other classical artificial life problems, such as robot cooperation in a closed ecosystem (Steels 1994) and action selection (Tyrrell 1993,1994). We hope that the comparative study of the results and conclusions for each of these problems will teach us a few lessons on regulation.

From the pure cooperation point of view, we are currently studying the rationality of the adaptive tit-for-tat behavior, in the quest of the full adaptive or even the Eldorado behavior. We also plan to conduct experiments with different payoff matrices, in an effort to generalize the model. Finally, we are preparing experiments on continuous games.

References

Ashby, W.R. (1960) *Design for a brain - The origin of adaptive behaviour.* 2nd revised edition, Chapman & Hall.

Axelrod, R., and Hamilton, W.D. (1981) The evolution of cooperation. *Science* 211:1390-96.

Axelrod, R. 1984. *The evolution of cooperation.* Basic Books.

Axelrod, R., and Dion, D. (1988) The further evolution of cooperation. *Science* 242:1385-90.

Beaufils, B., Delahaye, J.-P., and Mathieu, P. (1996) Our meeting with gradual: A good strategy for the iterated prisoner's dilemma. *Proceedings of the Artificial Life V Conference*, Nara, Japan.

Feldman, M.W., and Thomas, E.A.C. (1987) Behavior-dependent contexts for repeated plays of the prisoner's dilemma II: Dynamical aspects of the evolution of cooperation. *Journal of Theoretical Biology* 128:297-315.

Fogel, D. (1993) Evolving behaviors in the iterated prisoner's dilemma. *Evolutionary Computation* 1:77-97.

Milinski, M. (1987) Tit for tat in sticklebacks and the evolution of cooperation. *Nature* 325:433-435.

Nowak, M., and Sigmund, K. (1992) Tit for tat in heterogeneous populations. *Nature* 355:250-53.

Nowak, M., and Sigmund, K. (1993) A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner's dilemma game. *Nature* 364:56-58.

Steels, L. (1994) A case study in the behavior-oriented design of autonomous agents. *Proceedings Simulation of Adaptive Behavior 1994*, 445-452.

Stanley, E.A., Ashlock, D., and Tesfatsion, L. (1994) Iterated Prisoner's Dilemma with Choice and Refusal of Partners. *Artificial Life III*, Addison-Wesley.

Tyrrell, T. (1993) The use of hierarchies for action selection. *Adaptive Behavior* 1:387-420.

Tyrrell, T. (1994) An evaluation of Maes' bottom-up mechanism for behavior selection. *Adaptive Behavior* 2:307-348.

Tzafestas, E. (1995a) Beyond cooperation and competition : Explorations with a quantitative tit-for-tat model. *LAFORIA Research Report 95/34.*

Tzafestas, E. (1995b) *Vers une systémique des agents autonomes : Des cellules, des motivations et des perturbations.* Ph.D. diss., LAFORIA-IBP, Univ. Pierre et Marie Curie, Paris.

Tzafestas, E. (1998) Tom Thumb Robots Revisited : Self-Regulation as the Basis of Behavior. *Proceedings of the Artificial Life VI Conference,* San Diego, CA.

van Gelder, T., and Port, R. (1995) It's about time : An overview of the dynamical approach to cognition. *Mind as Motion : Explorations in the dynamics of cognition,* by T. van Gelder and R. Port, Cambridge, Mass.: MIT Press.