

# A Tit for Tat Negotiation Strategy for Real-Time Bilateral Negotiations

Tim Baarslag, Koen Hindriks, and Catholijn Jonker

**Abstract.** We describe the strategy of our negotiating agent, Nice Tit for Tat Agent, which reached the finals of the 2011 Automated Negotiating Agent Competition. It uses a Tit for Tat strategy to select its offers in a negotiation, i.e.: initially it cooperates with its opponent, and in the following rounds of negotiation, it responds in kind to the opponent's actions. We give an overview of how to implement such a Tit for Tat strategy and discuss its merits in the setting of closed bilateral negotiation.

## 1 Introduction

This paper presents a new negotiation strategy, called the Nice Tit for Tat Agent, which we developed and entered into the Second Automated Negotiating Agent Competition<sup>1</sup> (ANAC2011). ANAC is a tournament between a set of negotiating agents which perform closed bilateral negotiation using the alternating offers protocol. The negotiation environment consists of multi-issue scenarios, and is closed in the sense that there is uncertainty about the opponent's preferences.

Our negotiation strategy is based on the principle of Tit for Tat: cooperating on the first move and then mirroring whatever the other player did in the preceding round. Thus, Tit for Tat is a strategy of cooperation based on reciprocity [1].

Tit for Tat has been applied and found successful in many other games, including the Iterated Prisoner's Dilemma game. It is considered to be a very robust strategy, mainly because of the following three features:

- i. It is never the first to defect (i.e., it plays nice as long as the the opponent plays nice as well);

---

Tim Baarslag · Koen V. Hindriks · Catholijn M. Jonker  
Delft University of Technology, The Netherlands  
e-mail: {T.Baarslag, K.V.Hindriks, C.M.Jonker}@tudelft.nl

<sup>1</sup> <http://mmi.tudelft.nl/anac>

- ii. It can be provoked into retaliation by a defection of the opponent;
- iii. However, it is forgiving after just one act of retaliation.

In this paper, we discuss how to implement such a Tit for Tat strategy in the setting of closed bilateral negotiation.

## 2 A Tit for Tat Negotiation Strategy

We first give a broad overview of the Tit for Tat negotiation strategy. The ideas behind and details of implementation are given below.

The Nice Tit for Tat agent plays a tit for tat strategy with respect to its own utility. The agent will initially cooperate, then respond in kind to the opponent's previous action, while aiming for the Nash point of the negotiation scenario. After each move by the opponent, it updates its Bayesian opponent model to make sure it responds with a nice move to a concession by the opponent.

### 2.1 *Four Ways of Repaying the Favor*

When implementing a Tit for Tat strategy, the first question that needs answering is: in what way should an agent reciprocate? In closed bilateral negotiation, there are two utility functions in play, one of which is unknown to the other agent. Therefore, there are two different actions from the opponent that could be considered 'nice':

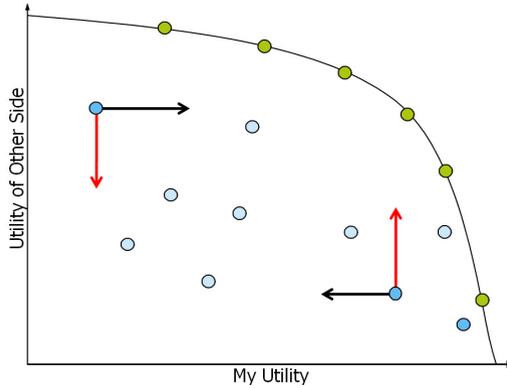
- i. The opponent concedes according to its own utility function;
- ii. The opponent offers more utility to the agent, according to the agent's utility function.

Both actions can be reciprocated by the agent by again choosing one of the two options. Therefore, there are four different ways for a Tit for Tat agent to reciprocate (see Fig. 1).

For the Tit For Tat agent, we have elected to reciprocate according to the agent's own utility function. That is, when the opponent makes a bid that is better for the agent, in return, the agent will produce a bid of lower utility for itself. There is a good reason to do so: the agent's own utility function is known to the agent. The other three ways to reciprocate depend on the utility information provided by an opponent model, which is inherently unreliable.

### 2.2 *Reciprocating and Making Nice Moves*

The basic idea of the Tit for Tat strategy is to reciprocate in terms of the agent's own utility as described in Sec 2.1, but it also takes into account the utility of its opponent. To do so, it constructs an opponent model using Bayesian learning [4, 5].



**Fig. 1** Four different ways to reciprocate

Using this opponent model, a naive implementation of the Tit for Tat strategy would go as follows:

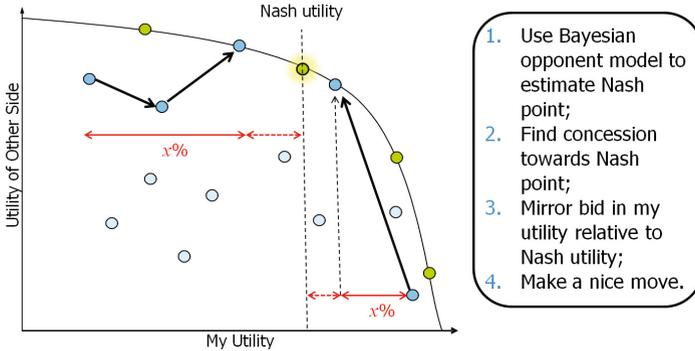
- i. Measure the opponent’s concession in terms of the agent’s own utility function;
- ii. Mirror this bid as described in Sec 2.1, sacrificing the same amount as is offered by the opponent;
- iii. Make the offer as attractive as possible for the opponent using the Bayesian opponent model.

However, this implementation would lead to an agent strategy that is far too nice: if the opponent yields  $\frac{1}{2}$  utility to the agent, then it would respond in kind by sacrificing the same amount. In other words: the agent would be satisfied with deals of  $\frac{1}{2}$  utility. Given that many domains of ANAC have win–win outcomes with utilities much higher than that, this is clearly a suboptimal approach.

Therefore, in our version of the agent, it generally aims for more than half, depending on the negotiation scenario. Instead, it uses the opponent model to make an estimate of the location of the *Nash point* of the negotiation scenario, and then aims for this outcome (see Fig. 2). For example: if the opponent has made an offer that is 70% on the way to the Nash point, the agent will respond in kind by approaching from the other side, making an offer that is 30% away from the Nash point.

### 2.3 Acceptance Strategy

The preceding sections focused on the bidding strategy of the Tit for Tat agent. However, there is a second component of a negotiator’s strategy that can also highly influence the outcome of a negotiation, namely its acceptance strategy. In [2], several acceptance conditions are defined that are designed to perform well in conjunction



**Fig. 2** The four steps taken by the Tit for Tat bidding strategy to determine the next offer

with an arbitrary bidding strategy. The Nice Tit for Tat agent is equipped with a particular type of acceptance condition called  $AC_{combi}$ , which the paper shows to work better than the majority of more simple generic conditions.

The basic idea behind  $AC_{combi}$  is as follows: in case the bidding strategy plans to propose a deal that is worse than the opponent's offer, it has reached a consensus with the opponent and thus accepts the offer. However, if there still exists a gap between the two offers and time is short,  $AC_{combi}$  waits for an offer that is not expected to improve in the remaining time.

### 3 Conclusion and Future Work

In this paper, we have provided an overview of the strategy of the Nice Tit for Tat Agent, which has participated in the finals of the ANAC2011 competition.

We have designed a new negotiation strategy based on the well-known Tit for Tat principle. We have covered both the procedure that our agent follows to select a bid that reciprocates the play by its opponent, and the way it chooses to accept a certain negotiation outcome.

As has been found in post-tournament analysis of the ANAC results [3], the Tit for Tat agent was the only agent in the tournament to match the behavior of the opponent (which is to be expected). This means it plays tough against hardheaded negotiators, but it also means it plays nice against strategies that concede easily. It is observed in the paper that this approach might not be as successful in negotiation as in some other games, because it does not exploit the conceding strategies enough to reach the top rankings. More research is required to find this delicate balance between cooperative and competitive behavior of a negotiating agent.

## References

1. Axelrod, R.: *The Evolution of Cooperation*. Basic Books (1984)
2. Baarslag, T., Hindriks, K., Jonker, C.: Acceptance conditions in automated negotiation. In: *Proceedings of The Fourth International Workshop on Agent-based Complex Automated Negotiations, ACAN 2011* (2011)
3. Baarslag, T., Hindriks, K., Jonker, C.: Towards a Quantitative Concession-Based Classification Method of Negotiation Strategies. In: Kinny, D., Hsu, J.Y.-J., Governatori, G., Ghose, A.K. (eds.) *PRIMA 2011. LNCS*, vol. 7047, pp. 143–158. Springer, Heidelberg (2011)
4. Hindriks, K.V., Tykhonov, D.: Opponent modelling in automated multi-issue negotiation using bayesian learning (2008)
5. Zeng, D., Sycara, K.: Bayesian learning in negotiation. *International Journal of Human Computer Systems* 48, 125–141 (1998)