The Function of Time in Cooperative Negotiations*

Sarit Kraus
Graduate School of Library Studies and
Dept. of Computer Science
Hebrew University, Jerusalem, 91904 Israel
e-mail: sarit@cs.huji.ac.il

Jonathan Wilkenfeld
Dept. of Government and Politics
University of Maryland
College Park, MD 20742
e-mail:wilkenfeld@umd2.umd.edu

Abstract

Work in distributed artificial intelligence (DAI) has, since its earliest years, been concerned with negotiation strategies. which can be used in building agents that are able to communicate to reach mutually beneficial agreements. In this paper we suggest a strategic model of negotiation that takes the passage of time during the negotiation process itself into consideration. Changes in the agent's preferences over time will change their strategies in the negotiation and, as a result the agreements they are willing to reach. We will show that in this model the delay in reaching agreements can be avoided.

Introduction

Research in distributed artificial intelligence (DAI) is concerned with how automated agents can be designed to interact effectively. One important capability that could aid inter-agent cooperation is negotiation; agents could be built that are able to communicate their respective desires and compromise to reach such mutually beneficial agreements.

Work in DAI has been concerned with negotiation strategies [Davis and Smith 1983; Georgeff 1983; Malone *et al.* 1988; Durfee 1988; Rosenschein and Genesereth 1985; Sathi and Fox 1989; Conry *et al.* 1988; Zlotkin and Rosenschein 1990] which can be used in building agents that are able to communicate to reach mutually beneficial agreements. Sycara ([Sycara 1987]), using case-based reasoning, and Kraus et al. ([Kraus *et al.* 1991]) modeled negotiations from a cognitive standpoint.

One of the main criticisms of using negotiations as a way of reaching mutual benefit is that negotiation is a costly and time-consuming process and, consequently, it may increase the overhead of coordination (see [Bond and Gasser 1988]). In the presence of time constraints, planning and negotiation time should be taken into consideration The negotiation may be either about job sharing or resource allocation. In both cases we want to prevent the agents from spending too much time on negotiation and therefore not keeping to their timetable for satisfying their goals.

In this paper we suggest a strategic model of negotiation that takes the passage of time during the negotiation process itself into consideration. Changes in the agent's preferences over time will change their strategies in the negotiation and, as a result the agreements they are willing to reach. We will show that the delay in reaching agreements can be avoided.

Following [Rosenschein and Genesereth 1985; Zlotkin and Rosenschein 1990; Kraus and Wilkenfeld 1990a] we examine negotiation using game theoretic techniques with appropriate modifications to fit artificial intelligence situations. We will focus primarily on works in game theory and economics that have studied the effect of time preferences on the negotiation process, following the classic paper by Rubinstein ([Rubinstein 1982]). Comparing our work to that of Zlotkin and Rosenchein, [Zlotkin and Rosenschein 1990] we make no assumption about the protocol the agents use for negotiations. Also, our model takes the passage of time during the negotiation process itself into consideration, which in turn influences the outcome of the negotiations and avoids delays in reaching an agreement.

Initial Setting

Two autonomous agents A and B have a common goal they want to satisfy as soon as possible. In order to satisfy any goal, costly actions must be taken and an agent cannot satisfy the goal without reaching an agreement with the other agent. Each of the agents wants to minimize its costs, i.e., prefers to do as little as possible. We note that even though the agents have the same goal (under our simplified assumptions), there is actually a conflict of interests. The agents try to reach an agreement over the division of labor. We assume that each step in the negotiation takes time, and the agents have preferences for reaching agreements in different

*A preliminary version of this paper was presented at the 10th International Workshop on DAI. This work was partially completed while the first author was at the Institute for Advanced Computer Studies and Dept. of Computer Science, University of Maryland, College Park.

time periods.

We make the following assumptions:

1. Full information - each agent knows all relevant information including the other agent's preferences for the different outcomes over time.
2. The agents are rational - they will behave according to their preferences.
3. Commitments are enforced - if an agreement is reached both sides are forced to follow it.
4. Assumptions (1)-(3) are common knowledge.

We will also assume that there are no other agents in the environment which can help them and that any division of the work that is needed to satisfy the goal is possible. We will relax the last assumptions later. In [Kraus and Wilkenfeld 1991] we examine the case of N agents, where $N \geq 3$.

Example 1 *Two agents must paint a wall. One has the brush and the other has the paint. In order to paint the wall both the paint and the brush are needed. Painting is a costly operation, and any delay in the painting is also costly to both sides.*

The Structure of Negotiations

Our strategic model of negotiations is a model of Alternative Offers.[1] We utilize modified definitions from [Osborne and Rubinstein 1990].

Definition 1 Agreement:
An agreement is a pair (s_A, s_B), in which s_i is agent i's portion of the work[2] needed to satisfy the agents' goal. The set of possible agreements is

$$S = \{(s_A, s_B) \in \mathcal{R}^2 : s_A + s_B = 1 \text{ and } s_i \geq 0, \text{ for } i = A, B\}$$

The agents' preferences over S are opposed. Each agent prefers to do less rather than more. That is, agent A prefers $s^1 \in S$ to $s^2 \in S$ if and only if $s_i^1 < s_i^2$.

The negotiation procedure is as follows. The agents can take actions only at certain times in the set $\mathcal{T} = \{0, 1, 2 ...\}$. In each period $t \in \mathcal{T}$ one agent, say i, proposes an agreement, and the other agent (j) either accepts the offer (Y) or rejects it (N). If the offer is accepted, then the negotiation ends, and the agreement is implemented (i.e. each of the agents does its part of the job). After a rejection, the rejecting agent then has to make a counter offer and so on. There are no rules which bind the agents to any previous offers and there is no limit on the number of periods.

Definition 2 Negotiation Strategies:
Let F be the set of all sequences of functions $f = \{f^t\}_{t=0}^{\infty}$, where $f^0 \in S$, for t even $f^t : S^t \rightarrow S$, and for t odd $f^t : S^{t+1} \rightarrow \{Y, N\}$ (S^t is the set of all sequences of length t of elements in S and Y, N are defined above). F is the set of all strategies of the agent who starts the bargaining. Similarly, let G be the set of all strategies of the agent who, in the first move, has to respond to the other agent's offer; that is, G is the set of all sequences of functions $g = \{g^t\}_{t=0}^{\infty}$ such that for t even $g^t : S^{t+1} \rightarrow \{Y, N\}$ and for t odd $g^t : S^t \rightarrow S$.

Let $\sigma(f, g)$ be a sequence of offers in which A starts the bargaining and adopts $f \in F$, and B adopts $g \in G$. Let $L(f, g)$ be the length of $\sigma(f, g)$ (where the length may be infinite). Let $La(f, g)$ be the last element of $\sigma(f, g)$ (if there is such an element). We present a formal definition for the outcome of the negotiation.

Definition 3 Outcome of the negotiation:
The outcome function of the negotiation is defined by

$$\mathcal{P}(f, g) = \begin{cases} D & \text{if } L(f, g) = \infty \\ (La(f, g), L(f, g) - 1) & \text{otherwise} \end{cases}$$

Thus, the outcome (s, t) where $s \in S$ is interpreted as the reaching of agreement s in period t and the symbol D indicates a perpetual disagreement.

We note here that by defining an outcome to be either a pair (s, t) or D, we have made a restrictive assumption about the agents' preferences. We assume that agents care only about the nature of the agreement, and the time at which the outcome is reached, and not about the sequence of offers and counteroffers that leads to the agreement, i.e., no "decision-regret" (see [Raiffa 1982]).

The last component of the model is the preference of the agents on the set of outcomes. Each agent has preferences over agreements reached at various points in time.[3] The *time preferences* and the preferences between agreements are the driving force of the model.

Formally, we assume that agent $i = A, B$ has a preference relation $\succeq_i$ on the set $\{S \times \mathcal{T}\} \cup \{D\}$.

Here we will concentrate on two utility functions that yield preference relations under the assumption that an agent wants to maximize its own utilities.

Definition 4 Utility function with time constant discount rate
Let $(s, t) \in S \times \mathcal{T}$ be an outcome of the negotiation, then the $Utility_i\{(s, t)\}$ where $i = A, B$ is defined to be $\delta_i^t(1 - s_i)$, where $0 < \delta_i < 1$, and $Utility_i\{D\} = -\infty$.

Definition 5 Utility function with a constant cost of delay
Let $(s, t) \in S \times \mathcal{T}$ be an outcome of the negotiation then the $Utility_i'\{(s, t)\}$ where $i = A, B$ is defined to be $1 - s_i - c_i t$ where $c_i > 0$ and $Utility_i'\{D\} = -\infty$.

Both of the above utility functions capture the gains of the agents as the difference between the work they

[1] See [Osborne and Rubinstein 1990] for a detailed review of the bargaining game of Alternating Offers.

[2] A similar definition can be given concerning a division of resources.

[3] The way these preferences are determined by the agents is beyond the scope of this paper. They can either be given to the agent with the specifications of its task, or it may be decided by the agent itself after analyzing his goals

agreed to perform and the whole work that is needed, and have a discount over time.

How will a rational agent choose his strategy for the negotiation? A useful notion is the Nash Equlibrium ([Nash 1953; Luce and Raiffa 1957]). A pair of strategies (σ, τ) is a Nash Equilibrium if, given τ, no strategy of A results in an outcome that A prefers to the outcome generated by (σ, τ) and, similarly, to B given σ. If there is a unique equilibrium, and if it is known that an agent is designed to use this strategy, no agent will prefer to use a strategy other than these ones.

However, the use of Nash Equilibrium is not an effective way of analyzing the outcomes of the models of Alternating Offers since it puts few restrictions on the outcome and yields too many equilibria points. (See ([Rubinstein 1982] for the proof). Therefore, we will use the stronger notion of **(subgame) perfect equilibrium** (P.E.) (see [Rubinstein 1982]) which requires that the agents' strategies induce an equilibrium at any stage of the negotiation, i.e., in each stage of the negotiation, assuming that an agent follows the P.E. strategy, the other agent does not have a better strategy than to follow its own P.E. strategy. So, if there is a unique perfect equilibrium, and if it is known that an agent is designed to use this strategy, no agent will prefer to use a strategy other than this one in each stage of the negotiations. The following theorem shows that there exists a unique P.E. which ends the negotiation after the first period. This unique solution is characterized by a pair of agreements x^* and y^*, that satisfies: (1) Agent A is indifferent between "y^* today" and "x^* tomorrow," and (2) Agent B is indifferent between "x^* today" and "y^* tomorrow." When a unique pair of x^* and y^* satisfies this statement, there exists a unique P.E. [Rubinstein 1982].

Theorem 1 *Suppose agent A starts the negotiations. Let*

$$x^* = \left(\frac{\delta_B(1-\delta_A)}{1-\delta_A\delta_B}, \frac{1-\delta_B}{1-\delta_A\delta_B}\right) \quad y^* = \left(\frac{1-\delta_A}{1-\delta_A\delta_B}, \frac{\delta_A(1-\delta_B)}{1-\delta_A\delta_B}\right)$$

$(\widehat{f}, \widehat{g})$ is a subgame perfect equilibrium of the strategic model of Alternative Offers where the agents' utility function is defined in definition 4 iff
$\widehat{f}^t(s^0, ..., s^{t-1}) = x^*$ for all $(s^0, ..., s^{t-1}) \in S^t$, if t is even, and

$$\widehat{f}^t(s^0, ..., s^t) = \begin{cases} Y & if s_1^t \leq y_1^* \\ N & if s_1^t > y_1^* \end{cases}$$

if t is odd. The strategy $\widehat{g}$ of agent B has the same structure; the roles of x^* and y^* are reversed, the words "odd" and "even" are interchanged, and each subscript A is replaced by B. The outcome is that A proposes x^* in period 0 and B immediately accepts this offer.

Proof: Similar to the proof in [Osborne and Rubinstein 1990] with some small modifications. ∎

Even though the structure of the strategic model of Alternating Offers allows negotiation to continue indefinitely, in the unique subgame perfect equilibrium it terminates immediately. Similar results can be obtained using the utility function of constant delay defined in Definition 5, when $c_A \neq c_B$. These results demonstrate our idea that allowing time into the negotiation process can lead to an efficient negotiation.

Other elements that influence the outcome of the negotiation include the patience of the agents. If A_1's losses over time are greater than A_2's then he will do more of the work. In addition, the agent who starts the negotiation has an advantage over the other agent (for example if both agents have the same rate of delay δ then the first one will do only $\delta/(1+\delta)$ of the job and the other will do $1/(1+\delta)$). A simple way to avoid this asymmetry in the model is the following (see [Osborne and Rubinstein 1990]): at the beginning of each period each agent is chosen with probability 1/2 (independently across periods) to be the one to make the first offer. We return now to the agents from Example 1, and demonstrate the above results.

Example 2 *We denote the agent with the paint by P and the agent with the brush by B. Let an agreement be a pair (s_P, s_B) where s_i is the agreed portion of the wall agent i will paint and $s_P + s_B = 1$.*

Suppose agent B and agent P have the following utility function correspondingly: $U_B\{(s,t)\} = 0.9^t(1-s_B)$ and $U_P\{(s,t)\} = 0.8^t(1-s_P)$. If agent B is designed to use the unique P.E. then agent P should use it too. Suppose agent P starts the negotiation. He will offer $((0.8 * 0.1)/(1-0.9*0.8), 0.2/(1-0.9*0.8))$ which is approximately $(0.286, 0.714)$ and agent B will accept the offer immediately.

Finite Set of Feasible Agreements

Until now we have assumed that the agents can divide the work between them, in any way that they have agreed upon. Unfortunately, this cannot usually be done. If two agents need to carry blocks or to deliver packages, this work can be divided only in a discrete manner and usually in a finite number of possible agreements.

The strategic model of Alternative Offers is useful here also, if the preferences of the agents satisfy similar requirements as in the continuity case. An additional requirement is that the the length of a single period is fixed[4].

Here, we will demonstrate this case using the following example. Suppose there are only three possible agreements the agents can reach $\{a, b, c\}$. Let $\preceq_i$ $i = A, B$ denote the preferences of the agents over the possible outcomes which satisfy the following assumptions:

Disagreement is the worst outcome: For every n$t \in \mathcal{T}$ and $s \in \{a, b, c\}$, $(s,t) \succ_i D$ where $i \in \{A, B\}$

Conflict of Interests: $(a,0) \succ_A (b,0) \succ_A (c,0)$, $(c,0) \succ_B (b,0) \succ_B (a,0)$.

[4]The discrete case is not usually discussed in the game theory community, but see [Muthoo 1989]

Monotonicity in Time: time is valuable to both sides, i.e., for every $t_1, t_2 \in \mathcal{T}$, $i \in \{A, B\}$, and $s \in \{a, b, c\}$, if $t_1 < t_2$ $(s, t_1) \succ_i (s, t_2)$.

Stationarity: preferences between (s_1, t_1) and (s_2, t_2) depend only on s_1, s_2 and the differences between t_1 and t_2, i.e. for every $s_1, s_2 \in \{a, b, c\}$, $t_1, t_2, t_3 \in \mathcal{T}$ $i \in \{A, B\}$ $(s_1, t_1) \succ_i (s_2, t_1)$ iff $(s_1, 0) \succ_i (s_2, 0)$ and if $(s_1, t_1) \succeq_i (s_2, t_1 + t_2)$ then $(s_1, t_3) \succeq_i (s_1, t_3 + t_2)$.

In addition $(a, 1) \sim_A (b, 0)$[5] and $(b, 1) \sim_A (c, 0)$. $(a, 0) \sim_B (b, 1), (b, 0) \sim_B (c, 2)$.

Assuming that agent A starts the negotiations, the strategies from Theorem 1, $(\widehat{f}, \widehat{g})$ where $x^* = a$ and $y^* = b$, is a perfect equilibrium. We may also conclude that when the domain is discrete and finite, and only a finite number of possible agreements are feasible, delay in reaching cooperation may be avoided by incorporating time into the model and by using the notion of subgame Perfect Equilibrium.

The agents can opt out

Until now we assumed that the agents must continue the negotiation since disagreement was the worst outcome to both sides. Let us consider the case in which one of the agents has the ability to opt out of the negotiation. For example, suppose agent A has another goal he may satisfy (usually with lower priority); he can benefit from doing a different job than painting the wall; or, in the case in which agents negotiate in order to allocate resources and the agents can reach an agreement with another agent for another type of resource allocation. The threat of leaving the negotiation may influence the outcome in some cases.

We assume that there are W units of the work that should be divided by two agents. The set of possible agreements, S, includes all the pairs $(s_A, s_B) \in \mathcal{N}^2$ where $s_A + s_B = W$. We modify the negotiation strategies (Definition 2) such that if agent i receives an offer from his partner he can opt out of the negotiation (O), in addition to accepting the offer (Y) or rejecting it (N). $\sigma(f, g)$, $L(f, g)$, $La(f, g)$ and the outcome of the negotiations are defined as in the previous sections, but $La(f, g)$, which is the last element of $\sigma(f, g)$, can be either $s \in S$ or O. Thus the outcome (O, t) is interpreted as one of the agents opting out of the negotiation at period t. We note that the length of the time periods is fixed. The agents' preferences in this case are over agreements reached at various points in time, and over opting out at various points in time.

The conditions on the preference relations of the agents are similar to those of the previous section. First we assume that the least-preferred outcome is disagreement (D).

(A0) For every $s \in S$ and $t \in \mathcal{T}$, $(s, t) \succ_i D$ and $(O, t) \succ_i D$ (Disagreement is the worst outcome).

[5] $(s_1, t_1) \sim_i (s_2, t_2)$, $i \in \{A, B\}$ indicates that agent i is indifferent between the two outcomes.

The next two conditions (A1), (A2) concern the behavior of $\succ_i$ on $S \times \mathcal{T}$, i.e. concerning agreements reached in different time periods. Condition (A1) requires that among agreements reached in the same period, agent i prefers smaller numbers of units s_i.

(A1) if $r_i < s_i$, then $(r, t) \succ_i (s, t)$.

The next assumption greatly simplifies the structure of preferences among agreements. It requires that preferences between (s_1, t_1) and (s_2, t_2) depend only on s_1, s_2 and the differences between t_1 and t_2.

(A2) For all $r, s \in S$, $t, t_1, t_2, \delta \in N$ and $i \in \{A, B\}$, $(r, t_1) \succeq_i (s, t_1 + \delta)$ iff $(r, t_2) \succeq_i (s, t_2 + \delta)$ (Stationarity). We note that assumption (A2) does not hold for O.

We will consider the case in which any agent has a number $c_i > 0$ $i \in \{A, B\}$ such that:

(A3) $(s, t_1) \succeq_i (\bar{s}, t_2)$ iff $s_i + c_i * t_1 \leq \bar{s}_i + c_i * t_2$.

We also assume that both agents prefer to opt out sooner rather than later. Formally:

(A4) If $t_1 < t_2$ then $(O, t_1) \succ_i (O, t_2)$, $i \in \{A, B\}$.

We do not make any assumption concerning the preferences of an agent for opting out versus an agreement. This enables us to consider different types of cases of opting out. For example, in the "wall painting" case, opting out may be giving up the goal, buying a brush or covering the wall with wallpaper. Formally, there is no fixed $s \in S$ such that for every $t \in \mathcal{T}$, $(s, t) \sim (O, t)$ as in [Shaked and Sutton 1984].

Let us define the "outcome" of opting out as follows:

Definition 6 *For every $t \in \mathcal{T}$ and $i \in \{A, B\}$, let $Pos_i^t = \{s^t \mid (s^t, t) \succeq_i (O, t)\}$. If Pos_i^t is not empty we define $\widehat{s^i}^t = min_{\succeq_i}\{Pos_i^t\}$, otherwise we define $\widehat{s^A}^t = (-1, W+1)$ and $\widehat{s^B}^t = (W+1, -1)$.*

We would like now to introduce two additional assumptions that will ensure that an agreement might be reached.

(A5) For every $t \in \mathcal{T}$ if $\widehat{s^i}^t \geq 0$ then $(\widehat{s^j}^t, t) \succ_i (\widehat{s^i}^t, t)$ and if $\widehat{s^i}^{t+1} \geq 0$ then $(\widehat{s^i}^{t+1}, t+1) \succ_j (\widehat{s^j}^t, t)$ where $i, j \in \{A, B\}$ and $i \neq j$.

Assumption (A4) ensures that if there are some agreements agent i prefers in the next period over opting out, then there is at least one of those agreements that agent j also prefers over opting out in this period.

An additional assumption is necessary to ensure that an agreement is possible at least in the first period.

(A6) $\widehat{s^i}_i^0 \geq 0$

$\widehat{s^i}^0$ is the worst agreement for agent i in period 0 which is still better than opting out. So, the requirement that this agreement will be at least zero, which is in S, ensures that there exists at least one agreement agent i prefers over opting out.

In [Kraus and Wilkenfeld 1990b] we have proved that under the above assumptions, if there exists a period when one of the agents will prefer opting out over any agreement and the game has not ended in prior periods, then an agreement will be reached in the period prior to this period. This result is the basis to our main result in this case, which is described in the following theorem.

Theorem 2 *Let $(\hat{f},\hat{g})$ be a P.E. of a model satisfying A0-A6 such that $\widehat{s^i}^t_i - \widehat{s^j}^t_i < 2c_i + c_j$, $i,j \in \{A,B\}, i \neq j$. If B offers first then $P(\hat{f},\hat{g}) = ((\widehat{s^B}^1_A - 1 + c_A, W - (\widehat{s^B}^1_A - 1 + c_A)), 0))$ and if A offers first $P(\hat{f},\hat{g}) = ((W - (\widehat{s^A}^1_B - 1 + c_B), \widehat{s^A}^1_B - 1 + c_B), 0))$*

Proof: The proof of this theorem and the proof of Theorem 3 appear in [Kraus and Wilkenfeld 1990b].

Example 3 *Suppose the area to be painted can be divided into 20 sub-pieces only, i.e. $W = 20$. And let us assume that $c_P = 2$ and $c_B = 3$. We also assume that the agent with the paint, P, prefers opting out over painting more than 14 pieces in the first period, i.e., $\widehat{s^P}^0 = (14,6)$, and $\widehat{s^P}^1 = (13,7)$ and $\widehat{s^P}^2 = (15,5)$.*

Agent B prefers opting out over painting more than 11 pieces in the first period, i.e, $\widehat{s^B}^0 = (11,9)$ and $\widehat{s^B}^1 = (9,11)$ and $\widehat{s^B}^2 = (8,12)$.

If B starts the negotiations then he will paint 10 pieces and agent A will do 10. If agent P starts then agent B will do 9 and agent P will do 11. In both cases the negotiations will end after the first period.

Time is valuable only to one side

Suppose one of the agents does not lose as time goes on and even gains at least in the early stages of the negotiation. For example, the agents need to reach an agreement on sharing a resource, but one of the agents continues to use this resource until the agreement is reached. One may suspect that the agent who gains over time will try to delay reaching an agreement. Nevertheless, if the other agent can opt out of the negotiation, agreement can be reached without a delay with conditions similar to what the losing agent can gain from opting out. (see [Kraus and Wilkenfeld 1990a]).

Since we consider here the case of sharing resources we assume that the object is desirable, i.e. condition (A1) is as follows:

(A1) if $r_i > s_i$, then $(r,t) \succ_i (s,t)$;

We also modify condition (A3).

Each agent has a number c_i $i \in \{A,B\}$ such that:

(A3) $(s,t_1) \succeq_i (\bar{s},t_2)$ iff $(s_i + c_i * t_1) \geq (\bar{s}_i + c_i * t_2)$. We assume that agent A gains over time ($c_A > 0$) and that agent B loses over time ($c_B < 0$), i.e., agent B prefers to obtain any given number of units sooner rather than later, while agent A prefers to obtain any given number of units later rather than sooner.

Furthermore, we assume that agent B prefers to opt out sooner rather than later and vice versa for agent A. Formally,

(A4) if $t_1 < t_2$ $(O,t_1) \succ_B (O,t_2)$ and $(O,t_2) \succ_A (O,t_1)$.

We also modify condition (A5).

(A5) For every $t \in \mathcal{T}$ $(\widehat{s^B}^t, t) \succ_B (\widehat{s^B}^{t+1}, t+1)$ and if $\widehat{s^B}^t_A \geq 0$ then $(\widehat{s^B}^t, t) \succ_A (O, t+1)$.

Assumption (A4) ensures that if there are some agreements agent B prefers over opting out, then there is at least one of those agreements that agent A also prefers over opting out in the next period.

Our main results are summarized in the following theorem.

Theorem 3 *Let $(\hat{f},\hat{g})$ be a P.E. of a model satisfying A0-A6. Suppose agent B offers first and $\widehat{s^B}^t_A - \widehat{s^B}^{t-1}_A \leq c_A$. If $|c_B| \geq c_A + 1$, then $P(\hat{f},\hat{g}) = ((\widehat{s^B}^1_A + 1 + c_A, \widehat{s^B}^1_B - 1 - c_A), 0)$ If $|c_B| < c_A + 1$, then $P(\hat{f},\hat{g}) = ((\widehat{s^B}^1_A, \widehat{s^B}^1_B), 1)$.*

If A is the first agent then $P(\hat{f},\hat{g}) = ((\widehat{s^B}^0_A, \widehat{s^B}^O_B), 0)$.

The Application of the Theory in Building Autonomous Agents

How can one use the above theoretical results in building agents capable of acting and negotiating under time constraints and complete information?

We note that in each of the cases we have investigated, the perfect-equilibrium strategies are determined by parameters of the situation. For example, in the case in which the agents have utility functions with time constants discount rate the strategies are determined by those discount factors (δ_i). Or, for example, in the case in which the agents can opt out and they have constant delays (c_i), the strategies depend on the constant delays and the worst agreement for a player which is still better for him than opting out in period one ($\hat{s^i}^1$).

So, one can supply agents with the appropriate strategies for each of the cases we have dealt with. When the agent participates in one of those situations, he will need to recognize which type of situation it is. Assuming the agent is given the appropriate arguments about the situation it is involved in (i.e. in the case the agents have utility functions with time constants what is the value of δ_i), he can construct the exact strategy for its specific case and use it in the negotiations. Since we provide the agents with unique perfect equilibrium strategies, if we announce it to the other agents in the environment, the other agents can not do better than to use their similar strategies.

Conclusion and Future Work

We have demonstrated how the incorporation of time into the negotiation procedure contributes to a more efficient negotiation process. We show that in different cases this model, together with the assumptions of complete information and that the agents' strategies induce an equilibrium in any stage of the negotiation, may result in the agent being able to use negotiation strategies that will end the negotiation without a delay. We suggest that these results are useful in particular in situations with time constraints. We are in the process of using this model in developing agents that will participate in crisis situations where time is an important issues.

The most obvious outstanding question concerns the relaxation of the assumption of complete information. In many situations the agents do not have full information concerning the other agents. Several works in game theory and economics have considered different versions of the model of Alternative Offers with incomplete information. (See for example, [Rubinstein 1985; Osborne and Rubinstein 1990; Chatterjee and Samuelson 1987]). We are in the process of modifying those results for use in DAI environments.

Acknowledgement

We would like to thank K. Arrow, R. Aumann, B. Bueno de Mesquita, J. Oppenheimer, A. Rubinstein and P. Young for helpful suggestions about game theory and negotiation. We would also like to thank J. Hendler and D. Perlis for helpful discussions.

References

Bond, A. H. and Gasser, L. 1988. An analysis of problems and research in DAI. In Bond, A. H. and Gasser, L., editors, *Readings in Distributed Artificial Intelligence.* Morgan Kaufmann Publishers, Inc., San Mateo, California. 3–35.

Chatterjee, K. and Samuelson, L. 1987. Bargaining with two-sided incomplete information: An infinite horizon model with alternating offers. *Review of Economic Studies* 54:175–192.

Conry, S. E.; Meyer, R. A.; and Lesser, V. R. 1988. Multistage negotiation in distributed planning. In Bond, A. H. and Gasser, L., editors, *Readings in Distributed Artificial Intelligence.* Morgan Kaufmann Publishers, Inc., San Mateo, California. 367–384.

Davis, R. and Smith, R.G. 1983. Negotiation as a metaphor for distributed problems solving. *Artificial Intelligence* 20:63–109.

Durfee, E. H. 1988. *Coordination of Distributed Problem Solvers.* Kluwer Academic Publishers, Boston.

Georgeff, M. 1983. Communication and interaction in a multi-agent planning. In *Proc. AAAI-83*, Washington, D.C. 125–129.

Kraus, S. and Wilkenfeld, J. 1990a. An automated strategic model of negotiation: Extended abstract. In *AAAI-90 Workshop on Reasoning in Adversarial Domains*, Boston.

Kraus, S. and Wilkenfeld, J. 1990b. The function of time in cooperative negotiations. Technical Report UMIACS TR 90-131 CS TR 2547, Institute for Advanced Computer Studies, University of Maryland.

Kraus, S. and Wilkenfeld, J. 1991. Negotiations over time in a multi agent environment: Preliminary report. In *Proc. of IJCAI-91*, Australia. To appear.

Kraus, S.; Ephrati, E.; and Lehmann, D. 1991. Negotiation in a non-cooperative environment. *Journal of Experimental and Theoretical Artificial Intelligence.* Accepted for publication.

Luce, R. D. and Raiffa, H. 1957. *Games and Decisions.* John Wiley and Sons.

Malone, T. W.; Fikes, R. E.; and Howard, M. T. 1988. Enterprise: A marketlike task schedule for distributed computing environments. In Huberman, B. A., editor 1988, *The Ecology of Computation.* North Holland.

Muthoo, A. 1989. A note on bargaining over a finite number of feasible agreements. Unpublished paper, Department of Economics, University of Bristol.

Nash, J. F. 1953. Two-person cooperative games. *Econometrica* 21:128–140.

Osborne, M. J. and Rubinstein, A. 1990. *Bargaining and Markets.* Academic Press Inc., California.

Raiffa, H. 1982. *The Art and Science of Negotiation.* Harvard University Press.

Rosenschein, J. and Genesereth, M.R. 1985. Deals among rational agents. In *Proc. of IJCAI-85*, Los Angeles, California. 91–99.

Rubinstein, A. 1982. Perfect equilibrium in a bargaining model. *Econometrica* 50(1):97–109.

Rubinstein, A. 1985. A bargaining model with incomplete information about preferences. *Econometrica* 53(5):1151–1172.

Sathi, A. and Fox, M. S. 1989. Constraint-directed negotiation of resource reallocations. In r, L. Gasse and Huhns, M. N., editors, *Distributed Artificial Intelligence, Volume II.* Pitman/Morgan Kaufmann, London. 163–194.

Shaked, A. and Sutton, J. 1984. Involuntary unemployment as a perfect equilibrium in a bargaining model. *Econometrica* 52(6):1351–1364.

Sycara, K.P 1987. *Resolving Adversarial Conflicts: An Approach to Integrating Case-Based and Analytic Methods.* Ph.D. Dissertation, School of Information and Computer Science, Georgia Institute of Technology.

Zlotkin, G. and Rosenschein, J. 1990. Negotiation and conflict resolution in non-cooperative domains. In *Proceedings of AAAI-90*, Boston,MA. 100–105.