

Decommitment in Multi-resource Negotiation

(Short Paper)

Bo An
Dept. of Computer Science
University of Massachusetts,
Amherst, USA
ban@cs.umass.edu

Victor Lesser
Dept. of Computer Science
University of Massachusetts,
Amherst, USA
lesser@cs.umass.edu

Kwang Mong Sim
Dept. of Computer Science
Hong Kong Baptist University,
Hong Kong
bsim@comp.hkbu.edu.hk

ABSTRACT

This paper presents the design and implementation of negotiation agents that negotiate with other entities for acquiring multiple resources. In our approach, agents utilize a time-dependent negotiation strategy in which the reserve price of each negotiation issue is dynamically determined by 1) the likelihood that negotiation will not be successfully completed (*conflict probability*), 2) the *expected agreement price* of the issue, and 3) the *expected number of final agreements*. Results from a series of experiments indicate that on average, our negotiation strategy achieved higher average utility than traditional negotiation strategies.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent Systems

General Terms

Algorithms, Design

Keywords

Multi-resource negotiation, negotiation agents, heuristics

1. INTRODUCTION

This paper investigates automated negotiation in resource allocation among resource providers (sellers) and consumers (buyers), where consumer agents may require multiple resources to successfully complete their tasks. Therefore, consumer agents may need to engage in multiple negotiations. If the multiple negotiations are not all successful, consumers gain nothing. The negotiation problem in this paper has the following two features: 1) When acquiring multiple resources, a resource consumer agent only knows the reserve price available for the entire set of resources, i.e., the highest price the agent can pay for all the resources, rather than the reserve price of each separate resource. 2) Agents can decommit from tentative agreements at the cost of paying a penalty.

Because resource providers and consumers may have different goals, preferences, interests, and policies, the problem of negotiating an optimal allocation of resources within a group of agents has been found to be intractable both in computation [1] and communication [2]. The multi-resource negotiation studied in this paper is even more complex due to decommitment. An agent's bargaining

Cite as: Decommitment in Multi-resource Negotiation (Short Paper), Bo An, Victor Lesser and Kwang Mong Sim, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. 1553-1556.

Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

position in each round is determined by many factors like market competition, deadline, current agreement set, trading partners' proposals, and market dynamics. During each round of negotiation, an agent has to make many decisions and there are many possible choices for each decision. Thus, it is difficult to construct an integrated framework in which all these factors are optimized concurrently. Rather than explicitly modeling those inter-dependent factors and then determining each agent's best decisions by the intractable combined optimization, this work tries to connect those inter-dependent factors and develops a set of heuristics to approximate agents' decision making during negotiation. In our approach, agents utilize a time-dependent negotiation strategy in which the reserve price of each negotiation issue is dynamically determined by 1) the likelihood that negotiation will not be successful (*conflict probability*), 2) the *expected agreement price* of the issue, and 3) the *expected number of final agreements* given the set of tentative agreements made so far.

2. NEGOTIATION MECHANISM

We make the following assumptions: 1) Agents have incomplete information about others. 2) A consumer agent negotiates over multiple negotiation issues in parallel and, for each negotiation issue, the agent concurrently negotiates with its trading partners.

All the analysis in this paper is from the perspective of a randomly selected buyer a . Let $\mathcal{I} = \{I_1, I_2, \dots, I_l\}$ be the set of negotiation issues of a and τ be the negotiation deadline. Let a negotiation period of a be denoted by t , $t \in \{0, 1, \dots, \tau - 1\}$. On issue I_j , a has a set $\mathcal{TP}_j^t$ of trading partners at round t . Also, a has a set $\mathcal{CP}_j^t$ of trading competitors on negotiation issue I_j at round t . $\phi_{a \rightarrow s}^t$ is the proposal of a to its trading partner $s \in \mathcal{TP}_j^t$ at round t . $\phi_{s \rightarrow a}^t$ is the proposal of agent s to a at round t . RP and IP are the reserve price and the desirable price of a before negotiation, respectively. IP_j is a 's initial proposal for negotiation issue I_j , i.e., $\phi_{a \rightarrow s}^0$, and it follows that $\sum_j IP_j = IP$. RP^t is a 's reserve price for all negotiating issues $\mathcal{I}^t$ at round t . Once a tentative agreement about I_j becomes a final agreement, a doesn't need further negotiation about I_j . Therefore, $\mathcal{I}^t \subseteq \mathcal{I}^{t-1} \subset \mathcal{I}$.

An agent can decommit from an agreement within λ rounds after the agreement has been made. Assume a makes an agreement Ag about issue I_j with agent s at round $\text{Tm}(Ag) = t$ and the agreement price is $\text{Pr}(Ag)$. Assume a decommits from the agreement Ag at round t' where $t' - \text{Tm}(Ag) \leq \lambda$. The penalty of the decommitment is defined by $\rho(\text{Pr}(Ag), t, t', \lambda)$. This work assumes that penalty 1) increases with time and agreement price and 2) the maximum penalty is less than the agreement price. Therefore, if an agent makes unnecessary agreements for a resource, it will decommit from these unnecessary agreements.

If the two parties decommit at the same time, they don't need to

pay a penalty to each other. An agreement made in the bargaining process is called a *tentative* agreement and it becomes a *final* agreement if neither party decommits from the agreement in the λ rounds after the agreement was made.

From the perspective of the whole negotiation, all the negotiation issues of a are dependent in the sense that a 's utility from the whole negotiation depends on the agreements on all the issues. a tries to make agreements for all the issues and a gains nothing if it fails to make an agreement on one issue, no matter how many and how good the other agreements are. The utility function of a when a makes at least one final agreement for each issue is defined as:

$$\mathcal{U} = RP - \sum_{I_j \in \mathcal{I}} \sum_{Ag \in \mathcal{FAG}_j^{\tau+\lambda}} \text{Prc}(Ag) + \sum_{t=0}^{\tau+\lambda} (\rho_{in}^t - \rho_{out}^t)$$

where $\tau + \lambda$ is the maximum period that a is involved in negotiation and decommitment, $\mathcal{FAG}_j^{\tau+\lambda}$ is the set of final agreements on negotiation issue I_j at time $\tau + \lambda$, ρ_{out}^t is the penalty a pays to other agents at round t , and ρ_{in}^t is the payment of penalty a receives from other agents at round t . If a fails to make a final agreement for at least one issue, a gains nothing and the utility is defined as:

$$\mathcal{U} = - \sum_{I_j \in \mathcal{I}} \sum_{Ag \in \mathcal{FAG}_j^{\tau+\lambda}} \text{Prc}(Ag) + \sum_{t=0}^{\tau+\lambda} (\rho_{in}^t - \rho_{out}^t)$$

Negotiation consists of a *bargaining stage* and a *decommitment stage*. A pair of buyer and seller agents bargain by making proposals in alternate rounds. At each round, one agent makes a proposal first, then the other agent has three choices in the bargaining stage: 1) accept the proposal, 2) reject the proposal, or 3) make a counter proposal. This work assume that buyers always propose first to sellers at each round of negotiation. Bargaining between two agents terminates 1) when an agreement is reached or 2) with a conflict when one of the two agents' deadline is reached or one agent quits the negotiation. An agent has the opportunity to decommit from the agreement within λ rounds after the agreement has been made and the decommitting agent pays the penalty to the other party involved in the decommitted agreement.

3. NEGOTIATION STRATEGIES

3.1 An overview of negotiation strategies

Algorithm 1 gives an overview of a 's strategy during the bargaining stage and decommitment stage. At round $t = 0$, a needs to make an initial proposal IP_j to each trading partner s . During each later round ($t > 0$), a will always first update its information structures (see Algorithm 2). First, if another agent revokes an agreement, then remove the agreement from the tentative agreement set. Second, if another agent sends a rejection proposal, then this corresponding negotiation thread terminates. If another agent accepts a proposal, then add the agreement into tentative agreement set. If one tentative agreement becomes a final agreement for the issue I_j as the negotiation moves to a new round, then a will decommit from all tentative agreements about I_j , stop all negotiation threads for I_j , and remove I_j from $\mathcal{I}^t$.

Next a computes τ_j^t for each issue $I_j \in \mathcal{I}^t$ and generates a proposal $\phi_{a \rightarrow s}^t$ to each trading partner $s \in \mathcal{TP}_j^t$ (Section 3.3). If $\phi_{a \rightarrow s}^t < \phi_{s \rightarrow a}^{t-1}$, then a sends the proposal to s directly. Otherwise, it adds $\langle \phi_{s \rightarrow a}^{t-1}, t \rangle$ into $\mathcal{TAG}_j^t$, which is a 's set of tentative agreements about issue I_j at round t .

For issue I_j , a checks whether the current set of agreements are more than necessary. If the current set of agreements is larger

Algorithm 1 Negotiation Strategy of Agent a

Data Structure: Tentative agreement set $\mathcal{TAG}_j^t$, final agreement set $\mathcal{FAG}_j^t$, sellers' proposal set for each issue I_j at round t .

Output: Final agreement set $\mathcal{FAG}_j^t$ for each I_j

- 1: Initial proposing: Let $t = 0$ and propose IP_j to every trading partner s about I_j .
 - 2: **repeat**
 - 3: $t++$;
 - 4: $\mathcal{I}^t = \mathcal{I}^{t-1}$;
 - 5: $\mathcal{TAG}_j^t = \mathcal{TAG}_j^{t-1}$, $\mathcal{FAG}_j^t = \mathcal{FAG}_j^{t-1}$ for $I_j \in \mathcal{I}^t$;
 - 6: **Step 1: initialization**
 - 7: **Step 2: deadline calculation**
 - 8: **Step 3: proposal generation**
 - 9: **Step 4: meet the agreement number constraint**
 - 10: **Step 5: meet the budget constraint**
 - 11: **Step 6: send left proposals**
 - 12: **until** 1) $t \geq \tau + \lambda$, or 2) $|\mathcal{FAG}_j^t| > 0$ for each I_j , or 3) $|\mathcal{TAG}_j^t| = 0$ for some I_j at $t \geq \tau_j^t$
-

than needed, then a recursively removes agreement Ag which minimizes the penalty needed for decommitting from that agreement. If $Ag \in \mathcal{TAG}_j^{t-1}$, then a decommits the agreement; otherwise, a sends the agent a proposal worse than $\phi_{s \rightarrow a}^{t-1}$.

Additionally, a checks whether the remaining agreement set violates the budget constraints. If this is true, then a recursively removes Ag that maximizes the expected utility of the remaining agreement set. If $Ag \in \mathcal{TAG}_j^{t-1}$, a decommits the agreement; otherwise, then a sends the agent a less desirable proposal.

Finally, if there are some agreements $\mathcal{TAG}_j^t$ but not in $\mathcal{TAG}_j^{t-1}$, a sends an *accept* proposal to those corresponding agents.

The whole negotiation process will terminate if 1) the deadline is reached, or 2) a makes a final agreement for each issue I_j , or 3) $|\mathcal{TAG}_j^t| = 0$ for some I_j at $t \geq \tau_j^t$, which means it no longer makes any sense for a to make any other agreements.

3.2 Different deadlines for different issues

The intuition behind using different negotiation deadlines for different issues is based on the following scenario: a makes an agreement about a scarce resource I_j before the deadline approaches. However, the other party to the agreement later decommits from the agreement. Then, the whole negotiation fails as it's difficult for agent a to get another agreement for the scarce resource I_j and thus a needs to pay the penalty for its other agreements. To avoid the situation happening, we can reduce the deadlines of scarce resources to increase the likelihood that we have a final agreement for those resources in place before the negotiation deadline.

The scarcity of a resource is evaluated based on the competition situation of the negotiation over issue I_j and s 's satisfaction about the agreement Ag . The competition situation of an agent is determined by the probability that it is being (not being) considered as the most preferred trading partner [4]. An agent's preferred trading partner refers to the one makes the best offer to the agent. a has $\mathcal{CP}_j^t$ competitors and $\mathcal{TP}_j^t$ partners. The probability that a is not the most preferred trading partner of any trading partner is $\mathcal{CP}_j^t / (\mathcal{CP}_j^t + 1)$. The probability of the agent a not being the most preferred trading partner of all the trading partners is

$$C_j^t = \left(\frac{\mathcal{CP}_j^t}{\mathcal{CP}_j^t + 1} \right)^{\mathcal{TP}_j^t}$$

C_j^t measures the scarcity of resource I_j at t . The bigger the

Algorithm 2 Initialization

```

1: for each  $I_j \in \mathcal{I}^t$  do
2:   for each  $s \in \mathcal{TP}_j^{t-1}$  do
3:     if  $\phi_{s \rightarrow a}^t = \text{"revoke } Ag \text{"}$  then
4:       remove  $Ag$  from  $\mathcal{TAG}_j^t$ 
5:     else
6:       if  $\phi_{s \rightarrow a}^t = \text{"reject"}$  then
7:         remove  $s$  from  $\mathcal{TP}_j^t$ 
8:       end if
9:     else
10:      if  $\phi_{s \rightarrow a}^t = \text{"accept"}$  then
11:        add  $\langle \phi_{s \rightarrow a}^{t-1}, t \rangle$  into  $\mathcal{TAG}_j^t$ 
12:      end if
13:    end if
14:  end for

15: for each  $Ag \in \mathcal{TAG}_j^t$  do
16:   if  $t - \text{Tm}(Ag) > \lambda$  then
17:     remove  $Ag$  from  $\mathcal{TAG}_j^t$  and add it to  $\mathcal{FAG}_j^t$ 
18:   end if
19: end for
20: if  $|\mathcal{FAG}_j^t| > 0$  then
21:   decommit from all agreements in  $\mathcal{TAG}_j^t$ , stop all negoti-
     ation threads for  $I_j$ , and remove  $I_j$  from  $\mathcal{I}^t$ .
22: end if
23: end for

```

value of C_j^t , the scarcer the resource I_j . If resource I_j is scarce and the other resources are sufficient, it's reasonable to decrease I_j 's deadline in order to get a "firm" negotiation result about negotiation issue I_j in order to decrease the probability that the whole negotiation fails due to the failure of the negotiation about issue I_j . However, if all the desired resources are scarce, it may be not necessary to decrease the deadline of all the resources. In other words, whether to decrease the deadline of the resource I_j may not depend on the absolute scarcity of the resource, but rather its "relative scarcity". The relative scarcity of the resource I_j is defined as the ratio of the I_j 's scarcity measure to the harmonic mean of the scarcity measure of all the resources:

$$RC_j^t = \frac{C_j^t}{\frac{|\mathcal{I}^t|}{\sum_{I_k \in \mathcal{I}^t} \frac{1}{C_k^t}}} = \frac{C_j^t \sum_{I_k \in \mathcal{I}^t} \frac{1}{C_k^t}}{|\mathcal{I}^t|}$$

Given the relative scarcity of each resource $I_j \in \mathcal{I}^t$, the deadline of issue I_j at time t is given as follows

$$\tau_j^t = \begin{cases} \tau & \text{if } RC_j^t < 1 \\ (RC_j^t + \frac{1}{\tau} - 1)^\epsilon & \text{if } RC_j^t \geq 1 \end{cases}$$

where $\epsilon < 0$. If the resource I_j is not scarce as compared to most resources, its deadline will remain the same. Otherwise, i.e., $RC_j^t \geq 1$, its deadline τ_j^t is smaller than τ and it can be found that τ_j^t will decrease with the increase of RC_j^t . That is, a scarcer resource will have a shorter deadline.

3.3 Generating proposals

Assume that a is negotiating with s about issue I_j . Then, a 's proposal to s at round t is given by:

$$\phi_{a \rightarrow s}^t = IP_j + (RP_j^t - IP_j)\delta_j^t$$

where RP_j^t is agent a 's current reserve price associated with negotiation issue I_j at round t and δ_j^t is agent a 's concession rate with

respect to negotiation issue I_j at round t , which is given by

$$\delta_j^t = T(t, \tau_j^t, \lambda) = (t/\tau_j^t)^\varepsilon$$

With infinitely many values of ε , there are infinitely many possible strategies in making concessions with respect to the remaining time. However, they can be classified into: 1) *Linear*: $\varepsilon = 1$, 2) *Conciliatory*: $0 < \varepsilon < 1$, and 3) *Conservative*: $\varepsilon > 1$ [4]. Before using any strategy, a needs to decide its reserve price RP_j^t . To calculate RP_j^t , we consider three factors: 1) the conflict probability χ_j^t , which measures the probability that a 's negotiation on resource I_j at round t will run into a conflict. 2) expected agreement price ϖ_j^t of issue I_j , and 3) the expected number $\psi(\mathcal{TAG}_j^t)$ of final agreements, which is based on the decommitment probability of each agreement.

RP_j^t is defined as:

$$RP_j^t = RP^t \frac{\chi_j^t \varpi_j^t \gamma(\mathcal{TAG}_j^t)}{\sum_{j=1} \chi_j^t \varpi_j^t \gamma(\mathcal{TAG}_j^t)}$$

where $RP^t = RP - \sum_{I_j \in \mathcal{I}} \sum_{Ag \in \mathcal{FAG}_j^t} \text{Prc}(Ag) + \sum_{t=0}^{t-1} (\rho_{in}^t - \rho_{out}^t)$ is agent a 's reserve price for all issues at round t .

Therefore, an agent's reserve price for resource I_j at round t will change over time during negotiation and it can be found that RP_j^t increases with the conflict probability χ_j^t , expected agreement price ϖ_j^t , and the expected number $\psi(\mathcal{TAG}_j^t)$ of final agreements.

4. EXPERIMENTATION

Experiments were carried out to study and compare the performance of our buyer agents (*HBAs*, heuristic-based buyer agents) with General buyer agents (*GBAs*). *GBAs* also use a time-dependent strategy in which the reserve price of issue I_i is determined by considering the distribution of the reserve price of issue I_i . *GBAs* make only one tentative agreement for each issue and don't decrease the deadline of a scarce negotiation issue. Each seller agent in the experiments randomly choose a negotiation strategy from the set of alternations outlined in [3]: the time-dependent function (linear, conceder, conservative) and the behavior-dependant function.

4.1 Experimental settings

In the experiments, agents were subjected to different market types, deadlines, and number of resources to acquire. Market type depends on the probability of the agent being a buyer (or a seller). The deadline of an agent is randomly selected from [15, 70] as we found that: 1) for very short deadline (< 15), very few agents could complete deals, and 2) for deadlines > 70 , there was little or no difference in the performance of agents. Hence, for the purpose of experimentation, a deadline between the range of 18 – 25 (respectively, 35 – 45 and 60 – 70) is considered as short (respectively, moderate and long). Each buyer may have different number of resources to acquire through negotiation. The number of resources each job (or task) needs is randomly selected from 1 to 15 and 1 – 3 (respectively, 7 – 9 and 13 – 15) is considered as lower range (respectively, mid-range and upper range). The value of ε (eagerness) is randomly generated from [0.1, 10] as it was found that when $\varepsilon > 10$ (respectively, $\varepsilon < 0.1$), there was little or no difference in performance of agents.

We normalize the utilities $U \in [U_{\min}, RP]$ of different experiments into the same space [0, 1] where $U_{\min} > 0$ is highest penalty the buyer may pay when the negotiation fails and RP is the reserve price. The normalized utility u' of an experiment is $u' = (U + U_{\min}) / (RP + U_{\min})$.

For each environment, more than 200 runs of experiments and their normalized utilities were averaged. We chose different de-




Figure 1: Deadline and average utility.




Figure 2: Number of resources and average utility.

commitment deadlines and penalties functions. For the empirical results presented here, $\lambda = 6$ is chosen as the decommitment period and the penalty function is $0.1 \times \text{Prc}(Ag) \times ((t' - t)/\lambda)^{1/2}$.

4.2 Observations

Observation 1: The experimental results in Fig. 1 show that the negotiation results become more favorable with the increase of the deadline for both *HBAs* and *GBAs*. Given the same deadline, *HBAs* achieved higher average utilities than *GBAs*. For very short (respectively, very long) deadlines, the average utilities of *HBAs* were comparatively not much higher than that of *GBAs*. With short (respectively, long) deadlines, both types of agents have equally insufficient (respectively, sufficient) time to optimize their agreements.

Observation 2: From Fig.2, we can find that, as the number of resources to acquire increased, the average utilities of both *HBAs* and *GBAs* decreased. We can also find that the advantage of *HBAs* over *GBAs* increases with the increase of number of resources, which corresponds to the feature of *HBAs* that *HBAs* make adaptive control over the multi-resource negotiation.

Observation 3: It can be observed from Fig. 3 that with different level of market competitions, *HBAs* always generate higher average utilities than *GBAs*. Additionally, when the competition is pretty high (e.g., the buyer-seller ratio is 10 : 1), the average utilities of the two types of agents are pretty close. That is because in a market high competition, all the agents have little chance of mak-




Figure 3: Market competition and average utility.

ing agreements, and thus it's very hard to find good agreements to satisfy all the resource agreements. We can also find that the average utilities of the two types of agents are also close when the competition is pretty low (e.g., the buyer-seller ratio is 1 : 10). In this case, and *HBAs* did not significantly outperform *GBAs*.

5. CONCLUSIONS

The contributions of this paper include: 1) To avoid the risk of the "collapse" of the whole negotiation due to failing to get some scarce resources, negotiation agents have the flexibility to adjust the deadline for different resources based on market competition, which allows agents to response to uncertainties in resource planning. 2) Each agent utilizes a time-dependent strategy in which the reserve price of each negotiation issue is dynamically determined by considering (*conflict probability*), *expected agreement price*, and *expected number of final agreements*. 3) As agents are permitted to decommit from agreements, an agent can make more than one agreement for each issue and the maximum number of agreements is constrained by the market situation and budget.

Finally, a future agenda of this work includes: 1) This work assumes that a buyer gains nothing if it fails to make agreements for all the issues, which can be relaxed so that the buyer gets some utility for the agreements for part of the negotiation issues. 2) This work assumes that the penalty is determined prior to negotiation. Negotiation protocol would become more flexible if agents also can negotiate over penalty. 3) Negotiation problem will become more complex if there are dependencies between negotiation issues.

6. REFERENCES

- [1] P. E. Dunne. Extremal behaviour in multiagent contract negotiation. *Journal of Artificial Intelligence Research*, 23:41–78, 2005.
- [2] U. Endriss and N. Maudet. On the communication complexity of multilateral trading: Extended report. *Journal of Autonomous Agents and Multi-Agent Systems*, 11(1):91–107, 2005.
- [3] P. Faratin, C. Sierra, and N. R. Jennings. Negotiation decision functions for autonomous agents. *Int. Journal of Robotics and Autonomous Systems*, 24(3-4):159–182, 1998.
- [4] K. M. Sim and C. Y. Choi. Agents that react to changing market situations. *IEEE Trans. on Systems, Man and Cybernetics, Part B*, 33(2):188–201, 2003.