

Strategic agents for multi-resource negotiation

Bo An · Victor Lesser · Kwang Mong Sim

Published online: 20 May 2010
© The Author(s) 2010

Abstract In electronic commerce markets where selfish agents behave individually, agents often have to acquire multiple resources in order to accomplish a high level task with each resource acquisition requiring negotiations with multiple resource providers. Thus, it is crucial to efficiently coordinate these interrelated negotiations. This paper presents the design and implementation of agents that concurrently negotiate with other entities for acquiring multiple resources. Negotiation agents in this paper are designed to adjust (1) the number of tentative agreements for each resource and (2) the amount of concession they are willing to make in response to changing market conditions and negotiation situations. In our approach, agents utilize a time-dependent negotiation strategy in which the reserve price of each resource is dynamically determined by (1) the likelihood that negotiation will not be successfully completed (*conflict probability*), (2) the *expected agreement price* of the resource, and (3) the *expected number of final agreements*. The negotiation deadline of each resource is determined by its relative scarcity. Agents are permitted to decommit from agreements by paying a time-dependent penalty, and a buyer can make more than one tentative agreement for each resource. The maximum number of tentative agreements for each resource made by an agent is constrained by the market situation. Experimental results show that our negotiation strategy achieved significantly more utilities than simpler strategies.

Keywords Automated negotiation · Negotiation strategy · Multi-resource

B. An (✉) · V. Lesser
Department of Computer Science, University of Massachusetts, Amherst, MA 01003, USA
e-mail: ban@cs.umass.edu; boancqu@gmail.com

V. Lesser
e-mail: lesser@cs.umass.edu

K. M. Sim
Department of Information and Communications, Gwangju Institute of Science and Technology,
Gwangju, South Korea
e-mail: kmsim@gist.ac.kr

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. This is a simple form of multi-linked negotiation where the resources are independent but are interrelated. Resources are independent in the sense that there is no dependence between different resources, i.e., acquiring one resource doesn't constrain how the other resources are acquired. However, from the perspective of the overall negotiation, resources are dependent as an agent's utility is determined by the sum of the cost of acquiring all the required resources and whether it obtains all the required resources. The negotiation problem in this paper has the following three features:

1. When acquiring multiple resources, a 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. In practice, given a plan and its resource requirements, an agent can easily decide the reserve price for all the resources in that plan based on the overall worth of the task. However, it is difficult (even impossible) for a resource consumer to understand how to set the reserve price for each separate resource. In fact, we show experimentally that it is undesirable to set a fixed reserve price for an individual resource prior to beginning negotiations.
2. Agents can decommit from tentative agreements at the cost of paying a penalty. Decommithment allows agents to profitably accommodate new tasks arriving or new negotiation events. If these events make some existing contracts less profitable or infeasible for an agent, that agent can decommit from those contracts [40].
3. Negotiation agents are assumed to have incomplete information about other agents, for example, a buyer agent knows the distribution of the reserve price of a seller agent and the number of trading competitors. However, an agent's negotiation status (the set of proposals it has received) and negotiation strategy are its private information. For strategic reasons, a negotiation agent won't disclose such information during negotiation. During negotiation, negotiation agents can quit negotiation at any time, even without notifying their trading partners. When a buyer acquires multiple resources, it concurrently negotiates with sellers to reach agreements for all the resources.

The negotiation problem considered in this paper can be motivated from various contexts. As an example, consider the negotiation management component [1] for Collaborating, Autonomous Stream Processing systems (CLASP) [7], which has been designed and prototyped in the context of System S project [19] within IBM Research to enable sophisticated stream processing. There are multiple sites running the System S software, each with their own administration and goals. Each site may only have limited processing capabilities, so cooperation among these sites can frequently be of mutual benefit. Consider the situation that a site receives a job. After planning [36], the site finds that using only its local resources, it cannot satisfy all resource requirements of the plan. Then, the site negotiates with other sites to acquire resources needed using its negotiation management component [1]. For each resource, there can be multiple providers and the site concurrently negotiates with different resource providers to construct agreements for these resources. The plan can be executed if and only if *all* resource requirements are satisfied. Therefore, while making a proposal to a trading partner for one resource, the site needs to consider the dynamically changing

negotiation environments (e.g., the number of sites requiring the same resource) and the negotiation situation of other negotiations for the same resource and for other resources.

Currently, there are limited techniques based on auctions or independent negotiations over single resources for performing the assembly of multiple resources required by a task. A centralized approach such as reverse combinatorial auctions [9,32] requires a controlling agent (the auctioneer) for determining which agents receive which resources based on the bids submitted by individual agents. However, the auctioneer may face significant computational overload due to a large number of bids with complex structure. Assume that each buyer runs a reverse combinatorial auction, each seller may participate in multiple auctions as there are multiple buyers requiring its resource. It's difficult for each seller to derive its optimal bids for all the concurrent auctions. An alternative approach is that each buyer (seller) submits its resource requirement (supply) to a super agent and the super agent runs auctions for all the buyers (sellers). However, it may be difficult to find such an auctioneer agent that selfish agents can trust and can comply with the decisions made by the auctioneer. Moreover, in dynamic environments where resource supply and demand arrive randomly, it is very difficult for the auctioneer to decide when to run auctions [3]. In our distributed approach, allocations emerge as the result of a sequence of distributed negotiations and each selfish agent acts on behalf of itself. An agent can negotiate with other agents when needed. The distributed model is also more suitable for the situation when the needed resources are from multiple electronic marketplaces, and more natural in cases where resources belong to different selfish agents and finding optimal allocations may be (computationally) infeasible. We feel it is key that the acquisition of multiple resources necessary is seen as an integrated process in which the results/status of any one negotiation affects all other negotiations.

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 terms of the amount of computation [11] and communication needed [12]. The multi-resource negotiation studied in this paper is even more complex due to the possibility of agents' decommitting from previously made agreements. An agent's bargaining position in each round is determined by many factors such as market competition, negotiation deadlines, current agreement set, trading partners' proposals, and market dynamics. During each round of negotiation, an agent has to make decisions on how to proceed with each negotiation thread and there are many possible choices for each decision based on a variety of factors. Thus, it is difficult to construct an integrated framework in which all these factors are optimized concurrently. Rather than explicitly modeling these inter-dependent factors and then determining each agent's best decisions by an intractable combined optimization, this work tries to connect those inter-dependent factors indirectly and develops a set of heuristics to approximate agents' decision making during negotiation. The distinguishing feature of negotiation agents in this paper is their flexibility; they can adjust (1) the number of tentative agreements for each resource and (2) the amount of concession by reacting to (i) changing market conditions, and (ii) the current negotiation status of all concurrently negotiating threads. In our approach, agents utilize a time-dependent negotiation strategy in which the reserve price of each resource is dynamically determined by (1) the likelihood that negotiation will not be successful (*conflict probability*), (2) the *expected agreement price* of the resource, and (3) the *expected number of final agreements* given the set of tentative agreements made so far. The negotiation deadline of each resource is determined by both its scarcity and the overall deadline for the entire negotiation. A buyer agent can make more than one tentative agreement for each resource and the maximum number of tentative agreements is constrained by the market situation in order to avoid the agent's making more agreements than necessary.

Our work here is connected to several lines of research in agent-mediated negotiation including multi-issue negotiation (e.g., [14–16, 24–26, 45, 48]), one-to-many negotiation [4, 5, 8, 29, 30, 35], negotiation strategies (e.g., [13, 22, 41, 42, 44]), and decommitment (e.g., [1, 31, 40]) (please see Sect. 5 for details). This paper presents the first design of negotiation agents in dynamic and uncertain environments in which (1) a consumer negotiates for multiple resources and its negotiation fails if it fails to get some resources, and (2) agents can choose to decommit from existing agreements within a fixed period. This research is intellectually challenging because of both the complex interactions among concurrent negotiations for multiple resources and the uncertainty associated with the outcome of these negotiations. This research provides a deep understanding of the influence of sophisticated negotiation mechanisms on individual agents' performance in dynamic environments, and hence contributes to the construction of effective problem-solving approaches in open environments. The proposed approach can be used for designing negotiation agents in many practical applications like service composition [34], Grid resource management [43], and supply chain [50].

The remainder of this paper is organized as follows. Section 2 introduces the multi-resource negotiation problem. Section 3 presents agents' negotiation strategies. Section 4 reports experimental results and presents an analysis of the properties of our model. Section 5 summarizes related work, and Section 6 concludes this paper.

2 Negotiation mechanism

2.1 Assumptions

We make the following assumptions about agents' knowledge and strategies:

- (1) Agents have incomplete information about each other. The assumption of incomplete information is intuitive because in practice, agents have private information, and for strategic reasons, they do not reveal their strategies, constraints, or preferences. In [37, p. 54], it was noted that the strategy of a trading agent corresponds to its internal program, and extracting the true internal decision process would be difficult. Moreover, when selfish agents have competing interests, they may have incentive to deviate from protocols or to lie to other agents about their preferences. This paper assumes that (1) agents know the number of trading partners and competitors and (2) the distributions of trading partners' reserve price. The assumption that the number of trading partners is known is less restrictive or similar to the assumptions in most related work (e.g., [15, 27, 29–31]). We consider both assumptions are realistic in practice. For example, consider the streaming processing system CLASP [7], each resource provider (consumer) always posts its resource supply (requirement). Further, the distribution of trading partners' reserve prices can be learned as a result of repeated interaction with agents in the marketplace. We explored the sensitivity of these assumptions in the experiment section.
- (2) A consumer agent negotiates over multiple resources in parallel and, for each resource, the agent concurrently negotiates with its trading partners. Given that the buyer doesn't know how to appropriately set the reserve price of each of its resources, one approach that requires no prior knowledge of the marketplace about current resource scarcity and expected competition of a specific resource is for a consumer to negotiate over all the resources in parallel. For each resource, there are multiple trading partners and the agent concurrently negotiates with all the trading partners. Therefore, each negotiation thread




Fig. 1 Buyer *a*'s multi-resource negotiation problem

of one resource has multiple concurrently existing outside options. Generally, a buyer obtains more desirable negotiation outcomes when it negotiates concurrently with all the sellers in competitive situations in which there is information uncertainty and there is a deadline for the negotiation to complete [29,30]. Additionally, inefficiency may arise in sequential negotiation when considering the overall time cost to complete all the necessary negotiations [16].

2.2 The negotiation problem

All the analysis in this paper is from the perspective of a randomly selected buyer *a* (see Fig. 1). Let $\mathcal{I} = \{I_1, I_2, \dots, I_l\}$ be the set of resources needed by *a* and τ be *a*'s negotiation deadline. Let a negotiation period of *a* be denoted by $t, t \in \{0, 1, \dots, \tau - 1\}$. For resource I_j, a has a set \mathcal{TP}_j^t of trading partners (sellers) at round t . Also, *a* has a set \mathcal{CP}_j^t of trading competitors (buyers) for resource 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 seller agent s to *a* at round t . *RP* and *IP* are the reserve price (maximum amount of money *a* can spend) and the desirable price of *a* before negotiation begins, respectively. IP_j is *a*'s initial proposal price for resource 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 resources \mathcal{I}^t at round t . Once a tentative agreement (defined below) for I_j becomes a final agreement, *a* doesn't need further negotiation about I_j . Therefore, $\mathcal{I}^t \subseteq \mathcal{I}^{t-1} \subseteq \mathcal{I}$.

An agent can decommit from an agreement within λ rounds after the agreement has been made. Assume *a* makes an agreement *Ag* about resource I_j with agent *s* at round $Tm(Ag) = t$ and the agreement price is $Prc(Ag)$. Assume *a* decommits from the agreement *Ag* at round t' where $t' - Tm(Ag) \leq \lambda$. The penalty of the decommitment is defined by $\rho(Prc(Ag), t, t', \lambda)$. This paper assumes that (1) penalty functions are nonnegative, continuous, and nondecreasing 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. An example of such a penalty function is $0.1 \times Prc(Ag) \times ((t' - t)/\lambda)^\zeta$ where $\zeta > 0$.

Penalties could be different from one resource to another resource. 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. Agent *a* needs to fulfil all its final agreements, i.e., *a* needs to pay for all final

agreements, even through it needs only one final agreement for each resource. a tries to make agreements for all its resources and a gains nothing if it fails to make an agreement for any resource in I , no matter how many and how good the agreements for other resources are. In other words, a requires a set of resources and only receives a positive utility if it acquires all of them, and zero otherwise. This assumption makes sense in some practical domains like some supply chain or Grid applications where the failure of one step (or one sub-task) will result in the failure of the whole task. The utility function of a when a makes at least one final agreement for each resource is defined as:

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

where $\tau + \lambda$ is the maximum period that a was involved in negotiation and decommitment, $\mathcal{FAG}_j^{\tau+\lambda}$ is the set of final agreements for resource I_j at $\tau + \lambda$, ρ_{out}^t is the penalty a pays to other agents at t when it decommits, and ρ_{in}^t is the payment of penalty a receives from other agents at t if they decommit.

If a fails to make a final agreement for at least one resource, a gains nothing and its utility is defined as:

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

In this case, a does not get the value RP since its task cannot be completed and thus its utility may be negative. Its only “income” in this case is the penalty received from its trading partners.

2.3 The negotiation protocol

As agents can choose to decommit from agreements, negotiation consists of a *bargaining stage* and a *decommitment stage* for each negotiation thread. This work adopts the well known alternating offers protocol (see [38, p. 100]) so that a pair of buyer and seller agents in a negotiation thread bargain by making proposals to each other. 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. For ease of analysis, this work assumes that buyers always propose first to sellers during negotiation. Many buyer-seller pairs can bargain simultaneously since each pair is in a negotiation thread. If the seller accepts the proposal of the buyer, negotiation terminates with a tentative agreement. If the seller rejects the proposal of the buyer, negotiation terminates with no agreement. If the seller makes a counter proposal, bargaining proceeds to another round and the buyer can accept the proposal, reject the proposal, or make a counter proposal. Bargaining between two agents terminates (1) when an agreement is reached or (2) with a conflict (i.e., no agreement is made) when one of the two agents’ deadline is reached or one agent quits the negotiation. After a tentative agreement is made, an agent has the opportunity to decommit from the agreement and the decommitting agent pays the penalty to the other party involved in the decommitted agreement.

2.4 The negotiation strategy

An agent's negotiation strategy is a function from the negotiation history to its actions at each negotiation round [37]. An agent a 's negotiation strategy can be represented as a sequence of functions $f_a = \{f_a^t\}_{t=0}^\infty$, where f_a^t is a 's strategy at round t . As the agent is negotiating for multiple resources and there are multiple negotiation threads for each resource, the agent's negotiation strategy f_a^t specifies for the agent what to do at round t for each of the active negotiation thread. For each trading partner s , the agent a has four choices: (1) accept the proposal by s , (2) reject the proposal by s , (3) make a counter proposal to s in the bargaining stage, or (4) decommit from the agreement between a and s in the decommitment stage.

A strategy profile $\mathcal{F} = (f_a, f_{\mathcal{TP}}, f_{\mathcal{CP}})$ is a collection of strategies, one for each agent, where $f_{\mathcal{TP}}$ and $f_{\mathcal{CP}}$ are the strategies for a 's trading partners and trading competitors, respectively. Let $\mathfrak{S} : \mathcal{F} \rightarrow \mathcal{O}$ be a social choice function which determines the negotiation result given the negotiation strategies \mathcal{F} of all the agents. Given the strategy profile of all the agents, game theory has been widely applied in analyzing the equilibria of bargaining models (e.g., Nash equilibria, Sub-game perfect equilibria, Sequential equilibria) [33]. The analytic complexity of equilibrium analysis increases rapidly when more elements (e.g., deadline, outside options, bargaining costs, market competition) and more agents are included in the model. As a result, in most models, only one or two elements are considered. For example, Rubinstein [38] studies a two-player sequential bargaining game in which bargaining cost is considered. The latest advance in computing sequential equilibrium strategies only considers a bilateral bargaining model in which one agent has incomplete information about the deadline of the other agent [18]. We take a set of elements into account, for example, deadline, outside option, market competition, multiple resources, and decommitment. In addition, we are not assuming that agents have complete information about the factors considered in our framework, which makes agents' reasoning even more difficult. Therefore, we feel that it is impractical to formally model the complex interaction that occurs between the bargaining and decommitment nor the interaction among multiple resources in the framework.

If we assume that each agent has information, which could be a probabilistic distribution, about other agents's strategies (i.e., $f_{\mathcal{TP}}$ and $f_{\mathcal{CP}}$), the optimization problem of agent a is to find the optimal negotiation strategy f_a^* from the set \mathcal{F}_a of possible negotiation strategies:

$$f_a^* = \operatorname{argmax}_{f_a \in \mathcal{F}_a} u_a(\mathfrak{S}(f_a, f_{\mathcal{TP}}, f_{\mathcal{CP}}))$$

where $u_a(\mathfrak{S}(f_a, f_{\mathcal{TP}}, f_{\mathcal{CP}}))$ is a 's utility of the negotiation result $\mathfrak{S}(f_a, f_{\mathcal{TP}}, f_{\mathcal{CP}})$. Agent a 's optimization problem at each negotiation round t can be formulated as a Markov Decision Process (MDP) $\langle S, A, P, R \rangle$ where the state set S can be characterized by the market situation (e.g., the number of buyers or sellers, the agreement set of each buyer or seller), action set A consists of all the actions each agent can choose (e.g., a counter-proposal including the price, or decommitment decision), transition function P is determined by agents' negotiation strategies and the change of market with time, reward function R is based on the utility each agent can gain from a specific state. As the action space A is infinite, solving the MDP problem could be computationally intractable [6]. Moreover, as stated before, it's impractical to assume that agents have information about other agents' negotiation strategies. For strategy or privacy reasons, an agent is unwilling to broadcast its decisions.

Given that (1) it's hard (even impossible) to compute agents' equilibrium strategies, and (2) it's not appropriate to assume that a knows other agents' negotiation strategies, this paper presents a set of heuristics for agents to make negotiation decisions at each negotiation round. The set of heuristics consider many relevant issues such as the risk that their negotiation partners may decommit (and therefore the fact that ideally a buyer needs to secure more than one

Algorithm 1 Negotiation Strategy of Agent a

Data Structure: Tentative agreement set $\mathcal{T}\mathcal{AG}_j^t$, final agreement set $\mathcal{F}\mathcal{AG}_j^t$, sellers' proposal set for each resource I_j at round t .

Output: Final agreement set $\mathcal{F}\mathcal{AG}_j^t$ 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: $\mathcal{T}\mathcal{AG}_j^t = \mathcal{T}\mathcal{AG}_j^{t-1}$, $\mathcal{F}\mathcal{AG}_j^t = \mathcal{F}\mathcal{AG}_j^{t-1}$ for $I_j \in \mathcal{I}^t$;
- 6: **Step 1: initialization (Algorithm 2)**
- 7: **Step 2: deadline calculation (Sect. 3.2)**
- 8: **Step 3: proposal generation (Sect. 3.3)**
- 9: **Step 4: meet the agreement number constraint (Sect. 3.4)**
- 10: **Step 5: send left proposals**
- 11: **until** 1) $t \geq \tau + \lambda$, or 2) $|\mathcal{F}\mathcal{AG}_j^t| > 0$ for each I_j , or 3) $|\mathcal{T}\mathcal{AG}_j^t| = 0$ for some I_j at $t \geq \tau_j^t$

agreement for any given resource), the competition that buyers face from other buyers, uncertainty about the reserve prices of their trading partners, multiple opportunities of reaching an agreement, the set of available tentative agreements, deadline, and negotiation history.

3 Heuristics based strategies

Agent a has l resources to acquire, and for each resource, a conducts multi-threaded negotiation with a set of trading partners. For each negotiation thread associated with the acquisition of a resource, a needs to decide (1) what is its proposal during the bargaining stage and (2) when and whether to decommit from an agreement in the decommitment stage.

3.1 An overview of negotiation strategies

Algorithm 1 gives an overview of a 's strategy during the bargaining stage and the decommitment stage (Table 1).

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 decommits from an agreement, then remove the agreement from the tentative agreement set. Second, if another agent sends a message indicating rejection of the current proposal, the corresponding negotiation thread terminates. If another agent accepts a proposal, then add the agreement to the tentative agreement set. If one tentative agreement becomes a final agreement (no decommitment allowed) for the resource 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 .¹

¹ The only additional value that can be achieved by keeping alive any future negotiation is the possibility that a trading partner is likely to decommit. In this case, it would be profitable to delay decommitment and thus the agent does not need to pay the decommitment penalty but receives the penalty from its trading partner. However, since each buyer does not know whether a seller will decommit from an agreement and the penalty increases with time, the buyer may have to pay a higher penalty if it has to decommit before the unnecessary tentative agreement becomes a final agreement. We evaluated the benefit of delaying decommitment through

Table 1 Symbols used in this paper

τ	Deadline of agent a
I_j	Resource j
RP	Reserve price for all resources
IP	Desirable price for all resources
IP_j	Initial proposal for resource I_j
τ_j^t	Deadline of agent a for resource I_j at round t
\mathcal{I}^t	The set of resources at round t
\mathcal{TP}_j^t	The set of partners (sellers) about I_j at round t
\mathcal{CP}_j^t	The set of competitors (buyers) about I_j at round t
$\phi_{a \rightarrow s}^t$	a 's Proposal to s at round t
\mathcal{P}_j^t	a 's Trading partners' proposals about I_j at round t
RP^t	a 's Reserve price for all negotiating resources at round t
RP_j^t	a 's Reserve price for resource I_j at round t
\mathcal{TAG}_j^t	a 's Set of tentative agreements for resource I_j at round t
\mathcal{FAG}_j^t	a 's Set of final agreements for resource I_j at round t
$\text{Prc}(Ag)$	Price of the agreement Ag
$\text{Tm}(Ag)$	Time when the agreement Ag was made
ρ_{out}^t	The penalty a pays to other agents at round t
ρ_{in}^t	The payment of penalty a receives at round t
C_j^t	The scarcity of resource I_j at t
RC_j^t	The relative scarcity of resource I_j at t
δ_j^t	The concession rate with respect to resource I_j at round t
χ_j^t	The conflict probability of the negotiation for I_j at t
ω_j^t	The expected agreement price of resource I_j at t
$\omega_s^t(Ag)$	The probability of s 's decommitting from Ag at t
$\varphi(\mathcal{TAG}_j^t)$	the expected number of final agreements given \mathcal{TAG}_j^t
$\gamma(\mathcal{TAG}_k^t)$	model how $\varphi(\mathcal{TAG}_j^t)$ affects the offering price

Next a computes the negotiation deadline τ_j^t for each resource $I_j \in \mathcal{I}^t$ (Sect. 3.2) and generates a proposal $\phi_{a \rightarrow s}^t$ to each trading partner $s \in \mathcal{TP}_j^t$ (Sect. 3.3). If $\phi_{a \rightarrow s}^t < \phi_{s \rightarrow a}^{t-1}$ (i.e., s 's last proposal is not acceptable), then a sends the proposal $\phi_{a \rightarrow s}^t$ to s directly. Otherwise, it adds $\langle \phi_{s \rightarrow a}^{t-1}, t \rangle$ into tentative agreement set \mathcal{TAG}_j^t .

For resource I_j , a checks whether the current set of agreements are sufficient. If the current set of agreements is more than needed, a recursively removes agreements from the tentative agreement set (Sect. 3.4). Assume that Ag needs to be removed and the trading partner in the agreement Ag is seller s . If $Ag \in \mathcal{TAG}_j^{t-1}$, then a decommits from the agreement. If Ag is not in \mathcal{TAG}_j^{t-1} , the agreement Ag has been just added to \mathcal{TAG}_j^t by a at time t but the seller involved in the agreement hasn't received the "accept" message from a . Although a doesn't intend to make the agreement Ag and a can quit the negotiation with s , it's better for a to continue the negotiation with s and try to get better agreements than an agreement in the current tentative agreement set \mathcal{TAG}_j^t . Therefore, a removes Ag from \mathcal{TAG}_j^t and sends

Footnote 1 continued

experimentation and found that delaying decommitment did not increase the buyer's average utility. In the current implementation, we do not take this into account.

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  $\phi_{s \rightarrow a}^t = \text{"decommit from } Ag \text{"}$  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 negotiation threads for  $I_j$ , and remove  $I_j$  from  $\mathcal{I}^t$ .
22: end if
23: end for

```

s a proposal with lower price than the price in the agreement Ag . Finally, if an agreement Ag is contained in \mathcal{TAG}_j^t but is not in \mathcal{TAG}_j^{t-1} , then a sends an *accept* proposal to the corresponding seller involved in the agreement Ag .

The overall negotiation process will terminate if (1) the deadline is reached, or (2) a makes a final agreement for each resource 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.

This work assumes that a buyer agent always offers the same price to all trading partners of one resource. Formal analysis of concurrent negotiation [2] suggests that it is an agent's dominant strategy to make the same offer to all trading partners. While this paper considers more complex negotiation, it is still intuitive to not make price discrimination proposals for the same resource. While making an offer, a buyer hopes that the offer would be accepted. If there are two offers which have the same probability of being accepted, the buyer will choose the offer with the lower price.

3.2 Different deadlines for different resources

The number of buyers and sellers for different resources varies. A resource is easy to obtain if the number of sellers is much larger than the number of buyers. In contrast, if there are more buyers and less sellers, the resource is relatively difficult to obtain since the resource seems "scarce" in terms of the ratio of supply to demand. The intuition behind using different negotiation deadlines for different resources is based on the following scenario: a makes an agreement about a scarce resource I_j before the deadline approaches. However, the other party involved in the agreement later decommits from the agreement. Then, the overall negotiation fails as it's difficult for agent a to get another agreement for the scarce resource I_j in

the remaining time and thus a needs to pay the penalty for its other agreements. To decrease the possibility of this 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 overall negotiation deadline. In other words, we would like to quickly secure one final agreement for a scarce resource. On one hand, by decreasing one resource’s artificial deadline, a is inclined to make larger concessions to its trading partners and thus its probability of making a final agreement for the resource increases. On the other hand, if it’s difficult for a to make a final agreement for one resource, a can know this earlier. Thus a can pay less decommitment penalties by decommitting from agreements earlier as penalties increase with time. However, the determination of this virtual deadline for scarce resources is a dynamic process which can either decrease or increase the deadline as conditions change in the future.

The scarcity of a resource I_j is evaluated based on the competition situation of the negotiation over resource I_j . A negotiator’s bargaining “power” is affected by the number of competitors and trading alternatives. Multiple options give a negotiator more “power” since the negotiating party needs not pursue the negotiation with any sense of urgency. The competition situation of an agent is determined by the probability that it is considered as the most preferred trading partner [44]. An agent’s preferred trading partner refers to the one who makes the best proposal to the agent. a has \mathcal{CP}_j^t competitors and \mathcal{TP}_j^t partners. While it’s impossible for a to compute exactly the probability that it is considered as the most preferred trading partner since a doesn’t know other agents’ negotiation strategies, the probability can be approximated in the following way. 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 approximated by

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

C_j^t measures the scarcity of resource I_j at t . With more trading partners, it is relatively less difficult to acquire the resource and C_j^t will decrease. With more trading competitors, it is relatively more difficult to acquire the resource and C_j^t will increase.

If resource I_j is scarce and the other resources are not scarce, it’s reasonable to decrease I_j ’s deadline in order to decrease the probability that the overall negotiation fails due to the failure of the negotiation about resource I_j . However, if all the desired resources are scarce, it may not be 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:

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

Using harmonic mean, the scarcer resource dominates the deadline calculation, which is close to the practice. Given the relative scarcity of each resource $I_j \in \mathcal{I}^t$, the deadline of resource I_j at time t is given as follows

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

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

3.3 Generating proposals

Since bargaining is fundamentally time-dependent [13,22], agents utilize a time-dependent strategy when making concessions. Assume that a is negotiating with s about resource 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 of resource I_j at round t and δ_j^t is agent a 's concession rate with respect to resource I_j at round t , which is given by

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

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*: $\varepsilon > 1$ [44]. ε reflects an agent's mental state about its eagerness for finishing the negotiation earlier [13,22]. Before making proposals, 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 aspiration level of the current negotiation for resource I_j , (2) expected agreement price ϖ_j^t of resource I_j , and (3) the expected number $\varphi(\mathcal{T}AG_j^t)$ of final agreements based on the estimation of the decommitment probabilities of the current tentative agreement set. Function $\gamma(\mathcal{T}AG_k^t)$ is used to model the effect of the expected number $\varphi(\mathcal{T}AG_k^t)$ of final agreements.

RP_j^t is defined as:

$$RP_j^t = RP^t \frac{\chi_j^t \varpi_j^t \gamma(\mathcal{T}AG_j^t)}{\sum_{I_k \in \mathcal{I}} \chi_k^t \varpi_k^t \gamma(\mathcal{T}AG_k^t)}$$

where $RP^t = RP - \sum_{I_j \in \mathcal{I}} \sum_{Ag \in \mathcal{F}AG_j^t} \text{Prc}(Ag) + \sum_{t=0}^{t-1} (\rho_{in}^t - \rho_{out}^t)$ is agent a 's reserve price for all resources at round t , i.e., the maximum amount of money that it can spend to acquire all the remaining resources. We can see that the reserve price RP_j^t increases with the increase of the conflict probability χ_j^t and expected agreement price ϖ_j^t . If the current negotiation for resource I_j seems difficult, a needs to set a higher reserve price for resource I_j . Similarly, a needs to set a higher reserve price for resource I_j if the expected agreement price for resource I_j is high. Later we will show that $\gamma(\mathcal{T}AG_j^t)$ decreases with the increase of $\varphi(\mathcal{T}AG_j^t)$. Thus, the reserve price RP_j^t decreases with the increase of the expected number $\varphi(\mathcal{T}AG_j^t)$ of final agreements, which is intuitive as buyers don't need to set a higher reserve price for a resource I_j when a has already made enough tentative agreements for I_j .

Conflict probability χ_j^t : Suppose that at round t , a 's last proposal $\phi_{a \rightarrow s}^{t-1}$ generates a utility of v_a for itself and v_s for s , and its trading partner s 's proposal $\phi_{s \rightarrow a}^{t-1}$ generates a utility of w_s for itself and w_a for a . Since a and s are utility maximizing agents, $v_a > w_a$ and $v_s < w_s$. If a accepts s 's last proposal, then it will obtain w_a with certainty. If a insists on its last proposal and (1) s accepts it, a obtains v_a and (2) s does not accept it, a may be subjected

to a conflict utility c_a . c_a is the worst possible utility for a (i.e., a 's utility in the absence of an agreement with s). If s does not accept a 's last proposal, a may ultimately have to settle with lower utilities (the lowest possible being the conflict utility), if there are changes in the market situation in subsequent cycles. For instance, a may face more competitions in the next or subsequent cycles and may have to ultimately accept a utility that is lower than w_a (even c_a). If the subjective probability of obtaining c_a is p_c (conflict probability) and the probability that a achieving v_a is $1 - p_c$, and if a insists on holding its last proposal, a will obtain a utility of $(1 - p_c)v_a + p_c c_a$. Hence, a will find that it is advantageous to insist on its last proposal only if

$$(1 - p_c)v_a + p_c c_a \geq w_a$$

i.e., $p_c \leq (v_a - w_a)/(v_a - c_a)$ [41,42,44]. The maximum value of $p_c = (v_a - w_a)/(v_a - c_a)$ is the highest probability of a conflict that a may encounter in which $v_a = RP_j^t - \phi_{a \rightarrow s}^{t-1}$ and $w_a = RP_j^t - \phi_{s \rightarrow a}^{t-1}$. p_c is a ratio of two utility differences. While $v_a - w_a$ measures the cost of accepting the trading agent's last proposal, $v_a - c_a$ measures the cost of provoking a conflict. $v_a - c_a$ represents the range of possible values of utilities between the best case utility and the worst case (conflict) utility.

If there is no tentative agreement for resource I_j , i.e., $|TAG_j^t| = 0$, the worst case utility c_a is 0. If $|TAG_j^t| > 0$, a can use one of its tentative agreements as the finally agreement and c_a is defined as

$$\max_{Ag \in TAG_j^t} \left(RP_j^t - \text{PrC}(Ag) - \text{Pnt}(TAG_j^t - Ag, t, \lambda) \right)$$

where $\text{Pnt}(TAG, t, \lambda)$ is an estimation of the penalty a needs to pay while decommitting from the set of agreements TAG . $\text{Pnt}(TAG, t, \lambda)$ is defined as

$$\sum_{Ag \in TAG} \frac{\sum_{t'=t}^{\text{Tm}(Ag)+\lambda} \rho(\text{PrC}(Ag), \text{Tm}(Ag), t', \lambda)}{\text{Tm}(Ag) + \lambda - t + 1}$$

in which any agreement $Ag \in TAG$ can be decommitted at any time before the decommitment stage expires.

Aggregated probability of conflict: Let p_c^i be the conflict probability of a with any of its trading partner s and w_a^i be a 's utility if it accepts s 's proposal, then the aggregated conflict probability of a with all of its trading partners about I_j is given as follows [41,42,44]:

$$\chi_j^t = \prod_{i=1}^{|TP_j^t|} p_c^i = \prod_{i=1}^{|TP_j^t|} \frac{v_a - w_a^i}{v_a - c_a} = \frac{\prod_{i=1}^{|TP_j^t|} (v_a - w_a^i)}{(v_a - c_a)^{|TP_j^t|}}$$

Expected agreement price ϖ_j^t : Different resources have different ranges of agreement prices. For example, you may need to spend \$20,000 for a car but only need \$500 for a bike. Therefore, it's necessary to consider a resource's expected agreement price ϖ_j^t while determining the reserve price of the resource. ϖ_j^t is computed based on agent a 's estimation of the reservation price of a trading partner. The estimation is characterized by a probability distribution $F_s(\cdot)$, where $F_s(y)$ denotes the probability that the reservation price of a trading partner s is no greater than y . $F_s(y)$ is identical and independent across all sellers.² This probability distribution is the prior belief of the buyer. For simplicity, let $F_j(y) = F_s(y)$ denote the

² Our model can also be extended to allow $F_s(y)$ to be different for different trading partners.

probability that the reservation price of any trading partner $s \in \mathcal{TP}_j^t$ is no greater than y . The probability density function of $F_j(y)$ is denoted by $f_j(y)$. The desirable price IP_j for resource I_j is simply computed by considering sellers’ reserve price for resource I_j : $IP_j = \int_{-\infty}^{\infty} f_j(y)ydy$.

Let $F_j^k(y)$ be the probability distributions of the k th highest maximum reserve price. The probability density function of $F_j^k(y)$ is denoted by $f_j^k(y)$. $F_j^1(y)$ is equal to the product of the probabilities that the maximum reserve price is less than or equal to y in each thread. $F_j^2(y)$ is equal to $F_j^1(y)$ plus the probability that the highest maximum reserve price is greater than y , and the second highest maximum reserve price is less than or equal to y . These probabilities can be calculated by the following formulas:

$$\begin{aligned}
 F_j^1(y) &= (F_j(y))^{|\mathcal{TP}_j^t|} \\
 F_j^2(y) &= F_j^1(y) + C_{|\mathcal{TP}_j^t|}^1 (1 - F_j(y))^{2-1} (F_j(y))^{|\mathcal{TP}_j^t|-1} \\
 F_j^k(y) &= F_j^{k-1}(y) + C_{|\mathcal{TP}_j^t|}^{k-1} (1 - F_j(y))^{k-1} (F_j(y))^{|\mathcal{TP}_j^t|-k+1}
 \end{aligned}$$

The corresponding probability density functions are:

$$\begin{aligned}
 f_j^1(y) &= |\mathcal{TP}_j^t| (F_j(y))^{|\mathcal{TP}_j^t|-1} \\
 f_j^2(y) &= f_j^1(y) - C_{|\mathcal{TP}_j^t|}^1 f_j(y) (F_j(y))^{|\mathcal{TP}_j^t|-1} + C_{|\mathcal{TP}_j^t|}^1 \\
 &\quad \times (|\mathcal{TP}_j^t| - 1) f_j(y) (1 - F_j(y))^{2-1} (F_j(y))^{|\mathcal{TP}_j^t|-2} \\
 f_j^k(y) &= f_j^{k-1}(y) - C_{|\mathcal{TP}_j^t|}^{k-1} (k - 1) f_j(y) \\
 &\quad \times (1 - F_j(y))^{k-2} (F_j(y))^{|\mathcal{TP}_j^t|-k+1} \\
 &\quad + C_{|\mathcal{TP}_j^t|}^{k-1} f_j(y) (1 - F_j(y))^{k-1} (F_j(y))^{|\mathcal{TP}_j^t|-k+1}
 \end{aligned}$$

We provide a heuristic approach to estimate the expected agreement price for resource I_j . When the number of trading partners is less than the number of trading competitors, the agreement price follows the highest maximum reserve price distribution. Otherwise, the agreement price follows a lower reserve price distribution. This is also the case with less trading competitors. The intuition behind the heuristic is as follows. Consider the single-shot negotiation between buyers and sellers in which buyers make offers first and then sellers decide whether to accept or not. If there is no competitors, the equilibrium offer of the buyer a is sellers’ lowest reserve price. If there is one competitor, the equilibrium offer of the buyer a is the second lowest reserve price. In the same way, if there are $|\mathcal{CP}_j^t|$ competitors, the equilibrium offer is the $(|\mathcal{CP}_j^t| + 1)$ th lowest reserve price, i.e., $(|\mathcal{TP}_j^t| - |\mathcal{CP}_j^t|)$ th highest reserve price. Since in our model buyers don’t know sellers’ exact reserve prices, distributions are used instead. Formally, ϖ_j^t is given as follows:

$$\varpi_j^t = \begin{cases} \int_{-\infty}^{\infty} f_j^{|\mathcal{TP}_j^t|-|\mathcal{CP}_j^t|}(y)ydy & \text{if } |\mathcal{TP}_j^t| > |\mathcal{CP}_j^t| \\ \int_{-\infty}^{\infty} f_j^1(y)ydy & \text{if } |\mathcal{TP}_j^t| \leq |\mathcal{CP}_j^t| \end{cases}$$

where \bar{y} is the upper bound of the possible reserve price for resource I_j . The above estimation is “conservative” in the sense that we assume that agent a is less competitive than its trading competitors.

$\gamma(\mathcal{TAG}_j^t)$ models how the current set \mathcal{TAG}_j^t of agreements will affect agent a 's reserve price for resource I_j at round t . a will set a lower reserve price if it has made more agreements. Since current agreements may be decommitted in the future. Rather than considering the number $|\mathcal{TAG}_j^t|$ of agreements having already made, it's more prudent to use the expected number of final agreements, which can be computed based on the decommitment probabilities of agreement set \mathcal{TAG}_j^t . The decommitment probability of an agreement $Ag \in \mathcal{TAG}_j^t$ between a and s is approximated by considering the competition situation of negotiation over resource I_j and s 's satisfaction about the agreement Ag .

The competition situation of negotiation over resource I_j is evaluated by the probability that the agent s is not the most preferred trading partner is $[(\mathcal{TP}_j^t - 1)/\mathcal{TP}_j^t]^{CP_j^t+1}$ [41,42,44]. s 's satisfaction about the agreement Ag is estimated by the probability that the agreement is no worse than the trading partner's reserve price. The price of the agreement $Ag \in \mathcal{TAG}_j^t$ is $\text{Prc}(Ag)$, s 's satisfaction about the agreement Ag is $F_j(\text{Prc}(Ag))$.

Hence, the approximation of the probability of s 's decommitting from agreement $Ag \in \mathcal{TAG}_j^t$ is defined as:

$$\omega_s^t(Ag) = \vartheta \times \left(1 - \left(\frac{\mathcal{TP}_j^t - 1}{\mathcal{TP}_j^t} \right)^{CP_j^t+1} \right) (1 - F_j(\text{Prc}(Ag)))$$

For the tentative agreement set \mathcal{TAG}_j^t , the expected number of final agreements is $\varphi(\mathcal{TAG}_j^t) = \sum_{Ag \in \mathcal{TAG}_j^t} (1 - \omega_s^t(Ag))$. Given $\varphi(\mathcal{TAG}_j^t)$, buyer a can determine how it will affect the reserve price about resource I_j at round t . $\gamma(\mathcal{TAG}_j^t)$ decreases with the increase of $\varphi(\mathcal{TAG}_j^t)$ and can be defined as:

$$\gamma(\mathcal{TAG}_j^t) = \frac{1}{(1 + \varphi(\mathcal{TAG}_j^t))^2}$$

3.4 Maximum number of final agreements

Since trading partners may decommit from agreements, a may need to make more than one tentative agreement for resource I_j . Then, how many agreements are enough for the resource I_j ? For an agreement Ag between a and a trading partner s , s may be inclined to decommit if there are many buyers requesting the resource. On the other hand, s may be inclined to decommit if the agreement price is not favorable from s 's perspective. Here we provide an approach to decide the maximum number of agreements a can make on resource I_j at round t based on the expected number of final agreements. Given the expected number $\varphi(\mathcal{TAG}_j^t)$ of final agreements about resource I_j at t , a needs to decide whether the tentative agreements is enough or insufficient. If \mathcal{TAG}_j^t is more than needed, a may decommit from some agreements. If the agreement set is insufficient, a will make more agreements if the negotiation deadline hasn't approached. This work assumes that a only needs to make one final agreement for each resource. Therefore, by intuition, the most favorable result for agent a is that a makes exactly one final agreement for each resource.

As a only needs one final agreement about resource I_j , if $\varphi(\mathcal{TAG}_j^t) \gg 1$, only part of the final agreements will be used by a , which corresponds to the tentative agreement set $\mathcal{TAG} \subset \mathcal{TAG}_j^t$. Maintaining the tentative agreement set \mathcal{TAG} is better than maintaining

the tentative agreement set \mathcal{TAG}_j^t as in the later case, a needs to pay more for redundant agreements. Therefore, it's better for a to decommit from some agreements in \mathcal{TAG}_j^t .

Let φ_j^t be the satisfactory number of final agreements about resource I_j at t which represents the upper bound of the number of final agreements needed. Before the deadline is reached, a has the opportunity to make more agreements and thus reach one final agreement. Thus, the satisfactory number of final agreements about resource I_j at $t < \tau$ is 1, $\varphi_j^t = 1$. After the negotiation deadline, a will determine whether to decommit from any agreement \mathcal{TAG}_j^t for resource I_j at round $\tau \leq t < t + \lambda$. Is it the best option for a to set the satisfactory number of final agreements about resource I_j at t be 1? Consider the following scenario, at t , the expected number of final agreements for resource I_j is 1 and the expected number of final agreements about any other resource is close to 0, which implies that the negotiation about other resources has a very high failure probability. If a sets φ_j^t to be 1, it's with very high probability that a would need to decommit from all its agreements. Therefore, a will not set a high φ_j^t value if $\varphi(\mathcal{TAG}_k^t)$ is small for another resource I_k . On the other hand, a will try to increase the probability of making one final agreement for each resource as it's desirable for a to make one final agreement for each resource. Concerning above, φ_j^t is defined as:

$$\varphi_j^t = \begin{cases} 1 & \text{if } t < \tau \\ \min_{I_k \in \mathcal{I}'} \varphi(\mathcal{TAG}_k^t) & \text{if } \tau \leq t < t + \lambda \end{cases}$$

If $\sum_{Ag \in \mathcal{TAG}_j^t} (1 - \omega_s^t(Ag)) < \varphi_j^t$, a needs to make more agreements as the expected number of agreements is less than φ_j^t . If $\sum_{Ag \in \mathcal{TAG}_j^t} (1 - \omega_s^t(Ag)) > \varphi_j^t$, a needs to decommit from some agreements. Let the set of tentative agreement set after removing unnecessary agreements be \mathcal{TAG} . The optimization problem of computing \mathcal{TAG} is given by

$$\min_{\mathcal{TAG}} \sum_{Ag \in \mathcal{TAG}_j^t - \mathcal{TAG}} \rho(\text{Pr}(Ag), \text{Tm}(Ag), t, \lambda)$$

where \mathcal{TAG} satisfies $\sum_{Ag \in \mathcal{TAG}} (1 - \omega_s^t(Ag)) \leq \varphi_j^t$.

Theorem 1 *The optimization problem of removing redundant tentative agreements is NP-complete.*

Proof We show that the problem is NP-complete by formulating the problem as a 0–1 Knapsack problem, which is well known to be NP-complete.

Formal definition of 0–1 Knapsack problem: There is a knapsack of capacity $c > 0$ and N items. Each item has value $v_i > 0$ and weight $w_i > 0$. Find the selection of items ($\delta_i = 1$ if selected, 0 if not) that fit, $\sum_{i=1}^N \delta_i w_i \leq c$, and the total value, $\sum_{i=1}^N \delta_i v_i$, is maximized.

The set of tentative agreements $\mathcal{TAG}_j^t = \{Ag_1, \dots, Ag_N\}$ can be treated as items. The value of each item Ag_i is defined as the penalty if a decommits from the agreement, i.e., $v_i = \rho(\text{Pr}(Ag_i), \text{Tm}(Ag_i), t, \lambda)$. The weight of each item Ag_i is defined as the probability that Ag_i will not be decommitted by a 's trading partner, i.e., $w_i = 1 - \omega_s^t(Ag_i)$. The capacity of the knapsack is defined as $c = \varphi_j^t$. $\delta_i = 1$ implies that Ag_i will be not decommitted by agent a .

The constraint of the optimization problem can be rewritten as the exact constraint $\sum_{i=1}^N \delta_i w_i \leq c$ of the 0–1 Knapsack problem. The optimization formula can be rewritten as

$$\min \sum_{i=1}^N (1 - \delta_i) v_i = \sum_{i=1}^N v_i + \min \sum_{i=1}^N -\delta_i v_i$$

Algorithm 3 Decommit from unnecessary agreements**Input:** Tentative agreement set \mathcal{TAG}_j^t .**Output:** Tentative agreement set \mathcal{TAG}_j^t satisfying the constraint of the maximum number of final agreements.

```

1: Sort all the tentative agreements  $\mathcal{TAG}_j^t$  by decreasing ratio of  $\rho(\text{Pr}(Ag_i), \text{Tm}(Ag), t, \lambda)$  to  $1 - \omega_s^t(Ag)$ .
2: Set  $\mathcal{TAG} = \emptyset$ ,  $i = 1$ , and  $Ag$  be the  $i$ th agreement in  $\mathcal{TAG}_j^t$ .
3: while  $\sum_{Ag' \in (\mathcal{TAG} + Ag)} (1 - \omega_s^t(Ag')) \leq \varphi_j^t$  do
4:   Add  $Ag$  into  $\mathcal{TAG}$ ;
5:    $i++$ , and let  $Ag$  be the  $i$ th agreement in  $\mathcal{TAG}_j^t$ ;
6: end while
7: return  $\mathcal{TAG}$ 

```

which is equivalent to

$$\max \sum_{i=1}^N \delta_i v_i$$

Thus, the optimization problem can be formulated as a 0–1 Knapsack problem and it's \mathcal{NP} -complete.

A simple greedy approximation algorithm is used to compute the set of agreements which will not be decommitted by a (Algorithm 3) [10]: first sort all the tentative agreements \mathcal{TAG}_j^t by decreasing ratio of penalty to probability that an agreement will not be decommitted by a 's trading partners, then greedily pick agreements in this order (starting from the first agreement) until when adding a new agreement will violate the constraint of the maximum expected number of final agreements. For a removed agreement $Ag \in \mathcal{TAG}_j^t$, a decommits from the agreement; otherwise, a sends the agent s a proposal worse than $\phi_{s \rightarrow a}^t$.

4 Empirical evaluation and analysis

In this section, we first detail the methodology for analyzing the performance of the developed negotiation strategies. We then proceed to the actual empirical study of the proposed strategies. Finally, some properties of our negotiation strategies are analyzed.

4.1 The methodology

To evaluate the performance of negotiation agents, a simulation testbed consisting of a virtual e-Marketplace, a society of trading agents and a controller was implemented using JAVA. The controller generates agents, randomly determines their parameters (e.g., their roles as buyers or sellers, set of resources they provide or acquire, initial prices, reserve prices, deadlines), simulates the entrance of agents to the virtual e-Marketplace, and handles message passing and payment transfer.

4.1.1 Agent design

While there has been a lot of research in agent-mediated negotiation [20, 25, 28], most work focuses either on bilateral multi-issue negotiation (e.g., [14–16, 24–26, 48]) or single issue

one-to-many negotiation (e.g., [4, 5, 8, 29, 30, 35]). One exception is [45] which studies concurrent one-to-many negotiations for multiple resources. But in [45], an agent is assumed to know the reserve price of each resource. Given that there is no existing negotiation agents dealing with our multi-resource negotiation problem, for comparison reason, we implemented three other types of buyers based on existing techniques for single resource negotiation and negotiation with decommitment: (1) *TDA*s using a time-dependent strategy, (2) *MTDA*s using a market based time-dependent strategy, and (3) *ACMA*s using an adaptive commitment management strategy detailed in [45]. Experiments were carried out to study and compare the performance of our buyer agents (*HBAs*, heuristic-based buyer agents) with *TDA*s, *MTDA*s, and *ACMA*s.

*TDA*s, *MTDA*s and *ACMA*s adopt the strategy suggested by Nguyen and Jennings [31] and make at most one tentative agreement for each resource. *TDA*s, *MTDA*s and *ACMA*s use the same approach to determine the reserve price of each resource and use existing single resource negotiation strategies for the negotiation for each resource. The reserve price of resource I_j of each *TDA* (or *MTDA* and *ACMA*) is determined by considering the distribution of the reserve price of resource I_j . Specifically, the reserve price of resource I_j is proportional to its average reserve price. That is,

$$RP_j^t = RP^t \frac{\int_{-\infty}^{\infty} f_j(y)}{\sum_{i=1}^l \int_{-\infty}^{\infty} f_i(y)}$$

where l is the number of resources (i.e., issues) to acquire.

Similar to *HBAs*, *TDA*s, *MTDA*s and *ACMA*s generate proposals using a time-dependent negotiation decision function [13], which is widely used for designing negotiation agents (e.g., [4, 5, 13, 16, 30, 31, 41, 42, 44, 45]). However *TDA*s, *MTDA*s and *ACMA*s adopts different concession making strategies, i.e., they take different ε values. As *HBAs*, *TDA*s adopt the linear concession strategy, i.e., $\varepsilon = 1$. In contrast, *MTDA*s take market competition into account when making proposals. An *MTDA*'s parameter ε for concession making is adjusted in the following way: while the number of sellers are less than the number of buyers, an *MTDA* chooses the conciliatory concession strategy by setting $\varepsilon < 1$. Otherwise, an *MTDA* uses the conservative or linear concession strategy by setting $\varepsilon \geq 1$. *MTDA*s' adaptive concession making strategy based on market competition has been shown to make *minimally sufficient* concessions in single resource negotiation [42]. *ACMA*s use the adaptive commitment management strategy used in [45] for each single resource negotiation. Specifically, *ACMA*s use a fuzzy decision making approach for deriving adaptive commitment management strategy profiles of buyers. The value of ε of a resource is determined dynamically at each round using fuzzy rules.

Each seller agent in the market randomly chooses a negotiation strategy from the set of alternations outlined in [13]: the time-dependent function (linear, conceder, conservative) and the behavior-dependent function (e.g., tit-for-tat). Each seller agent can only make at most one tentative agreement and it will decommit from an agreement if and only if it can benefit from the decommitment.

4.1.2 Experimental settings

In the experiments, agents were subjected to different market densities, market types, deadlines, number of resources to acquire or sell, and supply/demand ratio of each resource (see Table 2). Both market density and market type depend on the probability of generating an agent in each round and the probability of the agent being a buyer (or a seller). When the num-

Table 2 Experimental settings

Input data	Possible values		
	Favorable	Balanced	Unfavorable
Market type			
Supply/demand	10:1, 5:1, 2:1	1:1	1:2, 1:5, 1:10
Market density	Sparse	Moderate	Dense
No. of agents	6–35	36–65	66–95
Deadline	Short	Moderate	Long
T_{\max}	10–30	35–55	60–80
Resources/job	Lower range	Mid-range	High range
l	1–3	4–6	7–9

ber of agents are in the range of 6–35 (respectively, 36–65 and 66–95), the market is sparse (respectively, moderate and dense). The lifespan of an agent in the e-market, i.e., its deadline, is randomly selected from [10, 80]. The range of [10, 80] for deadline was adopted based on experimental tuning and agents' behaviors. In our experimental setting, we found that: (1) for a very short deadline (<10), very few agents could complete deals, and (2) for a deadlines longer than 80, there was little or no difference in the performance of agents. Hence, for the purpose of experimentation, a deadline between the range of 10–30 (respectively, 35–55 and 60–80) 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 9, where 1–3 (respectively, 4–6 and 7–9) is considered as lower range (respectively, mid-range and upper range). The value of ε (eagerness) is randomly generated from [0.1, 8] as it was found that when $\varepsilon > 8$ (respectively, $\varepsilon < 0.1$), there was little or no difference in performance of agents.

Each resource's demand (i.e., the number of buyers who want to buy the resource) may not be equal to its supply (i.e., the number of sellers who want to sell the resource). If one buyer is negotiating for multiple resources, there are two situations: (1) All the resources have the same supply/demand ratio. From a buyer agent's perspective, for a favorable (respectively, an unfavorable) market, the supply is much higher (respectively, lower) than the demand. (2) The resources have different supply/demand ratios. Then the range and variance of resources' supply/demand ratios will affect agents' performance. All our discussions of supply/demand ratio implicitly assume that the supply/demand ratio of each resource is randomly chosen.

There are four kinds of buyers (i.e., *HBA*, *TDA*, and *MTDA*, *ACMA*) and different kinds of sellers. The number of buyers (or sellers) of each kind is decided in a random way. Without loss of generality, we assume that, there is at least one agent for each kind of agent.

4.1.3 Performance measure

We use a number of performance measures in the experiments (Table 3). Analyzing agents' utility can provide insights into how effective a strategy is. Since negotiation outcomes of each agent are uncertain (i.e., there are two possibilities: eventually reaching a consensus or not reaching a consensus), it seems more prudent to use expected utility for all runs (rather than expected utility for all successful runs) as a performance measure. For ease of analysis, agent a 's utility u_a (defined in Sect. 2.2) is normalized in each experiment in the following way: $u'_a = u_a / |RP_a - IP_a|$, which implies that $u'_a \leq 1$ if not considering the penalty a

Table 3 Performance measure

Success rate	$R_{suc} = N_{success}/N_{total}$
Expected utility	$U_{exp} = (\sum_{i=1}^{N_{total}} U_i)/N_{total}$
Agreement per resource	$AG_{aver} = \frac{\sum_{i=1}^{N_{total}} \sum_{j=1}^{IS_i} A_i^j}{\sum_{i=1}^{N_{total}} IS_i}$
Rate of recovery from decommitment	$RR_{aver} = \frac{SD_{total}}{D_{total}}$
Message per resource	$M_{aver} = \frac{\sum_{i=1}^{N_{total}} \sum_{j=1}^{IS_i} M_i^j}{\sum_{i=1}^{N_{total}} IS_i}$
N_{total}	Total number of runs
$N_{success}$	No. of runs that reached consensus
U_i	Utility of the i th run
IS_i	The number of resources in the i th run
A_i^j	The number of tentative agreement for resource j in the i th run
M_i^j	The number of messages for resource j in the i th run
D_{total}	The number of runs in which one resource’s tentative agreements were all decommitted
SD_{total}	The number of runs in which negotiation is successful after one resource’s tentative agreements were all decommitted

received from sellers. This normalization is the same for agents with different strategies. It was pointed out in [20,47] that in addition to optimizing agents’ overall utility, enhancing the success rate is also an important evaluation criterion for designing negotiation agents.

In addition to the expected utility and success rate, it’s necessary to compare the number of messages sent and received by each buyer during negotiation. As the number of resources each buyer is acquiring may be different at each time, it’s intuitive to compare the number of messages sent or accepted for each resource. As an agent may make more than one tentative agreement for each resource, measuring the average number of tentative agreements for each resource is also important. During negotiation, it’s possible that all of one agent’s tentative agreements for one resource are decommitted by its trading partners and thus an agent’s ability to recover from such situation is extremely important. Therefore, we also record and compare the number of cases where an agent makes a final agreement after all its tentative agreements for one resource are decommitted by its trading partners.

4.1.4 Results

A “matched-pair” study was conducted to evaluate the performance of *HBAs*, as compared with *TDA*s, *MTDA*s, and *ACMA*s. At the beginning of each run (experiment), the controller of the testbed will generate all the agents and set the parameters of all the agents according to the experimental setting, e.g., the number of agents, the supply/demand ratio of each resource, etc. Among all the buyers, there are some *target* buyers, one for each negotiation strategy we want to compare. All the target agents at each run have the same properties. For example, when we want to compare the performance of *HBAs* with *TDA*s, *MTDA*s, and *ACMA*s, we create one target *HBA*, one target *TDA*, one target *MTDA* and one target *ACMA*, which have the same properties (e.g., the set of resources to acquire, the reserve price, the initial price) except that they use different negotiation strategies. Then all the agents negotiate

and compete with each other. At the end of this experiment, the controller will record the experimental results for each target agent, which will be averaged and analyzed on a large number of runs.

Extensive stochastic simulations were carried out for all the combinations of market density, market type and other agents' characterizations. All the values of different performance measures were averaged based on more than 10^6 runs. In addition, we tried different decommitment deadlines and penalties functions. Even though experiments were carried out for all the situations, due to space limitations, only representative results are presented in this section. For the empirical results presented in this section, the market is of moderate density, $\lambda = 4$ is chosen as the decommitment period and the penalty function is $0.06 \times \text{Prc}(Ag) \times ((t' - t)/\lambda)^{1/2}$. $\lambda = 4$ is chosen based on the value of negotiation deadline. The shortest negotiation deadline is 10 in our experiments and setting a decommitment period shorter than negotiation deadline is reasonable. As in [1, 31], we choose a penalty function in which the penalty increases with the contract price and the period between agreement making and decommitting. The multiplier 0.06 in the penalty function is chosen to make the decommitment penalty smaller than the contracting price. In the sensitivity analysis section (Sect. 4.2.6), we discussed the effect of changing the decommitment period and the penalty function. We also found that the confidence interval for each reported value is not wider than 0.001, which is negligible as compared with each value. Therefore, the experimental results can be trusted, which is mainly due to the large sample size.

4.2 Observations

4.2.1 Observation 1

HBA agents use three heuristics: Heuristic 1 (Sect. 3.2) is used to decide the deadline for each resource; Heuristic 2 (Sect. 3.3) is used to make a proposal for each resource in which the reserve price of each resource is dynamically chosen based on current market dynamics; Heuristic 3 (Sect. 3.4) is used to decide the number of tentative agreements to be made for each resource. Is it possible that a buyer in fact can get better negotiation performance by just using one or two heuristics? To verify that agents can get better negotiation performance by using all three heuristics simultaneously, we also compare the performance of *HBA*s with a special kind of buyers (called *HBA*-s here) which only use part of the heuristics used by *HBA*s. When a *HBA* doesn't use heuristic 2, it will use *MTDA*s' strategy to make proposals. When a *HBA* doesn't use heuristic 3, it makes at most one tentative agreement for each resource. *HBA-1*s are *HBA*s which don't use heuristic 1 and *HBA-12*s are *HBA*s which don't use heuristic 1 and heuristic 2. *HBA-123*s are equivalent to *MTDA*s. Table 4 shows the performance of *TDA*s, *MTDA*s, *ACMA*s, *HBA*s, and different types of *HBA*-s which only use part of *HBA*s's three heuristics.

From column 2 of Table 4, we can find that *HBA*s gain a higher expected utility U_{exp} than agents using other strategies. We also found that *HBA*-s get higher utilities than *ACMA*s, *MTDA*s, and *TDA*s. In addition, heuristic 2 seems more important than the other two heuristics. *HBA-2*s' expected utility is lower than that of *HBA-1*s and *HBA-3*s. The average utility of *HBA*-s when *HBA*-s don't use heuristic 2 is $(0.111 + 0.135 + 0.087)/3 = 0.111$. The average utility of *HBA*-s when *HBA*-s don't use heuristic 1 is $(0.153 + 0.135 + 0.144)/3 = 0.144$. The average utility of *HBA*-s when *HBA*-s don't use heuristic 3 is $(0.144 + 0.144 + 0.087)/3 = 0.125$. Therefore, *HBA*-s will get lower utility when they don't use heuristic 2, as compared with not using either heuristic 1 or heuristic 3. In the same way, we can conclude that

Table 4 Experimental results for 10^6 runs (performance measures are defined in Table 3)

Strategy	U_{exp}	R_{suc}	AG_{aver}	RR_{aver}	M_{aver}
HBA	0.206	0.59	1.34	$\frac{478}{1356} = 0.35$	86
HBA-1	0.153	0.58	1.27	$\frac{597}{2389} = 0.25$	91
HBA-2	0.111	0.50	1.33	$\frac{835}{3134} = 0.27$	89
HBA-3	0.144	0.43	0.63	$\frac{3052}{8945} = 0.34$	88
HBA-12	0.135	0.47	1.23	$\frac{2933}{9578} = 0.31$	92
HBA-13	0.144	0.42	0.63	$\frac{2704}{9362} = 0.29$	88
HBA-23	0.087	0.35	0.59	$\frac{3171}{10489} = 0.30$	84
ACMA	0.033	0.27	0.59	$\frac{3737}{13347} = 0.28$	85
MTDA	0.021	0.25	0.57	$\frac{4423}{15584} = 0.28$	84
TDA	0.019	0.25	0.72	$\frac{9200}{33459} = 0.27$	86

heuristic 3 is more important than heuristic 1. However, the above observations are based on the averaged results in *all* scenarios and they don't suggest that the heuristic 1 is more important than the other two heuristics in every specific scenario. When the supply/demand ratio of all the resources has a large variance, the average utility of *HBA*-s when *HBA*-s don't use heuristic 1 (respectively, heuristic 2 and heuristic 3) is 0.101 (respectively, 0.107 and 0.114), which implies that heuristic 1 is more important than the other two heuristics in this specific context. For all the values in columns 2 and 3, a *t*-test analysis with confidence level 95% was carried out and the difference between every two different values for the same performance is significant.

Column 3 of Table 4 shows that *HBA*s have higher success rates R_{suc} than agents using other strategies and *HBA*-s have higher success rates than *ACMA*s, *MTDA*s, and *TDA*s. In addition, heuristic 3 is more important than the other two heuristics from the perspective of achieving a higher success rate. This observation is intuitive since without using heuristic 3, each buyer makes only one tentative agreement and its probability of making a final agreement will be low if one or more trading partner decommits from an agreement. For the same reason, from column 4 of Table 4, we can see that *HBA*s have the highest number AG_{aver} of tentative agreements for each resource. *HBA*-s using heuristic 3 have more tentative agreements than *HBA*-s not using heuristic 3 which make at most one tentative agreement for each resource.

*HBA*s' number of runs in which all tentative agreements are decommitted is lower than all other kinds of buyers (see column 5 of Table 4). The recovery rate RR_{aver} of *HBA*s is also higher than the recovery rate of other kinds of buyers. For example, the recovery rate RR_{aver} of *HBA*s is $\frac{478}{1356} = 0.35$ indicating that there were 1356 situations in which all the tentative agreements for one resource were decommitted and 476 of the situations in which the agent made a final agreement. This observation corresponds with the intuition that *HBA*s are good at organizing and balancing the multi-resource negotiation. It's not surprising that *HBA*s will send more messages during negotiation as it may make more than one tentative agreement for each resource, which is mainly due to the use of heuristic 3. However, *HBA*s' average number M_{aver} of messages transferred for each resource is less than 3% higher than that of all other kinds of agents.

4.2.2 Observation 2

Our negotiation strategy uses the estimation of sellers' probability of decommitment. The decommitment probability is an approximation of the real probability, which is unknown to the buyer. It's impossible to justify our estimated "probabilities" with theory without making strong assumptions about knowing other agents' private information. Moreover, a seller's probability of decommitting from a tentative agreement is determined by many factors, e.g., its deadline, reserve price, its negotiation situation, which is unknown to a buyer.

Here we use an empirical approach to verify the accuracy of *HBAs*' estimation of decommitment probabilities. More specifically, *HBAs*' estimation of decommitment probabilities are compared with their trading partners' *real* decommitting actions during negotiation. Assume *HBAs* made n predictions $\langle \omega_s^t(Ag_1), \dots, \omega_s^t(Ag_n) \rangle$ for tentative agreements $\langle Ag_1, \dots, Ag_n \rangle$ throughout all the experiments in which $\omega_s^t(Ag_i)$ is a *HBA*'s predicted probability that its trading partner s will decommit from the agreement Ag_i . Then *HBAs*' accuracy of predicting decommitting probabilities is given by:


$$AP = \frac{\sum_{1 \leq i \leq n} AP(\omega_s^t(Ag_i))}{n}$$

where

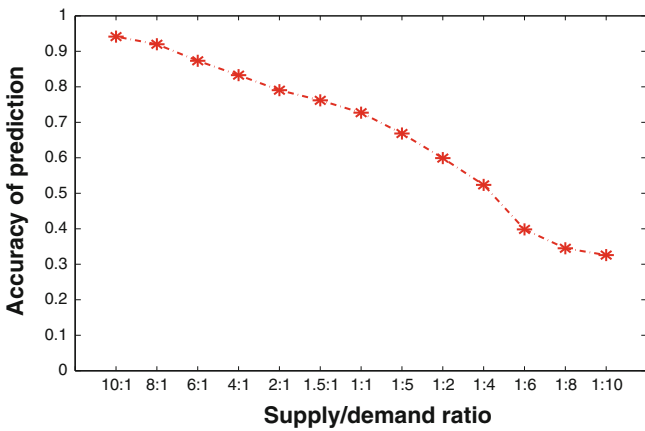
$$AP(\omega_s^t(Ag_i)) = \begin{cases} \omega_s^t(Ag_i) & \text{if } s \text{ decommits from } Ag_i \\ 1 - \omega_s^t(Ag_i) & \text{otherwise} \end{cases}$$

The average prediction accuracy in more than 10^6 runs is 0.774. Fig. 2 shows the factors affecting the prediction accuracy. First, the prediction accuracy increases with the increase of *HBAs*' deadlines (Fig. 2a). This result is intuitive as, with the increase of deadline, negotiation agents have longer time to interact with other agents. Then agents have a better understanding of the market and thus agents can make more precise predictions. Second, the prediction accuracy decreases with the increase of supply/demand ratio when all resources have the same supply/demand ratio (Fig. 2b). When the supply/demand ratio is low, *HBAs* face high pressure of competition and decommitment is more likely to happen for each tentative agreement. As a consequence, it's more difficult to make a precise prediction. Finally, the prediction accuracy changes little with the change of the number of resources (Fig. 2c). This observation is also intuitive as, a seller's decommitment decision is only affected by the agreement price, its reserve price and market competition. It has nothing to do with the negotiation status of other resources.


Our function of decommitment probability is based on our intuitions about which factors affect agents' decision to decommitment. The parameter $\vartheta = 0.68$ is a parameter of the function for computing trading partners' decommitment probabilities, which is based on experimental tuning. With the experimental tuning, we were able to get 77.4% accuracy averaged over all environments. However, it is unclear to us whether we can get a better result considering that *HBAs* do not know other agents' strategies nor their exact reserve prices. On a more positive note, our heuristic function performs in ways that would be expected. For instance, when a *HBA*'s uncertainty reduces (e.g., change the distribution of sellers' reserve prices), it gains higher prediction accuracy. A reasonable prediction approach should have the property that the prediction accuracy increases with the decrease of uncertainty which our approach does. Although we can reduce uncertainty in the market and thus get higher prediction accuracy, our experiments will become less interesting. In addition, it's impractical to assume that agents have (almost) complete information about others.



(a) Prediction accuracy and deadline



(b) Prediction accuracy and market competition



(c) Prediction accuracy and the number of resources

Fig. 2 Prediction accuracy of HBAs

Sim *et al.* [45,46] also proposed a function for evaluating a trading partner's decommitting probability, which are used by *ACMAs* and achieved an average 38% accuracy in all the scenarios. Although the function in [45] appears to be simpler as it only considers the prices of the proposals it has received, it is noted that [45] did not make the assumption that an agent has knowledge of the number of competitors. In contrast, our function takes both market competition and the trading partner's satisfaction of agreements based on each agent's knowledge about (1) the number of trading competitors and (2) the reserve price of each trading partner.

4.2.3 Observation 3


The experimental results in Fig. 3 show that: (1) Negotiation results become more favorable with the increase of the deadline for all kinds of buyers. With short (respectively, long) deadlines, different kinds of agents have equally insufficient (respectively, sufficient) time to optimize their agreements. (2) Given the same deadline, *HBAs* achieved higher utilities than *ACMAs*, *MTDAs*, and *TDA*s. (3) The advantages of *HBAs* over *MTDAs* and *TDA*s decreases when the market becomes more favorable.

Experimental results in Fig. 4 indicate that the success rate of *HBAs* are always higher than that of *ACMAs*, *MTDAs*, and *TDA*s. However, this advantage decreases when the market become more favorable. In addition, with the increase of deadline, agents' success rates have a large increase at the beginning and slightly decrease when the deadlines are long. When agents have long deadlines, agents have more time to bargain with other agents and seek good agreements with the increase of deadlines. Since agents use time-dependent strategies, buyers with longer deadlines are inclined to make less concessions at each time as agents will prefer to propose their reserve prices when their deadlines approach. Thus, buyers will become more patient and will not accept proposals which are not favorable enough while considering their future opportunities to make better agreements. Therefore, buyers with longer deadlines will fail to make agreements with some sellers, especially sellers with shorter deadlines. Although buyers' success rates decrease with the increase of deadlines when deadlines are relatively long, buyers' utilities increase with the increase of deadlines. This is because buyers will set higher expectation about the agreements with the increase of deadlines. Thus, the agreements made by buyers with longer deadlines are more favorable as compared with agreements made by buyers with shorter deadlines.


4.2.4 Observation 4

From Fig. 5 we can see that, as the number of resources to be acquired increases, the utilities of all kinds of agents decrease. That is because, with the increase of the number of resources each agent acquires, it's harder to manage all the negotiations and the probability that the overall negotiation fails increases, which directly correlates with the decreased success rates in the strategies explored here. *HBAs* always achieved higher utilities than *ACMAs*, *MTDAs*, and *TDA*s.

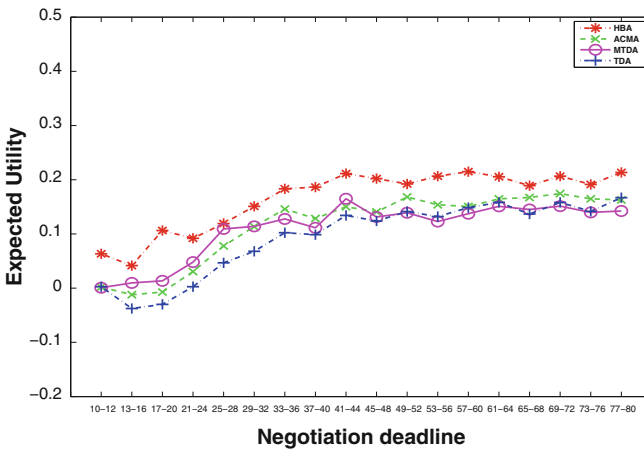
Experimental results in Fig. 6 indicate that the success rate of *HBAs* are always higher than that of *ACMAs*, *MTDAs*, and *TDA*s. However, this advantage decreases when agents have longer deadlines as in this case, all agents have enough time to negotiate for agreements. Agents' success rate decreases significantly as a small number of resources (e.g., 1 or 2). With more resources, it's more difficult for buyers to manage and establish agreements for all resources because of the difficulties of managing all the negotiation threads.



(a) Unfavorable market




(b) Balanced market

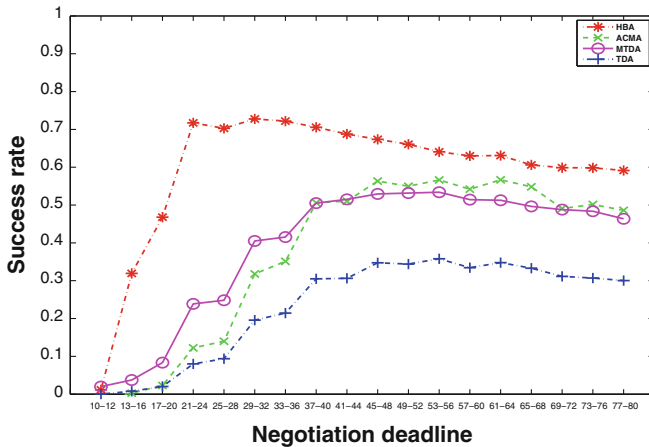


(c) Favorable market

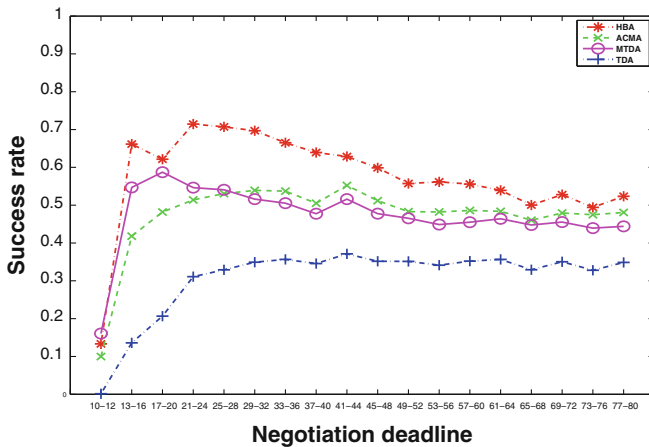
Fig. 3 Deadline and expected utility



(a) Unfavorable market

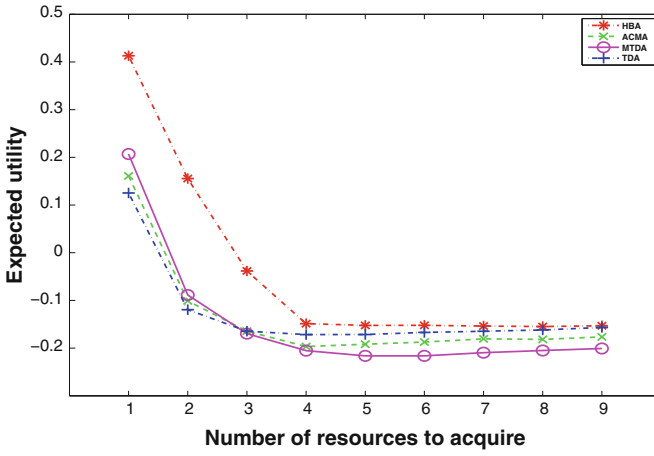


(b) Balanced market




(c) Favorable market


Fig. 4 Deadline and success rate



(a) Short deadline




(b) Moderate deadline




(c) Long deadline


Fig. 5 Number of resources and expected utility



(a) Short deadline



(b) Moderate deadline



(c) Long deadline

Fig. 6 Number of resources and success rate

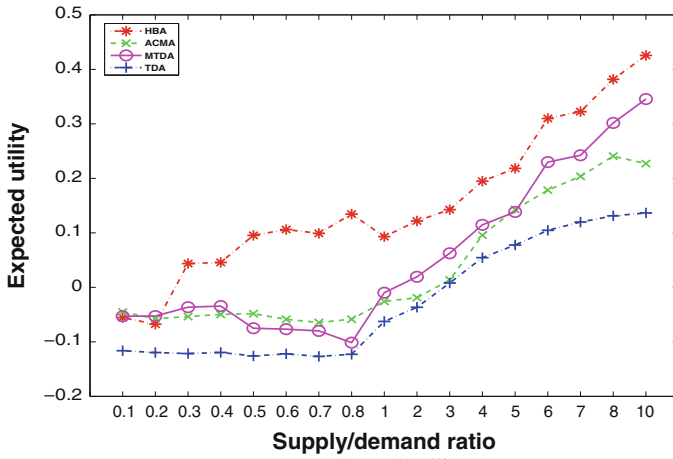
4.2.5 Observation 5

It can be observed from Figs. 7 and 8 that *HBAs* always get higher utilities (respectively, success rates) than *ACMAs*, *MTDAs*, and *TDA*s when all resources have the same supply/demand ratios. Additionally, when the supply/demand ratio is high (e.g., 10), the average utilities of the three types of agents are close especially in the long deadline case since agents have many choices and can easily switch from one agreement to another agreement, i.e., there is limited space to optimize the agreements. The advantage of *HBAs* in success rate decreases when agents have longer deadlines. Since buyer agents with different strategies compete with each other, it is possible that one strategy achieved much better negotiation results than another strategy in a specific market. Due the strategic interaction among agents, one strategy may achieve a good performance in only certain markets. In Fig. 7 we can see that when the ratio is in the range 0.5–0.7, *MTDAs* achieved very low utilities as compared with the utilities when the ratio less than 0.5 or higher than 0.7. When the ratio is in the range 0.5–0.7, *HBAs* achieved higher utility than that when the ratio less than 0.5 or higher than 0.7. When the supply/demand ratio is very low (e.g., 0.2–0.4), it is difficult for an agent to get agreements, thus all different strategies achieved low utilities. When the supply/demand ratio are slightly low (e.g., 0.5–0.7), some *HBAs* may make agreements for all required resources. An *MTDA* can also make agreements for some of its resources using its market-driven concession strategy. However, since *MTDAs* are lacking of the ability of coordinating their negotiation for multiple resources. They often can only satisfy part of their resources. Therefore, when the whole negotiation fails, an *MTDA* either pays a lot of penalties to decommit from its agreements or pays for some final agreements which have not been decommitted. Accordingly, *MTDAs* often get negative utilities. When *MTDAs* decommit from agreements, *HBAs* have a better chance to make new agreements in this situation. The experimental results also show that when the supply/demand is in the range 0.5–0.7, *MTDAs* made more agreements (including both tentative and final) than *TDA*s and *ACMAs* but the success rate of *MTDAs* is not higher than that of *TDA*s and *ACMAs*. When the market is almost balanced (e.g., the supply/demand is in the range of 0.8–1), it is easier for *MTDAs* to make agreements which can satisfy their resource requirements and their utilities are much higher than that when the supply/demand is in the range 0.5–0.7.


4.2.6 Sensitivity analysis

We also did additional experiments to explore how sensitive are our experimental results to changes of the parameters of our experimental environments or assumptions about our negotiation model.


- (1) With the increase of penalty, the average utility of agents including *HBAs* decreases. For example, when we double the penalty fee, the average utility of *HBAs* is decreased by 7%. The main reason is that with a higher penalty, a buyer is more likely to commit to an early agreement, which may have a low utility value. When the penalty fee is low, a buyer will decommit from an early agreement and make a new agreement with a higher utility value. Similarly, each seller is also more likely to stick to an early agreement when the penalty is high. *HBAs* always have better performance than other types of buyers when using different penalty functions.
- (2) With the increase of decommitment period λ , the average utility of agents including *HBAs* decreases. For instance, when we set a decommitment period $\lambda = 6$ instead of 4, the average utility of *HBAs* is decreased by 8%. With a longer decommitment period, the



(a) Short deadline




(b) Moderate deadline

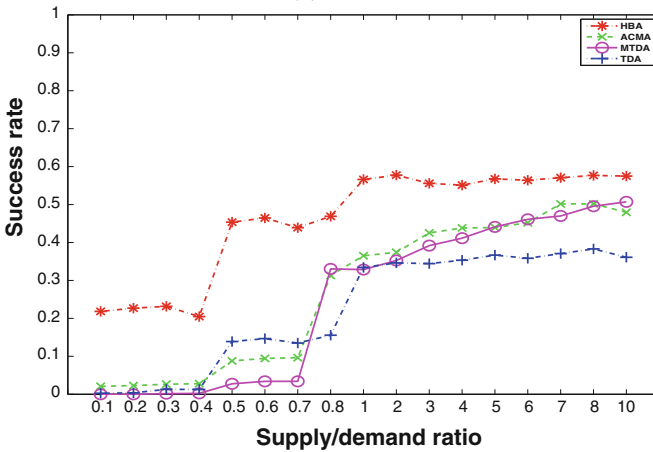


(c) Long deadline


Fig. 7 Supply/demand ratio and expected utility



(a) Short deadline



(b) Moderate deadline



(c) Long deadline

Fig. 8 Supply/demand ratio and success rate

- probability that an agreement will be decommitted will increase, and thus the probability that a buyer will get a final agreement decreases. However, the advantage of *HBAs* over other types of buyers increases with the increase of the decommitment period λ as buyers like *ACMAs*, *MTDAs*, and *TDA*s make at most one tentative agreement for each resource.
- (3) When agents have more accurate information about other agents, agents including *HBAs* achieved better performance. This paper assume that a buyer knows the probability distribution of sellers' reserve prices and the number of competitors. We find that that the accuracy of this information does have an effect on agents' negotiation performance. When a buyer's knowledge becomes less accurate, its utility decreases. For example, when the believed number of competitors is less than half of the actual number of competitors, the average utility of *HBAs* is 7% lower than that of *HBAs* knowing the actual number of competitors. However, even with this level incorrect information, *HBAs* still achieved better performance than other types of agents.
 - (4) While keeping the supply/demand ratio of each resource constant, market density has little effect on agents' performance. In a moderate density market, agents' average utilities are 2% lower than that in a market of dense density and are 1% higher than that in a market of sparse density.

4.3 Analysis of properties

Typically, agents use a *monotonic* concession protocol by insisting on their previous proposals or raising/reducing their proposals monotonically until an agreement is reached. In a dynamic negotiation environment, market competition and agents' evaluation may change over time, protocols that are not monotonic may achieve higher average utilities. Negotiation agents in this paper make a proposal based on market situation and the negotiation situations of other threads. Therefore, the proposed negotiation protocol is not monotonic.

In a favorable market, there are fewer competitors and more trading partners. Hence, an agent has stronger bargaining power and doesn't need to make large concessions. In an unfavorable market, an agent experiences more competition, and it may attempt to make more concessions. With respect to competition, an agent strives to avoid making large concessions in favorable markets or making too large concessions in unfavorable markets. Additionally, when the expected number of final agreements is high, an agent is inclined to make less concession as it only needs one final agreement.

Property 1 *Agents will make less concession with the increase of the expected number of final agreements when the worst possible utility doesn't increase.*

Take the resource I_j for example. The number of agreements has no effect on the expected agreement price ϖ_j^t . As the worst possible utility doesn't increase, the conflict probability χ_j^t will not increase. $\gamma(\mathcal{T}AG_j^t)$ will decrease with the expected number of final agreements $\varphi(\mathcal{T}AG_j^t)$. Therefore, the reserve price of resource I_j will decrease and thus agents will make less concession.

Property 2 *Agents will make less (respectively, more) concession with the increase of the number of trading partners (respectively, competitors).*

Take resource I_j for example. The number of trading partners has no effect on $\gamma(\mathcal{T}AG_j^t)$. With the increase of trading partners, χ_j^t will not increase and ϖ_j^t will also not increase.

Thus, the reserve price of resource I_j will not increase and thus agents make less concessions. Similarly, with the increase of trading competitors, ϖ_j^t will decrease. Thus, the reserve price of resource I_j will increase and thus agents will make more concession.

Property 3 *When competition is high and penalty is very low, agents may make agreements with all the trading partners.*

Take resource I_j for example. The decommitment probability increases with the increase of competition. As the penalty is very low, an agent with more tentative agreements won't pay too much penalty when it has to decommit from some tentative agreements. An extreme situation is that the agent can even make agreements with all the trading partners.

From Properties 2 and 3 we can learn that the market competition places an important role on deciding the amount of concessions and the number of tentative agreements. With respect to competition, a negotiation agent decides the maximum number of agreements. In a favorable market, there are fewer competitors and more trading partners. Hence, an agent doesn't need to make many agreements (concessions, respectively). In an unfavorable market, an agent's bargaining power decreases as it experiences more competition, and it may attempt to make more agreements (concessions, respectively) as its trading partners are more likely to decommit from agreements.

One possible strategy is to make agreements later and thus potentially a buyer will pay less decommitment penalties given that the penalty will increase with time. However, "delaying" agreements will also increase the probability that the whole negotiation fails. In addition, generally a buyer will increase its offering price gradually and it is possible that it can get some resources with a cheap price in the early negotiation stages. While taking the "delaying" strategy, the buyer will miss those cheap resources and buy expensive resources in a later time. Another disadvantage of delaying agreements is that the buyer may fail to get all resources when one seller decommits from agreements when the deadline is approaching. In our model, no agent can decommit from an agreement after a fixed time period based on when the agreement was made. Accordingly, making agreements earlier can potentially avoid negotiation's "collapsing" at the last minute. We examined agents' performance when they choose delaying agreements and found that such strategic "delaying" do not improve agents' performance.

As a result of this extensive empirical analysis, we have verified that the negotiation strategy for multi-resource acquisitions is both very effective in comparison to existing approaches and behaves in a consistent and appropriate manner as important characteristics of the marketplace are varied.

5 Related work

Automated negotiation is an important research area encompassing economics, game theory, computer science, and artificial intelligence, and has widely applied in many domains like electronic commerce, grid computation, and service composition. The literature of automated negotiation and negotiation agents forms a very large collection and space limitations preclude introducing all of them here. For a survey on negotiation agents, see [20], [25], and [28], respectively. The rest of this section discusses related work on multi-issue negotiation, one-to-many and many-to-many negotiation, negotiation strategy, organizational negotiation, and decommitment.

Multi-issue negotiation: There are two different definitions of a negotiation *issue* in the literature. In papers like [14,24,26], an issue is an attribute (e.g., price, quality,

delivery time) of a resource. In this case, multi-issue negotiation is bilateral. An issue can also be treated as a resource as in [16,45,46] and in this case, a buyer can negotiate with multiple sellers for each resource. If a seller has multiple resources, such multi-resource negotiation could be bilateral and each resource can be treated as an attribute. Multi-issue negotiation is more complex and challenging than single-issue negotiation as the solution space is multi-dimensional and it's often difficult to reach a *Pareto-efficient* solution [25]. Almost all the work on multi-issue negotiation focuses on bilateral negotiation and a variety of learning and searching methods are used, e.g., case-based reasoning [48], similarity criteria based search [14], decentralized search [24,26]. There usually exist different types of negotiation procedures [16] like *package deal*, *simultaneous negotiation*, and *sequential negotiation*. For sequential negotiation, agents need to decide a negotiation *agenda* (order of negotiation issues) [15]. Fatima *et al.* [16] studied different procedures for bilateral multi-issue negotiation and show that the package deal is the optimal procedure. Different from related work on bilateral multi-issue negotiation, this paper studies multi-resource negotiation where resources are provided by multiple agents and thus an agent is negotiating with multiple trading partners.

One-to-many and many-to-many negotiation: In many situations, an agent has an opportunity to make an agreement with more than one trading partners. An agent may also face the competition from agents of the same type, e.g., a buyer in negotiation faces competition from other buyers. Even if an agent interacts with many agents, an agent can pursue only one negotiation at a time in some models. An agent has to terminate a current negotiation in disagreement first, and then pursue a more attractive outside alternative. This kind of model is called bilateral negotiation with *outside options* [27,33]. However, the presumption that an agent can pursue only one negotiation at a time appears to be restrictive. In one-to-many negotiation [4,5,8,29–31,35,45,46], an agent can concurrently negotiate with multiple trading partners and an agent's proposal to one trading partner is affected by the status of its negotiation with other trading partners. In this paper, each agent concurrently negotiates with multiple trading partners for multiple resources and an agent's proposals to each trading partner depends on the negotiation with all the trading partners.

Concurrent negotiations: Sim *et al.* [45,46] proposed a coordination strategy for multi-resource negotiation where an agent can negotiate with multiple agents as in this paper. Each buyer in [45,46] knows the reserve price of each resource in advance and the buyer just needs to decide the concession strategy for each one-to-many negotiation for one resource; however, it did not assume that consumer agents know the number of competing consumers. In contrast, each buyer in this work is assumed to only know the value of its high level task, i.e., the reserve price of all resources required for the high level task. This paper proposes a set of heuristics for dynamically determining the reserve price of each resource based on the status of all negotiations. Furthermore, a buyer in [45,46] only makes one tentative agreement but in this paper, a buyer can make more than one tentative agreement.

Negotiation strategy: There has been a long history in game theory literature of developing techniques to find Nash equilibria and its refinements, e.g. sequential equilibria [23] for different bargaining games. However, due to the high analytic complexity of equilibrium analysis and lack of appropriate information, bounded rational agents instead are often constrained to play predefined tactics. Kraus *et al.* [22] proposed a strategic model in which the passage of time was taken into account. It has been shown that if agents use sequential equilibrium strategies, negotiation will end rapidly. To build more flexible and sophisticated negotiation agents, Faratin *et al.* [13] devised a negotiation model that defines a range of Negotiation Decision Functions (*NDFs*) for generating (counter-)proposals based on time,

resource, and behaviors of negotiators. Sim *et al.* [41,44] consider other factors, such as competition, trading alternatives, and differences of negotiators, and propose *market-driven agents* (MDAs) which can make minimally sufficient concessions. Game theoretical analysis [42] shows that the strategies of MDAs are in sequential equilibrium and market equilibrium for bilateral and multilateral negotiations. Like MDAs, negotiation agents in this paper make negotiation decisions that take into account market dynamics and negotiation status of all negotiation threads for all resources.

Organizational negotiation: Zhang *et al.* have studied a number of sophisticated negotiation problems in organizational contexts [49,50]. Automated negotiation becomes increasingly complex and difficult as (1) agents are large-grained and complex with multiple goals and tasks, (2) agents often have more negotiation tasks and organizational relationships among heterogeneous agents, (3) negotiation process is tightly interleaved with agents' negotiation, scheduling and planning processes. Zhang *et al.* [49,50] focus more on the coordination (a good "fit") of multiple negotiation tasks in organization context and they do not address agents' bargaining strategy in complex negotiation environments. In contrast, our work investigates how agents make concessions in dynamic negotiation environments where agents have multiple resources to negotiate.

Leveled commitment contracts: Sandholm *et al.* propose leveled-commitment contracts [40] in which the level of commitment is set by decommitting penalties. However, they only study the two-player game and they don't investigate agents' bargaining strategies with decommitment from agreements. In addition, the problem setting in [40] is far from the real-world settings. In the negotiation management system for CLASP [1], resource consumers can decommit from agreements made before at the cost of paying a penalty. However, the focus in [1] is only on the scheduling problem. This paper focuses on agents' negotiation strategies given that agents can decommit from agreements. Nguyen and Jennings [30,31] provide and evaluate a commitment model for concurrent negotiation. However, the maximum number of tentative agreements is determined prior to negotiation. In our work, the maximum number of tentative agreements is determined by the current market situation and will change dynamically during negotiation. In addition, our work studies a multi-resource negotiation problem, rather than single resource negotiation as in [30,31]. Furthermore, Nguyen and Jennings [30,31] make very restrictive assumptions about agents' available information, e.g., each agent is assumed to have knowledge about (1) other agents' negotiation strategies, (2) its negotiation success rate when it adopts certain strategy, and (3) its payoff when it adopts certain strategy. In this work, we assume that each agent has no knowledge about negotiation outcomes.

Combinatorial auctions: In combinatorial auctions [9,32], a large number of items are auctioned concurrently and bidders are allowed to express preferences on bundles of items. In contrast, in combinatorial reverse auctions, a buyer is to buy goods or services from many competing sellers. Combinatorial reverse auction [39] has some similarities to the problem studied in this paper. One difference is that we assume the agents negotiate over price of a single resource in which the buyer also submits proposals to sellers, but in combinatorial reverse auctions, only sellers submit bids to the buyer and the buyer determines the winning bids. While there is two-sided competition, market mechanisms like double auction can be used for resource allocation. The double auction is one of the most common exchange institutions where both sellers and buyers submit bids which are then ranked highest to lowest to generate demand and supply profiles. Some double auction mechanisms (e.g., BBDA [17]) have been applied to trading in markets. As mentioned in Sect. 1, auction mechanisms have some limitations (e.g., computational intractability, trustworthiness, significant computational overload). Furthermore, for the dynamic resource allocation problem, it is

very difficult for the auctioneer to decide when to run auctions. In our distributed approach, an agent can negotiate with other agents when needed. Our distributed negotiation approach seems more natural, more robust, and can accommodate decommitment.

6 Conclusions

This paper presents the design and implementation of negotiation agents that negotiate for multiple resources where agents don't know the reserve price of each resource and are allowed to decommit from existing agreements. The contributions of this paper include: (1) To avoid the risk of the "collapse" of the overall negotiation due to failing to acquire some scarce resources, negotiation agents have the flexibility to adjust the deadline for different resources based on market competition, which allows agents to respond to uncertainties in resource planning. (2) Each agent utilizes a time-dependent strategy in which the reserve price of each resource 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 resource and the maximum number of agreements is constrained by the market situation. (4) An extensive set of experiments were carried out and the experimental results show that each of the proposed heuristics contributes to improve agents' performance and our proposed approach achieved better negotiation results than representative samples of existing negotiation strategies. Additionally, our heuristics always perform in a consistent and appropriate manner in different markets.

The experimental results showed that negotiation agents with our negotiation strategy, i.e., *HBAs*, achieved better negotiation results (higher expected utilities and higher success rates) than *ACMAs*, *MTDAs*, and *TDA*s which are based on existing approaches for single resource negotiation in the literature. Moreover, it's better for *HBAs* to use all the three heuristics together as each heuristic has different features. The heuristic for proposal creation seems more important than the other two heuristics. From our experimental results we can see that, when the negotiation environment is either very "tough" (i.e., short deadline, high competition, and more resource to negotiate) or very "favorable" (i.e., long deadline, less competition, and less resource to negotiate), *HBAs* did not significantly outperform *MTDAs* and *TDA*s. That is because in a "tough" market, all the agents have little opportunity for making individual agreements, and thus it's very hard to find a good set of agreements that satisfy all the resource requirements. In contrast, in a very "favorable" market, agents can easily make good agreement set. It is in the middle ground that you see the significant advantage of our approach.

Finally, future research directions of this work include: (1) This paper assumes that a buyer gains nothing if it fails to make agreements for all the resources, which can be relaxed so that the buyer gets some utility for the agreements for part of the resources. In addition, the negotiation problem will become more complex as we consider interdependencies between resources [49, 50]. (2) At the present stage, the decommitment penalty is determined prior to negotiation. In the future work, we will treat the penalty as a first class attribute when agents can negotiate over price and decommitment penalty. (3) While this paper assumes that agents are selfish, it would be interesting to investigate agents' negotiation strategies for multi-resource negotiation in cooperative (or semi-cooperative) environments (e.g., cooperative sensor networks [21]) in which agents are optimizing some system-level objectives (e.g., social welfare). (4) Another future research topic is to develop new complex negotiation approaches for the formation of automated virtual agent enterprises (VAE), which is

formed to meet a specific objective or to provide a specific service. Achieving this objective or service involves performing a series of tasks that require repeated negotiations among VAE members. Thus, designing effective negotiation mechanisms is crucial to the formation and operation of the VAE.

Acknowledgements This work was supported primarily by the Engineering Research Centers Program of the National Science Foundation under NSF Cooperative Agreement No. EEC-0313747. Any Opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect those of the National Science Foundation. Financial support for Kwang Mong Sim for this work was supported by the Korea Research Foundation Grant funded by the Korean Government (MEST) (KRF-2009-220-D00092). A preliminary report of this work appeared in *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2008, pp.1553–1556. This article contains additional analysis and extensive experimental results not present in the conference version.

References

1. An, B., Douglis, F., & Ye, F. (2008). Heuristics for negotiation schedules in multi-plan optimization. In *Proceedings of the seventh international joint conference on autonomous agents and multi-agent systems* (pp. 551–558).
2. An, B., Gatti, N., & Lesser, V. (2009, September). Extending alternating-offers bargaining in one-to-many and many-to-many settings. In *Proceeding of the 9th IEEE/WIC/ACM international conference on intelligent agent technology (IAT)* (pp. 423–426).
3. An, B., Lesser, V., Irwin, D., & Zink, M. (2010). Automated negotiation with decommitment for dynamic resource allocation in cloud computing. In *Proceedings of the ninth international joint conference on autonomous agents and multi-agent systems* (pp. 981–988).
4. An, B., Sim, K. M., & Lesser, V. (2007). Evolving the best-response strategy to decide when to make a proposal. In *Proceedings of the 2007 IEEE congress on evolutionary computation* (pp. 1035–1042).
5. An, B., Sim, K. M., Tang, L., Li, S., & Cheng, D. (2006). A continuous time negotiation mechanism for software agents. *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, 36(6), 1261–1272.
6. Bernstein, D. (2005). *Complexity Analysis and Optimal Algorithms for Decentralized Decision Making*. PhD thesis, University of Massachusetts, Amherst.
7. Branson, M., Douglis, F., Fawcett, B., Liu, Z., Riabov, A., & Ye, F. (2007). CLASP: Collaborating, autonomous data stream processing systems. In *Proceedings of the ACM/IFIP/USENIX 8th international middleware conference* (pp. 348–367).
8. Brazier, F., Cornelissen, F., Gustavsson, R., Jonker, C., Lindeberg, O., Polak, B., & Treur, J. (2002). A multi-agent system performing one-to-many negotiation for load balancing of electricity use. *Electronic Commerce Research and Applications*, 1(2), 208–224.
9. Cramton, P., Shoham, Y., & Steinberg, R. (Eds.). (2006). *Combinatorial Auctions*. Cambridge: MIT Press.
10. Dantzig, G. (1957). Discrete variable extremum problems. *Operations Research*, 5, 266–277.
11. Dunne, P. E. (2005). Extremal behaviour in multiagent contract negotiation. *Journal of Artificial Intelligence Research*, 23, 41–78.
12. Endriss, U., & Maudet, N. (2005). On the communication complexity of multilateral trading: Extended report. *Journal of Autonomous Agents and Multiagent Systems*, 11(1), 91–107.
13. Faratin, P., Sierra, C., & Jennings, N. R. (1998). Negotiation decision functions for autonomous agents. *International Journal of Robotics and Autonomous Systems*, 24(3–4), 159–182.
14. Faratin, P., Sierra, C., & Jennings, N. R. (2002). Using similarity criteria to make issue trade-offs in automated negotiations. *Artificial Intelligence*, 142(2), 205–237.
15. Fatima, S. S., Wooldridge, M., & Jennings, N. R. (2004). An agenda-based framework for multi-issue negotiation. *Artificial Intelligence*, 152(1), 1–45.
16. Fatima, S. S., Wooldridge, M., & Jennings, N. R. (2006). Multi-issue negotiation with deadlines. *Journal of Artificial Intelligence Research*, 27, 381–417.
17. Fudenberg, D., Mobius, M., & Szeidl, A. (2007). Existence of equilibrium in large double auctions. *Journal of Economic Theory*, 127(1), 550–567.
18. Gatti, N., Giunta, F. D., & Marino, S. (2008). Alternating-offers bargaining with one-sided uncertain deadlines: An efficient algorithm. *Artificial Intelligence*, 172(8–9), 1119–1157.

19. Jain, N., Amini, L., Andrade, H., King, R., Park, Y., Selo, P., & Venkatramani, C. (2006). Design, implementation, and evaluation of the linear road benchmark on the stream processing core. In *Proceeding of the 25th ACM SIGMOD international conference on management of data* (pp. 431–442).
20. Jennings, N. R., Faratin, P., Lomuscio, A. R., Parsons, S., Sierra, C., & Wooldridge, M. (2001). Automated negotiation: Prospects, methods and challenges. *International Journal of Group Decision and Negotiation*, 10(2), 199–215.
21. Krainin, M., An, B., & Lesser, V. (2007, November) An application of automated negotiation to distributed task allocation. In *Proceedings of the 2007 IEEE/WIC/ACM international conference on intelligent agent technology* (pp. 138–145).
22. Kraus, S., Wilkenfeld, J., & Zlotkin, G. (1995). Multiagent negotiation under time constraints. *Artificial Intelligence*, 75(2), 297–345.
23. Kreps, D., & Wilson, R. (1982). Sequential equilibria. *Econometrica*, 50(4), 863–894.
24. Lai, G., Li, C., & Sycara, K. (2006). Efficient multi-attribute negotiation with incomplete information. *Group Decision and Negotiation*, 15, 511–528.
25. Lai, G., Li, C., Sycara, K., & Giampapa, J. (2004, December) Literature review on multi-attribute negotiations. Technical Report CMU-RI-TR-04-66, Robotics Institute, Carnegie Mellon University.
26. Lai, G., Sycara, K., & Li, C. (2008). A decentralized model for automated multi-attribute negotiations with incomplete information and general utility functions. *Multiagent and Grid Systems*, 4(1), 45–65.
27. Li, C., Giampapa, J. A., & Sycara, K. (2006). Bilateral negotiation decisions with uncertain dynamic outside options. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Special Issue on Game-Theoretic Analysis and Stochastic Simulation of Negotiation Agents*, 36(1), 31–44.
28. Lomuscio, A. R., Wooldridge, M., & Jennings, N. R. (2003). A classification scheme for negotiation in electronic commerce. *International Journal Group Decision Negotiation*, 12(1), 31–56.
29. Nguyen, T. D., & Jennings, N. R. (2003). A heuristic model of concurrent bi-lateral negotiations in incomplete information settings. In *Proceedings of the eighteenth international joint conference on artificial intelligence* (pp. 1467–1469).
30. Nguyen, T. D., & Jennings, N. R. (2004). Coordinating multiple concurrent negotiations. In *Proceedings of the 3rd international conference on autonomous agents and multi-agent systems* (pp. 1064–1071).
31. Nguyen, T. D., & Jennings, N. R. (2005). Managing commitments in multiple concurrent negotiations. *Electronic Commerce Research and Applications*, 4(4), 362–376.
32. Nisan, N., Roughgarden, T., Tardos, E., & Vazirani, V. V. (Eds.). (2007). *Algorithmic Game Theory*. Cambridge: Cambridge University Press.
33. Osborne, M. J., & Rubinstein, A. (1990). *Bargaining and markets*. New York: Academic.
34. Preist, C., Bartolini, C., & Bye, A. (2003). Agent-based service composition through simultaneous negotiation in forward and reverse auctions. In *Proceedings of the 4th ACM conference on electronic commerce* (pp. 55–63).
35. Rahwan, I., Kowalczyk, R., & Pham, H. H. (2002). Intelligent agents for automated one-to-many e-commerce negotiation. *Australian Computer Science Communications*, 24(1), 197–204.
36. Riabov, A., & Liu, Z. (2005). Planning for stream processing systems. In *Proceedings of the twentieth national conference on artificial intelligence* (pp. 1205–1210).
37. Rosenschein, J., & Zlotkin, G. (1994). *Rules of Encounter: Designing conventions for automated negotiation among computers*. Cambridge, MA: MIT Press.
38. Rubinstein, A. (1982). Perfect equilibrium in a bargaining model. *Econometrica*, 50(1), 97–109.
39. Sandholm, T. (2007). Expressive commerce and its application to sourcing: how we conducted \$35 billion of generalized combinatorial auctions. *AI Magazine*, 28(3), 45–58.
40. Sandholm, T., & Lesser, V. (2001). Leveled commitment contracts and strategic breach. *Games and Economic Behavior*, 35(1–2), 212–270.
41. Sim, K. M. (2002). A market-driven model for designing negotiation agents. *Computational Intelligence, Special Issue in Agent Technology for E-Commerce*, 18(4), 618–637.
42. Sim, K. M. (2005). Equilibria, prudent compromises, and the “waiting” game. *IEEE Transaction on Systems, Man and Cybernetics, Part B*, 33(3), 712–724.
43. Sim, K. M. (2006). G-commerce, market-driven g-negotiation agents and grid resource management. *IEEE Transactions on Systems, Man and Cybernetics, Part B*, 36(6), 1381–1394.
44. Sim, K. M., & Choi, C. Y. (2003). Agents that react to changing market situations. *IEEE Transaction on Systems, Man and Cybernetics, Part B*, 33(2), 188–201.
45. Sim, K. M., & Shi, B. (2008). Adaptive commitment management strategy profiles for concurrent negotiations. In *Proceedings of the first international workshop on agent-based complex automated negotiations* (pp. 16–23).
46. Sim, K. M., & Shi, B. (2010). Concurrent negotiation and coordination for grid resource co-allocation. *IEEE Transaction on Systems, Man and Cybernetics, Part B*, 40(3), 753–766.

47. Sim, K. M., & Wang, S. (2004). Flexible negotiation agent with relaxed decision rules. *IEEE Transaction on Systems, Man and Cybernetics, Part B*, 34(3), 1602–1608.
48. Sycara, K. (1991). Problem restructuring in negotiation. *Management Science*, 37(10), 1248–1268.
49. Zhang, X., Lesser, V., & Abdallah, S. (2005). Efficient management of multi-linked negotiation based on a formalized model. *Journal of Autonomous Agents and MultiAgent Systems*, 10(2), 165–205.
50. Zhang, X., Lesser, V., & Wagner, T. (2006). Integrative negotiation among agents situated in organizations. *IEEE Transactions on System, Man, and Cybernetics: Part C*, 36(1), 19–30.