

Concurrent bi-lateral negotiation in agent systems

Thuc Duong Nguyen and Nicholas R. Jennings
Department of Electronics and Computer Science
University of Southampton, UK.
Email: {tdn01r, nrj}@ecs.soton.ac.uk*

Abstract

Bi-lateral negotiations represent an important class of encounter in agent-based systems. To this end, this paper develops and evaluates a heuristic model that enables an agent to participate in multiple, concurrent bi-lateral encounters in competitive situations in which there is information uncertainty and deadlines.

1 Introduction

Automated negotiation is a key form of interaction in agent-based systems. Given its wide variety of objectives, such negotiations exist in many different forms including one-to-one, one-to-many and many-to-many [4]. Generally speaking, however, the latter two cases are dealt with using some form of auction protocol (be it single-sided or double-sided, respectively), while the former are often tackled using some form of heuristic method (see section 4 for more details). Here we focus on the one-to-one case, in which one agent wants to purchase a service¹ from another. Moreover, we consider competitive situations in which the agents have *no a priori knowledge about the preferences of their opponents*. In such cases, the agents exchange proposals, representing acceptable solutions, until either an agreement is reached or the negotiation terminates with a failure.

To date, one of the inherent drawbacks of bi-lateral negotiation models is that the agent has to a priori identify a single partner to interact with. While this is acceptable if there is only one provider of the desired service, it is inefficient in an uncertain setting if there are multiple providers of the service that each have different characteristics. In this multiple provider case, there are two alternatives: (1) negotiate sequentially with all the providers or (2) negotiate concurrently with them. The former has the disadvantage that it may result in lengthy negotiation encounters, but has

the advantage that it is comparatively easy to use the outcome of one negotiation to dictate behavior in subsequent ones. The latter case has the advantage of taking less time, but the disadvantage that coordinating behaviors among the various negotiation threads is more difficult². Since we are interested in situations in time-constrained domains (such as e-commerce and grid computing), we concentrate on the concurrent case and develop a coordinated bidding model in which the various negotiation threads *mutually influence* one another [5]. By mutually influence, we mean that the progress and agreement in one negotiation thread is used to alter the behavior of the agent in another thread for the same service. For example, having obtained a good deal in one thread, the agent may adopt a tougher stance in its other threads, to see if it can get an even better deal than the one it already has³.

This paper advances the state of the art in the following ways. Firstly, most existing models of bi-lateral negotiation do not deal with concurrent encounters. Secondly, those models that do provide facilities for dealing with the concurrent case (eg [6]) tend to adopt rigid strategies and coordination mechanisms that are not generally applicable (see section 4 for more details). In contrast, our model allows the agent to apply tactics and strategies without restriction. Finally, we systematically evaluate our model to show its effectiveness in a variety of negotiation situations and to highlight the relative advantages and disadvantages of such concurrent mechanisms.

The remainder of the paper is structured in the following manner: section 2 details the model and section 3 evaluates it. Section 4 relates the model to current work in the field and, finally, section 5 presents the conclusions.

²Such coordination is necessary to ensure the agent does not end up with multiple agreements for the service when only one is required.

³This model differs from a one-to-many auction in that it allows *direct* interaction between the agent requiring the service and the providers offering it. This ability to exchange unmediated counter-offers enables the participants to indicate their preferences and constraints directly to one another.

*This work is sponsored by BT.

¹A service is here viewed as an abstract representation of the capability of an agent.

2 The concurrent negotiation model

The agent that wishes to purchase the service is called the *buyer* and the agents that offer the service are called the *sellers*. Service agreements (contracts) are assumed to be multi-dimensional (covering issues such as *price*, *quality*, *quantity*, etc.). The buyer has a deadline t_{max} by when it must conclude its negotiations for the service. Similarly, each seller i has its own negotiation deadline $t_{i,max}$. All the agents have their own preferences about the service and this information is private (as are the strategies the individual agents follow). The agents follow an alternating sequential protocol [3], in which the illocutions are *offer* (a proposal made by one agent to the other), *counter-offer* (a counter proposal from an agent in response to a proposal it received), *accept* (accept a proposed offer), *finalize* (secure a deal with the chosen seller), *decline* (reject the previously accepted offer) and *withdraw* (terminate the negotiation thread). The difference between an *accept* and *finalize* is necessary in this work to deal with the problem of concurrent encounters. If the buyer accepts an offer from a seller then this is viewed as binding on the seller (for a specified period of time that is assumed to be longer than t_{max}). However, it is not binding on the buyer. Thus, the buyer may accept several offers from multiple sellers in any one negotiation episode. However, when it has completed all the negotiations, the buyer will *finalize* one of the accepted deals with one of the sellers and *decline* the others (thus freeing them from their commitment to the proposal). This two phase process is necessary so that the buyer can use accepted deals as a base line for the subsequent concurrent negotiations⁴.




Figure 1. System architecture

In more detail, the model for the buyer agent consists of two main components: a *coordinator* and a number of *negotiation threads* (see figure 1). The negotiation threads deal directly with the various sellers (one per seller) and are responsible for deciding what counter-offers to send to them

⁴This model is obviously biased in favor of the buyer and future work will look at relaxing this constraint so that sellers can also renege on deals.

and what proposals to accept. For maximum flexibility, we assume that the buyer agent may adopt different strategies in each of its threads. Thus the agent may choose to adopt a tough strategy in some of its threads in order to see if it can obtain a good deal and a more concessionary strategy in other threads in order to maximize its chances of having some acceptable deals. The relative mix of such strategies causes the agents to have different operational characteristics and we report on our initial evaluation to this end in section 3. We adopt separate semi-independent threads for reasons of modularity and coherence. The alternative of having every single negotiation move centrally coordinated and intercepting all the bids received from all the sellers, is viewed as a computational bottleneck for the time-constrained environments we are targeting.

Each negotiation thread inherits the preferences from the main buyer agent, including the acceptable ranges of values for each negotiation issue, the deadline of the negotiation and the current reservation value (the lowest utility value of an offer that the agent considers acceptable). The coordinator decides the negotiation strategies for each thread (details of how it does this are given in section 2.1). After each round⁵, the threads report back their status to the coordinator. If a thread reaches a deal with a particular seller, it terminates its negotiation. Based on the coordination schema it is using (see section 2.1 for more details), the coordinator will then notify all other negotiation threads of the new reservation value and it may change the negotiation strategy of some of them. The detailed working of the two components are described below.

2.1 The coordinator

The coordinator is the most important component of the *buyer agent*. It is responsible for coordinating all the negotiation threads and choosing an appropriate negotiation strategy for each thread.

As a first step, the coordinator acts like a blackboard for shared information about the ongoing encounters. It receives the current status from the various negotiation threads (including the proposal's values), keeps a list of agreements reached and notifies the threads about any changes in their reservation values. Second, and more importantly, the coordinator decides the negotiation strategy for each thread initially and whether this should change over time to reflect any agreements that have been made to date.

To ground our model, at this time we consider S to be composed of the class of time dependent strategies advocated in [1] for bi-lateral negotiations in uncertain environments with time constraints. These strategies fall into three

⁵A round consists of the exchange of one offer and one counter-offer between the buyer and all the sellers. At this time we assume the sellers respond in approximately the same time frame although in the future we will consider what happens if they have widely varying response times.

categories, namely: *conceder* (S_c), *linear* (S_l) and *tough* (S_t) where $S = S_c \cup S_l \cup S_t$. All of the strategies start with the same initial value that is generated in relation to the deadline and the reservation value. The *conceder* strategy quickly lowers its value until it reaches its reservation value. The *linear* strategy drops to its reservation value in a steady fashion. Finally, the *tough* strategy keeps its value until the deadline approaches, then it rapidly drops to its reservation value.

In his empirical analysis of the behavior of negotiating agents that adopted these strategies, Faratin showed that if it is possible to approximate the type of the opponent then the agent can alter its strategy to be more effective. Given this observation, the coordinator attempts such a classification. Specifically, at time t : $2 < t \leq t_{max}$, called the *analysis time*, the coordinator tries to determine if a given seller is a *conceder* or a *non-conceder*. In particular, assume O_j^i is the value of the offer that seller agent i made at time j : $0 \leq j \leq t$. Then seller i is considered a *conceder* if $\forall k \in [2, t] : \frac{O_k^i - O_{k-1}^i}{O_{k-1}^i - O_{k-2}^i} > \alpha$ where α is the threshold value set on concessionary behavior. There is a similar characterization of non-conceders and if the agent falls into neither category, it is judged not classified.

Let the set of *conceder* and *non-conceder* agents be represented by A^c and A^n , respectively. Now, given the set of strategies $S = \{S_c, S_l, S_t\}$ and the set of agents $A = \{A^c, A^n\}$, the coordinator changes the strategy for each negotiation thread based on the type of the agent it is negotiating with, in order to try and obtain better outcomes. Agents belonging to the set A^c are willing to concede in order to end up with agreements. Therefore, if the agent negotiates toughly with some of them (keeping its offer consistent), it may obtain a deal that has higher value than if it continues negotiating in its present manner. However, if the agent negotiates in this way with all the agents, it may not obtain any deals at all. Therefore, for reasons of balance, the agent will negotiate in a tough manner with a subset of the agents in A^c , specifically with a percentage ($P_t^c\%$) of them. For the remainder of the agents in A^c , the strategy remains unchanged. Similarly, if the agent believes a particular agent is in the set A^n then in order to make sure it obtains a deal with some of them, it makes some of its own strategy more conciliatory. Specifically, for the agents belonging to the set A^n , a fixed percentage ($P_c^n\%$) of them will have their behavior made conciliatory, while the remainder have their strategies unchanged. There is no change to agents whose behavior cannot be classified.

2.2 The negotiation threads

An individual negotiation thread is responsible for dealing with an individual seller agent on behalf of the buyer. Each such thread inherits its preferences from the buyer

agent and has its negotiation strategy specified by the coordinator.

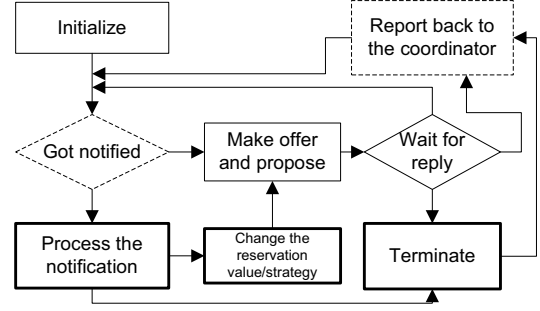


Figure 2. A single negotiation thread

In each thread (see figure 2), there are three main sub-components; namely *communication* (represented by the dotted lines), *process* (represented by the bold lines) and *strategy* (represented by the normal lines). The *communication* subcomponent is responsible for communicating with the coordinator. Before each round, it checks for incoming messages from the coordinator and if there are any, it passes them to the *process* subcomponent. After each round, it reports the status of the thread (including proposed proposals and the deal's value if an agreement is reached) back to the coordinator. The *process* subcomponent processes messages from the *communication* subcomponent. This can either be changing the reservation value or changing the strategy. The *strategy* subcomponent is responsible for making offers/counter-offers, as well as deciding whether or not to accept the offer made by the seller agent. It uses the reservation value as the basis for deciding whether to accept the seller's offer; in this case any offer with a value greater than this is accepted, otherwise a counter-proposal is made (unless the deadline has passed in which case a decline is sent).

3 Empirical evaluation

Having defined the model, the next step is to evaluate it. Given the aims of our work, we are interested in operational performance and so we decided to evaluate it empirically. In particular, we pose a number of hypotheses and evaluate them in different types of environment. Here, the environment is defined by the number of participating agents (between 5 and 30), the number of negotiation issues in the proposals (between 1 and 8) and the deadlines for the buyer and the sellers (between 5 and 30). The ranges for these parameters are chosen based on our experience of typical values in the domains in which we are interested.

The experiment's dependent variables are the time taken to reach an overall agreement, the number of proposals exchanged in coming to this agreement and the relative per-

formance improvement of the concurrent model over the sequential model. Given the utility value for the outcome of a particular negotiation in a given environment for the sequential model (U_s) and the one obtained by the concurrent model (U_c), the performance improvement of the concurrent model is calculated as *performance improvement* = $\frac{(U_c - U_s)}{U_s}$. Clearly, if this value is negative then the sequential model has outperformed the concurrent one.

Our concurrent model is compared against a sequential negotiation model based on [1]. In this model, all the agents' preferences, as well as the allocation of the strategies, are drawn from the same distributions as the concurrent ones. The only difference with [1] is that if the buyer agent reaches an agreement of value p in negotiation i , then in all subsequent negotiations, p will be its new reservation value.

We now turn to the specific hypotheses.

Hypothesis 1 *The time to complete the negotiation is less for the concurrent model than for the sequential one.*

EVALUATION. Figure 3 shows the percentage of time saved by performing the negotiation concurrently compared with sequentially. As can be seen, this saving is proportionate to the number of participating agents. The reason for this is that by negotiating concurrently the time consumed for all the threads is not more than the largest deadline of the sellers and the buyer. For each agent, it is only allowed to continue the negotiation until its deadline at which point it must stop. Thus, for each negotiation thread i , the longest period it is allowed to continue is $\min(t_{max}, t_{i_{max}})$. As a result, our model's negotiation will stop at the latest period $t = \min(t_{max}, \max(t_{1_{max}}, t_{2_{max}} \dots t_{n_{max}}))$. On the other hand, when negotiating sequentially, the buyer needs to wait for the first negotiation to finish, then it can start the second negotiation and so on. In general, each negotiation could last until the deadline $\min(t_{max}, t_{i_{max}})$. Thus, the overall negotiation could finish at $t = \min(t_{max}, t_{1_{max}}) + \min(t_{max}, t_{2_{max}}) + \dots + \min(t_{max}, t_{n_{max}})$ in the worst case.




Figure 3. Percentage of time saved

Hypothesis 2 *The number of proposals that are made in the concurrent model is less than the number in the sequential one.*

EVALUATION. Figure 4 shows the number of proposals saved (as a percentage) by performing the negotiation concurrently compared with sequentially. As can be seen, the average number of proposals made in the concurrent model is less than the number in the sequential model and this difference increases in proportion to the number of participating agents. From our experiments, we observe that the total number of agreements reached in the concurrent model is more than in the sequential model (greater than 50% more). This means, in the concurrent model, more negotiations terminate before their deadlines are passed. Thus, the total amount of time it takes to complete all the negotiations in the concurrent model is less than in the sequential model. As the number of proposals made in each negotiation is relative to the time it takes to complete the negotiation, the amount of proposals made in the concurrent model is less.




Figure 4. Percentage saving in the number of proposals sent

Hypothesis 3 *To realize the benefits of concurrent negotiations, the buyer agent's deadline must not be too short.*

EVALUATION. Figure 5 shows the performance of the concurrent model with different values of the buyer's deadline. This shows the longer the buyer's deadline, the better the performance improvement. This occurs because if the buyer's deadline is short (less than 15 units in this case), the time when one negotiation thread reaches an agreement is necessarily close to the deadline. Thus it has little effect on the other negotiation threads. On the other hand, if the buyer's deadline is longer, once a negotiation thread finishes with an agreement, it can be used to influence the other threads. Hence giving better deals.

Hypothesis 4 *The final agreements reached by the concurrent model have, on average, higher or equal utility for the buyer than those of the sequential model (assuming the deadline is not too short).*




Figure 5. Performances versus deadline

EVALUATION. When a thread negotiates with a seller, it tries to find an offer that is acceptable to that particular seller. This is achieved by exploring the space of agreements in some way. In the sequential model, each thread starts with a reservation value p_s , which can either be the initial preference value (if this is the first encounter or no agreement has yet been reached) or the value of the best previous agreement reached in an earlier negotiation. Moreover, this value remains unchanged until the particular negotiation thread finishes. On the other hand, in the concurrent model, each thread starts with a reservation value p_c ($p_c = p_s$ at $t = 0$), which is the initial preference value of the agent. This value may then be changed during the course of the negotiation as a result of an agreement obtained in another thread. This, in turn, narrows the space of agreements for the buyer to only those that have a higher utility value than the current reservation value. Hence, if the buyer reaches an agreement, the utility value of this agreement will be greater than the one it already has. Assuming all the threads have sufficient deadlines, whenever an agreement is reached, the search space of all the concurrent negotiation threads will be reduced simultaneously. Thus, on average, the buyer strives to reach a higher utility value for a greater proportion of the negotiation time with more sellers than it does in the sequential model. This means the performance of the concurrent model is often better than the sequential one.

However, in some cases, by narrowing the space of agreements, no intersection with the seller's space of agreements may be found in the concurrent model. Therefore, the agents will not be able to reach an agreement and so the utility value of the final agreement is reduced. In these cases, the overall performance of the concurrent model will be less than its sequential counterpart. Our experiments' results indicate that in all environmental settings, on average, the results of the concurrent model are better than the sequential, ranging marginally from 1.5% to 2.5% depending on the number of participating sellers (see figure 6).

Hypothesis 5 *Changing the strategy in response to the*




Figure 6. The improvement rate

agent's assessment of the ongoing negotiation is equal or better than not doing so.

Hypothesis 6 *To improve the performance of the concurrent model, the analysis time should be moderately early (to have time to have some effect) but not too early (so it is reasonably accurate).*

EVALUATION. To evaluate these hypotheses, we varied the *analysis time* (see section 2.1) relative to the buyer's deadline (see figure 7). As can be seen, the best performance improvement is obtained when the sellers are analyzed about a third of the way into the negotiation period. This is sufficiently near the beginning to be able to have an effect on the rest of the negotiations, but sufficiently far into the encounter to make a reasonable approximation about the type of the negotiation opponent. With respects to hypothesis 5, the outcome with analysis time equal to 100% is equivalent to an agent that does not change its strategy during the encounter. As can be seen, this leads to poor outcomes and so changing (at any time) is not worse, in most cases, is beneficial.




Figure 7. Performance versus analysis time

Hypothesis 7 *When dealing with sellers in A^c , the tougher the buyer negotiates the better the overall outcome it obtains.*

EVALUATION. To evaluate this hypothesis, we varied P_t^c (see section 2.1) through all possible values. Figure 8 shows the more tough the agent is, the better the outcome it obtains. This is because when dealing with a *conceding* seller, if the buyer keeps its offer consistent, as the deadline approaches, the seller will quickly lower its proposal value close to its reservation value (if it has a deadline shorter than that of the buyer thread⁶). Thus, if an agreement is reached at this point, its utility value for the buyer will be higher than that obtained if the buyer adopts any other strategy.




Figure 8. Performance versus degree of response toughness

4 Related work

In the field of bi-lateral negotiation, there are a number of proposed models, each with their relative advantages and disadvantages. For example, [3] ensures agreements are Nash equilibria and is guaranteed to terminate after a fixed period of time. However, the model assumes (unrealistically) that the agents' preferences are completely revealed. In contrast, [1] uses a heuristic approach that operates in the presence of information uncertainty but his results are analyzed empirically rather than theoretically (as per our model). Similarly, [2] does not require the agents to completely reveal their preferences. However, it does assume there are a fixed number of preference sets of agent's types with known probabilities of occurring. In all cases, however, these models do not deal with concurrent negotiation encounters.

The only extant model that deals explicitly with concurrent encounters is [6]. Like us, they use the concept of sub-negotiators and a coordinating agent. However, by considering the negotiation as a distributed constraint satisfaction problem, they limit the range of possible negotiation strategies that can be used. In contrast, our model can incorporate strategies without restriction. Moreover, our model has

a well defined coordination mechanism that has been empirically evaluated (their model does not give such detail).

5 Conclusions and future work

This paper has developed a heuristic model for managing concurrent negotiations in time-constrained settings where agents have no prior knowledge of their opponents or their types. Through empirical evaluation, we showed how the model leads to better deals, more quickly than its sequential counterpart. We also highlighted the importance of the time when the opponents' negotiation strategies are classified and on the response to this assessment in terms of the degree of toughness adopted.

There are, however, a number of areas that still require further work. First, the means by which negotiation opponents are classified as being conceder or non-conceder needs to be refined so that this monitoring can be made on an ongoing basis (rather than as a one-off decision). Second, we need to allow for the possibility of sellers decommitting and then having these commitment changes feed-back into the buyer's negotiation behavior. Third, we wish to extend the implementation so that the coordinator and the negotiation threads also embody fundamentally different models of bi-lateral negotiation (e.g. based on constraint-satisfaction, game-theory or any other method that is likely to be effective). In this case, the key challenge is in designing the coordinator so that it can select, monitor and modify these strategies in line with the agent's overarching negotiation objectives.

References

- [1] P. Faratin. *Automated Service Negotiation Between Autonomous Computational Agents*, PhD Thesis. Queen Mary College, London, England, 2001.
- [2] S. S. Fatima, M. Wooldridge, and N. R. Jennings. Negotiation strategies for agents with incomplete information. *Proc. 8th Int. Workshop on Agent Theories, Architectures and Languages (ATAL)*, Seattle, USA, pages 53–68, 2001.
- [3] S. Kraus. *Strategic Negotiation in Multi-Agent Environments*. MIT Press, Cambridge, USA, 2001.
- [4] A. R. Lomuscio, M. Wooldridge, and N. R. Jennings. A classification scheme for negotiation in electronic commerce. *Int Journal of Group Decision and Negotiation*, 12(1), 2003.
- [5] T. D. Nguyen and N. R. Jennings. A heuristic model for concurrent bi-lateral negotiations in incomplete information settings. *Proc. 18th Int. Joint Conf. on AI*, 2003.
- [6] I. Rahwan, R. Kowalczyk, and H. H. Pham. Intelligent agents for automated one-to-many e-commerce negotiation. *Twenty-Fifth Australian Computer Science Conference (ACSC2002)*, 4:197–204, 2002.

⁶If the seller has a longer deadline then no agreement will be reached.