Coordinating multiple concurrent negotiations*

Thuc Duong Nguyen and Nicholas R. Jennings School of Electronics and Computer Science University of Southampton, UK.

E-mail: {tdn01r, nrj}@ecs.soton.ac.uk

Abstract

To secure good deals, an agent may engage in multiple concurrent negotiations for a particular good or service. However for this to be effective, the agent needs to carefully coordinate its negotiations. At a basic level, such coordination should ensure the agent does not procure more of the good than is needed. But to really derive benefit from such an approach, the agent needs the concurrent encounters to **mutually influence** one another (e.g. a good price with one opponent should enable an agent to negotiate more strongly in the other interactions). To this end, this paper presents a novel heuristic model for coordinating multiple bilateral negotiations. The model is empirically evaluated and shown to be effective and robust in a range of negotiation scenarios.

1. Introduction

Automated negotiation is a key form of interaction in agentbased systems and such negotiations exist in many different forms [5]. In this paper, we focus on one such form, namely one-to-many negotiations in service-oriented contexts. Here, a service is simply viewed as an abstract representation of an agent's capability (this view is now widespread in a range of domains that we are targeting for our work, including the web, the grid, pervasive computing and e-business). Thus, one agent is seeking to provision a single service (described by multiple attributes, such as cost, time, quality, etc.) from a number of potential providers. Now this type of encounter is usually handled via some form of single-sided (reverse) auction protocol. However, in [6], we introduced multiple, concurrent bilateral negotiations as an alternative. Our approach offers a number of advantages over its more traditional counterpart (especially in the time-constrained environments that motivate our work).

First, in most reverse auctions, the buyer is only allowed to select an agreement from the set proposed by the sellers. On the other hand, the buyer in our approach can also send proposals and counter-proposals. For multidimensional agreements, this two way communication is important because it allows the buyer to provide an indication of the areas of the search space where it would like to see the agreements lie. This, in turn, should lead to more efficient negotiations because the agents can focus on the relevant areas of the space more rapidly. Moreover, separate bi-lateral encounters (threads) enable the buyer agent to deploy different strategies in each. Thus, in some cases, it will adopt a very tough bargaining stance in order to try and obtain a high value agreement, while in other cases it can adopt a strategy that is more likely to lead to an agreement (but perhaps one that is not as valuable). This variability means negotiation can be tailored to the individual opponents (e.g. some opponents may be known to be desperate to obtain a deal), rather than derived implicitly though the competition of the sellers (as happens in the traditional auctions). Also, the agreement reached in one thread can be used to influence negotiation behavior in other threads. This gives the buyer additional strategic information (and hence bargaining power) that can be exploited to obtain better deals.

Second, the time at which an agreement is reached in the multiple concurrent negotiation case can be reduced. For auctions that do not have deadlines, the end time is indeterminate which is unacceptable for our time-constrained domain. In auctions where there is a deadline, no agreement can be reached before this time. On the other hand, by using multiple concurrent negotiations, deals are likely to be available before the overall deadline and if these are deemed satisfactory the agent may decide to terminate other negotiations (perhaps sacrificing some potential gain) in order to take benefit from the agreed deal more quickly.

The downside of the concurrent negotiation approach is that the agent has to coordinate its bidding behavior across multiple negotiation threads. At a basic level, such coordination needs to ensure the agent does not procure more of the good than is needed. However to really derive ben-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

^{*} This work is sponsored by BT Exact.

efit from such an approach, the agent needs the concurrent encounters to *mutually influence* one another. Thus a good deal reached in one negotiation thread should enable it to negotiate more strongly in remaining threads because it already has a good deal that it can fall back on. In contrast, this coordination in auctions is enforced by the auction protocol which, in a sense, serializes the bidding and the buyer's decision making problem.

Against this background, a number of coordination techniques have been proposed (see section 4 for more details). Generally speaking, however, these approaches suffer from a number of shortcomings. First, most techniques deal with the various sellers in a homogeneous way (i.e. the behavior of the buyer agent is fixed throughout the negotiation process, regardless of the agent it is dealing with). This rigid approach is unlikely to be effective in open and dynamic environments because not all the participating sellers are likely to be similar. They will typically come from various sources and have different objectives. Thus, some will be desperately searching for an agreement, whereas others will just be trying to improve their current positions. Consequently, in order to be effective, the buyer agent needs to be flexible in its bargaining behavior and change its negotiation strategy to fit with its prevailing context. Second, the mutual influence capability of the existing techniques is limited. It can only ensure that the threads will not accept any agreement with a value lower than what has already been achieved. While this may be adequate to secure a deal, it does not allow the buyer to obtain the best possible agreement. Third, existing techniques tend not to take into account available information about participating sellers (e.g. their typical negotiation trends, or whether they seem to be desperate to obtain a deal or adopt a tough stance). This information is sometimes available in a social system of agents [10], [4] and if known should be used to enable the buyer agent to reach better deals.

To this end, in this paper we report on a novel coordination model that removes the aforementioned shortcomings. Specifically, we develop a heuristic approach to coordinating multiple concurrent negotiations. In so doing, this work advances the state of the art in the following ways. First, prior to the negotiation episode, the buyer selects its negotiation strategy for the various threads based on its belief about the available service providers (this belief is represented as a probability distribution over the different types of providers, see section 2.1). Second, we classify the sellers according to their specific behaviors during the encounter and, consequently, adapt the agent's negotiation behavior based on these classifications. Third, the results from a successfully terminated thread can be used to influence other ongoing threads. Finally, we empirically evaluate our model against the other main approaches advocated in the literature and show that it can outperform them in a broad range

of negotiation situations.

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

2. The concurrent negotiation model

This work builds upon the basic model outlined in [6]. In particular, the main contribution of this paper is in enhancing the coordinator component which was very simple in the initial proposal (see section 4 for more details). Before we can focus on this component, however, we first need to recap the basic architecture of our model. The agent that wishes to purchase the service is called the *buyer* and the agents that are capable of providing the service are called the sellers. Service agreements (contracts) are assumed to be multi-dimensional. The buyer has a hard deadline $t_{b_{max}}$ by when it must conclude its negotiations for the service. Similarly, each seller α has its own (private) negotiation deadline $t_{\alpha_{max}}$. All agents have their own preferences about the service and this information is private. Each agent has a range of strategies (S) that it can adopt¹ and its choice of strategy is also private information. Each thread follows a Sequential Alternating Protocol [9] where at each step an agent can either accept the offer from the opponent, propose its counter-offer, or opt out of the negotiation (typically if its deadline is reached).



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

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Given the time-constrained nature of our encounters, the types of strategy that we consider are the time-dependent family introduced in [3]. These can be broadly divided into three classes: the *conceder* strategy quickly lowers its value until it reaches its reservation (minimum acceptable) value. The *linear* strategy drops to its reservation value in a steady fashion. Finally, the *tough* strategy keeps its value until the deadline approaches and then it rapidly drops to its reservation value.

and what proposals to accept. Each 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 round², the threads report back their status to the coordinator. If a thread reaches a deal with a particular seller, it terminates its negotiation. The coordinator will then notify all other negotiation threads of the new reservation value and it may change the negotiation strategy for 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*. It is responsible for coordinating all the negotiation threads and choosing an appropriate negotiation strategy for each thread.

Before starting a negotiation, the coordinator considers the available information about the types of the sellers that are in the environment. In our case, we consider that seller agents can be of the following types: *conceder* (i.e. they are willing to concede in the search for deals) or *non-conceder* (i.e. they tend to negotiate in a tough manner). The set of available agent types is denoted as A_{types} : $A_{types} = \{con, non\}$. This information is represented as a probability distribution over the agent types. Such information may be based on past experiences, obtained from a trusted third party, or from a system of referrals [4]. If no such information is available, all agents are assumed to be unknown.

There are two further sources of information that aid the coordinator's decision making: the *percentage of success matrix (PS)* and the *pay off matrix (PO)*. The former measures the chance of having an agreement as the outcome of the negotiation when the buyer applies a particular strategy to negotiate with a specific type of the seller (e.g. when applying a tough strategy with a non-conceder seller, the average chance of reaching an agreement is 15%). The latter measures the average utility value of the agreement reached in similar situations (e.g. when applying a tough strategy with a conceder seller, the average utility value of the agreement, once reached, is 0.7). The values of these matrices are initially set to a common value to avoid bias³ and they are updated after all the negotiation threads finish (by averaging the values over a sufficient number of encounters, vari-

ances in deadlines and reservation values in the different encounters can be largely ignored).

Given this information, the coordinator calculates the probability of the first seller (a randomly picked agent from those that will be negotiated with for the service in question) being of a specific type. Based on this, the agent calculates the expected utility of applying the various strategies at its disposal for this particular seller and selects the one that maximizes this value. Formally, the expected utility $EU(\lambda)$ for strategy $\lambda \in S$ is calculated as:

$$EU(\lambda) = \sum_{a \in A_{types}} PS(\lambda, a) PO(\lambda, a) P(a), \qquad (1)$$

where P(a) is the probability that the seller agent is of type a and PS and PO are the values in the corresponding matrices, respectively. After finishing with the first seller, the coordinator uses a Bayesian update function to update the probability distribution of the agent types and continues on with the second seller. This process is repeated until the coordinator finishes allocating the strategies to all the negotiation threads.

To illustrate this in more detail, consider the following example. Assume the coordinator has the following data prior to negotiation:

- there are 100 participating sellers (n = 100). The set of sellers is $A_s = \{\alpha_1, \alpha_2, \dots, \alpha_{100}\}$ and this set is composed of two types of sellers: $A_s = A_{con} \bigcup A_{non}$, namely *conceder* and *non-conceder*.
- there are three available negotiation strategies that it can select for a given thread: $S = S_c \bigcup S_l \bigcup S_t$, namely *conceder*, *linear* and *tough*.
- the probability that the first seller is a conceder, $P(\alpha_1 \in A_{con})$ or $P(A_{con})$, is 0.45 and the probability that the first seller is a non-conceder, $P(\alpha_1 \in A_{non})$ or $P(A_{non})$, is 0.55.
- the values of the matrices *PS* and *PO* are:

PS	A_{con}	A_{non}	PO	A_{con}	A_{non}
S_c	0.35	0.75	S_c	0.5	0.4
S_l	0.25	0.28	S_l	0.35	0.4
S_t	0.6	0.15	S_t	0.7	0.65

Based on this information, the values for the EU functions are calculated, using equation (1), as follows:

$$EU(S_c) = 0.35 * 0.5 * 0.45 + 0.75 * 0.4 * 0.55 = 0.2438,$$

$$EU(S_l) = 0.25 * 0.35 * 0.45 + 0.28 * 0.4 * 0.55 = 0.1010,$$

$$EU(S_t) = 0.6 * 0.7 * 0.45 + 0.15 * 0.65 * 0.55 = 0.2426.$$

As can be seen, strategy S_c will be chosen for the first thread. In other words, the highest expected utility is achieved when considering seller α_1 as a non-conceder.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

² A round consists of the exchange of one offer and one counter-offer between the buyer and all the sellers.

³ Naturally these matrices could have differential values if the appropriate domain information was available in a particular context.

The probability distribution $P(A_{con})$ is then updated using Bayes rule as:

$$P(A_{con}|A_s \setminus \alpha_1) = \frac{P(A_s \setminus \alpha_1|A_{con})P(A_{con})}{P(A_s \setminus \alpha_1)}$$
$$= \frac{1 \cdot P(A_{con})}{\frac{n-1}{n-1}} = \frac{0.45}{0.99} = 0.4545$$

Since there are only two types of seller: conceder and non-conceder, we have $P(A_{non}) = 1 - P(A_{con})$. Hence:

$$P(A_{non}|A_s \setminus \alpha_1) = 1 - P(A_{con}|A_s \setminus \alpha_1) = 0.5455$$

Here, $P(A_{con}|A_s \setminus \alpha_1)$ is the probability that the second seller is a conceder. Similarly, $P(A_{non}|A_s \setminus \alpha_1)$ is understood as the probability that the second seller is a non-conceder. Again, the values for the EU functions for the second seller are calculated, using equation 1, as follows:

$$\begin{split} & \text{EU}(S_c) = 0.35 * 0.5 * 0.4545 + 0.75 * 0.4 * 0.5455 = 0.2432, \\ & \text{EU}(S_l) = 0.25 * 0.35 * 0.4545 + 0.28 * 0.4 * 0.5455 = 0.1009, \\ & \text{EU}(S_l) = 0.6 * 0.7 * 0.4545 + 0.15 * 0.65 * 0.5455 = 0.2441. \end{split}$$

Now, S_t will be chosen as the strategy for the second thread. The coordinator continues in this vein, until it finishes allocating the strategies to all the threads.

The other task of the coordinator is to classify the sellers during negotiation and to change the negotiation strategies for the threads. Specifically, the buyer attempts to characterize the sellers, based on the utility value of their proposals, into the sets A_{con} , A_{non} . Thus, at time $t: 2 < t \leq t_{b_{max}}$, called the *analysis time*, the coordinator tries to determine if a given seller is a *conceder* or a *non-conceder*. In particular, assume $U(\alpha, \tau)$ is the utility value of the offer that seller agent α made at time $\tau: (1 \leq \tau \leq t)$, according to the buyer agent's preferences. Then seller α is considered a *conceder* if $\forall \tau \in [3, t]: \frac{U(\alpha, \tau) - U(\alpha, \tau - 1)}{U(\alpha, \tau - 1) - U(\alpha, \tau - 2)} > \theta$ where θ is the threshold value set on concessionary behavior. If this condition is violated, seller α is considered a *non-conceder*.

Now, given the set of strategies S and the set of classified seller agents A_s , the coordinator changes the strategy for each negotiation thread based on the type of the agent it believes it is negotiating with. Specifically, for each agent $\alpha \in A_s$, the coordinator selects the strategy $\lambda \in S$ that provides the maximum expected utility and applies it to the corresponding thread, using equation (1), with $P(i \in A_{turner}) = \begin{cases} 1 & \text{if } \alpha \text{ is of type } j \end{cases}$

$$P(j \in A_{types}) = \begin{cases} 1 & ij \ \alpha \ is \ of \ typ \\ 0 & otherwise \end{cases}$$

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 subcomponents; namely communication (represented by the dotted lines), process (represented by the bold lines) and strategy. 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 back to the coordinator. The process subcomponent deals with 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 outlined the model, the next step is to evaluate its effectiveness. In this work, *empirical evaluation* is used as the method of measurement for a number of reasons. First, because our model is heuristic in nature, it is difficult to make meaningful theoretical predictions. Second, there are a number of internal variables that control the behavior of the model, as well as external variables that define the environment in which the model is being used. These variables are interrelated and need to be considered in a broad range of situations. Empirical techniques allow us to manipulate these variables, conduct the experiments and analyze the results.

In more detail, we use the *exploratory studies* evaluation technique [2]. With this method, *general hypotheses* are formed to express the intuitions about the causal factors within the model. The *experiments* are then conducted

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

and generate the results that either support these hypotheses or go against them. In our evaluation, the *independent* variables are given in table 1 and the *dependent* ones are listed in table 2.

Variables	Descriptions	values
n	the number of seller agents	[1,30]
m	the number of negotiation issues	[1,8]
$t_{\alpha_{max}}$	the negotiation deadlines of agent α	[5,30]
$x_{j_{min}}^{\alpha}$	minimum value for issue j for agent α	[0, 20]
$x_{j_{max}}^{\alpha}$	maximum value for issue j for agent α	[30, 50]
w_j^{lpha}	the weight of issue j for agent α	$\frac{1}{m}$

Table 1. The independent variables.

Since there are an infinite number of possible environments, selecting a finite subset of these is necessary to assess the performance of the model. To this end, the number of seller agents (n) and the number of negotiation issues (m) reflect typical values for our target domains. An agent α 's preference for issue j is represented by the tuple $\{x^{\alpha}_{j_{min}}, x^{\alpha}_{j_{max}}, w^{\alpha}_{j}\}.$ The tuple $[x^{\alpha}_{j_{min}}, x^{\alpha}_{j_{max}}]$ is an interval independent variable, whose scale is infinite. To simplify the analysis, therefore, we assume all issues have the same domain of values and we randomly set the value for $x^{\alpha}_{j_{min}}$ to be in the interval [0, 20] and $x^{\alpha}_{j_{max}}$ to be in the interval [30, 50]. The values for w_j^{α} are set to give all issues equal importance. The negotiation deadline for each agent is an ordinal independent variable, whose value is randomly chosen, ranging from 5 (very short deadline) to 30 (long deadline).

The seller agents in this evaluation are characterized by three independent variables whose values are set in the following manner:

- the values' domain for the set of negotiation issues: These domains are randomly generated (from the same distribution as the buyer agents' values) so that each domain intersects with the corresponding domain of the buyer's preference. For example, if the buyer's value domain for an issue j is $[x_{j_{min}}^b, x_{j_{max}}^b]$ then the corresponding value domain for seller α will be generated as $[x_{j_{min}}^\alpha, x_{j_{max}}^\alpha]$ that satisfies $x_{j_{min}}^b \leq x_{j_{min}}^\alpha \leq x_{j_{max}}^\alpha$.
- *the negotiation strategy*: Each seller is assigned a random strategy selected from a predefined set of alternations (as outlined in [3]). This set is composed of timedependant functions (like conceder, boulware and linear) and behavior-dependant tactics (such as tit-for-tat in its various forms).
- *the negotiation deadline*: The deadline for each seller is generated from the same distribution as for the buyer.

We benchmark our model (noted as eCN - for e-commerce Concurrent Negotiations) against the optimal solution (optimal) and three different controls, namely desperate (D), patient (P) and optimized pa*tient (OP)*. The optimal mechanism operates in a perfect information situation in which the agents know the preferences and strategies of other agents. Given this, the buyer agent is able to find the Pareto optimal agreement for each thread if such an agreement exists. If no such agreement exists, the utility of that thread is considered to be 0. The individual agreement that maximizes the buyer agent's utility is then selected as the optimal solution. The other three controls are based on the theoretical work of [8], which is the only other extant model that deals explicitly with concurrent encounters. Basically, D terminates all the negotiations whenever an agreement is found in any one thread, P waits until all the negotiations finish and then selects the highest value agreement as the final answer, and **OP** extends **P** in that whenever an agreement is found, its value is broadcast to all other ongoing threads so that they will not accept a lower value agreement.

Variables	Descriptions
U	the utility value of the final agreement
Ν	the number of successful negotiations

Table 2. The dependent variables.

After each experiment, we measure the utility value of the final agreement for the buyer (U). In our evaluation, the utility of an offer $X = \{x_1, x_2 \dots x_m\}$ to an agent α is calculated as:

$$U(X) = \sum_{j=1}^{m} w_j^{\alpha} \cdot \frac{x_j - x_{j_{min}}^{\alpha}}{x_{j_{max}}^{\alpha} - x_{j_{min}}^{\alpha}}$$
(2)

We also measure the number of agreements reached (as a percentage) during the whole negotiation encounter (N). In all cases, the results are gathered from a series of experiments in different environment settings. Each experiment consists of 2000 runs and the results are averaged and put through a regression test to ensure that all differences are significant at the 99% confidence level.

We now turn to the specific hypotheses.

Hypothesis 1 *Our model will achieve more and higher utility agreements than the controls.*

To evaluate this hypothesis, we average the utilities achieved with varying numbers of seller agents and varying deadlines. The results of our model, the controls and the optimal are displayed in figure 3. As can be seen, our model is between 6-8% better than the closest control and between 11-21% lower than the optimal.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.



Figure 3. Final utility value for varying numbers of sellers.

Amongst the controls, D has the worst utility (since it terminates whenever an agreement is reached), P has a better utility (since it waits until all the negotiations finish and selects the best deal), and OP provides the best utility (since it is an improved version of P). Consequently, from now on, we will only focus on OP as the main point of comparison.

Fundamentally, our model differs from the others in the way that the buyer agent behaves both prior to and during the negotiation process. Unlike the controls, in which the strategy employed by the buyer stays constant throughout the negotiation episode, each negotiation thread in our model will change its strategy if it believes there is a benefit in so doing. Recall from section 2.1, our agent selects its initial strategies based on its beliefs about the opponents that it is likely to encounter⁴. Specifically, as the sellers have different objectives, they are likely to behave differently. Some desperately want to sell their services, while others will only agree to a deal if it will benefit them more than what they already have. Since we do not know the exact characteristic of each seller, our initial strategy selection is not guaranteed to be accurate. However, we overcome this problem by reclassifying the sellers during the negotiation process (based on their actual behaviors rather than the general market beliefs). Based on this classification, some of the threads change their strategies. In some cases, this flexible behavior of the buyer agent helps it finds high value agreements that would not have been found otherwise. In our experiments, for example, we found that in 30-40% of threads where an agent changes its strategy an agreement is reached where one would not have been possible without the change. Consequently, this increases the buyer's utility and leads to an improvement in the model's performance. The improvement in utility is particularly marked when the buyer can recognize a conceder seller and can negotiate in a very tough manner to obtain a high value deal.

However, it is not always beneficial for an agent to change its strategy. This is particularly the case when it leads to conflict between the buyer and the seller and no agreement can be reached. Thus, in 12% of threads where an agent changed its strategy, no agreement was reached, whereas without such a change an agreement would have been found (although this may not have been the agent's overall agreement).

No of sellers	5	10	15	20	25	30
eCN	1418	1615	1710	1762	1802	1830
OP	1389	1593	1690	1744	1776	1804
optimal	1686	1827	1887	1908	1928	1946

Table 3. Number of successful negotiations.

In terms of agreements made, as can be seen from table 3, our model produces more agreements than the others. This improvement can also be explained by the adaptive nature of our strategy selection. Compared to the controls, the number of times our strategy selection leads to a conflict is lower than the number of times it leads to an agreement. This, in turn, leads to a modest increase in the number of successful negotiations.

Hypothesis 2 To realize the benefits of our model, the buyer agent's deadline cannot be too short.

Figure 4 shows the difference in the performance of our model compared to **OP** with different values of the buyer's deadline. As can be seen, the longer the deadline, the better the performance of our model. This is because the buyer's deadline affects our classification of the sellers which is a key deciding factor for our improved performance over other controls (see hypothesis 1). Recall from section 2.1, the buyer categorizes the sellers according to their proposals' utility values. Thus, if the deadline is too short, the data gathered is insufficient for the buyer to detect a meaningful pattern in the behavior of a particular seller. Thus, the classification of sellers is inaccurate and so the model performs poorly. On the other hand, given an adequate negotiation deadline (above 10 in this case), the buyer will have more data to analyze the sellers' behaviors. This, in turn, improves the accuracy of the classification process and, eventually, leads to better deals.

In our experiments, most of the final agreements are obtained from threads that adapt their strategies after the classification process (this occurred in about 75% of all the tests). This figure further explains why accurately classifi-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

⁴ Naturally, these beliefs may not necessarily be true. Therefore we examine the effect of the accuracy of this initial selection on the model's performance in hypothesis 4.



Figure 4. Buyer's performance for varying deadlines.

cation of the sellers plays such an important role in improving the model's performance.

Hypothesis 3 *The larger the number of participants, the closer the utility produced by our model is to the optimal.*

Here, we measure the differences in the results obtained by our model with the optimal as the number of participants increases. The results with respect to utility and number of agreements are displayed in table 4.

No of sellers	5	10	15	20	25	30
U(<i>eCN</i>)	0.56	0.61	0.64	0.66	0.68	0.70
$U(eCN) t_{b_{max}} \ge 15$	0.62	0.67	0.70	0.72	0.75	0.76
U(optimal)	0.76	0.77	0.78	0.79	0.80	0.81
N(eCN)	1418	1615	1710	1762	1802	1830
$N(eCN) t_{b_{max}} \ge 15$	1499	1672	1753	1797	1830	1860
N(optimal)	1686	1827	1887	1908	1928	1946

Table 4. Buyer's performance with varyingnumbers of sellers.

As can be seen, the greater the number of participating sellers, the closer our result is to the optimal. Specifically, the gap between the results decreases from 18% to 9% as the number of sellers increases from 5 to 30. This is explained by the fact that the buyer only finalizes the deal with the seller that provides the highest value deal. Thus, as the number of sellers increases, so does the number of agreements reached by the threads. Since these agreements are used to influence other ongoing threads, the utility value of the final agreement will be improved. This is also the situation for the number of successful negotiations. With 5 sellers, the agent only succeeds in 40% of the encounters. However, this rate increases to nearly 82% when there are 30 participating sellers. Furthermore, if we only consider cases where the buyer has a sufficient deadline (larger than

10 units in this case), our results come even closer to the optimal (the gap decreases from 15% to 4%, whereas the success rate is increased from 52% to 87%). Again this is mainly due to the accuracy of our classification process (see hypothesis 2).

Hypothesis 4 *The more accurate the agent's information about the probability distribution of agent types, the better the performance of our model.*

To ensure our model can perform robustly in unpredictable environments, this set of experiments evaluates its reliance on the accuracy of information an agent holds about the market place. Specifically, we consider the degree to which the probability distribution P (defined in section 2.1) matches reality and what impact this has on the initial selection of negotiation strategies.



Figure 5. The accuracy of the belief versus the performances.

The initial selection of strategies is only part of the story since the buyer agent can reclassify the sellers and change its strategy. Nevertheless, it can be observed from figure 5^5 that the accuracy of this information does have an effect on the result of the process, albeit by a small figure (1-2%). In our experiments, about 9% of the agreements were reached in threads before the sellers' classification occurs and some of these initial agreements become the final solutions at the end of the encounter process. Thus, the aforementioned small improvement was made by improving these early agreements.

⁵ Here, the unknown plot corresponds to the case where P has equal values throughout, 50% to the case where half of the values in P are correct, and 100% is where P reflects the actual strategies of the sellers. OP does not use P in its decision making and optimal also operates with the correct values for the sellers' strategies.

4. Related work

Most of the existing work in the area of one-to-many negotiations uses some form of single-sided auction. However while this has many advantages in a range of scenarios, it also has a number of shortcomings as discussed in section 1. Thus, a number of researchers have considered multiple concurrent negotiations as an alternative.

In AutONA [1], concurrent negotiation is used in the domain of operational procurement. In more detail, the buyer agent in AutONA seeks to buy a good that consists of multiple purchase units from several sellers. It tries to divide the quantity up among the sellers and for each seller it generates offers using a belief system over price (a probability distribution over prices per unit, parameterized by the properties each unit may have). The negotiation finishes when an acceptable solution is found. AutONA has been empirically evaluated and is shown to secure good outcomes. However, this approach is not suitable in our context because we want to procure an indivisible good (a service) from a single provider whereas AutONA focuses on dividing the good among all the potential sellers. If we wanted to extend our model to deal with procuring multiple instances of a particular service then AutONA's strategy could be exploited to complement our own.

On the other hand, the theoretical model presented in [8] has the same objective as our work. They also use the same basic concepts of sub-negotiators and a coordinating agent. However, their coordination mechanisms are limited (see the D, P and OP strategies in section 3) and we have shown that our coordination mechanism is significantly more effective.

In our previous work, we developed a simple coordinator for this target environment and we showed the benefits of our approach over a series of sequential negotiations [6]. The work described in this paper, however, extends this basic coordination mechanism in two ways: (1) the initial strategy selection is based on the buyer's beliefs about the potential providers (previously we did not take into account this information) and (2) the mutual influence capability is improved (previously we changed the strategies at the analysis time in a randomized way).

5. Conclusions and future work

This paper has developed a novel heuristic model for managing concurrent negotiations in time-constrained settings. Through empirical evaluation, we showed how the model leads to good deals. This is especially true when there are large numbers of participants and when the agent is given sufficient time to reach a deal. We also showed that our coordination mechanism outperforms the existing mechanisms proposed in the literature, both in terms of the number of deals that are made and the utility obtained from them. Our model is also currently being used in a number of real world applications to form and maintain coalitions in business and e-science virtual organizations [7] and in an internal project of BT concerned with logistics planning.

For the future, there are a number of ways in which our model can be improved. First, we want to extend our negotiation model so that the participating providers can also renege from deals (as it stands, this capability is only available to the buyer agent). This, in turn, will allow the agents to behave in a more flexible and efficient manner. Second, we believe the solution found by our model can be still further improved by incorporating some form of reinforcement learning into the seller classification process. This will help increase the accuracy of this process and will, in turn, result in higher performance of the model.

References

- A. Byde, M. Yearworth, K. Y. Chen, and C. Bartolini. Autona: A system for automated multiple 1-1 negotiation. In *Proceedings of the 2003 IEEE International Conference on Electronic Commerce*, pages 59–67, Newport Beach, CA, USA, 2003. IEEE Computer Society.
- [2] P. Cohen. *Empirical Methods for Artificial Intelligence*. MIT Press, Cambridge, Massachusetts, 1995.
- [3] P. Faratin, C. Sierra, and N. Jennings. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems*, 24(3-4):159–182, 1997.
- [4] S. E. Lander and V. R. Lesser. Sharing meta-information to guide cooperative search among heterogeneous reusable agents. *IEEE Trans. Knowl. Data Eng.*, 9(2):193–208, 1997.
- [5] A. R. Lomuscio, M. Wooldridge, and N. R. Jennings. A classification scheme for negotiation in electronic commerce. *Int. J. of Group Decision and Negotiation*, 12(1):31–56, 2003.
- [6] T. D. Nguyen and N. R. Jennings. A heuristic model for concurrent bi-lateral negotiations in incomplete information settings. In *Proceedings of the 18th Internatinal Joint Conference on AI*, pages 1467–1469, Acapulco, Mexico, 2003.
- [7] T. J. Norman, A. Preece, S. Chalmers, N. R. Jennings, M. Luck, V. D. Dang, T. D. Nguyen, V. Deora, J. Shao, A. Gray, and N. Fiddian. Agent-based formation of virtual organisations. *Int. J. Knowledge Based Systems (to appear)*, 2004.
- [8] I. Rahwan, R. Kowalczyk, and H. H. Pham. Intelligent agents for automated one-to-many e-commerce negotiation. *Twenty-Fifth Australian Computer Science Conference*, 4:197–204, 2002.
- [9] A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, 1982.
- [10] B. Yu and M. P. Singh. Searching social networks. In Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multiagent Systems, pages 65–72, Melbourne, Australia, 2003.