

Managing commitments in multiple concurrent negotiations

Thuc Duong Nguyen

*MLB1, PP12, B62 Orion building (B62-MH)
Austral Park, Martlesham Heath
Ipswich IP5 3RE, UK*

Nicholas R. Jennings

*School of Electronics and Computer Science
University of Southampton
Southampton SO17 1BJ, UK*

Abstract

Automated negotiation by software agents is a key enabling technology for agent mediated e-commerce. To this end, this paper considers an important class of such negotiations — namely those in which an agent engages in multiple concurrent bilateral negotiations for a good or service. In particular, we consider the situation in which a buyer agent is looking for a single service provider from a number of available ones in its environment. By bargaining simultaneously with these providers and interleaving partial agreements that it makes with them, a buyer can reach good deals in an efficient manner. However, a key problem in such encounters is managing commitments since an agent may want to make intermediate deals (so that it has a definite agreement) with other agents before it gets to finalize a deal at the end of the encounter. To do this effectively, however, the agents need to have a flexible model of commitments that they can reason about in order to determine when to commit and to decommit. This paper provides and evaluates such a commitment model and integrates it into a concurrent negotiation model.

Key words: agents, automated negotiation, commitments, concurrent negotiation

Email addresses: duong.nguyen@bt.com (Thuc Duong Nguyen),
nrj@ecs.soton.ac.uk (Nicholas R. Jennings).

1 Introduction

Automated negotiation is a key form of interaction in agent-based systems and such negotiations exist in many different forms (1). 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 (2). In more detail, 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. Traditionally, this type of encounter is handled via some form of single-sided (reverse) auction protocol. However, in previous work, we introduced multiple, concurrent bilateral negotiations as an alternative (3), (4). 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 multi-dimensional contracts, 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. Furthermore, the buyer in our approach can deploy different strategies when bargaining with different types of providers. 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 through 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.

Despite these advantages, however, the negotiation protocol in our previous approach was somewhat unrealistic since it was heavily biased in favor of the buyer. Thus, during the negotiation, the buyer agent could make a number of

intermediate deals with various sellers (where each such deal is a temporary agreement with a specific seller). Then, when its deadline is reached, the buyer selects the most profitable deal as the *final agreement* and declines others. It can operate in this way because these intermediate deals are assumed to be binding on the sellers (meaning they are not allowed to renege from deals once committed) but not on the buyer. Although these unbreakable commitments make it easier for the buyer to achieve good deals, it is highly disadvantageous for the sellers. Thus, in order to be applicable in realistic negotiation situations, the model needs to be extended so that it can deal with situations in which the providers can also renege from deals. To deal with this situation, we develop a commitment manager and an associated reasoning model that enables the agent to behave in a flexible and efficient manner.

To date, a number of commitment models have been developed, each with its own advantages and disadvantages (see section 4 for more details). We base our model on the notion of *leveled commitment contracts* (5) in which an agent can decommit (for whatever reason) simply by paying a decommitment fee to the other agent. In so doing, our work advances the state of the art in the following ways. First, it allows the participating agents to be able to renege from deals whenever they deem appropriate, simply by paying a decommitment fee to their counterparts. Since the providers are no longer forced to be tied to their commitments, they have greater freedom in their behaviors. Second, the agents in our model have different deliberation mechanisms for various penalty levels, thus, they can flexibly perform in a wide variety of e-marketplaces. Finally, our commitment model allows the buyer agent to have a trade-off between the number of agreements it makes and their utility values. This capability helps it to effectively select different commitment strategies according to its purchasing objectives.

The remainder of the paper is organized as follows: section 2 details our new bargaining model and section 3 presents the initial experimental results. Section 4 relates the model to current work in the field and, finally, section 5 presents the conclusions.

2 The negotiation model

The foundation of this work is the concurrent negotiation model outlined in (4). Building on this, the main contribution of this paper is the integration of the ability to reason about commitment and decommitment for the intermediate agreements. Before we can focus on this new ability, however, we first need to recap the basic architecture of our model.

In more detail, 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. This deadline is also the time when the service will be performed by the chosen seller. 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 negotiation thread (bargaining with a particular seller) follows a Sequential Alternating Protocol (7) where at each step an agent can either accept the offer from the opponent, propose its counter-offer, renege from its commitment or opt out of the negotiation (typically if its deadline is reached).

In more detail, the model for the buyer agent consists of three main components: a *coordinator*, a number of *negotiation threads* and a *commitment manager* (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. The coordinator decides the negotiation strategies for each thread. After each round, the threads report back their status to the coordinator. If a thread reaches a deal with a particular seller, it terminates its negotiation and waits until the deadline $t_{b_{max}}$ is reached. 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 commitment manager, which is newly introduced in this work, handles any issue that is related to commitment and decommitment. It is involved when a thread needs to decide whether or not to accept a proposed offer (it makes the decision based on the buyer's current commitment and its commitment strategy; see section 2.3 for more detail) or when a seller decides to renege from a committed deal (it updates its status accordingly). The result of the commitment manager (either accept or reject) will be passed through the coordinator for cross checking with other threads before getting back to the calling thread. The detailed working of the three components are described in the following subsections.

2.1 The coordinator

The coordinator is responsible for coordinating all the negotiation threads and choosing an appropriate negotiation strategy for each thread. Before starting

¹ Given the time-constrained nature of our encounters, the types of strategy that we consider are the time-dependent family introduced in (6). 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.




Fig. 1. System architecture.

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 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} : $types = \{con, non\}$. This information is represented as a probability distribution over the agent types, which may be based on past experiences, obtained from a trusted third party, or from a system of referrals (8). If no such information is available, all agents are assumed to have a uniform distribution.

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. The latter measures the average utility value of the agreement reached in similar situations.

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:

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

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 (see (4) for more details).

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 $U(\alpha, t')$ is the utility value of the offer that seller agent α made at time t' : ($1 \leq t' \leq t$), according to the buyer agent's preferences. Then seller α is considered a *conceder* if $\forall t' \in [3, t]: \frac{U(\alpha, t') - U(\alpha, t'-1)}{U(\alpha, t'-1) - U(\alpha, t'-2)} > \theta$ 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:

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

2.2 The negotiation threads

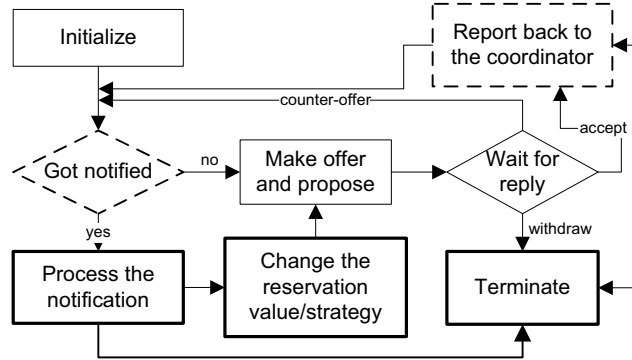


Fig. 2. A single negotiation thread.

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. The structure of a negotiation thread is presented in figure 2. Specifically, each thread is composed of three subcomponents, 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 and the commitment manager. 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 (by cooperating with the commitment manager).

2.3 The commitment manager

Each time the buyer and a seller α decide to agree on an intermediate deal with utility value $U(\alpha, t)$ (according to the buyer's preferences), this deal is binding on both agents. If either of them decides to break the contract, it has to pay a decommitment fee (ρ) to its opponent. This fee is dynamically calculated as a percentage of the utility of the deal² and is also based on the time when the contract is broken³. To this end, the function to calculate the decommitment fee at time $t < t_{b_{max}}$ is chosen as follows:

$$\rho(t) = U(\alpha, t) \times \left(\rho_0 + \frac{t - t_\alpha}{t_{b_{max}} - t_\alpha} \times (\rho_{max} - \rho_0) \right) \quad (2)$$

where t_α is the contract time, when the deal is agreed upon, ρ_0 is the initial penalty (the fee to pay if the deal is broken at contract time, t_α) and $\rho_{max} \geq \rho_0$ is the final penalty (the fee if the deal is broken at execution time, $t_{b_{max}}$).

By means of an illustration, consider the following example. Assume the buyer's deadline ($t_{b_{max}}$) is 10, the initial penalty (ρ_0) is 5% and the final

² Traditionally, there are two ways of calculating the decommitment fee, namely *fixed value* (all contracts have the same fixed decommitment fee that is decided prior to the negotiation) and *percentage of contract value* (the fee is defined as a percentage of the utility value of the contract). It has been empirically demonstrated that the latter type allows the agents to be more flexible in deliberating about their behaviors and enables them to gain a higher utility value than the former (9). Consequently, we use the percentage of contract value in our model.

³ This factor is incorporated to discourage the agent from dropping its commitment towards the end of the negotiation (where it is more difficult to draft in a replacement). Consequently, the later an agent decommits, the more it has to pay.

penalty (ρ_{max}) is 10%; a deal with the expected utility value 0.58 was made at time 6. At time 9, if the buyer wants to decommit, by (2), it has to pay:

$$\begin{aligned}\rho(t) &= 0.58 \times \left(0.05 + \frac{9-6}{10-6} \times (0.10 - 0.05) \right) \\ &= 0.58 \times 0.0875 \\ &= 0.05075\end{aligned}$$

Since the buyer agent now has to pay a fee every time it breaks a contract, it cannot simply just agree on all deals and, later, select the highest value deal as the final agreement (as it did in the original version of our model). Thus, when presented with a potential agreement from a specific seller, the buyer has to decide whether it should take this deal or reject it. In some cases, by rejecting this agreement and, later on, committing to another deal, the buyer will gain a better utility value (see section 3 for more details). To capture this, when presented with a contract $\phi(\alpha)$ that has utility value of $U(\alpha, t)$ from seller α at time t , the buyer will accept $\phi(\alpha)$ as an intermediate deal (and renege on its current commitment, if one exists) if all of the following conditions are satisfied:

- (1) If it already has a commitment with another agent α' at time $t_{\alpha'}$ and this deal has not been broken, the utility gained by taking this new offer must be greater than that of the current deal, after having paid the decommitment fee. This means $U(\alpha, t) > U(\alpha', t_{\alpha'}) + \rho(t)$.
- (2) The degree of acceptance (μ) for $\phi(\alpha)$ must be over a predefined threshold (τ). This threshold specifies how the buyer should accept the offers, whether it is *greedy* (tends to accept any possible deal) or *patient* (only deals that provide a certain expected utility value will be accepted). μ is calculated by comparing the utility value of $\phi(\alpha)$ with the predicted utility value of the next set of contracts from other sellers, also taking into account the relation between the current time and the buyer's deadline. Specifically, the formula for calculating μ is:

$$\mu(\phi(\alpha)) = \frac{U(\alpha, t) - \rho(t)}{\max\{U_{exp}(\alpha_i, t) \mid \alpha_i \in A_s \setminus \alpha\}} \times \frac{t}{t_{b_{max}}}, \quad (3)$$

where $\rho(t)$ is the decommitment fee that the buyer has to pay if it has already committed to a deal with another seller (if it has not, $\rho(t)$ is considered to be 0) and $U_{exp}(\alpha_i, t)$ is the predicted utility of the next proposal from seller α_i . The value of $U_{exp}(\alpha_i, t)$ is calculated as:

$$U_{exp}(\alpha_i, t) = U(\alpha_i, t) + \frac{d_U(t, t-1)}{d_U(t-1, t-2)} \times |d_U(t, t-1)|, \quad (4)$$

where $d_U(t_1, t_2)$ is the distance, in terms of utility value, between two offers of seller α_i at time t_1 and t_2 : $d_U(t_1, t_2) = U(\alpha_i, t_1) - U(\alpha_i, t_2)$.

To illustrate the operation of the commitment manager in more detail, consider the following example. There are 4 participating sellers, the buyer's deadline ($t_{b_{max}}$) is 6, the initial penalty (ρ_0) is 10%, the final penalty (ρ_{max}) is 20% and the threshold (τ) is 0.8. The buyer has committed on a deal with seller 4 at time 2 with the expected utility value of 0.21. The utility values of previous offers from all the sellers are displayed in table 1.

<i>agent</i>	t=1	t=2	t=3
α_1	0.03	0.12	0.16
α_2	0.01	0.04	0.10
α_3	0.1	0.19	0.23
α_4	0.11	0.21	-

Table 1
Utility values of the offers.

At time 3, the buyer has to decide whether it will accept the offer $\phi(\alpha_3)$ from seller α_3 . Since it is already committed to a deal with α_4 , if it wants to take $\phi(\alpha_3)$, it will have to pay a decommitment fee to α_4 . By (2), the fee it has to pay is:

$$\rho(3) = 0.21 \times \left(0.1 + \frac{3-2}{6-2} \times (0.2 - 0.1) \right) = 0.026$$

As can be seen, $U(\alpha_3, 3) < U(\alpha_4, 2) + \rho(3)$, so the first condition is violated. Thus, the buyer will reject $\phi(\alpha_3)$ and remain with its commitment with α_4 .

<i>agent</i>	t=1	t=2	t=3	t=4	t=5
α_1	0.03	0.12	0.16	0.28	0.4
α_2	0.01	0.04	0.10	0.30	0.26
α_3	0.1	0.19	0.23	0.31	0.36
α_4	0.11	0.21	-	-	-

Table 2
Utility values of the offers (cont.).

At time 4, however, seller α_4 decides to renege on its current deal and pay the decommitment fee to the buyer. According to equation (2), it has to pay:

$$\rho(4) = 0.21 \times \left(0.1 + \frac{4-2}{6-2} \times (0.2 - 0.1) \right) = 0.0315$$

As can be seen, this decommitment from α_4 leaves the buyer with no agreement. Now, at time 5, the buyer has to decide if it should take up the offer from α_1 (table 2 shows the utility values of the offers from all the sellers). Since it has no intermediate agreement, the first condition is satisfied. To evaluate the second condition, the buyer first calculates the value for $U_{exp}(\alpha_1, 5)$ and $U_{exp}(\alpha_2, 5)$ using (4):

$$U_{exp}(\alpha_2, 5) = 0.26 + \frac{-0.04}{0.2} \times 0.04 = 0.252$$

$$U_{exp}(\alpha_3, 5) = 0.36 + \frac{0.05}{0.08} \times 0.05 = 0.391$$

The value of $\mu(\phi(\alpha_1))$ is then calculated, using equation (3), as:

$$\mu(\phi(\alpha_1)) = \frac{0.4}{0.391} \times \frac{5}{6} = 0.852$$

This time, since $\mu(\phi(\alpha_1)) > \tau$, the buyer will commit to this deal. It keeps on bargaining in this way until its deadline is reached. If, at that time, there is an intermediate deal that has not been broken, this deal is selected as the final agreement. If, however, no such deal exists, the negotiation is considered unsuccessful and terminated without an agreement.

Up until this point, we have only considered the situation where the buyer agent commits to a maximum of one intermediate contract at one time. However, it is possible for the buyer to commit to more than one contract at any one time and then, later, select the best one and decommit from the others. This represents a cautious approach and avoids the risks associated with committing to only one contract which is then revoked near the deadline, leaving the agent with insufficient time to find a replacement. The downside of this approach, however, is that if the sellers do not renege, the buyer may end up paying a significant part of its utility value as the penalty fee. Nevertheless, in some cases, it may be beneficial for the buyer to consider the option of having multiple commitments during the bargaining process. To capture this, assume that the maximum number of commitments that the buyer will hold at any one time is $\omega \geq 1$ and let $\Omega = \{C_i \mid i \in [1, \omega]\}$ be the set of contracts that the buyer is currently committed to: $|\Omega| \leq \omega$. Here, $C_i = \{\phi', t'\}$ is the contract that consists of an intermediate deal ϕ' that has been agreed at time t' . Assuming that at time t , there is an offer ϕ from seller k that has $\mu(\phi(k)) > \tau$ (see equation 3⁴) and $U(\phi(k), t) > U(\phi', t_k) + \rho(t) \quad \forall C(\phi', t_k) \in \Omega$ then this

⁴ If $|\Omega| < \omega$, $\rho(t)$ is considered to be 0. If not, $\rho(t)$ is the penalty the buyer will have to pay to break from the contract $C(\phi', t_k) \in \Omega$ that has the minimum value

offer will then be accepted by the buyer. This, in turn, means the following steps will be taken:

- (1) If Ω is not full (i.e. $|\Omega| < \omega$), $C(\phi, t)$ will be added to Ω : $\Omega = \Omega \cup C(\phi, t)$
- (2) If Ω is full (i.e. $|\Omega| = \omega$) select $C(\phi', t_{k'}) \in \Omega$ that has the minimum value of $U(\phi', t_{k'})$ and decommit from that contract. Then, $C(\phi, t)$ will be added to Ω : $\Omega = \Omega \cup C(\phi, t)$

Now if a seller k that has a contract $C(\phi(k), t_k)$ decides to withdraw its commitment, that contract will be subtracted from Ω : $\Omega = \Omega \setminus C(\phi(k), t_k)$. Then at the end of the bargaining process, if there is more than one contract in Ω , the buyer simply selects the one that has the highest utility value as the final agreement and decommits from all the others.

3 Empirical evaluation

Having introduced the commitment manager, the next step is to evaluate its effects on the model. We choose *empirical evaluation* 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.

3.1 Experimental setup

In more detail, we use the *exploratory studies* evaluation technique (10). With this method, general hypotheses are formed to express the intuitions about the causal factors within the model. The experiments are then conducted and generate the results that either support these hypotheses or go against them. In our evaluation, the independent variables are given in table 3 and the dependent ones are listed in table 4.

Apart from the control variables described in table 3, other control variables are selected as per (11). Specifically, the number of seller agents (n) is set in the range of [1, 30] and the number of negotiation issues (m) is set in the range of [1, 8]. An agent α 's preference for issue j is represented by the tuple $\{x_{jmin}^\alpha, x_{jmax}^\alpha, w_j^\alpha\}$. The tuple $[x_{jmin}^\alpha, x_{jmax}^\alpha]$ is an interval independent variable, $\overline{\text{of } U(\phi', t_{k'})}$.

<i>Variables</i>	<i>Descriptions</i>	<i>values</i>
ρ_0	the initial penalty fee	[5,100]
ρ_{max}	the final penalty fee ($\rho_{max} \geq \rho_0$)	[5,100]
τ	the μ threshold	[0,1.5]
ω	the number of concurrent commitments	[1,4]

Table 3

The independent variables.

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 $x_{j_{min}}^\alpha$ to be in the interval $[0, 20]$ and $x_{j_{max}}^\alpha$ 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 50 (long deadline). The penalty fee (both initial and final) is also an ordinal independent variable, whose value is randomly chosen, ranging from 5% (small) to 100% (equal to the value of the contract). Similarly, the τ threshold is either 0 (meaning the buyer is greedy and will commit to any intermediate deal that it can get hold of) or 0.5 (meaning the buyer is patient and will only engage on a deal that provides high expected utility value)⁵.

The seller agents in this evaluation are characterized in a similar fashion to ones set up in our previous experiments (4). Specifically, they 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 $x_{j_{min}}^b \leq x_{j_{min}}^\alpha \leq x_{j_{max}}^b \leq x_{j_{max}}^\alpha$.
- *the negotiation strategy*: Each seller is assigned a random strategy selected from a predefined set of alternations (as outlined in (6)). This set is composed of time-dependant 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.

The only difference is that now if a seller has committed to a deal, it has a chance of being made an outside offer with the utility value of 1.0 (which is the highest possible utility value). Thus, there is a probability that it will

⁵ Future work will investigate in more detail how this value affects the outcome of the model.

decommit. To this end, we consider three types of sellers:

- *loyal*: once a seller has committed to an intermediate deal, it will not renege from it.
- *loose*: a seller always breaks a committed deal if it is presented with a better option.
- *partial*: if a seller finds a better option, it will break a committed deal with a percentage of probability (as per (12)). In this experiment, we set this percentage to be 50%, meaning that half of the time a seller finds a better deal, it will renege and half of the time it will stay with its current deal.

<i>Variables</i>	<i>Descriptions</i>
U	the utility value of the final agreement
N	the number of successful negotiations
D	the number of decommitments made by buyer

Table 4

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:

$$U(X) = \sum_{j=1}^m w_j^\alpha \cdot \frac{x_j - x_{jmin}^\alpha}{x_{jmax}^\alpha - x_{jmin}^\alpha}$$

We also measure the number of agreements reached at the end of the negotiation encounter (N) and the average number of decommitments that the buyer made (D). In all cases, the results are gathered from a series of experiments in different environment settings. Each experiment consists of 1000 runs and the results are averaged and put through a regression test to ensure that all differences are significant at the 99% confidence level.

3.2 *Experimental hypotheses*

We now turn to the specific hypotheses.

Hypothesis 1 *When dealing with loose or partial sellers, the higher the penalty fee is, the lower the number of final agreements reached by the buyer.*

To evaluate this hypothesis, we measure the number of final agreements achieved with varying types of seller agents (see figure 3). As can be seen, the number of final agreements reached by the buyer is dramatically reduced as the penalty




Fig. 3. Number of successful negotiations for varying penalty fee.

fee is increased. Specifically, when dealing with loose sellers, around 97% of the negotiations are successful when the penalty fee is 5%. As the penalty fee increases to 100%, this success rate drops down to only 84%. Similarly, the figures when dealing with partial sellers are 98% and 92%, respectively. This decreasing trend is explained by the deliberation mechanism of the buyer. Specifically, assume that the buyer has already made a commitment with seller k and now it is presented with another offer from seller k' . If it decides to take this new offer from k' , it will have to pay k a decommitment fee ρ . As the penalty fee is increased, so is ρ . Thus, in some cases, the buyer cannot afford to take this new offer and it has to stay with its commitment to k . Later on, if k decides to break its commitment, the buyer is left with no intermediate agreement. As such, there may not be enough time for the buyer to find another replacement deal and, thus, no final agreement can be reached. On the other hand, if the buyer can take the offer from k' , the probability that k' will renege is less than that of k . Thus, a final agreement can be reached.

Another observation is that the more loyal the seller is, the greater the number of final agreements that the buyer makes. This difference is caused by the probability of the sellers breaking their commitments. Since a loyal seller never reneges, once it has committed, its contract is kept until either it is declined by the buyer or it is selected as the final agreement. Therefore, once an intermediate deal is reached, a final agreement is always guaranteed to exist. However, this is not the case for the other types of sellers. Once they have committed, it is not guaranteed that they will actually stay faithful with their commitments. If a seller breaks a contract, the buyer has to find a replacement. If it fails to do so, no final agreement will be achieved. Thus, the less loyal the sellers are, the fewer chances there are for the buyer to reach a final agreement.

Hypothesis 2 *The higher the penalty fee, the lower the utility of the final agreement gained by the buyer.*

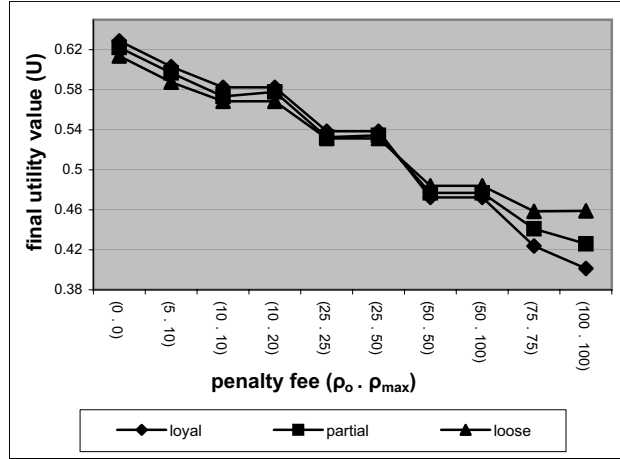


Fig. 4. Final utility value for varying penalty fee.

As can be seen from figure 4, this trend is true for all seller types. Specifically, when dealing with loose sellers, the average utility of the final agreement for the buyer drops from 0.61 to 0.46 when the penalty fee goes from 5% to 100%. The corresponding figures for partial and loyal sellers are 0.62 to 0.43 and 0.63 to 0.40, respectively. The reason for this decrease in the final utility value is that the higher penalty fees mean more chance that the buyer will commit to an early agreement (and stay with this commitment until either its deadline is reached or the corresponding seller decides to renege). These early commitments by the buyer have two main effects. First, such agreements tend to have lower utility value for the buyer, compared to the contracts that are offered at a later stage (the buyer cannot afford to take these contracts due to high decommitment fees). Second, once that commitment is later broken, the buyer will have to find a replacement. Even if it is successful in finding one, since there is not much time for bargaining, the utility value of this newly found agreement is likely to be less than that of the previous deal. Consequently, the utility gained by the buyer is reduced.

Furthermore, with increasing penalty fee, the more loyal the seller, the lower the value of the final agreement gained by the buyer (see figure 4). The reason for this observation is because the buyer benefits from the decommitment fee gained when a seller reneges from a committed deal. As per our experimental setup, loose sellers decommit more often than partial sellers and loyal sellers never renege. Thus, as the penalty fee increases, the buyer will benefit more when dealing with less loyal sellers.

Hypothesis 3 *The buyer decommits less frequently as the penalty fee increases.*

Figure 5 shows the average number of decommitments made by the buyer for varying penalty fees and different seller types. Since the buyer's deliberation includes the decommitment fee it has to pay if it want to replace its current




Fig. 5. Number of buyer’s decommitments for varying penalty fee.

intermediate deal (see equation 2), the less it has to pay, the more favorable it will be to take up a better deal. Thus, even when a seller offers an intrinsically higher value contract than the current deal it has, the buyer may be better off sticking with its existing commitment in order to avoid paying a hefty fine. This is why the buyer almost never reneges when the penalty fee is close to 100%.

Hypothesis 4 *The more patient the buyer, the higher the utility for the final agreement. However, the chance of having a final agreement is reduced.*

We start by looking at the performance of the model with a number of different values for the degree of acceptance (specifically $\tau \in [0.1, 1.5]$). For simplicity, we fix the penalty fee value at ($\rho_0 = 5, \rho_{max} = 10$) and assume we are dealing with *partial* sellers. These values are chosen just to give us an idea of how the results could potentially be and a detailed analysis will follow. The results are displayed in figure 6.

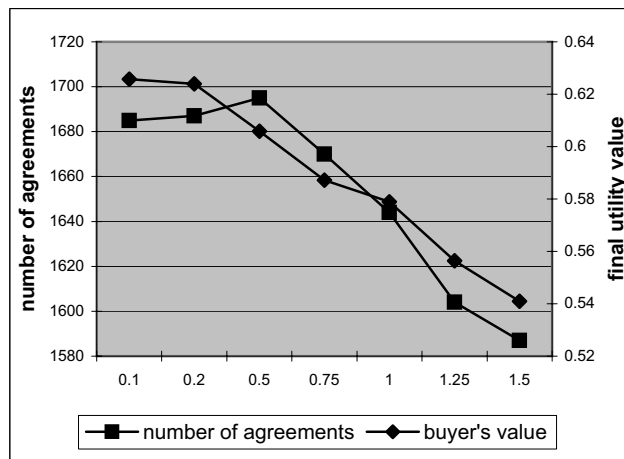


Fig. 6. Performance vs degree of acceptance.

As can be seen, as the value of τ increases, the utility value for the final agreement decreases. This is because in a particular negotiation, if the buyer tends to ignore the current offer from the seller, in favor of a higher value one at a later time, there is a possibility that a high value offer will not be forthcoming (e.g. the seller may run out of time or be at the limit of its reservation values). Thus, towards the end of the encounter it will have to settle for a lower value deal (because this is better than no deal). This, in turn, puts a downward trend on the final utility value achieved.

On the other hand, the number of final agreements reached increases as the value of τ increases up to 0.5, then it decreases. Now, since we are dealing with partial sellers, if they are presented with a better outside offer, they have the chance to renege and may leave the buyer with no agreement at hand. When the value for τ is small (less than 0.5 in this case), the buyer tends to take up any offers that are available to it at an early time. Later, when the seller that is sharing the commitment with the buyer decides to back down, there might not be enough time for the buyer to recover from this loss and thus it might end up with no agreement at the end of the encounter. However, if it is too strict on accepting intermediate deals, it also risks the chance of having obtained no deal at all. This is the situation when the value of τ increases past 0.5. From figure 6, it can be seen that by setting value of τ at around 0.5, the buyer will achieve the highest number of final agreements with a reasonably good final utility value.

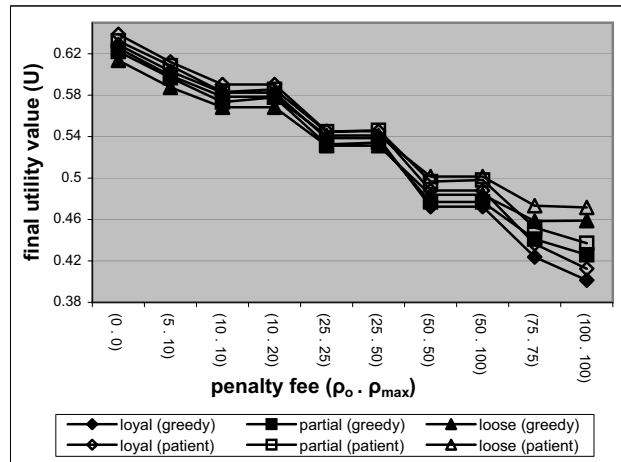


Fig. 7. Final utility value for varying penalty fee.

We extend the aforementioned result by comparing the results of having two different values for τ : *greedy* ($\tau = 0$) and *patient* ($\tau = 0.5$) in the experiments with different penalty values, as well as different seller types. Recall, the greedy buyer will commit to any offer that it can take (if it is more beneficial than the one it currently has, taking into account the decommitment fee it will have to pay). In contrast, the patient buyer will only commit to an offer that has significantly greater value (compared with the one that it currently has). As it

only accepts higher value contracts compared to its counterpart, the patient agents' final agreements always have higher utility value than those of the greedy agent (see figure 7).

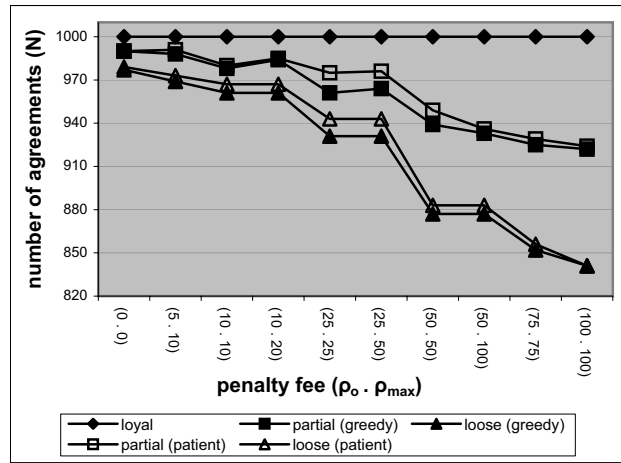


Fig. 8. Number of successful negotiations for varying penalty fee.

Now, not only does the patient agent gain higher utility value, the number of successful agreements achieved is also higher than or, at least, equal to that of the greedy agent (see figure 8). The reason for this is related to the way an intermediate agreement is accepted by the buyer. The greedy buyer accepts a higher number of intermediate agreements than its counterpart⁶. Thus, its chance of having an agreement reneged upon is higher than that of the patient agent. In some cases, this decommitment limits the chance of the buyer of having an agreement at the end of the negotiation. Consequently, the patient agent will be able to reach more agreements than the greedy one at the end of the bargaining process.

However, even though it can gain better utility value than its greedy counterpart, the patient agent manages to get fewer agreements than its counterpart. This is because the patient agent only accepts a deal if the degree of acceptance (μ) of this deal is greater than a threshold (in this case, $\tau = 0.5$). Thus, not all the deals proposed by the sellers satisfy this condition. Indeed, in some cases, none of the proposed contracts satisfy this condition. This limits the chance of the buyer having an agreement at the end of the negotiation. On the other hand, the greedier the agent is, the higher the chance that an offer will be accepted. Consequently, the greedy agent will be able to reach more agreements than the patient one at the end of the bargaining process.

Hypothesis 5 *When dealing with loyal sellers, the buyer is better off committing to a maximum of one contract at any one time. For other seller types,*

⁶ The greedy buyer accepts any possible agreement, whereas the patient one only accepts agreements that have a significantly greater value compared with the one that it currently has.

the buyer should commit to a maximum of two contracts.




Fig. 9. Number of agreements vs buyer's maximum commitments.

Figure 9 shows the number of agreements obtained by the buyer at the end of the encounter when it varies the number of commitments it can hold at any one time (here $\omega \in [1, 4]$). As can be seen, when holding more than one commitment, the buyer increases its chance of reaching an agreement when dealing with non-loyal sellers. In particular, when ω is increased from 1 to 2, the buyer gains 0.9% more final agreements when dealing with partial sellers and 2% more when dealing with loose sellers. This improvement can be explained simply by looking at the behaviors of the sellers. As the sellers are not loyal, when presented with an outside offer, they may renege. If this happens near the end of the negotiation process and the buyer can only commit to a single contract, it will leave the buyer very little time to find an alternative (and in some cases it will not be able to do so). On the other hand, if the buyer is holding more than one contract and an agent reneges then it has something that it can fall back on and it is less vulnerable to being left with no agreement. For values of $\omega > 2$, however, the improvement is comparatively minor because when the buyer is committing to more than one contract, the chance that all the sellers renege is significantly reduced compared to the situation when the buyer can only have one commitment at a time. This, in turn, has a very slight impact on the number of agreements achieved.

The final utility achieved by the buyer with varying values for ω is displayed in figure 10. As can be seen, when $\omega > 2$, the utility gained is dramatically reduced (8% decrease when ω goes from 2 to 3 and 16% decrease when ω goes from 3 to 4). This is because if a buyer agent has more contracts at the end of the negotiation process, it will end up paying a significant penalty fee for breaking them. When $\omega = 2$, the situation is similar to that of dealing with loyal sellers. However, when negotiating with partial or loose sellers, $\omega = 2$ gives similar results and, in some cases, is better than setting ω to 1. The reason for this is because as the non-loyal seller agents can renege on their

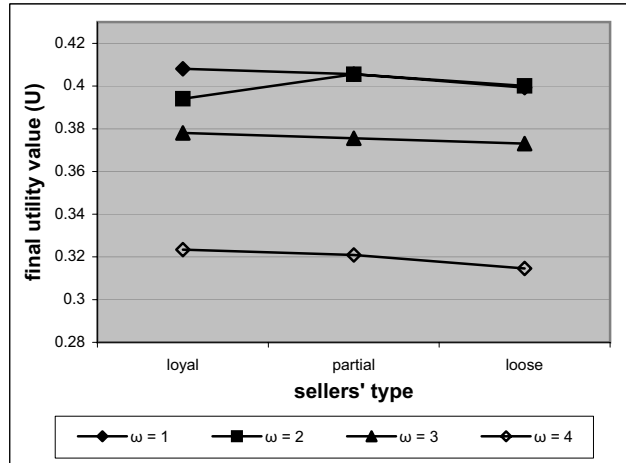


Fig. 10. Final utility value vs buyer's maximum commitments.

commitments and if they do so towards the end of the negotiation process, the buyer will gain additional penalty fees from those sellers and is still left with at least one intermediate contract in hand. Thus, at the end, it is still able to have the final agreement, but it does not have to pay a decommitment fee to any other seller agent.

As can be seen, when dealing with loyal sellers, the buyer does not necessarily need to have more than one commitment since it can be sure that the sellers will never renege from their deals. However, when dealing with partial or loose sellers, the situation is different. The greater the number of commitments it holds, the higher the number of final agreements it is likely to obtain. Nevertheless, the final utility value reached is decreased because it will have to pay a large amount of penalty fees to decommit from these commitments. To this end, the buyer is best setting ω to 1 when dealing with loyal sellers and ω to 2 when dealing with other types of seller to ensure that it will achieve highest possible utility value together with an acceptable number of final agreements.

4 Related work

Traditionally, once a contract is made in a negotiation, it is binding on all participants. Neither party can back out no matter what happens in the future (13), (14), (15). This is also the case for existing concurrent negotiation models (4), (16), (17). However, this view is very limiting for the agents and it may lead to irrational and inefficient behavior (12). As a result, a number of methods have been developed to overcome this limitation.

One of the first pieces of work in this area was the contract net protocol (18), where there is a possibility for a decommitment. Here, the contracter agent could send a termination message to cancel the contract, even when a part

of it has been fulfilled by the contractee. As the agents participating in a contract net are generally assumed to be cooperative, they do not mind losing their effort (even without any form of compensation). In a similar fashion, the role of commitment for cooperative agents was examined in the context of automated scheduling of meetings (19). In e-commerce settings, however, these models are inappropriate because the agents are not always cooperative and they seek to maximize their individual gains.

For self-interested agents, *contingency contracts* have been introduced as a method of allowing them to break commitments (20). In this case, an agent's commitment to a contract is made contingent on specific future events. Thus, if these specified contingencies arise, the agents are allowed to drop their commitments (21). However, there are a number of problems associated with this type of contract (5). First, not all possible future events are known to the agents beforehand, thus, they cannot always make optimal use of contingency contracts. Second, this type of contract is useful when the number of future events is small. If, however, this number increases, it is cumbersome or even impossible for all the events to be monitored. Furthermore, these events may not affect the original contract independently, they may have a combined effect on the value of the contract (13). As a result, this approach is not adopted in our work.

The most advanced work in the area, and also the basis to our work, is the *leveled commitment contracts* (LVC) (5). Our commitment manager is built upon the same basic intuition that any agent can freely decommit from a contract, for whatever reason they deem appropriate, by simply paying a decommitment fee to the other partner. However, our model is different in a number of important ways. First, the original LVC only covers a two person game. We have extended this to cover the multiple providers found in our target environment. Second, we do not just reason about decommitment, we also deliberate about when and how to make a commitment. Third, LVC require the agents to have information about the actual and alternative options of their opponents in order to be able to calculate the Nash equilibrium decommitment threshold. This assumption is unrealistic in practical scenarios and is not required in our model. Finally, unlike LVC (which typically assumes a fixed penalty for decommitting, regardless of the stage of the process at which the commitment is broken), our model takes the cost of ongoing commitment into account by introducing variable penalty contracts. Again, we believe this is more realistic for most real-world settings.

5 Conclusions and future work

This paper has introduced a commitment handling capability that can be applied in managing concurrent negotiations in time-constrained settings. This ability increases the flexibility and realism of the participating sellers and relaxes the previous unrealistic constraints we imposed (3; 4). Our empirical results have highlighted the fact that different penalty levels have different effects on the performance of the model. In addition, we show that the more patient the buyer is, the better the deal it will obtain. Nevertheless, if the buyer wants to secure more agreements, it should be greedier in making commitments. Our extended 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 (22) and in an internal project of BT concerned with logistics planning (23).

For the future, there are a number of ways in which our model can be improved. First, we would like to experiment with different strategies (e.g. alternative methods for the buyer to decide whether or not to accept an offer from a seller) for our buyer agent to see how they affect the final outcome of the model. Second, we would like to investigate different penalty levels for different agents, perhaps based on their negotiation histories. In particular, if an agent comes to the negotiation with poor reputation (e.g. it frequently reneged on its previous encounters), that agent should have to pay a higher penalty fee than one that has shown itself to be more trustworthy. Finally, we would like to improve the decision making of our agents so that they can make more accurate predictions about their opponents' decommitment strategies. This will, we believe, also increase the performance of the model.

References

- [1] A. R. Lomuscio, M. Wooldridge, N. R. Jennings, A classification scheme for negotiation in electronic commerce, *Int. J. of Group Decision and Negotiation* 12 (1) (2003) 31–56.
- [2] M. P. Singh, M. N. Huhns, *Service-Oriented Computing: Semantics, Processes, Agents*, John Wiley and Sons, Ltd., 2005.
- [3] T. D. Nguyen, N. R. Jennings, A heuristic model for concurrent bi-lateral negotiations in incomplete information settings, in: *Proceedings of the 18th International Joint Conference on AI, Acapulco, Mexico, 2003*, pp. 1467–1469.
- [4] T. D. Nguyen, N. R. Jennings, Coordinating multiple concurrent negotiations, in: *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multi Agent Systems, New York, USA, 2004*, pp. 1064–1071.

- [5] T. W. Sandholm, V. R. Lesser, Leveled commitment contracts and strategic breach, *Games and Economic Behavior* 35 (2001) 212–270.
- [6] P. Faratin, C. Sierra, N. Jennings, Negotiation decision functions for autonomous agents, *Robotics and Autonomous Systems* 24 (3-4) (1997) 159–182.
- [7] A. Rubinstein, Perfect equilibrium in a bargaining model, *Econometrica* 50 (1) (1982) 97–109.
- [8] S. E. Lander, V. R. Lesser, Sharing meta-information to guide cooperative search among heterogeneous reusable agents, *IEEE Trans. Knowl. Data Eng.* 9 (2) (1997) 193–208.
- [9] M. Andersson, T. Sandholm, Leveled commitment contracts with myopic and strategic agents, in: *In Proceedings of the Fifteenth National Conference on Artificial Intelligence, 1998*, pp. 38–44.
- [10] P. Cohen, *Empirical Methods for Artificial Intelligence*, MIT Press, Cambridge, Massachusetts, 1995.
- [11] T. D. Nguyen, N. R. Jennings, Reasoning about commitments in multiple concurrent negotiations, in: *Proceedings of the 6th International Conference on E-Commerce, Delft, The Netherlands, 2004*, pp. 77–84.
- [12] C. B. Excelente-Toledo, R. A. Bourne, N. R. Jennings, Reasoning about commitments and penalties for coordination between autonomous agents, in: *Proceedings of the 5th Int Conf on Autonomous Agents (Agents-2001), Montreal, Canada, 2001*, pp. 131–138.
- [13] J. S. Rosenschein, G. Zlotkin, *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*, MIT Press, Cambridge, USA, 1994.
- [14] P. Faratin, Automated service negotiation between autonomous computational agents, Ph.D. thesis, Queen Mary College, London, England (2001).
- [15] S. Kraus, *Strategic Negotiation in Multi-Agent Environments*, MIT Press, Cambridge, USA, 2001.
- [16] I. Rahwan, R. Kowalczyk, H. H. Pham, Intelligent agents for automated one-to-many e-commerce negotiation, *Twenty-Fifth Australian Computer Science Conference* 4 (2002) 197–204.
- [17] A. Byde, M. Yearworth, K. Y. Chen, C. Bartolini, Autona: A system for automated multiple 1-1 negotiation, in: *Proceedings of the 2003 IEEE International Conference on Electronic Commerce, IEEE Computer Society, Newport Beach, CA, USA, 2003*, pp. 59–67.
- [18] R. G. Smith, The contract net protocol: High-level communication and control in a distributed problem solver, *IEEE Transactions on Computers* 29 (12) (1980) 1104–1113.
- [19] S. Sen, E. Durfee, The role of commitment in cooperative negotiation, *International Journal on Intelligent and Cooperative Information Systems* 3 (1) (1994) 67–81.
- [20] H. Raiffa, *The Art and Science of Negotiation*, Harvard University Press, Cambridge, USA, 1982.
- [21] N. R. Jennings, *Commitments and conventions: The foundation of co-*

- ordination in multi-agent systems, *The Knowledge Engineering Review* 8 (3) (1993) 223–250.
- [22] T. J. Norman, A. Preece, S. Chalmers, N. R. Jennings, M. Luck, V. D. Dang, T. D. Nguyen, V. Deora, J. Shao, A. Gray, N. Fiddian, Agent-based formation of virtual organisations, *Int. J. Knowledge Based Systems* 17 (2-4) (2004) 103–111.
- [23] T. D. Nguyen, A heuristic model for concurrent bilateral negotiations in incomplete information settings, Ph.D. thesis, University of Southampton, Southampton, England (2005).