Adaptive, Confidence-Based Strategic Negotiations in Complex Multiagent Environments

Xin Li and Leen-Kiat Soh

Computer Science and Engineering
University of Nebraska-Lincoln
256 Avery Hall, Lincoln, NE 68588
{xinli, lksoh}@cse.unl.edu

Abstract

We have designed and implemented an adaptive, confidence-based negotiation strategy for conducting multiple, concurrent negotiations among agents in dynamic, uncertain, and real-time environments. Our strategy deals with how to assign multiple issues to a set of concurrent negotiations. When an agent is confident about a particular peer agent, it uses a packaged approach by negotiating on multiple issues with that peer in a single negotiation job. Otherwise, it uses a pipelined approach by negotiating with the peer one issue at a time in a sequence of negotiation jobs. Hence, the confidence of an agent's profile or view of other agents is crucial, and that depends on the environment in which the agents operate. Each initiating agent is also motivated to improve both the process and the outcome of its negotiations, taking into account factors such as time spent and messages sent during negotiation. Our experiments show that the adaptive, confidence-based negotiation strategy outperforms the purely pipelined or purely packaged strategy in a variety of aspects.

Introduction

Negotiation is a form of interaction among autonomous agents in which a group of agents with a desire to cooperate but with potentially conflicting interests seek to reach an agreement on a set of issues (Raifa 1982, Walton and Krabbe 1995, Kraus et al. 1998). From the perspective of specific applications (e.g., e-market), a negotiation issue is any good or service that one agent can provide to another. From the perspective of problem solving in multiagent systems, a negotiation issue is a (scarce) resource or capability. Given its ubiquity, a negotiation may address multiple issues or only one issue. In this paper we focus on multi-issue negotiations on task and resource allocation.

For cooperative problem solving, an agent that encounters a problem may initiate 1-to-many negotiations and concurrently negotiate with multiple agents over responding resources or capabilities. During negotiations, both the agent that initiates negotiations (i.e., the initiating agent) and each peer agent that responds to a negotiation request (i.e., the responding agent) need to make long-term strategies to conduct the negotiation processes and take short-term tactics to implement each negotiation step. In this paper, we focus on strategic negotiations instead of detailed negotiation tactics; specifically, we focus on the initiating agent's negotiation strategy on how to conduct concurrent 1-to-many negotiations with multiple issues.

The multiagent environments considered in this paper are complex: dynamic, uncertain, and real-time. The initial states that trigger the agents' negotiation process in the first place may dynamically change while the negotiation process is still going on. The negotiation outcome is uncertain and a negotiation carried out in the same manner may not always yield the same results at different times. Also, each negotiation is time-constrained. In such environments, agents that make good strategic decision-making are expected to conduct better negotiations in the long run than those that do not. However, the complex factors make strategic negotiations difficult.

To address the multi-issue strategic negotiations in complex environments, we propose an adaptive, confidence-based negotiation strategy (Soh and Li 2004). After identifying a set of capable peers for a particular task at hand, the initiating agent needs to decide how to conduct its multiple, concurrent negotiations. If the initiating agent is confident in a peer's consistency, then it uses a packaged approach; otherwise, a pipelined approach. In a pipelined approach, the two negotiating agents negotiate on one issue at a time. In a packaged approach, the two agents negotiate on multiple issues in one negotiation job. Using a pipelined approach, the initiating agent can switch remaining, yet-to-be-negotiated issues to other peers that have completed their negotiation jobs with the initiating agent; while using a packaged approach, no issues can be switched between agents, with the lack of flexibility.

There has been a lot of research work on both tactical negotiations—how to conduct a negotiation with the step-by-step negotiation tactics (e.g., Rosenschein and Zlotkin 1994, Kraus et al. 1998), and strategic negotiations—how to manage negotiations (e.g., Nguyen and Jennings 2004, Fatima et al. 2004). We focus on strategic negotiations to improve negotiation process and outcome. Nguyen and

Jennings (2004) propose a heuristic approach to coordinating multiple concurrent negotiations on a single issue. Since they do not consider multi-issue negotiations, there is no issue assignment work and no issue will be switched from the pipeline of one peer to another peer. Fatima et al. (2004) focus on bilateral multi-issue negotiations. They determine the equilibrium strategies for two negotiation procedures: issue-by-issue and package deal, based on the agents' preferences. Our work focuses on multi-issue negotiations similar to theirs, but our initiating agent manages concurrent negotiations rather than only one bilateral negotiation. The pipelined approach and packaged approach in our strategy are similar to their negotiation procedures: issue-by-issue and package deal. But we focus on possible issue switching rather than the negotiation order. And our agents decide negotiation strategies based on the confidence in other agents' behavior rather than the preferences.

We have designed and implemented the proposed negotiation strategy. Our experiments compare the performance of the proposed strategy with the purely pipelined strategy and the purely packaged strategy and show that the confidence-based one generally outperforms the others. Further, we also investigate how well our strategy works in improving the process and the outcome of negotiations in different environmental settings.

Confidence-Based Negotiation Strategy

The adaptive, confidence-based negotiation strategy is designed and implemented based on (Soh and Li 2004). The overall problem domain is multiagent coalition formation (Soh and Li 2003, Li and Soh 2004). When an agent does not have all resources or capabilities required for the accomplishment of a complex task (e.g., multi-sensor target tracking, disaster rescue), it may initiate a coalition formation process to form a coalition with some others and execute the task together. During coalition formation, the initiating agent ranks all the peer agents based on the current task requirement and the peer agents' past coalition formation behavior profiled, and then it negotiates with the top-ranked peer agents. In a real-time environment, agents need to form coalitions soon enough to meet the task requirements. Due to the uncertain and noisy characteristics in the communication, roles, and resources, it is possible that peers ranked high do not perform as expected. For example, the initiating agent may rank peer A_{j_1} as the best cooperator and expect to reach an agreement for its negotiation in a short time. However, A_{j_1} is busy and unable to entertain the negotiation request. Without a flexible management strategy, the initiating agent would have to wait until A_{j_1} is available, probably missing the time requirement. It is also possible that, because of the dynamic nature that we assume of our negotiation environment, the ranking of a peer by an agent may change during a negotiation and may thus require the agent to terminate the ongoing negotiation in favor of another peer. Thus there is a need for a concurrent negotiation strategy that is flexible, capable of adapting to the behavior of all the negotiating peers as well as the real-time observation of the negotiation activities. Although our original problem domain is coalition formation among agents, the strategy is designed as a generic one oriented to general problem domains to handle strategic negotiations in complex environments.

In each individual negotiation, an initiating agent may employ a pipelined negotiation approach, or a packaged negotiation approach. In the pipelined approach, the initiating agent, lacking knowledge on how the responding agent performs in coalition formation, negotiates only one issue (subtask) in each negotiation job. As one job completes, the agent subsequently negotiates other issues in succeeding negotiation jobs. If the peer cannot negotiate as expected, for example, the communication channel between them is busy, the remaining issues will be moved from the current pipeline to another peer's pipeline. This allows the agent to be *cautious* and *opportunistic* at the same time. In the packaged approach, the initiating agent, with confident knowledge of the past behaviors of the responding agent, packages multiple issues into one single negotiation job since it can anticipate how the responding agent is likely to perform based on its past consistent negotiation experience.

The proposed strategy is based on three assumptions: (1) the *Efficient Multi-Issue Negotiation* assumption, (2) the *Overlapping Capabilities* assumption, and (3) the *Additive Multi-Issue Negotiation Evidence* assumption. The first assumption basically indicates that by packaging issues into one single negotiation, an initiating agent is able to reduce the number of messages it has to send, which in turn reduces the computational and communication costs as well. This first assumption motivates an agent to prefer the packaged approach. The second assumption facilitates the pipelined approach. If the responding agents do not share common capabilities, then the pipelined approach will not have the possibility to exploit the responding agents of early-completing negotiation jobs. The third assumption assumes that to persuade a responding agent to agree with a set of issues, the amount of evidence information (or messages) that the initiating agent has to provide for all the issues in the packaged approach is the total amount of all the evidence information that it has to provide for each issue separately in the pipelined approach. This may not be always true: an agent may gain more leeway, establish more leverage, or lose grounds when negotiating on multiple issues with another agent, such as in a goods bartering scenario, depending on whether the initiating agent has an upper hand.

Figure 1 depicts how an agent with confidence-based negotiation strategy can outperform agents with either purely packaged or purely pipelined strategy in some scenarios. Suppose that there are three peer agents that respond to the initiating agent's negotiation requests: A_{j_1}, A_{j_2}, and A_{j_3}. Suppose that A_{j_1} and A_{j_2} are inconsistent peers and A_{j_3} is consistent. Using the purely packaged strategy (Figure 1(a)), an initiating agent would have to run the risk of having all three issues rejected as a package if

A_{j_1} is too busy right now. Using the purely pipelined approach (Figure 1(b)), an initiating agent would have to initiate an additional negotiation with A_{j_3} when in fact the initiating agent could have simply approached A_{j_3} with a package of two issues. Figure 1(c) subsequently illustrates that with a hybrid approach, the initiating agent can switch issues around to those peers that have completed their negotiation jobs with itself.

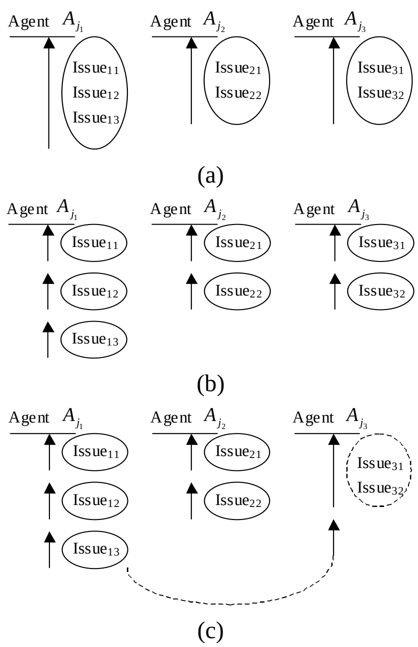

Figure 1. Strategic negotiations: (a) purely packaged, (b) purely pipelined, and (c) adaptive, confidence-based. The dashed oval in (c) means the negotiation job has finished.

To compute an agent's confidence in each of its peers, we use two categories of parameters: (1) those that are related to the negotiation process, e.g., the tardiness degree of the peer which indicates the time it takes for the peer to make decision to respond to a message from the initiating agent, and (2) those that are related to the negotiation outcome, e.g., the agreement degree of the peer which indicates the degree of its agreeing with the initiating agent's help request on an issue. The confidence value of the initiating agent A_i in a peer agent A_j, which is in the set of its all peer agents $\{A_{j_1}, A_{j_2}, ...\}$, is based on the variance of the observed parameters from A_i's viewpoint:

$$Confidence_{A_i}^{A_j} = \sum_{k=1}^{n} \frac{w_k}{1 + \sqrt{\sum_{\tau=1}^{l} (F_{A_i,A_j}^k(t_\tau) - \overline{F_{A_i,A_j}^k})^2 \Big/ (l-1)}}$$

where $F_{A_i,A_j}^k(t_\tau)$ is the perceived value of the kth parameter $F_{A_j}^k$ at time t_τ ($\tau \in [1, l]$), l is the window size ($l \geq 1$), $\overline{F_{A_i,A_j}^k}$ is the mean of all $F_{A_i,A_j}^k(t_\tau)$, w_k is the weight of F_{A_i,A_j}^k, and n is the number of parameters of A_j. Note that when l is equal to 1, the confidence value of A_i in A_j's consistency is 0. This makes sense since A_i has not yet interacted with A_j. It also means in the beginning an initiating agent prefers the pipelined approach.

Experiments and Results

We have implemented the adaptive, confidence-based negotiation strategy in a multiagent system. In our system, each agent has multiple overlapping capabilities and is capable of performing multiple tasks. When an agent encounters a task, it first analyzes whether it is able to solve the problem all by itself; if not, it initiates a coalition formation process. Each agent has 2+*N* threads: (1) a *core* thread to manage tasks, reason, and learn, (2) an *execution* thread for task simulation, and (3) *N negotiation* threads for concurrent negotiations with other agents.

In this paper, we report some experimental results aimed to quantitatively evaluate the adaptive, confidence-based negotiation strategy in complex environments by comparing its performance with the purely packaged and the purely pipelined strategy. Specifically, we compare these three different strategies from the perspectives of (1) the time and communication cost in negotiation process, and (2) the utility of negotiation outcome.

Experimental Settings

To evaluate the adaptive, confidence-based negotiation strategy, we conducted a series of experiments. In the experiments, we simulated a variety of environments, designed a set of experimental scenarios, and designed a series of tasks and agents.

Environment Simulation. In the experiments, we simulate three characteristics of the environment:

(1) *Dynamism Range*: This represents the range within which agents' dynamic behavior varies. We simulated a variety of dynamism ranges of the environment through setting different variation ranges of agents' *tardiness degrees*. Note that the tardiness degree of an agent indicates the time (measured in the number of time ticks) it takes for the agent to respond to a message from the initiating agent. Thus, this range describes a key characteristic of the negotiation *process*.

(2) *Uncertainty Range*: This represents the range within which agents' uncertain negotiation outcome varies. We simulated a variety of uncertainty ranges of the environment through setting different variation ranges of agents' *agreement degrees*. Note that the agreement degree of an agent indicates the degree of its agreeing with the initiating agent's help request on an issue. It describes negotiation *outcome*. If the responding agent agrees fully to the initiating agent's offer, then the utility of the outcome is 1.0. If the negotiation fails without a (partial) deal, then the utility is 0.0. If a partial deal is reached, then the utility of the outcome is proportional to the worth of the partial deal.

(3) *Real-Time Degree*: This represents the critical degree of the time constraint in negotiations. We simulated a

variety of real-time environments by setting a time limit on each negotiation process. Any negotiation that runs longer than its time limit is aborted, yielding 0.0 for the outcome utility. Table 1 shows the categories of real-time degrees. Real-time degree 1 denotes the non-real-time setting where all negotiations are able to complete within the time limit of 50,000 time ticks.

Real-Time Degree	Time Limit (ticks)
1	50,000
2	27,500
3	25,000
4	22,500
5	20,000
6	17,500
7	15,000

Table 1. Real-time environmental setting: different degrees of time limits set on negotiations.

Experimental Scenarios. We set up 6 experimental scenarios, as shown in Table 2, to compare the performance of the confidence-based strategy with the purely packaged and purely pipelined strategies. The first three scenarios, ES1, ES2, and ES3, serve as the baseline environments (non-real-time). We use these results to observe the negotiation process and outcome utility where the agents do not have to worry about time constraints. We use the last three scenarios, ES4, ES5, and ES6, to observe how the different strategies perform when time constraints play a role. In terms of the degrees of dynamism and uncertainty, ES1 and ES4 are considered to have low degrees, ES2 and ES5 average, and ES3 and ES6 high.

For example, in ES1, an agent may take up to $2RTT_{\min}$ to respond to a message, where $RTT_{\min}$ is the minimum round-trip time for a message to be transmitted between two agents. Also, the utility outcome of a negotiation ranges between 0.6 and 1.0. In ES3, an agent may take up to $3RTT_{\min}$ to respond to a message. Also, the utility outcome of a negotiation ranges between 0.2 and 1.0. Thus, the environment in ES3 is more dynamic and uncertain than that in ES1.

Scenario	Dynamism Range	Uncertainty Range	Real-Time
ES1	$[RTT_{\min}, 2RTT_{\min}]$	[0.6, 1.0]	No (RTD =1)
ES2	$[RTT_{\min}, 2.5RTT_{\min}]$	[0.4, 1.0]	No (RTD =1)
ES3	$[RTT_{\min}, 3RTT_{\min}]$	[0.2, 1.0]	No (RTD =1)
ES4	$[RTT_{\min}, 2RTT_{\min}]$	[0.6, 1.0]	Yes (RTD >1)
ES5	$[RTT_{\min}, 2.5RTT_{\min}]$	[0.4, 1.0]	Yes (RTD >1)
ES6	$[RTT_{\min}, 3RTT_{\min}]$	[0.2, 1.0]	Yes (RTD >1)

Table 2. Experimental scenarios. RTD = Real-Time Degree.

For our experiments, we run ES1, ES2, and ES3 once each. However, we run ES4, ES5, and ES6 six times each, with each run corresponding to a particular real-time degree greater than 1.

Tasks and Agents. We designed a series of coalition formation tasks (99 tasks) for a specific initiating agent. In each task there are 6 issues to be negotiated with other agents. We assume that the average number of evidence messages the initiating agent has to provide for each issue is the same, which is set as 2.

Since we aim to compare the performance of the three strategies in same contexts, we use the exactly same environments when each strategy is employed, including the same series of tasks, the same initiating agents, the same number of consistent agents, the same number of inconsistent agents, and the same confidence degree of an agent in a task.

Further, because we are concerned with the strategic negotiations, each agent's tardiness degree has the same average ($0.5RTT_{\min}$), and so does each agent's agreement degree (0.8). Depending on the particular experimental scenarios, however, the ranges of these values differ such that some agents behave more consistently than others.

Note also that due to the real-time constraints, it is possible that an agent never gets around to tackle all issues of a task. Thus, there will be issues that are not negotiated at the end of the time limit. For these issues, each will count as 0.0 towards the overall utility of the task. In this manner, we not only keep track of failed negotiations, but also take into account negotiations waiting in the pipelines that an agent fails to initiate.

Comparison on Cost in Negotiation Processes

Based on the efficient multi-issue negotiation assumption, the packaged negotiation approach can reduce the message and time cost needed in the negotiation processes. To examine whether the confidence-based strategy outperforms others in negotiation cost, we compare each strategy's average number of messages sent for negotiations per task, and average negotiation time per task in non-real-time environments (ES1, ES2, and ES3). We specify the negotiation time of a task as the time taken between when the agent receives the task as a problem and when the agent receives the outcome of the last negotiation. If the time limit comes first before all issues are negotiated, then all ongoing negotiations will be aborted, and remaining issues that are yet-to-be-negotiated will also be removed. In this case, the negotiation time of the task is the time limit allocated to the task. Figure 2 shows the results.

In non-real-time environments, the number of messages sent for negotiations in each task depends on (1) the number of issues in the task, (2) the number of evidence messages required by each responding agent, (3) the number of peer agents employed with the packaged or the pipelined approach, and (4) the issues' distribution pattern among responding agents. It is not influenced by the environment's dynamism and uncertainty degree. From Figure 2(a), we observe that the confidence-based strategy needed more messages than the purely packaged, but fewer messages than the purely pipelined. This result is as expected and it is due to that: (1) in our experimental setting, the average number of evidence messages the initiating agent has to provide for each issue is the same, and (2) a multi-issue negotiation has a lower message

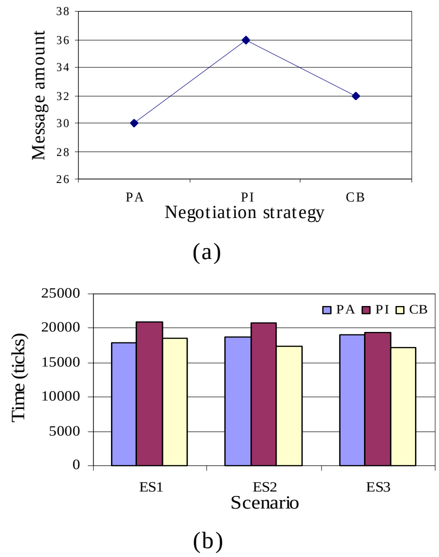

Figure 2. Comparison on negotiation cost: (a) message cost, and (b) time cost. PA = Purely Packaged, PI = Purely Pipelined, and CB = Confidence-Based.

overhead than multiple single-issue negotiations. In the confidence-based strategy, there are not only multi-issue negotiations but also single-issue negotiations. Thus, the confidence-based strategy used more messages than the purely packaged strategy, but fewer than the purely pipelined. We conclude that the confidence-based strategy can reduce the message cost in negotiation process compared with the purely pipelined strategy, but not so with the purely packaged.

From Figure 2(b), we have the following observations.

In each dynamic and uncertain environmental scenario, the average negotiation time per task, represented as $\overline{T}_{nego}$, in the purely pipelined strategy is always the highest. This is consistent with the efficient multi-issue negotiations assumption.

When the environment is more dynamic and uncertain (moving from ES1 to ES3), (a) $\overline{T}_{nego}$ for the purely packaged strategy increases, (b) $\overline{T}_{nego}$ for the purely pipelined strategy decreases, and (c) $\overline{T}_{nego}$ for the confidence-based strategy decreases. When the environment is more dynamic and uncertain (in ES2 and ES3), it is more likely for a peer to possibly respond more slowly. That results in more chances for some issues waiting in the pipelines to be moved from slow-progressing pipelines to the fast-progressing ones, if the latter have empty slots in the schedule of the pipelines. This thus results in the lower values of $\overline{T}_{nego}$ for the purely pipelined and the confidence-based strategies. On the other hand, in the purely packaged strategy, no issues can be moved. As a result, its $\overline{T}_{nego}$ increases.

The confidence-based strategy outperformed the purely packaged strategy in terms of $\overline{T}_{nego}$ in ES2 and ES3, but not so in ES1. This is because in ES1, the agents were faced with a less dynamic and less uncertain environment such that the agents actually should be performing mostly packaged negotiations. In a way, we have found a cutoff point below which the confidence-based strategy will not do as well as the purely packaged strategy.

The confidence-based strategy outperformed the purely pipelined strategy in terms of $\overline{T}_{nego}$ in all non-real-time scenarios. This is because the agents in the confidence-based strategy were able to save time using packaged negotiations, indicating that the environment was still consistent enough for the agents to make useful profiles of their interactions with other agents. Ultimately, we hypothesize that when the environment is extremely dynamic and uncertain, the confidence-based strategy will conform to the purely pipelined one.

We do not report the corresponding comparative study of the three strategies in the real-time experimental scenarios of ES4-ES6 because of the aborted negotiations and issues that were not even negotiated due to the imposed time limits. However, we did observe the same patterns and trends in the negotiations that were not aborted.

Comparison on Negotiation Outcome

To compare the strategies in terms of the quality of the negotiation outcome, we look at the utility as previously defined. Figure 3 shows the results of the three real-time experimental scenarios, ES4, ES5, and ES6. Each outcome utility is the average utility of all 99 tasks negotiated, where each task had 6 issues to be negotiated. Note that issues that did not have a chance to be negotiated would also be counted as a "failed" negotiation (utility = 0.0). Note also that Figure 3 does not include the results for non-real-time environments (real-time degree =1).

From Figure 3, we have the following observations.

As the real-time constraints became more stringent, the negotiation outcomes deteriorated. This is as expected. We also observe that as the environments became more dynamic and uncertain, the deteriorations started to take place noticeably at a smaller real-time degree. This is because the agents behaved in larger ranges of dynamism and uncertainty caused some negotiations to progress too slowly.

The confidence-based strategy was able to withstand the real-time constraints up to the 4th degree (time limit = 22,500 ticks) while the other strategies started to feel the pressure of real-time factors in the 2nd and 3rd degrees. This shows that (a) our confidence-based strategy is more robust than the other two strategies, and (b) the strategy is able to make good decisions on when to use packaged and pipelined approaches within a certain degree of real-time constraints.

The confidence-based strategy outperformed the other two strategies in almost all real-time scenarios in our experiments. This validates our confidence-based strategy, which adaptively combines the more opportunistic

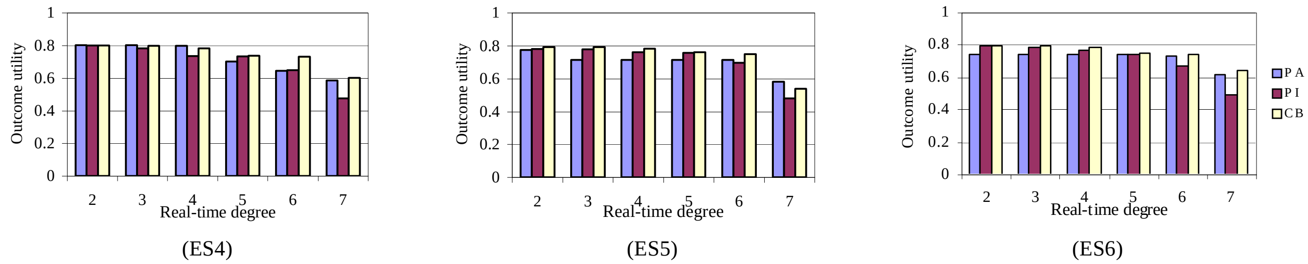

Figure 3. Comparison on negotiation outcome for ES4, ES5, and ES6, from less dynamic and uncertain to more dynamic and uncertain. PA = Purely Packaged, PI = Purely Pipelined, and CB = Confidence-Based.

pipelined approach and the more efficient packaged approach to conduct the multiple, concurrent negotiations. However, there was an exception to this claim: ES5, real-time degree = 7. In this scenario, the purely packaged strategy outperformed the confidence-based one. We are still investigating the reasons that might have caused this exception.

The purely pipelined strategy did not always outperform the purely packaged strategy in negotiation outcome. The purely packaged strategy actually outperformed the purely pipelined one when the real-time degree value was high (real-time degree = 6 or 7). When a negotiation task is highly time-constrained, the pipelined strategy may not be able to negotiate all issues. In that case, being opportunistic does not necessarily pay off: if the negotiations are not completed soon enough, the issues waiting in the pipelines will not get a chance to be moved from a slower pipeline to a faster one. Thus, the advantage of the pipelined strategy lessens. On the other hand, by packaging issues into one single negotiation, the packaged strategy can actually by chance complete negotiations within the imposed time limit for some tasks. Thus, we have found a cutoff point in terms of real-time degree above which the purely packaged strategy starts to outperform the purely pipelined strategy for the three different scenarios. Note that these numbers also roughly coincide with the observed real-time degree at which the confidence-based strategy started to deteriorate.

Combining the presented results, we see a key contribution of our experiments. In general, when the environment is more dynamic and uncertain, the performance of the purely packaged strategy deteriorates (e.g., in terms of the time cost). However, when the environment is highly constrained, the purely packaged strategy has a better chance than the purely pipelined one in completing tasks since it is more likely for the former to complete negotiating all issues in time while the latter even does not have time to be opportunistic.

Conclusions and Future Work

We have designed and implemented an adaptive, confidence-based negotiation strategy in the problem domain of multiagent coalition formation. This strategy addresses multi-issue strategic negotiations in complex environments. In our strategy, an agent chooses a combination of pipelined and packaged approaches to conduct multiple, concurrent negotiations with its peers based on the agent's confidence in its peer's consistency. Our experiments have found that the confidence-based strategy generally outperformed the purely packaged and purely pipelined strategies in a variety of uncertain, dynamic, and real-time environments. In the immediate future, we will conduct further experiments with larger dynamism and uncertainty ranges to locate the cutoff point beyond which the purely packaged or pipelined strategy outperforms the confidence-based strategy.

References

Fatima, S. S., Wooldridge, M., and Jennings, N. R. 2004. Optimal Negotiation of Multiple Issues in Incomplete Information Settings. In *Proc. AAMAS'04*, pages 1080-1087, New York City, NY, 2004.

Kraus, S., Sycara, K., and Evenchik, A. 1998. Reaching Agreements through Argumentation: A Logical Model and Implementation. *Artificial Intelligence*, 104(1-2):1-69.

Li, X. and Soh, L.-K. 2004. Learning-Based Multi-Phase Coalition Formation. In *Proc. WS on Coalition and Teams: Formation and Activity, at AAMAS'04*, pages 9-16, New York City, NY, 2004.

Nguyen, T. D. and Jennings, N. R. 2004. Coordinating Multiple Concurrent Negotiations. In *Proc. AAMAS'04*, New York City, NY, 2004.

Raiffa, H. 1982. *The Art and Science of Negotiation*. Cambridge, USA: Harvard University Press.

Rosenschein, J. S. and Zlotkin, G. 1994. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. Cambridge MA: MIT Press.

Soh, L.-K. and Li, X. 2003. An Integrated Multi-Level Learning Approach to Multiagent Coalition Formation. In *Proc. IJCAI'03*, pages 619-624, Acapulco, Mexico, 2003.

Soh, L.-K. and Li, X. 2004. Adaptive, Confidence-Based Multiagent Negotiation Strategy. In *Proc. AAMAS'04*, pages 1048-1055, New York City, NY, 2004.

Walton, D. N. and Krabbe, E. C. W. 1995. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. Albany, NY: SUNY Press.