Adaptive, Confidence-Based Multiagent Negotiation Strategy

Leen-Kiat Soh and Xin Li
Computer Science and Engineering
University of Nebraska-Lincoln
115 Ferguson Hall
Lincoln, NE 68588-0115 USA
{lksoh, xinli}@cse.unl.edu

Abstract

We propose an adaptive 1-to-many negotiation strategy for multiagent coalition formation in dynamic, uncertain, real-time, and noisy environments. Our strategy focuses on multi-issue negotiations where each issue is a request from the initiating agent to the responding agent. The initiating agent conducts multiple concurrent negotiations with responding agents and in each negotiation it employs (1) a pipelined, one-at-a-time approach, or (2) a confidence-based, packaged approach. In the former, lacking knowledge on the responding agent, it negotiates one issue at a time. In the latter, with confident knowledge of the past behavior of the responding agent, it packages multiple issues into the negotiation. We incorporate this adaptive strategy into a multi-phase coalition formation model (MPCF) in which agents learn to form coalitions and perform global tasks. The MPCF model consists of three phases: coalition planning, coalition instantiation and coalition evaluation. In this paper, we focus on the instantiation phase where the negotiations take place.

1. Introduction

In this paper we focus on multi-issue negotiations on task and resource allocation for multiagent coalition formation in dynamic, uncertain, real-time, and noisy environments.

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 [5, 14]. From the perspective of 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. A negotiation may address multiple issues or only one issue.

Coalition formation is an important method of cooperation among autonomous agents in multiagent environments [9]. In general, each autonomous agent only has incomplete views of the world and is incapable of performing specific global tasks all by itself. So some agents may form coalitions to allocate tasks among them to achieve the global goals. We have designed and implemented a multi-phase coalition formation model (MPCF) in which case-based reinforcement learning (CBRL) is applied to equip each agent with learning ability to form coalitions effectively and efficiently in dynamic, uncertain, real-time, and noisy environments.

We emphasize the *coalition instantiation* phase within the MPCF model. The model consists of three phases. The first phase generates a plan to form the intended coalition. The second phase carries out and instantiates the plan through candidate selection and multiple concurrent negotiations. Finally, the last phase evaluates the coalition formation process and quality of the coalition to learn to improve on future formation activity.

To address the characteristics of the environments, we propose an *adaptive*, *confidence*-based negotiation strategy for our coalition instantiation phase. In general, an agent is confident if it believes that it will succeed in carrying out its intended actions [3] or in our framework, if the agent believes that its knowledge of its peers is accurate. Based on the latter definition of confidence, our strategy specifies that a coalition-initiating agent employs (1) a *pipelined*, one-at-a-time approach, and (2) a confidence-based, *packaged* approach. In the former, the initiating agent, lacking knowledge on the responding agents, negotiates one issue at a time. As the negotiation processes complete, the agent subsequently negotiates other issues. In the latter, the initiating agent, with confident knowledge of the past behaviors of the responding agent, packages multiple issues into each negotiation. The initi-

AAMAS'04, July 19-23, 2004, New York, New York, USA.

ating agent is also capable of using both approaches in a hybrid, dealing with a mixed group of responding agents.

Note that our adaptive, confidence-based negotiation strategy is a meta-negotiation strategy. It manages how the negotiations should be conducted in terms of scheduling and assigning the different issues to different peers. This strategy does not deal with how each negotiation should proceed at each interaction step between an initiating agent and a responding agent. Our strategy is to integrate a pipelined approach with a packaged approach to benefit from the advantages of both and address the disadvantages of both. To determine which approach to use in our problem domain in a particular situation, we define *confidence* and use that to guide our decision making. This *confidence* is hinged upon how the agent profiles its peers' dynamic behavior and the environmental changes. Thus, an initiating agent is able to adaptively select different negotiation approaches to improve the request satisfaction and the cost effectiveness.

In the following, we first discuss some related work in negotiations and coalition formation. Then, we briefly present our coalition formation model. In Section 4, we propose and describe the adaptive, confidence-based negotiation strategy. Subsequently, we present some preliminary results before concluding.

2. Related Work

Although considerable research has been conducted either in coalition formation among self-interested agents (e.g., [7, 13]), or in coalition formation among cooperative agents (e.g., [8]), little work has been done in coalition formation among both self-interested and cooperative agents. Furthermore, there have been no attempts to study coalition formation among such agents in a dynamic, real-time, uncertain, and noisy environment, which is a typical real-world environment and in which a sub-optimal coalition needs to be formed in a real-time manner. Our research addresses negotiation strategy in coalition formation problems among both self-interested and cooperative agents operating in such environments.

To automate negotiation processes, a number of negotiation mechanisms have been proposed and studied. Rahwan and his colleagues [4] briefly classified them into: *game-theoretic* and *auction-based* mechanisms, *heuristic-based bargaining* mechanisms, and *argumentation-based* approaches. In game-theoretic analysis, researchers usually attempt to determine the optimal strategy by analyzing the interaction as a game between identical participants, and seeking its equilibrium (e.g., [6]). In cases where it is not possible to reach the optimal outcome due to resource limitations, dynamic environment or incomplete information, some heuristics have been devised. Heuristics are rules of thumb that produce "good enough" outcomes, and are mainly based on empirical testing and evaluation (e.g., [1]). Argumentation-based negotiations allow agents to exchange, in addition to proposals and indications of their acceptance or rejection, meta-information about them, such as the reasons for their proposal, and for accepting or rejecting the proposals (e.g., [3]).

At the lowest level, the step-by-step negotiation mechanism that we use in our framework is an integration of heuristic-based bargaining mechanism and argumentation-based mechanism [11]. In a dynamic, uncertain, real-time, and noisy environment in which each agent only has incomplete information about the environment and other agents, an initiating agent tries to achieve an agreement on the request to produce a good-enough and soon-enough outcome, and then evaluates the outcome. In our framework, we extend the above to a higher level, looking inter-coalition competition, where an agent can be a member of different coalitions simultaneously and can be subjected to requests for the same resources or capabilities that it has.

We focus here on the management of negotiations instead of how each negotiation is to be conducted [11]. Our coalition-initiating agent assigns and schedules negotiations to different peers based on the current environment and the confidence in the peers' behaviors or capabilities, derived from dynamic profiling of the peers.

3. Multi-Phase Coalition Formation Model

Our overall framework is based on a model called the Multi-Phase Coalition Formation (MPCF) model. The model consists of three phases: coalition planning, coalition instantiation and coalition evaluation, as depicted in Figure 1. In *coalition planning*, the agent applies case-based reasoning (CBR) to obtain a coalition formation plan. In *coalition instantiation*, the agent carries out the planned formation strategy through 1-to-many negotiations with coalition candidates. In *coalition evaluation*, the agent evaluates the coalition formation process, the formed coalition structure (if a coalition is successfully formed), and the coalition execution outcome (if the coalition is executed eventually) to determine the utility of the planned strategy and reinforces the strategy. In the following, we briefly outline the MPCF model.

In coalition planning, the coalition-initiating agent applies CBR to derive a specific coalition formation plan for the current problem based on a previous plan stored in the casebase. This can avoid building a coalition formation plan from scratch. Given a problem (a task) to solve, the agent retrieves from its casebase the best case of the highest similarity with the current problem and the highest utility. Based on the difference between the best case and the new problem, the agent adapts the case solution to compose a coalition formation plan which specifies the number of coalition candidates, the number of expected

AAMAS'04, July 19-23, 2004, New York, New York, USA.

coalition members, the time allocated for coalition instantiation, the allocation algorithm, and the number of messages recommended.

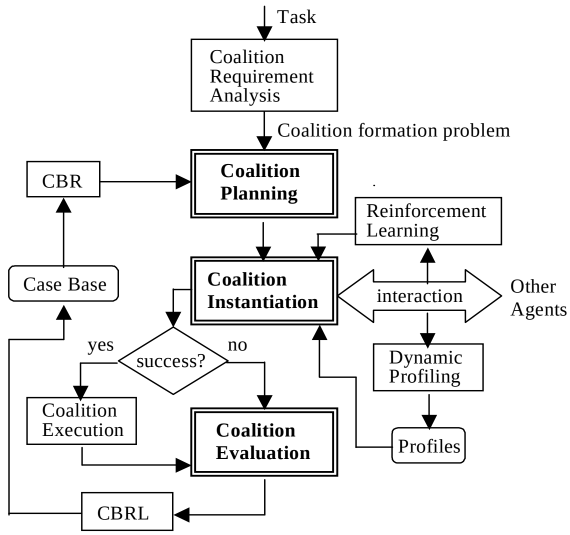

Figure 1. MPCF model

The coalition instantiation phase implements the coalition formation plan to form a coalition. At first, the coalition-initiating agent normalizes the task—dividing the task into separate execution units as different negotiation issues, computing the potential utilities of its peers, and ranking the peers based on the potential utilities. Then the agent concurrently negotiates with each selected peer agent on the set of subtasks in an attempt to form the intended coalition. Each negotiation is argumentative where the initiating agent attempts to persuade the responding agent to perform a task or provide a resource by providing support or evidence for its request [11].

The coalition evaluation phase provides the basis for an agent to improve its coalition formation plans. This phase evaluates both the coalition instantiation process (in terms of time spent, number of messages used, number of peers approached, etc.) and the execution outcomes of the subtasks agreed upon in the coalition (in terms of the number of subtasks performed by highly-capable peers, etc.). In general, a good plan is one that uses little computational and communication resources with successful instantiations and subsequent executions.

We employ an integrated case-based reinforcement learning strategy [10] to utilize the results of the evaluation phase to influence the coalition planning phase.

Our case-based reinforcement learning (CBRL) design is aimed at identifying the situation where a coalition formation plan was successful and reinforcing it. First, each agent has a casebase of coalition formation cases. A coalition formation case consists of a problem description, a solution, an outcome, and its utility. The *problem* description consists of an agent's external and internal environments and the task description. The *solution* part gives a coalition formation plan, outlining the number of peers to approach, the time needed to carry out the coalition, and the types of peers needed for a successful coalition. The *outcome* part indicates the coalition instantiation results among agents, subtasks' execution results, and the evaluation values of the actual coalition formation process. The *utility* indicates the quality of the case, specifically, the quality of the plan in addressing the coalition problem represented in the case. Coupling the evaluation and the problem description, the agent can learn a new coalition formation case to increase its coverage of cases or can update the original best case's utility using the evaluation result to *reinforce* the case.

As part of the coalition instantiation module, an agent maintains a dynamic profile of its peers in terms of the negotiation results and coalition history. A coalition-initiating agent uses the profile together with its knowledge of its peers' capabilities and of the current environment and problems at hand to compute the potential utility of each peer (coalition candidate) for a particular coalition. Using this potential utility, an agent reinforces its decision in selecting a candidate, in the following manner:

$$PU_{A_i,A_j}(s,a,t+1) \leftarrow (1-\beta) * PU_{A_i,A_j}(s,a,t) + \beta * C_{A_j}(A_i,t+1),$$

where A_j is a particular peer agent; the state s corresponds to the current coalition formation problem; the action a corresponds to coalition candidate selection; $PU_{A_i,A_j}(s,a,t)$ is the old potential utility of A_j and $PU_{A_i,A_j}(s,a,t+1)$ is the updated one; β is the learning rate ($0 \leq \beta \leq 1$); and $C_{A_j}(A_i,t+1)$ is the weighted sum of A_j's characteristic parameters as measured by A_i. With the above formula, an initiating agent prefers to approach peer agents that have been helpful and *coalition-worthy*.

4. Confidence-Based Negotiation Strategy

In this section, we describe our adaptive, confidence-based negotiation strategy. As discussed in Section 3, this strategy drives the coalition instantiation phase. Note that the coalition planning phase provides a plan, outlining the characteristics of peers to approach, and specifics about how the coalition should be formed. The coalition instantiation phase has to decide how to carry out the plan to form the intended coalition. We separate the plan from the instantiation to simplify the representation of the coalition formation problems and the case-based reasoning process.

Here we describe the motivations behind our confidence-based strategy. 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 during the candidate selection process do not perform as expected. For example, the initiating agent may rank peer A_j as the best candidate and expect to reach an agreement for its negotiation in a short time. However, A_j is busy and unable to entertain the negotiation request. Without a flexible management strategy, the initiating agent would have to wait until A_j is available, probably missing the time requirement. It is also possible that, because of the dynamic nature that we assume of our 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 we realize that there is a need for a management strategy that is flexible, capable of adapting to the profiled behavior of the peers as well as the real-time observation of the negotiation activities.

Figure 2 depicts the two approaches of our confidence-based strategy: (1) a *pipelined*, one-at-a-time approach, and (2) a confidence-based, *packaged* approach. In the former, the initiating agent, lacking knowledge on how the peers perform in coalition formation, negotiates one issue at a time, via multiple, concurrent negotiation processes. As the negotiation processes complete, the agent subsequently negotiates for other pending issues. This allows the agent to be *cautious* and *opportunistic* at the same time. In the latter, the initiating agent, with confident knowledge of the past behavior of the peers, packages multiple issues into each negotiation. The underlying principle of the above two approaches is that if the agent knows its peers well, then it is willing to package multiple issues into one negotiation for more efficient coalition formation. If it does not know its peers well, then it is willing to take a more cautious step—negotiating with one issue at a time, trying to avoid getting delayed with a particular peer with a package of issues. It is also more opportunistic, constantly monitoring the progress of multiple concurrent negotiations and switching issues to whichever peers that are capable and becoming available.

Our strategy is based on two assumptions:

- **The Efficient Multi-Issue Negotiation Assumption**. We assume that packaging multiple issues into a negotiation is more efficient than negotiating each issue one by one. Thus, an agent will prefer to perform multi-issue negotiations if it has confidence in its profiling of its peers.
- **The Overlapping Capabilities Assumption**. We assume that in the multiagent system, the agents have a substantial number of overlapping capabilities. That is, an agent is aware of numerous peers that can satisfy an issue. This assumption facilitates the pipelined approach. If there are only a few overlapping capabilities, the pipelined approach will not have the flexibility to be opportunistic. For example, if there is only one peer that knows how to perform task *T*, then the agent has no choice but to reserve *T* for the peer.

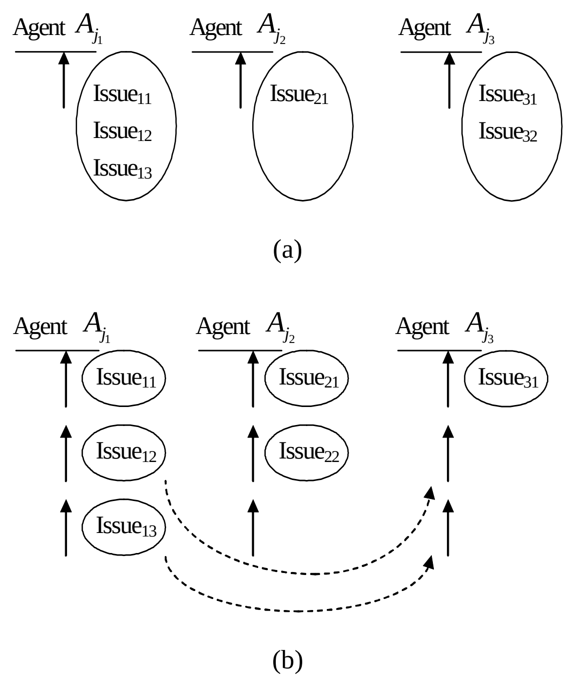

Figure 2. Confidence-based negotiation strategy: (a) packaged, (b) pipelined.

In the following, we further describe several aspects of our strategy: (1) confidence, (2) negotiation approach selection, (3) learning to manage negotiations, (4) inter-coalition competition, (5) intra-coalition competition, (6) contingency handling, and (7) case-based negotiation.

4.1. Confidence

Our negotiation strategy is based on the confidence that the coalition-initiating agent has in its peers. Specifically, we measure the confidence based on how consistent each peer's negotiation or coalition behavior is. A peer's negotiation or coalition behavior is based on the neighborhood profile that the agent maintains. The parameters profiled include the helpfulness of the peer indicating the satisfaction degree of requests to the peer, the agent's reliance on the peer in terms of the ratio of sending requests to the peer among all peers, the reliance of the peer on the agent in terms of the ratio of receiving requests from the peer among all peers, and other negotiation-derived parameters such as: (1) a *tardiness* degree indicating the communication delay between the agent and the peer, (2) a *hesitation* degree indicating how readily the peer is to agree to a request, (3) an *availability* degree of capability indicating whether the peer possesses the desired capability to solve task, and so on.

An agent computes its confidence value in a peer along the peer's many parameters as described above based on the standard deviations of the parameters. A parameter value with a small standard deviation means that the peer exhibits consistency in this particular parameter (such as tardiness degree). As a result, the agent has a high confidence—it can expect what the communication delay would be if it approaches the peer for a negotiation. The confidence value of the initiating agent A_i in a peer agent A_j 's k_{th} characteristic $C^k_{A_j}$ can be computed in the following formula:

$$Confidence_{A_i}^{C_{A_j}^k} = \frac{1}{1+\sqrt{\sum_{i=1}^{l}(C^k_{A_j,t_i} - C^k_{A_j,average})^2}}$$

where $C^k_{A_j,t_i}$ is the perceived value of $C^k_{A_j}$ at time t_i ($i\in [1,l]$) and $C^k_{A_j,average}$ is the average value of $C^k_{A_j,t_i}$ during the certain time period.

The composite confidence value of the initiating agent in a peer agent is simply a weighted sum of the confidence values of the peer's characteristics.

With the definition of confidence, an agent can decide whether to take a pipelined approach, or a packaged approach, or a hybrid one.

4.2. Negotiation Approach Selection

Based on the confidence that an agent has in its peers, it decides which approach to undertake. Figure 3 shows a hybrid approach.

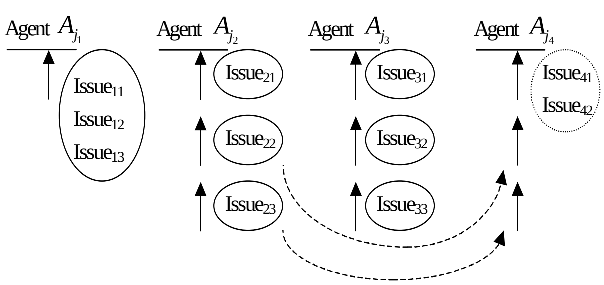

Figure 3. Negotiation approach selection. The dashed oval indicates a completed negotiation

In Figure 3, the initiating agent negotiates with four agents simultaneously for a coalition formation problem. It negotiates with peers A_{j_1} and A_{j_4} using the packaged approach and with peers A_{j_2} and A_{j_3} using the pipelined approach. These two approaches have different characteristics. The *pipelined* approach is more conservative in terms of not overloading a particular responding agent with multiple issues at one time. It is also more cautious, flexible and opportunistic since the initiating agent may dynamically redistribute an issue from the original pipeline to a "free and capable" pipeline. For example, when the pipeline of peer A_{j_2} is impeded due to communication delay, loss of communication channel, or agent faults, the pending issues waiting at the pipeline can be transferred to another "free and capable" pipeline (e.g., agent A_{j_4}). Under the Efficient Multi-Issue Negotiation Assumption, the *packaged* approach is more efficient and reduces communication cost. But it is also more risky and rigid. If Agent A_{j_1} proves to be abnormally slow in reaching a deal, and if all other peers have completed their respective negotiations, the initiating agent would not be able to transfer some of the issues in the package from A_{j_1} to the other peers.

Taking into consideration the advantages and disadvantages of both negotiation approaches, we propose the following basic approach selection principles: (1) package the issues when dealing with a peer that the agent has high confidence in, (2) pipeline the issues when dealing with a group of peers that the agent has low confidence in, (3) package the issues if the coalition is time-critical, and (4) pipeline the issues if the coalition is highly important (due to robustness of the pipelined approach).

Note that how to package issues and what issues to use to setup each pipeline to start the negotiations are outlined by the coalition formation plan generated via CBR.

4.3. Learning to Manage Negotiations

With the profiling mechanism and reinforcement learning in the MPCF model, a coalition-initiating agent dynamically profiles each responding agent to update the potential utility of each peer (as a coalition candidate) and reinforce their cooperation relationship. There are two additional learning scenarios associated with our adaptive, confidence-based negotiation strategy: (1) the coalition-initiating agent learns whether to package or pipeline the issues, and (2) the coalition-initiating agent learns whether and how to transfer an issue from one pipeline to another pipeline.

The first learning scenario measures the success rate of each approach in specific coalition formation tasks. If an approach leads to a successful formation and a good-quality coalition, then the agent will learn to factor in that piece of knowledge in its future selection of approach. If the pipelined approach was carried out without transferring any issues across the pipelines, then that means the pipelined approach had been unnecessary. The agent should learn to lean towards the packaged approach next time around when it encounters the same coalition task.

The second learning scenario is specific to the pipelined negotiation approach. In the pipelined approach, issues can be transferred to other peer agents. The transfer is opportunistic and depends on the current status of each negotiation. To decide whether and how to transfer the issues effectively and efficiently, the initiating agent needs to predict how the peers will act in the next step. A peer that is not busy now may become very busy in the next few moments. A communication link that is fast now may become very congested and slow next. Based on the confidence that an agent has in its peers and past "transfer" utility, it can learn to be more conservative and patient. This is the key to improve the quality of the coalition since the original assignment of subtasks (issues) to each peer is deemed to be the best plan that the agent has derived during its coalition planning phase. Thus, transferring the issues around would reduce the quality of the coalition, in hope of trading it off for a more efficient coalition formation process.

4.4. Inter-Coalition Competition

Conceptually, an agent may try to instantiate two coalitions overlapping each other temporally. For example, a task *T* comes in, and the agent starts a coalition formation process to handle *T*. While the agent still negotiates to deal with task *T*, another task *R* arrives and the agent also starts promptly a coalition formation to handle *R*. In this case, there are two sets of negotiations, possibly directed to some same peers. It is possible to have a pipeline going to one peer and yet another going to the same peer, with different sets of issues lined up. This is where inter-coalition competitions can occur. (Note that as a responding agent, a peer only needs to consider competing requests and does not have to worry about the coalitions.)

The problem with this is the overburdening of a particular peer, affecting both negotiations. In our computation of a peer's potential utility, we do take into account the current relationship that an agent has with each peer. And one of the parameters is whether the agent is now negotiating with each peer. This will serve as a guideline to alert the subsequent coalition formation processes to avoid approaching the same peer again, relieving some of the pressure from inter-coalition competitions.

4.5. Intra-Coalition Competition

In our problem domain, in addition to inter-coalition competition among overlapped coalition formation, there may be also conflicts among the requested capabilities by a same initiating agent (intra-coalition competition). In this case, the responding peer cannot decide which issue is offered in prior. In the packaged approach, the responder needs to check with the initiator to decide the offer order; while in the pipelined approach, the initiator may dynamically order the requests to multiple issues according to their priorities and the negotiation results on previous issues. So the initiator can make an endogenous agenda and conduct the negotiations following the agenda [2].

When the initiator makes the agenda, it can decide to negotiate with multiple responding agents on a same important issue. It also can decide to negotiate on urgent issues first. By focusing on these issues first, the agent realizes that it has a higher probability to have them agreed to by the peers than later. Compared with the packaged approach, the pipelined approach is more flexible since the initiator can change the agenda dynamically.

Furthermore, to deal with uncertain and dynamic environments, an agent may extraneously requests for more resources or services in anticipation of failed negotiations. As a result, algorithms that are greedy, worried and lazy have been proposed [12]. For our coalition instantiation process, these algorithms would introduce intra-coalition competition. For example, a worried algorithm [12] will prompt an agent to approach several peers for one unit of the same issue as an insurance policy. Once the agent secures an agreement from one of the peers, it has to immediately terminate the other negotiations. However, terminating negotiations incurs costs to the perception by the peers of the agent.

4.6. Contingency Handling

Ultimately, the coalition formation process is based on a plan and the plan may not work as expected during the coalition instantiation process. Focusing only on the packaged approach vs. the pipelined approach, we see that it is possible for the initiating agent to lose contact with a peer due to communication loss, agent faults, or extremely long delay. In such a case, the agent needs to have a contingency plan to persist with the coalition instantiation process.

In general, if the packaged negotiation approach is used, and a peer is found to be not responsive, then the agent breaks up the package and distributes the issues to the remaining peers in the coalition. For the pipelined approach, the response is more straightforward—simply transferring the issues to other pipelines. And, in the case where the agent is stuck with a particular issue negotiating with a non-responsive peer, the agent will duplicate another negotiation with another peer on the same issue and conduct the two negotiations in parallel. Whichever negotiation completes successfully first will prompt the agent to terminate the other.

AAMAS'04, July 19-23, 2004, New York, New York, USA.

4.7. Case-Based Negotiation

In our confidence-based negotiation strategy, the initiating agent conducts a case-based reflective argumentative approach, extended from [11]. In the original design, for each negotiation, the initiating agent first finds a specific local negotiation strategy through CBR. Then, it activates a thread to negotiate. The negotiation management module of the agent monitors various negotiation threads and changes the local negotiation strategies in real-time. The module will terminate all remaining negotiations once it finds out that it no longer can form a viable coalition. The module will terminate its redundant requests once it has secured agreements from successful negotiations.

To extend from [11], we see two approaches to represent the negotiation cases in the casebase. First, each case contains the local negotiation strategy for a number of issues. Second, each case contains the local negotiation strategy for only one issue.

The first approach is directly applicable to the packaged approach. With multiple issues captured in one case, the local negotiation strategy can provide guidelines on how to negotiate with a package of issues. However, there is a potential problem. For *n* issues, there are O(2^n) possible combinations among these issues. So the casebase could become intractable as the number of issues grows.

The second approach caters to the pipelined approach as this approach only deals with one issue in one negotiation. For *n* issues, O(*n*) cases are enough to cover the local negotiation strategies on one specific issue. So the casebase can be much smaller and more modular. However, these cases would fail to capture the additional stresses that multi-issue negotiations bring to the responding peers.

A solution combining the two approaches with a hierarchical organization of cases may address the benefits and weaknesses of the two approaches.

5. Experimental Results

We have implemented the MPCF model in a multi-agent system where the adaptive, confidence-based negotiation strategy is to be incorporated. Presently, our system performs 1-to-many negotiations, similar to the packaged approach. 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 3+*N* threads: (1) a *core* thread to manage tasks, reason, and learn, (2) a *communication* thread, (3) an *execution* thread for task simulation, and (4) *N negotiation* threads for concurrent negotiations with other agents.

In this paper, we report some preliminary experimental results aimed to study the communication cost difference between high-confident negotiation and low-confident negotiation scenarios. We use two specific characteristic parameters, *tardiness* degree (TD) and *hesitation* degree (HD), to measure the initiating agent's confidence to each peer agent.

We conducted three experiments as shown in Table 1. In ES1, all peers have the same tardiness degrees with the agent. However, each peer has a different hesitation degree. In ES2, all peers have the same hesitation degree but different tardiness. In ES3, all peers have the same hesitation and tardiness degrees.

Experiment Set	Peer Characteristics
ES1	Same TD, different HD
ES2	Different TD, same HD
ES3	Same TD, same HD

Table 1. Experiment sets

To specify the various hesitation degrees, we use the number of evidence messages needed for each agent to be persuaded during the argumentative negotiation. In ES1, we further created two sub-experiments, as shown in Table 2. In EXP1, we set the hesitation degree of each peer different but fixed. For example, agent A_2 will have a hesitation degree of 7; and so on. With this, we created a high-confident negotiation environment with a set of peers of different characteristics. In EXP2, the hesitation degree of each peer is fluctuant and based on a uniform distribution within a specified range. With this fluctuant HD setup, an agent was expected to profile its peers less successfully than with the fixed HD setup, which is a low-confidence negotiation environment.

	HD Distribution	Confidence
EXP1	Fixed: $HD_{A_i} = 9 - i$	High
EXP2	Fluctuant: $HD_{A_i} \in [9-i-1,\ 9-i+1]$	Low

Table 2. Sub-experiments in ES1

Figure 4 shows the agent A_1's approach frequencies to its peers. We define *approach frequency* as the number of times an agent approaches a peer agent for coalition formation over the duration of the entire experiment. From Figure 3, we see that in EXP1 agents A_7, A_8 and A_9 were approached significantly more frequently. This is because these agents had smaller HD values and agent A_1 is highly confident in them. As a result, they were approached more often as they are likely to improve the quality of the coalition formation process.

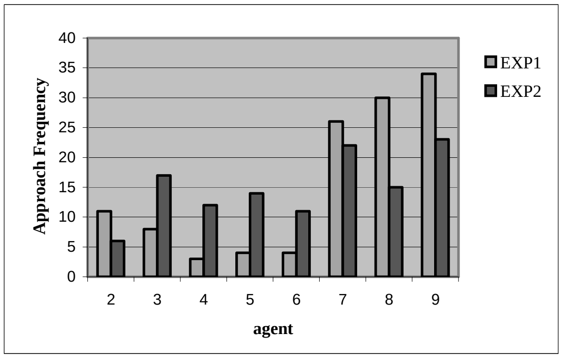

Figure 4. Approach frequencies to peer agents

6. Conclusions

In this paper, we have described an adaptive, confidence-based negotiation strategy in which a pipelined, one-at-a-time approach and a confidence-based, packaged approach can be employed adaptively to manage 1-to-many negotiations, within a Multi-Phase Coalition Formation model. Our objective to use such a strategy is to improve the request satisfaction and cost effectiveness of coalition formation in a dynamic, noisy, uncertain, and real-time environment.

Given the strategy, we expect an agent to behave in the following manner: (1) initially, an initiator prefers the pipelined negotiation approach since it has not yet obtained much information about other agents and thus is not confident enough; but as the time progresses, the packaged approach will be used more frequently than before; (2) for time-critical issues, the packaged approach is preferred while for high-importance issues, the pipelined approach is preferred; and (3) for the same set of negotiation issues, the adaptive, confidence-based negotiation strategy is supposed to have more issues being offered than the pure packaged approach, and be able to negotiate faster than the pure pipelined approach. We have discussed several key aspects such as confidence, selection, learning, inter- and intra-coalition competitions, contingency handling, and case-based negotiation.

Presently, we have implemented the coalition formation model, end-to-end, and partially some of the coalition instantiation steps. We have presented some preliminary results. Our future work is to completely incorporate the strategy into the multi-phase coalition formation model and conduct experiments to study the proposed strategy.

References

[1] S. S. Fatima, M. Wooldridge, and N. R. Jennings, "Multi-Issue Negotiation under Time Constraints", in *Proceedings of AAMAS'2002*, Bologna, Italy, 2002.

[2] R. Inderst, "Multi-Issue Bargaining with Endogenous Agenda", *Games and Economic Behavior*, 30: 64-82, 2000.

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

[4] I. Rahwan, P. McBurney, and L. Sonenberg, "Towards a Theory of Negotiation Strategy", in *Proceedings of the Fifth Workshop on Game Theoretic and Decision Theoretic Agents (GTDT-2003)*, Melbourne, Australia, 2003.

[5] H. Raiffa, *The Art and Science of Negotiation*, Cambridge, MA: Harvard University Press, 1982.

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

[7] S. Sen, and P. S. Dutta, "Searching for Optimal Coalition Structures", in *Proceedings of ICMAS'2000*, pp. 286-292, Boston, MA, 2000.

[8] O. Shehory, K. Sycara, and S. Jha, "Multi-Agent Coordination through Coalition Formation", in *Intelligent Agents IV: Agent Theories, Architectures and Languages*, *Lecture Notes in AI*, number 1365, pp. 143-154, Springer, 1997.

[9] O. Shehory, and S. Kraus, "Methods for Task Allocation via Agent Coalition Formation", *Artificial Intelligence*, 101: 165-200, 1998.

[10] L.-K. Soh, and X. Li, "An Integrated Multi-Level Learning Approach to Multiagent Coalition Formation", in *Proceedings of IJCAI'2003*, pp. 619-624, Acapulco, Mexico, 2003.

[11] L.-K. Soh, and C. Tsatsoulis, "Reflective Negotiating Agents for Real-Time Multisensor Target Tracking", in *Proceedings of IJCAI'2001*, pp. 1121-1127, Seattle, WA, 2001.

[12] L.-K. Soh, and C. Tsatsoulis, "Agent-Based Argumentative Negotiations with Case-Based Reasoning", in *Working Notes of the AAAI Fall Symposium Series on Negotiation Methods for Autonomous Cooperative Systems*, pp. 16-25, North Falmouth, MA, 2001.

[13] F. Tohme, and T. Sandholm, "Coalition Formation Processes with Belief Revision among Bounded Rational Self-Interested Agents", *Journal of Logic and Computation*, 9(6): 793-815, 1999.

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