

# Handling Uncertainty in Multi-Agent Commitments \*

Ping Xuan and Victor R. Lesser  
Department of Computer Science  
University of Massachusetts at Amherst  
Amherst, MA 01003  
{pxuan, lesser}@cs.umass.edu

UMass Computer Science Technical Report 1999-05

January 18, 1999

## Abstract

Commitments play a central role in multi-agent coordination. However, they are inherently uncertain and it is important to take these uncertainties into account during planning and scheduling. This paper addresses the problem of handling the uncertainty in commitments. We propose a new model of commitment that incorporates the uncertainty, the use of contingency analysis to reduce the uncertainty, and a negotiation framework for handling commitments with uncertainty.

**Keywords:** Multi-agent coordination, Multi-agent planning, Decision making under uncertainty

## 1 Introduction

Handling uncertainty is one of the major topics in artificial intelligence. In a multi-agent system, an agent needs to handle not only the uncertain aspects of its own activities, but also the uncertain perceptions about other agents' behaviors. In this paper we study the uncertainty associated with the commitments among the agents and techniques to reduce the uncertainty.

In a multi-agent system, each agent can only have a partial view of other agents' behavior. Therefore, in order to coordinate the agents' activities, the agents need to have a mechanism to bridge their activities based on the partial knowledge. *Commitments* has emerged, among many research groups [Cohen and Levesque, 1990, Castelfranchi 1993, Jennings 1993, Decker and Lesser, 1992], as the bridge for multi-agent coordination.

By definition, a commitment specifies a pledge to do a certain course of action [Jennings 1996]. A number of commitment semantics have been proposed, for example, the "Deadline" commitment  $C(T, Q, t_{dl})$  in [Decker and Lesser, 1992], means a commitment to do (achieve quality  $Q$  or above for) a task  $T$  at a time  $t$  so that it finishes before a specified

---

\*Effort sponsored by the Defense Advanced Research Projects Agency (DARPA) and Air Force Research Laboratory Air Force Materiel Command, USAF, under agreement number F30602-97-1-0249 and by the National Science Foundation under Grant number IIS-9812755 and number IRI-9523419. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Defense Advanced Research Projects Agency (DARPA), Air Force Research Laboratory, National Science Foundation, or the U.S. Government.

deadline,  $t_{dl}$ . When such a pledge is offered, the receiving agent can then do its own reasoning and planning based on this commitment, and thus achieves coordination between the agents.

However, there are a number of uncertainties associated with commitments. First, there is the question about whether or not the commitment can be fulfilled by the offering agent. Tasks may fail, for example, and thus cannot achieve the quality promised. Or, the results may be delayed and therefore cannot meet the deadline. Also, the task  $T$  itself may depend on some preceding actions, and there are uncertainties about those actions. Since the receiving agent depends on the predictable outcome of the commitment, this uncertainty must be considered. This type of uncertainty originates from the uncertainty of the underlying tasks. In this paper we propose the modeling of such uncertainty in terms of a distribution of the possible outcomes of a commitment, based on the statistical behavior of the tasks.

A second source of uncertainty comes from the agent decision/planning process. As we know, flexibility is needed in order for the agent to operate in a dynamic environment. Therefore, when an agent's beliefs and desires change, the agent should be able to change or revoke the commitments [Jennings 1996]. Hence, changes in the commitment can occur because of tasks not directly related to the fulfillment of the commitment. To the receiving agent, this can cause problems because its actions may depend on the honoring of the commitment in the offering agent. This aspect of uncertainty is inherent to the making of the commitment itself and not from possible underperformance of tasks. In this paper we take into account of this uncertainty by explicitly describing the possibility of *future* modification/revocation of the commitment. Contingency planning [Onder and Pollack 1997], a mechanism for handling uncertain failures, is used in this work in order to reduce uncertainty and plan for possible future events such as failure or de-commitment.

There is still another form of uncertainty caused by the partial knowledge of the offering agent regarding the agent who needs this commitment. Namely, how important or useful the commitment is to the receiving agent, and when the commitment would be not very useful to the receiving agent? To tackle this problem we define the *marginal gain or loss* [Sandholm 1993] value of commitment and use this value to decide how the agents perform their reasoning and planning.

For coordination to be successful when there are these forms of uncertainties, there must be structures that allow agents to interact predictably, and also flexibility for dynamic environment and imprecise viewpoints, in addition to the local reasoning capability [Durfee, Lesser, and Corkill 1989]. For this purpose, we propose a domain-independent, flexible negotiation framework for the agents to negotiate their commitments. Our work differs from the *conventions* and *social conventions* [Jennings 1993, Jennings 1996] in that our negotiation framework is domain-independent, and allows the agent to integrate the negotiation process in problem solving and dynamically reason about the local and social impact of changes of commitment, whereas conventions and social conventions define a set of rules for the agents to reconsider their commitments and ramifications to other agents when commitments change.

The rest of this paper is structured as follows. In Section 2 we discuss the modeling of commitments, focusing on the uncertainties we discussed above. Section 3 discusses the impact of uncertainty in commitments on planning and scheduling, in particular the use of contingency analysis. In Section 4 we discuss the negotiation framework for handling the commitments with uncertainty. Experimental results illustrating the strength of our approach, is provided in Section 5. We conclude with a brief summary in Section 6.

## 2 Uncertainty in Commitments

For the purpose of illustration, our discussion uses the TÆMS framework [Decker and Lesser, 1993] for modeling the agent task environment. This does not introduce loss of generality because TÆMS is domain-independent and capable of expressing complex task environments. In the TÆMS language, agents' problem solving processes are described in terms of *tasks* organized in a way to reflect the decomposition of a task into low-level actions as well as alternative ways to achieve the same high-level task. TÆMS explicitly quantifies the quality, cost, and duration (q,c,d) characteristics of *methods* (primitive tasks) in terms of discrete probability distributions. The *quality accumulation functions* (qaf), such as *min*, *max*, *sum*, etc., describe how lower-level task qualities translate to higher-level task qualities. Inter-task relationships such as *enables* ( $T_a$  enables  $T_b$  means  $T_a$  must have accomplished a positive quality before  $T_b$  can start,

essentially a precedence relationship), can be used to specify various complex interactions among tasks. In terms of reasoning and coordination using the TÆMS, our discussion will focus around the scheduling framework of Design-to-Criteria scheduling [Wagner, Garvey, and Lesser 1997] and the Generalized Partial Global Planning (GPGP) family of coordination mechanisms [Decker and Lesser, 1992].

Figure 1 shows an example of task structures in agents A and B. Note the relations “A2 enables B2” and “A4 enables B4”. They involve tasks in different agents, therefore are called *non-local effects* (NLE). The existence of NLEs drives the need of coordination.

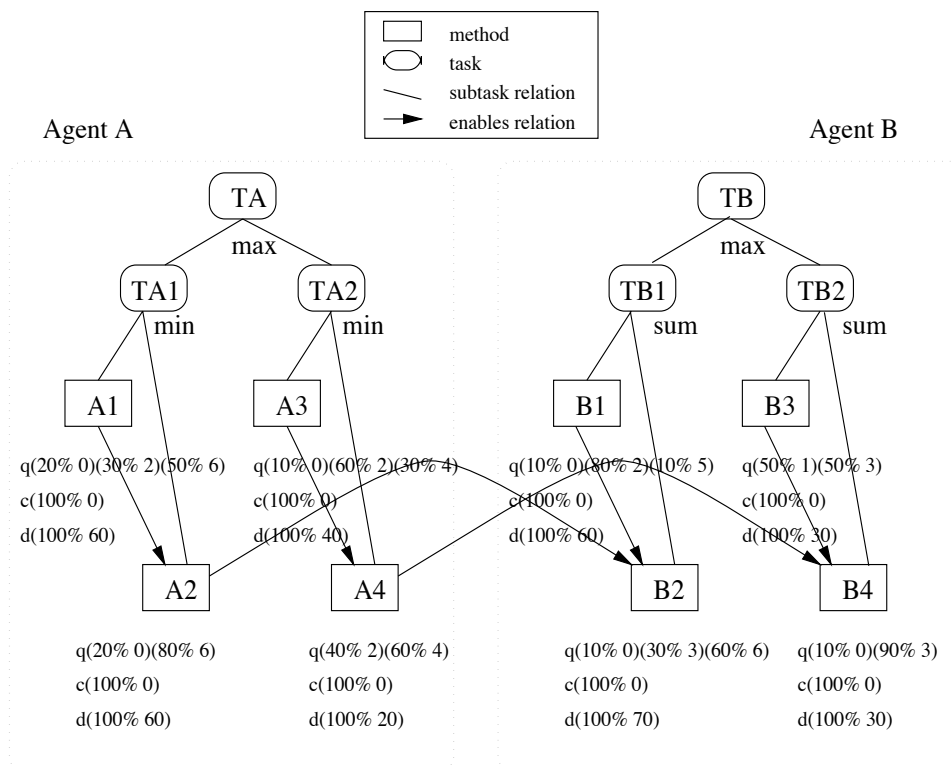


Figure 1: Example Task Structure

Assuming both agents try to maximize their quality outcome, and they both have a deadline of 160. Under Design-to-Criteria, initially A would select schedule (A1, A2) and B would select (B1, B2). Then, after the agents detect the NLE between A2 and B2, A would proactively pledge to complete A2 by time 120, with some estimated quality. We will show that, although this example uses a very simple task structure, it clearly demonstrates the importance of uncertainty.

The first step for incorporating uncertainty in commitments is to take into account the uncertainty of underlying tasks. In [Decker and Lesser, 1992], a commitment specifies only the *expected* quality of the committed task. However, expected values often do not provide sufficient information for effective coordination, especially when there are possible task failures. In our example, A2 has expected quality 4.8. But, there is a 20% chance that method A2 will fail ( $q=0$ ), and thus cannot enable B2. Furthermore, because A2 is enabled by A1, which also fails in 20% of the time, the result is that the commitment has only 64% chance of being useful to B. To address this problem, the commitment should specify a distribution of possible outcomes, i.e., “64% chance ( $t = 120$ )  $\wedge$  ( $q = 6$ ), 36% chance ( $t = 120$ )  $\wedge$  ( $q = 0$ )”. In general, uncertainty also exists similarly for other aspects of commitments, such as start time, finish time, duration, and cost, but for clarity we discuss only the quality aspect here. Thus, assuming no deadline constraints are violated because of variations in task durations, the calculation of this distribution (based on a given

plan/schedule) is straightforward. Let  $C(T)$  be a commitment about task  $T$ , and  $pred(T)$  for the predecessor tasks of  $T$  (tasks that enable  $T$ ), then the probability of the quality outcome equals  $x$  is,

$$p(q(C(T)) = x) = p(q(T) = x) \cdot p(q(pred(T)) > 0).$$

Next, because commitments are generally future-oriented, agents need to revise their speculations about the future and therefore also the decision making *over the time*. This introduces the uncertainty in decision making, in this case, the uncertainty about whether the agent respects or honors the commitment — in addition to the probabilistic outcome of commitments. For instance, we notice that an agent may de-commit its commitments during its problem solving process, when keeping the commitments is in conflict to its performance goal. As before, initially at time 0, agent  $A$  chooses the plan  $A1, A2$  and offers commitment about  $A2$  to  $B$ . However, at time 60, where  $A1$  completes, in the case that  $A1$  fails or has  $q = 2$ ,  $A$  would perform re-planning and select the alternative plan  $A3, A4$  because that plan has a better quality outcome now. Clearly,  $B$  should be able to know this information *at time 60* rather than to notice the commitment not in place at time 120. More interestingly, however, if we can specify *at time 0* that there is a possibility of de-commitment at a future time (60), then  $B$  can plan even further. On the other hand, if at time 60  $A1$  finishes with quality 6, then the quality outcome of the commitment has changed to a *better* distribution “80%  $q=6$  and 20%  $q=0$  at time 120”, because now  $q(A1)=6$ . It would also be helpful if this information can be updated to  $B$ . In other words, agent  $A$  can tell  $B$ , “right now I pledge to do  $A2$  before time 120. However, you may hear *more information* about the commitment at time 60.” The additional future information may be good (better distribution) or bad (de-commit). But the important thing is that the other agent,  $B$ , can make arrangements *ahead of time* to prepare for such information, hence better coordination.

One way to represent this uncertainty is to calculate  $p'(C(T), t)$ , the probability that  $C(T)$  will remain kept at time  $t$ . The exact calculation of  $p'$  depends on the knowledge of (1) when and what events will trigger re-scheduling, and (2) whether or not a future re-scheduling would lead to changes in commitments — in other words, prediction of future events, decisions, and actions. Obviously, for complex systems, those information could be computationally expensive (if not impossible) to get. To avoid this problem, we do not calculate  $p'$  directly, instead we focus on the first part — the events that may cause re-scheduling to change commitments, for example,  $A1$ 's possible failure (or low quality) at time 60. This occurs only 50% of the time, which mean in 50% of time re-schedule will not happen at time 60, therefore, 50% is a *lower bound* of  $p'(60)$ . It is implied that there is no change in the commitment before time 60, because of no re-scheduling, i.e.,  $p'(t < 60) = 1$ . To agent  $B$ , this implies that time 60 is a possible checkpoint for the commitment offered by  $A$ .

The checkpoints are easy to calculate by analyzing the schedule to see how failure or low performance of a method affects the performance goal of the agent. In the language of contingency analysis, the tasks in the *critical region* are critical to the agent performance (and/or commitment), and thus their potential low performance outcome events would become the checkpoint events. The event information may include: the time the event may occur, the task to be watched, the condition for re-scheduling (i.e., quality equals 0), and a lower bound for  $p'$ .

The third source of uncertainty comes from the partial knowledge the other agent, namely, how important is this commitment to others? To answer this question we first need to know how important this commitment is to *me*. By knowing this we can avoid bad coordination situations such as offering (and pay the high cost of honoring) a commitment that is of little value to the receiving agent, or in the contrary, canceling a commitment that is very important to the agent needs it for only little gain in local performance. To solving this problem we use the notion of *marginal cost* and define the *marginal loss* as the difference of agent performance without making the commitment and the one making the commitment. A zero marginal loss means the commitment is “free”, i.e., the offering agent would strive to do the same with or without making the commitment, such as the case of  $A$  offers commitment on  $A2$ . Like quality values, marginal loss values are also dependent on future outcomes, and can change over time. For example, the same commitment on  $A2$  would incur a marginal loss if  $A1$  finishes with quality 2, because in that case the alternative plan ( $A3, A4$ ) would have higher expected local quality. Similarly, we define *marginal gain* as the difference of agent performance when receiving the commitment and the one *without* receiving it. A marginal gain of zero indicates that the receiving agent is indifferent to the commitment.

Marginal gain/loss can be expressed in terms of the utility values (or distributions of utility values), in this case,

task qualities. However, we need to note that since agents may use different utility scales, a quality of 5 could have different meanings in different agents. We can use the relative *importance* to indicate how quality values in the other agent translate to the quality values in this agent. For simplicity we do not address the importance issue here any further, and assume the importance value of 1.0 is always used, i.e., the quality scales are the same in all agents. In order to evaluate the marginal gain/loss against a particular commitment, we simply compare the best-quality alternatives with and without the commitment. For example, if  $B$  can have a commitment of  $A2$  before time 90, then  $B$ 's average quality would be 6.15 (assuming the commitment does not fail). Without this commitment, however,  $B$ 's best alternative is  $(B1)$ , which only produces average quality 2.10. Hence, this commitment has a marginal gain of 4.05.

### 3 The Impact on Planning and Scheduling

Now that a commitment has uncertainty associated with it, agents can no longer regard a commitment as guaranteed, and assume the absence of failures. Therefore, planning and scheduling in an agent becomes harder. However, the benefit of using uncertainty comes from better understanding of the commitment in the agents and therefore more effective coordination. To achieve this, we also need to change the local scheduling/planning activities. Traditionally, when the uncertainty of commitment is overlooked and thus the commitment is *assumed* to be failure proof, re-scheduling is often performed *reactively* to handle the appearance of an unexpected failure that blocks the further execution. This type of reaction is *forced upon* rather than being planned ahead. In a time sensitive environment, it is often too late. Therefore, it is desirable that the agent has the capability of planning *in anticipation* of possible failures and know the options if failures do occur. This way, necessary arrangements can be made before the failure may occur, and also we save the effort of re-scheduling by adopting a planned-ahead action in case of failure.

To handle possible failure outcomes in commitment, we use *contingency analysis* in conjunction with the Design-to-Criteria scheduling. Due to space limitation, we cannot describe the details of contingency analysis here; details are available in [Wagner, Raja, and Lesser 1999]. In our approach, a failure in the commitment can be treated the same way as a failure in a local task. First, we analyze the possible task failures (or low quality outcomes) or commitment failures and identify alternatives that may improve the overall quality outcomes when failure occurs.

Through contingency analysis, the resulting *schedule* is no longer a linear sequence of actions, as it is with ordinary scheduling; rather it has a *branching* structure that specifies alternatives and the conditions for taking the alternatives. In Figure 2, (a) shows the linear schedules of agent  $A$  and  $B$ , and (b) shows the schedules with contingency. Clearly, the linear schedule only specifies the preferred path in the contingency schedule, where as a contingency schedule specifies a set of paths based on possible future outcome. Using contingency analysis, the value of a schedule is now computed based on this branching structure, and therefore is more accurate. To utilize this branching structure we need to monitor the progress of the execution and dynamically discover and analyze possible future branches, and therefore it is closely related to the monitoring of an anytime search process in the solution space (set of possible execution paths), such as the work of [Hansen and Zilberstein 1996].

Contingency analysis can also be used to handle uncertainty originated from changing/revoking the commitments. As mentioned before, we can identify the *critical regions* in the schedule that may have significant impact on the overall quality if a failure occur in the critical regions, thus leads to the discovery of checkpoints. On the other hand, once we have the checkpoint information regarding a commitment, we can make contingency schedules to specify a *recovery option*. Let  $T^\alpha$  indicate that task  $T$  has outcome  $\alpha$ , for example,  $T^F$  for failure of  $T$ ,  $T^2$  for  $q=2$ . Then we can specify a recovery option for  $(B1, B2)$  such as  $(B1^2, A1^F, B3)$  to indicate that when  $B1$  finishes with  $q=2$  and  $A1$  fails, the agent should run  $B3$ . This is a generalization of the previous case, since conceptually we can regard the failure of commitment as a type of de-commitment which comes at the same time as the finish time of the commitment.

The use of marginal gain/loss becomes very important in scheduling and coordination. Although in our modeling of commitments, changes or de-commitments are allowed (unlike the traditional case, where commitments are *assumed* to be fixed, that is, in the absence of failures), these changes are *social* rather than local. The introduction of marginal

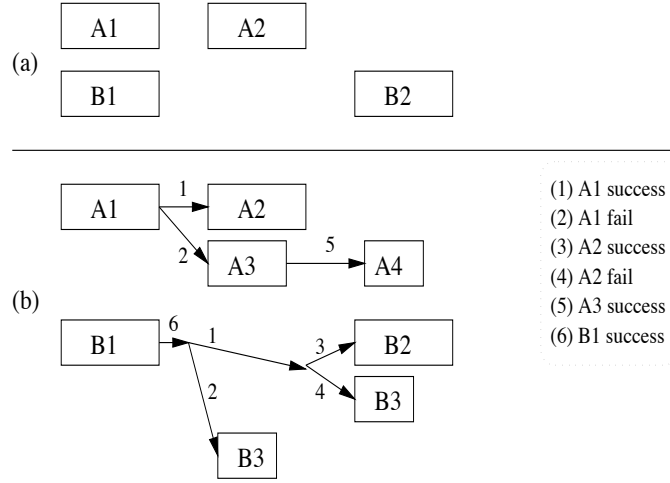


Figure 2: Schedules with contingency

gain/loss ensures that commitments are properly respected in a social context. If the overall utility of a multi-agent systems is the sum of the utilities in each agent, assuming the importance of activities in different agents is normalized, then only when the marginal gain is greater than the marginal loss, a commitment is socially worthwhile. Likewise, the commitment should be revoked only where the marginal loss is greater than the gain. The difference between marginal gain(s) and loss(es) becomes the utility of the commitment itself (which is different from the utility of the task being pledged). Therefore, the social utility of a schedule is the local utility of the schedule plus/minus the marginal gain/loss of the commitment received/offered. Note that marginal gain/loss also changes during the course of problem solving, therefore it needs to be re-evaluated when some tasks are finished.

## 4 The Negotiation Framework

In order to add flexibility to coordination, we also introduce a commitment negotiation framework that allows agents to interact with each other in order to achieve better coordination. This negotiation framework provide the following primitives for agent negotiation (here RA stands for the agent requesting/receiving the commitment, and OA for the agent offering the commitment):

- *request*: RA ask an agent to make a commitment regarding a task. Additional information includes the desired parameters of the commitment (task, quality, finish time, etc.) as well as the marginal gain information.
- *propose*: OA offers a commitment to one agent. Additional information includes the commitment content (with uncertainty associated) and possible marginal loss.
- *accept*: RA accept the term specified in OA's commitment.
- *decline*: RA chooses not to use OA's offer. This can happen when RA does not find the offer attractive but does not generate a counter proposal.
- *counter*: RA requests for a change in the parameters specified in the offered commitment, i.e., makes a counter-proposal. Changes may include better quality or quality certainty (i.e., a better distribution), different finish time, earlier/later possible checkpoints/re-schedule time
- *change*: OA makes changes to the commitment. The change may reflect the OA's reaction/compromise to RA's counter-proposal. Of course, the RA may again use the *counter* primitive to react to this modified commitment as necessary, until both sides reach consensus.
- *no-change*: If the OA cannot make a change to the commitment according to the counter-proposal, it may use this primitive to signal that it cannot make a compromise.

- *decommit*: OA cancels its offer. This may be a result of agent re-planning.
- *update*: both RA and OA can provide updated or more accurate information regarding a commitment, such as changes in marginal gain/loss, changes in the uncertainty profile of the commitment during the course of problem solving, etc.
- *fulfilled*: the task committed was accomplished by OA.
- *failure*: the commitment was failed (due to unfavorable task outcomes).

These primitives are used not only during the establishment of commitment, but also during the problem solving process. Therefore, they allow agents to negotiate and communicate their commitments dynamically during the problem solving period. The negotiation process help agents to be better informed about each other’s desires, intentions, and outcomes, therefore reduces the uncertainty in commitments and results in better coordination. For example, at time 0, if agent *A* offers *B* a commitment to complete *A2* before time 120, agent *B* can see that this commitment is useless and counter-propose agent *A* to commit on task *A4* before time 130. If such a commitment is offered with 100% certainty, the marginal gain is 2.6. However, agent *A* can only offer 90% certainty on *A4*, and such a commitment would cause a marginal loss of 0.72, which is acceptable to both agents. Clearly, the negotiation process helps the discovery of alternative commitments that leads to better social solutions. This is done by using marginal gain/loss information in negotiation. Without those information, agents’ coordination decisions would be based on local information only.

Under this framework, each agent can implement a *policy* using the primitives, which decides its communication protocol based on the negotiation strategy the agent will use to carry out the negotiation. The policy decides issues such as what parameters to choose when requesting/offering a commitment, how much effort (time and iterations) the agent is willing to spend on the negotiation, and how often the agent updates its commitments, etc. For example, an agent can choose to neglect counter-proposals if it cannot afford the planning cost or does not have the capability to reason about counter-proposals. The policies are often domain-dependent, and the reasoning of the policies is beyond the scope of this paper. A formal account of the reasoning models for negotiation to form a joint decision is provided in [Faratin, Sierra, and Jennings 1998]. In a general sense, negotiation can be viewed as a distributed search problem, and the policies reflect how the agents relax their constraints and search for compromises, such as the work of [Moehlman, Lesser, and Buteau 1992]. In this work, we use a simple policy that counter-propose only when the offered commitment brings no overall gain (i.e., marginal gain is less than marginal loss). If a counter-proposal cannot be found, the agent simply declines the commitment.

## 5 Experiments

In order to validate our approach, we implemented a generic agent that can work with a textual TÆMS input. We simulate two instances of such agent, *A* and *B*, to work on the task structures presented in Figure 1. We perform some comparisons to show how the handling of uncertainty improves coordination, and therefore improve overall performance. We assume that both agents have deadline 160, and both agents try to maximize quality outcomes.

First, we study the base case, where commitments do not carry uncertainty information, and no negotiation is used: in this case, one agent pro-actively offer a commitment to the other agent, using only expected quality and finish time. In Figure 3 we shows the distribution of the final quality outcomes for 200 runs. Three histograms for the quality of *A*, quality of *B*, and the sum of them are shown in this figure.

From the trace, we observed that *A*’s commitment about *A2* to finish by time 120 does not leave agent *B* with enough time to finish task *B2* by its deadline 160. However due to no negotiation, *B* cannot confirm that *A2* cannot arrive earlier, and they cannot discover an alternative commitment for task *A4*, since both agents found their best local alternative: (*A1*, *A2*) for *A* and (*B1*) for *B*.

In the second case, we add uncertainty information to the commitments. The commitment is still pro-active (with no negotiation), but the agents can use contingency planning to reduce the uncertainty in commitments. Figure 4




Figure 3: Base Case




Figure 4: Second Case: With Uncertainty

show the results for 200 runs. Here we can see some slight improvement of quality outcomes in both agents, but the similar pattern of histograms indicates that this has only minor impact on the scheduling. Due to no negotiation, the improvements are restricted to agent's local activities. For example, we notice that when  $A_1$  finishes with quality 2,  $A$  will choose to switch to plan  $(A_3, A_4)$  instead of continue to run  $A_2$  (therefore effectively de-commit its commitment) because now  $(A_3, A_4)$  has higher expected utility.

As the last case, we incorporate negotiation and using the marginal gain/loss information in commitment coordination. The results, shown in Figure 5, have very different patterns in the histograms. This indicates that the major changes in the agent's activities. We can see that now  $A$  has a relatively lower quality outcome than it does in the previous cases, but  $B$  has significant performance improvements. The overall result is that the sum of their qualities improved significantly. This is because the agents are able to find a better commitment between them (namely the commitment on  $A_4$ ) now. This commitment is social in that it helps to achieve overall better utility, although not all agents can have local gains. The following table shows the average quality outcomes in each case.

Average Q	A	B	sum(A,B)
Case 1	3.3	2.32	5.62
Case 2	3.49	2.67	6.16
Case 3	2.53	4.6	7.13

This also shows that the integration of all the mechanisms: negotiation, contingency planning, and marginal gains/losses is very important in effectively handling of uncertainties. These mechanisms handles different aspects of uncertainty, and they work together to achieve better coordination.






Figure 5: Third Case: Negotiation

## 6 Conclusion

In conclusion, we identified three sources of uncertainty inherent in commitments and discussed the ways to incorporate them into the modeling of commitments, and the mechanisms to handle the uncertainties, such as contingency analysis and negotiation. The goal of this work is to improve coordination effectiveness, and ultimately, to improve the overall utility of the multi-agent problem solving. Our results indicate that these mechanisms significantly improves coordination.

Obviously, our approach is computationally more expensive than previous approaches where uncertainties are not explicit, especially when the distributions propagate in the analysis, and when the number of contingent plan increases. One way to manage the complexity is to recognize that the analysis of possible future contingency plans can be an anytime process, and therefore we may trade off accuracy with the effort of analysis. Heuristics for effectively pruning the search space can also be applied.

The ability to handle uncertainty in commitments is especially important in a time-sensitive environment where agents cannot afford to re-schedule when failures occurs. Hence, by taking into account of the possibility of failure, this work also improves the reliability in problem solving. As the next step, we will address the issues related to the more general problem of fault-tolerance in multi-agent systems. Interesting issues may include: how to handle the uncertainty related to new (and possibly important) tasks arriving to an agent (which in turn may affect scheduling), the cost of dynamic monitoring, and the adaptive management of redundancies for fault-tolerance, etc.

## References

- [Castelfranchi 1993] C. Castelfranchi, "Commitments: from individual intentions to groups and organizations". AI and theories of groups & organizations: Conceptual and Empirical Research. Michael Prietula, editor. AAAI Workshop, 1993. Working Notes.
- [Cohen and Levesque, 1990] Philip R. Cohen and Hector J. Levesque. "Intention is choice with commitment." *Artificial Intelligence*, 42(3):213-261, 1990.
- [Decker and Lesser, 1992] Keith S. Decker and Victor R. Lesser, "Generalizing the partial global planning algorithm." *International Journal of Intelligent and Cooperative Information Systems*. 1(2):319-346. June 1992.
- [Decker and Lesser, 1993] Keith S. Decker and Victor R. Lesser, "Quantitative modeling of complex computational task environments." *Proceedings of the Eleventh National Conference on Artificial Intelligence*, 214-217, Washington, July 1993.
- [Durfee, Lesser, and Corkill 1989] E. H. Durfee, V. R. Lesser, and D. D. Corkill, "Trends in cooperative distributed problem solving". *IEEE Trans. Knowledge. Data Eng.* 1(1):63-83, 1989.
- [Faratin, Sierra, and Jennings 1998] P. Faratin, C. Sierra, and N. R. Jennings, "Negotiation Decision Functions for Autonomous Agents". *Int. Journal of Robotics and Autonomous Systems* 24 (3-4) 159-182, 1998.

- [Hansen and Zilberstein 1996] E. A. Hansen and S. Zilberstein. "Monitoring the Progress of Anytime Problem-Solving", Proceedings of the 13th National Conference on Artificial Intelligence, pp. 1229-1234, Portland, Oregon, 1996.
- [Jennings 1993] N. R. Jennings, "Commitments and Conventions: The Foundation of Coordination in Multi-agent Systems". The Knowledge Engineering Review, 8(3):223–250, 1993.
- [Jennings 1996] N. R. Jennings, "Coordination Techniques for Distributed Artificial Intelligence". Foundations of Distributed Artificial Intelligence, G.M.P. O'Hare and N.R. Jennings eds, John Wiley 1996.
- [Moehlman, Lesser, and Buteau 1992] Theresa A. Moehlman, Victor R. Lesser, and Brandon L. Buteau, "Decentralized Negotiation: An Approach to the Distributed Planning Problem", Group Decision and Negotiation, 2:161-191, Kluwer Academic Publishers 1992.
- [Onder and Pollack 1997] Onder, N. and Pollack, M., "Contingency Selection in Plan Generation". Proceedings of the Fourth European Conference on Planning (ECP'97), 1997.
- [Sandholm 1993] Sandholm, T., "An Implementation of the Contract Net Protocol Based on Marginal Cost Calculations". Eleventh National Conference on Artificial Intelligence (AAAI-93), Washington D.C., pp. 256-262, 1993
- [Wagner, Garvey, and Lesser 1997] Wagner, Thomas A., Garvey, Alan J. and Lesser, V.R. Criteria Directed Task Scheduling. In Journal for Approximate Reasoning—Special Scheduling Issue, Vol 19, pp. 91-118, 1998.
- [Wagner, Raja, and Lesser 1999] Thomas Wagner, Anita Raja, and Victor Lesser, "Modeling Uncertainty and its Implications to Design-to-Criteria Scheduling", Under review, also available as UMASS Department of Computer Science Technical Report TR-98-51. 1999.