Submitted to Cooperative Information Gathering

Using Agent Commitments as Planning Contexts

Ping Xuan
Department of Math and Computer Science
Clark University
950 Main Street
Worcester, MA 01610
pxuan@clarku.edu

Victor Lesser
Department of Computer Science
Computer Science Building
University of Massachusetts at Amherst
Amherst, MA 01003
lesser@cs.umass.edu

Abstract

A key challenge to planning in multiagent systems is to deal with uncertainty related to coordination, which means to manage the interdependencies between agent activities. Typically, a decentralized agent has only partial knowledge of the system and must deal with nondeterministic outcomes of both local and nonlocal actions. While there are many multiagent planning approaches, the issue of uncertainty in coordination has not been adequately addressed. Thus, in this paper we propose a framework that incorporates uncertainty into agent planning and coordination by using agent commitments as the multiagent context for agent local planning. A new, richer model of commitment is proposed to represent uncertainty and resolve interdependency. We then discuss the techniques to enhance planning and coordination under this framework, and demonstrate how to apply our framework to existing approaches and how it may improve problem solving.

1 Introduction

When decentralized agents cooperate in problem solving, the coordination of the interrelated tasks in different agents is a central issue. Multiagent coordination is the process of managing interdependencies between agent activities [28]. While agents may attain coordinated behaviors through a range of methods, such as assigning predefined organizational roles and obeying social rules and conventions, the most direct and effective method of coordination is through explicit communication. Here, by communication we mean the exchange of meta-level *control* messages rather than domain-level information. Control messages convey information concerning the structure and representation of the abstract problem solving process itself, such as the agent's intentions, the completion status of a task, or the layout of the schedule. Control messages resolve the uncertainty an agent has regarding coordination, and influence the agents' planning and decision-making actions. In contrast, domain-level communication refers to the transmission of information by the tasks themselves. For example, the transmission of input/output data during the execution of task modules.

The separation of control-level communication and domain-level communication is necessary in order to develop domain-independent coordination techniques. For the agents to make proper coordination decisions, the timing for communication and the content of the control messages are extremely important. Obviously, asking an agent to provide all meta-level information at all times is not only infeasible in all but most trivial cases, but also undesirable because a lot of the information may not be meaningful to any agent except the sending agent itself. For example, an agent may be distracted by too many unnecessary communication messages and its responsiveness could be hampered [17]. Typically, an agent does not need to know exact details about how another agent plans and executes all its actions, as long as the interrelationship between the two agents' tasks are properly managed. Thus, agents only need to provide control information relevant to the interrelationships[10, 11, 9]. Such information can serve as the planning context for an agent to take into consideration when planning its local actions toward solving the common goal. The more complete the multiagent planning context is, the less uncertainty an agent would have regarding the possible future course of actions, and the better coordination and planning can be achieved.

In this paper we examine the issue of when to communicate and what level of detail is needed in the control messages in order to establish a complete multiagent planning context. Our approach is to use *commitments* as the basis for multiagent planning context, describe the detail of the planning context in terms of uncertainty information in the commitments, and then extend and refine the model of commitments to incorporate the uncertainty. Commitments are the interface for agents to understand other agents' activities without having to know all details about their internal activities. In this regard commitments can also be viewed as distilled information: details that other agents cannot understand are filtered out, yet their consequences are incorporated in commitments so that other agents can account for their impact indirectly. Communication of commitments is intended for agents to establish mutual understanding of each other's control process without having to reveal everything the agents know about solving the problem. The dynamics of commitments encapsulates the dynamics of the agent's activities, therefore allows other agents to correctly characterize their behavior and the relation to each other. Ideally, if the model of commitments is detailed enough, the agents do not need to exchange more information other than commitments in order to establish a multiagent planning text. We will then discuss what techniques can be applied to take into account the richer model of planning context, and show how different levels of details can lead to different levels of coordination behavior. Obviously, a more detailed context would give the agents better understanding of the situation and hence better plans, but it would also require more communication (in terms of both the frequency of communication and the amount of the information).

Our assumption is that the agents, acting as autonomous problem solvers, have sophisticated local control and can intelligently formulate local plans. The agent is responsible for selecting alternative goals or plans and managing constraints such as task interdependencies and resource constraints. As such, when situations change, the agent can decide to switch to different goals and plans in order to achieve better overall performance for the whole system. Thus, the right context information is crucial to the quality of decision making.

Conceptually, we may model agent tasks in a hierarchical task network representation such as the TAEMS framework [11]. Such a representation may also model interrelationships of the tasks and the agent's utility functions. The agent's decision can then be viewed as a decision theoretic planning problem, or equivalently, an optimization problem - how to select the best actions so that the global utility can be maximized. However, since the agent typically only has a partial view of the system and has to deal with many sources of uncertainty, solving the problem optimally is infeasible in most cases. Here, we will assume that a two-leveled approach is used [9]: the problem is split into two subproblems: the agent local planning and agent coordination. While the agent local planning is primarily focused at task selection and sequencing, coordination is focused at ensuring that the local plans of the agents are compatible with each other. This way, a boundary between an agent's internal and external problem solving is defined. An agent would perform local optimization based on its own belief about the world and uses the coordination process as an interface with the external world. In other words, the coordination process needs to provide the right multiagent planning context to the planning system.

The planning context has to be able to represent uncertainty in coordination. For example, agent X's plan may rely on the successful completion of a task A in agent Y. The coordination process would ensure that Y's plan schedules A at the right time. However, since tasks often may have non-deterministic outcomes, there are chances that A may fail. Existing approaches often rely on conventions and/or fixed coordination mechanisms to handle this type of uncertainty when the failure occurs, but they lack flexibility and do not explicitly reason about these uncertainties and their consequences to overall performance of the plan. Also, agents may have to perform expensive re-scheduling operations to react to those unexpected events, which may affect both efficiency and performance. A good multiagent planning context would construct a partial view of the perceived system and establish a simplified, filtered model to represent the system external to the agent. Such encapsulation allows an agent to concentrate only on the information that may impact its decision making and removes unnecessary details, i.e., to know only it has to know.

This work is aimed at providing a complete and accurate encapsulation and handling the uncertainty related to coordination. We emphasize that our goal is not to propose yet another planning method or coordination mechanisms, but rather a framework that formalize the impact of the multiagent system toward the agent's local planning, i.e., to factor in the multiagent context into an agent's own belief. Then, a range of (single-agent) planning approaches can be used (with slight modifications) to create local plans. In this work we will be using the Design-to-Criteria (DTC) [43] scheduler and the Generalized Partial Global Planning (GPGP) [10] family of coordination mechanisms to illustrate how to apply our framework on top of them and how our framework enhances agents' ability to handle uncertainty.

As mentioned earlier, our approach is based on the use of agent commitments as the vehicle for specifying agent coordination activities. The focus of this research is to incorporate uncertainty into agent commitments [45]. Previous research on coordination [3, 5, 10, 24] has widely established that commitments are the central building blocks for coordination mechanisms. It manages interrelationships and serves as the interface between one agent's plan and the other agents' plans. We argue that the existing model of commitments does not accurately reflect the semantics behind the commitments and does not take into account the uncertainty. Thus, we extend the model of commitments, formally and quantitatively specify its semantics, and provide a more complete context. We present this richer model under the TAEMS task modeling framework [11], which uses a domain-independent representation for complex multiagent cooperation problems. We will show that the new model encapsulates uncertainty in nonlocal interrelationships, and we develop techniques to utilize the information and improve agents' decision making ability.

The types of information to be communicated as multiagent planning context may include organizational information, policy information, and state information. Organizational information reflects the roles of the agents, how tasks are distributed, and the structure and relationship of the tasks, and the objective of the agents. Policy information reflects the agent's problem solving strategy, for example its goals, intentions, plans, etc. State information is information about the status of the problem solving, for example current time and the quality of completed tasks. While some information such as organizational roles are typically long term properties that may be determined offline or communicated as static information at the beginning of problem solving, other information, in particular the state information, is inherently dynamic and therefore requiring a dynamic communication process. In this paper we will focus on the dynamic policy and state information and formalize this communication process, and connect it to the dynamics of the commitments, therefore provide an end-to-end view of the planning and execution process.

2 Related Work

From a distributed goal search perspective, the separation of agent local reasoning and coordination is a natural result of applying the divide-and-conquer methodology, for example in the works of Lesser [27], Durfee and Montgomery [15], etc. The problem can be visualized by a distributed goal search tree, as illustrated in the work of Lesser [27] and Jennings [24]. The global search space is then divided into several local search spaces and local agent reasoning is the process of each agent working on its local goals. In this view, local reasoning is essentially a planning and scheduling process. Single-agent planning methods can be applied to solve the local search problem. In fact, many distributed planning frameworks start from single agent planning systems and extend them to distributed ones, notably Corkill's work [8] extending the planner NOAH [32] and the work of desJardins and Wolverton extending the SIPE-2 planner [44] to form the DSIPE distributed planning system [12]. However, these extensions are often *ad hoc*. The lack of a common representation for defining multiagent planning context leads to difficulties in generalizing their approaches to other domains and limits their applicability.

Many other multiagent planning methods are not extensions of single agent planning approach but they too adopt the separation of planning and coordination. There, the local planning process is often uniquely designed to operate with the model of coordination used, but nonetheless is based on some form of local search process with nonlocal constraints. These constraints can be viewed as a form of multiagent planning context although they are not formalized nor complete. Examples include the Partial Global Planning (PGP) work by Durfee *et al*, where some form of greedy search techniques is used, and in Tambe's work on teamwork activities [37], where agents reason to find a locally optimal solution based on a set of predefined fully planned teamwork activity rules.

Among these planning approaches the one that comes closest to a formal definition of multiagent planning contexts is the Design-to-Criteria (DTC) scheduler [39, 40, 43, 41, 42] by Wagner *et al.* DTC is designed to work with the Generalized Partial Global Planning (GPGP) coordination framework and is a sophisticated heuristic planning and scheduling framework. DTC and GPGP are based on the domain-independent TAEMS [11] modeling framework and together they allow agents to cooperate under uncertainty and also adapt to different environment characteristics. In this paper we will discuss DTC in some detail and use DTC as to show how uncertainty handling of coordination activities affects local agent reasoning, and how our enriched model of agent commitments can be used in conjunction with agent planning and scheduling techniques.

Once the goals are distributed, an agent may apply any individual (single-agent) planning approach, provided that the proper multiagent context is represented in that approach. This is typically done by the use of commitments to decouple the subgoal interaction problem. A number of commitment semantics have been proposed, for example, the *Deadline* commitment $C(T, Q, t_{dl})$ in [10], means a commitment to do (achieve quality Q or above) a task T at a certain time t (decided by the agent who pledges this action) so that it finishes before a specified 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. This way, the commitment reduces nonlocal interrelationships to external constraints and therefore de-couple the agent plans. For example, if a task A to be performed by agent X depends on the successful completion of task B in a different agent Y (the "B enables A" interrelationship), a commitment by agent Y that offers to perform B (with parameters such as the time-frame of B's completion) would decouple the two agents plans and translate the interrelationship into local constraints.

Jennings [25] identified three of the most common mechanisms for managing the coordination process in DAI systems: organizational structuring, exchanging meta-level information, and multiagent planning. Jennings also points out that in all three cases, *commitments* (pledges to take a specific course of action) and *conventions* (means of monitoring commitments in changing circumstances) are the foundation of coordination. In this paper we extend the notion of commitments to be dynamic objects and the monitoring of commitments becomes an integral part of our model. Therefore, our model extends *conventions* from a set of predefined rules to dynamic decisions with uncertainty in commitments.

At an organizational structure level, the role each agent takes in the organization corresponds to a long-term, high-level commitment about the types of activities the agent will pursue. The problem solving is simplified because an agent can assume that all other agents will adhere to there roles and it also understands that all other agents will assume it will adhere to its role during problem solving. Norms [7, 13], conventions [24], and social laws [35] fall into this category. Organizational structures, such as hierarchies, can be viewed as implicit rules of coordination that specify the pattern of information and control relationships between agents [19], therefore are a form of long-term commitment as well. An example of using this type of coordination is the distributed vehicle monitoring testbed (DVMT) [26] where each agent has its own area of the search space, i.e., its sensor area.

Meta-level information exchange [19] involves agents sending each other control-level information about their current priorities and focus. The exchanged information provides some level of detail about one agent's goal search subtree, so that other agents can build a representation about this agent's activities. Such information acts as an influence to the agent local planning, not as control or a constraint. Examples include Durfee and Lesser's [16] Partial Global Planning (PGP), and later the domain-independent Generalized Partial Global Planning (GPGP) approach by Decker and Lesser [10]. In the next section we will examine GPGP in some detail. Commitments, as explicitly used in GPGP, are exchanged between the agents and serve as the way coordination exerts influence on local planning. An agent may offer a commitment to another agent to resolve a coordination relationship, but it is not mandated that the receiving agent must utilize this commitment. Also, agents may alter their local plans in an autonomous fashion and therefore commitments may be changed. These type of meta-level communication can be viewed as a medium-term source of knowledge regarding an agent's commitments.

Multiagent planning, in a narrow sense, means the process of producing a multiagent plan that requires all agents agreeing on all of their activities before they start acting. This applies to to both centralized multiagent planning and distributed multiagent planning, Centralized multiagent planning aims at building a central plan, either by a central planner or a central coordinator. For example, in the work of Cammarata *et al* [2] the agents (aircrafts) first choose one agent to produce a conflict-free plan for all of the agents. Such a plan often specifies the joint actions of agents. Alternatively, Georgeff [20] proposes another approach: first, each agent makes its own plan, then a central planning agent collects all the local plans, analyzes the conflicts among them, and finally, modifies them to resolve the conflicts (and at the same time specify the synchronization activities among them.) However, building a central plan involves strong assumptions such as deterministic actions or global knowledge of the system states.

With distributed multiagent planning, no agent has the global view and the plan is divided among the agents. Generally, agents' plans conflict with each other, so a plan synchronization process is needed. The agents' plans then would converge on identical plans by information exchange and conflict resolution. Some examples of distributed multiagent planning include Corkill's [8] distributed hierarchical planner based on NOAH, desJardins and Wolverton's [12] DSPIE planner, von Martial's [38] model for coordinating plans, and Conry et al's [6] multistage negotiation protocol for cooperatively resolving conflicts. Commitments are the foundation of this approach since it is required that the agents keep their pledges as specified in the final plan that all agents agree on. Only when the situation changes radically during execution would replanning be performed.

A common problem in these planning approaches is the issue of uncertainty in coordination. Many planning frameworks, for example the SharedPlans model of Grosz and Kraus [22], often depend on assumptions such as complete system knowledge and deterministic action. As a result they represent global plans and do not deal explicitly with uncertainty. The prevailing view of regarding coordination as identifying and managing constraints to the local search process in a distributed goal search handles uncertainty in the system in a limited way. For example, social conventions can be applied in case of a plan failure (e.g., a failed commitment) [24]. However, relying on fixed coordination mechanism may not be appropriate in all situations. For example, in the work of Castelfranchi *et al* [4], it is shown that in some cases norms may be violated deliberately in order to adhere to a more important goal. This would require a better representation of the planning context that enables the understanding of the specific situation and allows planning flexibility. A dynamic model is needed to capture the changes in the system and enable dynamic evaluation.

Sophisticated local planners (such as DTC) have the ability of handling uncertainty, but their ability is limited to the uncertainty in local actions. For example, an extension to the DTC scheduler deals with building contingency plans to handle uncertain events such as possible failures of local actions [42]. However, due to the lack of a formal representation for uncertainty related to coordination, such extension is still based on the static model of commitments. In this paper we will propose a dynamic model to represent a multiagent planning context that incorporates uncertainty in coordination. The idea of adding more structures to the plan, i.e., introducing contingency plans, is an important step toward understanding and exploiting uncertainty in coordination, and we will show how to expand the use of contingency planning [30] in DTC to respond to uncertainty in agent commitments.

An alternative approach is to model planning and coordination in an integrated decision theoretic framework that deals with uncertainty. Markov decision processes (MDPs) or partially observable Markov decision processes (POMDPs) are often used in representing single-agent planning problems. In these models, a plan is represented through a decision policy, which can also be viewed as a set of contingency plans. The concept of a multiagent planning context can be viewed as the information needed for an agent to expand its local view and construct a new single agent decision process. More recently, the use of these models on multiagent planning problems are being discussed [1, 46, 31]. There, the issue of communication is of central importance. A major problem with this approach is the complexity of the theoretic model - it is unlikely that efficient and practical solutions will be found, although there are heuristic approaches for special classes of those problems [21]. But, although these models typically use low level representations such as states and actions, high-level concepts such a commitments can find their interpretations in the solutions. Thus, these models may indeed provide theoretical underpinnings for our approach. More importantly, these concepts are very important in developing heuristics and approximation solutions[46], which are crucial to solving multiagent decision problems due to their complexity.

3 Commitments as Planning Contexts

In this section we examine the issue of representing multiagent planning context by looking at how existing approaches deal with planning and coordination. In Figure 1, we illustrate a two-agent cooperation problem using the TAEMS representation. The deadline for this overall task is time 160. In this example problem, there are two nonlocal interrelationships, namely the *enables* relationships "$A2$ enables $B2$" and "$A4$ enables $B4$". Also, each of the low-level tasks ($A1$ to $A4$ and $B1$ to $B4$) has nondeterministic outcomes. This is characterized through a discrete probability distribution of the outcome quality (i.e. utility) values. For example, q(20% 0)(80% 6) means 20% chance of having quality 0 and 80% chance quality 6. For the agents to maximize overall utility (which is the sum of the local utilities in each agent in this case), they need to take into consideration not only the local constraints (such as the local *enables* relationships, task deadlines, the quality accumulation functions that decide how utility propagates from low-level tasks to top-level tasks, and the uncertain outcomes), but also nonlocal interrelationships.

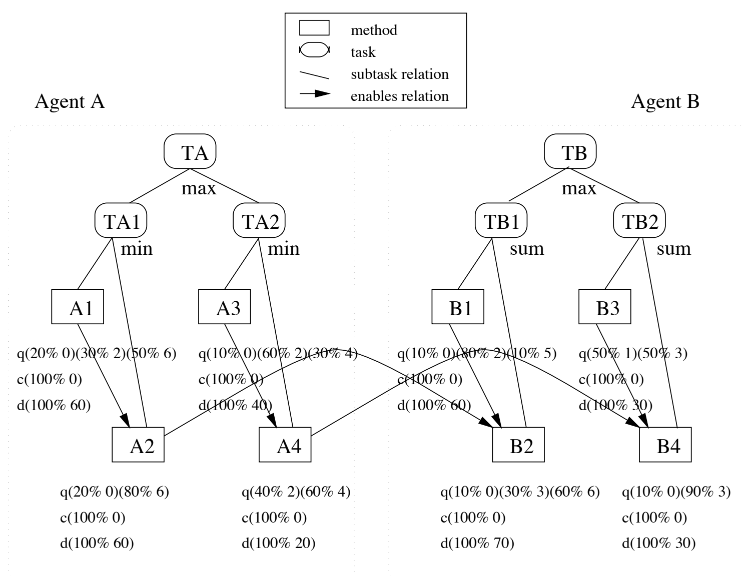

Figure 1: Example Task Structure

However, we need to note that what Figure 1 shows is the global view of the problem. Each agent's local view (i.e., the subjective view) contains only the local portion of the TAEMS task structure plus the knowledge of the nonlocal interrelationships (i.e., the names of the tasks involved in the nonlocal interrelationships). An agent does not know the quality and duration distribution of the remote tasks in those interrelationships. Figure 2 shows each agent's partial view.

Without any additional information about each other's plan, the agents may try to maximize

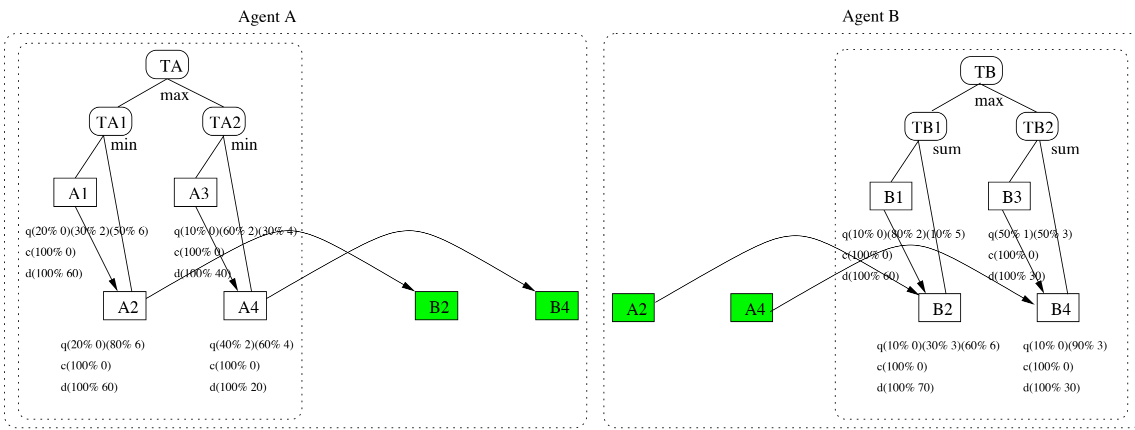

Figure 2: Local Views

their expected local utility. For example, by using the DTC scheduler, initially A would select schedule $(A1, A2)$ and B would select $(B1, B2)$. Then, agents would apply coordination mechanisms to handle the nonlocal relationships. For example, by using the GPGP coordination mechanism, when agent A detects the interrelationship between $A2$ and $B2$, A would send out a control message and proactively pledge to complete $A2$ by time 120, with some estimated quality. This basically provides a planning context for agent B. Agent B can understand this commitment and take it into consideration in its own planning because B can relate the commitment to the nonlocal interrelationship that enables $B2$. Alternatively, if A simply tells B its complete schedule, B would not be able to understand it completely because B knows nothing about the task $A1$, even though $A1$ is fairly important to task $A2$ because $A1$ enables $A2$.

This suggests that commitments may be the appropriate representation for describing the planning context since they provide some degree of plan encapsulation. However, such a simple context is often not sufficient for coordination at a finer and more effective level. For example, because $A2$ and $A1$ both may fail, the agents' plans need to change accordingly to react to those events. The agent can either reschedule when the events occur, or develop *contingency plans* before hand to specify what alternative plan to perform when the events occur. In Figure 3, (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, whereas a contingency schedule specifies a set of paths based on possible future outcomes. Using contingency analysis, the utility value of a schedule is now computed based on this branching structure, and therefore is more accurate.

Here, the events that trigger contingency plans may not always be reflected in the context provided by the simple commitments model. For example, the failure of $A1$ may cause B to select task $B3$, but the commitment about $A2$ says nothing about how $A1$ may affect $A2$. To address this problem, we need to view commitments as dynamic objects and specify how a commitment may evolve over time. For example, instead of telling B that $A1$ may fail at time 60, A may tell B that the commitment about $A2$ is inherently uncertain and may fail at time 60. The latter is more understandable to B since it does not require B to understand details about task $A1$, yet at the same time fully accounts for the effects of $A1$ with regard to the commitment. As such, the reasoning for building contingency plans in agent B can be simplified.

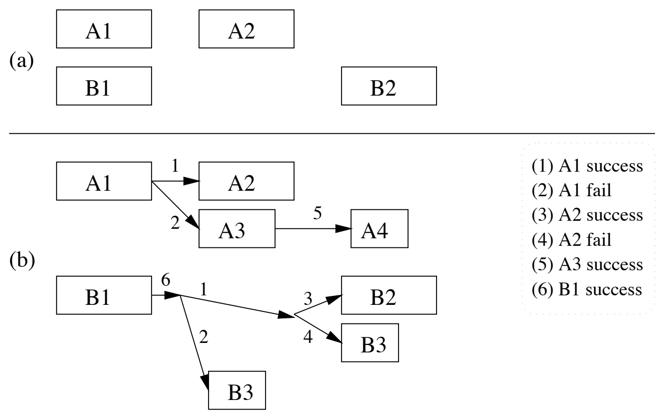

Figure 3: Schedules with contingency

This is just one of the types of uncertainty that has not been modeled in the existing model of commitments. This uncertainty reflects 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 being pledged 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. In other words, we describe commitments as dynamic objects, and specify their statistical guarantee semantics.

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 its commitments [25]. 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 originates from the existence of commitment itself, not from the underlying tasks. In other words, it is inherent to the making of the commitment itself and not from possible under-performance of tasks, which is already addressed as the first source of uncertainty. In this paper we take into account this uncertainty by explicitly describing the possibility of *future* modification/revocation of the commitment. Contingency planning [30], 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. Also, a number of approaches have been proposed to handle this particular type of uncertainty, such as using a leveled commitments contracting protocol [34] and using option pricing schemes for evaluating contracts [36].

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 whether keeping a commitment is beneficial to the social utility or not. To tackle this problem we define the *marginal gain or loss* [33] value of commitment and relate

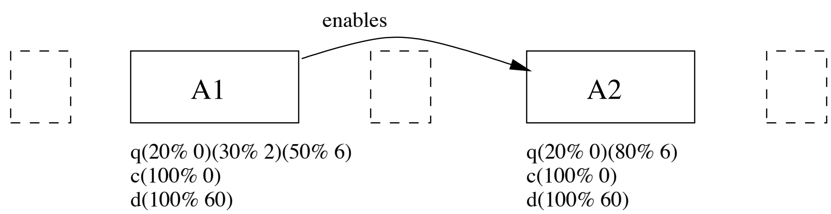

Figure 4: Uncertainty in Commitment

that to the social utility. Thus, the agents can reason at the social utility level instead of at the local utility level.

For coordination to be successful when there are these forms of uncertainties, there must be structures that allow agents to interact predictably and flexibly for dynamic environments and imprecise viewpoints, in addition to the local reasoning capability [14]. For this purpose, we propose a domain-independent, flexible framework for the agents to manage their commitments. Our work differs from the *conventions* and *social conventions* [24, 25] 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.

4 Incorporating Uncertainty in Commitments

The first step for incorporating uncertainty in commitments is to quantify the uncertain outcomes of underlying tasks. In [10], a commitment specifies only the *expected* quality (q) of the committed task. However, expected values often do not provide sufficient information for effective coordination, especially when there are possible task failures[1] For example, Figure 4 shows some tasks in the schedule of agent A. Suppose A offers a commitment about method $A2$ to the agent B, and assume $A2$ is to be enabled by another local method $A1$. In this case, $A2$ itself has an expected quality of 4.8 ((20% 0)(80% 6)). However, there is a 20% chance that method $A2$ will fail (q=0), and thus cannot be useful to B. Furthermore, because $A2$ is enabled by $A1$, which also fails in 20% of the time, the result is that the commitment has only a 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, if $C(T)$ is a commitment about task T, the outcome distribution of $C(T)$ (i.e., p(C(T)), or equivalently, the *actual* outcome of task T, p(T)) depends not only on the outcome distribution of T ($p^\circ(T)$, which does not take into account the effect of interrelationships), but also depends on the outcome distributions of the set of predecessor tasks of T, pred(T). A predecessor task of T is a task that either enables T, or has some other interrelationships with T that may change the outcome distribution of T. Obviously, the outcome of a predecessor task (i.e., p(M)) in turn depends on the outcomes of its own predecessors. In the simplest case, let us assume that the only source of uncertainty comes from method quality,

[1]It is assumed that q=0 means the method fails and in this situation any method enabled by it cannot proceed.

and only enables interrelationships exist, then, the probability of the quality outcome equals x

$$p(\mathrm{q}(\mathrm{C}(\mathrm{T}))\mathrm{=x}) = (\prod_{\mathrm{M}\in \mathrm{pred}(\mathrm{T})} p(\mathrm{q}(\mathrm{M})\mathrm{>}0)) \cdot p^{\circ}(\mathrm{q}(\mathrm{T})\mathrm{=x}) \tag{1}$$

Probability propagation of general cases that involves duration, cost, as well as other types of interrelationships can be similarly deduced. This way, an agent can tell the other agent about the outcome profile of the commitment without the need to reveal details of its schedule, which the other agent needs not to know.

Next, because commitments are 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. In our example in Figure 1, initially at time 0, agent A chooses the plan $A1, A2$ and offers commitment about $A2$ to B. However, at time 60, when $A1$ completes, in the case that $A1$ fails or has $q = 2$, A may replan and select an alternative plan that can produce a better (local) expected quality outcome. Clearly, B could be informed *at time 60* rather than waiting to notice that the commitment is 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 take into account that possibility and not heavily depend on this commitment. On the other hand, if at time 60 $A1$ finishes with quality 6, then the quality outcome of the commitment has *updated* 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 sent 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 about (1) when and what events will trigger re-scheduling (or cause the current plan to be interrupted, for example, due to the frequent arrival of more important tasks that cause the local agent to never complete its commitment), and (2) whether or not a future re-scheduling would lead to changes in commitments — in other words, information that allows a more accurate prediction of future events, decisions, and actions. These information could also be valuable for an agent to determine how much meta-level control (planning, scheduling, and re-scheduling, etc) information is associated with a scenario. Obviously, for complex systems, this information can 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 means 50% of the time re-scheduling 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 update point for the commitment offered by A.

The update points are calculated by analyzing the schedule to see at what times a failure or low performance of a method could seriously affect the performance goal of the agent in the future.

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 update point 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'. Such a model of commitments specifies not only the possible outcome of the commitments but also their dynamics. The agent receiving the commitment thus can view it as a simplified view of the offering agent's schedule that specifies only the relevant context and abstracts away the details it does not need to know.

The third source of uncertainty comes from the partial knowledge of the other agent, namely, how important this commitment is to others. To answer this question we first need to know how important this commitment is to this agent. By knowing this we can avoid bad coordination situations such as offering a commitment (and paying the high cost of honoring it by not rescheduling to achieve higher local quality) that is of little value to the receiving agent, or in the contrary, canceling a commitment that is very important to the agent that needs it for only little gain in local performance. To solve this problem we use the notion of *marginal cost* and define the *marginal loss* as the performance difference between the schedule without making the commitment and the schedule 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 in agent performance when receiving the commitment and when not 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 agents may use different utility scales. Thus, we use the relative *importance* to indicate how quality values in the other agent translate to the quality values in this agent. For example, agent A may believe that utility in agent B has importance 2.0, i.e., the utility in agent B equals twice the amount in A. Thus, it implies that a marginal gain of 5 in B can offset marginal loss of 10 in agent A. Clearly, a rational agent would try to maximize the value of its local utility plus marginal gain in other agents and minus the marginal loss due to the commitment it offered. 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 general, though, a simple importance rating is not enough to characterize an agent's utility model or the group utility function. a more complex model, such as the MQ model by Wagner [39], could be used. In order to evaluate the marginal gain/loss against a particular commitment, we simply compare the best-quality alternatives with and without the commitment, and use the difference as the marginal gain/loss.

As a result of the above discussion, Figure 5 shows the extended TAEMS specification of an example commitment, which pledges to do task $A2$. Such a specification of a commitment is derived from the schedule of the offering agent and communicated to the receiving agent. The commitment thus serves as the consensus between the agents: it accurately reflects its dynamics in the offering agent and provides a complete yet compact context for the receiving agent. It can be viewed as a generalization of the static, constraints-like commitment in terms of coordination, but more importantly, it provides a well contained multiagent planning context, which simplifies

```
(spec_uncertain_commitment
  (label com1) (from_agent agentA) (to_agent agentB)
  (task A2)
  (type deadline)
  (outcomes              ;; -- uncertain outcomes
    (o1
      (density 100%)
      (quality 6 64% 0 36%) (finish_time 120 100%)))
  (update                ;; -- list of possible update points
    (u1
      (lowerbound 50%) (update_time 60 100%)))
  (marginal_loss 0.0)    ;; -- no marginal cost to agent A
  ...
)
```

Figure 5: Commitment that incorporates uncertainties

coordination in uncertain environments and defines a clear interface for local scheduling. Different local scheduling methods can be used as long as the multiagent planning context is incorporated in those methods.

5 The Impact on Planning and Scheduling

Now that a commitment has uncertainty associated with it, agents can no longer regard a commitment as guaranteed, or assume the absence of failures. Therefore, planning and scheduling in an agent become 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 in 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 is capable of planning *in anticipation* of possible failures and knows the options if failures do occur. This way, necessary arrangements can be made before the failure may occur, and we also save the effort of re-scheduling by adopting a planned-ahead action in case of failure.

To handle possible failure outcomes in commitments, we use *contingency analysis* in conjunction with the Design-to-Criteria scheduling. The use of contingency analysis has been introduced into DTC to offer an alternative way of dealing with local task failures instead of rescheduling [42]. In our approach, a failure in a commitment will be treated the same as the way a failure in a local task is treated. First, we analyze the possible task failures (or low quality outcomes) or commitment failures and identify alternatives that may improve the overall quality outcome 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.

Contingency analysis can also be used to handle uncertainty originating 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 occurs in the critical regions, thus leading to the discovery of update points. On the other hand, once we have the update point information regarding a commitment, we can make contingency schedules to specify a *recovery option*. Let T^α indicate that task T has outcome α, 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 a commitment as a type of de-commitment that 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, at least in the absence of failures), these changes are *social* rather than local. The introduction of marginal gain/loss ensures that commitments are properly respected in a social context. If the overall utility of a multiagent system 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 is a commitment socially worthwhile. Likewise, the commitment should be revoked only where the marginal loss is greater than the gain or when the commitment cannot be honored due to task failures. 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.

6 The Communication/Negotiation Framework

In order to add flexibility to coordination, we also introduce a commitment communication framework that allows agents to interact with each other in order to achieve better coordination. This framework provides 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 asks 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 associated uncertainty) and possible marginal loss.
- *accept*: RA accepts the terms 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 (or later) possible update points/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 failed (due to unfavorable task outcomes).

These primitives are used not only during the establishment of commitment, but also during the problem solving process. The communication process is a manifestation of the lifecycle of a commitment. Agents use these primitives to negotiate and communicate their commitments dynamically during the problem solving period. The negotiation (and communication) process keeps agents better informed about each other's desires, intentions, and outcomes, and therefore reduces the uncertainty in commitments resulting in better coordination. In our case, using the example problem in Figure 1, the following negotiation steps may be taken:

1. At time 0, both agents perform local planning based on their partial views (see Figure 2) of the task structure. The locally optimal plan (using the Design-to-Criteria scheduler) for A is to perform $(A1, A2)$ (with expected utility 2.88), and the optimal for B is $(B1, B2)$ (with expected utility 6.15 - ignoring the enabling relationship for $B2$).
2. Agent A then *proposes* to B that it intends to offer a a commitment about task $A2$ before time 120. This type of proposal is based on one of the coordination mechanisms defined in GPGP - handling hard interrelationships - in which the agent that performs the enabling task proactively offers a commitment about the enabling task to other agent. Note that the other agent may or may not have the enabled task in its plan. The offering agent offer the commitment without this knowledge. Also, since the commitment is derived from A's current schedule, the marginal loss for A is 0.
3. Now B receives the proposed commitment. Since B's current plan is $(B1, B2)$, according to the deadline (160), $B2$ has to be enabled by time 90. Thus, B's reaction to the offer certainly depends on how much detail is included in the commitment. If only a simple deadline is used, B has no way of telling whether $A2$ could actually finish by time 90, and thus cannot accurately predict how useful the commitment would be. But if the commitment is specified according to the richer model we have proposed in this paper, B understands that it is impossible for $A2$ to finish by time 90. Thus, the commitment is useless.
4. Now B has a better assessment of its local schedule. Since $B2$ couldn't be enabled on time, the actual schedule would be simply $(B1)$, with expected utility 2.1. This is not locally optimal now, because an alternative plan $(B3, B4)$ would have better expected quality (4.7, ignoring

the *enables* relationship between $A4$ and $B4$ for now - which is the assumption of DTC). To achieve this expected utility, the alternative schedule would require a commitment from A to complete $A4$ with positive quality by time 130 *with 100% certainty*.

5. Thus, B would *counter-propose* to A - basically, telling A that the marginal gain for the commitment about $A2$ is 0, and the following commitment is desired instead: to finish $A4$ by time 130 with 100% certainty. The marginal gain for B would be 2.6 (4.7-2.1).

6. A receives this counter proposal, and it sees that there is an alternative schedule $(A3, A4)$ - expected utility 2.16, that could partially satisfy the counter proposal, but doing so would incur a 0.72 (2.88-2.16) marginal loss. The certainty factor would be 90%, not the 100% proposed, so A can estimate that the overall value of the commitment for the commitment about $A4$ would be $(90\% \times 2.6 - 0.72) = 1.62$, so the new commitment would bring this much increase to overall utility compared to the first proposed commitment.

7. A now proposes again, this time with the new commitment about $A4$, marginal loss 0.72. Under the new model, this commitment also specifies the performance distribution ((10% 0)(36% 2)(54% 4)) and the update points (at time 40, with at least 90% chance keeping the commitment at that time).

8. B receives the new proposal and checks again with current schedule $(B1)$. Still, the alternative schedule $(B3, b4)$ performs much better: the expected utility (assuming $B4$ enabled 90% of the times) is 4.43, which means the marginal gain for the commitment is 2.33. Thus, the value of the new commitment is 1.61.

9. B has no other alternatives, so B would *accept* this commitment. Thus, both agents adopt this new commitment.

The above steps does not take into account the contingency planning options in both agents (or the effect of rescheduling in DTC). Actually, if rescheduling or full contingency plans are used in the scheduler (see the analysis in the next section), the expected utility would be slightly greater.

Clearly, the negotiation process as discussed above helps the discovery of alternative commitments that leads to better social solutions. This is done by using marginal gain/loss information in negotiation. Without that 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 about 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 [18]. 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 [29]. However, a thorough discussion of complex negotiation frameworks is beyond the scope of this paper. 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. A more complex and recent work can be found in the MultiStep negotiation mechanism by Zhang and Lesser [47], which is also based on the use of marginal gains/costs but with a multi-dimensional utility function.

7 Experiments and Analysis

In order to validate our approach, we implemented a generic agent that can work with a textual TAEMS input. We simulate two instances of such agents, A and B, to work on the task structures presented in Figure 1. We use experimental data 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 offers a commitment to the other agent, using only expected quality and finish time. This is exactly the DTC/GPGP approach described in the example in the previous section (steps 1-3, but using the old commitment semantics): agents both apply the DTC scheduler and invoke the GPGP one-shot (no negotiation) coordination mechanism 4 (handling hard coordination relationships). In Figure 6 we show the percentage distribution of the final quality outcomes for 200 runs. From the left to the right, the figures show the quality result of agent A, agent B, and the sum, respectively.

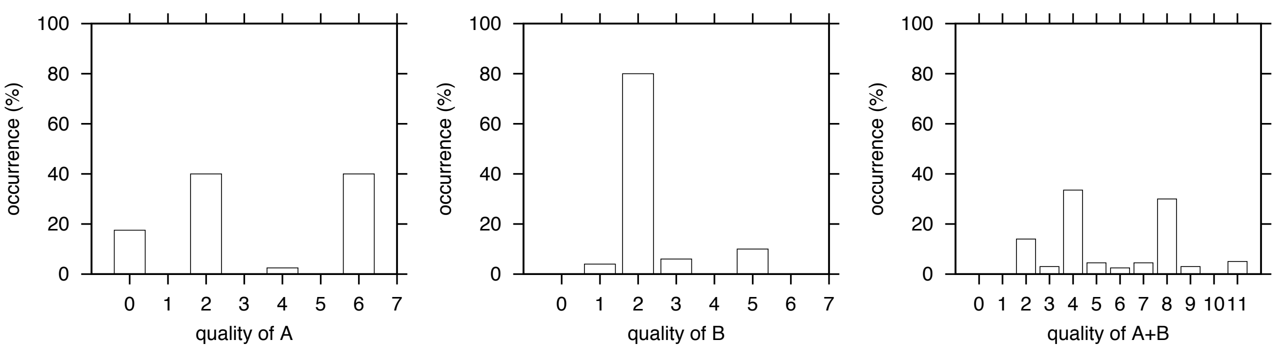

Figure 6: Base Case

From the trace, we observed that A's commitment to finish $A2$ 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, including the probabilistic outcomes and update points. The commitment is still pro-active (with no negotiation), but the agents can now use contingency planning to reduce the uncertainty in commitments. In this case, contingency branches for the DTC scheduler is added, which brings local improvements to the schedule. Also, the agents can now update the commitment status at runtime, which is an addition to the GPGP coordination mechanisms. Figure 7 shows the results for 200 runs.

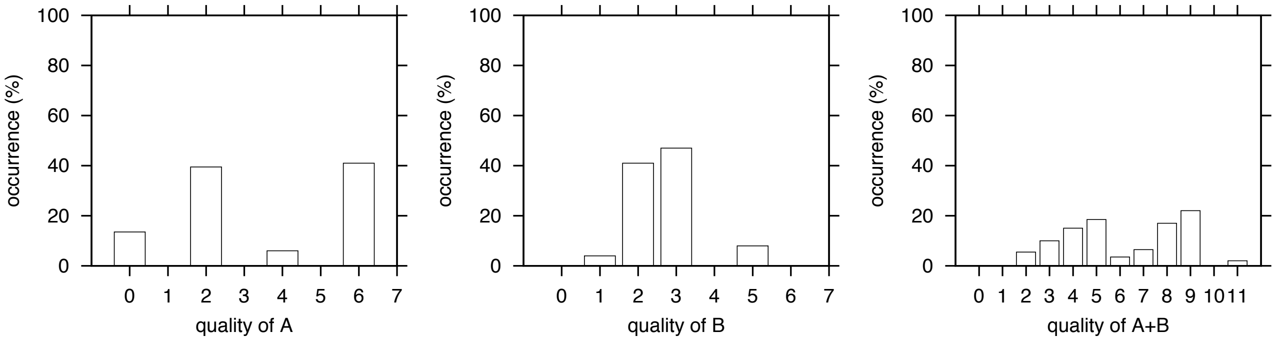

Figure 7: Second Case: With Uncertainty

Here we can see some slight improvement of quality outcomes in both agents compared to the base case, but the similar pattern of distributions suggests that this has only minor impact on the scheduling. Essentially, agents now produce contingency plans so that they can switch to a better alternative when an undesirable situation occurs. This can be viewed as an incremental improvement to the existing plan, thus it is not going to change the agent behavior pattern significantly. Due to no negotiation, the improvements are restricted to agent's local activities. For example, we notice that when $A1$ finishes with quality 2, A will choose to switch to plan $(A3, A4)$ instead of continuing to run $A2$ (therefore effectively de-committing from its commitment) because now $(A3, A4)$ has higher local expected (local) utility. Similarly, B does not need to wait at time 120 to know that the commitment is not coming - in fact, B would know much earlier - at time 0. Thus B would know that its schedule will fail and would switch to $B3$ when $B1$ has a low quality outcome.

As the last case, we incorporate negotiation and using the marginal gain/loss information in commitment coordination. The results, shown in Figure 8, have very different patterns in the distributions. 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 $A4$) now, through negotiation with marginal costs. This commitment is social in that it helps to achieve better overall utility, although not all agents have local gains.

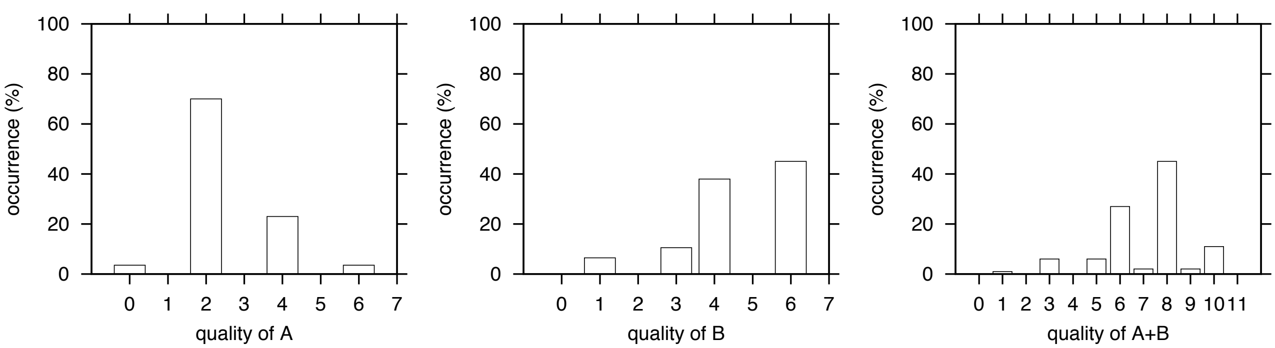

Figure 8: Third Case: Negotiation

Table 1 shows the average quality outcomes in each case: clearly, in case 2, we see that both

Average Quality	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

Table 1: Average Quality Values

A and B improves their average local utility, but in case 3 the average utility of A becomes lower and the average utility of B becomes much higher. This shows the different aspects of improvement our approach may bring to agent cooperation: first, the *incremental* improvement to local agent planning due to better uncertainty model of commitments and sophisticated planning (the second case). This can be seen as improving the existing schedule and therefore average local utility increases. The second aspect of improvement is on the selection of better plans by enabling the search of better social plans (using marginal cost/gain information during the establishment of agent commitments, as in the third case). In this case, agents may have different plans from the locally optimal plan, and one agent may have lower average utility but the average social utility is guaranteed to improve. This also shows that the integration of all the mechanisms: negotiation, contingency planning, and marginal gains/losses is very important in effectively handling the uncertainties. These mechanisms handle different aspects of uncertainty, and they work together to achieve better coordination.

The example problem in our experiments is a simple one. Indeed, we can extract the scheduling data from the simulation study and obtain the actual schedules used, and then perform quantitative analyses. Figure 9 shows the actual schedules for the three cases.

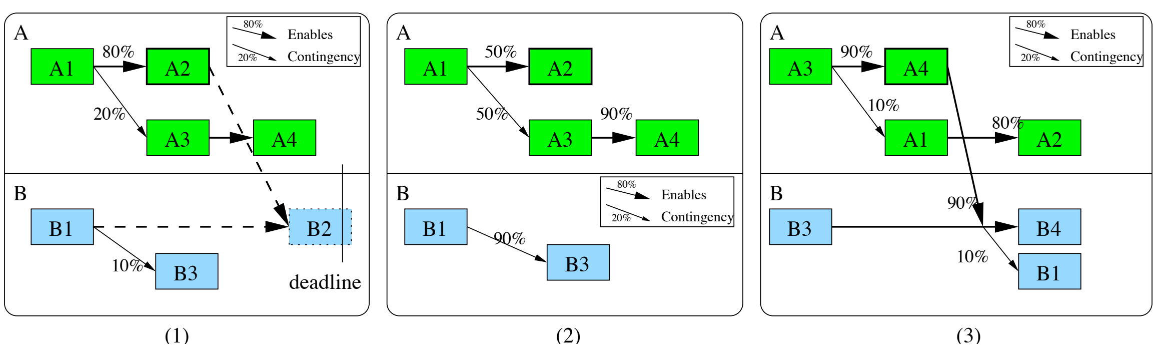

Figure 9: The Plans Used in Three Cases

The calculation of the expected qualities is straightforward according to the TAEMS model. The result is in Table 2. It confirms the simulation results in Table 1.

In both (1) and (2), agent A offers B a commitment about task A2. In (1), the original DTC/GPGP approach, agent A has an initial schedule of (A1,A2) but would change schedule to (A3,A4) if task A1 fails. Agent B's plan is to perform (B1,B2) and switch to B3 if B1 fails. However, because B has a tight deadline, the commitment about A2 cannot arrive in time to enable B2 and therefore cannot contribute to agent B.

In case 2, the schedules seem identical to those of case 1, but in fact the conditions for invoking

Expected Quality	A	B	sum(A,B)
Case 1	3.312	2.3	5.612
Case 2	3.48	2.7	6.18
Case 3	2.448	4.5	6.948

Table 2: Expected Quality (calculated)

the contingency plans are different. This can be verified by the probability numbers associated with the contingency branches. For example, in (1), there are only 20% chance (i.e., when A1 fails) that agent A would switch to (A3,A4). But in (2), the probability is 50%, meaning that the contingency plan will be invoked when A1 fails or A1 has quality 2. The reason is that the contingency plan has better expected local utility (its expected quality is 2.16 compared to 1.6 when A1 has quality outcome 2). This is an example of a local decommitment as discussed in the second source of uncertainty (section 4). The commitment still cannot be used by agent B, but with more information about the commitment, B knows that it should not expect the commitment to come in time. Thus, B would modify the contingency plan so that task B3 is invoked 90% of the time (when B1 has quality 0 or 2) instead of the 10% in the previous case (only when B1 fails, i.e., has quality 0).

Case 2 shows how to improve local scheduling by the use of contingency planning and dynamic updating. The changes in the schedules reflect local incremental improvements to the previous case, because the use of each contingency branch is based on calculated gains in expected local utility.

Case 3 involves a different commitment (A3) as a result of the negotiation process for establishing commitments using marginal cost information. Specifically, at the beginning of the negotiation, the agents start with the schedules in case 2. It can be calculated that the commitment about task A2 has 0 marginal loss to A (since A2 is already in A's best local schedule) and 0 marginal gain for B (because it has no effect due to the deadline.) However, a commitment about A3 would have different marginal values: the marginal loss to A is $3.48 - 2.448 = 1.032$, and the marginal gain for B is $4.5 - 2.7 = 1.8$. The net gain in total expected utility is $6.948 - 6.18 = 0.768$. Thus, B would counterpropose the commitment about A3 (arguing for a gain of 1.8 toward its local utility), and agent A would perform a new search and find that the new commitment would lose 1.032 to the local utility. Thus, agent A can conclude that the new commitment would lead to a 0.768 increase to the total utility compared to the schedules at the beginning of the negotiation. Therefore, both agents would agree on the new commitment, which lead to new schedules with better total utility.

Both case 2 and 3 reflect techniques that are monotonic, i.e., ensuring that the new schedules are better, although one emphasizes local improvements while the other emphasizes the social welfare. Of course, such improvements are based on a more complex representation and thus require more sophisticated reasoning ability. They would require more computation and also more communication, both in terms of the size of the information exchanged and the frequency of the exchanges.

For example, let us look at the amount of communication during execution. Assume each communication count as 1. In case 1, the only communication involved is to let agent B know that A2 has finished. Thus, the amount of communication is 0.8 (since there is 0.8 probability that A2 will be performed.) In case 2, agent A would also send an update when A1 finishes (with a 0.5 probability that A would not decommit). So the amount of communication in case 2 would be 1.5. In case 3, the commitment about A4 has one update point when A3 finishes, with a 0.9 probability that A would not decommit. So the amount of communication would be 1.9. Of course, this is a simple estimate. Communication may be reduced if the agents use some conventions such as not to communicate when a commitment is fulfilled. Also, this does not account for the communication needed for establishing the commitments. In fact, there may be a long communication process during the negotiation procedure in case 3. It should be noted that when a schedule is very complicated, updates about the commitments may be frequent. In those cases we may need to use simplifying techniques to reduce the update frequency. For example, we may use threshold conditions such as "the commitment's success rate drop below 50%" to decide if update communication is needed.

Another advantage for adopting a complete, self-contained commitment model is to reduce the need for re-scheduling. Typically re-scheduling is a very heavy-weight operation. Thus, although contingency analysis adds overhead to scheduling, it reduces the overall scheduling cost. This is particularly useful for sophisticated schedulers such as DTC, which already does a lot of computation about alternative plans and can perform contingency analysis quite easily.

8 Conclusion

In summary, we identified three sources of uncertainty inherent in commitments and discussed 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 formalize a model of multiagent planning context and to improve coordination effectiveness, and ultimately, to improve the overall utility of the multiagent problem solving. Our results indicate that these improve coordination. The features of this model include:

- The identification of three types of uncertainty associated with commitments: uncertainty due to uncertain task outcomes, uncertainty due to possible de-commitment actions, and uncertainty due to incomplete knowledge about other agents' activities. The effect of these uncertainties are quantitatively represented through a number of parameters such as the probabilistic distribution of possible outcomes, times and chances of possible update and de-commit events, and the marginal gain or loss of utility.

- Explicit reasoning at the global utility level: by introducing marginal gain and loss, an agent's local reasoning also leads to improvement of global utility. This is very important for the negotiation of commitments since it provides a direction, i.e., toward the increase of total expected utility.

- Quantification: the above parameters provide the necessary context for the agents to quantitatively reason about the impact of the commitments to their planning and scheduling, and together with the quantitative reasoning model in the agents, we can evaluate the problem solving performance and hence also the effectiveness of coordination mechanisms.

- Statistical guarantee semantics: due to the uncertainty involved, the guarantee implied in a commitment is not strict. Our model explicitly specifies the nature of the guarantee and therefore keeps the guess work out of the agent planning, and also simplifies the contingency planning problem since it specifies when contingency plans would be needed.

- The view of a commitment as a dynamic, evolving object: it may be modified, updated, or de-committed. During the problem solving, the events unfold and uncertainties are replaced by factual outcomes. Therefore, the probabilistic expectation about the commitment also changes over time when new information comes. Also, this dynamic model of commitments facilitates the monitoring process of multiagent cooperation since the changes in commitments indicate changes in the plan context, and therefore serve as possible events to be noted during monitoring.

- Proactive control: by explicitly introducing uncertainty into commitments, agents can make contingency plans that anticipate uncertain outcomes and/or take actions before the outcome arrives. This is an example of *proactive control* in agent planning, as compared to reactive control where an agent responds to an event after the event occurs.

- Forward looking ability: for an uncertain commitment, we are not only interested in knowing what may happen to it, but also interested in knowing *when* information about the commitment may arrive. This is important because the timing of such information may affect the choice of agent actions. For example, a person planning a trip sometimes will face the uncertainty of whether or not airplane tickets for a particular date would be available. Such information may affect the decision of whether or not to travel by train. However, the time when that information becomes available is also important because the time frame for buying train tickets is limited. This feature would allow the analysis of future events and thus be able to foresee and deal with potential planning issues many cycles ahead, a critical capability in many application domains.

This richer model of commitments is backward compatible with existing models. By incorporating uncertainty, this model offers a better interface between coordination and agent domain problem solving. Coordination mechanisms built around commitments can be separated from the domain. And, using commitments as multiagent planning contexts also allow the use of different scheduling approaches, as long as they are extended to take into account the multiagent planning context.

Our work so far has been focused on cooperative agents rather than self-interested agents. In cooperative multiagent systems, the agents' goal is to increase the expected total group utility. It needs to be pointed out that all three types of uncertainty mentioned here exist the same way for commitments in self-interested agents. However, for totally selfish agents, normally they would not exchange their marginal gains or losses, thus negotiation must be based on some other utility exchange model. It is interesting to note that self-interested behavior is not always desirable in agent societies, so a balanced model of individual and social utility may be used [23]. Also, many researchers are now looking at *deliberative agents* [4], where agents' social stance could be situation-specific. Again, in these cases an agent's utility model would also consider the society in which they act.

With the introduction of uncertainties in our model of commitments, 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 plans 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.

Acknowledgments

Effort sponsored by the Defense Advanced Research Projects Agency (DARPA) and Air Force Research Laboratory Air Force Materiel Command, USAF, under agreement number F30602-99-2-0525 and by the National Science Foundation under Grant number IIS-9812755. 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.

References

[1] C. Boutilier. Sequential optimality and coordination in multiagent systems. In *Proceedings of the Sixteenth International Joint Conferences on Artificial Intelligence (IJCAI-99)*, July 1999.

[2] S. Cammarata, D. McArthur, and R. Steeb. Strategies of cooperation in distributed problem solving. In *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, 1983.

[3] C. Castelfranchi. Commitments: from individual intentions to groups and organizations. In *AI and theories of groups & organizations: Conceptual and Empirical Research. Michael Prietula, editor. AAAI Workshop Working Notes.*, 1993.

[4] Cristiano Castelfranchi, Frank Dignum, Catholijn M. Jonker, and Jan Treur. Deliberate normative agents: Principles and architecture. In N.R. Jennings and Y. Lespérance, editors, *Intelligent Agents VI — Proceedings of the Sixth International Workshop on Agent Theories, Architectures, and Languages (ATAL-99)*, Lecture Notes in Artificial Intelligence. Springer-Verlag, Berlin, 2000.

[5] Philip R. Cohen and Hector J. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42(3):213–261, 1990.

[6] S. E. Conry, K. Kuwabara, V. R. Lesser, and R. A Meyer. Multistage negotiation for distributed constraint satisfaction. *IEEE Transactions on Systems, Man, and Cybernetics*, 21, 1991.

[7] R. Conte, R. Falcone, and G. Sartor. Agents and norms: How to fill the gap? In R. Conte, R. Falcone, and G. Sartor, editors, *Agents and norms, special issue, AI and Law*, volume 7, 1999.

[8] D. D. Corkill. Hierarchical planning in a distributed environment. In *Proceedings of the Sixth International Conference on Artificial Intelligence*, 1979.

[9] Keith Decker and Victor Lesser. Designing a family of coordination algorithms. In *Proceedings of the First International Conference on Multi-Agent Systems*, San Francisco, June 1995.

[10] Keith S. Decker and Victor R. Lesser. Generalizing the partial global planning algorithm. *International Journal of Intelligent and Cooperative Information Systems*, 1992.

[11] Keith S. Decker and Victor R. Lesser. Quantitative modeling of complex computational task environments. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, pages 214–217, 1993.

[12] M. desJardins and M. J. Wolverton. Coordinating a distributed planning system. *AI Magazine*, 20(4), 1999.

[13] F. Dignum. Autonomous agents and social norms. In *Proceedings of ICMAS'96 Workshop on Norms, Obligations, and Conventions*, 1996. Revised version in AI and Law, vol. 7, 1999.

[14] E. H. Durfee, V. R. Lesser, and D. D. Corkill. Trends in cooperative distributed problem solving. *IEEE Trans. Knowledge. Data Eng.*, 1989.

[15] E. H. Durfee and T. A. Montgomery. Coordination as distributed search in a hierarchical behavior space. *IEEE Transactions on Systems, Man, and Cybernetics*, 21, 1991.

[16] Edmund H. Durfee and Victor R. Lesser. Using partial global plan to coordinate distributed problem solvers. In *Proceedings of the TenthInternational Conference on Artificial Intelligence*, 1987.

[17] Edmund H. Durfee and Victor R. Lesser. Predictability versus responsiveness: Coordinating problem solvers in dynamic domains. In *Proceedings of the Seventh National Conference on Artificial Intelligence*, pages 66–71, 1988.

[18] P. Faratin, C. Sierra, and N. R. Jennings. Negotiation decision functions for autonomous agents. *Int. Journal of Robotics and Autonomous Systems*, 1998.

[19] L Gasser. Dai approaches to coordination. In N.M. Avouris and L. Gasser, editors, *Distributed Artificial Intelligence: Theory and Praxis*. Kluwer Academic Publishers, 1992.

[20] M. Georgeff. Communication and interaction in multi-agent planning. In A. H. Bond and L. Gasser, editors, *Readings in Distributed Artificial Intelligence*. Morgan Kaufmann, San Mateo, CA, 1988.

[21] C. V. Goldman and S. Zilberstein. Optimizing information exchange in cooperative multi-agent systems. In *Proceedings of the Second International Conference on Autonomous Agents and Multi-agent Systems*, Melbourne, Australia, 2003.

[22] B. J. Grosz and S. Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2):269–357, 1996.

[23] Lisa Hogg and Nick Jennings. Variable sociability in agent-based decision making. In N.R. Jennings and Y. Lespérance, editors, *Intelligent Agents VI — Proceedings of the Sixth International Workshop on Agent Theories, Architectures, and Languages (ATAL-99)*, Lecture Notes in Artificial Intelligence. Springer-Verlag, Berlin, 2000.

[24] N. R. Jennings. Commitments and conventions: The foundation of coordination in multi-agent systems. *The Knowledge Engineering Review*, 1993.

[25] N. R. Jennings. Coordination techniques for distributed artificial intelligence. In G.M.P. O'Hare and N.R. Jennings, editors, *Foundations of Distributed Artificial Intelligence*. John Wiley, 1996.

[26] V. Lesser and D. Corkill. The distributed vehicle monitoring testbed. *AI Magazine*, pages 15–33, 1983.

[27] V. R. Lesser. A retrospective view of fa/c distributed problem solving. *IEEE Transactions on Systems, Man, and Cybernetics*, 21, 1991.

[28] Thomas W. Malone and Kevin Crowston. Toward an interdisciplinary theory of coordination. Technical Report 120, Center for Coordination Science, MIT Sloan School of Management, 1991.

[29] 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, 1992.

[30] N. Onder and M. Pollack. Contingency selection in plan generation. In *Proceedings of the Fourth European Conference on Planning (ECP'97)*, 1997.

[31] D. Pynadath and M. Tambe. The communicative multiagent team decision problem: Analyzing teamwork theories and models. *JAIR*, 16:389–423, 2002.

[32] E. D. Sacerdoti. *A Structure for Plans and Behavior*. Elsevier, 1977.

[33] T. Sandholm. An implementation of the contract net protocol based on marginal cost calculations. In *Eleventh National Conference on Artificial Intelligence (AAAI-93)*, pages 256–262, Washington D.C., 1993.

[34] T. Sandholm and V. Lesser. Advantages of a leveled commitment contracting protocol. In *Thirteenth National Conference on Artificial Intelligence (AAAI-96)*, pages 126–133, 1996.

[35] Y. Shoham and M. Tennnenholtz. On social laws for artificial agent societies: off-line design. *Artificial Intelligence*, 73:231–252, 1995.

[36] Katia Sycara. Using option pricing to value commitment flexibility in multi-agent systems. Technical Report CMU-CS-TR-97-169, Carnegie Mellon University, 1997.

[37] M. Tambe. Towards flexible teamwork. *Journal of Artificial Intelligence Research*, 7(83-124), 1997.

[38] Frank von Martial. *Coordinating Plans of Autonomous Agents*. Springer-Verlag, 1992.

[39] Thomas Wagner. *Toward Quantified Control for Organizationally Situated Agents*. PhD thesis, University of Massachusetts at Amherst, 2000.

[40] Thomas Wagner, Alan Garvey, and Victor Lesser. Complex gola criteria and its application in design-to-criteria scheduling. In *Proceedings of thte Fourteenth National Conference on Artificial Intelligence*, 1997.

[41] Thomas Wagner and Victor Lesser. Towards ubiquitous satisficing agent control. In *1998 AAAI Symposium on Satisficing Models*, 1998.

[42] Thomas Wagner, Anita Raja, and Victor Lesser. Modeling uncertainty and its implications to design-to-criteria scheduling. Technical Report TR-98-51, UMASS Department of Computer Science, 1998.

[43] Thomas A. Wagner, Alan J. Garvey, and V. R. Lesser. Criteria directed task scheduling. *Journal for Approximate Reasoning: Special Scheduling Issue*, 19:91–118, 1998.

[44] D. E. Wilkins. *Practical Planning: Extending the Classical AI Planning Paradigm.* Morgan Kaufmann, 1988.

[45] Ping Xuan and Victor Lesser. Incorporating uncertainty in agent commitments. In *Intelligent Agents VI: Agents, Theories, Architectures and Languages (ATAL), Proceedings of The Sixth International Workshop on Agent Theories, Architectures, and Languages (ATAL-99), Lecture Notes in Artificial Intelligence 1757.* Springer-Verlag, 1999.

[46] Ping Xuan, Victor Lesser, and Shlomo Zilberstein. Communication decisions in multi-agent cooperation: Model and experiments. In *Proceedings of the Fifth International Conference on Autonomous Agent (AGENTS 01)*, pages 616–623, Montreal, Canada, 2001.

[47] Xiaoqin Zhang and Victor Lesser. Multi-linked negotiation in multi-agent system. In *Proceedings of the First International Joint Conference on Autonomous Agents And MultiAgent Systems (AAMAS 2002)*, 2002.