A Multi-Leveled Negotiation Framework *

Xiaoqin Zhang
Computer and Information Science Department
University of Massachusetts at Dartmouth
x2zhang@umassd.edu

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

Tom Wagner
Automated Reasoning Group
Honeywell Laboratories
Tom.Wagner@honeywell.com

Categories and Subject Descriptors:[I.2.11]: Distributed Artificial Intelligence, Intelligent agents

General Terms: Standardization, Performance

Keywords: layered negotiation, uncertainty

1. INTRODUCTION

Usually negotiation is structured as a single level process. Such negotiation can require a complicated reasoning process – particularly when the agent has multiple tasks and the tasks may be achieved in different ways, include sequencing constraints, and consume internal or external (shared) resources. Uncertainty in task execution may further complicate the negotiation process as behavior deviates from the expected. The deviation can cause re-negotiation over commitments or the adjustment of local activities so as to still meet the commitments. In this work negotiation is viewed as one of an agent's many interleaved activities – including scheduling, execution, and other negotiations. This view plus the complexities of negotiation mentioned earlier has led us to construct a two-level negotiation framework that makes the complexity inherent in this view more tractable. In this two-level negotiation framework, the negotiation process is performed at different abstraction levels to reduce the complexity of the search. An agent thus reasons about and negotiates over more important issues at the upper level, and then refines the rough commitments at the lower level in order to optimize its local plan and accommodate additional constraints and uncertainties. The focus of this work is on the decision-making process of negotiation, rather than the negotiation protocol or the language.

*This material is based upon work supported by the National Science Foundation under Grant No.IIS-9812755 and the Air Force Research Laboratory/IFTD and the Defense Advanced Research Projects Agency under Contract F30602-99-2-0525. 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, Air Force Research Laboratory/IFTD, National Science Foundation, or the U.S. Government.

AAMAS'03, July 14–18, 2003, Melbourne, Australia.
ACM 1-58113-683-8/03/0007.

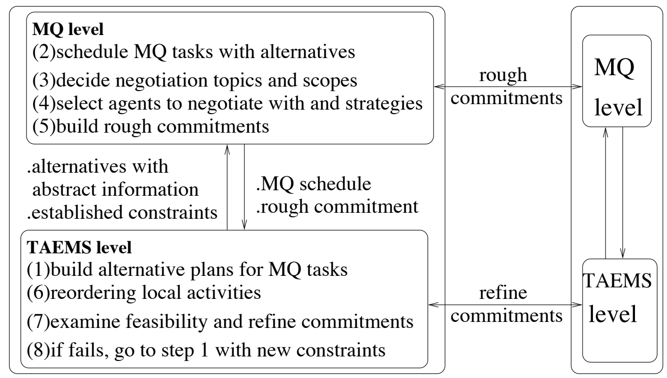

Figure 1: A two-level negotiation framework

There are two kinds of issues related to decision-making in negotiation: 1) Those issues, which have strong influence on local plan selection and involve utility transferred between agents, should be negotiated first at the upper level and rough commitments should be constructed for them. 2) Those issues, which have less influence on local plan selection and involve reasoning about the detailed structure of the low level activities, do not have to be directly reasoned about on the upper level and do not need to be decided on the upper level.

2. OVERVIEW OF BASIC IDEAS

The MQ model [5] describes the agent's organizational knowledge about task utility but it does not support detailed modeling of tasks and their interactions, and lacks of representation of the uncertainty characteristics and resource requirements of tasks. These details are represented using the TÆMS [1] task modeling language. The proper integration of the MQ and TÆMS models and reasoning processes enables agents to reason about both organizational level task value and to handle detailed feasibility, analysis, and implementation of tasks. These are the two frameworks that underlie and support this two-level negotiation.

In the two-level framework, an agent has an MQ level view of its local activities, which is a set of potential MQ tasks, each associated with certain $MQPS$ (MQs produced by this task) and $MQCS$ ((MQs consumed by this task)), which can be mapped into the agent's utility given the agent's current MQ state. For each MQ task T, there is a TÆMS task structure that describes the detailed activities for this task. Different plans to accomplish the MQ task T can be generated from the TÆMS task group TG by the DTC scheduler, and each plan has different quality, duration and cost characteristics that affect the $MQPS$ and $MQCS$ of the task T. This is the first step [step 1] shown in Figure 1, which describes the two-level negotiation framework.

The extended MQ scheduler generates a partial order schedule that indicates what tasks the agent should attempt to execute, what plans are used to execute these tasks, and the execution ordering. This schedule represents the agent's best choice about what activities it should do to maximize its local utility increase [step 2]. Based on these schedules, the agent can reason about the utility of a specific commitment. Negotiation on the MQ level is a multi-dimensional negotiation that includes the amount of the transferred MQ, the temporal constraints of the commitment and the quality constraints of the commitment [step3]. Also, the agent can select which agents to negotiate with and the appropriate negotiation strategy according to organizational relationships and the negotiation issues [step 4]. The MQ level negotiation builds rough (partial-specified) commitments for those issues that should or could be reasoned about the MQ level [step 5].

After building a local MQ schedule and rough commitments on the MQ level, the agent reorders its local activities on the TÆMS level [step 6]. In this reordering process, the agent optimizes its local schedule by taking advantage of the interrelationships among low-level tasks/methods. Also the agent verifies the feasibility of its local schedule given rough commitments from the MQ level and those additional constraints from the TÆMS level [step 7]. Negotiation on the TÆMS level involves refining those rough commitments as needed. If the agent can find a feasible local schedule by reordering and renegotiation on the TÆMS level, it can execute its local schedule and perform all of its commitments. If unexpected events cause conflict in the execution process, the agent needs to check if the conflict can be solved by refining any commitments. Otherwise, if the conflict can't be resolved given all current constraints, the agent needs to discard some commitments (decommits), establish other commitments on already scheduled local activities and go back to the MQ level to reschedule, and possibly result in constructing new commitments [step 8].

3. REWARD MODELS

Agents build rough commitments as a result of the MQ level negotiation. Future refinement as result of the lower level negotiation is possible given the range specified by the rough commitment. The refinement will affect the flexibility of the commitment and hence affect the value/cost of the commitment. Thus agents need to negotiate over the reward model which specifies how the refinement is related to the value of the transferred MQ. Because the flexibility is related to the value/cost of the commitment, the agents need to come to an agreement on how the latter refinement is related to the value of the transferred MQ. There are two possible models:

1. *Pre-paid flexibility* model. The contractee agent E pays $v1$ of MQ_i for the contractor agent R to perform task NL during any time period (not shorter than d) within $[t1, t2]$ as agent E requests. This agreement provides agent E with the freedom to further refine this commitment, and agent R agrees to accommodate any request from agent E within the pre-defined range. No matter what request agent E will make, or even if agent E does not make any further requests, agent R will receive $v1$ of MQ_i as decided in the rough commitment.
2. *Dynamic flexibility* model. The contractee agent E pays $v2$ of MQ_i for the contractor agent R to perform task NL within the range of $[t1, t2]$. If agent E requests a restriction on this range to $[t1 + x, t2 - y], (t2 - y) - (t1 + x) \geq d$, and if agent R could accept this request, agent E will pay $((x + y) * \beta + 1) * v2$ of MQ_i to agent B. Agent B would decide to accept this additional refinement request or not, according to its current problem-solving context. If agent R does not accept this request, it is still obliged to perform NL during $[t1, t2]$ and in turn is guaranteed to get $v2$ of MQ_i as the rough commitment defines.

These two models provide different degrees of freedom for the agents. The agents can choose a model according to the constraints and uncertainties of their local activities during the negotiation process.

4. SUMMARY AND RELATED WORK

This paper explores an alternative approach to negotiation in which an agent reasons about and negotiates over more important issues at the upper level (MQ level), and then refines the rough commitments at the lower level in order to optimize its local schedule and accommodate additional constraints and uncertainties of execution behavior. Multi-layered agent architectures have been proposed by other researchers, e.g. the InterRTaP [2] architecture based on BDI agent model , which is different from the utility-driven, quantitatively-reasoning agent control model in our work. DECAF [3] has also suggested a layered architecture based on separation along functional lines. However, these architectures have not addressed organizational concerns in the agent's goal selection process, as we do through the MQ framework, and none of them is focused on studying of the layered negotiation as our work does. Negotiation has been involved in both the goal selection and the action planning phases in multi-agent mission coordination[4]. These ideas are similar to the rough commitment built on the higher level in our work, though the negotiation issues in these work are much simpler and the systems are cooperative systems. hence no need to reason about local utility.

5. REFERENCES

[1] K. S. Decker and V. R. Lesser. Quantitative modeling of complex environments. *International Journal of Intelligent Systems in Accounting, Finance, and Management*, 2(4):215–234, Dec. 1993.

[2] K. Fischer, J. Muller, and M. Pisschel. Unifying control in a layered agent architecture. In *IJCAI95, Agent Theory, Architecture and Language Workshop 95*, 1995.

[3] K. D. John Graham. Towards distributed, environment centered agent framework. In *Intelligent Agents IV, Agent Theories, Architectures, and Languages*. Springer-Verlag, 2000.

[4] D. J. Musliner and K. D. Krebsbach. Multi-agent mission coordination via negotiation. In *Working Notes of the AAAI Fall Symposium on Negotiation Methods for Autonomous Cooperative Systems*, 2001.

[5] T. Wagner and V. Lesser. Evolving real-time local agent control for large-scale mas. In J. Meyer and M. Tambe, editors, *Intelligent Agents VIII (Proceedings of ATAL-01)*, Lecture Notes in Artificial Intelligence. Springer-Verlag, Berlin, 2002.