

Integrative Negotiation Among Agents Situated in Organizations

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*

Abstract—This paper addresses the problem of negotiation in a complex organizational context. An integrative negotiation mechanism is introduced, which enables agents to dynamically select a negotiation attitude based on the degree of external directedness. Experimental work explores the question of whether it always improves the organization’s social welfare to have an agent be completely externally-directed when negotiating and making choices. Results show that there are situations in which it is better for the organization if agents are *partially externally-directed* in their negotiations with other agents rather than *completely externally-directed*. The paper discusses the driving factors behind this unexpected result.

Keywords: integrative negotiation, motivation, group and organizational dynamics

I. INTRODUCTION

In Multi-Agent systems (MAS), agents negotiate over task allocation, resource allocation and conflict resolution problems. Until now almost all related work on negotiation can be categorized as falling into one of two general classes: negotiation in market-like systems and negotiation in distributed problem solving systems. In market-like systems, agents are self-interested and negotiate to maximize their own local utility [1], [2], [3], [4], [5], [6]; in distributed problem solving systems, agents negotiate to find a solution that increases their joint utility [7], [8], [9]. This latter approach is based on the assumption that full cooperation, at the local agent level, will lead to an overall increase in the social welfare of the system.

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 not withstanding 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.

Little work has been done to study negotiation between these two extreme cases.

When an agent is negotiating with other agents over task performance and/or resource consumption, it must explicitly reason about the *value* of performing/not-performing the tasks or allocating/not-allocating the resource. The way in which the value is computed differs depending on how the agent chooses to evaluate its negotiations with the other agents. We label an agent *completely self-directed* if it does not take into consideration how much utility the other agent can potentially gain when the local agent makes a commitment to complete the requested task as a result of negotiation. In other words, the agent is completely driven by the values that it locally attaches to task performance or resource consumption. In contrast, an agent is *completely externally-directed* if it sees the other agent’s gain as its own in a 1:1 fashion. Note that these labels identify *how* value or utility is computed and do not relate to the agent’s overall objectives. We distinguish the notion of agents being self-interested or cooperative from the notion of an agent being self-directed or externally-directed. We call an agent *self-interested* if its organizational goal is to maximize only its local utility and an agent is *cooperative* if it is intent on maximizing the overall social utility. Whereas the “direction” terminology defines how value is computed by the agent, the self-interested/cooperative terminology specifies the agent’s overall goal. With respect to negotiation, the degree of an agent’s self-interestedness/cooperativeness defines its meta-goal in terms of its overall relationship to the agent society, while the degree of self-directness/externally-directness defines the local mechanism used by an agent to help achieve its meta-goal. For instance, if the agent is cooperative and externally directed, it will work to maximize social welfare and will base its computations, during negotiation, on the values communicated to it by other agents. However, an agent who is cooperative and self-directed will also work to maximize social welfare, but based its own calculation/prediction of

social welfare. It does not consider how much the other agent would potentially gain as a result of a special commitment, because the other agent's potential gain is not thought by this agent as a reliable factor that indicates of social welfare, this agent rather leaves the information out of its consideration.

We feel that as the sophistication of multi-agent systems increases, MAS will be neither simple market systems where each agent is purely self-interested, seeking to maximize its local utility, nor distributed problem solving systems where all agents are completely cooperative working to maximize their joint utility. This will occur for the following reasons. First, agents from different separate organizational entities will come together to dynamically form virtual organizations/teams for solving specific problems that are relevant to each of their organizational entities [10]. How these agents work in their teams will often be dependent on the existence of both long-term and short-term relationships that are based on the goals of their underlying organizational entities. Secondly, even for agents from organizations with meta-goals that indicate self-interestedness, it might be beneficial for them to be partially externally-directed when they are in the situations where they will have repeated transactions with other agents from other organizational entities. Additionally, even agents working solely with agents of their own organizational entities will take varying negotiation attitudes in the spectrum of completely externally-directed to completely self-directed in order for the organization to best achieve its overall goal. The latter perspective is based on a *bounded-rational* argument: **it is not possible from a computational or communicational perspective for an agent to be fully cooperative**, because the agent needs to take into account the current and expected change in the utilities of all agents in the organization and the state of achievement of all organizational goals to be fully cooperative. Thus, it may be best for the organization to have agents being partially externally-directed in their local negotiation with other agents rather than being completely externally-directed in order to deal more effectively with the uncertainty of not having a more informed view of the state of the entire agent organization. We feel a similar argument can be made for self-interested agents. It may not always be advantageous for them to take the negotiation attitude of completely self-directed. Rather, in some context, the more external-directed attitude will lead to an increase in their own local utility.

Note that this work pertains to deliberate agents situated in an agent society where there are organizational relationships among agents. The agents can make choices about with whom to collaborate, how to negotiate, what to charge for services, etc. Further, the negotiation attitude will be dependent on the relationships among the negotiating parties and the particular negotiation issue, and the state of achievement of relevant organizational goals. In the experimental work reported in this paper, we are also assuming that agents are not acting in a hostile manner nor gaming the situation based on the meta-level information transferred among agents. However, we feel that by adding some additional mechanisms that allow an agent to adjust the character of the meta-level information that is exchanged, hostile/gaming agents can be handled within the

basic framework laid out in this paper.

Let's consider the supply chain example in Figure 1. There are different organizational relationships among agents. For instance, there is an agent (agent_IBM_2) who produces hard drives, belonging to the IBM Company. It provides hard drives for three different agents, with the following organizational relationships to it:

- 1) Agent_IBM_2 provides hard drives for the other agent (agent_IBM_1), which also belongs to IBM but assembles PCs.
- 2) Agent_IBM_2 provides hard drives to an NEC agent (agent_NEC), and as the transactions between them become more frequent and regular, they form a virtual organization based on the recent transactions.
- 3) Agent_IBM_2 occasionally provides hard drives for a distributor center (agent_DIS) based on a simple market-like mechanism.

When agent_IBM_2 negotiates with these three agents, it should use different negotiation attitudes that reflects the different relationships. For instance, when it negotiates with agent_IBM_1, it may need to be more externally-directed than it is towards the other two agents if its most important meta-goal is to increase the utility of IBM. However, even for the good of IBM's benefit, it may not be the best choice for agent_IBM_2 always to be completely externally-directed towards agent_IBM_1. Sometimes it may bring IBM more profit for agent_IBM_2 to provide hard drives to agent_DIS rather than to agent_IBM_1, if agent_IBM_1 is not certain whether it really needs the hard drive.

When agent_IBM_2 negotiates with agent_NEC, it may need to be more externally-directed than it is towards agent_DIS given the virtual organization it has formed with agent_NEC. The appropriate level of local cooperation depends on how important the utility increase of this virtual organization is to agent_IBM_2, how the goal to increase the utility of this virtual organization relates to its other goals, and how certain the information provided by agent_DIS compares to the information received from other sources. Also, as we noticed before, the formation of this virtual organization is dynamic; it may also disappear sometime later as the environment changes, so agent_IBM_2 should adapt its negotiation attitude dynamically too.

From the above examples we find it necessary to have a mechanism that supports agents choosing from among many different negotiation attitudes in the spectrum from completely self-directed to completely externally-directed, and easily switching from one attitude to another. The choice of negotiation attitude should depend on the agent's organizational goals, the current environmental circumstance, which agent it is negotiating with, and what issue is under negotiation. There should also be no requirement of a centralized controller that coordinates the agent's behavior.

So far, there has been no such negotiation mechanism which provides the above capabilities for agents (see related work in Section VI). In this paper, we introduce a negotiation mechanism which enables agents to construct negotiation attitudes in the spectrum from completely self-directed to completely externally-directed in a uniform reasoning framework called

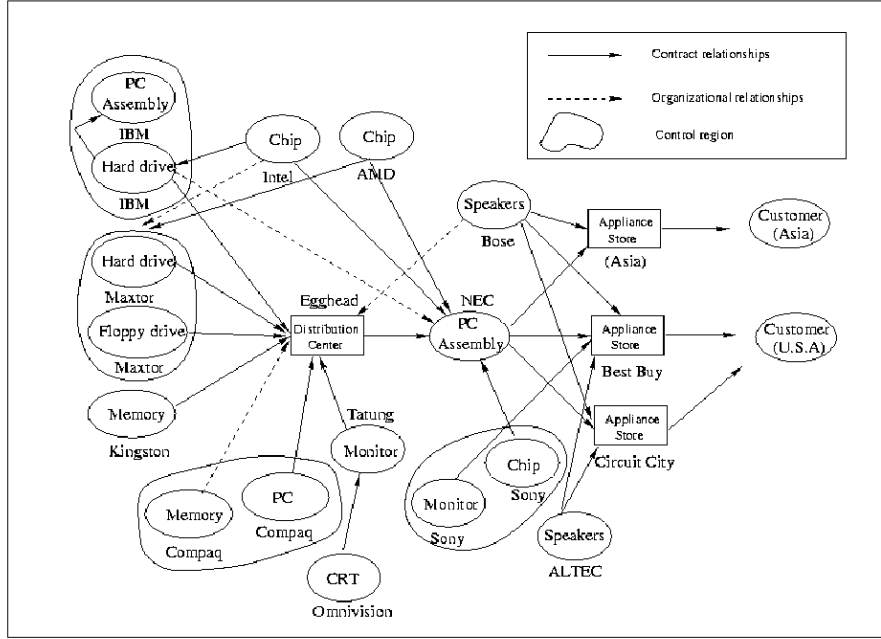


Fig. 1. Supply chain example

the Motivational Quantities (MQ) framework [11]. The MQ framework provides the agent with an appropriate utility model for quantitatively reasoning about how specific task allocation decision relates to satisfying its organizational goals. In the remainder of the paper, the MQ framework is reviewed in Section II. Section III describes the integrative negotiation mechanism. Section IV uses examples to explain the ideas more fully. Section V presents experimental results that explore how different negotiation attitudes affect the agent's performance and the social welfare of the overall system. Section VI discusses related work and Section VII concludes and identifies further work.

II. MQ FRAMEWORKS

The MQ framework [11] is an agent control framework that provides the agent with the ability to reason about which tasks should be performed and when to perform them. The reasoning is based on the agent's organizational concerns. The basic assumption is that agents are complex, with multiple goals related to the multiple roles they play in the agent society. The progress towards one goal cannot substitute for the progress towards another goal. Motivational Quantities (MQs) are used to represent the progress towards organizational goals quantitatively. Each agent has a set of MQs which it is interested in and wants to accumulate. Each MQ_i in this set represents the progress toward one of the agent's organizational goals. Each MQ_i is associated with a preference function (utility curve), U_{f_i} , that describes the agent's preference for a particular quantity of the MQ_i . The agent's overall utility is a function of the different utilities associated with the MQs it tracks: $U_{agent} = \gamma(U_i, U_j, U_k, \dots)$. The structure of function γ represents the agent's preference and emphasis on different organizational goals. The MQ framework thus provides an approach to compare the agent's different motivational factors through a multi-attribute function. Not all agents have the same

MQ set. If two agents need to construct a commitment through coordination or negotiation, and use MQ as an exchange medium, they need to have at least one MQ in common, or be willing to form one dynamically. Different agents may have different preferences for the same MQ.

MQs are consumed and produced by performing MQ tasks. The agent's overall goal is to select tasks to perform in order to maximize its local utility through collecting different MQs. This does not mean that the agent has to be "self-interested"; it only means that the agent selects its actions to contribute to its multiple organization goals. If "to help agent B" is one of the goals of agent A, then agent A will act in a cooperative manner with respect to agent B. If two or more agents have a goal in common and hence have the same MQ in common, they act as a group or a team working collaboratively toward this goal. MQ tasks are abstractions of the primitive actions that an agent may perform. The agent compares and selects tasks that are associated with different organizational goals. Each MQ task T_i has the following characteristics:

- Earliest start time (est), est_i . The performance of T_i before this time does not generate valid results.
- Deadline, dl_i . The accomplishment of T_i after this time does not generate valid results.
- MQ task T_i needs some process time to be accomplished, denoted as d_i .
- MQ task T_i produces certain quantities of one or more MQs, denoted as $MQPS$ (MQ production set). The production of MQs reflects the progress made in accomplishing the organizational goal associated with this specific MQ.
- MQ task T_i consumes certain quantities of one or more types of MQs, denoted as $MQCS$ (MQ consumption set). The consumption of MQs represents resources consumed by performing this task, or favors owed to other agents for subcontracting work.




Fig. 2. Motivational Quantities and Utilities

The MQ scheduler schedules current potential MQ tasks, and produces a schedule of a set of MQ tasks, specifying their start times, and finish times. The scheduler takes the following factors into consideration: the $MQPS$, $MQCS$, duration d_i , the earliest start time est_i and the deadline dl_i of each MQ task, and the agent's current accumulation of MQs . Notice that MQ is always being evaluated in the context of agent's current MQ accumulation state. For example, Figure 2 shows a single utility curve for a single MQ_i . The first one unit MQ_i brings the agent Q_1 units of utility U_i . After the agent has collected 2 units of MQ_i , the additional one unit of MQ_i brings the agent additional $(Q_3 - Q_2)$ units of utility U_i . $(Q_3 - Q_2)$ is not necessarily equal to Q_1 , they are calculated based on the utility curve associated with MQ_i .

The MQ framework provides the comparison of tasks that need to be performed for different reasons: for different organizational goals, for other agents to gain some financial benefit or favors in return, for cooperation with other agents, etc. It also supports different utility functions that relate the execution of tasks to the importance of organizational goals. The MQ framework is related to the work on joint intentions [12] and joint goals [13]. In this work, the agent reasons logically about the existence of joint goals (based on information exchanged and its local knowledge database) and then decides which activities to perform and how it should interact with the other agents under joint goals. However, this work differs from the MQ framework in the following way. The joint goal work does not address how the agent chooses from multiple candidate goals, or how the agent decides which activities to perform at a given time. Instead, it focuses on finding the existence of the joint goals. In contrast, the MQ work focuses on deciding which goals (or tasks) to perform, when to perform them and how to perform them from a quantitative perspective rather than from a logical one.

In summary, the motivational qualities (MQ) framework provides an agent with the capability to reason about different goals in an open, dynamic and large-scale MAS, hence the agent can evaluate a negotiation issue from an organizational perspective.

III. INTEGRATIVE NEGOTIATION

In a complex agent society, an agent will need to work with other agents from a variety of different organizational positions. For example, an agent from its own group, an agent

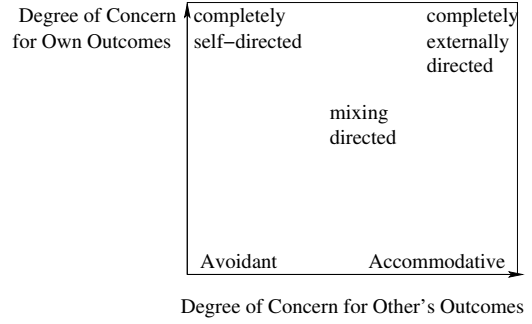


Fig. 3. The dual concern model

who has a higher position and thus more authority, an agent from a cooperative company, or an agent from a competing company and so forth. The agent's attitude toward negotiation is not just simply either self-directed or externally-directed, the agent needs to qualitatively reason about each negotiation session, and so it can choose an appropriate negotiation attitude.

Figure 3 describes this dual concern model. When the agent only attaches importance to its own outcome, its attitude toward negotiation is completely self-directed; when an agent attaches the same degree of importance to its own outcome as it does to the outcomes of the other agent, its attitude is completely externally-directed; when the agent attaches more importance to the outcomes of other agents and no importance to its own outcome, its attitude is accommodative; if the agent attaches no importance to any outcomes, its attitude is avoidant (the negotiation is not worth its time and effort). From this model, we find that there are potentially many options between the two extremes of completely self-directed and completely externally-directed. These other options depend on the importance the agent attaches to the increase of its own utility relative to the importance it attaches to the increase of the other agents' utility.

Let's use task allocation as an example of negotiation where for each task t that agent A wants agent B to complete, certain MQs are transferred from agent A to agent B if agent B agrees to complete the task. The conceptual model here is that agent B is motivated by the potential increase in its MQs to perform tasks for agent A (note that this does not convert the MQs to currency as not all agents may be interested in said MQs). We will start with a simple, abstract example. In this model, when agent B commits to accomplishing task t , based on a contract that is mutually agreed upon by the two agents (formed either dynamically or pre-defined), it is then obligated to perform the task, otherwise it may incur a penalty. When B successfully accomplishes t , the agreed upon amount of the MQ will be transferred from agent A to agent B. Note that agent B must actually decide whether or not it is interested in performing t . This evaluation is done via the MQ framework and the associated MQ scheduler. The evaluation uses agent B's preference for the MQ in question to determine the relative value of performing t for agent A compared to other candidate tasks agent B may have. This evaluation process, in turn, determines agent B's attitude toward the negotiation of task

t .

In terms of specifics, there are two types of MQs that could be transferred with the successful accomplishment of task t : *goal_related MQ* and *relational MQ*. These classes are conceptual and used to clearly differentiate motivations for task performance from attitudes toward negotiation issues – in reality, they are both simply MQs. *Goal_related MQs* are associated with an agent’s organizational goals, generally increases in MQ volume, and hence have positive benefits to the agent’s utility. Note that the agent’s designer determines which kinds of MQs the agent tracks (and is interested in), defines the agent’s preference for each via the utility functions discussed earlier, and determines how these relate to the agent’s organizational goals. When dealing with *goal_related MQs*, the agent collects MQs for its own utility increase. In this sense, agent B’s performance of task t is motivated by “self-interested” reasons if payment is via a *goal_related MQ*. For example, task t has 3 units of MQ_x transferred with it, and for agent B, the utility curve of MQ_x is: $u(x) = 2x$, that means, the utility of agent B will increase by 6 units by collecting 3 units of MQ_x through performing task t . Agent B decides whether to accept task t by reasoning about its value relative to the cost of the resources it will expend in the performance of t and the opportunities it will forgo by taking this task t . In this case, as the task doesn’t consume any MQs, the resource expenditure is time or in terms of opportunity cost. Because this reasoning process pertains to *goal_related MQs*, it is “self-directed” for the agent’s only concerns is its own utility increase.

Consider a modified case. Suppose that by having task t accomplished by agent B, agent A’s own utility increases by 20 units. If agent B takes this fact into consideration when it makes its decision about task t , agent B is externally-directed with agent A because agent B is also concerned about agent A’s outcome (in addition to its own). If we want agent B to consider A’s utility, we need to introduce another MQ designed to model B’s (revised) preference for A to have a utility increase also. To reflect the B’s attitude toward A’s outcome, we introduce a *relational MQ*, the preference for which represents how externally-directed agent B is with agent A concerning task t . Let $MQ_{ba/t}$ be the relational MQ transferred from agent A to agent B when agent B performs task t for agent A. Since $MQ_{ba/t}$ is a *relational MQ*, its only purpose is to measure the attitude of agent B towards agent A concerning task t . the utility of agent B toward problem solving, we will not consider the utility produced by any relational MQs such as $MQ_{ba/t}$. Likewise with agent A. When agent A transfers $MQ_{ba/t}$ to agent B, we will not tabulate the negative change in utility of agent A because the change in utility is not related to problem solving progress but is instead related to the transfer of a relational MQ. The reason for this approach is that in this paper our performance metric is social welfare as it is conventionally used, which is in terms of progress toward joint goals. From this view, the utility produced by a relational MQ can be seen as *virtual utility*. Though $MQ_{ba/t}$ produces virtual utility, is important because it carries the information of how important task t is




Fig. 4. different mapping functions of $MQ_{ba/t}$

for agent A¹ and makes it possible for agent to consider agent A’s outcome when it makes its own decisions. Actually, how $MQ_{ba/t}$ is mapped into agent B’s (virtual) utility, meaning utility that is not included in the social welfare computation² depends on how externally-directed agent B is with agent A.

Suppose that 20 units $MQ_{ba/t}$ are transferred with task t , representing the utility agent A gained by having agent B perform task t , Figure 4 shows four different functions for mapping $MQ_{ba/t}$ to agent B’s virtual utility.

Function a, b and c are linear functions: $U_a(MQ_{ba/t}) = k * MQ_{ba/t}$.

If $k = 1$ (a), $U_b(MQ_{ba/t}) = MQ_{ba/t} = U_a(t)$ ($U_a(t)$ denotes the utility agent A gained by transferring t), then agent B is completely externally-directed to agent A;

If $k > 1$ (b), $U_b(MQ_{ba/t}) > MQ_{ba/t} = U_a(t)$, then agent B is accommodative to agent A³;

If $k < 1$ (c), $U_b(MQ_{ba/t}) < MQ_{ba/t} = U_a(t)$, then agent B is partially externally-directed with agent A;

If $k = 0$, $U_b(MQ_{ba/t}) = 0$, then agent B is completely self-directed with respect to agent A. In this case, if agent A wants agent B to do task t , it needs to transfer another kind of MQ (the *goal_related MQ*) to agent B, agent B and agent A can negotiate about what type of *goal_related MQ* to transfer and how much of it should be transferred, regarding how and when agent B could accomplish task t . In the following examples and the experimental work, we assume that the type and amount of the transferred *goal_related MQs* are fixed and agents do not negotiate about them, so we can focus on demonstrating how the *relational MQ* works.

¹It is assumed that agents are honest and don’t lie about the importance of task t . We recognize that this assumption may not hold in all applications. It is worth noting, however, that it is actually difficult to lie effectively in the MQ framework because the agents do not necessarily know each other’s mapping function for relational MQs. Consider Figure 4. If agent A is interacting with agent B and agent A does not know which mapping function (a, b, c, d) that agent B is using, it will be difficult for agent A to know the impact that its local choices will have on agent B’s response.

²In remainder of the paper, we may omit the word “virtual” before utility, but we know that this *relational MQ* only maps into virtual utility that is not real utility. In the experimental work, neither the agent’s utility nor the social welfare includes the virtual utility from *relational MQ*.

³This function can be used to represent authority relationship between agents. When k is set to a very large number, agent A actually has authority over agent B - the task from agent A has high priority in agent B’s agenda. Another way to express the authority relationship in MQ framework is to use the *goal_related MQ*. A similar preference utility function like this one associated with a *goal_related MQ* can represent the authority. However, the difference is that there is no “real” utility transferred between agents in the first approach.

The mapping function could also be a nonlinear function (d) that describes a more complicated attitude of agent B to agent A, i.e., agent B being completely externally-directed with agent A until certain organizational goal is met indicated by the level of MQ , and then becoming self-directed. An agent can adjust the utility mapping function to reflect its relationship with another agent, which could be its administrator, colleague, friend, client or competitor. By adjusting some parameters in the mapping function, more subtle relationships could be managed. The agent could differentiate a friendly colleague from an unfriendly colleague, also it could draw distinctions between a best friend and an ordinary friend. The structure of the function reflects that for how long and to what extent the agent would like to be externally-directed.

Different from the *goal-related MQs*, which are built by the agent's designer and whose utility curves are not changing, the utility curves of the *relational MQs* can be adjusted by the agent dynamically to reflect its dynamic relationships with other agents. Additionally, the agent's attitude towards another agent could be "issue-specific"; given an agent could play multiple roles, there could be different issues negotiated between agents, and the agents should select different attitude according to what issue is negotiated. For example, for the colleague's request to contribute to a shared professional job and for the same colleague's request for a ride, even though both requests come from the same agent, the agent's attitude could be different.

By introducing this agent-oriented, issue-specific *relational MQ* into negotiation, the agent's attitude toward another agent concerning a specific issue can be represented as the utility curve associated with the *relational MQ*. This mechanism is called an *integrative negotiation* mechanism, which supports the agent's choosing a negotiation attitude of any type from completely self-directed to completely externally-directed. The agent's attitude towards a negotiation issue is affected by the utility mapping function of the transferred MQ with this issue. In the MQ framework, the MQ scheduler enables the agent to optimize its schedule and maximize its local utility. While the framework directly supports the concept of relational MQs and being motivated to cooperate on that basis, the use of MQ transference in this paper extends the MQ framework to interconnect the local scheduling problems of two or more agents in a dynamic fashion (based on the current context). Prior to this work, no meaningful work had been done in MQ transference or the implications of it.

How can an agent choose its attitude toward other agents in such a complex organization context? We are not planning to present a solution to this question in this paper, but we feel that the agent should dynamically adjust its attitude by analyzing the other party, the issue in negotiation and its current problem-solving status. In Section IV, we show that for a simple scenario the optimal attitude can be formally specified and for that scenario we can learn through local observation what are the best attitude. The following information should be considered in this decision making process: "Who is the other agent?" "How is its organizational goals related to mine?" "What is its objective?" "What is its relationship to me?" and so forth. Some of this information can be learned from

experience [14]. In [15], we presented a formalized analytical model and showed that the best negotiation attitude can be driven through the calculation based on this model and the available information of the environmental context.

IV. THE SCENARIO

In this section, we introduce a simple example of an agent society and show how the integrative negotiation mechanism works using the MQ framework. There are three agents in this society as shown in Figure 5.

- 1) The Computer-Producer Agent (*c*): receives *Purchase_Computer* tasks from an outside agent (which is not considered in this example). Figure 5 shows that to accomplish a *Purchase_Computer* task, the Computer-Producer Agent needs to generate an external request for hardware (*Get_Hardware* task), and also needs to ship the computer (*Deliver_Computer*) through a transport agent.
- 2) The Hardware-Producer Agent (*h*): receives *Get_Hardware* tasks from the Computer-Producer Agent, it also receives *Purchase_Parts* tasks from an outside agent.
- 3) The Transport Agent (*t*): receives *Deliver_Computer* tasks from the Computer-Producer Agent, it also receives *Deliver_Product* tasks from an outside agent.

In this example, every agent collects the same type of *goal-related MQ*: MQ_s . The utility curve for MQ_s is: $utility(x) = x$ and every agent uses this same function. Each task that the agent receives includes following information:

- Its earliest start time (est), the performance of the task before this time does not generate valid results.
- Its deadline (dl): the latest finish time for the task.
- Its reward (r): if the task is finished by the deadline, the agent will get reward r (which is r units of MQ_s).
- The early finish reward rate (e): if the agent can finish the task by time ft as it promised in the contract, it will receive an additional early finish reward. The reward sum is adjustable so that if the agent finishes even sooner, additional rewards are given. The relationship is expressed mathematically as: $\max(e*r*(dl-ft), r)$. The maximum additional reward is r so that the total reward possible for task performance, including both basic reward and additional reward, is $2*r$.

As Figure 6 shows, the Hardware-Producer Agent receives *Purchase_Parts* task from an outside agent with x units of MQ_s , where x is a random number varying from 2 to 10. The Computer-Producer Agent has long-term contract relationship with the Hardware-Producer Agent and the Transport Agent: its *Get_Hardware* task always goes to the Hardware-Producer Agent with a fixed reward of 3 units of MQ_s , and its *Deliver_Computer* task always goes to the Transport Agent with a fixed reward of 3 units of MQ_s . Every *Purchase_Computer* task comes to the Computer-Producer Agent with a reward of 20 units of MQ_s if it is finished by its deadline (the reward can be higher if the task is finished earlier, see the following example). The Computer-Producer Agent would have its local utility increased by 14 units after paying the reward to the

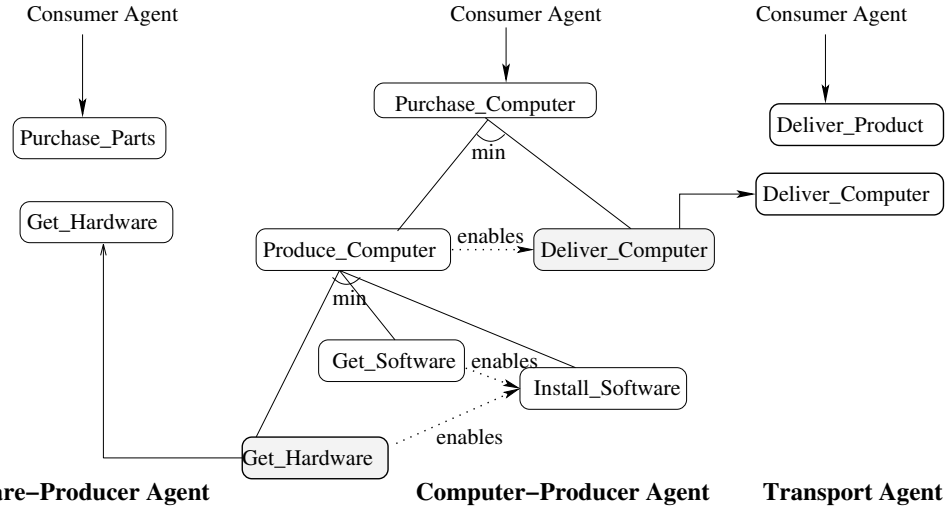


Fig. 5. Agent Society




Fig. 6. Tasks with Different MQs

Hardware-Producer Agent and the Transport Agent). Assume the tasks *Get_Hardware* and *Deliver_Computer* have the same importance, the accomplishment of each task would result in 7 units utility increase for the Computer-Producer Agent. This information is reflected by the 7 units of $MQ_{hc/t}$ transferred with task *Get_Hardware* and 7 units of $MQ_{tc/t}$ transferred with task *Deliver_Computer*. $MQ_{hc/t}$ ⁴ is a *relational MQ* introduced to reflect the relationship of the Hardware-Producer Agent with the Computer-Producer Agent concerning task t . The transferred $MQ_{hc/t}$ with the task represents the utility increase of the Computer-Producer Agent by having this task accomplished. How it is mapped into the Hardware-Producer Agent's virtual utility depends on the Hardware-Producer Agent's attitude towards the utility increase of the Computer-Producer Agent regarding task *Get_Hardware*. If the *Purchase_Computer* task could be finished earlier than its deadline, the Computer-Producer Agent could get more than 20 units reward. The extra utility increase could be estimated and reflected by more than 7 units transferred $MQ_{hc/t}$ or $MQ_{tc/t}$ to the other two agents. Suppose the Computer-Producer Agent receives the following task:

- Task name: *Purchase_Computer_A*
- Earliest-start-time: 10
- Deadline: 70
- Reward: 20 units $MQ_{\$}$
- Early finish reward rate: $e=0.01$

Through the reasoning of the MQ scheduler, the Computer-Producer Agent decides to accept it and finish it by time

40 (it leaves 4 units slack time) to earn extra early reward 6 $((70 - 40) * 0.01 * 20)$ units $MQ_{\$}$. Its local utility increases by 20 $(20+6)$, after paying the sub-contractor agents) units after the accomplishment of this task. Hence the following two task requests: *Get_Hardware_A* and *Deliver_Computer_A* are sent to the Hardware-Producer Agent and the Transport Agent respectively:

task name	Get_Hardware_A	Deliver_Computer_A
<i>est</i>	10	30
<i>deadline</i>	20	40
<i>reward</i>	3 units $MQ_{\$}$ 10 units $MQ_{hc/t}$	3 units $MQ_{\$}$ 10 units $MQ_{tc/t}$

In this example, we look at three different attitudes, for how the Hardware-Producer Agent negotiates with the Computer-Producer Agent over the task *Get_Hardware*. The different attitudes are specified in terms of a linear function: $U_{ha}(MQ_{hc/t}) = k * MQ_{hc/t}$.

- 1) $k=1$, the Hardware-Producer Agent is completely externally-directed.
- 2) $k=0.5$, the Hardware-Producer Agent is partially externally-directed.
- 3) $k=0$, the Hardware-Producer Agent is completely self-directed.

Now we can look at how these different attitudes affect the negotiation process of the Hardware-Producer Agent. Suppose there are two other tasks *Purchase_Parts_A* and *Purchase_Parts_B* received by the Hardware-Producer Agent besides task *Get_Hardware_A*, this results in the three tasks being sent to the MQ Scheduler (suppose the initial MQ set

⁴Similarly, $MQ_{tc/t}$ is a *relational MQ* that reflects the relationship of the Transport Agent with the Computer-Producer Agent concerning task t . Detailed discussion about it is omitted here.

is empty):

task name	<i>est</i>	deadline	process time	MQPS
Get_Hardware_A	10	20	10	$[MQ_s, 3]$ $[MQ_{hc/t}, 10]$
Purchase_Parts_A	10	30	10	$[MQ_s, 4]$
Purchase_Parts_B	10	20	10	$[MQ_s, 9]$

The decisions made by the agent depend on the attitude taken:

- If the Hardware-Producer Agent is completely externally-directed to the Computer-Producer Agent ($k = 1$), the best MQ schedule produced is:
[10, 20] Get_Hardware_A [20, 30] Purchase_Parts_A
the Hardware-Producer Agent will have 7 units utility increase after the accomplishment of this schedule.
- If the Hardware-Producer Agent is completely self-directed to the Computer-Producer Agent ($k = 0$), the best MQ schedule produced is:
[10, 20] Purchase_Parts_B [20, 30] Purchase_Parts_A
the Hardware-Producer Agent will have 13 units utility increase after the accomplishment of this schedule.
- If the Hardware-Producer Agent is partially externally-directed to the Computer-Producer Agent ($k = 0.5$), the best MQ schedule produced is the same as above. However, if the task *Purchase_Parts_B* comes with 6 units MQ_s instead of 9 units, then the best MQ schedule produced is:
[10, 20] Get_Hardware_A [20, 30] Purchase_Parts_A
the Hardware-Producer Agent will have 7 units utility increase after the accomplishment of this schedule. A similar reasoning process also applies to the Transport Agent.

The above example shows how an agent reacts in a negotiation process depends on its attitude towards the other agent regarding this issue, and also is affected by the other tasks on its agenda. The more externally-directed an agent is, the more it will sacrifice its own utility for the other agent's utility increase. This integrative negotiation mechanism enables the agent to manage and reason about different negotiation attitudes it could have with another agent regarding a certain issue.

V. EXPERIMENTAL RESULTS

The example in Section IV shows that an agent needs to sacrifice some of its own utility gain to be externally-directed with another agent. One important question is: can externally-directed agents improve *social welfare*?⁵ Another important question is: when should an agent be externally-directed and how externally-directed it should be? To explore these questions, the following experimental work⁶ was done

⁵Social welfare refers to the sum of the utilities of all the agent in the society, i.e. the sum of the utilities of the three agents: the Computer-Producer Agent, the Hardware-Producer Agent, and the Transport Agent. Social welfare is collected in the experiments just for us to compare different policies. It is never being used by individual agents in their local decision making processes, because this information is not available for them at all. For individual agent, the only available information besides its local information is the relational MQ from the other agent with whom it is negotiating.

⁶The experiments are performed in the MASS simulator environment[16], and the agents were built using the JAF agent framework[17]

based on the scenario described in Section IV. The Hardware-Producer Agent has a choice of three different attitudes toward the Computer-Producer Agent: completely externally-directed (C) ($k=1.0$), partially externally-directed (H) ($k=0.5$), and completely self-directed (S) ($k=0$), the Transport Agent has the same three choices, so there are 9 combinations: SS (both agents are completely self-directed), SC (the Hardware-Producer Agent is completely self-directed while the Transport Agent is completely externally-directed), SH (the Hardware-Producer Agent is completely self-directed while the Transport Agent is partially externally-directed), HS, HC, HH, CS, CH, CC. The data is generated by running 48 groups of experiments; in each group the agents work on the same incoming task set under the nine different situations. The tasks in each set for each group experiment are randomly generated with different rewards and deadlines within certain ranges.

Table I shows the comparison of each agent's utility and the social welfare under these different situations. The percentage numbers are the normalized utility numbers based on the utility gained when agent is completely self-directed. When both the Hardware-Producer Agent and the Transport Agent are completely externally-directed with respect to the Computer-Producer Agent (CC), the society gains the most social welfare. Even when both agents are only partially externally-directed (HH), the social welfare is still very good. However, when one agent is completely externally-directed and the other agent is completely self-directed (CS, SC), the social welfare does not improve much compared to the completely self-directed (SS) case⁷. The reason for this lack of significant improvement is that, in this example, to accomplish task *Purchase_Computer* requires that both the task *Get_Hardware* and the task *Deliver_Computer* needs are successfully completed. When one agent is completely externally-directed, it sacrifices its own utility, but task *Purchase_Computer* may still fail because the other agent does not cooperate on the subtask, thus the utility of the Computer-Producer Agent does not increase as expected, and the global utility does not improve. This happens when the completion of a task is spread over more than two agents – thus the information from the Computer-Producer Agent about its utility increase is only an estimation because it depends not only on task *Get_Hardware* for the Hardware-Producer Agent, but also relies on task *Deliver_Computer* for the Transport Agent. In this situation, if the Hardware-Producer Agent has no knowledge about the attitude of the Transport Agent (and what other tasks it will be receiving including their worth and frequency), then it may not be a good idea to be completely externally-directed towards the Computer-Producer Agent.

Table II shows the results of statistical significance (t-test) testing about the social welfare under the different cooperative situations. For example, the first line in Table II shows that with the 0.01 Alpha-level, we can reject the hypothesis H_0 that the difference between the social welfare when both agents are completely externally-directed and the social welfare when

⁷Results from t-test have shown that the difference of the social welfare between CC and SS, also between HH and SS, are statistically significant.

	Utility of Computer Producer Agent	Percentage	Utility of Hardware Producer Agent	Percentage	Utility of Transport Agent	Percentage	Social Welfare	Percentage
SS	218	1.000	575	1.000	856	1.000	1649	1.000
CC	842	4.08	415	0.72	766	0.90	2022	1.23
HH	587	2.84	493	0.86	806	0.94	1886	1.14
SC	301	1.41	587	1.02	798	0.93	1686	1.02
CS	469	2.24	364	0.63	839	0.98	1672	1.01
HS	390	1.87	467	0.81	845	0.99	1702	1.03
SH	292	1.36	585	1.02	815	0.95	1692	1.03
HC	632	3.06	500	0.87	772	0.90	1905	1.16
CH	761	3.68	405	0.70	802	0.94	1967	1.19

TABLE I
COMPARISON OF PERFORMANCE

Difference of Social Welfare	Number to Compare	Ho	Ha	Result	Alpha	p
CC - SS	330	=330	> 330	Reject Ho	0.01	0.008
HH - SS	180	=180	>180	Reject Ho	0.01	0.0008
SC - SS	0	=0	>0	Fail to reject Ho	0.01	0.0179
CS - SS	0	=0	>0	Fail to reject Ho	0.01	0.0965

TABLE II
RESULTS FROM STATISTICAL TESTS

	Utility of Hardware Producer Agent	Percentage	Social Welfare	Percentage
Completely Self-Directed	583	1.0	1679	1
Completely Externally-Directed	395	0.68	1887	1.13
Partially Externally-Directed	487	0.83	1831	1.09

TABLE III
UTILITY OF HARDWARE-PRODUCER AGENT AND SOCIAL WELFARE

both agents are completely self-directed is equal to 330⁸, compared to the hypothesis H_a that the difference between the social welfare when both agents are completely externally-directed and the social welfare when both agents are completely self-directed is greater than 330.

Table III shows the expected utilities of the Hardware-Producer Agent and the expected social welfare under the three possible situations: when the Hardware-Producer Agent is completely self-directed, completely externally-directed and partially externally-directed. When the Hardware-Producer Agent chooses one attitude, the Transport Agent may adopt one of the three different attitudes. For example, when the Hardware-Producer Agent chooses to be completely self-directed, the global situation could be SS, SC, or SH. The utility numbers in the table are the expected values of the utilities under these three different situations. Table III tells us that when a cooperative task involves more than two agents and when the other agents' attitudes are unknown, being completely externally-directed means sacrificing its own utility significantly and thus, at least in this scenario, is not a good idea.

We recognized that the above conclusion might relate to the parameters of the experiments. Table IV shows these parameters. For example, the third row of the table shows that the Hardware-Producer Agent receives two *Purchase_Parts* task every 15 time clicks, the reward for each *Purchase_Parts*

Agent	Task	Reward	Frequency every 15 time clicks	d
<i>c</i>	<i>Purchase_Computer</i>	20	1	16
<i>h</i>	<i>Get_Hardware</i>	3	1	7
<i>h</i>	<i>Purchase_Parts</i>	[2,10]	2	6
<i>t</i>	<i>Deliver_Computer</i>	3	1	6
<i>t</i>	<i>Deliver_Product</i>	[2,10]	2	7

TABLE IV
EXPERIMENT PARAMETERS

falls in the range of [2, 10], and the duration of the task is 6. Every *Purchase_Computer* task comes to the Computer-Producer Agent with a reward of 20 units of MQ_s , if it is finished by its deadline, the Computer-Producer Agent would have its local utility increased by 14 units (With the deduction of the 6 units of MQ_s transferred to the Hardware-Producer Agent and the Transport Agent). This information is sent to the Hardware-Producer Agent (and also the Transport Agent) by attaching 7 (14 divided by 2 agents) units of *relational MQ* ($MQ_{hc/t}$ for the Hardware-Producer Agent) with the task-announcing proposal. This information is taken into consideration by the MQ scheduler when the Hardware-Producer Agent makes its decision on this proposal. However, this information is not necessarily accurate because it is based on the assumption that the task *Produce_Computer* will be finished on time. Whether this assumption is appropriate

⁸330 is 20% of social welfare under the SS situation (1649), and 180 is 11% of social welfare under the SS situation.

depends on whether the Hardware-Producer Agent and the Transport Agent would accept the subcontracts and fulfill them on time. The uncertainty associated with this information comes from the uncertainty of the other contractor agent’s (the Transport Agent) decision, where the other contractor agent’s decision is based on the following issues:

- 1) The agent’s attitude toward the Computer-Producer Agent (how externally-directed it is); the more externally-directed it is, the more likely this subcontract will be accepted.
- 2) The outside offers the agent receives: how good they are, how frequent they are and how they affect the subcontract task. If the outside offer is not higher compared to the reward from the subcontract, or if they are not very frequent, or if they do not conflict with the subcontract task, the subcontract will be more likely to be accepted.

Because these issues are unknown by the Computer-Producer Agent and the Hardware-Producer Agent, the uncertainty associated with the information about the local utility increase cannot be resolved. This is why we make the statement at the beginning of this paper: *it is not possible from a computational or communicational perspective for an agent to be fully cooperative, because the agent needs to have complete global information to be fully cooperative*. Thus, it may be best for the organization to have agents being partially externally-directed in their local negotiation with other agents rather than being completely externally-directed in order to deal more effectively with the uncertainty of not having a more informed view of the state of the entire agent organization⁹. Generally, an agent should put appropriate weight on external information provided by other agents in an uncertain environment in order to deal with distraction. When there is more uncertainty related to the external information, an agent should be more self-directed, and it should be more externally-directed if the external information is more certain.

Additional experiments have been done using different parameters. Table V shows the social welfare under different conditions. When the rewards of outside offers fall into the range of [11, 19], for the best social welfare, both agents should be completely self-directed.

However, if there is no uncertainty or less uncertainty, it may be the best for the agent to be completely externally-directed or more externally-directed toward the group task in order to increase the social welfare. This does not mean the agent has to grant every subcontract of the group task, the decision also depends on the outside offer. If the outside offer is significantly better than the subcontract even with taking into consideration of the contractee agent’s utility increase, and if the contractor agent can only choose one between the subcontract of the group task and the outsider offer, the contractor agent will take the outside offer and drop the subcontract even if it

⁹This issue of distraction in a distributed interpretation system [18], [19] is caused by anonymous evaluation of the validity of locally generated hypothesis. The problem caused by subsequent integration into the reasoning of another agent is very similar to the issues described in the experiments. The solution to this problem in a distributed interpretation system is to modify local reasoning process to only partially explore the information received from another agent. This approach is similar in character to the idea suggested in this paper.

Reward from outside offer	SS	CC	HH
[2, 10]	1.0	1.23	1.14
[11, 19]	1.0	0.93	0.98

TABLE V
SOCIAL WELFARE USING DIFFERENT PARAMETERS

is completely externally-directed. And in fact, this choice increases the social welfare.

Based on the above experimental results, we feel there are at least two different ways for agents to choose the appropriate level of cooperation. One approach is that the agent who has more global view/knowledge can inform other agents about how likely the estimated utility increase will be true, and the other agents can adjust their cooperation levels based on the reliability of this information. Another approach is that the individual agent can learn from the past experience to adjust the level of cooperation.

VI. RELATED WORK

Glass and Grosz [20] developed a measure of social consciousness called “brownie points” (BP). The agent earns BP each time it chooses not to default a group task and loses BP when it does default for a better outside offer. The default of a group task may cause the agent to receive group tasks with less value in the future, hence reducing its long term utility. The agent counts BP as part of it overall utility beside the monetary utility. A parameter $BPweight$ can be adjusted to create agents with varying levels of social consciousness. This relates to our utility mapping function associated with the *relational MQ* which can be adjusted to reflect the agent’s different attitude in negotiation. However, the *relational MQ* is agent-oriented and issue specific, so the agent can model different attitudes towards each agent and negotiation issue. Additionally, the mapping function can be a nonlinear function and describe a more complicated attitude. Their work assumes there is a central mechanism controlling the assignment of group tasks according to agent’s rank (agent’s previous default behavior), which is not always appropriated for an open agent environment. Instead, in our assumption, agents are all independent and there is no central control in the society.

Axelrod [21] has shown stable cooperative behavior can arise when self-interested agents adopt a reciprocating attitude toward each other. The agent cooperates with another agent who has cooperated with it in previous interactions. The idea of the reciprocity is related to our work if the *relational MQ* is used bi-directionally between agents, agent A collect some *relational MQ* from agent B and in the future the accumulated *relational MQ* could be used to ask agent B do some work for it, in this way, the *relational MQ* actually works as a quantitative measure of reciprocity. Sen developed a probabilistic reciprocity mechanism [14] in which the agent K chooses to help agent J with certain probability p and p is calculated based on the extra cost of this cooperation behavior and how much effort it owes agent J because agent J has helped it before. There are two parameters in the formula

for calculating p which can be adjusted so that the agent can choose a specific cooperation level. However, this work assumes that cooperation always leads to aggregate gains for the group, and it was based on a known cost function - that is, they know how much extra it will cost them to do X for another agent. Neither of these two assumptions are necessary in our work. Also our work deals with more complex and realistic domains where tasks have real-time constraints and there are potentially complex interrelationships among tasks distributed across different agents.

Our experimental work has shown that even in a cooperative system it may not be the best for the social welfare to have agents be completely externally-directed. Similar result is also shown in [22], which uses a distributed constraint satisfaction model that is much different from the underlying model in this work. Vidal [23] has also studied the teaming and selflessness when using multi-agent search to solve task-oriented problems. His study also shows the fact that neither absolute selfishness nor absolute selflessness result in better allocations, and the fact that the formation of small teams usually leads to better allocations. This work explores a similar issue as in our work, however, it is in a relatively simplified domain and there is no complex interaction among agents. Other related work includes the cooperative negotiation work on task allocation [24], where the agents use the marginal utility gain and marginal utility cost to evaluate if it worth to accept a task contract in order to increase the global utility. However in this work, the agent acts as in a “completely-cooperative” mode and there is no choice on how cooperative it wants to be.

This paper is an extended version of [25]. Compared with the conference paper, this extended paper has the following improvements. In this paper, we introduce two new concepts “self-directed” and “externally-directed”, which are different from “self-interested” and “cooperative”. This paper provides a more complete description of the MQ framework. This paper also includes more experimental result. We performed additional experiments using different parameters, the results show that the best policy depends on the environmental context such as the outside offer, so it is important to have agents to dynamically choose the level of cooperation.

VII. CONCLUSION AND FUTURE WORK

We introduce an integrative negotiation mechanism that enables agents to interact over a spectrum of negotiation attitudes from completely self-directed to completely externally-directed in a uniform reasoning framework, namely the MQ framework. The agent can not only choose to be self-directed or externally-directed, but also can choose how externally-directed it wants to be. This provides the agent with the capability to dynamically adjust its negotiation attitude in a complex agent society. Introducing this mechanism in the agent framework also strengthens the capability of multi-agent systems to model human societies. Multi-agent systems are important tools for developing and analyzing models and theories of interactivity in human societies. There are many complicated organizational relationships in human society, and every person plays a number of different roles and

is involved in different organizations. A multi-agent system with this integrative negotiation mechanism is an ideal test-bed to model human society and to study negotiation and organization theories. Experimental work shows it may not be a good idea to always be completely externally-directed in a situation involving an unknown agent’s assistance; in that case, choosing to be partially externally-directed may be appropriate for both the individual agent and also for the society.

We recognize that the experimental results are scenario specific and they do not answer the question about how externally-directed an agent should be in a given situation. In [15], we presented an analytical model of the environment that enables the agent to predict the influence of its negotiation attitude on its own performance and also on the social welfare, hence to select the appropriate level of cooperation to balance its own utility achievement and the social welfare. We plan to develop learning techniques that enable an agent to learn from its previous interactions with other agents about how to adjust its negotiation attitude parameter.

REFERENCES

- [1] T. Sandholm and V. Lesser, “Advantages of a leveled commitment contracting protocol,” in *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, 1996, pp. 126–133.
- [2] —, “Issues in automated negotiation and electronic commerce: Extending the contract net framework,” in *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS95)*, 1995, pp. 328–335.
- [3] M. Andersson and T. Sandholm, “Leveled Commitment Contracting among Myopic Individually Rational Agents,” in *Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS98)*, Paris, France, 1998, pp. 26–33.
- [4] —, “Leveled Commitment Contracts with Myopic and Strategic Agents,” in *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, Madison, WI, July 1998, pp. 38–45.
- [5] —, “Time-Quality Tradeoffs in Reallocation Negotiation with Combinatorial Contract Types,” in *Proceedings of the Sixteenth National Conference on Artificial Intelligence*, Orlando, FL, 1999, pp. 3–10.
- [6] T. Sandholm and N. Vulkan, “Bargaining with deadlines,” in *National Conference on Artificial Intelligence (AAAI)*, 1999.
- [7] 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*, vol. 21, no. 6, Nov. 1992.
- [8] S. Lander and V. Lesser, “Understanding the role of negotiation in distributed search among heterogeneous agents,” in *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, 1993, pp. 438–444.
- [9] S. Sen and E. H. Durfee, “A formal study of distributed meeting scheduling,” *Group Decision and Negotiation*, vol. 7, pp. 265–289, 1998.
- [10] A. P. R. Eugenio Oliveira, *Book on European perspectives on AMEC*. Springer-Verlag, June 2000, ch. Agents advanced features for negotiation in Electronic Commerce and Virtual Organisations formation process.
- [11] T. Wagner and V. Lesser, “Evolving real-time local agent control for large-scale mas,” in *Intelligent Agents VIII (Proceedings of ATAL-01)*, ser. Lecture Notes in Artificial Intelligence, J. Meyer and M. Tambe, Eds. Springer-Verlag, Berlin, 2002.
- [12] P. Cohen and H. Levesque, “Intention is choice with commitment,” *Artificial Intelligence*, vol. 42, 1990.
- [13] N. R. Jennings and E. H., “Using joint responsibility to coordinate collaborative problem solving in dynamic environments,” in *Proceedings the Tenth National Conference on Artificial Intelligence*, 1992.
- [14] S. Sen, “Reciprocity: a foundational principle for promoting cooperative behavior among self-interested agents,” in *Proc. of the Second International Conference on Multiagent Systems*. Menlo Park, CA: AAAI Press, 1996, pp. 322–329.
- [15] J. Shen, X. Zhang, and V. Lesser, “Degree of Local Cooperation and its Implication on Global Utility,” *Proceedings of Third International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2004)*, July 2004. [Online]. Available: <http://mas.cs.umass.edu/paper/359>

- [16] B. Horling, R. Vincent, and V. Lesser, "Multi-agent system simulation framework," in *16th IMACS World Congress 2000 on Scientific Computation, Applied Mathematics and Simulation*. EPFL, August 2000.
- [17] B. Horling and V. Lesser, "A Reusable Component Architecture for Agent Construction," University of Massachusetts at Amherst, Computer Science Technical Report TR-98-30, May 1998.
- [18] V. R. Lesser and D. D. Corkill, "The distributed vehicle monitoring testbed," *AI Magazine*, vol. 4, no. 3, pp. 63–109, Fall 1983.
- [19] V. R. Lesser, "A retrospective view of FA/C distributed problem solving," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 21, no. 6, pp. 1347–1363, Nov. 1991.
- [20] A. Glass and B. Grosz, "Socially conscious decision-making," in *Proceedings of Agents 2000 Conference*, Barcelona, Spain, June 2000, pp. 217 – 224.
- [21] R. Axelrod, *The Evolution of Cooperation*. Basic Books, 1984.
- [22] H. Jung, M. Tambe, and S. Kulkarni, "Argumentation as distributed constraint satisfaction: applications and results," in *Proceedings of the Fifth International Conference on Autonomous Agents*, J. P. Müller, E. Andre, S. Sen, and C. Frasson, Eds. Montreal, Canada: ACM Press, May 2001, pp. 324–331.
- [23] J. M. Vidal, "The effects of cooperation on multiagent search in task-oriented domains," *Journal of Experimental and Theoretical Artificial Intelligence*, 2003, to appear.
- [24] X. Zhang, R. Podorozhny, and V. Lesser, "Cooperative, multistep negotiation over a multi-dimensional utility function multi-agent systems negotiation," in *Proceedings of the IASTED International Conference, Artificial Intelligence and Soft Computing (ASC 2000)*, 2000, pp. 136–142, <http://dis.cs.umass.edu/xqzhang/pub/tr00-02.ps>.
- [25] X. Zhang, V. Lesser, and T. Wagner, "Integrative negotiation in complex organizational agent systems," in *the Proceedings of the IEEE/WIC International Conference on Intelligent Agent Technology (IAT 2003)*, Halifax, Canada, Oct 13-16 2003, pp. 140 –146.