The Mediation Algorithm for Real Time Negotiation

Timothy Rauenbusch
Division of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138 USA
tim@eecs.harvard.edu

ABSTRACT

Mediation is presented as an algorithm to be used by a group of agents involved in a negotiation in which relevant information is vast and decentralized. It is an anytime algorithm that enables agents to incrementally reveal information in search of an optimal outcome. Allocation Improvement applies Mediation to the task allocation problem. Its performance is demonstrated on a negotiation problem involving multi-agent planning for intruder tracking.

Categories and Subject Descriptors

I.2.11 [**Computing Methodologies**]: Artificial Intelligence — *Multiagent Systems*

General Terms

Algorithms, Design, Experimentation

1. INTRODUCTION

The goal of the work reported in this paper is the development of a general-purpose negotiation method that can be used in multi-agent planning. The environments we consider are characterized by the following features: (1) relevant information is vast and decentralized and (2) negotiation time is limited since agents need to interleave planning with acting. A motivating application of the work is a group of agents working to form a plan to perform a group action as modeled by a SharedPlan [1].

In these applications, a group of agents G that has agreed to perform an action α. Before α can be performed, agents in the group must make several interrelated decisions, which include determining the responsibilities of each agent, the assignment of parameters relevant to performing α, and so on. Our analysis makes use of the notion of a *proposal*. Each proposal corresponds to one full specification of a way to perform α and encodes all of the decisions the agents need make in specifying how that action is to be performed. The *negotiation set* $\mathbf{P}$ denotes set of all possible proposals. The size of $\mathbf{P}$ is combinatorial in the number of decisions that must be made.

Each agent g has a *utility function*, u_g, that maps each element in $\mathbf{P}$ to a real value. There is a global objective function f that enumerates the desirability of elements of $\mathbf{P}$, given the utility function of each agent (e.g., social welfare). The goal of the agents is to find the best possible element of $\mathbf{P}$, according to this global objective function. In our model, we do not restrict the utility and objective functions that agents may have. Therefore, in real planning situations where $\mathbf{P}$ is large, it is not feasible for agents to compute and reveal their full utility function given their limited amount of time.

Mediation, presented below, provides a general method for agents to reveal their utility function incrementally over the course of a negotiation. As a result, it is able to find good, albeit possibly sub-optimal, solutions quickly.

Mediation is a parameterized algorithm that supports the use of domain-specific knowledge to enhance the performance of the search through an update procedure, if such knowledge is available. The Allocation Improvement algorithm is an example of the use of such knowledge in task allocation domains. It builds upon theoretical work by Sandholm [2] by providing a specific ordering of OCSM contracts for task reallocation. We provide initial experimental data to illustrate the effectiveness of Allocation Improvement in a particular task allocation domain.

We are currently working on providing update procedures specifically appropriate for agents involved in collaborative planning that enables agents to quickly find good plans for performing a group action. We are also working on extensions to the Mediation algorithm to support dynamic changes to the negotiation set during the negotiation.

2. MEDIATION ALGORITHM

The Mediation algorithm (Figure 1) implements an iterative hill climbing search through the negotiation space, storing the best proposal found so far as a variable labeled b (for *best proposal*). A member of G is chosen to act as mediator, and at each step this agent selects and communicates an element of $\mathbf{P}$, labeled c for *current proposal*, to the agents in the group. Each agent then responds with a message that is based on the proposal that was broadcast; msg_i denotes the message sent by agent i. The messages are combined to form a value, denoted $\textsc{value}(msg_1, msg_2, ..., msg_n)$. If that value is preferred to the value for the current b (based on the preference relation $\succ$), b is updated with the current proposal c.

The algorithm is anytime: it can be halted at any time and will return the best proposal found so far. The proposal stored in variable b is returned as the outcome when the procedure is terminated. Therefore, Mediation is applicable even if agents do not know in advance how much time they will have to negotiate.

With a choice of agent messages and VALUE function that correspond to agents' utility functions and objective function, Mediation implements a hill-climbing search in the space of proposals, with objective function f. We illustrate an appropriate choice of mes-

AAMAS'02, July 15-19, 2002, Bologna, Italy.

function MEDIATION **returns** an outcome
inputs: $\mathbf{P}$, G, UpdateProcedure
let $b \leftarrow \emptyset$, $b_{\text{val}} \leftarrow$ VALUE($\emptyset$)
loop
$c \leftarrow$ next value generated by UpdateProcedure
broadcast c to G
for each G_i in G
receive msg_i from G_i
$c_{\text{val}} \leftarrow \text{VALUE}(msg_1, msg_2, ..., msg_n)$
if ($c_{\text{val}} \succ b_{\text{val}}$) then
$b \leftarrow c$, $b_{\text{val}} \leftarrow c_{\text{val}}$
until (stop signal)
return b

Figure 1: Mediation Algorithm

Let p ←a random element of $\mathbf{P} \setminus \{\emptyset\}$; return p
for $i = 1 \ldots |T|$
for $t \leftarrow$ every set of tasks of size i
for a ←every possible assignment of agents in G to tasks in t
$q \leftarrow$ substitute a in p; return q
if $q_{\text{val}} \succ p_{\text{val}}$ in Mediation, then $p \leftarrow q$.

Figure 2: Allocation Improvement Update Procedure

sages and VALUE in the experimental domain below.

A parameter of Mediation is an update procedure that determines the element of $\mathbf{P}$ that is chosen for c at each step. Allocation Improvement is an example of an update procedure and is presented below.

3. ALLOCATION IMPROVEMENT

Allocation Improvement is an update procedure appropriate for agents involved in a negotiation over *task allocation*, in which $\mathbf{P}$ is of a particular structure. Specifically, each element of $\mathbf{P}$ corresponds to a different allocation of agents in G to a set of tasks T.

As shown in Figure 2, the first proposal p is chosen randomly from $\mathbf{P}$; it provides a context, from which subsequent proposals are generated. In subsequent iterations, the procedure returns proposals that result from making substitutions in p for i-tuples of tasks where i increases from 1 to $|T|$. Substitutions for each i-tuple of tasks is made sequentially with each permutation of agents in lexicographic order, while maintaining the allocations for the other tasks. p is always maintained to correspond with the best proposal seen so far (b from the mediation algorithm).

Sandholm [2] proves that the optimal allocation may only be found by allowing arbitrarily complex changes to the initial allocation. Allocation Improvement makes use of this result and the intuition that the primary determinant of proposal desirability is the assignment of each task to the agent that is best suited to perform it, while increased desirability associated with task interaction is a secondary concern. It specifies a search order over the set of all changes to the allocation, with the goal of finding good solutions quickly.

4. EXPERIMENTAL EVALUATION

The Mediation algorithm has been evaluated in a task allocation domain that is based on a real time, multi-sensor intruder tracking scenario. A group of sensor agents work together to perform a sequence of tracking tasks as an intruder moves through its region.

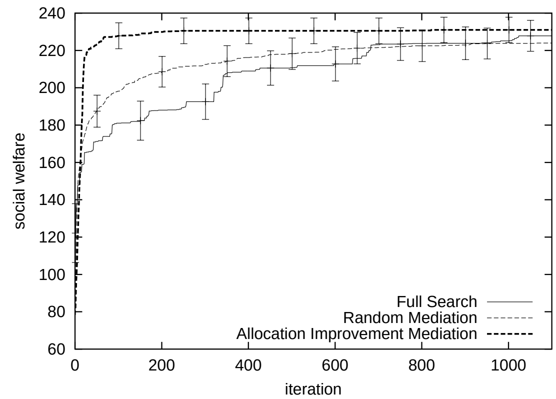

Figure 3: Comparison of Mediation algorithms in the 4-agent tracking domain

Benefits to the group are higher when a task is assigned to an agent that is as close as possible to the target. There is task interaction: there is a warm-up cost for an agent performing a task, but a set of consecutive tasks incur only one warm-up cost. Benefits and costs are quantified by a utility function known locally to each agent and the objective function is social welfare.

In msg_i, agent G_i reports the value associated with its quality and expected distance from the object, less the warm-up costs, for the set of tasks to which it has been allocated in the latest proposal. The VALUE function is chosen to be the sum of the values reported in each agent's message.

In the experiments, Allocation Improvement was compared to other instances of the Mediation algorithm using two different types of update rules. *Full Search* simply returns successive elements of $\mathbf{P}$ as they would be explored in a depth-first search. *Random Mediation* returns a random element of $\mathbf{P}$ at each iteration.

Each negotiation method was run on 100 problem instances with groups of 4 agents. Figure 3 shows the social welfare of proposal b after each iteration of the Mediation Algorithm, using the three different update rules, along with 95% confidence intervals. The results strengthen the hypothesis that Allocation Improvement is an effective strategy in this domain, particularly for agents in real time environments that may only be able to search through a small number of proposals before they require an outcome.

5. ACKNOWLEDGMENTS

The author thanks Barbara Grosz and Charles Ortiz for their guidance and support in developing the content and presentation of this work. The research described in this paper was partially supported by National Science Foundation grants IIS-9978343 and CDA-94-01024, and by the DARPA Autonomous Negotiating Teams Program, Contract F30602-99-C-0169.

6. REFERENCES

[1] Barbara Grosz and Sarit Kraus. The evolution of SharedPlans. In Michael Wooldridge and Anand Rao, editors, *Foundations and Theories of Rational Agency*, number 14 in Applied Logic Series, pages 227–262. Kluwer Academic Publishers, The Netherlands, 1999.

[2] Tuomas Sandholm. Contract types for satisficing task allocation: I theoretical results. In *AAAI 1998 Spring Symposium: Satisficing Models*, Stanford University, California, March 23-25 1998.