Learning to Select Negotiation Strategies in Multi-agent Meeting Scheduling

Elisabeth Crawford and Manuela Veloso*

Computer Science Department, Carnegie Mellon University,
Pittsburgh PA 15213, USA
{ehc, mmv}@cs.cmu.edu

Abstract. In this paper, we look at the Multi-Agent Meeting Scheduling problem where distributed agents negotiate meeting times on behalf of their users. While many negotiation approaches have been proposed for scheduling meetings, it is not well understood how agents can negotiate strategically in order to maximize their users' utility. To negotiate strategically, agents need to learn to pick good strategies for negotiating with other agents. We show how the *playbook* approach, introduced by [1] for team plan selection in small-size robot soccer, can be used to select strategies. Selecting strategies in this way gives some theoretical guarantees about regret. We also show experimental results demonstrating the effectiveness of the approach.

1 Introduction

Personalized software agents for meeting scheduling have the potential to reduce the daily cognitive load on computer users. Scheduling meetings can be a time consuming process requiring many email messages to be exchanged, and often existing meetings need to be moved to make room for new ones. Potentially, software agents can remove this burden entirely by communicating with each other to schedule meetings. Since user's have ownership of their own calendars, and private preferences about meeting scheduling, it makes sense to approach this problem in a distributed manner. Automated negotiation has been proposed as a method for multiple agents to reach agreement on meeting times. Negotiation approaches have many advantages over the open calendar approach taken by Microsoft Outlook (see [2] for a discussion).

Typically negotiation protocols feature a meeting initiator that proposes meeting times and collects the proposals of other participants. Consider, for instance, the following simplified protocol:

- while there is no intersection in proposals
  - the initiator proposes some times to the other agents
  - each agent proposes some times to the initiator

* Thanks to the reviewers and Michael Bowling for helpful comments and suggestions. This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. NBCHD030010. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Defense Advanced Research Projects Agency (DARPA), or the Department of Interior-National Business Center (DOI-NBC).

C. Bento, A. Cardoso, and G. Dias (Eds.): EPIA 2005, LNAI 3808, pp. 584–595, 2005.

In this context, a negotiation strategy is a set of rules for deciding what times to propose at each point in the process. The space of possible negotiation strategies is extremely large. Even if we restrict the space in some way, e.g. to strategies that offer a fixed number, x, of new times per negotiation round, there are still a huge number of options. In particular, there is a different strategy for each possible value of x, and then there are all the ways of combining these values of x with rules for deciding what particular times to offer. In developing software agents for meeting scheduling, we are faced with the problem of: (i) deciding which negotiation strategies agents should consider, and (ii) designing methods that agents can use to choose between these strategies when negotiating a particular meeting.

In order to most effectively satisfy user preferences, we would like our agents to adapt their behavior to each of the agents they negotiate with. There is a wide range of important ways in which agents can differ. For instance, agents can represent users of different importance and busyness, use very different negotiation strategies, and can have users with very different preferences. Clearly a strategy that works well for negotiating with one agent may work very poorly with another. Poor strategy choice can lead to meetings being scheduled at times the user does not like, or to the negotiation process taking a very long time. In general, we would like an agent to trade-off satisfying its user's preferences, with minimizing the length of the negotiations, in a way that maximizes its user's utility.

One method for deciding what strategy to use when negotiating with a particular agent is to use a model based approach that tries to construct a model of the agent and then based on this model select a strategy. There are a number of reasons why this approach would be difficult to use in practice. Firstly, obtaining an accurate enough model of another agent is a very difficult learning problem, since the only interaction agents have is through the exchange of times when they negotiate meetings. From this information, it is hard to make accurate conclusions about what times an agent prefers, how busy the agent is, what negotiation strategy it is employing etc. Secondly, to build a model of another agent, many training examples are required. It would be preferable if an agent was able to learn to negotiate, while actually negotiating.

In this paper, we show how an agent can learn online which strategies to use by observing its own rewards, as opposed to trying to model the other agents. Our approach is based on the idea of *plays* introduced by Bowling, Browning and Veloso [1]. Bowling *et al.* focus on the domain of robot soccer (small-size league) where they equip a team with a series of multi-agent plans called a *playbook*. The team plan to use at a given point in time is selected according to a no-regret learning algorithm. We show how we can apply these ideas to the problem of learning how to negotiate with different agents. Our experimental results demonstrate that this approach allows a learning agent to converge to sensible strategies for negotiation with different fixed strategy agents. We also show that an agent learning online using this approach can perform well in comparison to the best (in hindsight) fixed strategy.

2 Related Work

A variety of methods for reaching agreements on meeting times have been proposed in the last ten years, including negotiation based approaches, e.g. [3,4], Distributed