

Learning to Select Negotiation Strategies in Multi-Agent Meeting Scheduling

Elisabeth Crawford and Manuela Veloso

Computer Science Department, Carnegie Mellon University
5000 Forbes Ave
Pittsburgh, PA, 15213
{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 an agent needs to learn to pick good strategies for each agent. We show how the *playbook* approach introduced by (Bowling, Browning, & Veloso 2004) 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.

Introduction

Personalized software agents for meeting scheduling have the potential to reduce the daily cognitive load on computer users. Scheduling a meeting can be a time consuming process requiring many email messages to be composed and read, and often old meetings to be moved to make room for new meetings. Potentially, software agents can remove this burden entirely by communicating with each other to schedule meetings. Since users 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 (Crawford & Veloso 2004).

Typically negotiation protocols feature a meeting initiator that makes proposals for the meeting time and collects the proposals of all participants. Consider for instance the following simplified protocol:

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

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 each round. In developing software agents for meeting scheduling we are faced with the problem of deciding which negotiation strategies (taken from this huge space) agents should consider as well as designing methods that the agents can use to choose between these strategies when negotiating a particular meeting.

In order to most effectively satisfy user preferences when negotiating meetings we would like our agents to be able to adapt their behavior according to which agent they are negotiating with. There is a wide range of important ways in which agents can differ. For instance, agents can represent users of vastly different levels of importance and busyness, they can use very different negotiation strategies and can have users with very different preferences about meeting times. Clearly a strategy that works well for negotiating with one agent may work very poorly for 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 maximize their user's utility by scheduling meetings at their users' preferred times while minimizing the length of the negotiation process according to their user's preferences about this possible trade-off.

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. Learning how to negotiate is an intrinsically online problem and it is thus preferable for the learning to occur online. Furthermore this online aspect means there is a difficulty with statically mapping agent models to negotiation strategies. The best negotiation strategy to use can depend on changing aspects of the learning agent and the other agent, since as more meetings are added to the agents' calendars their behavior can change even if their underlying strategy does not. Online algorithms have the potential to adapt to these changing conditions.

In this paper we show how an agent can learn online which strategies to use when negotiating with different agents 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 (2004). 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 achieves a higher pay-off for its user than an agent using a fixed strategy or randomizing over all the strategies in the *playbook*.

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. (Jennings & Jackson 1995; Sen & Durfee 1998; Gonzalez & Santos), Distributed Constraint Reasoning (DCR) approaches (Pragnesh Jay Modi 2005) and market based approaches (Ephrati, Zlotkin, & Rosenschein 1994). In this section we describe work on the first two methods looking in particular at how user preferences are dealt with.

Sen and Durfee (1998) conducted a probabilistic and simulation based analysis of negotiation strategies. The basic framework they considered was:

1. Host announces meeting
2. Host offers some times
3. Agents send host some availability information
4. Repeat 2 and 3 until an intersection is found.

Similar protocols have been looked at by other researchers, for example, (Jennings & Jackson 1995), (Gonzalez & Santos) and (Garrido & Sycara 1995), while (Shintani, Ito, & Sycara 2000) looked at a more complex protocol. These negotiation approaches have handled user preferences for meeting times in quite different ways. Shintani *et al.* (Shintani, Ito, & Sycara 2000) propose a persuasion based approach to negotiation. The persuasion mechanism involves compromising agents adjusting their preferences so that their most preferred times are the persuading agent's

most preferred times. This method relies strongly on the cooperation of the agents in complying with this protocol. Sen *et al.* (Sen, Haynes, & Arora 1997) also look at user preferences and show how voting methods can be used (within a software agent) to evaluate different options for scheduling a meeting when users have multi-faceted and possibly competing preferences.

Garrido and Sycara (1995) and Jennings and Jackson (1995) take the approach of allowing agents to not only propose meeting times but also to quantify their preferences for proposals. The agent that is collecting the proposals then makes decisions about meeting times based on the reported utilities of all the meeting participants. This style of approach involves a lot of trust, since for the procedure to work well all the agents must report their preferences truthfully. While the approaches outlined are all concerned with user preferences they differ from the work described here in that we are interested in *how an agent can negotiate strategically in order to satisfy its user's preferences*.

Distributed Constraint Reasoning (DCR) approaches have also been applied to multi-agent meeting scheduling. For example Modi and Veloso (Pragnesh Jay Modi 2005) model the meeting scheduling problem according to the DCR paradigm and evaluate strategies for making *bumping decisions*. The way in which agent's decide when to *bump* (i.e., move an existing meeting to accommodate a new meeting) can have implications for the efficiency of the meeting scheduling process. A meeting that is bumped must be rescheduled and it is possible that in rescheduling a bumped meeting more meetings will be bumped. Intuitively, if the agents want the scheduling process to finish quickly, they should try to bump meetings that will be easy to reschedule. Similarly to the negotiation approaches the work on DCR has not focused on how agents can act strategically, rather the agents have been assumed to be cooperative.

Plays for Meeting Negotiation

In the context of small-size robot soccer (where an overhead camera and an off-board computer allow for coordinated team planning) Bowling *et al.* (2004) introduce the notion of a play as a team plan. Each play must assign a role to each of the robots e.g. one robot may be tasked with attempting to shoot, another with guarding the team's own goal and so forth. Each play also has an applicability condition that determines in what scenarios it applies and a termination condition that is used to determine when the play has been completed and a new play needs to be selected. For example an offensive play might be applicable whenever the ball is in the opponent's half of the field and be considered to have terminated when a goal is scored, or the applicability condition is violated.

The *playbook* captures all the plays that are available to the team. Bowling *et al.* provide a simple language that a human expert can use to add new plays. Typical playbooks consist of about 15 strategies. Playbooks are often compiled just prior to a match in order to reflect the most current knowledge the human experts have about the opponent team. During the course of a game, plays are weighted according to their level of success or failure and the play to

use at each decision point is selected based on these weights. The weights on plays are adapted in such away that regret about play selection goes to zero in the limit. We describe the process of weight adaptation in detail in the next section.

The meeting negotiation problem has a number of important features in common with small-size robot soccer. In both domains the space of available strategies is huge and it is not possible for agents to adapt online if they must consider the entire strategy space. Furthermore the environment in both domains is dynamic, the models of the ‘opponents’ are unknown and online learning is required for good performance. In this section we will discuss how we adapt the plays formulation to the problem of learning how to negotiate with different agents.

We can map the plays terminology from robot soccer to the meeting scheduling problem. The plays correspond to complete negotiation strategies, the opponent corresponds to the agent the learning agent is negotiating with, and the playbook is simply the set of negotiation strategies available to the learning agent. Unlike in robot soccer, in the meeting scheduling problem we are playing with multiple agents in succession or possibly playing with multiple agents at the same time and thus the learning agent must adapt strategy selection for each of the different agents it negotiates with simultaneously.

Each negotiation strategy has 4 elements, a rule for deciding applicability, a rule for deciding at each negotiation round what times to offer (if any) independent of the exact proposals received, a rule for deciding which times (if any) to offer based on the proposals received and a rule for deciding when to give up. Figure 1 shows an example strategy, Offer-k-b, that offers k new available times each round and after b rounds starts taking into account the proposals it has received. If necessary Offer-k-b will offer times that would require an already scheduled meeting to be bumped. Depending on the values of k and b this strategy can be very selfish and cause the negotiation to take a long time. As such, if the ‘opponent’ agent is very important the strategy is only applicable if the value of k is large and the value of b is small.

Each time a new meeting needs to be scheduled if the learning agent is an attendee it selects which strategy to use according to the adapted playbook for the initiator agent. If the learning agent is initiating the meeting it selects a possibly different negotiation strategy for communicating with each attendee according to the adapted playbook for that attendee. The learning agent considers the execution of a strategy to be complete when (i) the meeting it was selected to schedule has been added to the agent’s calendar and (ii) any meetings that the learning agent is involved in that have been bumped as a result of scheduling the new meeting have been rescheduled for new times. A strategy is also considered to have been completely executed if the learning agent has given up on scheduling the new meeting or on rescheduling a bumped meeting. Each time a strategy terminates the playbook weights are updated according to the success of the strategy.

In the next section we discuss how the learning agent adapts playbook weights for negotiating with each other

1. **APPLICABILITY:** if importance(other-agent) > very-important and ($k < 20$ and $b > 5$) return false; else return true.
2. **INDEPENDENT OFFER:** in any negotiation round offer my k most preferred, available, un-offered times.
3. **DEPENDENT OFFER:** if negotiation round > b . Apply the simple compromiser sub-strategy which works as follows:
 - If I am an attendee of the meeting, search for any times proposed by the initiator that I am available for but have not offered. If one or more such times exist offer my most preferred time. Else offer the time proposed by the initiator that contains the meeting with the fewest participants out of all such proposed times.
 - If I am the initiator rank all times proposed by other agents according to the number of agents that have proposed that time. Out of all the times with the highest number of proposals if any of these times are available, offer my most preferred such time, otherwise offer the un-available time containing the meeting with the fewest participants.
4. **ABANDON:** if negotiation round > 50 return true.

Figure 1: Offer-k-b negotiator

agent such that its regret about the selection of strategies goes to zero in the limit under certain conditions.

Adapting weights and selecting strategies for negotiating with different agents

For each ‘opponent’ agent the learning agent must learn which strategies to select. The learning algorithm has the following key components, (i) a rule for updating the weights on strategies in the playbook and (ii) a rule for selecting the strategy to apply based on these weights. Bowling *et al.* (2004) applied results in the literature on *experts* problems (also commonly referred to as *k-armed bandits* problems) to derive the rules required. We are able to use the same rules for adapting weights on negotiation strategies. In this section we briefly describe the approach and its basis in the experts literature. For a more complete treatment we refer to the reader to (Bowling, Browning, & Veloso 2004).

In the experts problem an agent chooses actions or options repeatedly based on the instructions it receives from a set of *experts*. Each time the agent needs to make a choice it selects which expert to listen to. In the traditional formulation once the action or option has been selected the agent receives a pay-off from that action. In addition the pay-offs it would have received had it followed the advice of each of the other experts is revealed. The performance of the agent

is measured by the notion of regret. Let the reward received from following the advice of expert i at choice point p be r_i^p . The regret of the agent after k choices have been made is given by the following formula:

$$\text{regret}_k = \max_{\text{over experts } i} \sum_{p=0}^k r_i^p - \sum_{p=0}^k r_{x_p}^p$$

where x_p denotes the expert the agent chose at choice point p . Regret is simply the award achievable by always asking the best expert minus the reward actually achieved. A desirable property of an experts algorithm is that average regret goes to zero as the number of choices approaches infinity. There exist algorithms for various formulations of the problem that achieve no-regret in the limit e.g. (Littlestone & Warmuth 1989; Auer *et al.* 1995)

Bowling *et al.* (2004) show how algorithms for selecting experts with no regret can be used to select plays. In the context of plays (and in the context of selecting strategies for negotiation) we need to use a different formulation of regret that takes into account the fact that not all plays (or strategies) are applicable at each choice point. This can be done by using the notion of *Sleeping Experts* developed by Freund *et al.* (1997). We say an expert is awake when it is applicable at a particular choice point and asleep otherwise. Following the notation used in (Bowling, Browning, & Veloso 2004) we let $a_i^p = 1$ if expert i is awake at choice point p and $a_i^p = 0$ otherwise. Then if $\Delta(n)$ is the set of probability distributions over all n experts. We get the following formula for sleeping regret (SR) after k choices:

$$SR_k = \left(\max_{x \in \Delta(n)} \sum_{p=1}^k \sum_{i=1}^n a_i^p \left(\frac{x(i)}{\sum_{j=1}^n x(j)a_j^p} \right) r_i^p \right) - \sum_{p=0}^k r_{x_p}^p$$

The first half of the formula simply quantifies the reward the agent could have received if the best possible distribution over awake experts had been selected at each choice point.

In the context of plays and negotiation strategies there is one final difficulty. Unlike in the traditional experts problem agents only find out the reward of the action they actually take. In order to account for this Bowling *et al.* (2004) combine elements of the Exp3 algorithm proposed by Auer *et al.* (Auer *et al.* 1995) (which handles the problem of unknown rewards) with the sleeping regret approach of (Freund *et al.* 1997). We describe their approach here and use it to adapt playbook weights for each ‘opponent’ agent and select the strategy to use according to these weights.

Let $R_i^k = \sum_{p=0}^k \hat{r}_i^p$. Where $\hat{r}_i^p = 0$ if i not selected at point p and $\frac{r_i^p}{Pr(x_p=i)}$ otherwise. We call the weight for strategy i at decision point p , w_i^p and we let $w_i^p = e^{R_i^p}$. The value $e^{R_i^p}$ is denoted as m_i^p and we refer to this value as the multiplier and use it to adjust the weights according to the reward received from carrying out the negotiation strategy (or play). The probability that the strategy chosen at point p denoted x_p is strategy i is given by the following equation:

$$Pr(x_p = i) = \frac{a_i^p w_i^p}{\sum_j a_j^p w_j^p}$$

Once strategy x_p has been executed and the reward $r_{x_p}^p$ received we update the weights as follows:

$$w_i^t = \hat{w}_i^p \cdot N_i^p$$

where $\hat{w}_i^p = w_i^{p-1}$ for i not selected, but for i selected:

$$\hat{w}_i^p = w_i^{p-1} (m_i^p)^{\frac{1}{Pr(x_p=i)}}$$

The N_i^p term is used to ensure that sleeping does not affect a strategy’s probability of being chosen. $N_i^p = 1$ if $a_i^p = 0$ and otherwise:

$$N_i^p = \frac{\sum_j a_j^p w_j^{p-1}}{\sum_j a_j^p \hat{w}_j^p}$$

To apply the approach to negotiation we need to decide how we are going to set the multipliers. The multipliers specify the degree to which the success or failure of a strategy affects the weight. We base the multipliers on a model of user utility. We let the utility a user derives from a negotiation strategy take into account three elements:

1. the user’s preference for the time-of-day (tod) the new meeting is scheduled for — $val(tod)$.
2. the increase (or decrease) in utility caused by other meeting in the calendar being moved, i.e. for all meetings that were moved we say the agent’s utility is increased by $\sum_{\text{moved}} val(tod_{\text{new}}) - \sum_{\text{moved}} val(tod_{\text{old}})$.
3. the number of negotiation rounds r required to schedule the new meeting and move any old meetings.

The user’s utility function is parametrized by two constants α and β which specify the relative importance of time-of-day valuations and negotiation cost. Formally a user’s utility for the outcome of a negotiation strategy is modeled as:

$$U(i) = \alpha(val(tod) + \sum_{\text{moved}} val(tod_{\text{new}}) - \sum_{\text{moved}} val(tod_{\text{old}})) - \beta r$$

We use the user’s utility function and highest time-of-day value to estimate the maximum possible utility a negotiation strategy can achieve. We then set the multiplier according to how the reward actually achieved relates to this maximum. The multiplier is set according to the first row of Table 1 that applies. Also note that if the negotiation strategy fails to schedule the new meeting, or to reschedule any bumped meetings, a failure has occurred. Currently we use a multiplier of 0.25 for this case.

The bounds on regret obtained by using the plays approach are strongest if the ‘opponent’ agent is using a fixed strategy and we assume that changes to the environment (i.e., the calendars) are not affecting the rewards. If the other agent is also learning, then in the terminology of (Auer *et al.* 1995), we are dealing with a *non-oblivious* adversary. As such, since the playbook approach builds on Exp3, the theoretical bounds on regret are weaker.

Evaluation

In this section we describe how we have evaluated the effectiveness of using a plays approach to select negotiation strategies.

$U(i) > 0.75 * maxU$	1.75
$U(i) > 0.5 * maxU$	1.5
$U(i) > 0.25 * maxU$	1.25
$U(i) > 0$	1
$U(i) > 0 - 0.25 * maxU$	0.75
$U(i) > 0 - 0.5 * maxU$	0.5
$U(i) > 0 - 0.75 * maxU$	0.25

Table 1: The multiplier used is given by the first row in the table for which the left hand entry evaluates to true

Communication Protocol

We have created a simulation environment consisting of a set of agents equipped with a common protocol for communicating about meetings. The protocol has three basic stages: a negotiation phase, in which agents exchange proposals, a pending stage, in which a time proposed by all the agents is agreed upon, and a confirmation stage, after which the meeting is entered into the agents’ calendars. Support is also provided for *bumping* (canceling and rescheduling) meetings. There are a number of different types of messages that the agents exchange:

- meeting time proposals
- requests to bump meetings
- cancellation notices for meetings
- pending requests for times — when a meeting initiator finds an intersection in proposals, it sends a pending request for one of the times in the intersection to each of the participants.
- pending responses — when an attendee receives a pending request it responds with either:
 - a pending acceptance and marks the meeting as pending, or
 - a pending rejection (if the time is pending for another meeting, we require that the agent reject the request).
- confirmation notices — sent out by the initiator when all attendees reply to a pending request with a pending acceptance.

Negotiation Strategies

We have implemented a number of negotiation strategies that comply with the protocol outlined. We use two of these strategies in our experiments in this paper. The first strategy — Offer-k-b was previously described (see Figure 1). This strategy is parametrized and hence it in fact covers a large number of distinct strategies. The second strategy we use is called Availability-Declarer. This strategy can be very useful in practice, particularly in situations where the agents are very busy. The Availability-Declarer strategy is outlined in Figure 2. The key feature of this strategy is that it offers all the available times in the first week straight away. In subsequent negotiation rounds it does the same for later weeks.

Preferences

We use a simple model of time-of-day preferences in our experiments. Each agent has a preference ordering over times

1. **APPLICABILITY:** if $importance(\text{other-agent}) \geq \text{moderately-important}$ return true.
2. **INDEPENDENT OFFER:** in the first round offer all available times for the current week, in second round offer all available times for the following week and so on until all available times up until the last possible time for the meeting have been offered.
3. **DEPENDENT OFFER:** if negotiation round > 5 , apply the simple compromiser sub-strategy described in Figure 1
4. **ABANDON:** if negotiation round > 50 return true.

Figure 2: Availability Declaring Negotiator

in the morning, times in the middle of the day and times in the afternoon. If for example, the agent prefers the morning, then the middle of the day, and then the afternoon, times in the morning are assigned a value of 3, times in the middle of the day a value of 2 and times in the afternoon, a value of 1.

Experiments and Results

We have conducted a preliminary evaluation of the effectiveness of using a plays based approach to select negotiation strategies. Each of the experiments described in this section consists of one learning agent and three fixed strategy agents of varying preferences and busyness. The learning agent has three strategies in its playbook — Availability-Declarer, Offer-10-5 and Offer-3-5. In the experiments we discuss, these strategies are always applicable.

In each experiment the agents schedule approximately 80 new two person meetings (we restrict our attention to two-person meetings to simplify the discussion). The learning agent is an attendee (not an initiator) of each of these 80 meetings. We show how the learning agent’s playbook weights converge to sensible strategies for each of the fixed strategy agents.

In our first experiment the learning agent has a preference for morning meetings, followed by midday meetings followed by afternoon meetings. The α and β values of the learning agent’s utility function are 4 and 0.1 respectively. The agent’s calendar is approximately 25% full when the experiment is started. Unlike the meetings we schedule in the testing phase the initial meetings in the calendar can involve any number of the agents.

Figure 3 shows how the learning agent’s playbook weights adapt for Agent2. Agent2 starts out with a similar number of initial meetings to the learning agent, uses the Availability-Declarer strategy and has the same time preferences as the learning agent. Figure 3 shows how the playbook weights quickly converge to towards the Availability-Declarer strategy. While the other two strategies are also likely to work well in this instance, the Availability Declarer strategy offers the possibility resolving the negotiation




Figure 3: Weights adaptation for Agent2

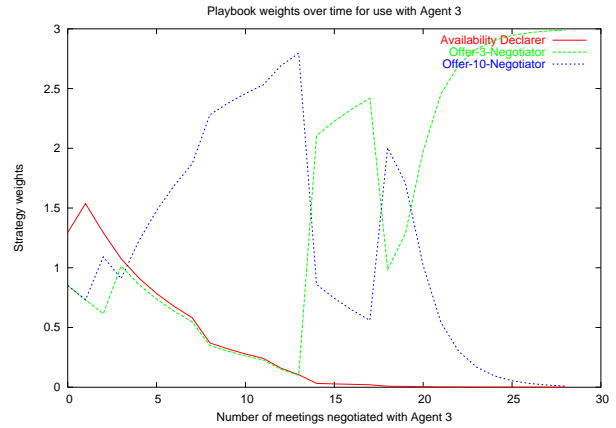


Figure 4: Weight adaptation for Agent3

faster. Since the learning agent and Agent2 have the same preferences there is no strategic advantage to the learning agent only releasing its availability slowly.

Figure 4 shows the weight adaptation for Agent3. Agent3 uses the Availability-Declarer strategy and starts out with a similar calendar density to the learning agent, but with opposite preferences. Agent3 most prefers afternoons, followed by the middle of the day, followed by the morning. Figure 4 shows that the learning agent quickly establishes that the Availability-Declarer strategy is less useful for negotiating with Agent3 than the Offer-10-5 and Offer-3-5 strategies. After about 25 meetings have been scheduled the weights converge on the Offer-3-5 strategy. Note that the Availability-Declarer strategy is a poor choice for use with Agent3. When both agents negotiate with this strategy, the initiator (always Agent3 in these experiments) is likely to quickly find a large intersection of available times. The initiator can choose its most preferred time in this intersection and since Agent3's and the learning agent's preferences clash, the time chosen will likely be bad for the learning agent. The learning agent has a clear strategic incentive to declare its available times more slowly and in order of preference. Since the learning agent's utility function rates achieving good times-of-day much higher than minimizing the number of negotiation rounds, it converges on the Offer-3-5 strategy rather than the Offer-10-5. This is despite the learning agent's calendar being quite full (93%), and hence mutually available slots fairly rare, by the time the experiment concludes.

Figure 5 shows the weight adaptation for Agent4. Agent4 has similar preferences to the learning agent, preferring mid-day times, then mornings, and then afternoons. Agent4 uses the Offer-10-5 negotiator and starts off with a dense initial calendar (about 80% full). Figure 5 shows that the learning Agent quickly determines that the Offer-3-5 strategy is not very effective when dealing with a very busy agent that has similar preferences. After approximately 15 meetings have been scheduled the learning agent converges on the Availability-Declarer strategy.

We ran the same experiment described above but with a

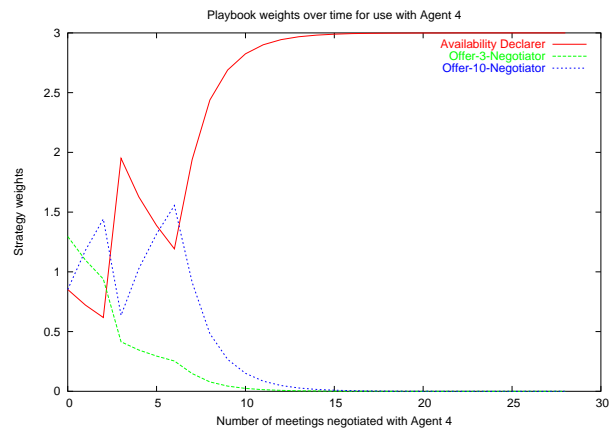


Figure 5: Weight adaptation for Agent4

different utility function for the learning agent and different initial calendars. The utility function we used had an α value of 4 and a β value of 1. Making this change caused the weights to converge on the Availability-Declarer strategy for each of the agents since the negative effect of negotiation length on utility was greatly increased.

We have also run preliminary experiments to evaluate how well the learning agent actually optimizes the user's pay-off. We found that using the learning agent resulted in significantly higher utility (approximately 15%) than randomizing over the strategies in the playbook. In fact in all the experiments we ran, the learning agent performed better than it would have had it used the best (in hindsight) fixed strategy.

Conclusions and Future Work

We introduced the idea of using a *playbook* approach for learning to select the best strategies for negotiating with different agents. The space of all possible negotiation strategies is huge, and as such it is not possible for an agent to learn how to negotiate in the complete space. The plays-based approach cuts the strategy space down to a set of strategies that are effective in different situations allowing an agent

to learn which of these strategies work best with different fixed-strategy agents. This approach provides some theoretical bounds on the regret the learning agent can experience. We have demonstrated experimentally that this approach holds a lot of promise for learning how to select negotiation strategies.

We are currently working on a thorough experimental evaluation of plays for multi-agent meeting scheduling. This evaluation will consider a wider variety of plays and fixed-strategies. In the future we also plan to explore the problem of learning how to negotiate in the presence of agents that are also learning.

Acknowledgments

Thanks to the anonymous reviewers and Michael Bowling for helpful comments and suggestions. This research is sponsored by the Department of the Interior (DOI) - National Business Center (NBC) and the Defense Advanced Research Projects Agency (DARPA) under contract no. NBCHC030029. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the DOI, NBC, DARPA or the U.S. government.

References

- Auer, P.; Cesa-Bianchi, N.; Freund, Y.; and Schapire, R. E. 1995. Gambling in a rigged casino: the adversarial multi-armed bandit problem. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science*, 322–331. IEEE Computer Society Press, Los Alamitos, CA.
- Bowling, M.; Browning, B.; and Veloso, M. 2004. Plays as effective multiagent plans enabling opponent-adaptive play selection. In *Proceedings of International Conference on Automated Planning and Scheduling (ICAPS'04)*. in press.
- Crawford, E., and Veloso, M. 2004. Opportunities for learning in multi-agent meeting scheduling. In *In the Proceedings of the AAAI 2004 Symposium on Artificial Multi-agent Learning, Washington, DC*.
- Ephrati, E.; Zlotkin, G.; and Rosenschein, J. S. 1994. A non-manipulable meeting scheduling system. In *Proc. International Workshop on Distributed Artificial Intelligence*.
- Freund, Y.; Schapire, R. E.; Singer, Y.; and Warmuth, M. K. 1997. Using and combining predictors that specialize. In *STOC '97: Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, 334–343. New York, NY, USA: ACM Press.
- Garrido, L., and Sycara, K. 1995. Multi-agent meeting scheduling: Preliminary experimental results. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS'95)*.
- Gonzlez, A., and Santos, J. A negotiation protocol for meeting scheduling based on a multiagent system.
- Jennings, N. R., and Jackson, A. J. 1995. Agent based meeting scheduling: A design and implementation. *IEE Electronics Letters* 31(5):350–352.

Littlestone, N., and Warmuth, M. K. 1989. The weighted majority algorithm. In *IEEE Symposium on Foundations of Computer Science*, 256–261.

Pragnesh Jay Modi, M. V. 2005. Bumping strategies for the private incremental multiagent agreement problem. In *AAAI Spring Symposium on Persistent Agents*.

Sen, S., and Durfee, E. H. 1998. A formal study of distributed meeting scheduling. *Group Decision and Negotiation* 7:265–289.

Sen, S.; Haynes, T.; and Arora, N. 1997. Satisfying user preferences while negotiating meetings. *Int. J. Hum.-Comput. Stud.* 47(3):407–427.

Shintani, T.; Ito, T.; and Sycara, K. 2000. Multiple negotiations among agents for a distributed meeting scheduler. In *Proceedings of the Fourth International Conference on MultiAgent Systems*, 435 – 436.