

Mechanism Design for Multi-Agent Meeting Scheduling Including Time Preferences, Availability, and Value of Presence

Elisabeth Crawford and Manuela Veloso

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

Abstract

In this paper we examine the benefits and limitations of mechanism design as it applies to multi-agent meeting scheduling. We look at the problem of scheduling multiple meetings between various groups of agents that arise over time. Each of the agents has private information regarding their time preferences for meetings. Our aim is to extract this information and assign the meetings to times in a way that maximises social welfare. We discuss problems with previous attempts to design incentive compatible (IC) and individually rational (IR) mechanisms for the task. We focus on the problem of determining when agents are available. In particular, we show that when agents with general valuation functions are asked to supply their availability for meeting times, there is no IC and IR mechanism. Given this impossibility result, we show how the likelihood of violating IR can be reduced through agents expressing their value for the presence of others at meetings. We also show how requesting agent preferences for entire schedules helps to eliminate IC problems.

1. Introduction

For many people, scheduling meetings is a daily and time consuming task. Scheduling multiple meetings with many participants, possibly in parallel, is cognitively demanding. As such, a number of researchers e.g [10, 7, 2] have looked into designing intelligent software agents that can remove this burden from computer users. In this paper we assume we have software agents that know the preferences of their user. We focus then, on the problem of designing a mechanism that enables such agents to effectively schedule meetings.

We view the multi-agent meeting scheduling problem (MAMS) as consisting of a set of agents A , a set of sequential time slots T and a set of meetings M that are requested over time. Each agent A has preferences about scheduling. We call agent i 's true preference function θ_{it} , where the domain of θ_{it} is a scheduling option (e.g a particular time, or

an entire schedule) and the range is a value representing the utility of this option to the agent. A solution to the MAMS problem consists of an assignment of a subset of the requested meetings M to time slots in T . Clearly there can be good and bad solutions to MAMS. When designing a system to solve the problem it is necessary to decide what, if any metric, is to be maximised.

The paper is organised into 5 sections. First we introduce some concepts from mechanism design and briefly discuss negotiation approaches to MAMS. In Section 3 we highlight IC problems with previous work. In Section 4 we concentrate on the problem of providing agents with an incentive to truthfully reveal their availability. We show that there is no mechanism that can solve MAMS when agents are asked to reveal their availability, satisfying IC and IR. Finally, in Section 5, we propose a method for reducing IR problems by introducing a process whereby, agents quantify the value of other agents to meetings.

2. Background

Mechanism design is concerned with defining the rules of a game such that some desirable outcome can be achieved. The participants of the game are assumed to act with *self-interest* according to their private preferences about the outcome [8].

The outcome desired by the mechanism designer depends largely on the specific application. However, the *social welfare maximising outcome* is a common choice. The social welfare maximising outcome is the outcome that maximises the sum of the agents' utilities. To compute the social welfare maximising outcome we need to know each agents utility for every possible outcome. These values are known only to each agent. As such, it is the job of the mechanism to get the agents to state these values truthfully.

Two common desiderata for mechanisms are incentive compatibility and individual rationality. A mechanism is *incentive compatible* (IC) if it is every agents dominant strategy to reveal their utility values truthfully. A mechanism is *individually rational* (IR) if agents cannot receive a negative pay-off from the mechanism. IC mechanisms normally

require some agents to make payments to the system and the system sometimes also distributes payments to agents. The clark tax mechanism [1] is a well known mechanism that generally achieves both IC and IR.

The *clarke tax mechanism* [1] works by charging each agent for the amount they influence the chosen meeting time. Each agent i states their money value for each possible outcome. The social welfare maximising outcome x is then selected. Agent i pays an amount equal to the sum of the other agents' utilities for the outcome that would have been chosen had they not participated, minus the sum of the other agents' utilities for x .

In a real world MAMS mechanism we cannot use money to decide on meeting times. Instead, we need to replace money with points as done by Ephrati et al [2]. The system can distribute points to the users periodically which they can use to bid for meeting times. These points take on some of the qualities of money as they are needed repeatedly by users to influence the scheduling of new meetings.

Much of the previous work on the MAMS problem has looked at using negotiation protocols between agents to find a solution e.g. [7, 10, 4]. The following, is a simple example of the kind of negotiation protocol looked at in previous work:

- Host announces meeting
- Iterate until host finds an intersection of times
 - Host proposes some times
 - Participants propose some times

These negotiation approaches assume that agents are to some extent truthful about their availability and utilities for times and/or are willing to compromise their preferences. We can say an agent is truthful in such a protocol if they offer all times they are available for eventually, and offer times in the order of their preference. Truthfulness however, is certainly not a dominant strategy. Suppose all the agents but Agent i are truthful. Then Agent i may be able to gain by just proposing her favourite times and hoping the other agents eventually agree. This lack of IC could result in problems. For instance, an intersection of available times may not be found, despite one actually existing or the process may iterate for a very long time.

3. Issues with existing mechanism design approaches to MAMS

A mechanism design approach to MAMS was first taken by Ephrati, Zlotkin and Rosenschein [2]. Ephrati et al, used a *clarke tax* mechanism to try and ensure IC. Three different approaches were taken to implementing the mechanism: meeting-oriented; calendar-oriented and schedule-oriented. The authors claim that each of these approaches is IC. However their proof did not account for the *repeated* application of the clark tax mechanism, rather it looked at a single step. We show by counter example that each of these approaches are not IC and we observe that achieving IC is not a trivial problem in MAMS.

3.1. Meeting-Oriented Approach

In the meeting-oriented approach proposed by [2] when each new meeting enters the system all agents are given more points with which to bid, then the agents attending the meeting bid over all time slots in the time horizon. The social choice is calculated, the tax applied and the meeting recorded in the system's copy of each agents' calendar. The time the meeting is scheduled for, is not considered available when scheduling future meetings for those agents. The authors claim that this is IC, however agents may have combinatorial preferences over meetings and times. Consider the following example. Agent 1 has a low utility for meetings on Tuesday, but wants to complete some work with some other agents by Wednesday that will require two meetings over two days. Now, if Agent 1 thinks the other agents want one of the meetings scheduled on Wednesday, it may be in Agent 1's best interest to overbid for time slots on Tuesday. This is because Agent 1 gets a high utility from both meetings happening by Wednesday. Suppose

| | Tue | Wed | Thu | Fri |
|---------------|-----|-----|-----|-----|
| Agent1 | 2 | 4 | 3 | 2 |
| Agent2 | 3 | 5 | 4 | 4 |
| Agent3 | 3 | 6 | 3 | 3 |
| Total Welfare | 8 | 15 | 10 | 9 |

Table 1. The Agents' true preferences

the agents' true preferences over the times without lookahead are those in Table 3.1. Let Agent 1's utility from the first meeting being scheduled on Tuesday *and* the second on Wednesday be 20 - which is more than her preference for those days on their own. Now, if we assume that the other two agents will not change their preferences when the new meeting enters the system, the first meeting will be scheduled on Wednesday and the second on Thursday. This will have a utility of 7 for Agent 1, and she will pay a tax of 0 for the first meeting and 0 for the second. Now suppose Agent 1 changes her bid for Tuesday to 10. The first meeting will be scheduled on Tuesday and Agent 1 will pay a tax of 5. The second meeting, will be scheduled on Wednesday, and agent 1 will again pay a tax of 0. So in total Agent 1 will get a payoff of 15 which is higher than the 7 points for bidding truthfully. Hence we have shown by counter-example that the meeting-oriented approach proposed in [2] is not IC.

3.2. Calendar-Oriented Approach

In the *calendar-oriented* approach proposed in [2] the agents express their preferences over a *complete time period once only*. As time passes new times are added to the period under consideration and the agents are asked for their

preferences. This approach does not allow agents to express preferences that are *meeting dependent*. Again this approach is not IC.

Consider the following scenario. Agent 1 has a high utility for back to back meetings. Suppose Agent 1 is considering how to allocate her preferences for 1,2 and 3pm on Monday. Agent 1 knows from experience that 2pm is a very popular time with many of the people she regularly needs to meet with, and as such she expects a meeting will be scheduled then. Agent 1 has a low utility for meeting at 1pm, but since she expects a meeting to be scheduled at 2pm perhaps she should overrate her preference for 1pm? Table 3.2 con-

| | 1pm | 2pm | 3pm | 4pm |
|--------|-----|-----|-----|-----|
| Agent1 | 3 | 4 | 2 | 2 |
| Agent2 | 2 | 3 | 2 | 4 |
| Total | 5 | 7 | 4 | 6 |

Table 2. True time preferences for Monday afternoon

tains Agents 1’s and Agent 2’s true preferences over meeting times on Monday afternoon. If we schedule 2 meetings with these preferences the first will be scheduled at 2pm and the second at 4pm and Agent 1 will have a pay off of 5. But, Agent 1’s back to back preference means the declaration in Table 3 will lead to a higher payoff. If Agent 1 reveals

| | 1pm | 2pm | 3pm | 4pm |
|--------|-----|-----|-----|-----|
| Agent1 | 5 | 4 | 2 | 2 |
| Agent2 | 2 | 3 | 2 | 4 |
| Total | 7 | 7 | 4 | 6 |

Table 3. Strategic declarations for Monday afternoon

the preferences in Table 3 she will receive 12 points in utility and pay a tax of 1 for the first meeting and a tax of 2 for the second resulting in a final payoff of 9 points. Hence we have shown by counter-example that the calendar-oriented approach proposed in [2] is not IC.

3.3. Schedule-Oriented Approach

The problems that have arisen in the above mechanisms have been caused by a lack of expressiveness in the bidding process. The lack of expressiveness has resulted in truthfulness not being a dominant strategy and hence the mechanisms are not IC. The final approach proposed in [2] allows the agents to bid over all possible schedules. The schedules are generated by taking every possible combination of meetings and time slots up to some time horizon. As each new meeting enters the system the agents are allowed to bid over all feasible schedule configurations. This approach has

more promise in terms of IC, however, there are still problems.

Suppose the time horizon in this method is two weeks. Further, suppose that the time a particular meeting is scheduled next month affects some agent’s preference for when a meeting is scheduled this month. This causes the same IC problem we saw for the meeting-oriented example. Thus, this mechanism is not IC either. The longer the time horizon the less likely this problem is to occur. As the time horizon increases the number of possible schedules grows exponentially and hence there is a trade-off between IC and complexity.

4. Towards an effective solution for MAMS

In this section we consider in more detail a variant of the most promising approach proposed in [2] - the schedule-oriented approach. We show that even if we ignore time horizon issues, this mechanism has IC problems. We discuss how this particular problem can be solved. We then consider the problem of incorporating the availability of agents for particular times into a mechanism. We show that there is no mechanism that can incorporate availability declarations from agents that is both IC and IR for MAMS. Finally, we discuss the implications of this impossibility result.

4.1. A schedule-oriented mechanism

The mechanism shown in Table 4.1 is a variation of the schedule-oriented mechanism described in [2]. The mechanism in [2] has a time frame parameter that specifies the number of time slots that are to be frozen as time progresses. By frozen we mean that any meeting scheduled cannot be changed. Clearly, if we allow time frames longer than one time slot, the agents cannot fully express their preferences and we will have IC problems. Thus, in the above mechanism we do not refer to a time frame. In fact, a meeting could be moved right up until the *current-time* if this would maximise social welfare. Unfortunately, despite the fact that allowing agents to fully express their preferences within some time horizon has greatly reduced IC problems, at least one issue remains.

4.1.1. Further IC issues with the schedule-oriented approach The following example exposes a further IC problem with the schedule-oriented approach. Suppose our time horizon in Table 4.1 is a week in advance. It is Monday, there is some meeting that will most likely be scheduled on Friday that Agent 1 wants to attend. Agent 1 has a strong preference for the meeting to be at 2pm. Now suppose Agent 1 expects 5 new meetings will enter the system between now and Friday. When each meeting enters, Agent 1 will have to bid over the possible schedules. The question, is whether or not Agent 1 should truthfully reveal, that he greatly prefers schedules where the meeting is held at 2pm.

Input: length of time slots, time horizon

procedure MAMS1:

1. Divide calendar into time slots t_1, t_2, \dots, t_h of length m
2. $current-time = t_1$
3. $horizon-time = t_h$
4. \forall agents i , $budget_i = x$
5. When agent requests meeting, do
6. \forall agents i , $budget_i+ = x_{new}$
7. notify all agents of meeting
8. notify agents if their presence is requested
9. generate all possible combinations of meetings
and times from current time to horizon time,
 $schedule_1, schedule_2, \dots, schedule_k$
10. \forall agents i , request valuation for each schedule
11. check \forall agents i and \forall schedules j :
12. $value_i(schedule_j) \leq budget_i$
13. select schedule that maximises social welfare
14. report schedule to agents
15. \forall agents i , $budget_i- = clarktax(i)$
16. when m units of time have passed
17. current time = next time slot
18. create new time slot at horizon, and reassign horizon time

Table 4. Variant of the schedule-oriented approach

If Agent 1 reveals this preference when each new meeting needs to be scheduled, his preferences are likely to strongly influence the outcome. Thus, he can expect to pay a high tax each time. Instead of paying the tax each time, Agent 1 would be better off pretending he does not have a strong preference for the time of the meeting until the 5th meeting enters the system. At this point he can declare his preference and thus only risk having to pay a high tax once. As such, it is not a dominant strategy for Agent 1 to reveal his true preferences at each iteration and hence the mechanism is not IC. The complicated nature of this example highlights that designing an IC mechanism for this problem is non-trivial.

4.1.2. A fix We show a way the problem described in the previous section can be removed. Suppose we fix a schedule for some time period in one round of bidding. This eliminates the ability of meetings to be moved, and thus the problem described in the previous section. However, since we also need agents to be able to fully express their preferences over the time period once the schedule has been fixed, no new meeting can be added into that time period. Consider the algorithm in Table 5

The difficulty with the mechanism in Table 5 is that the agents would have to think ahead about the meetings they want scheduled. The shorter the time period chosen, e.g 3 days, a week, a fortnight etc the greater the time horizon IC problem but the smaller the problem of having to think

Input: period-length, slot-length

procedure MAMS2:

1. Divide the calendar up into time periods p_1, p_2, \dots, p_n
2. Divide the time periods up into slots t_1, t_2, \dots, t_k
3. for current-period in p_1, p_2, \dots, p_n
4. while current-period not expired:
5. append requests for meetings in $p_{current+x}$ $x > 0$ to lists
6. \forall agents i , $budget_i+ = no-meetings * x_{new}$
7. notify agents of all meetings
8. notify agents of which meetings their presence is requested at
9. steps 9-15 of MAMS1

Table 5. Fixed schedule-oriented approach to MAMS

ahead. We believe MAMS2, to be IC, although we have not constructed a formal proof. The main focus of this paper is on the incorporation of availability declarations into mechanisms for MAMS, and that is where we now turn our attention.

4.2. Availability

An important missing element from the approaches discussed thus far, is a way of handling the unavailability of agents. The systems proposed in [2], did not allow agents to indicate they are unavailable for certain times. In the calendar-oriented and meeting-oriented approaches from [2], if agents had a meeting scheduled by the system for a particular time, they were considered unavailable at that time. This was the only way they could be considered unavailable. However, for a meeting scheduling system to achieve efficient results we need agents to be able to declare their unavailability. Users may have important commitments such as appointments with medical specialists that cannot be moved. Furthermore, there can also be implications for the efficiency of the organisation the user works for. Suppose for instance, that a user has a very important client presentation. It may be in both the user's and the company's interest if the user has no meeting immediately prior to the presentation so they can complete their preparations. As such, any practically implementable MAMS system must be able to get agents to reveal their availability as well as their time preferences.

We say that when an agent declares they are unavailable for a time, that they *veto* or *block* that time. In other words, their declaration means that no meeting of which they are a participant can be held at that time. Simply asking agents for the times they are unavailable is not an IC approach. If we used this scheme there would be a clear incentive for the agents to say they are only available at the times for which they have the highest utility. Even more so than with preferences, we need agents to report their availability truthfully.

4.3. Availability and Individual Rationality

In this section we show that there is no mechanism for MAMS with availability declarations that is both IC, IR and maximises social welfare when agents have general valuation functions.

4.3.1. Some Definitions Recall that a mechanism is *individually rational* (IR) if every agent is never worse off from participating in the mechanism. IR is a very desirable property in mechanism design since users do not have to reason about whether or not they should take part.

When applying mechanism design to MAMS we are assuming agents have *quasi-linear* preferences in order to circumvent the *Gibbard-Satterthwaite Impossibility Theorem* [3] [9]. Let $x \in X$ be the possible outcomes of the mechanism. Let p_i be the payment agent i is required to make to the system. Now we let $u_i(x)$ be agent i 's utility from the whole process if the mechanism selects x , and let $v_i(x)$ be i 's valuation for that outcome.

We say that agent i has *quasi-linear* preferences if $u_i(x) = v_i(x) - p_i$, i.e., i 's utility is simply her valuation for the chosen outcome, minus the payment she must make. The assumption of quasi-linear preferences is considered reasonable in markets and it is no less reasonable in the context of meeting scheduling. Despite the restriction of quasi-linear preferences, we allow the valuation functions of the agents for outcomes to be general.

The Groves family of mechanisms were developed in [6], [1] and [11] for the domain of quasi-linear preferences. Let $\theta_i \in \theta^r$ be the reported valuation function of each agent over X

The Groves mechanism selects the outcome, x^* such that,

$$x^*(\theta^r) = \underset{x \in X}{\operatorname{argmax}} \sum_{\forall i} v_i(x)$$

In other words, Groves mechanisms select the outcome that maximises social welfare.

The payment rule of the groves mechanism is,

$$p_i(\theta^r) = h_i(\theta_{-i}^r) - \sum_{j \neq i} v_j(x^*)$$

where $h_i : \theta_{-i}^r \rightarrow \mathfrak{R}$ is an arbitrary function on the reported valuations of every agent except i .

The Groves mechanism are unique in the domain of quasi-linear preferences in that they are the only direct mechanisms that are allocatively efficient, IC and implementable in dominant strategies for general valuation functions [5]. By the Revelation Principle this result extends to general mechanisms[3].

Lemma 4.1 *There exists no Groves mechanism for MAMS with availability declarations that maximises social welfare*

and simultaneously satisfies IC and IR for agents with general valuation functions.

Proof For notational convenience we will say that an agent veto's a time by submitting a bid of 0, and accepts a time by submitting a bid of greater than or equal to 1. For the mechanism to be IR we require $u_i(x^*) \geq 0 \forall$ agents i . In other words i 's payment p_i must not exceed i 's valuation for the chosen outcome. Now, for IR to hold $v_i(x^*) \geq h_i(\theta_{-i}^r) - \sum_{i \neq j} v_j(x^*)$. Notice that h_i only depends on the reported valuations of the other agents. Consider the situation where agent i 's valuation for the winning time is 3. Now, for the mechanism to be IR $h_i(\theta_{-i}^r) \leq \sum_{i \neq j} v_j(x^*) + 3$. Can we ensure this without reference to i 's valuations i.e. only with reference to the valuations of the other players?

Without the power of veto we can ensure we never charge i an amount that is more than IR by reference to the fact that the winning outcome is the social choice. For instance, in the clarke tax mechanism, $h_i(\theta_{-i}^r)$ is set to the sum of the valuations of the other agents for the outcome that would have been chosen had i not participated. Then $v_i(x^*) \geq h_i(\theta_{-i}^r) - \sum_{i \neq j} v_j(x^*)$, since otherwise x^* would not have been the social choice. Note that we can make no such conclusions here. Agent i 's ability to veto times means that the difference between the utility to the other agents of the winner without agent i may be much greater than the utility to the other agents of x^* . The only outcomes we can make conclusions about are those which i does not veto. However we cannot base our function $h_i(\theta_{-i}^r)$ on such instances since there may be only one time that agent i does not veto. More importantly this would allow agent i to significantly influence the function h_i and the IC of Groves mechanisms depends upon h_i being independent of agent i .

Thus, since there is nothing we can safely base h_i on, we must simply set it to some arbitrary number. To ensure we do not violate IR then, we must set $h_i(\theta_{-i}^r) \leq 1 + \sum_{i \neq j} v_j(x^*) \forall \theta_{-i}^r$. Since we said 1 is the smallest possible value i could have assigned to the winning time. There is a simple manipulation that shows that this mechanism is not IC. Suppose we have the valuations of two agents for two times as in Table 7. Now, x_2 is clearly the social choice. No-

| | x_1 | x_2 |
|--------|-------|-------|
| Agent1 | 4 | 2 |
| Agent2 | 1 | 8 |
| Total | 5 | 10 |

Table 6. Agents' true preferences over outcomes

tice that this outcome yields a payoff of at most 2 to Agent1. Clearly truth telling is not a dominant strategy since the payoff to Agent 1 of the lie ($x_1 = 4, x_2 = 0$) yields a payoff of

at least 3 given our IR payment rule. Hence in our meeting scenario with general valuation functions there is no Groves mechanism that is simultaneously IC and IR if we take unavailability into account.

Theorem 4.2 *There is no mechanism for MAMS with availability declarations that is both IC, IR and maximises social welfare when agents have general valuation functions.*

Proof This result follows directly from the uniqueness of Groves mechanisms and from the revelation principle.

The only possible way we can achieve IC and IR is by defining some restriction on the valuations of the agents. This is something we would like to explore in the future. However, the IC problems due to the possible complexity of agent preferences discussed with regards to Ephrati et al’s work [2] indicate the defining a sensible restriction may be very difficult if not impossible. We note that the result described here is not only applicable to the MAMS problem. It applies more generally to situations where it is necessary to allow agents to veto some outcomes.

4.4. Clarke Tax Mechanism for Meeting Scheduling Without IR

We saw in the previous section that it is not possible to have a meeting scheduling mechanism that is both IC and IR for agents with general valuation functions. However, suppose we ignored the problem of IR and used a clarke tax mechanism to encourage truthfulness. Table 8 shows example tax calculations for three agents under this mechanism.

| | sch_1 | sch_2 | sch_3 | sch_4 | sch_5 | Tax |
|--------|---------|---------|---------|---------|---------|-----|
| Agent1 | 1 | 4 | 0 | 0 | 0 | 11 |
| Agent2 | 5 | 3 | 4 | 9 | 4 | 0 |
| Agent3 | 2 | 4 | 9 | 9 | 3 | 1 |
| Total | 8 | 11 | 13 | 18 | 7 | |

Table 7. Agents’ true preferences, including availability

In the table a value of 0 indicates that the schedule is infeasible for that agent. When choosing the social welfare maximising outcomes these schedules are blocked off. Thus, sch_2 with value 11 is chosen. Now consider the tax that Agent 1 has to pay. If Agent 1 was not involved in the process, sch_4 would have been the social choice with value to the other players of 18. Without Agent 1, the value for sch_2 is only 7. To compute Agent 1’s tax we minus 7 from 18 giving a tax of 11 as displayed in the table. Agent 1 is paying a high tax because their unavailability is greatly affecting the outcome. Notice that the tax that Agent 1 pays is

significantly higher than her value for the chosen schedule. Thus, it was clearly not in Agent 1’s interest to participate in mechanism in this round. If we continue to schedule meetings and Agent 1’s unavailability continues to cause her to pay an amount of tax that is not IR she may end up with no points to express positive preferences with. We cannot limit the amount an agent spends to stop her points from becoming negative, because even if she only places bids of size 1 on times she is available for, her unavailability on its own can cause a large tax. This example demonstrates how the mechanism’s failure to ensure IR can reduce its effectiveness.

5. Reducing Veto Power

The key problem we have encountered so far, is each agent’s ability to block off times at which the meeting can occur. In this section, we propose a reduction of this power to veto times, that makes it less likely that in the scheduling of any meeting an agent pays more tax than the utility he/she receives.

In many instances of the meeting scheduling problem, when there are more than two participants, one person’s unavailability at a particular time does not completely block off that time. If the time is very particularly favoured by the other participants, or if the unavailable participant is not very important for the meeting, the meeting may still be scheduled. What we then aim to do is describe a mechanism that copes with this reduction of veto power.

We propose a mechanism where not only do agents express preferences over schedules, but also over attendees. Each agent specifies their preferences for schedules and the amount their utility for a schedule is reduced by the absence of every combination of the other participants from each meeting in the schedule. Thus, a schedule may be picked where some agents are not available for all meetings.

Recall the preferences in Table 8. Now suppose that Agent 1 loses 1 utility point each for the absence of the other two agents. Agent 2 loses 3 utility points from Agent 1’s absence and 2 from Agent 3’s. Agent 3 loses 2 from Agent 1’s absence and 1 from Agent 2’s. Suppose also, that when two agents are unavailable the utility for the remaining agent is always reduced by 10. When a combination of the agents is unavailable we reduce the utility of that schedule to the other agents by the amount they value the participation of the absent combination. Table 9 shows the reduction in utilities that occurs. To show how this works, let’s consider Agent 2. Now Agent 2’s original preference for sch_3 was 4, as shown in Table 8. But, Agent 1 cannot attend a meeting specified in this schedule, so Agent 2’s valuation for sch_3 is reduced by 3, which is the amount that Agent 2 values Agent 1’s presence at the meeting. The other reductions are computed similarly.

| | sch_1 | sch_2 | sch_3 | sch_4 | sch_5 | Tax |
|--------|---------|---------|---------|---------|---------|-----|
| Agent1 | 1 | 4 | 0 | 0 | 0 | 0 |
| Agent2 | 5 | 3 | 1 | 6 | 1 | 1 |
| Agent3 | 2 | 4 | 7 | 7 | 1 | 1 |
| Total | 8 | 11 | 8 | 13 | 2 | |

Table 8. True preferences, including value for others

We now choose the social welfare maximising schedule. Notice that the schedule chosen will be a schedule for which Agent 1 is unavailable - sch_4 since the absence of agent 1 did not outweigh the others values for that time. The tax is then computed using the normal clark tax formula. To show that this scheme maintains IC, if used once, we first need a formal expression for the tax that is imposed on the agents. We will define the tax as follows:

$$tax_i(\theta_i) = \sum_{j \neq i} v_j(x_{-i}^*) - v_j'(absent(x_{-i}^*)) \quad (1)$$

$$- \sum_{j \neq i} v_j(x^*) - v_j'(absent(x^*)) \quad (2)$$

where v_j' computes the loss in utility to agent j of the absence of unavailable agents for each meeting and x_{-i}^* is the outcome that would have been chosen had i not participated. The tax is thus, the value the other agents would have had for the social welfare maximising outcome had i not been a participant in the process at all, minus the other agents valuations for the outcome chosen.

Notice that agent i can have no effect on the first summand. Thus, i wants to report θ_i and her value for the presence of others such that the system choses x^* that maximises the following equation:

$$v_i(x^*) - v_i'(absent(x^*)) + \sum_{j \neq i} v_j(x^*) - v_j'(absent(x^*))$$

But this is exactly what the mechanism is trying to maximise since this will give the social welfare maximising outcome. Thus, agent i is best off reporting his/her true valuations.

So this system maintains IC if used once. Thus, we can use it in either MAMS1 or MAMS2 or some other mechanism to handle availability. This scheme can still result in outcomes that are not IR, however it will happen less often as demonstrated by the example.

6. Conclusions

The benefit of taking a mechanism design approach to MAMS is that it gives the possibility of maximising social welfare with no strategic thinking required by the agents (they can just report their true valuations). However, to achieve this goal, we need to design a mechanism that is IC and *practical*.

We have demonstrated IC problems with previous approaches and have shown ways to fix or reduce these difficulties. Furthermore, we have addressed the practical issue of incorporating availability declarations into a mechanism. We have proved that it is impossible to do this and maintain IC and IR when agents have general valuation functions. We have shown how the IR problem can be reduced, by allowing agents to express how much they value the attendance of each combination of participants at a meeting. However further progress needs to be made for mechanism design to be practical for MAMS. Our results demonstrate that making mechanism design work in real-world multi-agent systems is a theoretically challenging problem.

In the future we would like to further explore the problem of IC in MAMS. Also, the feasibility of any mechanism relies on the existence of software agents that can learn people's scheduling preferences. As such, exploring this learning task is an important direction for future research.

Acknowledgments

Thanks to Tuomas Sandholm and the Fall'03 CS 15-892 class for helpful feedback.

References

- [1] H. E. Clarke. Multi-part pricing of public goods. *Public Choice*, 11:17–33, 1971.
- [2] E. Ephrati, G. Zlotkin, and J. S. Rosenschein. A non-manipulable meeting scheduling system. In *Proc. International Workshop on Distributed Artificial Intelligence*, Seattle, WA, 1994.
- [3] A. Gibbard. Manipulation of voting schemes. a general result. *Econometrica*, 41:587–602, 1973.
- [4] A. Gonzalez and J. Santos. A negotiation protocol for meeting scheduling based on a multiagent system.
- [5] J. R. Green and J.-J. Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45:427–438, 1977.
- [6] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- [7] N. R. Jennings and A. J. Jackson. Agent based meeting scheduling: A design and implementation. *IEE Electronics Letters*, 31(5):350–352, 1995.
- [8] D. Parkes. Iterative combinatorial auctions: Achieving economic and computational efficiency, 2001.
- [9] M. Satterthwaite. Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.
- [10] S. Sen and E. H. Durfee. A formal study of distributed meeting scheduling. *Group Decision and Negotiation*, 7:265–289, 1998.
- [11] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.