Stable Group Scheduling

(Extended Abstract)

Hooyeon Lee
Stanford University
haden.lee@stanford.edu

Yoav Shoham
Stanford University
shoham@cs.stanford.edu

ABSTRACT

We consider the situation in which an organizer is trying to convene an event, and needs to decide on a *schedule*: a time slot and a set of invitees from a given set of agents. For each possible time slot, each agent has a single-peaked preference over the number of attendees at the event. Agent also has the outside option of not attending, which she prefers in some situations. The task of the organizer is to issue a maximum stable schedule – the invited agents prefer attending to not attending, the agents not invited do not regret not being invited, and the event has the maximum number of attendees subject to these stability requirements. We consider both the non-strategic and strategic cases. In the former, in which agents truthfully reveal their preferences, we provide a polynomial-time algorithm for determining whether a stable schedule exists, and if it does, determining the maximum such schedule. In the strategic case we provide a truthful mechanism for the case in which the preferences of the agents are monotonically increasing, and an impossibility result for the general case.

Categories and Subject Descriptors

J.4 [**Social and Behavioral Sciences**]: Economics

Keywords

Group Scheduling; Strategic Agents

1. INTRODUCTION

Imagine an event organizer trying to convene an event – for example, a fundraiser. She needs to choose a time for the event, and whom to invite among a set of possible invitees. Her goal is to maximize the number of attendees in order to maximize the total donations. The potential invitees, however, have their own preferences over both the time and the number of attendees at the event (but not the identities of attendees). For example, a given donor might want that there not be too few attendees so she doesn't feel the spotlight, but also that the event not be overly crowded. We allow the preferences of an invitee to vary across days, and also assume that an invitee always has the outside option of not attending, which she may prefer in some cases.

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

A *schedule* is a (time, invitee-set)-pair. A schedule is *stable* if all invitees prefer attending to not attending, and if no uninvited person wishes she had been invited. Stability is obviously desirable, but in general a stable schedule may not exist. This naturally raises the question of how hard it is to determine whether it does exist for a given setting, and if it does, what a stable schedule is with the maximum number of attendees. These questions take an extra meaning in the strategic case, in which agents may misreport their preferences. In this case, can the organizer incentivize the agents to disclose their true preferences, and if so what is the largest group that can be stably assembled? We provide a positive result for the non-strategic case in which agents are truthful. For the strategic case, we provide both a similar positive result when the agents' preferences are increasing, and an impossibility result for the general case.

2. RELATED WORK

Group scheduling is of tremendous practical importance, and much research has been devoted to it, including in Artificial Intelligence (AI). Jennings et al. [4] proposed the design of an agent-based meeting scheduling system, in which an autonomous agent negotiates with other agents on behalf of its human user. Mitchell et al. [6] and Maes [5] respectively focused on learning preferences of human users and reducing the amount of work needed by human users. Crawford and Veloso [1] approached to the negotiation-based group scheduling problem by training agents to learn about the strategies of other agents. Ephrati et al. [3] tackled incentive issues with strategic agents by proposing monetary-based scheduling systems, while assuming that availability of invitees is known; in our work we do not make this assumption.

The most closely related work of which we are aware is done by Darmann et al. [2]. There agents are assumed to have preferences over activities as well as number-of-participants. The authors define stability requirements and seek maximum solutions, where a solution is an assignment of agents to activities. We inherit from the work of Darmann et al. both the preference structure of the agents and the stability criterion. Beyond these commonalities, however, lie several differences. First, agents can be assigned to one of any number of activities in their work, whereas in our problem a single time slot must be selected for the event. Thus our framework can be viewed as a restriction of the solution space (from any number of activities to exactly one time slot). More dramatically, Darmann et al. only consider non-strategic agents, while in this work we consider both the non-strategic case and the strategic one.

3. FORMAL MODEL

Definition 1 (Setting)**.** An instance of the *Stable Group Scheduling Problem* (SGSP) is a tuple (N, M, P) where $N = \{a_1, a_2, \ldots, a_n\}$ is a set of n agents, $M = \{t_1, t_2, \ldots, t_m\}$ is a set of time slots, and P is a collection of preferences of agents $(P = (P_1, P_2, \ldots, P_n))$. For each agent a_i, P_i is a total preorder ($\succeq_i$) on the set of alternatives, $X = X_0 \cup \{x_\emptyset\}$, where $X_0 = (M \times \{1, 2, \ldots, n\})$ and $x_\emptyset$ is the outside option of not attending; for any alternative $x \in X$, $(t, k) \succeq_i x$ is interpreted as agent a_i weakly preferring attending the event at time t if k attendees are present (including herself) to the alternative x (and similarly for $x \succeq_i (t, k)$).

We set $A_i = \{(t, k) \in X_0 | (t, k) \succ_i x_\emptyset\}$ and say that agent a_i approves of all alternatives in A_i. When each agent is indifferent among all alternatives in A_i, we call the problem a *simple Stable Group Scheduling Problem* (s-SGSP).

Definition 2 (Schedule)**.** A *schedule* for an instance (N, M, P) is a pair (t, S_t) such that $t \in M$ and $S_t \subseteq N$, and is interpreted as the organizer chooses time t and invites a subset of agents, S_t. Note that $S_t = \emptyset$ is allowed in our definition. A schedule (t, S_t) is said to be *individually rational* if for every agent $a_i \in S_t$ it holds that $(t, |S_t|) \in A_i$. A schedule (t, S_t) is said to be *envy-free* if for every agent $a_i \notin S_t$ it holds that $(t, |S_t \cup \{a_i\}|) \notin A_i$. A schedule is *stable* if it is both individually rational and envy-free.

For each agent a_i, her preferences over schedules are easily induced by her preferences over alternatives.

Definition 3 (Single-peaked preferences)**.** Given an instance (N, M, P) of SGSP, the preferences of agent a_i are *single-peaked* (SPK) if for every fixed time slot $t \in M$ there exists an ideal number of attendees, $o_i^t \in \{1, \ldots, n\}$, such that for any $k_1 \leq k_2 \leq o_i^t$ it holds that $(t, k_2) \succeq_i (t, k_1)$ and for any $o_i^t \leq k_2 \leq k_1$ it holds that $(t, k_2) \succeq_i (t, k_1)$.

There are two important special cases of SPK-preferences: increasing preferences (INC-preferences) and decreasing preferences (DEC-preferences). Agent a_i is said to have an INC-preference with respect to time slot t if $o_i^t = n$ (this agent prefers maximizing attendees). Analogously agent a_i is said to have a DEC-preference with respect to t if $o_i^t = 1$. We assume that all agents have SPK-preferences with respect to all time slots; such an instance is called an SPK-instance of SGSP (and analogously for INC and DEC-instances).

4. RESULTS

4.1 The Non-Strategic Case

Theorem 1 (Easiness results for SPK-instance of SGSP)**.** Given an SPK-instance of SGSP, there exists an algorithm that terminates in polynomial time, and decides whether a stable schedule exists; if one exists, then the algorithm produces a maximum one.

4.2 The Strategic Case

Theorem 2. It is impossible to design a strategy-proof mechanism that finds a stable schedule, even if we restrict the problem instance space to DEC-instances of s-SGSP with just two agents ($|N| = 2$) and one time slot ($|M| = 1$).

It is also impossible to design a strategy-proof mechanism that finds a stable schedule, even if we restrict the problem instance space to INC-instances of SGSP with juts two agents ($|N| = 2$) and two time slots ($|M| = 2$).

We note that our impossibility results hold even for randomized mechanisms. While we have negative results for the general problem, a strategy-proof mechanism does exist for INC-instances of s-SGSP.

Theorem 3. There exists a (deterministic) strategy-proof mechanism that finds a maximum stable schedule in polynomial time, given an INC-instance of s-SGSP.

5. CONCLUSIONS

The main contribution of this work is twofold: the proposal of a formal model for the stable group scheduling problem (SGSP) and the answer to the question of how hard it is to determine whether a stable schedule exists. We showed easiness results for the non-strategic case, while showed impossibility results for the strategic case in general.

We plan to extend our work in several ways. One direction is to allow each agent to specify her preferences over time slots (in addition to the number of attendees). Another direction is to take the identities of attendees into account. In realistic situations an agent may have a set of constraints such as “I will not attend if some agent x attends” or “I will only attend if some agent y attends as well”. We believe such extensions will make our model more realistic and applicable in implementing multi-agent scheduling systems.

6. ACKNOWLEDGMENTS

This work was funded in part by the National Science Foundation (under grants IIS-0963478 and IIS-1347214), the U.S. Army (under grant W911NF1010250), and the Kwanjeong Educational Foundation.

7. REFERENCES

[1] E. Crawford and M. Veloso. Learning to select negotiation strategies in multi-agent meeting scheduling. *Progress in Artificial Intelligence*, pages 584–595, 2005.

[2] A. Darmann, E. Elkind, S. Kurz, J. Lang, J. Schauer, and G. Woeginger. Group activity selection problem. In *Internet and Network Economics*, pages 156–169. Springer, 2012.

[3] E. Ephrati, G. Zlotkin, and J. Rosenschein. A non-manipulable meeting scheduling system. In *Proceedings of the 13th international workshop on distributed artificial intelligence*, pages 105–125, 1994.

[4] N. R. Jennings, A. J. Jackson, and L. E. Ns. Agent-based meeting scheduling: A design and implementation, 1995.

[5] P. Maes. Agents that reduce work and information overload. *Communications of the ACM*, 37(7):30–40, 1994.

[6] T. Mitchell, R. Caruana, D. Freitag, J. McDermott, D. Zabowski, et al. Experience with a learning personal assistant. *Communications of the ACM*, 37(7):80–91, 1994.