Cost of Cooperation for Scheduling Meetings

Alon Grubshtein and Amnon Meisels

Acknowledgements The research was supported by the Lynn and William Frankel Center for Computer Science,and by the Paul Ivanier Center for Robotics.

Abstract Scheduling meetings among agents can be represented as a game - the Meetings Scheduling Game (MSG). In its simplest form, the two-person MSG is shown to have a price of anarchy (PoA) which is bounded by 0.5. The paper defines the *Cost of Cooperation* (CoC) for meetings scheduling games, with respect to different global objective functions. For an "egalitarian" objective, that maximizes the minimal gain among all participating agents, the CoC is non positive for all agents. This makes the MSG a *cooperation game*. The concepts are defined and examples are given within the context of the MSG. A game may be revised by adding a mediator (or with a slight change of its mechanism) so that it behaves as a cooperation game. Thus, rational participants can cooperate (by taking part in a distributed optimization protocol) and receive a payoff which will be at least as high as the worst gain expected by a game theoretic equilibrium point.

1 Introduction

Scheduling meetings between two or more people is a difficult task. Despite the advances offered by modern world electronic calendars these are often limited and serve as passive information repositories. As a result, a dedicated person is usually hired to handle this task. Previous attempts to automate the meetings scheduling

Alon Grubshtein and Amnon Meisels
Dept. of Computer Science
Ben Gurion University of the Negev
Beer-Sheva, 84105, Israel
e-mail: {alongrub,am}@cs.bgu.ac.il

problem (MSP) employ one of two extreme approaches: the cooperative approach and the competitive one. Cooperative methods for solving MSPs perform a distributed search for an optimum of a global objective [5]. The competitive approach investigates game theoretic equilibria of suitable games and designs strategies for the competitive agents [1]. One previous attempt to combine the two approaches was introduced in [3], which empirically attempted to introduce selfishness into a cooperative protocol for the Meeting Scheduling Problem (MSP). Instead of searching for a solution which is socially optimal and potentially harmful for some agents, alternative global objectives are examined. Such solutions maintain an important goal - they provide acceptable gains to each one of the agents.

The present paper attempts to introduce cooperation into a self interested scenario by using a simplified meetings scheduling game (MSG). The use of game theory for studying self interested interactions provides a sound mathematical basis for the analysis of strategic situations under a set of (commonly) acceptable axioms (cf. [8]). A fundamental result of game theory is the existence of equilibrium points which define an action (or strategy) profile that is stable (i.e., no participant can gain by deviating from this profile). There are many definitions of various types of stable points, and unless otherwise specified we will assume the common "pure strategy Nash Equilibrium" (NE). A NE is simply a set of actions (assignments) for each agent in which no agent can gain by unilaterally changing its value [8].

Investigating the simple MSG, one shows that the NE of the MSG may be different than its optimal points. This is similar to routing games (cf. [11]). Consequently, the widely accepted efficiency measure, known as the *Price of Anarchy (Stability)* [4, 9, 12, 8, 11], for the simple MSG is different than unity. For scheduling meetings and its underlying game, we present user oriented criteria, and formulate a game property based upon it - the Cost of Cooperation (CoC). This property can motivate selfish users to cooperate by presenting *each one* of the players a guaranteed cooperative gain which is higher than its worst equilibrium value.

2 Meetings Scheduling by Agents

The Meeting Scheduling Problem (MSP) involves a (usually large) set of users with personal schedules and a set of meetings connecting these users. It is an interesting and difficult problem from many aspects. In the real world, agents have a dual role. They are expected to be both self interested (i.e. seek the best possible personal schedule) but also cooperative (otherwise meetings will not be scheduled).

MSP users have personal calendars which are connected by constraints that arise from the fact that multiple groups of the overall set of users need to participate in meetings. Meetings involve groups of users and each meeting includes at least two participants. Meetings can take place in different places and carry a different significance for their participants. All participants have their personal schedules and need to coordinate their meetings. Each user is represented by an agent and the set of agents coordinate and schedule all meetings of all users. The final result is a set of updated personal schedules for all agents/users (a global calendar). One uses utilities to represent the significance that users attach to meetings and to certain features of the resulting schedule, as it appears in their personal calendars.

Two very different approaches were studied for solving MSPs. The first employs Distributed Constraints Satisfaction (DCSPs) and Optimization (DCOPs) [5, 6, 14]. In these studies, agents are expected to be *fully cooperative* and follow the protocol regardless of their private gains. The outcome of such protocols is a solution whose **global utility** is optimal. The global utility does not necessarily account for the quality of the personal schedules of the participants.

An alternative approach is offered by researchers in the field of *Game theory*. Here, agents are rational, *self-interested* entities which have different (and often conflicting) utility functions (cf. [8]). A large share of the game theoretic research related to MSPs emphasizes the underlying mechanism of the interaction [2, 1, 10]. The basic assumption of these studies is that the gain of agents from their scheduled meetings can be represented by some universal currency. Moreover, the assumption accepts a uniform exchange rate for some monetary means and *unscheduled* meetings. These fundamental assumptions seem very unrealistic for scheduling meetings among people. However, if one is ready to accept this monetary model many game theoretic results and mechanisms can be applied.

The approach of the present study examines simple game theoretic mechanisms to restrict or predict agents behaviors.

3 The meetings scheduling game (MSG)

Imagine a pair of self-interested users, each with her own preferences about her personal schedule. Many agreements need to be reached in order to generate a viable and consistent schedule. One way to enforce agreements is to pose the problem as a game with rules that enforce schedules that are then accepted by the participants. In a way, this is analogous to routing by a player that accepts the delay as a given result of the route it selected. One can think of this as the mechanism of the game.

We begin by defining a simple MSG. The game includes two players coordinating a meeting. The meeting can be scheduled in the morning or in the evening. Each player places her bid for scheduling by specifying her preferences (or expected gains) from each schedule. This list of preferences takes the form of a tuple $b_i = \langle x,y\rangle \in \{0,1,...B\} \times \{0,1,...B\}$ for the two time slots, where the first value (x) corresponds to the evening time slot and the second number (y) corresponds to the morning time slot. A higher number indicates a higher gain from having the meeting at the corresponding time slot and therefore a stronger desire to meet at that time slot. Once both bids are presented, a time slot is selected by the game's mechanism according to its decision rule, and the participants are informed of the decision. There are many possible decision rules, but for now we use the "utilitarian" approach [8] - the time slot selected is the one with the highest sum of bids on it (or when sums are equal, outcomes have equal probabilities). In other words, the selected time slot has the highest total gain for the players. The payoff of each participant is defined by the player's utility function (preference) for a time-slot.

We limit ourselves to the case where players are never indifferent to the results, i.e., they will always prefer some time slot over the other. Any bid of the form $\langle x,x\rangle$ is thus excluded (it results in a scheduling based on the opponent's preference).

The first approximation of the MSG assumes that when the meeting is held at a time that is not preferred by a player, her gain from it will equal zero. We will later remove this assumption. We distinguish between the payoffs of players when the meeting is held at the player's most desired time, and say that player 1's payoff in such a case is m, while player 2's payoff for an analogous case is k. This enables us to treat users that are essentially different.

Figure 1 depicts the simplest possible MSG for two players. The descriptive power of a strategy is limited to two values of preference, either 0 or 1 (preferred). In this example player 1 prefers to have the meeting in the morning. When player 2 also prefers the morning time-slot both players have the same dominant strategy - $\langle 0,1\rangle$. That is, bidding $\langle 0,1\rangle$ will *always* yield a higher payoff for both players.

When player 2 prefers the evening time-slot her dominant strategy becomes $\langle 1,0\rangle$ (player 1's dominant strategy remains $\langle 0,1\rangle$). This can immediately be translated into an equilibrium point resulting from playing the dominant strategy. When the desires of both players are the same, the equilibrium payoff is m,k, and when these desires conflict the equilibrium payoff is $\frac{m}{2},\frac{k}{2}$ (i.e., expected values).

	$\langle 0,1\rangle$	$\langle 1,0\rangle$
$\langle 0,1\rangle$	m,k	$\frac{m}{2},\frac{k}{2}$
$\langle 1,0\rangle$	$\frac{m}{2},\frac{k}{2}$	$0,0$

(a)

	$\langle 0,1\rangle$	$\langle 1,0\rangle$
$\langle 0,1\rangle$	$m,0$	$\frac{m}{2},\frac{k}{2}$
$\langle 1,0\rangle$	$\frac{m}{2},\frac{k}{2}$	$0,k$

(b)

Fig. 1: Payoff matrices of a simple MSG. Player 1 prefers a morning meeting, and player 2 preferences can be either the same (a), or the opposite (b)

If a NE is the expected outcome of an interaction, a valid question to ask would be "How globally efficient or inefficient is this solution?". This question is the focus of a large body of work on "the Price of Anarchy" (PoA) [8, 9, 4, 12, 11]. The PoA is a measure of the inefficiency of an equilibrium point, and is defined as the ratio between the cost of the *globally worst* NE, to the cost of the *globally optimal* solution. Studies of the price of anarchy originate with routing games (cf. [11]).

It is interesting to see that the same question can be formulated for MSGs. What is the PoA for MSGs ? In reality, a player does not know her opponent's private preferences and this uncertainty should be reflected in our analysis. However, we will consider the above simple example as two different games for the sake of simplicity. In the first game the desires of both opponents align and in the unique equilibrium point the payoffs are m,k. The global utility in this case is $m+k$. Clearly this is also the optimal solution when considering a cooperative mechanism. Hence, $PoA=1$.

When the desires of the opponents conflict, the global gain in the unique equilibrium point is $\frac{m+k}{2}$. This is not necessarily the optimal solution: when $m < k$ the optimal solution's global utility is k. In this case the PoA moves further away from unity as k increases, [1] and we have: $PoA = \frac{m}{2k} + \frac{1}{2}$. That is, *the Price of Anarchy for the simplest MSG* is bounded by $\frac{1}{2}$.

This behavior is also expected in more realistic form of the two-players two time-slots MSG in which the players may express their strategies in terms of multiple preference values. In such a case, the limit on the maximal preference of each player - the value B - is greater than 1. In such cases the MSG has at least one equilibrium point and one can find the worst possible value that it can have. Next, comes the formulation of these results in the form of lemmas and the outlines of their proofs.

Lemma 1. *Every participant in the MSG described above has one (possibly weak) dominating strategy, or bid, which is composed of the value B for the preferred time slot, and* 0 *for the remaining time slot.*

The proof is simple and requires showing that if an agent assigns B to the preferred time slot and 0 otherwise (i.e., each player's action would be either $\langle B,0\rangle$ or $\langle 0,B\rangle$), her opponent can only force a draw, which will result in a fair toss of coin. If both players assign this action, no unilateral deviation will result in a higher payoff.

Lemma 2. *There are only two possible equilibrium outcome values to the above MSG,* (m,k) *and* $(\frac{m}{2},\frac{k}{2})$

By examining all five possible outcome values, one can easily rule out assignments which lead to $(0,0)$, $(m,0)$ or $(0,k)$. Thus we are left with (m,k) and $(\frac{m}{2},\frac{k}{2})$.

From these two simple lemmas, it is clear that when preferences coincide, the worst equilibrium payoff is (m,k). In this case the price of anarchy (PoA) is 1 [11]. In fact, this is the only equilibrium value due to the dominant strategies. When the preferences for time slot of the two players are in conflict, the worst equilibrium payoff of both players is $(\frac{m}{2},\frac{k}{2})$. If (without loss of generality) $m < k$, then the PoA decreases to a value lower than 1 as k increases (and is bounded below by $\frac{1}{2}$).

We now proceed to generalize this game, and allow players to assign any non-negative gain when their less preferred time slot is selected by the game mechanism. More specifically, we define for player one the payoff for a non optimal time slot as $0 \le x \le m$, and for player two $0 \le y \le k$, as demonstrated in the example in figure 2a.

The lemmas hold in this case (the tie's payoff is slightly revised). This can easily be understood by noticing that the best response to an opponent's strategy basically remains the same. Just as before, ties which result in a random selection by the mechanism (coin flip) and a value of $\frac{m+x}{2}$ or $\frac{k+y}{2}$, are always preferred over losing (x or y respectively).

As a result, when the preferences of the two agents are the same, the worst NE (also the best) has a value of $m+k$ which is also optimal - leading to a PoA of 1. When preferences contradict, the value of the stable point is $\frac{m+x}{2} + \frac{k+y}{2}$. This is not

[1] Unlike the routing minimization problem [11, 12], MSP is a maximization problem, hence the PoA is always smaller than 1.

	$\langle 0,1\rangle$	$\langle 0,2\rangle$	$\langle 1,0\rangle$	$\langle 1,2\rangle$	$\langle 2,0\rangle$	$\langle 2,1\rangle$
$\langle 0,1\rangle$	m,y	m,y	$\frac{m+x}{2},\frac{k+y}{2}$	m,y	x,k	$\frac{m+x}{2},\frac{k+y}{2}$
$\langle 0,2\rangle$	m,y	m,y	m,y	m,y	$\frac{m+x}{2},\frac{k+y}{2}$	m,y
$\langle 1,0\rangle$	$\frac{m+x}{2},\frac{k+y}{2}$	m,y	x,k	$\frac{m+x}{2},\frac{k+y}{2}$	x,k	x,k
$\langle 1,2\rangle$	m,y	m,y	$\frac{m+x}{2},\frac{k+y}{2}$	m,y	x,k	$\frac{m+x}{2},\frac{k+y}{2}$
$\langle 2,0\rangle$	x,k	$\frac{m+x}{2},\frac{k+y}{2}$	x,k	x,k	x,k	x,k
$\langle 2,1\rangle$	$\frac{m+x}{2},\frac{k+y}{2}$	m,y	x,k	$\frac{m+x}{2},\frac{k+y}{2}$	x,k	x,k

(a)

	$\langle 0,1\rangle$	$\langle 0,2\rangle$	$\langle 1,0\rangle$	$\langle 1,2\rangle$	$\langle 2,0\rangle$	$\langle 2,1\rangle$
$\langle 0,1\rangle$	m,y	m,y	$\frac{m+x}{2},\frac{k+y}{2}$	m,y	$\frac{m+x}{2},\frac{k+y}{2}$	m,y
$\langle 0,2\rangle$	m,y	m,y	$\frac{m+x}{2},\frac{k+y}{2}$	m,y	$\frac{m+x}{2},\frac{k+y}{2}$	m,y
$\langle 1,0\rangle$	$\frac{m+x}{2},\frac{k+y}{2}$	$\frac{m+x}{2},\frac{k+y}{2}$	x,k	x,k	x,k	x,k
$\langle 1,2\rangle$	m,y	m,y	x,k	m,y	x,k	$\frac{m+x}{2},\frac{k+y}{2}$
$\langle 2,0\rangle$	$\frac{m+x}{2},\frac{k+y}{2}$	$\frac{m+x}{2},\frac{k+y}{2}$	x,k	x,k	x,k	x,k
$\langle 2,1\rangle$	m,y	m,y	x,k	$\frac{m+x}{2},\frac{k+y}{2}$	x,k	x,k

(b)

Fig. 2: Payoff matrices for a more general MSG with different mechanisms - "utilitarian" (left) and "egalitarian" (right). Also note that B=2, and that player 1 prefers a morning meeting, and player 2 prefers an evening meeting.

necessarily optimal. For example, if $\frac{m-x}{2} < \frac{k-y}{2}$ than the optimal result can be $x+k$, resulting in

$$PoA = \frac{m+y}{2(x+k)} + \frac{1}{2}$$

Combining these results we reach several intermediate conclusions. The first is that when players have contradicting preferences, the PoA is bounded below by $\frac{1}{2}$, and depends on the relationship between the players payoffs. The second is that when the two players have the same preferences, the PoA is 1. Our analysis also indicate that the PoA in the MSG depends on several factors:

- The agents' private payoffs (i.e. m,k,x,y). This motivates the use of a mechanism which maps payoffs to a uniform or universal scale which can further bound this value. Such a scale can be the quality of a schedule as described in [3]. Indeed, when the individual payoffs are equal, the PoA becomes unity.
- The *mechanism* of the MSG. Figure 2 depicts two possible mechanisms: a "utilitarian" mechanism (Figure 2a), and an "egalitarian" one (Figure 2b) which selects a time slot that maximizes the minimal bid on it. This leads to a different game, with a different NE.
- The PoA may also be affected by the designer's perception of optimality. For example, given the MSG of figure 2a, solutions yielding the gain $\langle x,k\rangle$ are optimal when considering the "utilitarian" approach (maximal combined payoffs). However, if one's perception of optimality is that it maximizes the lowest gain (e.g.,"egalitarian"), then the optimal solution gain for this MSG (Figure 2a) is $\langle \frac{m+x}{2},\frac{k+y}{2}\rangle$. Optimizing different objective functions can produce substantially "better" results [3].

The MSG toy example can become more realistic along several dimensions:

1. The maximal preference bid (B) - while this value increases the assignment space of each player, we have shown that it does not affect the behavior of the two-players, two-time-slots game.
2. The number of participants. Although not presented here, adding more players does not change the fact that these have a (possibly weak) dominant strategy.

3. The number of time slots. Adding more time slots requires several modifications to the game described above (for example, in case of a tie, will the mechanism randomly pick any of the possible time slots? just those that were tied? Can two time slots receive the same bid value? etc). Even after adding the required adjustment, finding the PoA will still rely on the utility that each agent associates with a time slot. In the case of three time slots, if each agent's utility from a meeting will be x for the least preferred time slot, z for the most preferred time slot, and y for the remaining time slot, one can state that if $x+z \geq 2y$ then a (weak) dominant strategy exists for the players (under some set of rules).
4. The number of meetings. Note that the natural scenario of incremental addition of meetings defines a different game which is an iterated game.

4 The Cost of Cooperation

Following this detailed tour of the Price of Anarchy for MSGs, an interesting question arises: what exactly is the meaning of the PoA for MSGs? In other words, what is the meaning of "anarchy" in the context of MSG. After all, "anarchy" for an MSG is the natural situation because agents want their best possible personal calendar. The "cooperative" method of optimizing some global goal is not very natural. Why would self interested users wish to optimize some global objective ? There is no sense of "public savings" or "minimization of overall cost", as in the case of routing games. Players in a MSG are only interested in cooperation as a form of mediation between them. This implies that examining efficiency in terms of some global value does not fit the agent's point of view.

The present paper proposes a different measure. Instead of addressing the "social" cost, the focus here is to quantify the cost (or gain) of a single agent, from participating in any cooperative protocol.

Definition 1. An agent's *Cost of Cooperation*(CoC) with respect to $f(x)$ is defined as the difference between the lowest gain the agent can have in any of the NEs of the underlying game (if any, otherwise zero) and the lowest gain an agent receives from a (cooperative) protocol's solution x that maximizes $f(x)$.

The intuition behind this definition is quite simple. A user about to interact with other users can *expect that the combined spontaneous action taken by all participants will result in a NE*. Since there can be more than one such stable points the user anticipates the worst case for herself, i.e., the personal lowest gaining NE. For her to join a cooperative protocol (and use a multi user distributed optimizing protocol), she uses knowledge about her possible payoffs. This information is provided by the CoC. When the CoC value is negative, the user knows that she *stands to gain* from joining the cooperative protocol.

As a simple example, consider the famous Prisoners' Dilemma game as depicted in Figure 3. In this problem, a single stable action profile exists - $\langle D,D\rangle$. When both participants select this action, each has a gain of 1. In the optimal solution that maximizes the overall gain (e.g., the *sum of all personal gains*) - $\langle C,C \rangle$ - the gain for each agent is 4. This corresponds to a CoC of -3 for each participant, which implies that both agents gain from following this cooperative solution.

	Cooperate	Defect
Cooperate	4, 4	0, 6
Defect	6, 0	1, 1

Fig. 3: The Prisoners' Dilemma game matrix

The $f(x)$-CoC for a set of users is defined as a tuple, in which every element CoC_i corresponds to the CoC of agent A_i. Defining the CoC for an entire system can be used to motivate the participation of users in the cooperative protocol. A negative CoC (i.e., gain), provides a self-interested user with a strong incentive to join the cooperative protocol. One may even venture to call this *Individual Rationality* [8].

4.1 Cooperation Games

Let us define special strategic situations in which the CoC vector of all participants has non positive values.

Definition 2. A game is an $f(x)$-*Cooperation game*[2] if there exists a solution to the game in which the CoC (with respect to $f(x)$) of all agents is non positive.

For example, the simple MSG of Figure 2a, is a cooperation game with respect to a Max-Min objective function (i.e., maximizing the minimal gain).

Given a general MSG, one may want to change it into a cooperation game (if it is not one already). This can be achieved by defining the cooperative goal (e.g., the objective function) in a specific manner, by changing the mechanism of the interaction itself, or by adding an interested party to the interaction. Some work in the direction of the latter, was recently reported in [13, 7]. *Mediators* are introduced as parties wishing to influence the choice of action (e.g., strategy) of participants which are not under their control. Mediators cannot enforce payments by the agents or prohibit strategies, but can influence the outcome of an interaction by offering to play on behalf of some (or all) of the participants. By doing so a mediator commits to a pre-specified behavior [13].

An interesting property of Routing mediators [13] is that they are capable of possessing information about the actions taken by agents. In a two player strategic situation, one may use the following routing mediator to generate a stable solution: the revised game includes mediated actions and each agent may either play its game as before, or let the mediator play for it. When the mediator plays for an agent, it will

[2] Not to be confused with cooperative games

always choose to assign the bid which will result in the minimal payoff to the agent's *opponent*. If both agents use the mediator, the mediator assigns the interaction which results in the lowest value that is at least as high as any value that results from the play of the agent's opponent.

Two important attributes of the new mediated interaction are the existence of at lease one, new, pure strategy NE (the assignment $\langle Med, Med \rangle$), and the existence of an action profile improving the gains of all agents (i.e. it is now a cooperation game). For example, consider the two agents interaction depicted in figure 4a. In this interaction, two agents must choose between playing U or D, L or R. The payoffs values are not specified, but their order is: $A_1 > A_2 > ... > A_8$.

	L	R
U	$\underline{A_6}, A_7$	$A_3, \underline{A_5}$
D	$A_8, \underline{A_2}$	$\underline{A_1}, A_4$

(a)

	L	R	Med
U	$\underline{A_6}, A_7$	$A_3, \underline{A_5}$	A_6, A_7
D	$A_8, \underline{A_2}$	$\underline{A_1}, A_4$	$A_8, \underline{A_2}$
Med	$\underline{A_6}, A_7$	$A_3, \underline{A_5}$	$\underline{A_3}, \underline{A_5}$

(b)

Fig. 4: A two agent strategic situation with ordinal payoffs - with and without a mediator

Our former restriction to pure strategies now results in a scenario which has no stable points. This means that in every solution one player stands to gain from deviating. The underlined values in figure 4a represent the best response that each agent has, in view of her opponents strategy. For example, if the first agent selects the U strategy when her opponent plays L, the second agent will change its assignment to R (resulting in D by the first agent and L again by player two). Going back to our former line of reasoning, this implies that an agent does not know what to expect from such an interaction, and any gain is plausible. However, by adding our previously described mediator, this situation is improved, as depicted in figure 4b.

When playing for the columns agent, the mediator "threatens" the row player - it always picks the move L, with the worst outcome for her (A_6, A_8 instead of the possible A_3, A_1). The opposite is true for the row player. Here if the mediator is selected it plays U so that it guarantees the worst outcome for the columns player (A_7, A_5). However, when both players use the mediator, it selects the action profile $\langle U, R \rangle$. This assignment is used because it is a valid assignment for which the column's payoff is A_5 (which is not worse than A_7 and A_5), and the row's payoff is A_3 (better than either A_6 or A_8). The end result is a transformation of the original game (by introducing a mediator) into a cooperation game. That is, if agents choose to participate in a utilitarian optimization protocol (cooperate) the end result will be $\langle D, R \rangle$, with a payoff of A_1 to agent one and A_4 to agent two. This improves the expected gain of *each one of the agents* from playing the game in Figure 4b.

5 Discussion

The paper discusses two opposite extreme approaches that are inherent to many multi agent scenarios - cooperation and competition. In order to investigate these a a simple scheduling problem was formulated in the form of a simple game and analyzed. The maximal preference bid and general payoffs are incrementally added and their impact on the stable points and the PoA of the interactions are examined. Our analysis leads naturally to a new measure that quantifies the cost/gain to agents from participating in a cooperative protocol. The Cost of Cooperation (CoC) as defined in the paper further defines a game property that is dubbed "Cooperation game". Participants in a cooperation game may be better off cooperating than playing (selfishly) out the game.

References

1. E. Crawford and M. Veloso. Mechanism design for multi-agent meeting scheduling. *Web Intelligence and Agent Systems*, 4(2):209–220, 2006.
2. E. Ephrati, G. Zlotkin, and J.S. Rosenschein. A non manipulable meeting scheduling system. In *Proc. Intern. Workshop on Distributed Artificial Intelligence*, Seatle, WA, 1994.
3. A. Gershman, A. Grubshtein, A. Meisels, L. Rokach, and Roie Zivan. Scheduling meetings by agents. In *Proc. 7th Intern. Conf. on Pract. & Theo. Automated Timetabling (PATAT 2008)*, Montreal, August 2008.
4. Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *in Proc. of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, pages 404–413, 1999.
5. R. T. Maheswaran, M. Tambe, E. Bowring, J. P. Pearce, and P. Varakantham. Taking dcop to the real world: Efficient complete solutions for distributed multi-event scheduling. In *Proc. 3rd Intern. Joint Conf. on Autonomous Agents & Multi-Agent Systems (AAMAS-04)*, pages 310–317, NY, New York, 2004.
6. J. Modi and M. Veloso. Multiagent meeting scheduling with rescheduling. In *Proc. 5th workshop on distributed constraints reasoning DCR-04*, Toronto, September 2004.
7. Dov Monderer and Moshe Tennenholtz. Strong mediated equilibrium. *Artificial Intelligence*, 173(1):180–195, 2009.
8. N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
9. Christos Papadimitriou. Algorithms, games, and the internet. In *STOC '01: Proc. of the 33rd annual ACM symposium on Theory of computing*, pages 749–753, 2001.
10. Adrian Petcu, Boi Faltings, and David Parkes. M-DPOP: Faithful distributed implementation of efficient social choice problems. *Journal of AI Research (JAIR)*, 32:705–755, 2008.
11. Tim Roughgarden. *Selfish Routing and the Price of Anarchy*. The MIT Press, 2005.
12. Tim Roughgarden and Éva Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, 2002.
13. Ola Rozenfeld and Moshe Tennenholtz. Routing mediators. In *IJCAI*, pages 1488–1493, Hyderabad, India, 2007.
14. R. J. Wallace and E. Freuder. Constraint-based multi-agent meeting scheduling: effects of agent heterogeneity on performance and privacy loss. In *Proc. 3rd workshop on distributed constrait reasoning, DCR-02*, pages 176–182, Bologna, 2002.