UDC 519.876.3, DOI:10.2298/CSIS100118012G

Cost of Cooperation for Scheduling Meetings

Alon Grubshtein and Amnon Meisels

Dept. of Computer Science
Ben Gurion University of the Negev
P.O.B 653, Beer-Sheva, 84105
Israel
{alongrub,am}@cs.bgu.ac.il

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 PoA bound provides a measure on the efficiency of the worst Nash Equilibrium in social (or global) terms. The approach taken by the present paper introduces the *Cost of Cooperation* (CoC) for games. The CoC is defined with respect to different global objective functions and provides a measure on the efficiency of a solution *for each participant* (personal). Applying an "egalitarian" objective, that maximizes the minimal gain among all participating agents, on our simple example results in a CoC which 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.
Although not all games are cooperation games, 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. 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.

Keywords: Multi-Agent Systems, Cooperation, Meeting Scheduling.

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 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 [8]. The competitive approach investigates game theoretic equilibria of suitable games and designs strategies for the competitive agents [2]. One previous attempt to combine the two approaches was introduced in [4], which empirically attempted to introduce selfishness into a cooperative protocol for the Meetings Scheduling

Problem (MSP). Instead of searching for a solution which is socially optimal and potentially harmful for some agents, alternative global objectives were examined. Such solutions maintain an important goal - they provide acceptable gains to each one of the agents [4].

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. [12]). 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" (PSNE, or simply NE). A NE is a set of actions (assignments), one for each agent, in which no agent can gain by unilaterally changing its value [12].

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. [15]). Consequently, the widely accepted efficiency measure, known as the *Price of Anarchy (Stability)* [7, 13, 16, 12, 15], is different than unity for the simple MSG. 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 of such an interaction is a set of updated personal schedules for all agents/users (a global calendar). In this work we use utilities to represent the significance that users attach to meetings and to certain features of the resulting schedule.

Two very different approaches were studied for solving MSPs. The first employs Distributed Constraints Satisfaction (DCSPs) and Optimization (DCOPs) [8, 10, 19]. In these studies, agents are expected to be *fully cooperative* and follow the protocol (algorithm) regardless of their private gains. The outcome of such protocols is a solution whose **global utility** is optimal (or consistent in the case of a DCSP). The global utility does not necessarily account for the quality of the personal schedules of the participants. Indeed, due to the combinatorial nature of the MSP it is often the case that a globally optimal solution includes a low quality solution for at least one participant [4].

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. [12]). A large share of the game theoretic research related to MSPs emphasizes the underlying mechanism of the interaction [3, 2, 14]. 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. One important example is that mechanisms for discouraging cheating, that are known for combinatorial auctions, can possibly be used for agents that schedule meetings [14].

The approach of the present study examines simple game theoretic mechanisms to restrict or predict agents behaviors in an inherently cooperative environment.

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.

An equilibrium point is defined with respect to a given game, so 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 bids 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 morning time slot and the second number (y) corresponds to the evening time slot. A higher number indicates a higher preference towards having the meeting at the corresponding time slot (a stronger desire to meet at that time slot). Note that the bids made by the players do not necessarily represent a player's real preference. That is, players may act strategically and attempt to "deceive" their opponent into an action which may result in a better or stable personal outcome. As we shall soon demonstrate, in the case of our simple example, bidding "truthfully" is a dominating action (the action that results in the best payoff for every action taken by the opponent).

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 [12] - 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 (reported) gain for the players. The end 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 "rows" player (player 1) and the "columns" player (player 2). The outcome of a joint bid is composed of two values: the left hand value represents player 1's gain while the right hand value represents player 2's gain from the joint action. 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 - $\langle 1, 0\rangle$. That is, bidding $\langle 1, 0\rangle$ will *always* yield a higher payoff for both players. Such a strategy profile is said to be stable since no player has an incentive to deviate from her bid.

When player 2 prefers the evening time-slot her dominant strategy becomes $\langle 0, 1\rangle$ (player 1's dominant strategy remains $\langle 1, 0\rangle$). 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).

If a NE, or a stable point, 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) [12, 13, 7, 16, 15]. 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

p_1 \ p_2	$\langle 0,1\rangle$	$\langle 1,0\rangle$
$\langle 0,1\rangle$	$0,0$	$\frac{m}{2},\frac{k}{2}$
$\langle 1,0\rangle$	$\frac{m}{2},\frac{k}{2}$	m,k

(a)

p_1 \ p_2	$\langle 0,1\rangle$	$\langle 1,0\rangle$
$\langle 0,1\rangle$	$0,k$	$\frac{m}{2},\frac{k}{2}$
$\langle 1,0\rangle$	$\frac{m}{2},\frac{k}{2}$	$m,0$

(b)

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

NE, to the cost of the *globally* optimal solution. Studies of the price of anarchy originate with routing games (cf. [15]).

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. If one is to examine the overall satisfaction of a solution then the global utility in this case is $m + k$. Clearly this is also the optimal solution's utility when considering a cooperative mechanism (optimizing the sum of gains depicted in the payoff bi-matrix). 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 forms 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. The resulting 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 dominating strategy, or bid, which is composed of the value B for the preferred time slot, and 0 for the remaining time slot.*

[1] Unlike the routing minimization problem [15, 16], MSP is a maximization problem, hence the PoA is always smaller than 1. This may lead to some confusion with readers already familiar with the PoA but is consistent with previous work

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 at most 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, and noting that a player can always impose a draw, one can easily rule out unstable 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. In fact, this is the only equilibrium value due to the dominance of 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 2.

p_1 \ p_2	$\langle 0,1\rangle$	$\langle 0,2\rangle$	$\langle 1,0\rangle$	$\langle 1,2\rangle$	$\langle 2,0\rangle$	$\langle 2,1\rangle$
$\langle 0,1\rangle$	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

Fig. 2: Payoff matrix for a more general MSG with B=2. Player 1 prefers an evening meeting, and player 2 prefers a morning meeting.

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 (which is 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 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 conclusion is that when the two players have the same preferences, the PoA is 1. Our analysis also indicates that the PoA for 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 [4]. Indeed, when the individual payoffs are equal, the PoA becomes unity.
- The *mechanism* of the MSG. While Figure 2 depicts an MSG with a "utilitarian" mechanism an "egalitarian" one which selects a time slot that maximizes the minimal bid is depicted in Figure 3. The different mechanisms leads to different games, with different stable solutions.
- The PoA may also be affected by the *designer*'s perception of optimality. For example, given the MSG of Figure 2 and $\frac{m-x}{2}<\frac{k-y}{2}$, the solutions yielding the gain $\langle x, k\rangle$ are optimal when considering the "utilitarian" approach (maximal combined payoffs), and the resulting PoA is the same as described above. If, however, one's perception of optimality is that it maximizes the lowest gain (e.g.,"egalitarian"), then the optimal solution gain for this MSG (Figure 3) is $\langle \frac{m+x}{2}, \frac{k+y}{2}\rangle$ and the PoA becomes unity. Optimizing different objective functions can produce substantially "better" results [4].

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 still relies on the utility that each agent associates with a time slot.

p_1 \ p_2	$\langle 0,1\rangle$	$\langle 0,2\rangle$	$\langle 1,0\rangle$	$\langle 1,2\rangle$	$\langle 2,0\rangle$	$\langle 2,1\rangle$
$\langle 0,1\rangle$	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

Fig. 3: Payoff matrix for the same MSG as the one depicted in figure 2 with a different decision rule (an "egalitarian" mechanism).

4. The number of meetings. Note that the natural scenario of incremental addition of meetings defines a different game which is an iterated game.

p_1 \ p_2	$\langle 0,0,1\rangle$	$\langle 0,1,0\rangle$	$\langle 1,0,0\rangle$	$\langle 1,0,1\rangle$	$\langle 1,1,0\rangle$	$\langle 0,1,1\rangle$
$\langle 0,0,1\rangle$	x	$\frac{x+z}{2}$	$\frac{x+y}{2}$	x	$\frac{x+y+z}{3}$	x
$\langle 0,1,0\rangle$	$\frac{x+z}{2}$	z	$\frac{y+z}{2}$	$\frac{x+y+z}{3}$	z	z
$\langle 1,0,0\rangle$	$\frac{x+y}{2}$	$\frac{y+z}{2}$	y	y	y	$\frac{x+y+z}{3}$
$\langle 1,0,1\rangle$	x	$\frac{x+y+z}{3}$	y	$\frac{x+y}{2}$	y	x
$\langle 1,1,0\rangle$	$\frac{x+y+z}{3}$	z	y	y	$\frac{y+z}{2}$	z
$\langle 2,1\rangle$	x	z	$\frac{x+y+z}{3}$	x	z	$\frac{x+z}{2}$

Fig. 4: Payoff matrix of the row player (p_1) in a three time slot max-sum MSG. The player's ordinal preferences are $evening > morning > noon$

As an example of a simple change consider a three time slot MSG, with a Max-Sum mechanism. Let us assume that users are allowed to add any bid representing different preferences (that is, bids which have at least two different values such as $\langle B,0,0\rangle$). Let us also assume that in the case of a tie over several time slots, the mechanism uniformly selects one of the tied slots. Figure 4 depicts the gains of the row player (p_1) in such a game. The player prefers to have the meeting in the evening, and if that is not possible she would rather have it in the morning. She least prefers to have it at noon. That is, the agent's utility from a meeting will be z for the least preferred time slot, x for the most preferred time slot, and y for the remaining time slot.

It can be shown that if both of the following conditions are satisfied for such an MSG:

- $x + z > 2y$
- $x + y > 2z$

Then, the bid assigning B to the preferred time slot and 0 to the remaining time slots (denoted b) is a (possibly weak) dominant strategy for the players.

Proof. Similarly to our previous lemmas, one can see that any player playing the suggested bid will divert the game towards some sort of tie involving her most preferred time slot, or towards the time slot itself (her opponent will never exceed the value B on any of her other time slots). Following our requirements on payoffs we can sort the possible outcomes:

$$z < \frac{y+z}{2} < y < \frac{x+y+z}{3} < \frac{x+z}{2} < \frac{x+y}{2} < x$$

Since the highest outcomes are those involving the required time slot, to complete our proof one needs only to show that for any given bid by the opponent, playing the suggested bid will result in the highest outcome (or at least as high as any other outcome). This, however, is immediate by the bids structure: If a three slots tie exist, any bid other than b will break the tie resulting in a worst outcome for the player (one of its less preferred time slot). Similarly, if a two slots tie exists then changing b will either result in a three slots tie (which produces a lower quality solution for the player), or worst yet, it will result in a less preferred "pure" outcome.
Thus, playing b is always a (possibly weak) dominant strategy in the MSG described above.

One can think of a "real" MSG as a graphical game [6], in which each participant is connected with others who have at least one common meeting with her. While this formulation provides a more detailed view of a large scale interaction, it provides little theoretical insights. Furthermore, finding an optimum of such a graphical game is a task which received little attention until only recently [5].

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 MSGs. 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 an 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 5. 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. If one examines the socially 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.

p_1 \ p_2	Cooperate	Defect
Cooperate	4, 4	0, 6
Defect	6, 0	1, 1

Fig. 5: The Prisoners' Dilemma game matrix

We define the $f(x)$-CoC for a set of users 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 think of this as a form of *Individual Rationality* [12].

4.1. Cooperation Games

Let us define special strategic situations in which the CoC vector of all participants has non positive values (we shall refer to these as CoC solutions).

Definition 2. *A game is a* Cooperation game[2] *if there exists a solution (strategy profile) to the game in which each player's payoff is as high as its worst equilibrium payoff. That is, there exists a solution for which the CoC (with respect to some $f(x)$) of all agents is non positive.*

Note that if there exists a solution to the game in which each player's payoff is as high as its worst equilibrium payoff, the objective function one may optimize will simply accept solutions in which the gains of each agent are higher than the worst NE gains, and reject all other solutions.

As an example consider the simple MSG of Figure 2. It is a cooperation game. If a cooperative mechanism optimizing the Max-Min objective function is used, all agents are guaranteed to have a non positive CoC (although this scenario may be treated as a degenerate one - all CoC values are zero).

Our definition of a cooperation game implies that any game with at least one PSNE is also a cooperation game (one can always take one of the worst NEs and present it as a solution with non positive CoC). However, we will usually focus on games in which the CoC is negative for all players.

It is important to note that a CoC solution does not imply stability. This can easily be understood by examining the prisoners' dilemma in Figure 5. An all-non-positive CoC solution in this case corresponds to the tuple $\langle C, C \rangle$. However, game theory does not treat this solution as a viable one since each player has a strong incentive to deviate from it. If on the other hand, players are treated as agents which follow a protocol (may not deviate from it) one need not worry about the stability of the solution. In this case, the agents' underlying users simply accept the cooperative framework since it guarantees a solution which is at least as good as that achieved by "playing out" the interaction *for each and every one of them*.

Figure 6 provides a visual representation of these ideas. In this figure, the space of possible outcomes of a given game is drawn. Each dot represents a possible outcome of the game along with the payoff each participant associates with that outcome. The social choice solution which maximizes the overall gain to participants is indicated in the bottom right corner ("Max-Sum"), while the set of all pareto efficient solutions is indicated by a line connecting these. A non positive CoC solution is any outcome situated in the colored area (it is easy to see why there always exists at least one pareto efficient solution within this area).

Although every game admits at least one (possibly mixed) stable point some games are not cooperation games. Consider for example the game in Figure 7. In this game, there exists only one stable solution (the mixed strategy $\langle \frac{1}{3}, \frac{1}{2} \rangle$). The expected outcome of each player in this mixed solution results in a payoff of 3 to each participant. However, there is no outcome with the desired property of non positive CoC to both players in this game.

Given a general game, one may want to change it into a cooperation game (if it is not one already). This can be achieved by changing the mechanism of

[2] Not to be confused with cooperative games

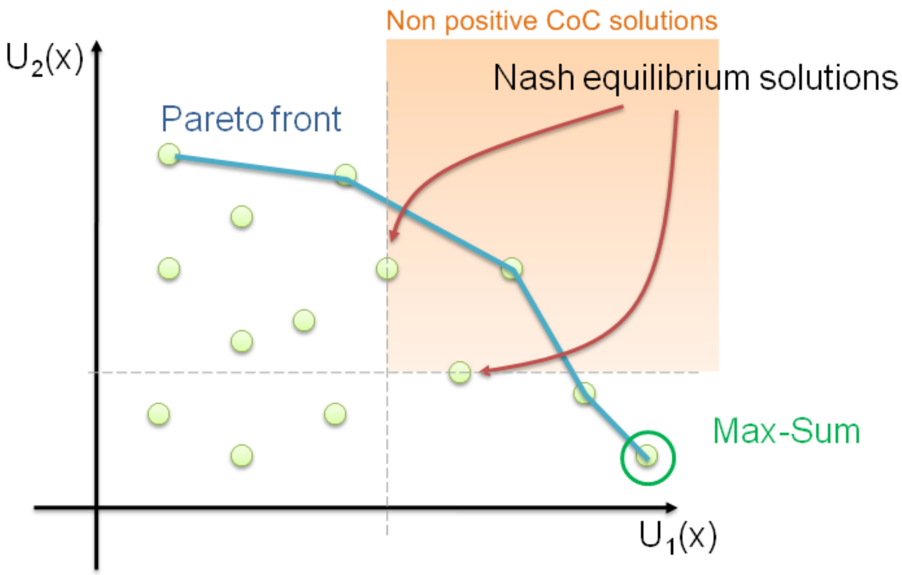

Fig. 6: A cooperation game

p_1 \ p_2	Left	Right
Up	2, 5	4, 1
Down	6, 1	0, 3

Fig. 7: A simple game. This game is not a cooperation game - the worst (expected) equilibrium outcome for each player is 3 and $\frac{7}{3}$, and no outcome can satisfy a non positive CoC tuple.

the interaction itself, or by adding an interested party to the interaction (that is, modify the interaction to introduce at least one PSNE). Some work in the direction of the latter approach that uses mediators was recently reported in [17, 11]. *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 [17].

An interesting property of Routing mediators [17] is that they are capable of possessing information about the actions taken by agents. While the authors of [11] apply mediators for the sake of a strong equilibrium, in our two player strategic situation, one may use the following routing mediator to generate a simple PSNE. The mediated game includes mediated actions and each agent may either play its game as before, or let the mediator play for it.

The following mediator is then used: the mediator may receive a message from the user asking it to play for her or not. When the mediator plays for an agent, it will 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.A player may always choose not to use the mediator (apply one of the original actions available to her).

p_1 \ p_2	Left	Right	Med
Up	2, 5	4, 1	2, 5
Down	6, 1	0, 3	0, 3
Med	6, 1	4, 1	4, 1

Fig. 8: A mediated version of the game in Figure 7

Figure 8 depicts the result of applying the above mediator to the game of Figure 7. In the new mediated game there are three new PSNE - $\langle Med, Left\rangle$, $\langle Med, Right\rangle$ and $\langle Med, Med\rangle$. Furthermore, since the new worst gains for the participant are 4 to player 1, and 1 to player two, the solution $\langle Down, Left\rangle$ has a non positive CoC tuple (player 1 gain's from the interaction is higher than her "expected" NE gain, while player 2's gain is not lower than her expected worst gain). The proposed mediator will always add at least one PSNE:

Theorem 1. *When applying the mediator described above to a two player game, the transformed interaction will include at least one PSNE in the strategy profile $\langle M, M\rangle$.*

Proof. The mediator described above always uses the lowest value to the opponent. Thus, the "M" column (row) is based on the lowest value of the row (column) player. Once the values are inserted to each cell but the last - $\langle M, M\rangle$ - one needs to show that there exists a joint action which has a left hand value as high as the highest value in all the left hand values of the "M" column, and a right hand value as high as the highest right hand value in the "M" row.

Let r be the row in the "M" column with the highest left hand value, and let us denote that value with x. Similarly, let c be the column in the "M" row with the highest right hand value - y. By the construction of the "M" column we know that x is the lowest value in row r. The same is true for y - according to the construction of row "M" we know that it is the lowest value in the column c. Examining the values in the intersection of row r and column c we conclude that the left hand value must be at least as high as x and the right hand value must be at least as high as y.

Since x and y are the highest values in the "M" column and row (respectively), the left (right) hand value of $\langle r, c\rangle$ is higher than all other values in the "M" column (row). As a result, adding these values to the cell indicated by $\langle M, M\rangle$, we receive a PSNE in that joint interaction.□

4.2. CoC solutions

As mentioned in the previous section, a CoC solution does not imply stability and should not be confused with solution concepts such as Aumann's strong Nash equilibrium [1]. A strong Nash equilibrium strives to achieve a solution which is deviation proof by any possible coalition of players. While such a concept has many desired properties it requires agents to rationalize over outcomes, does not provide any guarantees to individual players (except having lower gains in case of a deviation) and rarely occurs in most games. A CoC solution on the other hand is unstable in the game theoretic sense but as demonstrated in the previous section often exists in a multi agent interaction.

The approach taken by the present paper introduces cooperation, or a cooperative protocol, to a self interested interaction. This approach is different than that taken by previous works. Specifically, the attempt to introduce a party to a given game is extensively discussed in [11]. Although the mediator presented by Monderer and Tennenholtz is limited, its presence alters the initial interaction (as demonstrated). The changes introduced by a mediator alter the very course of the game. While a simple MSG includes two phases: action selection and payoff collection the introduction of a mediator can be translated as an additional interim phase. In this phase, the players wishing to employ the mediator notify it, a calculation is made by the mediator, and this information is sent onward to the game mechanism along with the rest of actions selected by those players who have not used the mediator. The game's mechanism then selects the time of the meetings. Using a cooperative protocol on the other hand is equivalent to stating that the players have no knowledge on how to choose their actions and are only interested in a solution which provides some guarantees to the quality of the solution. As such, they request the game mechanism to make its choice for them.

Providing guarantees to participants was also discussed in [18]. There, the author defines a mixed strategy t as a C competitive safety level strategy if the ratio between the expected payoff of the agent in a Nash equilibrium to its expected payoff in t is bounded by C. While we share the idea of providing a guarantee to participating agents, our approach examines a cooperative, multi agent approach which does not require stability. This is in contrast to the competitive safety level strategy introduced in [18] which provides a sound recommendation to each one of the participants separately.

Finally, it should be noted that the concept of CoC defines a subspace of the original search space. In this subspace (colored area of Figure 6) there may be several CoC solutions. That is, one may use existing solution concepts (such as the pareto front) over this area and still provide a strong incentive for players to join.

5. Related work

Research into distributed cooperative optimization has seen a growing interest in the past decade. Specifically, the large body of work on distributed algorithms for Distributed Constraint Optimization Problems (DCOPs) (cf. [9]) makes this approach suitable and widely accepted for formulating and solving many diverse MAS problems ([4, 20, 8]). DCOP agents exchange messages pretaining to their state and their valuation of it. The end goal of the agents is a globally optimal solution (a set of all agents' states). These agents are cooperative and will refrain from changing their state even if some consider it sub optimal (i.e., can improve their personal gain by changing their local state). A recent paper introduced a variation of DCOPs - Asymmetric DCOPs (ADCOP) [5] - which may be even better suited to problems such as the one discussed in our work. ADCOPs form a richer framework for cooperative optimization and allows search in game-like structures such as the one presented in Figures 4 and 8.

Cooperative optimization provides a rich set of algorithms for solving many problems, however, it is deemed ill fitting to many problems in which agents are mainly interested in their personal gain. In such cases, Game Theory is invoked (cf. [12]). It allows one to mathematically define a stable solution to an interaction involving a set of rational self interested agents. However, there has been very little work on distributed algorithms for finding such stable solutions [6].

The present work attempts to draw upon the advantages of both of these approaches. Although it considers a set of cooperative agents which may be thought of as a coalition (of all players), it is quite different from the coalitions discussed by Aumann's strong Nash equilibrium [1]. A strong Nash equilibrium strives to achieve a solution which is deviation proof by any possible coalition of players, but rarely exist. CoC solutions, on the other hand, provide no such guranatee, but are often present in many problems. Moreover, unlike other cooperative optimization approaches, CoC solutions maintain some form of Individual Rationality (IR).

An interesting approach is provided by Monderer and Tennenholtz in [11], who propose to alter a given interaction through a trusted third party - the mediator. Agents may communicate with the mediator which in turn provides some form of limited cooperation. However, unlike the present work which attempt to provide guarantees to participants, their work attempts to enrich the set of situations where stability against deviations by coalitions may be obtained.

Tennenholtz also attempts to provide guarantees to participants in [18]. Examining the ratio between the expected payoff in a NE to the expected payoff from a mixed strategy, a C competitive safety level strategy is defined. However, the end goal of this analysis provides a recommendation to each self interested participant. The method proposed by the present study assumes a cooperative protocol and is not limited to stable solutions.

The cooperative approach to MSPs attempts to maximize (minimze) a given objective function [19, 8, 10] and does not attempt to provide any guarantee on the outcome (or provide IR). In contrast, the competitive approach to MSPs attempts to design the mechanism of the interaction so that some stable solution is acheived [3, 2]. It assumes agents will interact in a selfish manner and end up in a stable state, although how agents may do so is not discussed and is often ignored. The method of the present paper is not limited to stable solutions, apply known cooperative algorithms, and provide some subjective guarantee to each participant.

6. Discussion

The present paper discusses two opposite extreme approaches that are inherent to many multi agent scenarios - cooperation and competition. In order to investigate these, a simple scheduling problem was formulated in the form of a simple game and analyzed. The maximal preference bid and general payoffs were incrementally added and their impact on the stable points and the PoA of the interactions were examined.

The analysis of a simple MSG 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 present paper further defines a game property that is termed “Cooperation game”. Participants in a cooperation game may be better off cooperating in a search protocol for an optimal solution than playing (selfishly) out the game.

These ideas may be applied to a large range of problems where participants are competitive but may need to reach some agreement. One may consider voting games, for example. In voting games, each agent has prefrences over the set of outcomes and a set of actions (votes for candidate solutions). A candidate solution is selected by the game mechanism based on a pre defined rule. The candidate solution which maximizes the overall utility (often referred to as the “social choice”) may also be one in which some agents receive a very low utility. Being self interested, agents will always choose to deviate from such solutions. On the other hand, stable solutions may result in a rather low overall satisfaction (utility). A CoC solution can be used to maximize the overall utility in a manner which assures agents’ gains are at least as high as those acheived in a fully competitive environment.

In the future our ideas can be used to identify classes of problems for which a simple objective function guarantees a non positive CoC for all players. Thus, one may use a distributed protocol for MAS and achieve better personal results than either letting participants play out the game, or search for a NE for them. We hope that through the CoC ideas we may be able to devise new objective functions which can provide guarantees on a wide set of problems.

7. Acknowledgment

The research was supported by the Lynn and William Frankel Center for Computer Sciences at Ben-Gurion University and by the Paul Ivanier Center for Robotics Research and Production Management.

References

1. Aumann, R.J.: Subjectivity and correlation in randomized strategies. Journal of Mathematical Economics 1(1), 67–96 (March 1974)
2. Crawford, E., Veloso, M.: Mechanism design for multi-agent meeting scheduling. Web Intelligence and Agent Systems 4(2), 209–220 (2006)
3. Ephrati, E., Zlotkin, G., Rosenschein, J.: A non manipulable meeting scheduling system. In: Proc. Intern. Workshop on Distributed Artificial Intelligence. Seatle, WA (1994)
4. Gershman, A., Grubshtein, A., Meisels, A., Rokach, L., Zivan, R.: Scheduling meetings by agents. In: Proc. 7th International Conference on Practice and Theory of Automated Timetabling (PATAT 2008). Montreal (August 2008)
5. Grubshtein, A., Grinshpoun, T., Meisels, A., Zivan, R.: Local search for distributed asymmetric optimization. In: Proc. of the 9th International Conference on Autonomous Agents and Multiagent Systems(AAMAS 2010). Toronto, Canada (May 2010)
6. Kearns, M.J., Littman, M.L., Singh, S.P.: Graphical models for game theory. In: UAI '01: Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence. pp. 253–260. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (2001)
7. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: in Proc. of the 16th Annual Symposium on Theoretical Aspects of Computer Science. pp. 404–413 (1999)
8. Maheswaran, R.T., Tambe, M., Bowring, E., Pearce, J.P., Varakantham, P.: 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). pp. 310–317. NY, New York (2004)
9. Meisels, A.: Distributed Search by Constrained Agents: Algorithms, Performance, Communication. Springer Verlag (2007)
10. Modi, J., Veloso, M.: Multiagent meeting scheduling with rescheduling. In: Proc. 5th workshop on distributed constraints reasoning DCR-04. Toronto (September 2004)
11. Monderer, D., Tennenholtz, M.: Strong mediated equilibrium. Artificial Intelligence 173(1), 180–195 (2009)
12. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: Algorithmic Game Theory. Cambridge University Press (2007)
13. Papadimitriou, C.: Algorithms, games, and the internet. In: STOC '01: Proc. of the 33rd annual ACM symposium on Theory of computing. pp. 749–753 (2001)
14. Petcu, A., Faltings, B., Parkes, D.: M-DPOP: Faithful distributed implementation of efficient social choice problems. Journal of AI Research (JAIR) 32, 705–755 (2008)
15. Roughgarden, T.: Selfish Routing and the Price of Anarchy. The MIT Press (2005)
16. Roughgarden, T., Tardos, É.: How bad is selfish routing? J. ACM 49(2), 236–259 (2002)
17. Rozenfeld, O., Tennenholtz, M.: Routing mediators. In: IJCAI. pp. 1488–1493. Hyderabad, India (2007)
18. Tennenholtz, M.: Competitive safety analysis: robust decision-making in multi-agent systems. Journal of Artificial Intelligence Research 17(1), 363–378 (2002)
19. Wallace, R.J., Freuder, E.: 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. pp. 176–182. Bologna (2002)
20. Zhang, W., Xing, Z., Wang, G., Wittenburg, L.: Distributed stochastic search and distributed breakout: properties, comparishon and applications to constraints optimization problems in sensor networks. Artificial Intelligence 161:1-2, 55–88 (January 2005)

Alon Grubshtein is a PhD student in the department of Computer Science at Ben-Gurion University working under the supervision of Prof. Amnon Meisels. His research is focused on competitive and cooperative agents and Disrtibuted Constraint Reasoning. Alon Grubshtein is currently employeed in Deutsche Telekom's Labarotories at Ben-Gurion University as a junior researcher.

Amnon Meisels has been a faculty member of the department of Computer Science at Ben-Gurion University for the last 20 years. Served as head of the Computer Science dept. from 2004 to 2006. The main field of research of Amnon Meisels is Constraints Processing and in the last 7 years mostly Distributed Constrained Search. His book "Distributed Search by Constrained Agents", was published by Springer Verlag in January 2008. Amnon Meisels has been part of the organizing committee of the series of workshops on distributed constraints reasoning (DCR series) every year since 2003. Important research milestones of the DisCSP group at BGU that is headed by Prof. Meisels in the last 7 years include: (a) *Asynchronous forward-checking Concurrent search - algorithms for DisCSPs*, (b) *Asynchronous ordering heuristics for DisCSP algorithms*, (c) *Concurrent run-time measures for experimental evaluation of DisCSP algorithms*, (d) *Privacy preserving search methods for DisCSPs*, and (e) *Asynchronous Forward-bounding for DisCOPs - establishing the phase transition*

Received: January 18, 2010; Accepted: May 3, 2010.