Cooperation and Competition When Bidding for Complex Projects: Centralized and Decentralized Perspectives

Skowron, Piotr; Rzadca, Krzysztof; Datta, Anwitaman

2017

Skowron, P., Rzadca, K., & Datta, A. (2017). Cooperation and Competition When Bidding for Complex Projects: Centralized and Decentralized Perspectives. IEEE Intelligent Systems, 32(1), 17-23.

https://hdl.handle.net/10356/84096

https://doi.org/10.1109/MIS.2017.4

Cooperation and Competition when Bidding for Complex Projects: Centralized and Decentralized Perspectives

Piotr Skowron
University of Oxford
Email: piotr.skowron@cs.ox.ac.uk

Krzysztof Rzadca
University of Warsaw
Email: krzadca@mimuw.edu.pl

Anwitaman Datta
Nanyang Technological University
Email: anwitaman@ntu.edu.sg

Introduction

Consider a complex project: involved, intricate, and consisting of many varied yet interrelated parts. The successful completion of such a project requires coordinated cooperation of a number of experts—people and companies—often organized as teams of subcontractors [1]. For instance, in the construction industry, to build an apartment building (a rather standard endeavor), typically, 30 to 40 individual sub-contractors are involved in 100 to 150 separate activities [2].

Assigning sub-tasks of a complex project to subcontractors is common. In the UK, the proportions of construction employees employed by sub-contractors in years 1983–1998 has grown by 20% [3]. In the UK, between 2008 and 2011, the number of freelancers increased by 12%; in Australia in 2012, 17.2% of the workforce were self-employed (8.5% as independent contractors). These are only a few examples of a growing tendency to develop projects by employing many specialized sub-contractors instead of a single company.

Nevertheless, it is not clear how to organize the market both for the issuer of the project (in this paper, called the client) and the subcontractors (later called the agents). Interaction between the agents applying for the employment in a project and the client is captured by the *hiring a team* problem [4]–[6]. The agents have private costs of participating in the project and may have different skills, thus only certain teams are able to complete the project on time. The client organizes an auction in which individual agents place their bids, i.e., their required salaries. After collecting the bids, the client selects the cheapest feasible team, i.e., the set of agents able to complete the project on time with the lowest total bid.

We generalize the hiring a team problem by exploring two organizations of the market. The original approach corresponds to the *centralized setting*: the agents communicate only with the client by issuing their required salaries (also referred to as bids) and it is the client's responsibility to select an appropriate team. Our main contribution lies in considering a different organization of the market, where *the agents first form teams and then bid for the project as consolidated groups* rather than as individuals. Since the organization of the agents into teams is not managed by a central entity, we refer to this setting as *decentralized*. To the best of our knowledge this formulation of the problem is novel and leads to a new class of games. This formulation has a natural interpretation: a client may not want to coordinate a project and to deal with individual subcontractors, but instead expects that the subcontractors coordinate among themselves and propose a bid for completing the whole project.

Additionally, we generalize the hiring a team problem by considering two types of agents' compensation. In the *project salary model* (corresponding to the original approach) the agents are paid for their work irrespectively of the contributed effort. We propose a new payment model, the *hourly salary*, in which the agents are paid for the time spent working on the project.

Throughout this paper we assume that we are given an oracle that, for a given team, can determine whether this team is feasible, i.e., whether it can successfully complete the project. In particular, given a budget and the requested agents' salaries, the oracle can be asked to find a feasible team or to find the cheapest feasible team. In the online appendix [7] we show how to create such an oracle for a concrete scheduling model (which moreover generalizes a commodity auction); we also determine its exact complexity.

Our approach generalizes two models: commodity auctions [8] and path auctions [9] (see the online appendix [7]). In a commodity auction, there is a set of items $I = \{i_1, i_2, \dots, i_q\}$ and agents owning certain subsets of I. A team is feasible if the agents have together all the items from I. A commodity auction can be mapped to our problem by considering that I is a set of independent activities; an agent owning a subset corresponds to an agent having skills to complete these activities. In a path auction, there is a graph G with two distinguished vertices: a source s and a target t. The agents correspond to the vertices in the graph; some vertices are connected by edges. A team is feasible if the participating agents form a path from s to t.

Since we consider teams of agents with sufficient skills, our model resembles cooperative skill games [10] and coalitional resource games [11]. These games, however, consider the stability of the *grand coalition* and interaction between its members. In contrast, our approach is to *expose the competition between multiple teams*. Thus, we do not apply the typical cooperative game theory concepts and, instead, model agents' cooperation and competition as a non-cooperative game (see the online appendix [7] for a detailed comparison). Our approach is thus closer to endogenous formation of coalitions [12]–[14].

In this paper we identify and formalize a new class of coalition games, which are the extensions of the hiring a team games. We propose concepts to characterize the stability of the winning teams and study their computational complexity. Table 1 summarizes our results. All the proofs omitted from the main text are provided in the online appendix [7].

<table>
<tr><th></th><th>Concept</th><th>existence</th><th>checking</th><th>finding</th></tr>
<tr><td rowspan="4">decentraliz.</td><td>rigor. strongly winning team</td><td>Not always</td><td>$O(n^2 \cdot \text{FCFT})$</td><td>$O(n^5 log(nv)\text{FCFT})$ † ρ</td></tr>
<tr><td>strongly winning team</td><td>Not always</td><td colspan="2">open problem</td></tr>
<tr><td>weakly winning team</td><td>Always</td><td colspan="2">$O(n^5 log(nv)\text{FCFT})$ † ρ</td></tr>
<tr><td>auction winning team</td><td>Always</td><td>$O(\text{FFT})$</td><td>$O(v \cdot \text{FFT})$</td></tr>
<tr><td rowspan="2">centralized</td><td>winning team (with asking salaries)</td><td>N/A</td><td colspan="2">$O(\text{FCFT})$</td></tr>
<tr><td>Strong Nash Equilibrium</td><td>Always†
Not always ⋆</td><td>$O(\text{FCFT})$</td><td>$O(n^3 log(nv)\text{FCFT}))$ † ρ</td></tr>
</table>

Table 1: *Summary of our results. "Existence" denotes whether a team/equilibrium always exists. "Checking" gives the complexity of checking whether a given team satisfies the definition. "Finding" gives the complexity of finding a team/equilibrium. FFT and FCFT are the complexities of the problems FFT and FCFT, respectively. The symbols † (respectively, ⋆) denotes that a result is valid only in the project (respectively, hourly) salary model. The symbol ρ denotes that a result is valid only if the salaries of the agents are rational numbers. Whenever one of the symbols †,⋆ or ρ is provided, it means that the problem for the other cases is still open.*

A Complex Project as a Game: a Formal Model

We consider a game in which a client (an issuer) submits a single complex project. The client has a certain valuation v of the project that is the maximal price that she is able to pay for completing the project.

There is a set $N = \{1,2, \ldots n\}$ of n agents. For each agent i, we define $\phi_i^m > 0$ to be the agent's *minimal salary* for which i is willing to work. This minimal salary may correspond to the agent's personal cost of participating in the project. The agent prefers to work for ϕ_i^m than not to work (and then to work for higher salary). The value ϕ_i^m is private to the agent—neither the issuer nor the other agents know ϕ_i^m.

A subset of the agents' population N forms a team to work on the project; the paper's core contribution is on how this process should be organized. A team $\mathcal{C}$ is a triple $\langle N_\mathcal{C}, \phi_\mathcal{C}, c_\mathcal{C} \rangle$ consisting of: the set of participating agents $N_\mathcal{C} \subseteq N$; a salary function $\phi_\mathcal{C}: N_\mathcal{C} \to \mathbb{N}$ assigning salaries to member agents; and the total cost of the team $c_\mathcal{C} \in \mathbb{N}$—the total amount of money earned by the participants of $\mathcal{C}$. Salaries are discrete (not only money is discrete, but also it is common in real-world auctions to specify a minimal difference between two successive bids). However, to derive some computational results, in some clearly marked places, we assume that the salaries can be rational numbers.

The same team may organize the work of its members on the project in various ways with varying efforts from participants. To capture this property, we introduce a notion of a *schedule*, $\sigma_\mathcal{C}: N_\mathcal{C} \to \mathbb{N}$, that assigns to each member of a team the amount of time this agent needs to spend on the project. Of course, there may exist many schedules for a single team. We expand the discussion on the notion of schedule in the online appendix [7].

We consider two models of agents' compensation. Let $\phi_{\mathcal{C}}^{tot}(i)$ denote the total amount of money agent i gets in team $\mathcal{C}$ (naturally, $c_{\mathcal{C}} = \sum \phi_{\mathcal{C}}^{tot}(i)$). In the *project salary* model $\phi_{\mathcal{C}}^{tot}(i)$ is equal to the salary of the agent $\phi_{\mathcal{C}}(i)$ (and thus does not depend on the amount of work assigned to that agent). In the *hourly salary* model $\phi_{\mathcal{C}}^{tot}(i)$ is equal to the product of the salary $\phi_{\mathcal{C}}(i)$ and the time t_i during which i processes her part of the project (t_i is known from the schedule).

In the project salary model the agents are interested in earning as much money as possible. The hourly salary model represents agents who are interested in having the highest possible hourly wage; thus, e.g., an agent prefers to work $t_i = 1$ time unit with a salary $\phi_i = 3$ to working $t_i = 2$ time units with a salary $\phi_i = 2$.

Different schedules might result in different completion times of the project. If the schedule results in a completion time that is satisfactory for the client, we say that the schedule is *feasible*. For some teams there might not exist a feasible schedule (e.g., if the members lack certain skills). We assume that there is an oracle that can answer whether a given schedule is feasible or not. This very general setting can be instantiated by providing a concrete oracle. For instance, in the online appendix [7] we show that by appropriately specifying the oracle, our results can be applied to commodity auctions and to path auctions. We also show there how to replace the general oracle with a concrete scheduling model.

A team $\mathcal{C}$ is *feasible* iff (i) the asking salaries are no-lower than the minimal salaries, $\phi_{\mathcal{C}}(i) \geq \phi_i^m$; and there exist a feasible schedule such that: (ii) the project budget is not exceeded ($c_{\mathcal{C}} \leq v$), and (iii) the cost $c_{\mathcal{C}}$ of the team $\mathcal{C}$ is consistent with the salaries $\phi_{\mathcal{C}}$. Specifically, in the project salary model $c_{\mathcal{C}} = \sum \phi_{\mathcal{C}}(i)$. In the hourly salary model $c_{\mathcal{C}} = \sum t_i \phi_{\mathcal{C}}(i)$.

A team $\mathcal{C}$ is *cheaper* than $\mathcal{C}'$ if it has a strictly lower cost $c_{\mathcal{C}} < c_{\mathcal{C}'}$ or if it has the same cost, but it is preferred by a deterministic tie-breaking rule ‚$\prec$‚, $N_{\mathcal{C}} \prec N_{\mathcal{C}'}$ (for the sake of concreteness we assume that ‚$\prec$‚ is the lexicographic order in which a team is represented by a concatenation of the sorted list of the names of its members).

Throughout this paper we use the FIND FEASIBLE TEAM (FFT) and FIND CHEAPEST FEASIBLE TEAM (FCFT) problems.

Problem 1 An instance of Find Feasible Team (FFT) consists of a project (with a budget v) and the set N of the agents with (known) minimal required salaries ϕ_i^m. The question is to find some feasible team or to claim there is no such. In the Find Cheapest Feasible Team (FCFT) we ask for the cheapest feasible team.

Centralized Formation of Teams

In the centralized model agents submit their asking salaries ϕ_i directly to the client. The client, having the asking salaries, wants to form the cheapest feasible team. We first show that this problem reduces to FFT, the problem of finding a feasible team. Then, we analyze the optimal bidding strategies of agents.

Proposition 1 The problem FCFT can be solved in time $O((\log(v) + n)\mathrm{FFT})$, where FFT is the complexity of the problem FFT. Having the asking salaries of the agents, the problem of finding the winning team can be solved in time $O(\mathrm{FCFT})$, where FCFT is the complexity of the problem FCFT.

The agents may behave strategically and manipulate their asking salaries to maximize their payoffs. We model this problem as a strategic game. An action of agent i is her asking salary $\phi_i \geq \phi_i^m$. The payoff of i is ϕ_i iff i is a member of the cheapest feasible team; otherwise the payoff of i is 0.

Interestingly, in the project salary model, there exist sets of vectors of actions which are stable against collaborative strategies of the agents. We recall that a vector of the agents' actions is a Strong Nash Equilibrium (SNE) if no subset of the agents can change its actions so that all the deviating agents obtain strictly better payoffs.

For each subset of the agents $N' \subseteq N$, by $\mathcal{C}^O(N')$ we denote the cheapest feasible team using only the agents from N' (if there is no feasible team consisting of the agents from N', the team $\mathcal{C}^O(N')$ does not exist).

Theorem 2 In the project salary model, if there exists a feasible team then there exists a Strong Nash Equilibrium. In every SNE, the set of the agents who get positive payoffs is the set of agents forming the cheapest feasible team, $N_{\mathcal{C}^O(N)}$.

Proof: Let $N^O = N_{\mathcal{C}^O(N)}$ be the set of the agents participating in the cheapest feasible team. We say that the action ϕ_i of the agent i is *minimal* if and only if $\phi_i = \phi_i^m$ We show how to construct the asking salaries ϕ_i^O of the agents from N^O that, together with the minimal actions of the agents outside N^O, form a Strong Nash Equilibrium. A sketch of the proof is as follows. We show the set of linear inequalities for the variables $\phi_i, i \in N^O$. Let us denote the maximal values of ϕ_i which satisfy the inequalities as ϕ_i^O (maximal in the sense that if we increase any value ϕ_i^O, then the new values will not satisfy all the inequalities any more). We show that the actions ϕ_i^O of the agents from N^O, together with the minimal actions of the agents outside of N^O, form an SNE and that the set of the solutions ϕ_i^O that satisfy all the inequalities is nonempty.

The first inequality states that the values ϕ_i must lead to a feasible solution:

$$\sum_{i \in N^O} \phi_i \leq v.$$

Next, as $\mathcal{C}^O$ is the cheapest feasible team, for each feasible team $\mathcal{C}'$ ($N^O \neq N_{\mathcal{C}'}$) such that $N^O \prec N_{\mathcal{C}'}$, $\mathcal{C}^O$ must have (weakly) lower cost:

$$\sum_{i \in N^O - N_{\mathcal{C}'}} \phi_i \leq \sum_{i \in N_{\mathcal{C}'} - N^O} \phi_i^m .$$

For a $\mathcal{C}'$ preferred over $\mathcal{C}^O$ ($N^O \neq N_{\mathcal{C}'}$ and $N_{\mathcal{C}'} \prec N^O$), $\mathcal{C}^O$ must have strongly lower cost:

$$\sum_{i \in N^O - N_{\mathcal{C}'}} \phi_i < \sum_{i \in N_{\mathcal{C}'} - N^O} \phi_i^m .$$

First, if the values ϕ_i^O satisfy the above three inequalities and the agents outside of N^O play their minimal actions, then the agents from N^O will get positive payoffs. If they did not get the positive payoffs, it would mean that there exists a feasible cheaper team $\mathcal{C}'$. However, the three inequalities imply that the agents from $N^O - N_{\mathcal{C}'}$ induce the lower total cost than the total cost of the agents from $N_{\mathcal{C}'} - N^O$; this ensures that agents N^O with actions ϕ_i^O form a cheaper team than $\mathcal{C}'$.

Next, we show that no set of agents $N_{\mathcal{C}'}$ can make a collaborative action $\bar{\phi}$, after which the payoff for all $N_{\mathcal{C}'}$ agents will be greater than previously. By contradiction, assume that there exists such a set of agents $N_{\mathcal{C}'}$ and such an action $\bar{\phi}$. First we consider the case when the payoff of some agent $i \notin N^O$ would change. This means that after $\bar{\phi}$ there would be a new cheapest feasible team $\mathcal{C}'$, where $i \in N_{\mathcal{C}'}$. However, we know that the total cost of the agents from $N^O - N_{\mathcal{C}'}$ is lower than the total cost of the agents from $N_{\mathcal{C}'} - N^O$. This means that $\mathcal{C}'$ cannot be cheaper than the team consisting of the agents from N^O. Finally, consider the case when only the payoffs of the agents from N^O change (and thus $N_{\mathcal{C}'} \subseteq N^O$). However, if the strict subset of N^O could form a feasible team, then $\mathcal{C}^O(N)$ would not be the cheapest. Thus, $N_{\mathcal{C}'} = N^O$. This means that every agent from N^O must have played a higher action (and others must have not changed their actions). Since ϕ_i^O were maximal, this means that after the action $\bar{\phi}$ some inequality, for some feasible team $\mathcal{C}''$, would not hold any more. Thus, we infer that $\mathcal{C}''$ is cheaper than $\mathcal{C}'$.

To check that there always exists a solution, we see that the definition of N^O ensures that the values $\phi_i^O = \phi_i^m$ satisfy all inequalities.

Finally, by contradiction we prove the N^O is formed by the same agents as forming the cheapest team. Assume that the set of the agents that get positive payoffs in some SNE is $N' \neq N^O$. However, if the agents from $(N^O - N')$ play their minimal actions, then the team consisting of the agents from N^O would be cheaper than the team consisting of the agents from N'. Thus, the agents from $(N^O - N')$ can deviate, getting better payoffs. This completes the proof.□

Interestingly, there is no analogous result for hourly salary model (see the online appendix [7]). The proof of Theorem 2 is constructive, but it requires considering all feasible teams and, so, leads to potentially high computational complexity. Finding an efficient algorithm for the problem of finding Strong Nash Equilibria in the project salary model is open. On the other hand, if the salaries of the agents can be rational numbers, we can find the salary function in SNE by a polynomial reduction to the FCFT problem. This result is particularly meaningful if the salaries have high granularity; rounding such a rational solution gives an integral solution which is nearly perfect.

Proposition 3 In the project salary model, if the salaries are rational, then finding a Strong Nash Equilibrium can be solved in time $O(n^3 log(nv)\text{FCFT}))$, where FCFT is the complexity of the problem FCFT. Checking whether a given vector of the asking salaries $\langle \phi_i \rangle, i \in N$ is a Strong Nash Equilibrium can be solved in time $O(\text{FCFT})$, where FCFT is the complexity of the problem FCFT.

Decentralized Formation of Teams

If the agents can communicate and coordinate their strategies, they form teams and bid for the project as consortiums. We propose the concept of a (rigorously) strongly winning team, in which no subset of agents can successfully deviate. We show how to characterize (rigorously) strongly winning teams and how to reduce the problem of finding them to the FCFT problem. We show that the strongly winning teams may not exist, and so we introduce the concept of a weakly winning team. We prove that a weakly winning team always exists (provided that there is a feasible team). We demonstrate how to reduce the problem of finding weakly winning teams to the FCFT problem.

We model the behavior of the agents as a strategic game. Agent i's action is a triple $\langle N_{\mathcal{C}}, \phi_{\mathcal{C}}, b_{\mathcal{C}} \rangle$. Intuitively, such an action means that the agent i decides to enter the team $\mathcal{C} = \langle N_{\mathcal{C}}, \phi_{\mathcal{C}}, b_{\mathcal{C}} \rangle$. The payoff of the agent is equal to $\phi_{\mathcal{C}}(i)$ if (i) $\mathcal{C}$ is feasible, (ii) each agent $j \in N_{\mathcal{C}}$ agrees to participate in $\mathcal{C}$ (i.e., they all play $\mathcal{C}$, and their payoffs are consistent with the bid of the team $b_{\mathcal{C}}$), and (iii) there is no feasible cheaper team $\mathcal{C}'$ such that all the agents from $N_{\mathcal{C}'}$ agree to participate in $\mathcal{C}'$. Otherwise, the payoff of i is 0.

Strongly Winning Teams

As the payoffs depend on whether the others agree to cooperate, rather than the Nash Equilibrium, the Strong Nash Equilibrium (SNE) should be used. In the following definition we propose an even more stable equilibrium concept—the Rigorously Strong Nash Equilibrium (RSNE), which requires that no subset of agents can deviate such that each agent gets a payoff *at least as good* as its payoff before deviating (instead of SNE's strictly better). Our approach is motivated by cautious agents. In an SNE, the agents have no incentive to deviate if they get the same payoff; however they also have no incentive not to deviate. Yet, any deviation will result in a serious payoff loss for some agents (changing their payoffs from a positive ϕ to zero). A cautious agent will prefer not to be exposed to the possibility of such a loss.

Definition 1 The vector of actions π is a *Rigorously Strong Nash Equilibrium (RSNE)* iff there is no subset of agents $N_{\mathcal{C}}$ such that the agents from $N_{\mathcal{C}}$ can make a collaborative action $\mathcal{C}$ after which the payoff of each agent i from $N_{\mathcal{C}}$ would be at least equal to her payoff under π and the payoff of at least one agent $i \in N$ would improve.

A RSNE requires that the payoff of at least one agent $i \in N$ must change as we treat as equivalent the teams with the same payoffs. For instance, in a game with three agents, a, b and c, if the team $\{a, b\}$ gets a positive payoff, it does not matter whether c plays $\langle \{c\}, v + 1 \rangle$ or $\langle \emptyset, v + 1 \rangle$: in both cases all payoffs are the same (recall that v is the client's maximal budget for the project).

Below we introduce additional definitions that help characterize the RSNE in our games.

Definition 2 A feasible team $\mathcal{C}$ is *explicitly endangered* by a team $\mathcal{C}'$ if (i) $\mathcal{C}'$ is feasible, (ii) $N_{\mathcal{C}} \cap N_{\mathcal{C}'} = \emptyset$ and (iii) $\mathcal{C}'$ is cheaper than $\mathcal{C}$. A feasible team $\mathcal{C}$ is *implicitly endangered* by a team $\mathcal{C}'$ if (i) $\mathcal{C}'$ is feasible, (ii) $N_{\mathcal{C}} \cap N_{\mathcal{C}'} \neq \emptyset$ and each agent from $N_{\mathcal{C}} \cap N_{\mathcal{C}'}$ gets in $\mathcal{C}'$ at least as good a salary as in $\mathcal{C}$, and (iii) either $N_{\mathcal{C}} \neq N_{\mathcal{C}'}$ or $\phi_{\mathcal{C}} \neq \phi_{\mathcal{C}'}$.

If there are agents belonging to both teams ($N_{\mathcal{C}} \cap N_{\mathcal{C}'} \neq \emptyset$), we do not consider the total cost of the alternative team $\mathcal{C}'$, as the decision whether $\mathcal{C}'$ will be formed depends solely on the agents from $N_{\mathcal{C}} \cap N_{\mathcal{C}'}$: if they decide to form $\mathcal{C}'$, $\mathcal{C}$ will not be formed, thus the client won't be able to choose between $\mathcal{C}$ and $\mathcal{C}'$.

A feasible team $\mathcal{C}$ is *(rigorously) strongly winning* iff there is a (Rigorously) Strong Nash Equilibrium in which the agents from $N_{\mathcal{C}}$ get positive payoffs $\phi_{\mathcal{C}}$. The following theorem relates endangerment (Definition 2) and a winning team.

Theorem 4 The team $\mathcal{C}$ is rigorously strongly winning if and only if $\mathcal{C}$ is not explicitly nor implicitly endangered by any team.

The result in Theorem 4 stated for RSNEs transfers to SNEs after a slight modification of the payoffs. It is sufficient to assume that an agent playing an empty team receives slightly higher payoff than if she plays a non-empty losing team. In other words, this modification associates some small costs with the preparation of a bid by the agents. Hereinafter, whenever we mention a strictly winning team we assume that the agents incur such costs. To state the result for SNEs we also need to use the definition of a team $\mathcal{C}$ being *strictly implicitly endangered* by $\mathcal{C}'$. This definition differs from being implicitly endangered only by not requiring the agents from $N_{\mathcal{C}} \cap N_{\mathcal{C}'}$ to have at least as good payoffs, but strictly better payoffs in $\mathcal{C}'$ than in $\mathcal{C}$.

Theorem 5 If there are small but positive costs of preparing the offer by the agents then the team $\mathcal{C}$ is strongly winning if and only if $\mathcal{C}$ is not explicitly nor strictly implicitly endangered by any team.

Theorems 4 and 5 lead to a simple brute-force algorithm for checking whether the team $\mathcal{C}$ can be a part of some RSNE. It is sufficient to check whether for each set of agents $N' \subseteq N_{\mathcal{C}}$ there exists a payoff function $\phi_{\mathcal{C}}$ and a cost c such that $\mathcal{C}$ is explicitly or implicitly endangered by $\langle N', \phi, c \rangle$ (such a condition can be checked by enumerating the payoff functions which assign to each agent his or her minimal salary, the salary that he or she obtains in $\mathcal{C}$, or the next higher salary). Below, we characterize RSNEs in the project salary model even more precisely.

Lemma 6 In the project salary model, the set of agents participating in a rigorously strongly winning team is the same as the set of agents participating in the cheapest feasible team.

Lemma 7 In the project salary model the bid of a strongly winning team is equal to the maximal allowed price v.

Lemmas 6 and 7 show that the problem of finding a strongly winning team reduces to the problem of finding a feasible team. The problem, thus, becomes an optimization problem; the strategic behavior of agents has no impact (see the online appendix [7]). An RSNE (and even an SNE) may not exist in some instances.

Proposition 8 Both in the project salary and in the hourly salary model, there may not exist a strongly winning team even though there exists a feasible team.

Proof: Consider a project with budget $v = 5$; and three identical agents a, b, c with minimal salaries $\phi_i^m = 2$ (in the hourly salary model, assume that each agent spends exactly 1 time unit on the project); a team of any two agents is feasible (able to complete the project on time and within the budget).

For the sake of contradiction assume there exists a team $\mathcal{C}$ that gets positive payoffs. Without loss of generality we assume that $N_{\mathcal{C}} = \{a, b\}$. At least one of the agents, let us say a, has to get salary at most equal to 2.5. However, the agents a and c, with the salaries equal to 3 and 2 respectively, can form a feasible team in which both a and c get better payoffs.□

Weakly Winning Teams

Proposition 8 suggests that a notion of a strongly winning team is too restrictive. The team $\{a, c\}$ can profit by deviating, e.g., by playing $\phi(a) = 3$ and $\phi(c) = 2$. But a should not be willing to deviate, as $\{a, c\}$ with payoffs $\phi(a) = 3$ and $\phi(c) = 2$ too is not stable (for instance, the team $\{b, c\}$ can play $\phi(b) = 2$ and $\phi(c) = 3$, and successfully deviate from $\{a, c\}$). In the above example no team strongly wins, even though intuitively there are teams that would agree to work. Thus, we propose a weaker notion of a winning team.

Definition 3 A feasible team $\mathcal{C}$ is *weakly winning* if it is not explicitly endangered by any team and for each feasible team $\mathcal{C}'$ such that $\mathcal{C}$ is implicitly endangered by $\mathcal{C}'$, there exists a feasible team $\mathcal{C}''$ such that $\mathcal{C}'$ is explicitly or implicitly endangered by $\mathcal{C}''$.

Proposition 9 There exists a weakly winning team if and only if there exists a feasible team.

Proposition 10 In the project salary model, if the salaries of the agents can be rational numbers, the problem of finding a weakly winning team and the problem of checking whether a team $\mathcal{C}'$ is weakly winning can be solved in time $O(n^5 log(nv)$FCFT).

Finding an efficient algorithm for the same problem with discrete salaries is still an open question.

Mechanism Design

In this section we analyze two mechanisms that a client can use to find a winning team: the first one sets the project's budget v; the second one uses a first-price auction.

First, we show that if the client is allowed to change the budget v there exists a simple mechanism (based on a binary search) ensuring the existence of a strongly winning team.

Theorem 11 If there exists a feasible team, then there exists a budget v^O for which there exists a strongly winning team. The problem of finding such a v^O can be solved in time $O(\log(v) \cdot \text{FFT})$.

In the second approach we use the first-price auction in which teams participate. In a standard first-price auction, an item's price starts from some minimal value (the least preferred outcome for the owner of the item). Bidders place bids for the current price. The asking price is gradually increased until there are no further bids; the last bidder wins. Similarly, in our proposed auction, the auction starts from the original budget v (the least preferred outcome for the client); the asking price gradually *decreased*. Teams place bids for the current asking price (as in the standard first-price auction, multiple bids for the same asking price are not allowed). The auction stops when no feasible team bids lower than the current asking price. This procedure leads to the concept of an auction-winning team.

Definition 4 A team $\mathcal{C}$ is *auction-winning* iff there is no feasible team $\mathcal{C}'$ such that $b_{\mathcal{C}'} < b_{\mathcal{C}}$ and for each agent $i \in N_{\mathcal{C}} \cap N_{\mathcal{C}'}$ the agent gets better salary in $\mathcal{C}'$, $\phi_{\mathcal{C}'}(i) \geq \phi_{\mathcal{C}}(i)$.

Proposition 12 The problem of checking whether a feasible team $\mathcal{C}$ is auction-winning can be solved in time $O(\text{FFT})$. The problem of finding an auction-winning team can be solved in time $O(v \cdot \text{FFT})$.

Conclusions

We presented a new class of coalitional games that model cooperation and competition for employment in a complex project. Our games extend and relate to a number of well-known problems, such as coalition formation, coalitional auctions, auctions for sharable items, etc. We considered two market organizations. In a centralized market, the winning team is selected by the client based on bids from individual agents; the agents are strategic about the salaries they request. In a decentralized market, the already-formed teams bid for the project, thus the agents are strategic both regarding their salaries and regarding their cooperation partners.

We proposed concepts of stability for each of our models and we showed how to reduce the problem of finding a winning team to the problem of finding a feasible one, for which we assumed we have an oracle with known complexity.

To instantiate our abstract model, in the online appendix [7] we show how to solve a scheduling problem in which the project is a set of independent tasks and the agents have certain skills in processing them (represented by unrelated processing speeds).

Acknowledgements: The authors thank Marcin Dziubiński and Piotr Faliszewski for their helpful comments and Edith Elkind for the fruitful discussion on the related literature.

This research has been partly supported by the Polish National Science Center grants Sonata (UMO-2012/07/D/ST6/02440) and Preludium (UMO-2013/09/ N/ST6/03661) and by European Research Grant ERC-StG 639945.

Abstract

To successfully complete a complex project, be it a construction of an airport or of a backbone IT system, agents (companies or individuals) must form a team having required competences and resources. A team can be formed either by the project issuer based on individual agents' offers (centralized formation); or by the agents themselves (decentralized formation) bidding for a project as a consortium—in that case many feasible teams compete for the contract. We investigate rational strategies of the agents (what salary should they ask? with whom should they team up?). We propose concepts to characterize the stability of the winning teams and study their computational complexity.

keywords: game theory, cooperative game theory, coalition formation, equilibria, skill games, scheduling, co-opetition

References

[1] T. M. Williams, "The need for new paradigms for complex projects," *International Journal of Project Management*, vol. 17, no. 5, pp. 269–273, 1999.

[2] K. D. Walsh, A. Sawhney, and H. H. Bashford, "Cycle-time contributions of hyper-specialization and time-gating strategies in us residential construction," in *Proceedings of 11th annual conference on lean construction-2003*, 2003, pp. 390–397.

[3] D. J. Edwards, "Accident trends involving construction plant: An exploratory analysis," *Journal of Construction Research*, vol. 04, no. 02, pp. 161–173, 2003.

[4] N. Chen and A. R. Karlin, "Cheap labor can be expensive," in *Proceedings of soda-2007*, 2007, pp. 707–715.

[5] A. Iwasaki, D. Kempe, Y. Saito, M. Salek, and M. Yokoo, "False-name-proof mechanisms for hiring a team," in *Proceedings of wine-2007*, 2007, pp. 245–256.

[6] N. Chen, E. Elkind, N. Gravin, and F. Petrov, "Frugal mechanism design via spectral techniques," in *Proceedings of focs-2010*, 2010, pp. 755–764.

[7] P. Skowron, K. Rzadca, and A. Datta, "Cooperation and competition when bidding for complex projects: Centralized and decentralized perspectives," arXiv.org, arXiv:1402.2970v3 [cs.GT], Sep. 2016.

[8] A. McCabe, "Frugality in set-system auctions," PhD thesis, University of Liverpool, 2012.

[9] N. Nisan and A. Ronen, "Algorithmic mechanism design (extended abstract)," in *Proceedings of stoc-1999*, 1999, pp. 129–140.

[10] Y. Bachrach, D. C. Parkes, and J. S. Rosenschein, "Computing cooperative solution concepts in coalitional skill games," *Artificial Intelligence*, vol. 204, no. 0, pp. 1–21, 2013.

[11] M. Wooldridge and P. E. Dunne, "On the computational complexity of coalitional resource games," *Artificial Intelligence*, vol. 170, no. 10, pp. 835–871, Jul. 2006.

[12] S. Hart and M. Kurz, "Endogenous formation of coalitions," *Econometrica: Journal of the Econometric Society*, pp. 1047–1064, 1983.

[13] F. Bloch, "Sequential formation of coalitions in games with externalities and fixed payoff division," *Games and Economic Behavior*, vol. 14, no. 1, pp. 90–123, 1996.

[14] D. Ray and R. Vohra, "A theory of endogenous coalition structures," *Games and Economic Behavior*, vol. 26, no. 2, pp. 286–336, 1999.

Authors' Bios

Piotr Skowron is a research assistant at the Department of Computer Science, University of Oxford, UK. He received a PhD in Computer Science in 2015 from the University of Warsaw, Poland. His PhD thesis entitled "Resource Allocation in Selfish and Cooperative Distributed Systems" won runner-up of the 2015 IFAAMAS Victor Lesser Distinguished Dissertation Award, 2016. His current research concerns multi-agent systems and computational social choice. He is interested in problems related to finding proportional representations for given populations of agents, and in applying optimization techniques to variety of social choice problems. He is interested in linking various ideas from game theory, resource allocation, approximation algorithms and computational social choice. Additionally, his research relate to scheduling and load balancing algorithms for distributed systems, resource allocation and applying game theoretical solutions to distributed systems.

Krzysztof Rzadca is an assistant professor at Faculty of Mathematics, Informatics and Mechanics, University of Warsaw, Poland. He received a PhD in Computer Science in 2008 jointly from Institut National Polytechnique de Grenoble (INPG), École Nationale Supérieure d'Informatique et de Mathématiques Appliquées de Grenoble (ENSIMAG), France and from Polish-Japanese Institute of Information Technology (PJIIT), Poland. His research focuses on resource management and scheduling in large-scale distributed systems, such as computational clouds or data centers. He is interested in ensuring efficiency and fairness in presence of selfish users and weak central control. He is an IEEE and an ACM member.

Anwitaman Datta is an associate professor in the School of Computer Science and Engineering at Nanyang Technological University in Singapore. His areas of interest are broadly in distributed systems, cybersecurity and social network (analysis). He received his Phd from the School of Computer and Communication Sciences at École Polytechnique Fédérale de Lausanne in Switzerland.