

Ecole polytechnique fédérale de Zurich Politecnico federale di Zurigo Federal Institute of Technology at Zurich

Institute of Theoretical Computer Science Peter Widmayer Matúš Mihalák Andreas Bärtschi

## Algorithmic Game Theory HS 2013

Exercise sheet 1

(This exercise is not compulsory but you will get feedback on submission.)

In these exercises, we talk about agents possessing some private information (such as the valuation of a painting), and about mechanisms, which ask the agents for this information, in order to compute something on this input, and decides the payments of the agents. These algorithms are publicly known, and the agents can thus speculate (if they want) about lying. If  $v_i$  denotes the valuation of agent *i*, and  $p_i$  the payment of agent *i* by the mechanism, the utility of player *i* is then  $u_i(a) = v_i(a) - p_i$ , where *a* denotes the result computed by the mechanism. In the example of auctioning a painting, *a* would denote who is the winner, and  $v_i(a)$  would then be zero, if *a* is not the winner, and it would be the valuation of agent *i* for the painting, if *i* was the winner. We say that the mechanism is *truthful* if for every lie  $v'_i$  and the correspondingly computed results *a'* and payments  $p'_i$  it holds that  $v_i(a) - p_i \ge v_i(a') - p'_i$ , i.e., the utility of agent *i* when not lying is at least as good as her utility when lying. We will sometimes denote the valuation of the painting by  $v_i$ .

## EXERCISE 1.1:

Consider a sealed-bid auction of one item (such as a painting) to n agents. Each agent has a valuation  $v_i$  for the painting, which is a private information (known only to the agent itself). Every agent i, i = 1, ..., n, sends a bid  $b_i$  to the auctioneer. The auctioneer decides the winner and the payment of the winner. At the end, the utility  $u_i$  of an agent is: zero, if the agent is not a winner;  $v_i - p_i$  if i is the winner and  $p_i$  is the payment of the agent for the item. In a reasonable auction, the item is given to the agent with the highest bid (the winner). We have seen in the lecture that the first-price auction, where the winner pays its bid (i.e., the highest bid), is not truthful (i.e., there can be cases where lying improves an agent's utility). However, the second-price auction, where the winner pays the second highest bid, is truthful.

- a) Is the third-price auction, i.e., the auction where the winner pays the third highest bid, truthful? Explain your answer.
- b) Consider the second-price auction with the following twist: in case the agents with the two highest bids have the same value (i.e., they bid the same), consider the payment of the winner (one of the two agents with the highest bid) to be the third highest bid. Is this auction truthful?
- c) Is the second-price auction truthful also when there is only one agent?

## EXERCISE 1.2:

Recall that in the multicast cost-sharing problem, a source in a network wants to send data in a multicast fashion to a set P of users. To transmit data over a link incurs certain costs, and every user i in the network is willing to pay a certain amount  $v_i$  to receive the data. This amount is called the valuation of a user. We consider the multicast problem on a tree network, where the root of a tree is identified with the source, the edges with the network links, and the vertices with the users (i.e. there is exactly one user at each vertex). The users are interested in the service, but at the same time they do act selfishly, i.e. they want to maximize their utility (kind of "happiness")  $v_i - p_i$ , where  $p_i$  is the payment that user i needs to pay and which the mechanism decides. Also, a user never accepts to pay more than its valuation.

The decision about who will be served and for what price is made by considering the "bids" of the users to pay for the service. The source then computes the set of serviced users  $R \subseteq P$  and also the payments for every user in R.

Let the tree T below be an instance of such a multicast cost sharing problem. Each link l is marked with its cost  $c_l$  and each node i with its valuation  $v_i$ , respectively.



a) Compute the optimum broadcast tree T(R), maximizing the profit (or worth)

$$W(v) = \sum_{i \in R} v_i - \sum_{l \in T(R)} c_l$$

of the source (for the values given in the figure), where  $R \subseteq P$  is the set of (to-be-computed) receivers.

- b) Consider the case where the payments are set to the values of the bids that every user reports to the root. If you were the user marked with an X in the tree T and you would want to receive the broadcast but lower your costs, what would you do? What valuation would you tell?
- c) Consider in our example the payment scheme in which each selected user  $i \in R$  pays the average of the total cost, i.e.,  $\frac{\sum_{l \in T(R)} c_l}{|R|}$ . Is such a payment truthful in our example? Why?
- d) Apply the "truthful" payments from the lecture to our example and state what value each user has to pay. Recall that the payment of user *i* is defined by his "optimum speculation", and in our setting is defined (for player *i*) by  $p_i = \sigma_i(v) \cdot v_i (W(v) W(v|^i \ 0))$ , where  $\sigma_i(v)$  is an indicator function that is 1 if  $i \in R$  and 0 otherwise, and  $(v|^i x)$  denotes the vector  $(v_1, ..., v_{i-1}, x, v_{i+1}, ..., v_n)$ .
- e) In this particular example, are the "truthful" payments covering the cost of the submission to R? (In other words, are these payments fulfilling the budget-balance requirement?)
- f) In the lecture, we have seen an example of budget-balance mechanism (the so called Moulin mechanism): The mechanism works in rounds. Initially, R – the set of the served people – is set to contain everyone. In every round, each agent  $i \in R$  submits a bid  $\tilde{v}_i$ ; the cost of the underlying tree T(R) is shared equally among the agents in R, i.e., each agent is asked to pay  $p := \sum_{l \in T(R)} c_l/|R|$ ; agents for which  $\tilde{v}_i < p$  are removed from R and a new round begins; if no agent is dropped from R, the mechanism stops and the agents in R are served at the price p each. Apply this mechanism to our example on the figure.

## EXERCISE 1.3:

Imagine that the city of Zurich is planing to build a bridge over the Zurich-lake. There are 10 municipalities around the lake that have an interest in this project and each one offers to pay a certain amount of Swiss Franks to contribute to the project. If the total amount of money offered covers the cost to build the bridge, it is built, but otherwise not. The city is willing to implement the process as a truthful mechanism so that no municipalities do not cooperate). Remember that this is achieved through "marginal contribution" which means that a municipality i does not pay anything if its valuation  $v_i$  does not change the decision of building the bridge, while in the other case i has to pay only the amount that is still needed to cover the building costs if only the valuations of the other municipalities are considered.

- a) If the bridge would cost 1.000.000 Swiss Franks, give an example in which
  - i) every municipality would pay 0 Swiss Franks, while the decision to build the bridge would still be positive.
  - ii) the total contribution of the municipalities would be budget-balanced, i.e. the payments of the municipalities would cover the cost of building the bridge (1.000.000 Swiss Franks).
- b) Consider the following payment scheme for the above social-choice problem: Order the agents according to their bids, smallest first. In this order, starting from the agent with the smallest bid, charge every agent i the bid of the agent i 1 (and, obviously, charge nothing the first player). Is this payment scheme truthful?