Information Processing Letters 113 (2013) 67–73

Contents lists available at SciVerse ScienceDirect

Information Processing Letters

www.elsevier.com/locate/ipl

Sponsored search, market equilibria, and the Hungarian Method ☆

Paul Dütting [a,*], Monika Henzinger [b], Ingmar Weber [c]

[a] École Polytechnique Fédérale de Lausanne (EPFL), Station 14, CH-1015 Lausanne, Switzerland
[b] University of Vienna, Faculty of Computer Science, Währinger Straße 29/6.32, A-1090 Vienna, Austria
[c] Yahoo! Research Barcelona, Avinguda Diagonal 177 (8th Floor), E-08018 Barcelona, Spain

ARTICLE INFO

Article history:
Received 25 July 2011
Received in revised form 21 September 2012
Accepted 26 November 2012
Available online 29 November 2012
Communicated by B. Doerr

Keywords:
Algorithms
Matching markets
Budget limits
Envy freeness
Incentive compatibility

ABSTRACT

Matching markets play a prominent role in economic theory. A prime example of such a market is the sponsored search market. Here, as in other markets of that kind, market equilibria correspond to feasible, envy free, and bidder optimal outcomes. For settings without budgets such an outcome always exists and can be computed in polynomial-time by the so-called Hungarian Method. Moreover, every mechanism that computes such an outcome is incentive compatible. We show that the Hungarian Method can be modified so that it finds a feasible, envy free, and bidder optimal outcome for settings with budgets. We also show that in settings with budgets no mechanism that computes such an outcome can be incentive compatible for all inputs. For inputs in general position, however, the presented mechanism—as any other mechanism that computes such an outcome for settings with budgets—is incentive compatible.

1. Introduction

In a matching market n bidders have to be matched to k items. A prime example of such a market is the sponsored search market, where bidders correspond to advertisers and items correspond to ad slots. In this market each bidder has a per-click valuation v_i, each item j has a click-through rate α_j, and bidder i's valuation for item j is $v_{i,j} = \alpha_j \cdot v_i$. More generally, each bidder i has a valuation $v_{i,j}$ for each item j. In addition, each item j has a reserve price r_j. A mechanism is used to compute an outcome (μ, p) consisting of a matching μ and per-item prices p_j. The bidders have quasi-linear utilities. That is, bidder i's utility is $u_i = 0$ if he is unmatched and it is $u_i = v_{i,j} - p_j$ if he is matched to item j at price p_j. The valuations are private information and the bidders need not report their true valuations if it is not in their best interest to do so.

Ideally, the market should be in equilibrium. In the context of matching markets this typically means that the outcome computed by the mechanism should be *feasible*, *envy free*, and *bidder optimal*. An outcome is feasible if all bidders have non-negative utilities and if the price of all matched items is at least the reserve price. It is envy free if it is feasible and if at the current prices no bidder would get a higher utility if he was assigned a different item. It is bidder optimal if it is envy free and if the utility of every bidder is at least as high as in every other envy free outcome. Another requirement is that the mechanism should be incentive compatible. A mechanism is incentive compatible if each bidder maximizes his utility by reporting truthfully no matter what the other bidders report.

For matching markets of the above form a bidder optimal outcome always exists [2], can be computed in polynomial time by the so-called Hungarian Method [3], and every mechanism that computes such an outcome is incentive compatible [4]. The above model, however, ignores the fact that in practice bidders often have budgets. Concrete examples include Google's and Yahoo's ad auction. Budgets are also challenging theoretically as they lead to

☆ A preliminary version of this paper appeared in Dütting et al. (2010) [1].

* Corresponding author.
E-mail address: paul.duetting@epfl.ch (P. Dütting).

0020-0190/$ – see front matter
http://dx.doi.org/10.1016/j.ipl.2012.11.006

discontinuous utility functions and thus break with the quasi-linearity of the original model without budgets.

In our model each bidder can specify a maximum price for each item. If bidder i specifies a maximum price of $m_{i,j}$ for item j, then he cannot pay any price $p_j \geqslant m_{i,j}$. Hence the utility of bidder i is $u_i = 0$ if he is unmatched, it is $u_i = v_{i,j} - p_j$ if he is matched to item j at price $p_j < m_{i,j}$ (strict inequality), and it is $u_i = -\infty$ otherwise.[1] As before an outcome is feasible if all bidders have non-negative utilities and if the price of all matched items is at least the reserve price. It is envy free if it is feasible and if at the current prices no bidder would get a higher utility if he was assigned a different item. It is bidder optimal if it is envy free and if the utility of every bidder is at least as high as in every other envy free outcome.

For this model we show that the Hungarian Method can be modified so that it always finds a bidder optimal outcome in polynomial time. We also show that no mechanism that computes such an outcome is incentive compatible for all inputs. For inputs in general position, i.e., inputs with the property that in a certain weighted multigraph defined on the basis of the input no two walks have exactly the same weight, our mechanism—as any other mechanism that computes a bidder optimal outcome—is incentive compatible [5]. All our results can be extended to more general (but still linear) utility functions.

A similar problem was previously considered by [6]. Their model differs from our model in several ways: (1) The utility u_i of bidder i is $u_i = 0$ if he is unmatched, it is $u_i = v_{i,j} - p_j$ if he is matched to item j at price $p_j \leqslant m_{i,j}$ (weak inequality), and it is $u_i = -\infty$ otherwise. (2) The reserve prices $r_{i,j}$ may depend on the bidders and the items. (3) An outcome is envy free if it is feasible and if for all bidders i and all items j either (a) $u_i \geqslant v_{i,j} - \max(p_j, r_{i,j})$ or (b) $p_j \geqslant m_{i,j}$. For these definitions they showed that for inputs in general position (a) a bidder optimal outcome always exists, (b) a bidder optimal outcome can be computed by a (rather complicated) mechanism in polynomial time, and (c) this mechanism is incentive compatible. For inputs that are not in general position a bidder optimal outcome may not exist as the following example shows.[2,3]

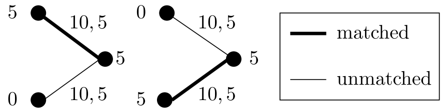

Fig. 1. Bidders are on the left side and items are on the right side of the graphs. The numbers next to the bidder indicate his utility, the numbers next to the item indicate its price. The labels along the edge show valuations and maximum prices. Matched edges are bold, while unmatched edges are thin.

Example 1. There are two bidders and one item. The valuations and maximum prices are as follows: $v_{1,1} = 10$, $v_{2,1} = 10$, and $m_{1,1} = m_{2,1} = 5$. While $\mu = \{(1,1)\}$ with $p_1 = 5$ is "best" for bidder 1, $\mu = \{(2,1)\}$ with $p_1 = 5$ is "best" for bidder 2. With our definitions a bidder optimal outcome is $\mu = \emptyset$ with $p_1 = 5$. See Fig. 1.

The sponsored search market was considered by [7], who proved the existence of a unique feasible, envy free, and Pareto efficient outcome. They also presented an incentive compatible mechanism to compute such an outcome in polynomial time. Their model, however, is less general than the model studied here as (1) the valuations must be of the form $v_{i,j} = \alpha_j \cdot v_i$, and (2) the maximum prices are per-bidder, i.e., for each bidder i there exists m_i such that $m_{i,j} = m_i$ for all j, and are required to be distinct.

Matching markets with more general, non-linear utility functions were studied in [8,9,5]. In [8] we proved the existence of a bidder optimal outcome for general utility functions with multiple discontinuities. In [9] a polynomial-time mechanism for *consistent* utility functions with a single discontinuity was given. In [5] we presented a polynomial-time mechanism for piece-wise linear utility functions with multiple discontinuities.

To summarize: (1) We show how to modify the Hungarian Method in settings with budgets so that it finds a bidder optimal outcome in polynomial time. (2) We show that in settings with budgets no mechanism that computes a bidder optimal outcome can be incentive compatible for all inputs. (3) We show how to extend these results to more general (but still linear) utility functions.

2. Problem statement

We are given a set I of n bidders and a set J of k items. We use letter i to denote a bidder and letter j to denote an item. For each bidder i and item j we are given a *valuation* $v_{i,j}$, for each item j we are given a *reserve price* r_j, and for each bidder i and item j we are given a *maximum price* $m_{i,j}$. We assume that the set of items contains a *dummy item* j_0 for which all bidders have a valuation of zero, a reserve price of zero, and a maximum price of ∞.[4]

We want to compute an *outcome* (μ, p) consisting of a matching $\mu \subseteq I \times J$ and per-item prices $p = (p_1, \ldots, p_k)$. We require that (a) every bidder i appears in exactly one bidder-item pair $(i, j) \in \mu$ and that (b) every non-dummy

[1] While requiring $p_j \leqslant m_{i,j}$ seems to be more intuitive, it has the disadvantage that the infimum envy prices may not be envy free themselves: There are three bidders and one item. All bidders have a valuation of 10 and the first two bidders have a maximum price of 5. Then any price $p \leqslant 5$ is not envy free because all bidders would prefer to be matched, and any price $p > 5$ is not bidder optimal because a slightly lower price would still be envy free.

[2] An input is in general position if in the weighted, directed, and bipartite multigraph with one node per bidder i, one node per item j, and one node for the dummy item j_0 and forward edges from i to j with weight $-v_{i,j}$, backward edges from j to i with weight $v_{i,j}$, reserve-price edges from i to j with weight $v_{i,j} - r_{i,j}$, maximum-price edges from i to j with weight $m_{i,j} - v_{i,j}$, and terminal edges from i to j_0 with weight 0 no two walks that start with the same bidder, alternate between forward and backward edges, and end with a distinct edge that is either a reserve-price edge, a maximum-price edge, or a terminal edge have the same weight.

[3] The example is not in general position because the walk that consists of the maximum-price edge from bidder 1 to item 1 and the walk that consists of the forward edge from bidder 1 to item 1, the backward edge from item 1 to bidder 2, and the maximum-price edge from bidder 2 to item 1 have the same weight.

[4] Reserve utilities, or *outside options* o_i, can be modelled by setting $v_{i,j_0} = o_i$ for all i.