



Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs


Truthful unit-demand auctions with budgets revisited



Monika Henzinger, Veronika Loitzenbauer*

University of Vienna, Faculty of Computer Science, Währinger Straße 29, 1090 Vienna, Austria

ARTICLE INFO

Article history:

Received 5 November 2013

Accepted 20 January 2015

Available online 23 January 2015

Communicated by P. Spirakis

Keywords:

Matching market

Ascending auction

Budget constraint

Incentive compatibility

Envy-freeness

Randomized algorithm

ABSTRACT

We consider auctions of indivisible items to unit-demand bidders with budgets. This setting was suggested as an expressive model for single sponsored search auctions. Prior work presented mechanisms that compute bidder-optimal outcomes and are truthful for a *restricted* set of inputs, i.e., inputs in so-called general position. This condition is easily violated. We provide the first mechanism that is truthful in expectation for *all* inputs and achieves for each bidder no worse utility than the bidder-optimal outcome. Additionally we give a complete characterization for which inputs mechanisms that compute bidder-optimal outcomes are truthful.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

When a user searches for a specific term in a web search engine, related advertisements are displayed on the search results page. The advertisements are assigned by an auction such that each advertiser receives at most one slot. The advertisers may have different preferences among the slots and budgets that limit the amount of money they can spend. The auctioneer may have a *reserve price* under which she is not interested in selling a slot. This is modeled as follows: The advertisers are called *bidders*, the slots correspond to *items*, and the budgets are modeled as *maximum prices* per bidder and item. Each bidder i can specify a *valuation* $v_{i,j}$ for each item j . A mechanism computes the prices p of the items as well as the assignment μ of the bidders to the items. The preferences of a bidder i are modeled by utility functions $u_{i,j}(p_j)$ such that his utility if he is assigned item j at a price p_j is $u_{i,j}(p_j) = v_{i,j} - p_j$ if the price is lower than his maximum price and $-\infty$ if the price is equal to or higher than his maximum price. A bidder has a utility of zero if he is not assigned any item; thus bidder i only accepts an item j if $u_{i,j}(p_j) \geq 0$. For more details on the expressiveness of quasi-linear utility functions with budgets in sponsored search see [1].

Search engine providers want to satisfy their customers as well as avoid fluctuations in the prices. This corresponds to a bidder-optimal and stable assignment of bidders to items. An outcome (μ, p) is *stable* if a competitive equilibrium is reached. In a *competitive equilibrium* no bidder would prefer a different item or being unmatched to the one he is matched to under the current prices, i.e., every bidder is *envy-free*, and additionally the prices of all unmatched items are equal to their reserve prices. In a *bidder-optimal* outcome each bidder obtains his best utility among all envy-free outcomes. To simplify the bidding for the advertisers as well as to be able to compute an envy-free outcome with respect to the true values of the bidders, bidders should obtain their best possible utility if they reveal their true preferences to the mechanism. This property is called *truthfulness* or *incentive-compatibility*.

* Corresponding author.

E-mail address: veronika.loitzenbauer@univie.ac.at (V. Loitzenbauer).




Fig. 1. An input for which no bo-mechanism is truthful. The edges are labeled with the valuations the bidders report for the items. The maximum prices of both bidders are 5. The displayed prices are the minimum envy-free prices for the given inputs, respectively. In the left graph the bidders report their true values. In the right graph the first bidder reports a wrong valuation of 1 for the first item. Thus an envy-free outcome with respect to the reported valuations exists already at the initial prices of zero and both bidders obtain a higher utility. Note that the utility gain could be arbitrarily high.

A large body of prior work on this and related problems exists. We summarize below only the most closely related work, see [2] for a more complete overview.

A bidder-optimal outcome exists for all strictly monotonically falling and locally right-continuous utility functions [3] and, thus, for the model used in this work. We say that a mechanism that computes a bidder-optimal outcome for every given input is a *bo-mechanism*. By definition, the utilities of the bidders in a bidder-optimal outcome are unique. Thus one bo-mechanism is incentive-compatible for a given input if and only if all bo-mechanisms are incentive-compatible for that input. Note that we distinguish between incentive-compatibility for a specific input and incentive-compatibility for all inputs.

Without budgets every bo-mechanism is incentive-compatible and its outcome is stable [4,5]. The inclusion of budgets into the model implies discontinuities in the utility functions, which destroys these desirable properties in general [1,6,7]. Aggarwal et al. [1] were among the first who added budget constraints to quasi-linear utility functions. For inputs in *general position*, i.e., certain non-degenerate inputs, Aggarwal et al. provided an incentive-compatible mechanism that computes a bidder-optimal stable outcome in polynomial time.¹ Aggarwal et al. state in [1], with (v, m, r) being the input to the auction:

In essence, any auction (v, m, r) can be brought into general position by arbitrarily small (symbolic) perturbations. In practice this assumption is easily removed by using a consistent tie-breaking rule.

We provide an example that shows that neither a deterministic nor a randomized tie-breaking rule, as suggested above, leads to an incentive-compatible mechanism for *all* inputs. As for the undisturbed input, the gain from lying can be arbitrarily high. Instead, we use further randomization, this time of the prices, and give a mechanism based on randomized tie-breaking and on randomized price extraction that is truthful in expectation. However, as shown in [1,7], there are degenerate inputs for which no bo-mechanism is incentive-compatible, even if the outcome is stable [3]. Hence, our randomized mechanism cannot be a bo-mechanism. Our mechanism builds upon the results of Dütting et al. in [8], who showed that a modification of the Hungarian Method [9] computes (in polynomial time) a bidder-optimal envy-free outcome for every given input, i.e., is a bo-mechanism. If the input is in general position, this mechanism is incentive-compatible and the outcome is a competitive equilibrium [7].

General position is a quite restrictive condition on the input. Intuitively, it forbids that any two maximum prices can be reached simultaneously during an ascending price mechanism. For example, “symmetric inputs”, i.e., inputs where two bidders input exactly the same valuations and budgets, violate the general position condition. However, in practice such inputs can easily arise. Consider the example with symmetric bidders in Fig. 1. Both bidders have the same budget and prefer the first item over the second. If a bidder-optimal outcome is computed for the true values, the most desirable item is not sold and both bidders have a utility of zero. Furthermore, through lying one of the bidders can ensure that all desirable items are sold. Thus no bo-mechanism can be incentive-compatible. A good outcome in this situation would be to assign the most desirable item with equal probability to each bidder and to charge prices so that the expected utility of each bidder is at least the utility that the bidder could achieve through lying. This is exactly what our algorithm does.

More formally, we improve upon the known results in three ways. The requirement that the input is in general position is a sufficient but not necessary condition for the existence of a truthful bo-mechanism. Specifically, there exist inputs that are not in general position but for which the Modified Hungarian Method (and thus every bo-mechanism) is incentive-compatible. Furthermore, there exists no polynomial-time algorithm that determines whether an input is in general position. Our first contributions are a generalization of the general position condition called *rematch condition* that excludes fewer inputs and a polynomial time algorithm to determine whether an input fulfills the rematch condition. Thus our condition provides new insights on when a bo-mechanism cannot guarantee incentive-compatibility for a given input. We use these insights then in our second and third contributions.

¹ Chen et al. [6] defined *weak* and *strong* stability, where only for the latter a stable outcome is equal to a competitive equilibrium. The definition of a stable matching in [1] corresponds to weak stability. In their model, the utility of a bidder is set to a negative value if the price of the item strictly exceeds the maximum price of this bidder for the item. On the contrary, in [8] as well as in the definition used in this work, a utility becomes negative as soon as the maximum price is reached. For the utility functions used in [8] and in this work weak and strong stability coincide.

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات