

Bidder Optimal Assignments for General Utilities^{*}

Paul Dütting, Monika Henzinger, Ingmar Weber

Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland
{paul.duetting,monika.henzinger,ingmar.weber}@epfl.ch

Abstract. We study the problem of matching bidders to items where each bidder i has general, strictly monotonic utility functions $u_{i,j}(p_j)$ expressing her utility of being matched to item j at price p_j . For this setting we prove that a bidder optimal outcome always exists, even when the utility functions are non-linear and non-continuous. Furthermore, we give an algorithm to find such a solution. Although the running time of this algorithm is exponential in the number of items, it is polynomial in the number of bidders.

1 Introduction

In two-sided matching markets buyers are to be matched to items and the seller receives a monetary compensation from the buyers. Such markets have been studied for several decades [1, 2]. They have seen a surge of interest with the spread of sponsored search auctions [3, 4], where advertisers (bidders) are competing for the available advertising slots (items). Our research is also motivated by sponsored search but our results are not specific to this setting in any way.

A solution is *bidder optimal* if it gives each bidder the highest possible utility. For the case where each bidder i has a utility function $u_{i,j}(p_j)$ for item j that drops linearly in the price p_j it has long been known how to find such solutions [5, 6]. Recently, an algorithm was presented that also copes with per-bidder-item reserve prices (where a certain minimum price $p_j \geq r_{i,j}$ is required if bidder i is matched to item j) and per-bidder-item maximum prices (where bidder i can pay no more than $m_{i,j}$ for item j) [7]. This algorithm requires the input to be in “general position”, which e.g. requires that all reserve and maximum prices are different. Our results do not require this assumption.

In [8–10] the *existence* of bidder optimal solutions was shown for general, strictly monotonic utility functions, as long as the utility functions are *continuous*. However, no algorithm was given to find such a solution. For the special case of piece-wise linear functions an algorithm was presented in [11, 12], where the arguments used to prove termination leads to a time bound exponential in the number of items. The authors then argue that arbitrary continuous functions can be uniformly approximated by such piece-wise linear functions. However,

^{*} This work was conducted as part of a EURYI scheme award (see <http://www.esf.org/euryi/>).

neither the approximation accuracy nor the running time are analyzed and the relationship between the two is unclear. Also, such a uniform approximation does not exist for *discontinuous* utility functions.

Our main contributions are two-fold. First, we prove that even for general, strictly monotonic discontinuous utility functions a bidder optimal outcome always exists. Second, we give an algorithm to find such an outcome. The running time of this algorithm is polynomial in the number of bidders n but exponential in the number of items k . However, for sponsored search k can be viewed as constant or, at least, $k \ll n$.

Our model includes per-bidder-item reserve prices $r_{i,j}$ and also per-bidder reserve utilities o_i , usually referred to as *outside options*. We can also model per-bidder-item max prices $m_{i,j}$ with a discontinuous drop of $u_{i,j}(p_j)$ at the price $p_j = m_{i,j}$. Other settings that can be modeled are interest rates where, e.g., up to a certain point a bidder can still pay from her own pocket but for higher prices she has to borrow money from a bank, leading to a faster drop in utility. Similarly, settings where the bidder is “risk averse” and loses utility faster for higher prices due to an associated higher variance can be modeled.

Note that both for our and related previous results the utility function is a function of the price only. Other inter-item dependencies cannot be modeled. Such dependencies include a drop in utility if some other bidder gets a particular item or a positive utility only if a bidder gets a particular set of items. For a survey concerning such combinatorial auctions we refer the reader to [13].

2 The Assignment Problem

The problem is defined as follows: 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. Each bidder i has a *utility function* $u_{i,j}(p_j)$ for each item j expressing her utility of being matched to item j at price p_j . We assume that (i) the utility functions $u_{i,j}(\cdot)$ are strictly monotonically decreasing and (ii) for the outside options o_i (defined below) there exist *threshold values* $\bar{p}_{i,j}$ s.t. $u_{i,j}(\bar{p}_{i,j}) \leq o_i$. We do *not* assume that $u_{i,j}(\cdot)$ is (globally) continuous, but we do require that (iii) it is locally right-continuous, i.e. that $\forall x : \lim_{\epsilon \rightarrow 0^+} u_{i,j}(x + \epsilon) = u_{i,j}(x)$.¹

We want to compute a *matching* $\mu \subseteq I \times J$, where we require *all* bidders to be matched. To ensure that this is possible, even if $k < n$, we introduce symbolic *dummy items*. If bidder i is matched to a dummy item j then this represents that i is actually not matched. In the *proofs of existence* of bidder optimal solutions we will assume that there are n dummy items, one for each bidder, in our *algorithm* we directly deal with the case of an unmatched bidder as the running time would suffer significantly from an increase of the (small) number of items k to $n + k$. To distinguish between a match to a dummy and a real item we call the latter case *properly matched*. We use $\mu(i)$ to denote the

¹ At the end of Section 3 we show that all three of these requirements are necessary for a bidder optimal solution to exist.

item that is matched to bidder i in μ and $\mu(j)$ to denote the bidder, if any, that is matched to item j . Note that we do *not* require all items to be matched.

We say that a matching μ with prices $p = (p_1, \dots, p_k)$ is *feasible* if (i) $p_j \geq r_{i,j}$ for all $(i, j) \in \mu$, where $r_{i,j}$ is a *reserve price*, and (ii) $u_{i,\mu(i)}(p_{\mu(i)}) \geq o_i$, where o_i is an *outside option*.² As mentioned, we have one (conceptual) dummy item j for each bidder i s.t. $u_{i,j}(x) = o_i - x$ and $u_{i',j}(x) = o_{i'} - 1 - x$ for all $i' \neq i$, as well as $r_{i,j} = 0$ and $r_{i',j} = o_{i'}$ for all $i' \neq i$.³ We say that the outcome (μ, p) is *stable* if for all $(i, j) \in I \times J$ we have $u_{i,j}(p_j) \leq u_{i,\mu(i)}(p_{\mu(i)})$, i.e. each bidder gets an item which, at the prices of the outcome, is one of her first choices. We will also refer to prices p as feasible and/or stable if a corresponding feasible and/or stable matching μ exists. The triple $(u_{i,j}(\cdot), r_{i,j}, o_i)$ along with the implicitly assumed sets I and J constitute the *input* of the assignment problem. Note that our definition of stability is slightly *stronger* than the definition in [7] in the sense that our stability implies their stability, which also involves the $r_{i,j}$.

Our goal is to find a *bidder optimal* solution. We say that (μ, p) is bidder optimal if it is both feasible and stable and for every other feasible and stable (μ', p') we have that $u_{i,\mu(i)}(p_{\mu(i)}) \geq u'_{i,\mu'(i)}(p'_{\mu'(i)})$ for all i .

Given the input $(u'_{i,j}(\cdot), r_{i,j}, o_i)$ let (μ', p') be the outcome of a mechanism and analogously for $u''_{i,j}(\cdot)$. We say that the mechanism is *truthful* if for every bidder i with utility functions $u_{i,1}(\cdot), \dots, u_{i,k}(\cdot)$ and any two input matrices of utility functions u' and u'' with $u'_{i,j}(\cdot) = u_{i,j}(\cdot)$ for all j and $u'_{k,j}(\cdot) = u''_{k,j}(\cdot)$ for all $k \neq i$ and all j we have that $u_{i,\mu'(i)}(p'_{\mu'(i)}) \geq u_{i,\mu''(i)}(p''_{\mu''(i)})$ for all i . Note that this definition only involves the utility functions and not $r_{i,j}$ or o_i . We assume that the $r_{i,j}$, which are a property of the sellers, cannot be falsified by the bidders. Furthermore, it is easy to see that misreporting the o_i is not beneficial to i . Overreporting can only lead to a missed chance of being assigned an item and underreporting can lead to a utility below the true outside option.

3 Existence of a Bidder Optimal Solution

Theorem 1. *For any input $(u_{i,j}(\cdot), r_{i,j}, o_i)$ to the assignment problem there exists a bidder optimal outcome (μ^*, p^*) .*

So far this result was only known for *continuous* utility functions [8–10]. Our proof consists of the following steps: Lemma 1, which can be proven by contradiction, shows that lowest feasible and stable prices are sufficient for bidder optimality. Lemma 2 shows that any two feasible and stable outcomes (μ, p) and (μ', p') can be combined to give a new feasible and stable solution with prices $\min(p, p')$. Lemma 3 then asserts that the infimum p^* over all feasible and stable prices, if it exists, corresponds to a feasible and stable outcome (μ^*, p^*) . Finally,

² The second part of the feasibility definition is often referred to as *individual rationality* of the bidders. Similarly, the reserve price condition can be referred to as individual rationality of the sellers with bidder-dependent payoffs.

³ The intuition is that in any feasible and stable outcome the price of item j can be as low as zero and so bidder i will have a utility of at least o_i .

Lemma 4 finishes the proof as it gives the existence of at least one feasible and stable outcome, establishing the existence of an infimum by Lemma 2.

Lemma 1. *If (μ^*, p^*) is feasible and stable and $p_j^* \leq p_j$ for all j and any (μ, p) that is feasible and stable, then (μ^*, p^*) is bidder optimal.*

Lemma 2. *Any two outcomes (μ, p) and (μ', p') which are feasible and stable for the input $(u_{i,j}(\cdot), r_{i,j}, o_i)$ can be combined s.t. there exists a matching $\hat{\mu}$ which together with the prices $\hat{p} = \min(p, p')$ is feasible and stable for the same input.*

Even though the setting in [10] is for *continuous* utility functions and for a slightly weaker definition of stability involving both bidder *and* seller, their proof of this particular lemma (their Lemma 2) goes through unchanged. Concretely, it is shown that each bidder i gets item $\hat{\mu}(i)$ at a price $\hat{p}_{\hat{\mu}(i)}$ corresponding to $u_{i, \hat{\mu}(i)}(\hat{p}_{\hat{\mu}(i)}) = \max(u_{i, \mu(i)}(p_{\mu(i)}), u_{i, \mu'(i)}(p'_{\mu'(i)}))$. In other words, the two outcomes are stitched together in the best possible way for all bidders.

Lemma 2 implies that if there are any feasible and stable prices then there are unique infimum prices $p^* = \inf\{p : p \text{ are feasible and stable prices}\}$. It remains to show that (i) p^* corresponds to a feasible and stable outcome (μ^*, p^*) (Lemma 3) and (ii) that there is at least one feasible and stable outcome (Lemma 4).

Lemma 3. *If there exists a feasible and stable outcome (μ, p) matching all bidders, then there exists a feasible and stable outcome (μ^*, p^*) matching all bidders with lowest prices. I.e. no other feasible and stable outcome (μ', p') matching all bidders can have $p'_j < p_j^*$ for any j .*

Our proof uses the following definitions: Let $F_p \subseteq I \times J$ be the *first choice graph* at prices p which contains an edge from bidder i to item j if and only if $j \in \operatorname{argmax}_{j'} u_{i,j'}(p_{j'})$. Note that (μ, p) is *stable* if and only if $\mu \subseteq F_p$. Let $\tilde{F}_p \subseteq F_p$ denote the subset of *feasible* edges (i, j) where $p_j \geq r_{i,j}$ and $u_{i,j}(p_j) \geq o_i$.⁴ For $i \in I$ and $j \in J$ we define $F_p(i) = \{j : \exists (i, j) \in F_p\}$ and $F_p(j) = \{i : \exists (i, j) \in F_p\}$. For $T \subseteq I$ and $S \subseteq J$ we define $\tilde{F}_p(T) = \cup_{i \in T} \tilde{F}_p(i)$ and $\tilde{F}_p(S) = \cup_{j \in S} \tilde{F}_p(j)$. We define $\tilde{F}_p(i)$, $\tilde{F}_p(j)$, $\tilde{F}_p(T)$, and $\tilde{F}_p(S)$ analogously. We call a (possibly empty) set $S \subseteq J$ *strictly overdemanded* for prices p wrt $T \subseteq I$ if (i) $\tilde{F}_p(T) \subseteq S$ and (ii) $\forall R \subseteq S, R \neq \emptyset : |\tilde{F}_p(R) \cap T| > |R|$. Using Hall's Theorem [14] one can show that a feasible and stable matching exists for given prices p if and only if there is no strictly overdemanded set of items S .

Proof. If we assume that there exists at least one feasible and stable outcome, then Lemma 2 shows that there exist unique infimum prices p^* .

For a contradiction suppose that there is no matching μ^* s.t. (μ^*, p^*) is feasible and stable. Then, by Hall's Theorem, there must be a set T of bidders s.t. $\tilde{F}_{p^*}(T)$ is strictly overdemanded for prices p^* wrt T .

In *any* feasible and stable outcome $(\hat{\mu}, \hat{p})$ we have $\hat{p}_j \geq p_j^*$ for all items j and, thus, the overdemand for the items in $\tilde{F}_{p^*}(T)$ can only be resolved if (i) at least

⁴ The second feasibility condition is redundant as if $u_{i,j}(p_j) < o_i$ then bidder i strictly prefers her dummy item whose price can always be assumed to be 0 in any bidder optimal solution. See Lemma 4.

one of the bidders $i \in T$ has a feasible first choice item $j \in J \setminus F_{p^*}(T)$ under \hat{p} or (ii) for some item $j \in F_{p^*}(T) \setminus \tilde{F}_{p^*}(T)$ we have that $\hat{p}_j \geq r_{i,j}$. Case (i) corresponds, for each pair $(i, j) \in T \times J \setminus F_{p^*}(T)$, to a price increase relative to p^* of $s_j^i = \inf\{x \geq 0 : u_{i,j}(p_j^* + x) \leq \max_{j' \in J \setminus F_{p^*}(T)} u_{i,j'}(p_{j'}^*)\}$, which is > 0 and contained in the set itself as $u_{i,j}(\cdot)$ is right-continuous.⁵ Case (ii) corresponds, for each pair $(i, j) \in I \times F_{p^*}(T) \setminus \tilde{F}_{p^*}(T)$, to a price increase relative to p^* of $f_j^i = r_{i,j} - p_j^*$, which is also > 0 . Let $\delta_j^i = \min(s_j^i, f_j^i)$ if $j \in F_{p^*}(i) \setminus \tilde{F}_{p^*}(i)$ and let $\delta_j^i = s_j^i$ otherwise. Then $\Sigma_\delta = \min_i \sum_j \delta_j^i$ is a *lower bound* on the sum of the price increases for *any* feasible and stable outcome $(\hat{\mu}, \hat{p})$.

Lemma 2, however, shows that for any $\epsilon > 0$ there exist feasible and stable prices p' s.t. $|p'_j - p_j^*| < \epsilon$ for all items j . For $\epsilon = \Sigma_\delta/|J|$ this gives a contradiction to the fact that the price increases corresponding to Σ_δ were required by *any* feasible and stable solution. We conclude that there exists at least one matching μ^* s.t. (μ^*, p^*) is feasible and stable. \square

Lemma 4. *For the assignment problem there are unique lowest stable prices p^* s.t. any other feasible and stable outcome (μ, p) has $p_j \geq p_j^*$ for all $j \in J$.*

Proof. By Lemma 3 we know that we only have to show the existence of *some* stable outcome (μ^*, p^*) . Set $p_j^* = \max_i(\bar{p}_{i,j})$ for the non-dummy items j (where the $\bar{p}_{i,j}$'s are the threshold values defined above) to ensure that all bidders have (at most) a utility of o_i and then match all bidders to dummy items at a price of 0. As no real item is matched and all utilities are o_i this is feasible. It is also stable as all real items have prices so high that no bidder strictly prefers them over a dummy item. \square

Finally, we show that all three conditions on the utility functions (see Section 2) are required to guarantee the existence of a bidder optimal solution:

Strict monotonicity: Consider a setting with three bidders and two items and the following utility functions: $u_{1,1}(x) = u_{3,2}(x) = 1 - x$, $u_{1,2}(x) = u_{3,1}(x) = -x$ and $u_{2,1}(x) = u_{2,2}(x) = 2$ if $x \leq 1$ and $u_{2,1}(x) = u_{2,2}(x) = 3 - x$ otherwise. All $r_{i,j} = o_i = 0$. Then one feasible and stable outcome is $\mu = \{(1, 1), (2, 2)\}$, $p = (0, 1)$ whereas another is $\mu = \{(2, 1), (3, 2)\}$, $p = (1, 0)$. In neither of the two settings can the price for the item with price 0 be lowered any further without upsetting stability. The first outcome is strictly preferred by bidder 1, whereas the second is strictly preferred by bidder 2.

Eventual drop in utility to o_i : Consider two bidders who both have the utility function $u_{i,1}(x) = 1/(1+x)$ for a single item. Again, $r_{i,1} = o_i = 0$. Then no matter how large p_1 is, both bidders will still strictly prefer the item over being unmatched.

Right continuity: Consider two bidders who both have the following utility function for a single item: $u_{i,1}(x) = 2 - x$ if $x \leq 1$ and $u_{i,1}(x) = -x$ otherwise. Then a price of $p_1 \leq 1$ will not be stable, as both bidders strictly prefer the item over being unmatched. So any stable price needs to satisfy $p_1 > 1$ and this

⁵ This no longer holds without the requirement of right-continuity as discussed at the end of this section.

set no longer contains its infimum. If we change the first condition of the utility function to $x < 1$, ensuring right-continuity, then the price $p_1 = 1$ is stable, even though the item cannot be assigned to either of the two bidders.

Truthfulness. If the reserve prices r_j depend only on the items, then Theorem 2 in [10] shows that any mechanism that computes a bidder optimal outcome is truthful. If the reserve prices $r_{i,j}$ also depend on the bidders, then this is no longer true: Consider a setting with two bidders and two items and the following utility functions, reserve prices, and outside options: $u_{1,1}(x) = 6 - x$, $u_{1,2}(x) = 5 - x$, $u_{2,1}(x) = 6 - x$, $u_{2,2}(x) = 6 - x$, $r_{1,1} = 2$, $r_{1,2} = 0$, $r_{2,1} = 1$, $r_{2,2} = 2$, and $o_1 = o_2 = 0$. The bidder optimal outcome is $\mu = \{(1, 1), (2, 2)\}$, $p = (2, 2)$. If Bidder 2 reports $u_{2,2}(x) = 0$, then the bidder optimal outcome is $\mu = \{(1, 2), (2, 1)\}$, $p = (1, 0)$. This gives Bidder 2 a strictly higher utility.

4 Algorithm for General Utilities

Here we present an algorithm which *directly* computes a bidder optimal outcome for general utility functions, as opposed to settling for a piecewise-linear approximation [11, 12]. Our algorithm assumes that computations of the “inverse utility function” $u_{i,j}^{-1}(x) = \min\{p : u_{i,j}(p) \leq x\}$ take constant time. If $u_{i,j}(\cdot)$ is continuous then $u_{i,j}^{-1}(\cdot)$ is indeed the inverse function. More generally, it is merely a one-sided inverse function satisfying $u_{i,j}^{-1}(u_{i,j}(p)) = p$.

Description of the Algorithm. Try out all possible matchings in which all bidders, but not necessarily all items, are matched. For a particular matching μ try all possible ways of ordering the (up to) k properly matched bidders. For each ordering initialize *lower bounds* on the prices $\forall j : b_j = \min_i r_{i,j}$. Execute the following steps for every properly matched bidder i according to the current ordering: Fix the price of item $\mu(i)$ to be $p_{\mu(i)} = \max(b_{\mu(i)}, r_{i,\mu(i)})$ and update $b_{\mu(i)} = p_{\mu(i)}$. If $u_{i,\mu(i)}(p_{\mu(i)}) < o_i$ then abort the check of the particular ordering.⁶ If there exists a previously considered bidder i' where $u_{i,\mu(i)}(p_{\mu(i)}) < u_{i',\mu(i')}(b_{\mu(i')})$ then also abort the check of the particular ordering.⁷ If neither of these two cases happens, update the vector of price bounds by setting $\forall j : b_j = \max(b_j, u_{i,j}^{-1}(u_{i,\mu(i)}(b_{\mu(i)})))$. Once all properly matched bidders have been considered, go through all bidders matched to dummy items. If for such a bidder i there exists a matched item $j = \mu(i')$ where $u_{i,j}(p_j) > o_i$ then abort. Otherwise, set $p_j = \max(p_j, u_{i,j}^{-1}(o_i))$ for all items j .⁸ After considering all matchings and orderings, output any pair (μ, p) corresponding to the lowest found prices p .

Lemma 5. *If the above algorithm terminates without aborting then the returned output (μ, p) is feasible and stable.*

⁶ The price $p_{\mu(i)}$ is now already too high to allow a feasible matching of i to $\mu(i)$.

⁷ We assumed that i' got her item at the current price bound $p_{\mu(i')} = b_{\mu(i')}$ but this price would no longer be stable wrt to i .

⁸ Note that these price updates can only affect unmatched items. For all matched items $\mu(i)$ the price was finalized when their bidder i was considered.

Proof. When a particular matched bidder i is considered the price $p_{\mu(i)}$ will not rise anymore until abortion or until a new matching/ordering is tried. During the consideration of i we update the prices of all other items to ensure that i 's utility is highest for item $\mu(i)$. As the prices for the *other* items might increase further this ensures that on termination we have $u_{i,j}(p_j) \leq u_{i,\mu(i)}(p_{\mu(i)})$ for all j . Prices are feasible as whenever we match a bidder i to an item j we set $p_j = \max(b_j, r_{i,j})$. Utilities are feasible as when considering i we only continue if i 's utility is non-negative and this does not change until termination. As for unmatched bidders, we ensure that they have a utility of no more than o_i for any real item and so they are stable. \square

If one requires $p_j \geq b_j$ in addition to feasibility, then the b_j can simply be absorbed into the reserve price by setting $r'_{i,j} = \max(r_{i,j}, b_j)$. So whenever the algorithm updates the bounds b_j while considering bidder i this is conceptually a restart of the algorithm with a new input, where reserve prices are changed and bidder i and item $\mu(i)$ have been removed.

Lemma 6. *If we know that for the input $(u_{i,j}(\cdot), r_{i,j}, o_i)$ a feasible and stable outcome exists, then every bidder optimal outcome (μ, p) matches at least one bidder i to an item j at price $p_{\mu(i)} = r_{i,\mu(i)}$.*

Proof. Theorem 1 in [10], whose proof does not require continuity, shows that given a bidder optimal outcome there is no other feasible (but not necessarily stable) solution where all bidders are strictly better off. But if *all* prices of matched items j satisfied $p_j > r_{\mu(j),j}$ then by setting all prices for the matched items to $p_j = r_{\mu(j),j}$ *all* bidders would benefit as we assume that utility functions are strictly decreasing in the price and prices p_j are still feasible. \square

Lemma 6 is the main ingredient to reconstruct the lowest stable prices for a given matching. It lets us remove one bidder i at a time where at each step we can ensure that $p_{\mu(i)}$ is not affected by removals of bidders in future iterations.

Theorem 2. *If for μ currently being tried by the algorithm there exist feasible and stable prices then the algorithm will find the lowest stable prices p for μ .*

Proof. By induction on the number of bidders n . If $n = 1$ then in the bidder optimal outcome the (unique) bidder i gets item $\mu(i) \in \operatorname{argmax}_{j'} u_{i,j'}(\max(b_{j'}, r_{i,j'}))$ at price $\max(b_{\mu(i)}, r_{i,\mu(i)}) = r_{i,\mu(i)}$ and the price of all items $j' \neq \mu(i)$ must be at least $u_{i,j'}^{-1}(u_{i,\mu(i)}(p_{\mu(i)}))$ to guarantee stability for bidder i . These are also the prices the algorithm will return if it does not abort. Now suppose the result holds for all $t \leq n$ and we want to prove the claim for $n + 1$. By Lemma 6 we know that at least one item j is sold for the price $p_j = \max(b_j, r_{i,j})$ to some bidder i , where b_j is the current lower bound. As we try out all possible orderings, this bidder i will also be selected first in an ordering and hence obtain item j at price p_j . To ensure that bidder i does not prefer a different item we must have increased the price bounds of all items $j' \neq j$ to $b_{j'} = \max(b_{j'}, u_{i,j'}^{-1}(u_{i,j}(p_j)))$ for *any* stable solution in which bidder i gets item j at price p_j . These are exactly the lower price bounds the algorithm ensures. Given these lower price bounds

and the fact that prices only increase, bidder i is stable for the prices computed for this ordering. Hence we can remove bidder i and are left with a new instance with n bidders and one less item. By induction, we find the lowest feasible and stable prices for this problem. As (i) we try out *all* possible orderings of bidders and (ii) for the correct ordering we obtain the lowest feasible and stable prices p respecting the initial lower bound $b_j = \min_i r_{i,j}$ for all j , we obtain the theorem. \square

Theorem 2 together with Lemma 1 shows that the algorithm outputs a bidder optimal outcome.

Running Time. There are $O(k!(n+k)^k)$ different matchings of n bidders to up to k items and there are $O(k!) = O(k^k)$ permutations of the up to k properly matched bidders. Computing the price updates for a given matching-bidder pair takes time $O(nk)$. Hence the overall running time is $O((n+k)^k \cdot k^{2k+1} \cdot n)$.

References

1. Shapley, L.S., Shubik, M.: The assignment game: The core I. *Journal of Game Theory* **29** (1972) 111–130
2. Roth, A.E., Sotomayor, M.A.O.: Two-sided matching: A study in game-theoretic modeling and analysis. Cambridge University Press (1990)
3. Edelman, B., Ostrovsky, M., Schwarz, M.: Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review* **97**(1) (2007) 242–259
4. Lahaie, S., Pennock, D.M., Saberi, A., Vohra, R.V.: Sponsored Search Auctions. In: *Algorithmic Game Theory*. Cambridge University Press (2007) 699–716
5. Leonard, H.B.: Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy* **91** (1983) 461–479
6. Demange, G., Gale, D., Sotomayor, M.: Multi-item auctions. *Journal of Political Economy* **94**(4) (1986) 863–72
7. Aggarwal, G., Muthukrishnan, S., Pál, D., Pál, M.: General auction mechanism for search advertising. In: *World Wide Web Conference*. (2009) 241–250
8. Crawford, K., Knoer: Job matching with heterogenuos firms and workers. *Econometrica* **49** (1981) 437–450
9. Quinzii, M.: Core and competitive equilibria with indivisibilities. *International Journal of Game Theory* **13**(1) (1984) 41–60
10. Demange, G., Gale, D.: The strategy structure of two-sided matching markets. *Econometrica* **53**(4) (1985) 873–888
11. Alkan, A.: Existence and computation of matching equilibria. *European Journal of Political Economy* **5** (1989) 285–296
12. Alkan, A., Gale, D.: The core of the matching game. *Games and economic behavior* **2**(3) (1990) 203–212
13. Cramton, P., Shoham, Y., Steinberg, R.: *Combinatorial Auctions*. MIT Press (2006)
14. Hall, P.: On representatives of subsets. *Journal of the London Mathematical Society* **10** (1935) 26–30