General Auction Mechanism for Search Advertising

Gagan Aggarwal[*] S. Muthukrishnan[†]
Dávid Pál[‡] Martin Pál[§]

Keywords
game theory, online auctions, stable matchings

ABSTRACT
Internet search advertising is often sold by an online automated auction. Typically a fixed number of slots k is available, and have to be allocated among n advertisers each of whom desires to display an ad. Several mechanisms to price the slots and allocate them to advertisers have been studied, including variants of the Generalized Second Price (GSP) mechanism, as well as mechanisms from the Vickrey-Clarke-Groves (VCG) family that are designed to be truthful for profit maximizing bidders. Extensions of these mechanisms to account for things like position constraints and reserve prices have also been proposed. Many any of these auction mechanisms can be viewed as computing a *bidder-optimal stable matching* with suitably defined preferences of the auction participants. This allows us to apply the theory of stable matchings pioneered by Gale and Shapley [13] to search auctions.

In this paper, we define a general stable matching model with money in which many of the existing and new auctions can be expressed. We show that in this model, a bidder-optimal stable matching always exists (under a mild non-degeneracy assumption), and that a mechanism based on computing such matching is truthful. Importantly, we give an algorithm to compute a bidder-optimal matching in polynomial time of $O(nk^3)$. As a result, we obtain the first known, truthful mechanism for a variety of bidders.

1. INTRODUCTION
Search engine companies like Yahoo!, Google or MSN display advertisements on web pages with search results or various kind of content. Typically, the web page has a number of separately marked *slots* reserved for the ads, and an auction is run to determine the set of ads to be displayed out of a pool of eligible ads that match the search query or content of the page. Typically, each advertiser must submit an ad together with a bid ahead of time. The purpose of the auction is to determine an assignment of ads to slots, as well as determine how much each winning advertiser should pay for displaying her ad. The payment scheme may charge the advertiser each time the ad is shown, or only in the event that a user clicks on the ad. More involved schemes charge the advertiser only when the user performs a pre-specified action in response to the ad, such as making a purchase from the advertiser's online store.

GSP is a popular auction mechanism used by major search engines. It assumes that there is a natural ordering on the ad slots (such as top to bottom or left to right). It asks each advertiser to submit a bid (this can be a per-impression, per-click or per-action bid), ranks the advertisers in decreasing order of their bids, and assigns the top k advertisers to the k available slots.

From the point of view of a bidder, GSP enjoys simple semantics: each slot has a price (that depends on the bids of other bidders), and the bidder is simply assigned the highest slot whose price does not exceed her bid. If the bidder is a "maximum price" bidder, in that her goal is simply to get the highest possible slot while making sure that her cost does not exceed a certain threshold, it is her dominant strategy to submit the threshold as her bid. This is true for any charging scheme, per impression, per click, or any generally defined action.

The basic GSP mechanism can be extended in various ways, for example by scaling each bid by a bidder-specific multiplier or by introducing a minimum price. These extensions are useful tools for the search engine: for example, giving a higher multiplier to better ads increases the overall ad quality, while giving a higher multiplier to ads likely to be clicked on in a per-click auction may improve revenue.

The class of VCG mechanisms [22, 7, 15] has been designed as a truthful mechanism for an important class of "profit maximizing" bidders. In the context of advertising auctions, we assume that a profit maximizing bidder i derives a certain (expected) value v_{ij} from her ad being placed in position j. The advertiser's profit is then equal to the (expected) profit from placing her ad, minus the (expected) payment she is charged for the slot. In a simpler model where the click probabilities of different slots are known to the auctioneer, the advertiser may simply submit her value of a click. The auctioneer then finds a maximum-weight assignment (which maximizes the overall value to all bidders), and determines payments according to a formula that incents the bidders to bid truthfully.

Our paper is motivated by the observation that the above mechanisms really compute a *bidder-optimal stable matching* between a

[*]Google, Inc., Mountain View, CA. gagana@google.com
[†]Google, Inc., New York, NY. muthu@google.com
[‡]David R. Cheriton School of Computer Science, University of Waterloo, Waterloo, ON, Canada. dpal@cs.uwaterloo.ca. Work done during summer 2007 internship at Google New York.
[§]Google, Inc., New York, NY. mpal@google.com

set of slots and a set of advertisers. We will define a model of bidder preferences and stable matchings with side payments in which this observation can be formalized. Indeed, proving that the outcomes of these mechanisms are bidder-optimal stable matchings in our model is not difficult. With this observation in mind, we set out to study a general scenario that accomodates the different set of bidders as well as the incorporates various features that search engines employ; this scenario not only includes mechanisms described above, but also other extensions.

In our model, each bidder specifies a maximum price she is willing to pay for an impression in a given slot, as well as a value of that slot to her (which may be greater than her willingness to pay for the slot). In addition, the search engine may specify a reserve (minimum) price for each bidder and each slot. We show that in this model, a bidder-optimal stable matching exists and can be computed in $O(nk^3)$ time, where n is the number of bidders and k the number of slots. Further, we show that the mechanism that elicits bidder's preferences and computes a bidder-optimal stable matching based on them is truthful in that for every bidder whose preferences can be expressed within our model, it is a dominant strategy to express her true preferences. In particular, this shows that there is a truthful auction mechanism for any combination of "maximum price" and "profit maximizing" bidders who can pay per imperssion, per click or per action, and can have constraints on the set of positions they may to appear in. Thus, what results is the first-known general auction mechanism that is truthful for a diverse set of bidders.

2. RELATED WORK

The GSP auction is the major vehicle for selling ads on the internet. It has been observed that although it is not truthful for "profit maximizing" bidders, it does have a Nash equilibrium that is efficient and its resulting prices are equal to VCG prices, see e.g. [11, 2]. A variant of GSP in which the bidder can specify the lowest (maximum) acceptable position has been proposed in [3], which also has a Nash equilibrium equivalent to a suitably defined VCG auction. The recent manuscript [12] explores the effect of adding minimum prices to GSP. The GSP mechanism assumes a fixed ordering on available slots. With increasingly complex web page layouts, this assumption is no longer universally valid; if there are several regions on the page where ads can be placed, it is not obvious that one slot or region should be universally preferred to another by all advertisers. Also, allowing the advertiser to choose the action she is charged for (impression vs click) destroys incentive properties of the auction, even if as is natural, one considers multiplying each per-click bid by the probability of a click). As long as different positions have different click probabilities, it is not difficult to come up with examples where switching from a per-click bid to per-impression bid or vice versa lowers the overall cost to the bidder while giving her the same or better slot.

The general class of VCG mechanisms follows from works [22, 7, 15]. For an overview of the VCG mechanism applied to sponsored search, see e.g. [1, 2]. VCG is a natural mechanism, but bidders may find it unintuitive to interpret the prices they are charged. Also, it does not directly support maximum-price bidders who naturally fit into GSP.

The stable matching model has been introduced by Gale and Shapley [13] in 1962 and has been studied extensively since then. The monograph [19] gives a great overview of important results in the area; we only mention themes that are directly relevant to our work. In the basic model introduced in [13], a set I of men is to be matched to a set J of women in a one to one fashion. Each man has a preference ordering on the set of women, and each woman has a preference ordering on the set of men. The goal is to find a matching that is stable in that there is no man and a woman in which the man would prefer the woman to his partner in the matching and vice versa. Gale and Shapley [13] give a "deferred acceptance" algorithm to compute a man-optimal stable matching.

The stable matching has been generalized to allow side payments between members of a matched pair. In such models, each participant has a preference relation on the set of possible potential (partner, payment) pairs. In models where the side payments are allowed to be arbitrary real numbers, the preference relation is often given by a set of utility functions, one for each man-woman pair expressing the man's preferences, and one woman's. A model with utility functions that are linear in money has been studied by [20, 18, 9]. It has long been known that in the linear utility model, a bidder-optimal stable matching is equivalent to VCG allocation and prices, [16, 6].

Arbitrary (non-linear) increasing continuous utility functions were considered in [8, 9, 4, 5]. These models crucially depend on the utility functions being continuous and defined on the whole set $\mathbb{R}$, an assumption we have to drop in our paper. The paper [8] shows that even in such a general model, there *exists* a bidder-optimal stable matching, but no algorithm was given to find it. Moreover they show that in a mechanism based on the man-optimal matching, it is weakly dominant for each man to reveal his true utility function.

Our paper builds heavily on this body of work. A feature that distinguishes our work is that in order to model "maximum price" bidders and reserve prices we need to introduce preferences that can not be expressed as continuous, strictly monotone utility functions. This seemingly innocent change introduces technical difficulties and makes the model harder to work with. Still, we are able to transfer the main structural results to our model. Under the assumption that the bidder preferences are in a "general position", we can still prove the existence of a bidder-optimal matching, and we give a very efficient algorithm to find it. The general position assumption can be lifted by adopting a suitable tie-breaking rule, which allows us to show that the stable matching mechanism is truthful. This gives us the first known truthful mechanisms for a variety of bidders.

3. THE MAX-VALUE MODEL

Our model consists of the set $I = \{1, 2, \ldots, n\}$ of bidders and the set $J = \{1, 2, \ldots, k\}$ of slots. We use letter i to denote a bidder and letter j to denote a slot.

Each bidder i has a *value* $v_{i,j}$ for each slot j how much is that slot worth to her, and a *maximum price* $m_{i,j}$ she is able and willing to pay for the slot. To motivate why $v_{i,j}$ and $m_{i,j}$ might be different, consider buying a house whose value you estimate significantly higher than your bank. While your value for the house is high, the amount of money your bank is willing to lend you is lower. Allowing the bidder to specify both a value and a maximum is also needed to model the GSP auction. In addition to bidder preferences, the seller specifies for each bidder i and each slot j a *reserve price* $r_{i,j}$.

For simplicity we assume that the reserve prices are known to the bidders in advance. For each i and each j we assume that $r_{i,j} \geq 0$, $v_{i,j} \geq 0$, $m_{i,j} \leq v_{i,j}$. If bidder i is interested in the slot j he specifies $m_{i,j} \geq r_{i,j}$. Otherwise, if bidder i has no interest in slot j he specifies negative $m_{i,j}$. We denote by v, m, r the $n \times k$ matrices with entries $v_{i,j}, m_{i,j}, r_{i,j}$ respectively. We refer to the triple (v, m, r) as an *auction instance* or simply *auction*.

We wish to find an assignment of slots to bidders, and compute how much each winning bidder should pay for her slot. To study strategic behavior of the bidders, we need to specify their relative preference for possible outcomes.

Bidder Preferences. We assume that each bidder is indifferent among various outcomes as long as her assigned slot (if any) and payment is the same. Let us define the utility of a bidder i who is offered a slot j at price p as follows. If $p \leq m_{i,j}$, we set $u = v_{ij} - p$. If $p > m_{i,j}$, we set $u = -1$. This utility, interpreted as a function of the price, is not continuous at $p = m_{i,j}$. If the bidder is not matched (at zero price), her utility is 0. Given a choice between slot j_1 at price $q_1 \leq m_{i,j_1}$ and slot j_2 at price $p_2 \leq m_{i,j_2}$, the bidder prefers the offer with higher utility, and is indifferent among offers that have the same utility. In particular, the bidder prefers to be not matched to being matched to a slot j at price that exceeds her maximum price m_{ij}. The bidder is indifferent between being matched with utility 0 and not being matched.

3.1 Stable Matching

We formalize the notion of a matching in the following definitions.

DEFINITION 1 (MATCHING). *A* matching *is a triple* (u, p, μ), *where* $u = (u_1, u_2, \dots, u_n)$ *is a non-negative* utility vector, $p = (p_1, p_2, \dots, p_k)$ *is a non-negative* price vector, *and* $\mu \subseteq I \times J$ *is a set of bidder-slot pairs such that no slot and no bidder occurs in more than one pair.*

If a pair $(i, j) \in \mu$, we say that bidder i is *matched* to slot j. We use $\mu(i)$ to denote the slot matched to a bidder i, and $\mu(j)$ to denote to denote the bidder matched to a slot j. Bidders i and slots j that do not belong to any pair in μ are said to be *unmatched.*

DEFINITION 2 (FEASIBLE MATCHING). *A matching* (u, p, μ) *is said to be* feasible *for an auction* (v, m, r), *whenever for every* $(i, j) \in \mu$,

$$p_j \in [r_{i,j}, m_{i,j}] \ , \tag{1}$$

$$u_i + p_j = v_{i,j} \ , \tag{2}$$

and for each unmatched bidder i is $u_i = 0$ *and for each unmatched slot j is* $p_j = 0$.

DEFINITION 3 (STABLE MATCHING). *A matching* (u, p, μ) *is* stable *for an auction* (v, m, r) *whenever for each* $(i, j) \in I \times J$ *at least one of the following inequalities holds:*

$$u_i + p_j \geq v_{i,j} \ , \tag{3}$$

$$p_j \geq m_{i,j} \ , \tag{4}$$

$$u_i + r_{i,j} \geq v_{i,j} \ . \tag{5}$$

A pair $(i, j) \in I \times J$ *which does not satisfy any of the three inequalities is called* blocking.

Geometric interpretation of inequalities (3), (4), (5) is explained in Figure 1. Note that if a bidder i is not interested in a slot j, then (4) is trivially satisfied.

A feasible matching does not have to be stable, and a stable matching does not have to be feasible. However, we will be interested in matchings that are both stable and feasible. More specifically, we will be interested in a particular matching (u^*, p^*, μ^*) that is stable, feasible, and is, with respect to each bidder's preferences, superior to any other feasible stable matching (u, p, μ). It is surprising that such a matching exists. Its existence for other simpler models, e.g., with continuous utility vs. price curves or other preference relations, is a core result of the theory of stable matchings.

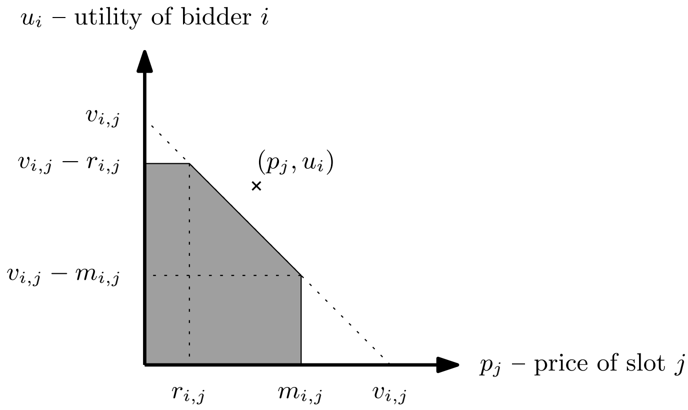

Figure 1: Matching is stable whenever for each bidder $i \in I$ and each slot $j \in J$ the point with coordinates (p_j, u_i) lies outside the gray region.

3.2 Our Results

One of our technical contributions is a proof of the existence of stable, feasible matching in our model.

THEOREM 4 (EXISTENCE OF BIDDER-OPTIMAL MATCHING). *If the auction* (v, m, r) *is in a "general position", it has a unique* bidder-optimal stable matching.

We defer the precise definition of general position to Definition 12. 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 propose an auction mechanism that, for any reserve prices r specified by the auctioneer, and any valutions v and maximum prices m specified by the bidders, computes the bidder-optimal stable matching (u^*, p^*, μ^*), assigns the slots to the bidders according to μ^* and charges the matched bidders prices p^* correspondingly. We call this mechanism the *stable matching mechanism*. We study this mechanism from the game-theoretic perspective and prove that the mechanism is truthful.

THEOREM 5 (TRUTHFULNESS). *The stable matching mechanism is a truthful mechanism for bidders in the max-value model. That is, submitting her true vectors* v_i *and* m_i *is a dominant strategy for each bidder i.*

Our final contribution is an algorithm that computes the bidder-optimal stable matching.

THEOREM 6. *There is an algorithm that finds the bidder-optimal stable matching in the max-value model in time* $O(nk^3)$. *Thus, there is a truthful mechanism for max-value bidders that can be implemented in this running time.*

Taken together, these results yield the first known truthful mechanism that is efficient to implement for all bidders who can be represented in our max-value model. This not only includes the well-known GSP or VCG and their variants by search engines, but much more.

4. MODELING ADVERTISING AUCTIONS

In this section, we will present examples of auction mechanisms commonly used in sponsored search. We will show how to model these mechanisms in our max-value model. In the next section we give examples of novel combined mechanisms that can be implemented in our model.

4.1 Existing Mechanisms

GSP pay-per-impression. In a Generalized Second Price auction, each advertiser i submits a single number b_i as her bid, which is the maximum amount she is willing to pay for displaying her ad. The auctioneer orders bidders in decreasing order of their bids, and assigns the first k advertisers to the k available slots in this order. The i-th allocated advertiser pays amount equal to the $(i+1)$-st bid for each impression.

GSP pay-per-click. An alternative is to charge the advertiser only in the event of a click on her ad. The bid b_i is interpreted as a maximum the advertiser is willing to pay for a click. Again, the advertisers are ordered by their per-click bid, and each allocated advertiser pays the next highest bid in the event of a click. In a quality-weighted variant, the ads are ordered by the product of their quality score q_i and bid b_i; the i-th advertiser pays $b_{i+1}\frac{q_{i+1}}{q_i}$ in the event of a click. Note that the expected cost per impression $b_{i+1}\frac{q_{i+1}}{q_i}\mathrm{ctr}_{i,i}$ depends not only on the next highest bid but also on the position, as long as the probability $\mathrm{ctr}_{i,j}$ of clicking on the ad i in position j depends on the position. Thus, there is no direct way to translate a per-click bid to a per-impression bid, without looking at the competitor's bids.

The VCG mechanism for profit-maximizing bidders. In a variant of the VCG mechanism considered e.g. in [2], each bidder i states her value V_i for a click. The auctioneer derives the expected value of each slot $v_{i,j} = V_i \cdot \mathrm{ctr}_{i,j}$ for that bidder by using an estimate $\mathrm{ctr}_{i,j}$ of the probability that the ad i would be clicked on if placed in position j. The auctioneer computes a maximum-weight matching in the bipartite graph on bidders and positions with $v_{i,j}$ as edge weights. The maximum weight matching μ^* gives the final allocation. For pricing, the VCG formula sets the price per impression of slot $j = \mu^*(i)$ to be $p_j = \sum_{k \in I \setminus \{i\}} v_{k,\mu'(k)} - v_{k,\mu^*(k)}$ where μ' is a maximum-weight matching with the set of bidders $I \setminus \{i\}$. Note that the per-impression price p_j can be translated to a per-click price by charging bidder i price $p_j/\mathrm{ctr}_{i,j}$ for each click. (Similar translation can be done for a generally defined user action other than a click, as long as the probability of the action can be estimated.)

For each of the above mechanisms, we define a corresponding type of bidder in the max-value model.

Max-per-impression bidder has a target cost per impression b_i. She prefers paying b_i or less per impression to any outcome where she pays more than b_i. Given that her cost per impression is at most b_i, she prefers higher (with lower index) position to lower position. Given a fixed position, she prefers paying lower price to higher price.

A max-per-impression bidder i can be translated into the max-value model by setting her $m_{i,j} = b_i$ for all positions $j \in J$, and setting her value $v_{i,j} = M(k+1-i)$ where M is a sufficiently large number ($M > b_i$ is enough).

Max-per-click bidder differs from a max-per-impression bidder in that she is not willing to pay more than b_i per click. We translate her per-click bid into our framework using predicted click probabilities: set $m_{i,j} = b_i \cdot \mathrm{ctr}_{i,j}$ for $i \in I$ and $v_{i,j} = M(k+1-i)$ where $M > b_i \max_j \mathrm{ctr}_{i,j}$.

Profit-maximizing bidder seeks the position and payment that maximizes her expected profit (value from clicks minus payment). If we assume that her value per click is V_i, such bidder is modeled by setting $v_{i,j} = m_{i,j} = V_i \cdot \mathrm{ctr}_{i,j}$.

We formalize the correspondence between the mechanisms and corresponding bidder types in the following theorem.

THEOREM 7. *The outcome (allocation and payments) of a (1) per-impression GSP, (2) per-click GSP, (3) VCG auction, respectively is a bidder-optimal stable matching for a set of (1) max-per-impression bidders, (2) max-per-click bidders, (3) profit-maximizing bidders, respectively.*

PROOF. Part (3) of the theorem has been observed by multiple authors including [16]. Chapter 7 of [19] as well as [6] discuss the relationship of the VCG mechanism for assignments and stable matchings.

We give a proof for part (1), per-impression GSP. The proof of part (2) for per-click GSP is similar. For simplicity, we assume that $n > k$ and all reserve prices are zero. Let $b_1 > b_2 > \cdots > b_n$ be the per-impression bids of the bidders. Without loss of generality, the bidders are ordered by decreasing order of their bids. (By the general position assumption, assume bids are distinct.)

Recall that we encode a max-per-impression bidder by setting $v_{i,j} = M(k-j+1)$ and $m_{i,j} = b_i$. The matching produced by the GSP auction is as follows: the matched pairs are $\mu = \{(1,1),(2,2),\ldots,(k,k)\}$, bidder's utilities $u_i = M(k-i+1) - b_{i+1}$ for $1 \le i \le k$, $u_i = 0$ for $i > k$, and prices $p_i = b_{i+1}$ for $i = 1,2,\ldots,k$.

It is easy to verify that the matching is feasible and stable according to Definitions 2 and 3.

First we show that any feasible matching in which the assignment is different from μ is not stable. Indeed, such a matching (u', p', μ') must have a bidder $i \le k$ such that i was not allocated a slot among the first i slots, and a slot $j \le i$ that is either unmatched or matched to some bidder $i' > i$.

From feasibility we have that $p_j = 0$ if slot j is unmatched and $p_j \le b_{i'}$ in case it is matched. In either case, $p_j < b_i$. Also, since bidder i is matched to some slot $j' > i$ (or unmatched), we know that $u'_i \le v_{i,j'} = M(k-j'+1)$. We now claim that (i,j) is a blocking pair. Since $v_{i,j} - u'_i \ge M[(k-j+1)-(k-j'+1)] \ge M$, inequalities (3) and (5) are violated, and since $p'_j < b_i$, inequality (4) is violated as well.

Now consider any matching with the assignment $\mu = \{(1,1),\ldots,(k,k)\}$. It is easy to verify that in order to be stable, it must be that $p_i \ge b_{i+1}$, otherwise the pair $(i+1,i)$ would be a blocking pair. Hence the matching with prices $p_i = b_{i+1}$ has the lowest possible prices and hence is bidder-optimal. $\square$

Minimum prices. Some search engines impose a minimum price r_i for each ad (for example, based on perceived quality of the ad). In GSP, only bidders whose bid is above the reserve price can participate. The allocation is in decreasing order of bids, and each bidder pays the maximum of her reserve price and the next bid. Minimum GSP prices are easily translated to the max-value model by setting $r_{ij} = r_i$ (if paying per impression) or $r_{ij} = r_i \cdot \mathrm{ctr}_{i,j}$ (if paying per click). Our model allows for separate reserve prices for different slots (e.g. higher reserve price for certain premium slots) that are not easily implemented in the GSP world.

4.2 New Auction mechanisms

Let us give a few examples of new auction mechanisms that are special cases of the max-value model.

GSP with arbitrary position preferences. Consider an advertiser i who wishes for her ad to appear only in certain slots. For example, [3] propose a GSP variant in which each bidder has the option to specify a prefix of positions $\{1,2,\ldots,\beta_i\}$ for some β_i she is interested in and exclude the remaining slots. Also, tools like Google's Position Preference allow the advertiser to specify arbitrary position intervals $[\alpha_i,\beta_i]$. We are however not aware of any published work that discusses more sophisticated position preferences. One would imagine that in the world of content advertising where there may be multiple areas designed for ads on a single page, having a richer language in which to express the preferences over slots would be beneficial to the advertiser. Such preferences are readily expressible in the max-value model. □

Combining click and impression bidders in GSP. Since both pay per click and pay per impression models are widely used in practice, it is useful to have a way of combining these two bidding modes. This can be easily done by computing a stable matching for a mixed pool of bidders. The following simpler approach is not appropriate, as it does not have the proper incentive structure.

Suppose we allow each bidder i to specify both a maximum price b_i, as well as a payment type $\tau_i \in \{\mathcal{I}, \mathcal{C}\}$. A naive combined auction orders bidders by decreasing b_i. Each advertiser with $\tau_i = \mathcal{I}$ is charged the next highest bid b_{i+1} for showing the ad. Each advertiser with $\tau_i = \mathcal{C}$ is charged b_{i+1} in the event that the user clicks on the ad. Note, this scheme gives advertisers a strong incentive to report $\tau_i = \mathcal{C}$ regardless of their true type (as long as the probability of user clicking is less than 1).

To offset this incentive, the auctioneer may introduce multipliers $0 < q_\mathcal{C} < 1$ and $q_\mathcal{I} = 1$ and set the effective bid of each bidder to be $b_i^{\text{eff}} = b_i q_{\tau_i}$. In the modified GSP auction, bidders are be sorted by their effective bid. Each bidder i who reports type $\tau_i = \mathcal{I}$ is charged b_{i+1}^{eff} for each impression, while each bidder reporting $\tau_i = \mathcal{C}$ is charged $b_{i+1}^{\text{eff}}/q_\mathcal{C}$ in the event of a click.

For any value of $0 < q_\mathcal{C} < 1$, there is a simple instance in which some bidder can gain by misreporting her type. Let ctr_1 and ctr_2 be the probability that an user will click on an ad in position 1 and 2 respectively. Assume this probability is the same for all ads, and that $\text{ctr}_1 > \text{ctr}_2$. Suppose that the first slot is won by a bidder of type $\mathcal{I}$, the second slot is won by a bidder of type $\mathcal{C}$, and that there is at least one more bidder with positive bid. If $q_\mathcal{C} > \text{ctr}_2$, the bidder in the second position can lower her overall cost while keeping the same position by reporting type $\mathcal{C}$ and keeping the same effective bid. On the other hand, if $q_\mathcal{C} < \text{ctr}_1$, bidder in the first position can lower her cost by reporting type $\mathcal{I}$, and adjusting her bid so that her effective bid stays the same. □

Diverse bidders. There are many types of bidders with different goals. Some like to think in terms of a maximum price per click or impression. Some prefer to target only certain positions (e.g. top of the page) for consistency or branding reasons. Others try to maximize their profit and are able to estimate the value of a specific user action. Each bidder may specify her goal in a language familiar to her. We are not aware of any prior research on auction mechanisms for such diverse set of bidders. □

5. ALGORITHM FOR COMPUTING THE BIDDER-OPTIMAL MATCHING

In this section we present an algorithm that for given auction (v, m, r) (in general position) computes the bidder-optimal stable matching. The algorithm starts with an empty matching $(u^{(0)}, p^{(0)}, \mu^{(0)})$ which is defined as follows. Utility of each bidder i is $u_i^{(0)} = B$, where B is a large enough number, such that $B > \max\{v_{i,j} \mid (i,j) \in I \times J\}$. Price of each slot j is $p_j^{(0)} = 0$. There are no matched pairs, i.e. $\mu^{(0)} = \emptyset$.

In each iteration, the algorithm finds an augmenting path, and updates the current matching $(u^{(t)}, p^{(t)}, \mu^{(t)})$ to the next matching $(u^{(t+1)}, p^{(t+1)}, \mu^{(t+1)})$. The algorithm stops when no more updates can be made, and outputs the current matching $(u^{(T)}, p^{(T)}, \mu^{(T)})$ at the end of the last iteration. We now describe an iteration in more detail. To do so, we introduce the concept of an update graph.

DEFINITION 8 (UPDATE GRAPH). *Given an auction (v, m, r), the* update graph *for a matching (u, p, μ) is a directed weighted bipartite multigraph with partite sets I and $J \cup \{j_0\}$, where j_0 is the* dummy *slot. The update graph consists of five types of edges. For each bidder i and each slot $j \in J$ there is*

- *a* forward edge *from i to j with weight $u_i + p_j - v_{i,j}$, if $p_j \in [r_{i,j}, m_{i,j})$;*
- *a* backward edge *from j to i with weight $v_{i,j} - u_i - p_j$, if $(i,j) \in \mu$,*
- *a* reserve-price edge *from i to j with weight $u_i + r_{i,j} - v_{i,j}$, if $u_i + r_{i,j} > v_{i,j}$ and $m_{i,j} > r_{i,j}$,*
- *a* maximum-price edge *from i to j with weight $u_i + m_{i,j} - v_{i,j}$, if $u_i + m_{i,j} > v_{i,j}$ and $m_{i,j} > r_{i,j}$,*
- *a* terminal edge *from i to j_0 with weight u_i if $u_i > 0$.*

An *alternating path* in the update graph starts with an unmatched bidder vertex i_0 with $u_{i_0} > 0$, follows a sequence of forward and backward edges, and ends with a reserve-price, maximum-price or terminal edge. We place the restriction that all vertices of the alternating path must be distinct, with the possible exception that the last vertex is allowed to appear once again along the path. The weight $w(P)$ of an alternating path P is the sum of weights of its edges.

Let $(u^{(t)}, p^{(t)}, \mu^{(t)})$ be a matching and $G^{(t)}$ be the corresponding update graph. A single iteration of the algorithm consists of the following steps.

1. If there is no alternating path, stop and output the current matching. Otherwise, let P be an alternating path in $G^{(t)}$ of minimum weight. Let $w^{(t)}(P)$ denote its weight, and let
$$P = (i_0, j_1, i_1, j_2, i_2, \ldots, j_\ell, i_\ell, j_{\ell+1}) \qquad \text{for some } \ell \geq 0 .$$
2. Let $d^{(t)}(i_0, y)$ be the length of the shortest path in $G^{(t)}$ from i_0 to any vertex y, using only forward and backward edges. If a vertex y is not reachable from i_0, $d^{(t)}(i_0, y) = \infty$.
3. Compute utility updates for each bidder $i \in I$. The vector $u^{(t+1)}$ gives the final utilities for the iteration.
$$u_i^{(t+1)} = u_i^{(t)} - \max\left(w^{(t)}(P) - d^{(t)}(i_0, i),\ 0\right) \tag{6}$$
4. Compute price updates for each slot $j \in J$.
$$p_j^{(t+)} = p_j^{(t)} + \max\left(w^{(t)}(P) - d^{(t)}(i_0, j),\ 0\right) \tag{7}$$
The final prices $p_j^{(t+1)}$ are equal to $p_j^{(t+)}$ with one exception. In case the last edge of P is a reserve-price edge, we set the price of slot $j_{\ell+1}$, the last vertex of P to be $p^{(t+1)} = \max(p^{(t+)}, r_{i_\ell, j_{\ell+1}})$.
5. Update the assignment $\mu^{(t)}$ along the alternating path P to obtain the new assignment $\mu^{(t+1)}$.

We have not specified how should the set of assignment edges be updated. Before we do that, let us state two invariants maintained by the algorithm.

(A1) The matching $(u^{(t)}, p^{(t)}, \mu^{(t)})$ is stable for the auction (v, m, r).

(A2) For every matched pair $(i, j) \in \mu^{(t)}$, $u_i^{(t)}$ and $p_j^{(t)}$ satisfy (1) and (2).

An important consequence of invariant (A1) is that forward edges have non-negative weight. Indeed, it can be easily checked that a forward edge with a negative weight would be blocking pair. Invariant (A2) guarantees that backward edges have zero weight. Similarly, invariant (A2) implies that the weight of every backward edge must be zero. Finally, each reserve-price, maximum-price and terminal edges has non-negative weight by definition.

LEMMA 9. *All edge weights in each update graph $G^{(t)}$ are non-negative.*

With non-negative edge weights, single-source shortest paths can be computed using Dijkstra's algorithm in time proportional to the square of the number of vertices reachable from the source. Since no unmatched vertex is reachable from any other vertex, there are at most $2k$ reachable vertices at any time, thus the shortest alternating path P and distances $d^{(t)}(i_0, y)$ can be computed in time $O(k^2)$.

Finally, let us deal with updating the assignment μ. Since the alternating path alternates between using forward (i.e. non-matching) and backward (i.e. matching) edges, a natural move is to remove all the matching edges of P and replace them by non-matching edges of P. Care must be taken however to take into account the special nature of the last edge of P as well as the fact that the last vertex of P may be visited twice. We consider three cases:

Case 1: P ends with a terminal edge, i.e. $j_{\ell+1}$ is the dummy slot. Flip matching and non-matching edges along the whole length of P. Bidder i_ℓ ends up being unmatched, and for $x = 0, 1, \ldots, \ell-1$, bidder i_x will be matched to slot j_{x+1}.

Case 2: P ends with a maximum-price edge. Consider two subcases:

(a) $j_{\ell+1} = j_\ell$. This means that the price bidder i_ℓ was matched to reached his maximum price. Flip matching an non-matching edges along P. This leaves bidder i_ℓ unmatched, and for $x = 0, 1, \ldots, \ell - 1$ bidder i_x is matched with slot i_{x+1}.

(b) Otherwise, the maximum price was reached on a non-matching edge. Keep the matching unchanged. That is, $\mu^{(t+1)} = \mu^{(t)}$.

Case 3: P ends with a reserve-price edge. This is the most complex case. Consider three subcases:

(a) Item $j_{\ell+1}$ is unmatched in $\mu^{(t)}$. This case increases the size of the matching. For $x = 0, 1, \ldots, \ell$, match bidder i_x with slot j_{x+1} .

(b) Item $j_{\ell+1}$ is matched in $\mu^{(t)}$ and the reserve price $r_{i_\ell, j_{\ell+1}}$ offered by bidder i_ℓ does not exceed the current price $p_{j_{\ell+1}}^{(t+)}$ of the slots. Keep the matching unchanged, that is, $\mu^{(t+1)} = \mu^{(t)}$.

(c) Item $j_{\ell+1}$ is matched in $\mu^{(t)}$ to some bidder $i_{\ell+1}$ and $r_{i_\ell, j_{\ell+1}} > p_{j_{\ell+1}}^{(t+)}$. If P is a path, that is, if P does not visit slots j_{i_ℓ} twice, we simply unmatch bidder $i_{\ell+1}$, and flip matching and non-matching edges of P. (This keeps the size of the matching the same, as bidder i_0 gets matched and bidder $i_{\ell+1}$ unmatched.)

If P visits $j_{\ell+1}$ twice, it must be that $j_{\ell+1} = j_d$ for some d. Note that it is not the case that $d = \ell$, since this would mean that i_ℓ was matched to $j_{\ell+1}$. This is impossible because the reserve price on this edge has been reached just now. This way, the end of P forms a cycle with at least 2 bidders and 2 slots. We flip the matching and non-matching edges along the cycle, but leave the rest of P untouched. This leaves bidder i_x matched to slot j_{x+1}, for $x = d, d + 1, \ldots, \ell$.

6. ANALYSIS OF OUR ALGORITHM

In this section we show that the algorithm from Section 5 indeed computes a bidder-optimal stable matching. It is not obvious that a stable matching even exists for any auction instance (v, m, r). Our algorithm provides a constructive proof of this fact. An alternate proof of existence can be done using limit arguments and the deferred acceptance algorithm on a sequence of discretizations of bidder's preference lists. The details are deferred to the full version of this paper.

LEMMA 10. *The matching $(u^{(T)}, p^{(T)}, \mu^{(T)})$ computed by the matching algorithm is feasible and stable.*

PROOF. Stability follows directly from invariant (A1). Feasibility follows from invariant (A2) and the fact that since there are no alternating paths, it must be that $u_i^{(T)} = 0$ for every unmatched bidder i. □

We shall prove the invariants later in this section. While a feasible stable matching always exists, there may not always be a bidder-optimal matching, as the following example shows. Consider the case of a single slot and two bidders with identical maximum bids. There are two stable matchings. In each matching, the slot is allocated to one of the bidders at maximum price. Each matching is preferred by one bidder over the other, hence there is no matching preferred by both of them.

Fortunately it turns out that the example above is degenerate, and that a bidder-optimal matching exists for every non-degenerate, or "general position" auction. To make this precise, we need the following two definitions.

DEFINITION 11 (AUCTION GRAPH). *The* auction graph *of an auction (v, m, r) is a directed weighted bipartite multigraph with partite sets I and $J \cup \{j_0\}$, where j_0 is the* dummy *slot. The auction graph contains five types of edges. For each bidder i and each slot $j \in J$ there exist*

- *a* forward edge *from i to j with weight $-v_{i,j}$,*
- *a* backward edge *from j to i with weight $v_{i,j}$,*
- *a* reserve-price edge *from i to j with weight $r_{i,j} - v_{i,j}$,*
- *a* maximum-price edge *from i to j with weight $m_{i,j} - v_{i,j}$,*
- *a* terminal edge *from i to j_0 with weight* 0.

DEFINITION 12 (GENERAL POSITION). *An auction (v, m, r) is in* general position *if for every bidder i, no two alternating walks in the auction graph that start at bidder i, follow alternating forward and backward edges and end with a distinct edge that is either a reserve-price, maximum-price or terminal edge, have the same weight.*

Any auction (v, m, r) can be brought into general position by a symbolic perturbation. In the algorithm implementation, this can be also achieved by breaking ties lexicographically by the identity of the final edge of the walk.

The next section is pretty technical. It establishes invariants needed to prove Theorem 6.

6.1 Invariants

Besides invariants (A1) and (A2) introduced in Section 5, we claim three more invariants.

(A3) Each unmatched slot has zero price.

(B1) if a bidder i is interested in slot j and $u_i^{(t)} + m_{i,j} = v_{i,j}$, then $(i,j) \notin \mu^{(t)}$.

(B2) If a bidder i is interested in a slot j and $u_i^{(t)} + r_{i,j} = v_{i,j}$, then $(i,j) \in \mu^{(t)}$ or $p_j^{(t)} \geq r_{i,j}$.

All the five invariants are proved by induction on t. Invariants (B1) and (B2) are technical and we omit their proofs in this version of the paper. However, we use them in the induction step to prove the first three invariants. Both (B1) and (B2) rely on the general position assumption.

PROOF OF THE INVARIANTS. The base case, $t = 0$, is readily verified. Invariant (A1) follows from that $u_i^{(0)} = B$ for all $i \in I$, $p_j^{(0)} = 0$ for all $j \in J$, and hence (3) is satisfied. Invariants (A2) and (A3) hold trivially.

Let us prove that $(u^{(t+1)}, p^{(t+1)}, \mu^{(t+1)})$ satisfies (A3). Note that $p^{(t+1)} \geq p^{(t)}$. The slots matched in $\mu^{(t)}$ remain matched in $\mu^{(t+1)}$, at most one additional slot is matched in $\mu^{(t+1)}$. The remaining slots are not reachable from i_0 in $G^{(t)}$, since for any such slot j, $p_j^{(t)} = 0$ and for any $i \in I$, $r_{i,j} > 0$ by the general position assumption, thus there is no forward edge to j. Hence the price of any such slot j remains zero.

Let us prove that $(u^{(t+1)}, p^{(t+1)}, \mu^{(t+1)})$ satisfies (A1). We consider three cases for any pair $(i,j) \in I \times J$:

Case 1: $p_j^{(t)} \in [r_{i,j}, m_{i,j})$. $(u^{(t)}, p^{(t)}, \mu^{(t)})$ is stable by the induction hypothesis and hence $u_i^{(t)} + p_j^{(t)} \geq v_{i,j}$. If $d^{(t)}(i_0, i) \geq w^{(t)}(P)$, then $u_i^{(t+1)} = u_i^{(t)}$ and $p_j^{(t+1)} \geq p_j^{(t)}$, thus $u_i^{(t+1)}$ and $p_i^{(t+1)}$ satisfy (3).

On the other hand, if $d^{(t)}(i_0, i) < w^{(t)}(P)$, then

$$u^{(t+1)} = u_i^{(t)} - (w^{(t)}(P) - d^{(t)}(i_0, i)) \ , \quad (8)$$
$$p_j^{(t+1)} \geq p^{(t+)} \geq p_j^{(t)} + (w^{(t)}(P) - d^{(t)}(i_0, j)) \ . \quad (9)$$

Since from i to j there is a forward edge in $G^{(t)}$,

$$d^{(t)}(i_0, j) \leq d^{(t)}(i_0, i) + (u_i^{(t)} + p_j^{(t)} - v_{i,j}) \ . \quad (10)$$

We add (8) to (9), subtract (10), and we get that $u_i^{(t+1)}$ and $p_j^{(t+1)}$ satisfy (3).

Case 2: $p_j^{(t)} \geq m_{i,j}$. Since $p_j^{(t+1)} \geq p_j^{(t)}$, (4) holds for $p_j^{(t+1)}$. (This case applies also if i is not interested in j.)

Case 3: $p_j^{(t)} < r_{i,j}$ and i is interested in j. $(u^{(t)}, p^{(t)}, \mu^{(t)})$ is stable by the induction hypothesis and hence $u_i^{(t)}$ satisfies (5). If $d^{(t)}(i_0, i) \geq w^{(t)}(P)$, then $u_i^{(t+1)} = u_i^{(t)}$ and hence $u_i^{(t+1)}$ also satisfies (5).

On the other hand, if $d^{(t)}(i_0, i) < w^{(t)}(P)$, then

$$u_i^{(t+1)} = u_i^{(t)} - (w^{(t)}(P) - d^{(t)}(i_0, i)) \ . \quad (11)$$

We claim that in $G^{(t)}$ there is reserve-price edge from i to j and thus

$$w^{(t)}(P) \leq d^{(t)}(i_0, i) + (u_i^{(t)} + r_{i,j} - v_{i,j}) \ . \quad (12)$$

To prove the existence of the reserve-price edge we show that $u_i^{(t)} + r_{i,j} > v_{i,j}$. The non-strict inequality holds since $u_i^{(t)}$ satisfies (5). The strictness follows since, by the induction hypothesis, $(u^{(t)}, p^{(t)}, \mu^{(t)})$ satisfies (A2) and (B2) .

By subtracting (12) from (11) we get that $u^{(t+1)}$ satisfies (5).

First, let us prove that $(u^{(t+1)}, p^{(t+)}, \mu^{(t)})$ satisfies (A2). Consider any pair $(i,j) \in \mu^{(t)}$. In $G^{(t)}$ there is a backward edge from j to i. By induction hypothesis, $(u^{(t)}, p^{(t)}, \mu^{(t)})$ satisfies (A2) and hence the backward edge has zero weight. Hence

$$d^{(t)}(i_0, i) = d^{(t)}(i_0, j) \ . \quad (13)$$

Therefore, from the updates (6), (7) follows $u_i^{(t+1)} + p_j^{(t+)} = u_i^{(t)} + p_j^{(t)}$ and hence (1) remains to hold.

If $w^{(t)}(P) \leq d^{(t)}(i_0, i)$, then $p_j^{(t+)} = p_j^{(t)}$ and thus (2) remains satisfied by $p_j^{(t+)}$. On the other hand, if $w^{(t)}(P) > d^{(t)}(i_0, i)$, then by the update (7) for prices

$$p_j^{(t+)} = p_j^{(t)} + (w^{(t)}(P) - d^{(t)}(i_0, j)) \ . \quad (14)$$

We also claim that there exists maximum-price edge from i to j and thus

$$w^{(t)}(P) \leq d^{(t)}(i_0, i) + (u_i^{(t)} + m_{i,j} - v_{i,j}) \ . \quad (15)$$

To prove the existence of the maximum-price edge we show that $u_i^{(t)} + m_{i,j} > v_{i,j}$. The non-strict inequality holds since $p_j^{(t)} \leq m_{i,j}$ and thus $u_i^{(t)} + m_{i,j} \geq u_i^{(t)} + p_j^{(t)} = v_{i,j}$ since by the induction hypothesis $(u^{(t)}, p^{(t)}, \mu^{(t)})$ satisfies (A2). Strictness follows since, by the induction hypothesis, $(u^{(t)}, p^{(t)}, \mu^{(t)})$ satisfies (B1).

Summing (13), (15), (14) and canceling common terms gives $p^{(t+)} \leq (u_i^{(t)} + p_j^{(t)} - v_{i,j}) + m_{i,j} = m_{i,j}$, where $u_i^{(t)} + p_j^{(t)} - v_{i,j} = 0$ follows from the induction hypothesis. Hence, since $p^{(t+)} \geq p^{(t)} \geq r_{i,j}$, (2) remains to hold for $p_j^{(t+)}$.

Finally, let us prove that $(u^{(t+1)}, p^{(t+1)}, \mu^{(t+1)})$ satisfies (A2). For any pair $(i,j) \in \mu^{(t)} \cap \mu^{(t+1)}$ we have already done it, since $p_j^{(t+1)} = p_j^{(t+)}$. It remains to consider pairs in $\mu^{(t+1)} \setminus \mu^{(t)}$. Let $P = (i_0, j_1, i_1, \ldots, j_\ell, i_\ell, j_{\ell+1})$ be the alternating path used to obtain $\mu^{(t+1)}$ from $\mu^{(t)}$. Any pair $(i,j) \in \mu^{(t+1)} \setminus \mu^{(t)}$ is an edge lying P and has the form $(i,j) = (i_x, j_{x+1})$. We consider two cases.

Case 1: $x < \ell$. In this case $(i,j) = (i_x, j_{x+1})$ is a forward edge and has weight $u_i^{(t)} + p_j^{(t)} - v_{i,j}$, and since it lies on a minimum-weight path,

$$d^{(t)}(i_0, j) = d^{(t)}(i_0, i) + (u_i^{(t)} + p_j^{(t)} - v_{i,j}) \ . \quad (16)$$

Since $w^{(t)}(P) \geq d^{(t)}(i_0, i)$ and $w^{(t)}(P) \geq d^{(t)}(i_0, j)$, the updated quantities are

$$u_i^{(t+1)} = u_i^{(t)} - (w^{(t)}(P) - d^{(t)}(i_0, i)) \ , \quad (17)$$
$$p_j^{(t+1)} = p_j^{(t)} + (w^{(t)}(P) - d^{(t)}(i_0, j)) \ . \quad (18)$$

The equality (1) for $u_i^{(t+1)}$ and $p_j^{(t+1)}$ follows by summing (17), (18) and subtracting (16).

Let us verify that $p_j^{(t+1)}$ satisfies (2). Since (i,j) is a forward edge, $p_j^{(t)} \in [r_{i,j}, m_{i,j})$. By the induction hypothesis $(u^{(t)}, p^{(t)}, \mu^{(t)})$ is stable, thus $u_i^{(t)} + p_j^{(t)} \geq v_{i,j}$, hence $u_i^{(t)} + m_{i,j} > v_{i,j}$ and consequently in $G^{(t)}$ there is a maximum-price edge from i to j of weight $u_i^{(t)} + m_{i,j} - v_{i,j}$. Therefore

$$w^{(t)}(P) \leq d^{(t)}(i_0, i) + u_i^{(t)} + m_{i,j} - v_{i,j} \ . \quad (19)$$

We add (18) to (19) and from that we subtract (16), we cancel common terms and we have $p_j^{(t+1)} \le m_{i,j}$. The verification of (2) for $p_j^{(t+1)}$ is finished by observing that $p_j^{(t+1)} \ge p_j^{(t)} \ge r_{i,j}$.

Case 2: $x = \ell$. Since we assume that $(i,j) = (i_\ell, j_{\ell+1})$ belongs to $\mu^{(t+1)} \setminus \mu^{(t)}$, it can be neither a terminal edge nor a maximum-price edge, and thus it must be a reserve-price edge and has weight $u_i^{(t)} + r_{i,j} - v_{i,j}$. By the same argument $p_j^{(t+)} \le r_{i,j}$, hence $p^{(t+1)} = r_{i,j}$ and clearly satisfies (2). Observe that

$$u^{(t+1)} = u^{(t)} - (w^{(t)}(P) - d^{(t)}(i_0, i)) ,$$
$$w^{(t)}(P) = d^{(t)}(i_0, i) + (u_i^{(t)} + r_{i,j} - v_{i,j}) .$$

Subtracting the two equations shows that $u_i^{(t+1)}$ and $p_j^{(t+1)}$ satisfy (1). ☐

6.2 Running Time

We bound the number of iterations by $O(nk)$ in the claim below. Since each iteration can be implemented in time $O(k^2)$, this gives us overall running time $O(nk^3)$.

LEMMA 13. *The matching algorithm finishes after at most $n(2k+1)$ iterations.*

PROOF. Consider the number of edges in the update graph. Initially, the graph $G^{(0)}$ has at most nk reserve-price, nk maximum-price and n terminal edges. We claim that in each iteration, the number of edges in the update graph is reduced by one. Since the algorithm must stop when there are no more edges left, this bounds the total number of iterations.

Consider an iteration t of the algorithm. We claim that in the alternating path $P = (i_0, j_1, i_1, \dots, j_\ell, i_\ell, j_{\ell+1})$, the last edge $(i,j) = (i_\ell, j_{\ell+1})$ will not appear in the update graph $G^{(t+1)}$. This is easily verified by considering three cases:

Case 1: If (i,j) is a terminal edge, then $w^{(t)}(P) = d^{(t)}(i_0, i) + u_i^{(t)}$ and hence $u_i^{(t+1)} = u_i^{(t)} - (w^{(t)}(P) - d^{(t)}(i_0, i)) = 0$.

Case 2: If (i,j) is a maximum-price edge, then $w^{(t)}(P) = d^{(t)}(i_0, i) + (u_i^{(t)} + m_{i,j} - v_{i,j})$ and hence $u_i^{(t+1)} + m_{i,j} = u_i^{(t)} - (w^{(t)}(P) - d^{(t)}(i_0, i)) + m_{i,j} = v_{i,j}$.

Case 3: If (i,j) is a reserve-price edge, then $w^{(t)}(P) = d^{(t)}(i_0, i) + (u_i^{(t)} + r_{i,j} - v_{i,j})$ and hence $u_i^{(t+1)} + r_{i,j} = u_i^{(t)} - (w^{(t)}(P) - d^{(t)}(i_0, i)) + r_{i,j} = v_{i,j}$.

The utilities never increase and the prices never decrease throughout the algorithm, thus the edge $(i_\ell, j_{\ell+1})$ does not appear in any update graph $G^{(t')}$ for any $t' > t$. ☐

6.3 Bidder Optimality

LEMMA 14. *Let (v, m, r) be an auction in general position, and let (u', p', μ') be any feasible stable matching. Then in any iteration t of the matching algorithm, we have that $u'_i \le u_i^{(t)}$ for all $i \in I$ and $p'_j \ge p_j^{(t)}$ for all $j \in J$.*

Without loss of generality assume that (u, p, μ) is such that there does *not* exist a pair $(i,j) \in \mu$ such that $p_j = m_{i,j}$. If there was such a pair, then we can decrease prices of some of the items and increase utilities of some of the bidders such that $p_j < m_{i,j}$. This is possible because of the general position assumption. See full version of the paper.

We prove Lemma 14 by induction on t. The base case, $t = 0$, trivially holds true, since by feasibility of (u', p', μ'), $p'_j \ge 0$ for all $j \in J$ and $u'_i \le B$ for all $i \in I$. In the inductive case, assume that $u^{(t)} \ge u'$ and $p^{(t)} \le p'$. We first prove that

PROPOSITION 15. $u^{(t+1)} \ge u'$ *and* $p^{(t+)} \le p'$.

We look "continuously" at updates (6) and (7). For that purpose we define for each $i \in I$ a continuous non-increasing function $u_i(x)$,

$$u_i(x) = u_i^{(t)} - \max\left(x - d^{(t)}(i_0, i),\ 0\right) ,$$

and for each $j \in J$ a continuous non-decreasing function $p_j(x)$,

$$p_j(x) = p_j^{(t)} + \max\left(x - d^{(t)}(i_0, j),\ 0\right) .$$

Clearly, $u^{(t+1)} = u(w^{(t)}(P))$ and $p^{(t+)} = p(w^{(t)}(P))$. To prove that $u^{(t+1)} \ge u'$ and $p^{(t+)} \le p'$, suppose by contraction that there exists $y \in [0, w^{(t)}(P)]$ such that either $u_i(y) < u'_i$ for some $i \in I$ or $p_j(y) > p'_j$ for some $j \in J$. We choose infimal such y. Clearly, $u(y) \ge u'$, $p(y) \le p'$ and $y < w^{(t)}(P)$. Consider the sets

$$I' = \{i \in I \mid u_i(y) = u'_i \text{ and } d^{(t)}(i_0, i) \le y\} ,$$
$$J' = \{i \in J \mid p_j(y) = p'_j \text{ and } d^{(t)}(i_0, j) \le y\} .$$

CLAIM 16. *Each slot $j \in J'$ is matched in $\mu^{(t)}$ to some $i \in I'$.*

PROOF OF THE CLAIM. Let $j \in J'$. If j was unmatched, then either $d^{(t)}(i_0, j) = w^{(t)}(P)$ or $d^{(t)}(i_0, j) = \infty$; however both options contradict the choice of y and that $j \in J'$. Thus j is matched to some $i \in I$, hence in $G^{(t)}$ there is a backward edge from j to i and thus $d^{(t)}(i_0, i) = d^{(t)}(i_0, j)$ and therefore $u_i(y) + p_j(y) = v_{i,j}$. Further, invariants (A2) and (B1) imply that $p_j^{(t)} \in [r_{i,j}, m_{i,j})$. Consequently, there is a maximum-price edge from i to j, $w^{(t)}(P) \le d^{(t)}(i_0, i) + (u_i^{(t)} + m_{i,j} - v_{i,j})$, and hence $p'_j = p_j(y) < p_j^{(t+)} = p^{(t)} + (w^{(t)}(P) - d^{(t)}(i_0, j)) \le m_{i,j}$. Therefore $p'_j \in [r_{i,j}, m_{i,j})$, and since (u', p', μ') is stable, $u'_i + p'_j \ge v_{i,j}$ and hence $u_i(y) = v_{i,j} - p_j(y) = v_{i,j} - p'_j \le u'_i$. On the other hand, by infimality of y, $u_i(y) \ge u'_i$. Thus $i \in I'$. ☐

CLAIM 17. *Each bidder $i \in I'$ is matched in μ' to some $j \in J'$.*

PROOF OF THE CLAIM. Since in $G^{(t)}$ there is a terminal edge from i to the dummy slot, $w^{(t)}(P) \le d^{(t)}(i_0, i) + u_i^{(t)}$. Hence

$$\begin{aligned} u'_i = u_i(y) &= u_i^{(t)} - (y - d^{(t)}(i_0, i)) \\ &> u_i^{(t)} - (w^{(t)}(P) - d^{(t)}(i_0, i)) \ge 0 , \end{aligned}$$

and thus bidder i is matched in μ' to some slot $j \in J$.

By feasibility of (u', p', μ'), $p'_j \in [r_{i,j}, m_{i,j}]$. By the assumption made at the beginning $p_j \ne m_{i,j}$. Therefore in $G^{(t)}$ there is a forward edge from i to j and thus

$$d^{(t)}(i_0, j) \le d^{(t)}(i_0, i) + (u_i^{(t)} + p_j^{(t)} - v_{i,j}) . \tag{20}$$

Clearly, since $i \in I'$,

$$u_i(y) = u_i^{(t)} - (y - d^{(t)}(i_0, i)) . \tag{21}$$

By the price update rule

$$p_j(y) \ge p_j^{(t)} + (y - d^{(t)}(i_0, j)) . \tag{22}$$

We add (21) to (22) and subtract from that (20) and we obtain

$$p_j(y) \ge v_{i,j} - u_i(y) .$$

Hence, since by feasibility of (u', p', μ'), $u'_i + p'_j = v_{i,j}$, we have

$$p_j(y) \ge v_{i,j} - u_i(y) = v_{i,j} - u'_i = p'_j .$$

Recalling that $p(y) \le p'$ we see that $p_j(y) = p'_j$.

Subtracting (21) from (20) and cancelling common terms we have

$$d^{(t)}(i_0, j) \le y + (u_i(y) + p_j^{(t)} - v_{i,j}) \ .$$

We upper-bound the right side of the inequality using that $u_i(y) = u'_i$, $p_j^{(t)} \le p_j(y)$ and $u'_i + p'_j = v_{i,j}$ and we have

$$d^{(t)}(i_0, j) \le y + (u'_i + p'_j - v_{i,j}) = y \ .$$

Thus $j \in J'$. □

From the two claims it follows that $|I'| = |J|'$ and that $\mu^{(t)}$ bijectively matches I' with J'. In particular $i_0 \notin I'$. Choose $j \in J'$ with smallest $d^{(t)}(i_0, j)$. Consider the minimum-weight path in $G^{(t)}$ from i_0 to j which uses only forward and backward edges. The vertex on the path just before j is a bidder $i \notin I'$. Clearly, $y \ge d^{(t)}(i_0, j) > d^{(t)}(i_0, i)$ and hence $u_i(y) < u'_i$. There is a forward edge from i to j, thus $p_j^{(t)} \in [r_{i,j}, m_{i,j})$ and also $u_i(y) + p_j(y) = v_{i,j}$, and hence (*) $u'_i + p'_j < v_{i,j}$. Since in $G^{(t)}$ there is a maximum-price edge from i to j, $p'_j = p_j(y) < m_{i,j}$, which together with (*) contradicts stability of (u', p', μ'). This proves Proposition 15.

To prove Lemma 14 it remains to show that $p^{(t+1)} \le p'$. This amounts to show that if $(u^{(t+1)}, p^{(t+1)}, \mu^{(t+1)})$ was obtained from $(u^{(t)}, p^{(t)}, \mu^{(t)})$ by updating along an alternating path P of which the last edge, $(i, j) = (i_\ell, j_{l+1})$, was a reserve-price edge and $p_j^{(t+)} < r_{i,j}$, then

$$r_{i,j} \le p'_j \ . \qquad (23)$$

Since (u', p', μ') is stable, either $u'_i + p'_j \ge v_{i,j}$ or $p'_j \ge m_{i,j}$. In former case, (23) follows from that $u_i^{(t+1)} = v_{i,j} - r_{i,j}$, Proposition 15 and that (u', p', μ') is stable. In latter case, (23) follows since the presence of the reserve-price edge from i to j guarantees that $m_{i,j} > r_{i,j}$.

The discussion thus far completes the proof of Theorem 4.

7. INCENTIVE COMPATIBILITY

In this section we will prove Theorem 5. A mechanism based on computing men-optimal stable matching has been shown to be truth-revealing in several contexts. For the basic stable matching problem without payments, a concise proof can be found in [17]. For the case of continuous utilities, a proof was given in [8]. Our proof for the max-value model mimics the overall structure of its predecessors. First, we show that there is no feasible matching in which every single bidder would be better off than in the bidder-optimal matching. (Note that if an agent or set of agents were to successfully lie about their preferences, the mechanism would still output a matching that is feasible with respect to the true preferences.) This property is known as weak Pareto optimality of the bidder-optimal matching.

LEMMA 18 (PARETO OPTIMALITY). *Let (v, m, r) be an auction in general position and let (u^*, p^*, μ^*) be the bidder-optimal matching. Then for any matching (u, p, μ) that is feasible for (v, m, r), there is at least one bidder $i \in I$ such that $u_i \le u_i^*$.*

Second, we show that every feasible matching is either stable, or has a blocking bidder-slot pair that involves a bidder who is not better off in this matching than in the bidder-optimal matching. Versions of the following lemma appear in [14, 10, 19]. The original statement in a model without money is attributed to J. S. Hwang.

LEMMA 19 (HWANG'S LEMMA). *Let (u, p, μ) be a matching that is feasible for an auction (v, m, r) in general position and let (u^*, p^*, μ^*) be the bidder-optimal matching for that auction. Let*

$$I^+ = \{i \in I \mid u_i > u_i^*\} \ .$$

If I^+ is non-empty, then there exists a blocking pair $(i, j) \in (I - I^+) \times J$.

Theorem 5 directly follows from Lemma 19. In fact, the lemma implies the following stronger statement.

THEOREM 20. *There is no way for a bidder or a coalition of bidders to manipulate their bids in a way such that every bidder in the coalition would strictly benefit from the manipulation.*

PROOF. Suppose there is a coalition I^+ of bidders that can benefit from submitting false bids. Let (v, m, r) be an auction that reflects the true preferences of all bidders, and let (v', m', r) be an auction that reflects the falsified bids. Note that $v'_i = v_i$ and $m'_i = m_i$ except for bidders $i \in I^+$.

Let (u, p, μ) be the bidder-optimal stable matching for the auction (v', m', r). First observe that the matching (u, p, μ) must be feasible for the true auction (v, m, r). This is because for each bidder $i \in I - I^+$, the feasibility constraints are the same in both auctions. For bidders $i \in I^+$, we need to verify that $p_j \le m_{i,j}$ whenever $(i, j) \in \mu$. This follows because the true bidder-optimal matching (u^*, p^*, μ^*) respects maximum prices, and any outcome that respects maximum prices is preferred over an outcome that doesn't.

Since (u, p, μ) is feasible, we can apply Lemma 19 and conclude that there is a pair (i, j) with $i \in I - I^+$ that is blocking for the auction (v, m, r). □

The rest of this section is devoted to the proofs of Lemmas 18 and 19.

PROOF OF LEMMA 18. For the sake of contradiction, suppose that there is a feasible matching (u, p, μ) such that $u_i > u_i^*$ for all $i \in I$. Note that every bidder must be matched in μ, since $u_i > u_i^* \ge 0$.

For each bidder $i \in I$, consider the slot $j = \mu(i)$ matched to bidder i in the matching μ. Since the pair (i, j) is not blocking for the bidder-optimal matching (u^*, p^*, μ^*), it must be that $p_j^* > p_j$. In particular, the existence of μ implies that there must be n slots with positive prices in the bidder-optimal matching μ^*, and that these slots are matched in μ as well.

In the matching algorithm of Section 5, if a slot ever becomes matched to a bidder, it stays matched to some bidder throughout the algorithm. Thus before the last iteration, at most $n - 1$ slots have positive prices. Suppose the last iteration, iteration $T - 1$, increases the size of the matching to n, and let j be the last slot to be matched. Let $i' = \mu(j)$ be the bidder matched to j in the hypothetical matching μ.

Let P be the shortest alternating path found in Step 1 of the last iteration of the matching algorithm. Recall that the first vertex of the path is denoted by i_0 and $w^{(T-1)}(P)$ denotes its length. If P ends with the reserve-price edge (i, j), it must be that i and j are matched in both μ and μ^* at the same reserve price, contradicting our assumption that $u_i > u_i^*$.

On the other hand, if P does not end with the reserve-price edge (i, j), we show that there is a shorter alternating path P' that does include this edge, which again leads to a contradiction. From Step 3 of the last iteration we have $u_i^{(T-1)} - u_i^* = w^{(T-1)}(P) - d^{(T-1)}(i_0, i)$. Let s be the length of the reserve price edge (i, j);

recall from Definition 8 that $s = u_i^{(T-1)} + r_{i,j} - v_{i,j}$. Now consider the alternating path P' that consists of the shortest path from i_0 to i followed by the reserve price (i, j) edge. We have

$$w^{(T-1)}(P) - w^{(T-1)}(P') = u_i^{(T-1)} - u_i^* - s = v_{i,j} - r_{i,j} - u_i^* .$$

Since $u_i^* < u_i \le v_{i,j} - r_{i,j}$, this difference is positive and hence P' must be a shorter alternating path than P. □

PROOF OF LEMMA 19. Without loss of generality assume that (u, p, μ) is such that there does *not* exist a pair $(i, j) \notin \mu$ such that $u_i + r_{i,j} = v_{i,j}$. If there was such a pair, then we can decrease prices of some of the items and increase utilities of some of the bidders such that $u_i + r_{i,j} > v_{i,j}$. (This is possible because of the general position assumption. See full version of the paper.) The set I^+ would only grow by such operation.

Let us denote by $\mu(I^+)$, $\mu^*(I^+)$ the set of slots matched to bidders in I^+ in matching respectively μ, μ^*. We consider two cases:

Case 1: $\mu(I^+) \ne \mu^*(I^+)$. For any $i \in I^+$ we have $u_i > u_i^* \ge 0$ and hence each bidder in I^+ is matched in μ to some slot. There exists a slot $j \in \mu(I^+)$, $j \notin \mu^*(I^+)$. Let $i = \mu(j)$. Since $i \in I^+$, $u_i > u_i^*$.

We argue that $p_j < p_j^*$: By the general position assumption $p_j^* \ne m_{i,j}$, and hence by feasibility of (u, p, μ), $p_j \in [r_{i,j}, m_{i,j})$ and $u_i + p_j = v_{i,j}$. Hence $u_i^* + p_j^* \ge v_{i,j}$. Therefore $p_j^* \ge v_{i,j} - u_i^* > v_{i,j} - u_i = p_j$.

In particular, j is matched in μ^* to some i', and by the choice of j, $i' \notin I^+$. Thus $u_{i'} \le u_{i'}^*$. By feasibility of (u^*, p^*, μ^*), $p_j^* \in [r_{i',j}, m_{i',j}]$ and $u_{i'}^* + p_j^* = v_{i',j}$. By the assumption on (u, p, μ) that we made at the beginning of the proof, $u_{i'} \ne v_{i',j} - r_{i',j}$.

Now, it is not hard to see that (i', j) is blocking pair for μ. This is because

$$\begin{aligned} p_j &< p_j^* \le m_{i,j} \ , \\ u_{i'} &\le u_{i'}^* = v_{i',j} - p_j^* \le v_{i',j} - r_{i,j} \ \text{ and} \\ u_{i'} &\ne v_{i',j} - r_{i',j} \ , \\ u_{i'} + p_j &< u_{i'}^* + p_j^* = v_{i',j} \ . \end{aligned}$$

Case 2: $\mu(I^+) = \mu^*(I^+) = J^+$. Since $u_i > u_i^*$ for $i \in I^+$, by stability of (u^*, p^*, μ^*) it follows that $p_j < p_j^*$ for $j \in J^+$.

Consider a reduced auction (v', m', r') on the set of bidders I^+ and set of slots J^+. We set the reserve prices to reflect the influence of bidders in $I \setminus I^+$. More specifically, let $I' = \{i \in I \setminus I^+ \mid u_{i'}^* \ge v_{i',j} - r_{i',j}\}$. For every $i \in I^+$ and $j \in J^+$, we set

$$r'_{i,j} = \max\left(r_{i,j}, \max_{i' \in I'} \min(m_{i',j}, v_{i',j} - u_i^*)\right) .$$

We also set $v'_{i,j} = v_{i,j}$ and $m'_{i,j} = m_{i,j}$ except that if $m_{i,j} \le r'_{i,j}$ we set $m'_{i,j} = -1$. It is not hard to show that if v, m, r is in general position, then so is (v', m', r'), using the fact that each utility u_i^* was at some point set to be equal to the length of some alternating walk in the auction graph.

Now consider the matchings μ and μ^* restricted to the sets I^+, J^+. If the restricted μ is not feasible for (v', m', r'), it must be because $p_j < r_{i,j}$ for some position $j = \mu(i)$. This can only happen if $r'_{i,j} > r_{i,j}$ and hence $r'_{i,j} = \max(m_{i',j}, v_{i',j} - u_{i'}^*)$ for some bidder $i' \in I \setminus I^+$.

On the other hand, it is easy to check that the restricted matching μ^* is feasible, stable and bidder-optimal for the auction (v', m', r'). If the restricted μ is feasible for this auction, by Lemma 18, there is a bidder $i \in I^*$ such that $u_i \le u_i^*$. This however contradicts the definition of the set I^+. □

8. REFERENCES

[1] Gagan Aggarwal. *Privacy Protection and Advertising in a Networked World*. PhD thesis, Stanford University, 2005.

[2] Gagan Aggarwal, Ashish Goel, and Rajeev Motwani. Truthful auctions for pricing search keywords. In *ACM Conf on Electronic commerce*, pages 1–7, 2006.

[3] Gagan Aggarwal, S. Muthukrishnan, and Jon Feldman. Bidding to the top: VCG and equilibria of position-based auctions. In *WAOA*, 2006.

[4] Ahmet Alkan. Existence and computation of matching equilibria. *European Journal of Political Economy*, 5(2-3):285–296, 1989.

[5] Ahmet Alkan and David Gale. The core of the matching game. *Games and Economic Behavior*, 2(3):203–212, 1990.

[6] Sushil Bikhchandani and Joseph M. Ostroy. From the assignment model to combinatorial auctions. In *Combinatorial Auctions*. MIT Press, 2006.

[7] Edward H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.

[8] G. Demange and D. Gale. The strategy structure of two-sided matching markets. *Econometrica*, 53(4):873–888, 1985.

[9] Gabrielle Demange, David Gale, and Marilda Sotomayor. Multi-item auctions. *Journal of Political Economy*, 94(4):863–872, 1986.

[10] Gabrielle Demange, David Gale, and Marilda Sotomayor. A further note on the stable matching problem. *Discrete Applied Mathematics*, 16:217–222, 1987.

[11] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, March 2007.

[12] Eyal Even-Dar, Jon Feldman, Yishay Mansour, and S. Muthukrishnan. On the effect of minimum prices on position auctions. Unpublished manuscript, 2008.

[13] David Gale and Lloyd S. Shapley. College admissions and the stability of marriage. *Am Math Monthly*, 69(1):9–15, 1962.

[14] David Gale and Marilda Sotomayor. Some remarks on the stable marriage problem. *Discrete Applied Mathematics*, 11:223–232, 1985.

[15] Theodore Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.

[16] Herman B. Leonard. Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy*, 91:461–479, 1983.

[17] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, Septermber 2007.

[18] M. Quinzii. Core and competitive equilibria with indivisibilities. *International Journal of Game Theory*, 13(1):41–60, 1984.

[19] Alvin E. Roth and Marilda A. Oliveira Sotomayor. *Two-sided matching: A study in game-theoretic modeling and analysis*. Cambridge University Press, 1990.

[20] Lloyd S. Shapley and Martin Shubik. The assignment game i: The core. *Intl J. of Game Theory*, 1(1):111–130, 1971.

[21] Hal R. Varian. Position auctions. *International Journal of Industrial Organization*, 2006.

[22] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *J.of Finance*, 16(1):8–37, 1961.