On Multiple Keyword Sponsored Search Auctions with Budgets*

Riccardo Colini-Baldeschi[1,**], Monika Henzinger[2,***],
Stefano Leonardi[1,**], and Martin Starnberger[2,***]

[1] Sapienza University of Rome, Italy
[2] University of Vienna, Austria

Abstract. We study *multiple keyword* sponsored search auctions with budgets. Each keyword has *multiple ad slots* with a click-through rate. The bidders have additive valuations, which are linear in the click-through rates, and budgets, which are restricting their overall payments. Additionally, the number of slots per keyword assigned to a bidder is bounded. We show the following results: (1) We give the first mechanism for multiple keywords, where click-through rates differ among slots. Our mechanism is incentive compatible in expectation, individually rational in expectation, and Pareto optimal. (2) We study the combinatorial setting, where each bidder is only interested in a subset of the keywords. We give an incentive compatible, individually rational, Pareto optimal, and deterministic mechanism for identical click-through rates. (3) We give an impossibility result for incentive compatible, individually rational, Pareto optimal, and deterministic mechanisms for bidders with diminishing marginal valuations.

1 Introduction

In *sponsored search* (or *adwords*) auctions advertisers bid on *keywords*. Such auctions are used by firms such as Google, Yahoo, and Microsoft [11]. The search result page for each keyword contains multiple slots for ads and each bidder is assigned to a limited number of slots on a search result page. The slots have a click-through rate (CTR), which is usually decreasing by the position of the slot on the search result page. The true valuation of the bidders is private knowledge and is assumed to depend linearly on the CTR. Moreover, valuations are assumed

* The original source of publication is "Automata, Languages, and Programming, Czumaj, A., Mehlhorn, K., Pitts, A., Wattenhofer, R. (eds.), Lecture Notes in Computer Science, vol. 7392. Springer Berlin / Heidelberg (2012), `http://dx.doi.org/10.1007/978-3-642-31585-5_1`". The original publication is available at `www.springerlink.com`.

** This work was partially supported from ERC Starting Grant 259515: Practical Approximation Algorithms.

*** This work has been funded by the Vienna Science and Technology Fund (WWTF) through project ICT10-002 and by the University of Vienna through IK I049-N.

to be additive, i.e., the total valuation of a bidder is equal to the sum of his valuations for all the slots that are assigned to him.

A further key ingredient of an adwords auction is that bidders specify a budget on the payment charged for the ads, effectively linking the different keywords. The deterministic Vickrey auction [18] was designed to maximize social welfare in this and more general settings without budget restrictions. However, the introduction of budgets dramatically changes the nature of the problem. The Vickrey auction may charge more than the budget and is no longer feasible. Moreover, bidders might get assigned to slots even though their budget is arbitrary small and other bidders are interested in those slots. Thus, as was observed before [7,8], maximizing social welfare is not the right optimality criterion to use. In a seminal paper by Dobzinski et al. [7,8], they considered the multi-unit case with additive valuations, which in the sponsored search setting corresponds to each keyword having only one slot and all slots having identical CTR. They gave an incentive compatible (IC) auction based on Ausubel's ascending clinching auction [3] that produces a Pareto optimal (PO) and individually rational (IR) allocation if budgets are *public*. They also showed that this assumption is strictly needed, i.e., that no deterministic mechanism for *private* budgets exists if we insist on incentive compatibility, individual rationality, and on obtaining an allocation that is Pareto optimal. This impossibility result for *deterministic* mechanisms was strengthened for our setting to *public* budgets in Dütting et al. [9]. The question was open what optimality result can be achieved for *randomized* mechanisms. Due to the impossibility results for deterministic mechanisms it is unlikely that "strong" optimality criteria, such as bidder optimality, are achievable. Thus, the first question to study is whether Pareto optimality, which is a basic notion of optimality, can be achieved with *randomized* mechanisms. Note that if an allocation is Pareto optimal then it is impossible to make a bidder better off without making another bidder or the auctioneer worse off, and is therefore the least one should aim for.

Our Results. We give a positive answer to the above question and also present two further related results. Specifically, the paper contains the following three results: (1) *Multiple keywords with multiple slots:* We show that the multi-unit auction of Dobzinski et al. [7,8] can be adapted to study adwords auctions with multiple keywords having multiple slots, and budget limits for each bidder. We specifically model the case of several slots with different CTR, available for each keyword, and a bound on the number of slots per keyword (usually one) that can be allocated to a bidder. We first provide an IC, IR, and PO deterministic auction that provides a fractional allocation for the case of one keyword with divisible slots. Note that the impossibility result in [9] does not hold for divisible slots. In contrast, the impossibility result in [7,8] for multi-unit auctions applies also to this setting, and achieving IC, IR, and PO deterministic auctions is only possible if budgets are *public*. Thus, we restrict ourselves to the public budget case. Our auction is a one-shot auction, i.e., each bidder interacts only once with the auction. We then show how to probabilistically round this fractional allocation for the divisible case to an integer allocation for the indivisible case with multiple keywords (i.e., the adwords setting) and get an auction that is IC in expectation, IR in expectation, and PO.

(2) *Multiple keywords with combinatorial constraints and multiple slots:* So far we assumed that every bidder is interested in every keyword. In the second part of the paper we study the case that bidders are interested in only a subset of the keywords, i.e., bidders have a non-zero identical valuation only on a *subset* of the keywords. The valuations are additive and each bidder is assigned at most one slot for a given keyword. We restrict the model by allowing only identical slots for each keyword, i.e., we require that all slots have the same CTR. This setting extends the combinatorial one-slot per keyword model considered by Fiat et al. [12] to multiple slots. We present a variation of the clinching auction that is deterministic, IC, IR, and PO.

(3) Finally, we also study non-additive valuations, namely valuations with diminishing marginal valuations. Diminishing marginal valuation (also called submodular) functions are widely used to model auction settings with marginal utilities being positive functions that are non-increasing in the number of items already allocated to the bidders. We show that even in the multi-unit (one slot per keyword) case there is no deterministic, IC, IR, and PO auction for private diminishing marginal valuations and public budgets. This shows how budgets complicate mechanism design: For the non-budgeted version of this setting Ausubel [3] gave his deterministic mechanism.

Related Work. Ascending clinching auctions are used in the FCC spectrum auctions, see [16,4,3]. For a motivation of adwords auctions see [17] on Google's auction for TV ads.

We first compare our results with those of a recent, unpublished work by Goel et al. [14] that was developed independently at the same time. They studied IC auctions with feasible allocations that must obey public polymatroid constraints and agents with identical or separable valuations (see their Lemma 3.10) and public budgets. The problem of auctions with polymatroid constraints was first studied by Bikhchandani et al. [6] for unbudgeted bidders and concave utilities. The auction in [14] is an adaption of the ascending auction in [6] to the case of budgeted bidders. The polymatroid constraints generalize on one hand the the multi-unit case in [7,8] and the multiple slots with different CTRs model presented in this paper. On the other hand, the PO ascending auction in [14] only returns allocations for divisible items whereas in Sect. 4 of this paper we demonstrate that these allocations can be rounded to *allocations for indivisible items* if we allow the auction to yield incentive compatibility in expectation. In Sect. 5, we present an IC, IR, and PO *deterministic* auction with feasible *allocations of indivisible slots* that obey matching constraints for the case of multiple identical slots.

There are three extensions of Dobzinski et al. [7,8]: (1) Fiat et al. [12] studied an extension to a combinatorial setting, where items are distinct and different bidders may be interested in different items. The auction presented in [12] is IC, IR, and PO for additive valuations and single-valued bidders (i.e., every bidder does not distinguish between the keywords in his public interest set). This is a special case of our combinatorial setting in Sect. 5 with multiple keywords but only one slot per keyword. (2) Bhattacharya et al. [5] dealt with private budgets, and gave an auction for one infinitely divisible item, where bidders cannot improve their utility by underreporting their budget. This leads to a randomized IC in expectation auction for one infinitely divisible item with both private valuations and budgets. (3) Several papers [1,2,10,13] studied *envy-free* outcomes that are bidder optimal, respectively PO, in an one-keyword adwords auction. In this setting they give (under certain conditions on the input) an IC auction with both private valuations and budgets.

Our impossibility result in Sect. 6 is related to two impossibility results: Lavi and May [15] show that there is no IC, IR, and PO deterministic mechanism for indivisible items and bidders with *monotone* valuations. Our result for indivisible items is stronger as it applies to bidders with non-negative and diminishing marginal valuations. In [14] the same impossibility result for *divisible* items and bidders with monotone and concave utility functions was given. Note that neither their result nor ours implies the other.

2 Problem Statement and Definitions

We have n bidders and m slots. We call the set of bidders $I := \{1, \dots, n\}$ and the set of slots $J := \{1, \dots, m\}$. Each bidder $i \in I$ has a private *valuation* v_i, a public *budget* b_i, and a public *slot constraint* κ_i, which is a positive integer. Each slot $j \in J$ has a public *quality* $\alpha_j \in \mathbb{Q}_{\geq 0}$. The slots are ordered such that $\alpha_j \geq \alpha_{j'}$ if $j > j'$, where ties are broken in some arbitrary but fixed order. We assume in Sect. 3 and 4 that the number of slots m fulfills $m = \sum_{i \in I} \kappa_i$ as the general case can be easily reduced to this setting.

Divisible case: In the divisible case we assume that there is only one keyword with infinitely divisible slots. Thus the goal is to assign each bidder i a fraction $x_{i,j} \geq 0$ of each slot j and charge him a payment p_i. A matrix $X = (x_{i,j})_{(i,j) \in I \times J}$ and a payment vector p are called an *allocation* (X, p). We call $c_i = \sum_{j \in J} \alpha_j x_{i,j}$ the *weighted capacity* allocated to bidder i. An allocation is *feasible* if it fulfills the following conditions: (1) the sum of the fractions assigned to a bidder does not exceed his *slot constraint* ($\sum_{j \in J} x_{i,j} \leq \kappa_i \ \forall i \in I$); (2) each of the slots is fully assigned to the bidders ($\sum_{i \in I} x_{i,j} = 1 \ \forall j \in J$); and (3) the payment of a bidder does not exceed his budget limit ($b_i \geq p_i \ \forall i \in I$).

Indivisible case: We additionally have a set R of *keywords*, where $|R|$ is public. The goal is to assign each slot $j \in J$ of keyword $r \in R$ to one bidder $i \in I$ while obeying various constraints. An assignment $X = (x_{i,j,r})_{(i,j,r) \in I \times J \times R}$ where $x_{i,j,r} = 1$ if slot j is assigned to bidder i in keyword r, and $x_{i,j,r} = 0$ otherwise, and a payment vector p form an *allocation* (X, p). We call $c_i = \sum_{j \in J} \frac{\alpha_j}{|R|} (\sum_{r \in R} x_{i,j,r})$ the *weighted capacity* allocated to bidder i. An allocation is *feasible* if it fulfills the following conditions: (1) the number of slots of a keyword that are assigned to a bidder does not exceed his *slot constraint* ($\sum_{j \in J} x_{i,j,r} \leq \kappa_i \ \forall i \in I, \forall r \in R$); (2) each slot is assigned to exactly one bidder ($\sum_{i \in I} x_{i,j,r} =$

$1\ \forall j \in J, \forall r \in R$); and (3) the payment of a bidder does not exceed his budget limit ($b_i \geq p_i\ \forall i \in I$).

Combinatorial indivisible case: In the combinatorial case not all keywords are identical. Every bidder $i \in I$ has a publicly known set of interest $S_i \subseteq R$, and valuation v_i for all keywords in S_i and a valuation of zero for all other keywords. We model this case by imposing $x_{i,j,r} = 0\ \forall r \notin S_i$.

Note that in all cases the budgets are bounds on *total* payments across keywords and *not* bounds on prices of individual keywords.

Properties of the auctions: The *utility* u_i of bidder i for a feasible allocation (X, p) is $c_i v_i - p_i$, the *utility* of the auctioneer (or mechanism) is $\sum_{i \in I} p_i$. We study auctions that select feasible allocations obeying the following conditions: (*Bidder rationality*) $u_i \geq 0$ for all bidders $i \in I$, (*Auctioneer rationality*) the utility of the auctioneer fulfills $\sum_{i \in I} p_i \geq 0$, and (*No-positive-transfer*) $p_i \geq 0$ for all bidders $i \in I$. An auction that on all inputs outputs an allocation that is both bidder rational and auctioneer rational is called *individually rational (IR)*. A feasible allocation (X, p) is *Pareto optimal* (PO) if there is no other feasible allocation (X', p') such that (1) the utility of no bidder in (X, p) is less than his utility in (X', p'), (2) the utility of the auctioneer in (X, p) is no less than his utility in (X', p'), and (3) at least one bidder or the auctioneer is better off in (X', p') compared with (X, p). An auction is *incentive compatible* (IC) if it is a dominant strategy for all bidders to reveal their true valuation. An auction is said to be *Pareto optimal* (PO) if the allocation it produces is PO. A randomized auction is IC in expectation, IR in expectation, respectively PO in expectation if the above conditions hold in expectation. We show that our randomized mechanism for indivisible slots is PO in expectation and that each realized allocation is PO.

3 Deterministic Clinching Auction for the Divisible Case

3.1 Characterization of Pareto Optimality

In this section we present a novel characterization of PO allocations that allows to address the divisible case of multiple slots with different CTRs. Given a feasible allocation (X, p), a *swap* between two bidders i and i' is a fractional exchange of slots, i.e., if there are slots j and j' and a constant $\tau > 0$ with $x_{i,j} \geq \tau$ and $x_{i',j'} \geq \tau$ then a swap between i and i' gives a new feasible (X', p) with $x'_{i,j} = x_{i,j} - \tau$, $x'_{i',j'} = x_{i',j'} - \tau$, $x'_{i,j'} = x_{i,j'} + \tau$, and $x'_{i',j} = x_{i',j} + \tau$. If $\alpha_j < \alpha_{j'}$ then the swap increases i's weighted capacity. We assume throughout this section that $\alpha_j \neq \alpha_{j'}$ for $j \neq j'$, the general case requires a small modification presented in the full version of our paper. To characterize PO allocations we first define for each bidder i the set N_i of bidders such that for every bidder a in N_i there exists a swap between i and a that increases i's weighted capacity. Given a feasible allocation (X, p) we use $h(i) := \max\{j \in J | x_{i,j} > 0\}$ for the slot with the highest quality that is assigned to bidder i and $l(i) := \min\{j \in J | x_{i,j} > 0\}$ for the slot with the lowest quality that is assigned to bidder i. Now, $N_i = \{a \in I | h(a) > l(i)\}$ is the set of all the bidders a such that i could increase his weighted capacity (and a could decrease his weighted capacity) if i traded with a, for example, if i received part of a's share of slot $h(a)$. To model sequences of swaps we define furthermore $N_i^k = N_i$ for $k = 1$ and $N_i^k = \bigcup_{a \in N_i^{k-1}} N_a$ for $k > 1$. Since we have only n bidders, $\bigcup_{k=1}^{n} N_i^k = \bigcup_{k=1}^{n'} N_i^k$ for all $n' \geq n$. We define $\tilde{N}_i := \bigcup_{k=1}^{n} N_i^k \setminus \{i\}$ as the set of *desired (recursive) trading partners* of i. See Fig. 3.1 for an example with four bidders. The bidders a in $\tilde{N}_i$ are all the bidders such that through a sequence of trades that "starts" with i and "ends" with a, bidder i could increase his weighted capacity, bidder a could decrease his weighted capacity, and the capacity of the remaining bidders involved in the swap would been unchanged. Now let $\tilde{v}_i = \min_{a \in \tilde{N}_i}(v_a)$ if $\tilde{N}_i \neq \emptyset$ and $\tilde{v}_i = \infty$ else.

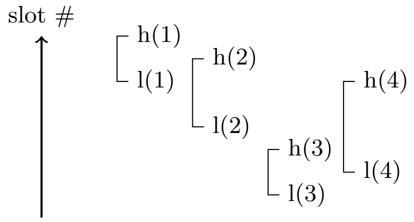

$N_1 = \{1,2\}$ $\quad N_2 = \{1,2,4\}$
$N_3 = \{1,2,3,4\}$ $\quad N_4 = \{1,2,3,4\}$

$N_1^2 = \{1,2,4\}$ $\quad N_2^2 = \{1,2,3,4\}$

$N_1^3 = \{1,2,3,4\}$

$\tilde{N}_1 = \{2,3,4\}$ $\quad \tilde{N}_2 = \{1,3,4\}$
$\tilde{N}_3 = \{1,2,4\}$ $\quad \tilde{N}_4 = \{1,2,3\}$

Fig. 1. Example of desired trading partners

Given a feasible allocation (X,p) we use $B := \{i \in I | b_i > p_i\}$ to denote the set of bidders who have a positive remaining budget. As we show below if for a given assignment we know $\tilde{v}_i$ for every bidder $i \in B$ then we can immediately decide whether the assignment is PO or not.

We say that a feasible allocation (X,p) contains a *trading swap sequence* (for short *trading swap*) if there exists a feasible allocation (X',p') and two bidders $u, w \in I$ such that

1. bidder w is a desired trading partner of u, i.e., $w \in \tilde{N}_u$,
2. for all $i \in I \setminus \{u,w\}$ it holds that the weighted capacity of i and the payment of i are unchanged by the swap, i.e., $\sum_{j \in J} \alpha_j x_{i,j} = \sum_{j \in J} \alpha_j x'_{i,j}$ and $p_i = p'_i$,
3. the weighted capacity of u increases by $\delta > 0$ and the weighted capacity of w decreases by δ, i.e., $\delta := \sum_{j \in J} \alpha_j (x'_{u,j} - x_{u,j}) = \sum_{j \in J} \alpha_j (x_{w,j} - x'_{w,j}) > 0$,
4. $v_u > v_w$, u pays after the swap exactly that amount more that w's weighted valuation decreases (i.e., $p'_u - p_u = v_w \delta$), and w pays exactly that amount less (i.e., $p_w - p'_w = v_w \delta$), and
5. u has a high enough budget to pay what is required by (X',p'), i.e., $b_u \geq p'_u$.

We say that the allocation (X',p') *results from the trading swap*. The existence of a trading swap is related to the $\tilde{v}_i$ of each bidder i with remaining budget.

Theorem 1. *A feasible allocation (X,p) contains no trading swaps if and only if $\tilde{v}_i \geq v_i$ for each bidder $i \in B$.*

The following theorem shows that the absence of trading swaps characterizes Pareto optimality. We will use exactly this fact to prove that the mechanism of the next section outputs a PO allocation.

Theorem 2. *A feasible allocation* (X, p) *is Pareto optimal if and only if it contains no trading swaps.*

Hence, the feasible allocation (X, p) is PO if and only if $\tilde{v}_i \geq v_i \ \forall i \in B$. This novel characterization of Pareto optimality is interesting, as the payment does not affect the values $\tilde{v}_i$, the payment only influences which bidders belong to B.

3.2 Multiple Keyword Auction for the Divisible Case

We describe next our deterministic clinching auction for divisible slots and show that it is IC, IR, and PO. The auction repeatedly increases a price "per capacity" and gives different weights to different slots depending on their CTRs. To perform the check whether all remaining unsold weighted capacity can still be sold we solve suitable linear programs. We will show that if the allocation of the auction did not fulfill the characterization of Pareto optimality given in Sect. 3.1, i.e., if it contained a trading swap, then one of the linear programs solved by the auction would not have computed an optimal solution. Since this is not possible, it will follow that the allocation is PO. A formal description of the auction is given in the procedures Auction and Sell. The input values of Auction are the bids, budget limits, and slot constraints that the bidders communicate to the auctioneer on the beginning of the auction, and information about the qualities of the slots. The auction is a so called "one-shot auction", the bidders are asked once for the valuations at the beginning of the auction and then they *cannot* input *any* further data.

The demand of the bidders for weighted capacity is computed by the mechanism based on their remaining budget and the current price. We assume throughout this section that $v_i \in \mathbb{N}_+$ and $b_i \in \mathbb{Q}_+$ for all $i \in I$.[3] The state of the auction is defined by the current price π, the next price π^+, the weighted capacity c_i that bidder $i \in I$ has clinched so far, and the payment p_i that has been charged so far to bidder i. We define the set of *active* bidders $A \subseteq I$ which are all those $i \in I$ with $\pi \leq v_i$, and the subset E of A of *exiting* bidders which are all those $i \in A$ with $\pi^+ > v_i$. The auction does not increase the price that a bidder $i \in I$ has to pay from π to π^+ for *all* bidders at the same time. Instead, it calls Sell each time before it increases the price for a single bidder. If the price that bidder $i \in A$ has to pay for weighted capacity is π then his demand is $d_i = \frac{b_i - p_i}{\pi}$. If the price he has to pay was already increased to π^+ then his demand is $d_i = \frac{b_i - p_i}{\pi^+} < \frac{b_i - p_i}{\pi}$. In this case, the demand corresponds to d_i^+, that is always equal to $\frac{b_i - p_i}{\pi^+}$. Different from the auction in [7,8,5] a bidder with $d_i = d_i^+$ is also charged the

[3] All the arguments go through if we simply assume that $v_i \in \mathbb{Q}_+$ for all $i \in I$ and there exists a publicly known value $z \in \mathbb{R}_+$ such that for all bidders i and i' either $v_i = v_{i'}$ or $|v_i - v_{i'}| \geq z$.

Algorithm 1 Clinching auction for divisible slots

procedure AUCTION$(I, J, \alpha, \kappa, v, b)$
  $A \leftarrow I;\ \pi \leftarrow 0;\ \pi^+ \leftarrow 1$
  $c_i \leftarrow 0,\ p_i \leftarrow 0,\ d_i \leftarrow \infty\ \forall i \in I$
  while $\sum_{i \in I} c_i < \sum_{j \in J} \alpha_j$ **do** ▷ unsold weighted capacity exists
    $E \leftarrow \{i \in A | \pi^+ > v_i\}$ ▷ bidders become exiting bidders
    for $i \in E$ **do**
      $(X, s) \leftarrow$ SELL$(I, J, \alpha, \kappa, c, d, i)$ ▷ sell to exiting bidder
      $(c_i, p_i, d_i) \leftarrow (c_i + s, p_i + s\pi, 0)$
    end for
    $A \leftarrow A \setminus E$ ▷ exiting bidders leave auction
    $d_i^+ \leftarrow \frac{b_i - p_i}{\pi^+}\ \forall i \in A$
    while $\exists i \in A$ with $d_i \neq d_i^+$ **do** ▷ bidders with price π exist
      $i' \leftarrow \min(\{i \in A | d_i \neq d_i^+\})$ ▷ select bidder with price π
      $(X, s) \leftarrow$ SELL$(I, J, \alpha, \kappa, c, d, i')$ ▷ sell to bidder
      $(c_{i'}, p_{i'}) \leftarrow (c_{i'} + s, p_{i'} + s\pi)$
      $d_{i'}^+ \leftarrow \frac{b_{i'} - p_{i'}}{\pi^+};\ d_{i'} \leftarrow d_{i'}^+$ ▷ increase bidder's price to π^+
    end while
    $\pi \leftarrow \pi^+;\ \pi^+ \leftarrow \pi^+ + 1$ ▷ increase price
  end while
  return (X, p)
end procedure

Algorithm 2 Determination of the weighted capacity that bidder i' clinches

procedure SELL$(I, J, \alpha, \kappa, c, d, i')$
  compute an optimal solution of the following linear program that is a vertex of the polytope defined by its constraints:
  minimize $\gamma_{i'}$
  s.t.: (a) $\sum_{i \in I} x_{i,j} = 1\quad \forall j \in J$ ▷ assign all slots
  (b) $\sum_{j \in J} x_{i,j} = \kappa_i\quad \forall i \in I$ ▷ slot constraint
  (c) $\sum_{j \in J} x_{i,j}\ \alpha_j - \gamma_i = c_i\quad \forall i \in I$ ▷ assign value to γ_i
  (d) $\gamma_i \leq d_i\quad \forall i \in I$ ▷ demand constraint
  (e) $x_{i,j} \geq 0\quad \forall i \in I, \forall j \in J$
  (f) $\gamma_i \geq 0\quad \forall i \in I$
  return $(X, \gamma_{i'})$
end procedure

increased price π^+ if he receives additional weighted capacity. Since our price is incremented by one in each round and is not continuously increasing as in prior work, this is necessary for proving the Pareto optimality of the allocation.

The crucial point of the auction is that it sells only weighted capacity s to bidder i at a certain price π or π^+ if it cannot sell s to the other bidders. It computes s by solving a linear program in Sell. We use a linear program as there are two types of constraints to consider: The slot constraint in line (b) of the LP, which constraints "unweighted" capacity, and the demand constraint in line (d) of the LP, which is implied by the budget limit and constraints weighted capacity. In the homogeneous item setting in [7,8,5] there are no slot constraints and the demand constraints are unweighted, i.e., $\alpha_j = 1\ \forall j \in J$. Thus, no linear program is needed to decide what amount to sell to whom.

For each iteration of the outer while-loop the auction first calls Sell for each exiting bidder i and sells him s for price π. This is the last time when he can gain weighted capacity. Afterward, he is no longer an active bidder. Next, it calls Sell for one of the remaining active bidders who has $d_i \neq d_i^+$. It sells him the respective s and increases his price to π^+. It continues the previous step until the price of each active bidder is increased to π^+. Then it sets π to π^+ and π^+ to $\pi^+ + 1$.

It is crucial for the progress and the correctness of the mechanism that there is a feasible solution for the linear program in Sell every time that Sell is called. This is proved in the full version. It follows that the final assignment X is a feasible solution of the linear program in Sell. Thus it fulfills conditions (1) and (2) for a feasible allocation. Condition (3) is also fulfilled as by the definition of the demand of a bidder, the auction guarantees that $b_i \geq p_i$ for all i. Thus, the allocation (X, p) computed by the auction is a feasible allocation. As no bidder is assigned weighted capacity if his price is above his valuation and the mechanism never pays the bidders, the auction is IR. As it is an increasing price auction, it is also IC.

Proposition 1. *The auction is individually rational and incentive compatible and the allocation (X, p) it outputs has only rational entries.*

We show finally that the allocation (X, p) our auction computes does not contain any trading swap, and thus, by Theorem 2 it is PO. The proof shows that every trading swap in (X, p) would lead to a superior solution to one of the linear programs solved by the mechanism. Since the mechanism found an optimal solution this leads to a contradiction.

Theorem 3. *The allocation (X, p) returned by our auction is Pareto optimal.*

4 Randomized Clinching Auction for the Indivisible Case

We will now use the allocation computed by the deterministic auction for *divisible* slots to give a randomized auction for multiple keywords with *indivisible* slots that ensures that bidder i receives at most κ_i slots for each keyword. The randomized auction has to assign to every slot $j \in J$ exactly one bidder $i \in I$ for each keyword $r \in R$. We call a distribution over allocations for the indivisible case *Pareto superior* to another such distribution if the expected utility of a bidder or the auctioneer is higher, while all other expected utilities are at least as large. If a distribution has no Pareto superior distribution, we call it *Pareto optimal.* The basic idea is as follows: Given the PO solution for the *divisible* case, we construct a *distribution over allocations* of the *indivisible* case such that the expected utility of every bidder and of the auctioneer is the same as the utility of the bidder and the auctioneer in the *divisible* case. To be precise, we do not explicitly construct this distribution but instead we give an algorithm that can sample from this distribution. The mechanism for the *indivisible* case would, thus, first call the mechanism for the *divisible* case (with the same input) and then convert the resulting allocation (X^d, p^d) into a representation of a PO distribution over allocations for the *indivisible* case. It then samples from this representation to receive the allocation that it outputs. During all these steps the (expected) utility of the bidders and the auctioneer remains unchanged. As the mechanism for the *divisible* case is IR and IC this implies immediately that the mechanism for the *indivisible* case is IR in expectation and IC in expectation. To show that the final allocation is PO in expectation and also PO ex post we use the following lemma.

Lemma 1. *For every probability distribution over feasible allocations in the indivisible case there exists a feasible allocation (X^d, p^d) in the divisible case, where the utility of the bidders and the auctioneer equals their expected utility using this probability distribution.*

Lemma 1 implies that any probability distribution over feasible allocations in the *indivisible* case that is Pareto superior to the distribution generated by our auction would lead to a feasible allocation for the *divisible* case that is Pareto superior to (X^d, p^d). This is not possible as (X^d, p^d) is PO. Additionally, each realized allocation is ex-post Pareto optimal: if in the *indivisible* case there existed a Pareto superior allocation to one of the allocations that gets chosen with a positive probability in our auction, then a Pareto superior allocation would exist in the *divisible* case. By the same argument as above this would lead to a contradiction.

We still need to explain how to use the PO allocation (X^d, p^d) for the *divisible* case to give a probability distribution for the *indivisible* case with expected utility for every bidder equal to the utility in the divisible case and how to sample efficiently from this distribution. Given an input for the indivisible case we use it *as is* as an input for the algorithm for the divisible case, ignoring the number of keywords. Based on the allocation (X^d, p^d) for the divisible problem we construct a matrix M' of size $|J| \times \lambda$, where λ is the least common denominator of all the $x^d_{i,j}$ values and where each column of M' corresponds to a feasible assignment for the indivisible one-keyword case. Note that the same assignment can occur in multiple columns of M'. The matrix M' is our representation of the distribution over allocations in the indivisible case. To sample from the distribution we pick for each $r \in R$ a column uniformly at random from the columns of M'. The r-th choice gives the assignment of bidders to the slots of keyword r. The payments are set equal to p^d. In the full version, we give the construction of M' such that after the above sampling step the expected weighted capacity allocation to bidder $i \in I$ equals $\sum_{j \in J} \alpha_j x^d_{i,j}$, i.e., its weighted capacity in the divisible case. Additionally, all of the slots are fully assigned to the bidders, and hence, the stated properties are fulfilled by the randomized auction.

5 The Combinatorial Case with Multiple Slots

We consider single-valued combinatorial auctions with multiple identical slots in multiple keywords. Every bidder $i \in I$ has valuation v_i on all keywords of his interest set S_i. All other keywords are valued zero. The interest sets S_i and the budgets b_i are public knowledge. We further restrict to the case where at most one slot per keyword is allocated to a single bidder (i.e., $\kappa_i = 1$). We require that at least m bidders are interested in each keyword, where m is the number of slots for a keyword.

In our auction, we extend the B-matching based approach and the concept of trading alternating paths in the bidder/keyword bipartite graph by Fiat et al. [12] for their single-slot per keyword setting to our multi-slot per keyword setting.

We characterize a feasible allocation (H, p) by a tuple $H = (H_1, H_2, \dots, H_n)$, where $H_i \subseteq S_i$ represents the set of keywords that are allocated to bidder i, and by a vector of payments $p = (p_1, p_2, \dots, p_n)$ with $p_i \le b_i$ for all $i \in I$. The utility of bidder i is defined by $u_i := v_i|H_i| - p_i$, and the utility of the auctioneer is $\sum_{i=1}^{n} p_i$. We base the allocation of the items in the clinching auction on B-matchings computed on a bipartite graph G with the union of keywords and bidders $(I \cup R)$ as vertex set and the preferences $\{(i,t) \in I \times R | t \in S_i\}$ as edge set. The vertices have degree constraints, which represent the demand constraints for the bidders and the number of unsold slots for the keywords. The B-matchings are the subgraphs of G, which fulfill the constraints, and have a maximal number of edges. The idea of the auction is to sell slots at the highest possible price such that all slots are sold and there exists no competition between bidders. We define the auction and give the proof of the following theorem in the full version.

Theorem 4. *The allocation (H^*, p^*) produced by the combinatorial clinching auction is incentive compatible, individually rational, and Pareto optimal.*

6 Impossibility for Diminishing Marginal Valuations

We assume in this section that we have multiple homogeneous indivisible items and bidders with private diminishing marginal valuations and public budgets. We show that there is no IC, IR, and PO deterministic mechanism for this case.

Bidder i's marginal valuation for obtaining a further item when k items are already assigned to him is $v_i(k+1)$. His valuation for obtaining $k+1$ items is therefore $\sum_{j=1}^{k+1} v_i(j)$. The marginal valuations have to fulfill $v_i(k) \geq v_i(k+1)$ for $k \geq 1$. The initial clinching auction in [3] was indeed proposed for the case of diminishing marginal valuations but without budget limits.

We use that the case of additive valuations, which was studied by Dobzinski et al. [7,8], is a special case of ours, and that they showed that their auction is the only IC, IR, and PO deterministic auction for that case. We study bidders with diminishing marginal valuations that report additive valuations in order to raise the price paid by the other bidders and consequently decrease their demand. A possible decrease of the price charged to the non-truth telling bidders follows.

Theorem 5. *There is no incentive compatible, individually rational, Pareto optimal, and deterministic mechanism for multiple homogeneous indivisible items and agents with private diminishing marginal valuations and public budget limits.*

References

1. Aggarwal, G., Muthukrishnan, S., Pál, D., Pál, M.: General auction mechanism for search advertising. In: WWW '09: Proceedings of the 18th international conference on World wide web. pp. 241–250. ACM (2009)
2. Ashlagi, I., Braverman, M., Hassidim, A., Lavi, R., Tennenholtz, M.: Position auctions with budgets: Existence and uniqueness. The B.E. Journal of Theoretical Economics 10(1) (2010)
3. Ausubel, L.M.: An efficient ascending-bid auction for multiple objects. American Economic Review 94(5), 1452–1475 (2004)
4. Ausubel, L.M., Milgrom, P.R.: Ascending auctions with package bidding. Frontiers of Theoretical Economics 1(1), 1019–1019 (2002)
5. Bhattacharya, S., Conitzer, V., Munagala, K., Xia, L.: Incentive compatible budget elicitation in multi-unit auctions. CoRR abs/0904.3501 (2009)
6. Bikhchandani, S., de Vries, S., Schummer, J., Vohra, R.V.: Ascending auctions for integral (poly)matroids with concave nondecreasing separable values. In: Teng, S.H. (ed.) SODA. pp. 864–873. SIAM (2008)
7. Dobzinski, S., Lavi, R., Nisan, N.: Multi-unit auctions with budget limits. In: FOCS. pp. 260–269. IEEE Computer Society (2008)
8. Dobzinski, S., Lavi, R., Nisan, N.: Multi-unit auctions with budget limits (2011), `http://ie.technion.ac.il/~ronlavi/papers/budget-constraints.pdf`
9. Dütting, P., Henzinger, M., Starnberger, M.: Auctions with heterogeneous items and budget limits (2012)
10. Dütting, P., Henzinger, M., Weber, I.: An expressive mechanism for auctions on the web. In: Srinivasan, S., Ramamritham, K., Kumar, A., Ravindra, M.P., Bertino, E., Kumar, R. (eds.) WWW. pp. 127–136. ACM (2011)
11. 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), 242–259 (2005)
12. Fiat, A., Leonardi, S., Saia, J., Sankowski, P.: Single valued combinatorial auctions with budgets. In: Shoham, Y., Chen, Y., Roughgarden, T. (eds.) ACM Conference on Electronic Commerce. pp. 223–232. ACM (2011)
13. Fujishige, S., Tamura, A.: A two-sided discrete-concave market with possibly bounded side payments: An approach by discrete convex analysis. Mathematics of Operations Research 32(1), 136–155 (2007)

14. Goel, G., Mirrokni, V.S., Leme, R.P.: Polyhedral clinching auctions and the adwords polytope. CoRR abs/1201.0404 (2012), to appear in 44th ACM Symposium on Theory of Computing (STOC 2012), New York, May 2012.
15. Lavi, R., May, M.: A note on the incompatibility of strategy-proofness and pareto-optimality in quasi-linear settings with public budgets - working paper. In: Chen, N., Elkind, E., Koutsoupias, E. (eds.) WINE. Lecture Notes in Computer Science, vol. 7090, p. 417. Springer (2011)
16. Milgrom, P.: Putting auction theory to work: The simulteneous ascending auction. Journal of Political Economy 108(2), pp. 245–272 (2000)
17. Nisan, N., Bayer, J., Chandra, D., Franji, T., Gardner, R., Matias, Y., Rhodes, N., Seltzer, M., Tom, D., Varian, H.R., Zigmond, D.: Google's auction for tv ads. In: Albers, S., Marchetti-Spaccamela, A., Matias, Y., Nikoletseas, S.E., Thomas, W. (eds.) ICALP (2). Lecture Notes in Computer Science, vol. 5556, pp. 309–327. Springer (2009)
18. Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. The Journal of Finance 16(1), 8–37 (1961)