Pricing Ad Slots with Consecutive Multi-unit Demand*

Xiaotie Deng[1,2,**], Paul Goldberg[3,***], Yang Sun[4], Bo Tang[2], and Jinshan Zhang[2,***]

[1] Department of Computer Science, Shanghai Jiaotong University, China
dengxiaotie@gmail.com
[2] Department of Computer Science, University of Liverpool, UK
{Bo.Tang,Jinshan.Zhang}@liverpool.ac.uk
[3] Department of Computer Science, Oxford University, UK
Paul.Goldberg@cs.ox.ac.uk
[4] Department of Computer Science, City University of Hong Kong, Hong Kong
sunyang@live.hk

Abstract. We consider the optimal pricing problem for a model of the rich media advertisement market, as well as other related applications. In this market, there are multiple buyers (advertisers), and items (slots) that are arranged in a line such as a banner on a website. Each buyer desires a particular number of *consecutive* slots and has a per-unit-quality value v_i (dependent on the ad only) while each slot j has a quality q_j (dependent on the position only such as click-through rate in position auctions). Hence, the valuation of the buyer i for item j is $v_i q_j$. We want to decide the allocations and the prices in order to maximize the total revenue of the market maker.

A key difference from the traditional position auction is the advertiser's requirement of a fixed number of consecutive slots. Consecutive slots may be needed for a large size rich media ad. We study three major pricing mechanisms, the Bayesian pricing model, the maximum revenue market equilibrium model and an envy-free solution model. Under the Bayesian model, we design a polynomial time computable truthful mechanism which is optimum in revenue. For the market equilibrium paradigm, we find a polynomial time algorithm to obtain the maximum revenue market equilibrium solution. In envy-free settings, an optimal solution is presented when the buyers have the same demand for the number of consecutive slots. We conduct a simulation that compares the revenues from the above schemes and gives convincing results.

Keywords: mechanism design, revenue, advertisement auction.

---

* The full version of the paper is available at: http://arxiv.org/abs/1308.1382

** Supported by the National Science Foundation of China (Grant No. 61173011) and a Project 985 grant of Shanghai Jiaotong University.

*** Supported by EPSRC project EP/K01000X/1 "Efficient Algorithms for Mechanism Design without Monetary Transfer".

B. Vöcking (Ed.): SAGT 2013, LNCS 8146, pp. 255–266, 2013.

1 Introduction

Ever since the pioneering studies on pricing protocols for sponsored search advertisement, especially with the generalized second price auction (GSP), by Edelman, Ostrovsky, and Schwarz [9], as well as Varian [16], market making mechanisms have attracted much attention from the research community in understanding their effectiveness for the revenue maximization task facing platforms providing Internet advertisement services. In the traditional advertisement setting, advertisers negotiate ad presentations and prices with website publishers directly. An automated pricing mechanism simplifies this process by creating a bidding game for the buyers of advertisement space over an IT platform. It creates a complete competition environment for the price discovery process. Accompanying the explosion of the online advertisement business, there is a need to have a complete picture on what pricing methods to use in practical terms for both advertisers and Ad space providers.

In addition to search advertisements, display advertisements have been widely used in webpage advertisements. They have a rich format of displays such as text ads and rich media ads. Unlike sponsored search, there is a lack of systematic studies on its working mechanisms for decision makings. The market maker faces a combinatorial problem of whether to assign a large space to one large rich media ad or multiple small text ads, as well as how to decide on the prices charged to them. We present a study of the allocation and pricing mechanisms for displaying slots in this environment where some buyers would like to have one slot and others may want several consecutive slots in a display panel. In addition to webpage ads, another motivation of our study is TV advertising where inventories of a commercial break are usually divided into slots of a few seconds each, and slots have various qualities measuring their expected number of viewers and the corresponding attractiveness.

We discuss three types of mechanisms and consider the revenue maximization problem under these mechanisms, and compare their effectiveness in revenue maximization under a dynamic setting where buyers may change their bids to improve their utilities. Our results make an important step toward the understanding of the advantages and disadvantages of their uses in practice. Assume the ad supplier divides the ad space into small enough slots (pieces) such that each advertiser is interested in a position with a fixed number of *consecutive* pieces. In modelling values to the advertisers, we modify the position auction model from the sponsored search market [9,16] where each ad slot is measured by the Click Through Rates (CTR), with users' interest expressed by a click on an ad. Since display advertising is usually sold on a per impression (CPM) basis instead of a per click basis (CTR), the quality factor of an ad slot stands for the expected impression it will brings in unit of time. Unlike in the traditional position auctions, people may have varying demands (need different spaces to display their ads) in a rich media ad auction for the market maker to decide on slot allocations and their prices.

We will lay out the the specific system parameters and present our results in the following subsections.

1.1 Our Modeling Approach

We have a set of *buyers* (advertisers) and a set of *items* to be sold (the ad slots on a web page). We address the challenge of computing prices that satisfy certain desirable properties. Next we describe the elements of the model in more detail.

- **Items.** Our model considers the geometric organization of ad slots, which commonly has the slots arranged in some sequence (typically, from top to bottom in the right-hand side of a web page). The slots are of variable quality. In the study of sponsored search auctions, a standard assumption is that the quality (corresponding to click-through rate) is highest at the beginning of the sequence and then monotonically decreases. Here we consider a generalization where the quality may go down and up, subject to a limit on the total number of local maxima (which we call *peaks*), corresponding to focal points on the web page. As we will show later, without this limit the revenue maximization problem is NP-hard.
- **Buyers.** A buyer (advertiser) may want to purchase multiple slots, so as to display a larger ad. Note that such slots should be *consecutive* in the sequence. Thus, each buyer i has a fixed *demand* d_i, which is the number of slots she needs for her ad. Two important aspects of this are
  - ◇ *sharp* multi-unit demand, referring to the fact that buyer i should be allocated d_i items, or none at all; there is no point in allocating any fewer
  - ◇ *consecutiveness* of the allocated items, in the pre-existing sequence of items.

  These constraints give rise to a new and interesting combinatorial pricing problem.
- **Valuations.** We assume that each buyer i has a parameter v_i representing the value she assigns to a slot of unit quality. Valuations for multiple slots are additive, so that a buyer with demand d_i would value a block of d_i slots to be their total quality, multiplied by v_i. This valuation model has been considered by Edelman et al. [9] and Varian [16] in their seminal work for keywords advertising.

Pricing Mechanisms. Given the valuations and demands from the buyers, the market maker decides on a price vector for all slots and an allocation of slots to buyers, as an output of the market. The question is one of which output the market maker should choose to achieve certain objectives. We consider two approaches:

- **Truthful Mechanism** whereby the buyers report their demands (publicly known) and values (private) to the market maker; then prices are set in such a way as to ensure that the buyers have the incentive to report their true valuations. We give a revenue-maximizing approach (i.e., maximizing the total price paid), within this framework.
- **Competitive Equilibrium** whereby we prescribe certain constraints on the prices so as to guarantee certain well-known notions of fairness and envy-freeness.

- **Envy-Free Solution** whereby we prescribe certain constraints on the prices and allocations so as to achieve envy-freeness, as explained below.

The mechanisms we exhibit are computationally efficient. We also performed experiments to compare the revenues obtained from these three mechanisms.

1.2 Related Works

The theoretical study of position auctions (of a single slot) under the generalized second price auction was initiated in [9,16]. There has been a series of studies of position auctions in deterministic settings [12]. Our consideration of position auctions in the Bayesian setting fits in the general one dimensional auction design framework. Our study considers continuous distributions on buyers' values. For discrete distributions, [4] presents an optimal mechanism for budget constrained buyers without demand constraints in multi-parameter settings and very recently they also give a general reduction from revenue to welfare maximization in [5]; for buyers with both budget constraints and demand constraints, 2-approximate mechanisms [1] and 4-approximate mechanisms [3] exist in the literature.

There are extensive studies on multi-unit demand in economics, see for example [2,6,10]. In an earlier paper [7] we considered sharp multi-unit demand, where a buyer with demand d should be allocated d items or none at all, but with no further combinatorial constraint, such as the consecutiveness constraint that we consider here. The sharp demand setting is in contrast with a "relaxed" multi-unit demand (i.e., one can buy a subset of at most d items), where it is well known that the set of competitive equilibrium prices is non-empty and forms a distributive lattice [11,15]. This immediately implies the existence of an equilibrium with maximum possible prices; hence, revenue is maximized. Demange, Gale, and Sotomayor [8] proposed a combinatorial dynamics which always converges to a revenue maximizing (or minimizing) equilibrium for unit demand; their algorithm can be easily generalized to relaxed multi-unit demand. A strongly related work to our consecutive settings is the work of Rothkopf et al. [14], where the authors presented a dynamic programming approach to compute the maximum social welfare of consecutive settings when all the qualities are the same. Hence, our dynamic programming approach for general qualities in Bayesian settings is a non-trivial generalization of their settings.

1.3 Organization

This paper is organized as follows. In Section 2 we describe the details of our rich media ads model and the related solution concepts. In Section 3, we study the problem under the Bayesian model and provide a Bayesian Incentive Compatible auction with optimal expected revenue for the special case of the single peak in quality values of advertisement positions. Then in Section 4, we extend the optimal auction to the case with limited peaks/valleys and show that it is NP-hard to maximize revenue without this limit. Next, in Section 5, we turn to the full information setting and propose an algorithm to compute the competitive equilibrium with maximum revenue. In Section 6, NP-hardness of envy-freeness for consecutive multi-unit demand buyers is shown. We also design a polynomial time solution for the special case where all advertisers demand the same number of ad slots. For simulations, we refer readers to read the full version of the paper.

2 Preliminaries

In our model, a rich media advertisement instance consists of n advertisers and m advertising slots. Each slot $j \in \{1, \ldots, m\}$ is associated with a number q_j which can be viewed as the quality or the desirability of the slot. Each advertiser (or buyer) i wants to display her own ad that occupies d_i consecutive slots on the webpage. In addition, each buyer has a private number v_i representing her valuation and thus, the i-th buyer's value for item j is $v_{ij} = v_i q_j$.

Throughout this thesis, we will often say that slot j is assigned to a buyer set B to denote that j is assigned to some buyer in B. We will call the set of all slots assigned to B the allocation to B. In addition, a buyer will be called a winner if he succeeds in displaying his ad and a loser otherwise. We use the standard notation $[s]$ to denote the set of integers from 1 to s, i.e. $[s] = \{1, 2, \ldots, s\}$. We sometimes use $\sum_i$ instead of $\sum_{i\in[n]}$ to denote the summation over all buyers and $\sum_j$ instead of $\sum_{j\in[m]}$ for items, and the terms $\mathrm{E}_{\mathbf{v}}$ and $\mathrm{E}_{v_{-i}}$ are short for $\mathrm{E}_{\mathbf{v}\in\mathbf{V}}$ and $\mathrm{E}_{v_{-i}\in V_{-i}}$.

The vector of all the buyers' values is denoted by $\boldsymbol{v}$ or sometimes $(v_i; v_{-i})$ where v_{-i} is the joint bids of all bidders other than i. We represent a feasible assignment by a vector $\boldsymbol{x} = (x_{ij})_{i,j}$, where $x_{ij} \in \{0, 1\}$ and $x_{ij} = 1$ denotes item j is assigned to buyer i. Thus we have $\sum_i x_{ij} \leq 1$ for every item j. Given a fixed assignment x, we use t_i to denote the quality of items that buyer i is assigned, precisely, $t_i = \sum_j q_j x_{ij}$. In general, when x is a function of buyers' bids $\boldsymbol{v}$, we define t_i to be a function of $\boldsymbol{v}$ such that $t_i(\boldsymbol{v}) = \sum_j q_j x_{ij}(\boldsymbol{v})$.

When we say that slot qualities have a single peak, we mean that there exists a peak slot k such that for any slot $j < k$ on the left side of k, $q_j \geq q_{j-1}$ and for any slot $j > k$ on the right side of k, $q_j \geq q_{j+1}$.

2.1 Bayesian Mechanism Design

Following the work of [13], we assume that all buyers' values are distributed independently according to publicly known bounded distributions. The distribution of each buyer i is represented by a Cumulative Distribution Function (CDF) F_i and a Probability Density Function (PDF) f_i. In addition, we assume that the concave closure or convex closure or integration of those functions can be computed efficiently.

An auction $M = (\boldsymbol{x}, \boldsymbol{p})$ consists of an allocation function $\boldsymbol{x}$ and a payment function $\boldsymbol{p}$. $\boldsymbol{x}$ specifies the allocation of items to buyers and $\boldsymbol{p} = (p_i)_i$ specifies the buyers' payments, where both $\boldsymbol{x}$ and $\boldsymbol{p}$ are functions of the reported valuations $\boldsymbol{v}$. Our objective is to maximize the expected revenue of the mechanism is $Rev(M) = \mathrm{E}_{\boldsymbol{v}}\left[\sum_i p_i(\boldsymbol{v})\right]$ under Bayesian incentive compatible mechanisms.

Definition 1. *A mechanism M is called* Bayesian Incentive Compatible *(BIC) iff the following inequalities hold for all i, v_i, v_i'.*

$$\mathrm{E}_{v_{-i}}[v_i t_i(\boldsymbol{v}) - p_i(\boldsymbol{v})] \geq \mathrm{E}_{v_{-i}}[v_i t_i(v_i'; v_{-i}) - p_i(v_i'; v_{-i})] \tag{1}$$

Besides, we say M is Incentive Compatible *if M satisfies a stronger condition that $v_i t_i(\boldsymbol{v}) - p_i(\boldsymbol{v}) \geq v_i t_i(v_i'; v_{-i}) - p_i(v_i'; v_{-i})$, for all $\boldsymbol{v}, i, v_i'$,*

To put it in words, in a BIC mechanism, no player can improve her *expected* utility (expectation taken over other players' bids) by misreporting her value. An IC mechanism satisfies the stronger requirement that no matter what the other players declare, no player has incentives to deviate.

2.2 Competitive Equilibrium and Envy-free Solution

In Section 5, we study the revenue maximizing competitive equilibrium and envy-free solution in the full information setting instead of the Bayesian setting. An outcome of the market is a pair $(\boldsymbol{X}, \boldsymbol{p})$, where $\boldsymbol{X}$ specifies an allocation of items to buyers and $\boldsymbol{p}$ specifies prices paid. Given an outcome $(\boldsymbol{X}, \boldsymbol{p})$, recall $v_{ij} = v_i q_j$, let $u_i(\boldsymbol{X}, \boldsymbol{p})$ denote the *utility* of i.

Definition 2. *A tuple $(\boldsymbol{X}, \boldsymbol{p})$ is a* consecutive envy-free pricing *solution if every buyer is consecutive envy-free, where a buyer i is consecutive envy-free if the following conditions are satisfied:*

- *if $X_i \neq \emptyset$, then (i) X_i is d_i consecutive items. $u_i(\boldsymbol{X}, \boldsymbol{p}) = \sum_{j \in X_i} (v_{ij} - p_j) \geq 0$, and (ii) for any other subset of consecutive items T with $|T| = d_i$, $u_i(\boldsymbol{X}, \boldsymbol{p}) = \sum_{j \in X_i} (v_{ij} - p_j) \geq \sum_{j \in T} (v_{ij} - p_j)$;*
- *if $X_i = \emptyset$ (i.e., i wins nothing), then, for any subset of consecutive items T with $|T| = d_i$, $\sum_{j \in T} (v_{ij} - p_j) \leq 0$.*

Definition 3. *(Competitive Equilibrium) We say an outcome of the market $(\boldsymbol{X}, \boldsymbol{p})$ is a* competitive equilibrium *if it satisfies two conditions.*

- *$(\boldsymbol{X}, \boldsymbol{p})$ must be consecutive envy-free.*
- *The unsold items must be priced at zero.*

We are interested in the revenue maximizing competitive equilibrium and envy-free solutions.

3 Optimal Auction for the Single Peak Case

The goal of this section is to present our optimal auction for the single peak case that serves as an elementary component in the general case later. En route, several principal techniques are examined exhaustively to the extent that they can be applied directly in the next section. By employing these techniques, we show that the optimal Bayesian Incentive Compatible auction can be represented by a simple Incentive Compatible one. Furthermore, this optimal auction can be implemented efficiently. Let $T_i(v_i) = \mathrm{E}_{v_{-i}}[t_i(\boldsymbol{v})]$, $P_i(v_i) = \mathrm{E}_{v_{-i}}[p_i(\boldsymbol{v})]$ and $\phi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$. From Myerson' work [13], we obtain the following three lemmas.

Lemma 1 (From [13]). *A mechanism $M = (x,p)$ is Bayesian Incentive Compatible if and only if:*
a) $T_i(x)$ is monotone non-decreasing for any agent i.
b) $P_i(v_i) = v_i T_i(v_i) - \int_{\underline{v}_i}^{v_i} T_i(z)dz$

Lemma 2 (From [13]). *For any BIC mechanism $M = (x,p)$, the expected revenue $\mathrm{E}_{\boldsymbol{v}}[\sum_i P_i(v_i)]$ is equal to the virtual surplus $\mathrm{E}_{\boldsymbol{v}}[\sum_i \phi_i(v_i)t_i(\boldsymbol{v})]$.*

We assume $\phi_i(t)$ is monotone increasing, i.e. the distribution is regular. Otherwise, Myerson's ironing technique can be utilized to make $\phi_i(t)$ monotone — it is here that we invoke our assumption that we can efficiently compute the convex closure of a continuous function and integration. The following lemma is the direct result of Lemma 1 and 2.

Lemma 3. *Suppose that x is the allocation function that maximizes $\mathrm{E}_{\boldsymbol{v}}[\phi_i(v_i)t_i(\boldsymbol{v})]$ subject to the constraints that $T_i(v_i)$ is monotone non-decreasing for any bidders' profile $\boldsymbol{v}$, any agent i is assigned either d_i consecutive slots or nothing. Suppose also that*

$$p_i(\boldsymbol{v}) = v_i t_i(\boldsymbol{v}) - \int_{\underline{v}_i}^{v_i} t_i(v_{-i}, s_i)ds_i \tag{2}$$

Then (x,p) represents an optimal mechanism for the rich media advertisement problem in single-peak case.

We will use dynamic programming to maximize the virtual surplus in Lemma 2. Suppose all the buyers are sorted in a no-increasing order according to their virtual values. We will need the following two useful lemmas. Lemma 4 states that all the allocated slots are consecutive.

Lemma 4. *There exists an optimal allocation x that maximizes $\sum_i \phi_i(v_i)t_i(\boldsymbol{v})$ in the single peak case, and satisfies the following condition. For any unassigned slot j, it must be that either $\forall j' > j$, slot j' is unassigned or $\forall j' < j$, slot j' is unassigned.*

Next, we prove that this consecutiveness even holds for all set $[s] \subseteq [n]$. That is, there exists an optimal allocation that always assigns the first s buyers consecutively for all $s \in [n]$. For convenience, we say that a slot is "out of" a set of buyers if the slot is not assigned to any buyers in that set. Then the consecutiveness can be formalized in the following lemma.

Lemma 5. *There exists an optimal allocation x in the single peak case, that satisfies the following condition. For any slot j out of $[s]$, it must be either $\forall j' > j$, slot j' is out of $[s]$ or $\forall j' < j$, slot j' is out of $[s]$.*

Since the optimal solution always assigns to $[s]$ consecutively (Lemma 5), we can boil the allocations to $[s]$ down to an interval denoted by $[l, r]$. Let $g[s, l, r]$ denote the maximized value of our objective function $\sum_i \phi_i(v_i)t_i(\boldsymbol{v})$ when we only consider first s buyers and the allocation of s is exactly the interval $[l, r]$. Then we have the following transition function.

$$g[s,l,r] = \max \begin{cases} g[s-1,l,r] \\ g[s-1,l,r-d_s] + \phi_s(v_s)\sum_{j=r-d_s+1}^{r} q_j \\ g[s-1,l+d_s,r] + \phi_s(v_s)\sum_{j=l}^{l+d_s-1} q_j \end{cases} \tag{3}$$

Our summary statement is as follows.

Theorem 1. *The mechanism that applies the allocation rule according to Dynamic Programming (3) and payment rule according to Equation (2) is an optimal mechanism for the banner advertisement problem with single peak qualities.*

4 Multiple Peaks Case

Suppose now that there are only h peaks (local maxima) in the qualities. Thus, there are at most $h-1$ valleys (local minima). Since h is a constant, we can enumerate all the buyers occupying the valleys. After this enumeration, we can divide the qualities into at most h consecutive pieces and each of them forms a single-peak. Then using similar properties as those in Lemma 4 and 5, we can obtain a larger size dynamic programming (still runs in polynomial time) similar to dynamic programming (3) to solve the problem.

Theorem 2. *There is a polynomial algorithm to compute revenue maximization problem in Bayesian settings where the qualities of slots have a constant number of peaks.*

Now we consider the case without the constant peak assumption and prove the following hardness result.

Theorem 3. *(NP-Hardness) The revenue maximization problem for rich media ads with arbitrary qualities is NP-hard.*

5 Competitive Equilibrium

In this section, we study the revenue maximizing competitive equilibrium in the full information setting. To simplify the following discussions, we sort all buyers and items in non-increasing order of their values, i.e., $v_1 \geq v_2 \geq \cdots \geq v_n$.

We say an allocation $\boldsymbol{Y} = (Y_1, Y_2, \cdots, Y_n)$ is efficient if $\boldsymbol{Y}$ maximizes the total social welfare e.g. $\sum_i \sum_{j \in Y_i} v_{ij}$ is maximized over all the possible allocations. We call $\boldsymbol{p} = (p_1, p_2, \cdots, p_m)$ an equilibrium price if there exists an allocation $\boldsymbol{X}$ such that $(\boldsymbol{X}, \boldsymbol{p})$ is a competitive equilibrium. The following lemma is implicitly stated in [11], for completeness, we give a proof below.

Lemma 6. *Let allocation $\boldsymbol{Y}$ be efficient, then for any equilibrium price $\boldsymbol{p}$, $(\boldsymbol{Y}, \boldsymbol{p})$ is a competitive equilibrium.*

By Lemma 6, to find a revenue maximizing competitive equilibrium, we can first find an efficient allocation and then use linear programming to settle the prices. We develop the following dynamic programming to find an efficient allocation. We first only consider there is one peak in the quality order of items. The case with constant peaks is similar to the above approaches, for general peak case, as shown in above Theorem 3, finding one competitive equilibrium is NP-hard if the competitive equilibrium exists, and determining existence of competitive equilibrium is also NP-hard. This is because that considering the instance in the proof of Theorem 3, it is not difficult to see the constructed instance has an equilibrium if and only if 3 partition has a solution.

Recall that all the values are sorted in non-increasing order e.g. $v_1 \geq v_2 \geq \cdots \geq v_n$. $g[s, l, r]$ denotes the maximized value of social welfare when we only consider first s buyers and the allocation of s is exactly the interval $[l, r]$. Then we have the following transition function.

$$g[s, l, r] = \max \begin{cases} g[s-1, l, r] \\ g[s-1, l, r-d_s] + v_s \sum_{j=r-d_s+1}^{r} q_j \\ g[s-1, l+d_s, r] + v_s \sum_{j=l}^{l+d_s-1} q_j \end{cases} \tag{4}$$

By tracking procedure 4, an efficient allocation denoted by $\boldsymbol{X}^* = (X_1^*, X_2^*, \cdots, X_n^*)$ can be found. The price $\boldsymbol{p}^*$ such that $(\boldsymbol{X}^*, \boldsymbol{p}^*)$ is a revenue maximization competitive equilibrium can be determined from the following linear programming. Let T_i be any consecutive number of d_i slots, for all $i \in [n]$.

$$\begin{aligned} \max \quad & \sum_{i \in [n]} \sum_{j \in X_i^*} p_j \\ s.t. \quad & p_j \geq 0 && \forall\, j \in [m] \\ & p_j = 0 && \forall\, j \notin \cup_{i \in [n]} X_i^* \\ & \sum_{j \in X_i^*} (v_i q_j - p_j) \geq \sum_{j' \in T_i} (v_i q_{j'} - p_{j'}) && \forall\, i \in [n] \\ & \sum_{j \in X_i^*} (v_i q_j - p_j) \geq 0 && \forall i \in [n] \end{aligned}$$

Clearly there is only a polynomial number of constraints. The constraints in the first line represent that all the prices are non negative (no positive transfers). The constraint in the second line means unallocated items must be priced at zero (market clearance condition). And the constraint in the third line contains two aspects of information. First for all the losers e.g. loser k with $X_k = \emptyset$, the utility that k gets from any consecutive number of d_k is no more than zero, which makes all the losers envy-free. The second aspect is that the winners e.g. winner i with $X_i \neq \emptyset$ must receive a bundle with d_i consecutive slots maximizing its utility over all d_i consecutive slots, which together with the constraint in the fourth line (winner's utilities are non negative) guarantees that all winners are envy-free.

Theorem 4. *Under the condition of a constant number of peaks in the qualities of slots, there is a polynomial time algorithm to decide whether there exists a competitive equilibrium or not and to compute a revenue maximizing revenue market equilibrium if one does exist. If the number of peaks in the qualities of the slots is unbounded, both the problems are NP-complete.*

Proof. Clearly the above linear programming and procedure (4) run in polynomial time. If the linear programming output a price $\boldsymbol{p}^*$, then by its constraint conditions, $(\boldsymbol{X}^*, \boldsymbol{p}^*)$ must be a competitive equilibrium. On the other hand, if there exist a competitive equilibrium $(\boldsymbol{X}, \boldsymbol{p})$ then by Lemma 6, $(\boldsymbol{X}^*, \boldsymbol{p})$ is a competitive equilibrium, providing a feasible solution of above linear programming. By the objective of the linear programming, we know it must be a revenue maximizing one.

6 Consecutive Envy-freeness

We first prove a negative result on computing the revenue maximization problem in general demand case. We show it is NP-hard even if all the qualities are the same.

Theorem 5. *The revenue maximization problem of consecutive envy-free buyers is NP-hard even if all the qualities are the same.*

Although the hardness in Theorem 5 indicates that finding the optimal revenue for general demand in polynomial time is impossible , however, it doesn't rule out the very important case where the demand is uniform, e.g. $d_i = d$. We assume slots are in a decreasing order from top to bottom, that is, $q_1 \geq q_2 \geq \cdots \geq q_m$. The result is summarized as follows.

Theorem 6. *There is a polynomial time algorithm to compute the consecutive envy-free solution when all the buyers have the same demand and slots are ordered from top to bottom.*

The proof of Theorem 6 is based on bundle envy-free solutions, in fact we will prove the bundle envy-free solution is also a consecutive envy-free solution by defining price of items properly. Thus, we need first give the result on bundle envy-free solutions. Suppose d is the uniform demand for all the buyers. Let T_i be the slot set allocated to buyer i, $i = 1, 2, \cdots, n$. Let P_i be the total payment of buyer i and p_j be the price of slot j. Let t_i denote the total qualities obtained by buyer i, e.g. $t_i = \sum_{j \in T_i} q_j$ and $\alpha_i = iv_i - (i-1)v_{i-1}$, $\forall i \in [n]$.

Theorem 7. *The revenue maximization problem of bundle envy-freeness is equivalent to solving the following LP.*

$$\begin{aligned} \textit{Maximize:} \quad & \sum_{i=1}^{n} \alpha_i t_i \\ s.t. \quad & t_1 \geq t_2 \geq \cdots \geq t_n \\ & T_i \subset [m], \quad T_i \cap T_k = \emptyset \quad \forall i, k \in [n] \end{aligned} \tag{5}$$

Through optimal bundle envy-free solution, we will modify such a solution to consecutive envy-free solution and then prove the Theorem 6.

7 Conclusion and Discussion

The rich media pricing models for consecutive demand buyers in the context of Bayesian truthfulness, competitive equilibrium and envy-free solution paradigm are investigated in this paper. As a result, an optimal Bayesian incentive compatible mechanism is proposed for various settings such as single peak and multiple peaks. In addition, to incorporate fairness e.g. envy-freeness, we also present a polynomial-time algorithm to decide whether or not there exists a competitive equilibrium or and to compute a revenue maximized market equilibrium if one does exist. For envy-free settings, though the revenue maximization of general demand case is shown to be NP-hard, we still provide optimal solution of common demand case. Besides, our simulation shows a reasonable relationship of revenues among these schemes plus a generalized GSP for rich media ads.

Even though our main motivation arises from the rich media advert pricing problem, our models have other potential applications. For example TV ads can also be modeled under our consecutive demand adverts where inventories of a commercial break are usually divided into slots of fixed sizes, and slots have various qualities measuring their expected number of viewers and corresponding attractiveness. With an extra effort to explore the periodicity of TV ads, we can extend our multiple peak model to one involved with cyclic multiple peaks. Besides single consecutive demand where each buyer only have one demand choice, the buyer may have more options to display his ads, for example select a large picture or a small one to display them. Our dynamic programming algorithm (3) can also be applied to this case (the transition function in each step selects maximum value from $2k + 1$ possible values, where k is the number of choices of the buyer).

Another reasonable extension of our model would be to add budget constraints for buyers, i.e., each buyer cannot afford the payment more than his budget. By relaxing the requirement of Bayesian incentive compatible (BIC) to one of approximate BIC, this extension can be obtained by the recent milestone work of Cai et al. [5]. It remains an open problem how to do it under the exact BIC requirement. It would also be interesting to handle it under the market equilibrium paradigm for our model.

References

1. Alaei, S.: Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. In: Proceedings of the 52nd IEEE Symposium on Foundations of Computer Science (FOCS), pp. 512–521 (2011)
2. Ausubel, L., Cramton, P.: Demand revelation and inefficiency in multi-unit auctions. In: Mimeograph, University of Maryland (1996)
3. Bhattacharya, S., Goel, G., Gollapudi, S., Munagala, K.: Budget constrained auctions with heterogeneous items. In: Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, New York, NY, USA, pp. 379–388 (2010)
4. Cai, Y., Daskalakis, C., Matthew Weinberg, S.: An algorithmic characterization of multi-dimensional mechanisms. In: Proceedings of the 43rd Annual ACM Symposium on Theory of Computing (2012)
5. Cai, Y., Daskalakis, C., Matthew Weinberg, S.: Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In: Proceedings of the 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science, FOCS 2012 (2012)
6. Cantillon, E., Pesendorfer, M.: Combination bidding in multi-unit auctions. C.E.P.R. Discussion Papers (February 2007)
7. Chen, N., Deng, X., Goldberg, P.W., Zhang, J.: On revenue maximization with sharp multi-unit demands. CoRR abs/1210.0203 (2012)
8. Demange, G., Gale, D., Sotomayor, M.: Multi-item auctions. The Journal of Political Economy, 863–872 (1986)
9. 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 (2007)
10. Engelbrecht-Wiggans, R., Kahn, C.M.: Multi-unit auctions with uniform prices. Economic Theory 12(2), 227–258 (1998)
11. Gul, F., Stacchetti, E.: Walrasian equilibrium with gross substitutes. Journal of Economic Theory 87(1), 95–124 (1999)
12. Lahaie, S.: An analysis of alternative slot auction designs for sponsored search. In: Proceedings of the 7th ACM Conference on Electronic Commerce, EC 2006, New York, NY, USA, pp. 218–227 (2006)
13. Myerson, R.B.: Optimal auction design. Mathematics of Operations Research 6(1), 58–73 (1981)
14. Rothkopf, M.H., Pekeč, A., Harstad, R.M.: Computationally manageable combinational auctions. Management Science 44(8), 1131–1147 (1998)
15. Shapley, L.S., Shubik, M.: The Assignment Game I: The Core. International Journal of Game Theory 1(1), 111–130 (1971)
16. Varian, H.R.: Position auctions. International Journal of Industrial Organization 25(6), 1163–1178 (2007)