THEORY OF COMPUTING, Volume 3 (2007), pp. 179–195
http://theoryofcomputing.org

Approximation Algorithms and Online Mechanisms for Item Pricing*

Maria-Florina Balcan† Avrim Blum†

Received: March 29, 2007; published: October 3, 2007.

Abstract: We present approximation and online algorithms for problems of pricing a collection of items for sale so as to maximize the seller's revenue in an unlimited supply setting. Our first result is an $O(k)$-approximation algorithm for pricing items to single-minded bidders who each want at most k items. This improves over work of Briest and Krysta (2006) who achieve an $O(k^2)$ bound. For the case $k = 2$, where we obtain a 4-approximation, this can be viewed as the following *graph vertex pricing* problem: given a (multi) graph G with valuations w_{ij} on the edges, find prices $p_i \geq 0$ for the vertices to maximize

$$\sum_{\{(i,j):w_{ij}\geq p_i+p_j\}} (p_i + p_j) .$$

We also improve the approximation of Guruswami et al. (2005) for the "highway problem" in which all desired subsets are intervals on a line, from $O(\log m + \log n)$ to $O(\log n)$, where m is the number of bidders and n is the number of items. Our approximation algorithms can be fed into the generic reduction of Balcan et al. (2005) to yield an incentive-compatible

*A preliminary version of this paper appeared in Proc. 7th ACM Conf. on Electronic Commerce (EC'06), pp.29-35, 2006.

†Supported in part by NSF grants CCF-0514922, CCR-0122581, and IIS-0121678.

ACM Classification: F.2, J.4

AMS Classification: 68W25, 68W20, 68Q32, 91B26

Key words and phrases: Combinatorial Auctions, Pricing Problems, Revenue Maximization, Approximation Algorithms, Online Optimization

Authors retain copyright to their work and grant Theory of Computing unlimited rights to publish the work electronically and in hard copy. Use of the work is permitted as long as the author(s) and the journal are properly acknowledged. For the detailed copyright statement, see http://theoryofcomputing.org/copyright.html.

© 2007 Maria-Florina Balcan and Avrim Blum

DOI: 10.4086/toc.2007.v003a009

auction with nearly the same performance guarantees so long as the number of bidders is sufficiently large. In addition, we show how our algorithms can be combined with results of Blum and Hartline (2005) and Kalai and Vempala (2003) to achieve good performance in the online setting, where customers arrive one at a time and each must be presented a set of item prices based only on knowledge of the customers seen so far.

1 Introduction

Consider the problem of a retailer trying to price its products to make the most profit. If customers had valuations over individual items only, then the problem of setting prices would be relatively easy: for each product i, the optimal price for that product is such that the profit margin p_i per item sold, times the number of customers who would buy at that price, is maximized. So, each item can be considered separately, and assuming the company knows its market well, the *computational* problem of setting prices is fairly trivial.

However, suppose that customers have valuations over *pairs* of items (e.g., a computer and a monitor, or a tank of gas and a cup of coffee), and will only purchase if the combined price of the items in their pair is below their value. In this case, we can model the problem as a (multi) *graph*, where each edge e has some valuation w_e, and our goal is to set prices $p_i \geq 0$ on the vertices of the graph to maximize total *profit*: that is,[1]

$$\text{Profit}(\mathbf{p}) = \sum_{e:w_e \geq \text{price}(e)} \text{price}(e),$$

where $\text{price}(e) = \sum_{i \in e} p_i$, and $\mathbf{p}$ is the vector of individual prices.

We call this the *graph vertex pricing* problem. More generally, if customers have valuations over larger subsets, we can model our computational problem as one of pricing vertices in a *hypergraph*, or in more standard terminology, the problem of pricing items in an unlimited-supply combinatorial auction with single-minded bidders. Guruswami et al. [15] show an $O(\log m + \log n)$-approximation for the general problem, where n is the number of items (vertices) and m is the number of customers (hyperedges). They also show that even the *graph* vertex pricing problem is APX-hard — and this is true even when all valuations are identical (if customers wanting just one item are allowed) or all valuations are either 1 or 2 (if such customers are not allowed). In related work, Hartline and Koltun [16] give a $(1+\varepsilon)$-approximation that runs in time exponential in the number of vertices, but that is near-linear time when the total number of vertices in the hypergraph is constant. Recently, Demaine et al. [10] have shown that it is hard to approximate the hypergraph vertex pricing problem within a factor of $\log^{\delta} n$, for some $\delta > 0$, assuming that NP $\not\subseteq$ BPTIME$\left(2^{n^{\varepsilon}}\right)$ for some $\varepsilon > 0$.

In this paper, we give a 4-approximation for the graph vertex pricing problem, and more generally we present an $O(k)$-approximation for the case of hypergraphs in which each edge has size at most k

[1] This formula corresponds to a model in which items have zero marginal cost to the retailer (digital goods) so that an item sold at price p_i generates profit p_i. Alternatively, if products have a fixed marginal cost, and we cannot sell them below cost (say, due to the presence of resellers), then we can think of p_i as the profit margin on item i and simply subtract our costs for the endpoints from each valuation w_e.

(i. e., all customers' valuations are over subsets of size at most k). The latter result improves over the work of Briest and Krysta [7] who give a bound of $O(k^2)$.

We also consider the highway problem studied in [15]. This problem is the special case of the hypergraph pricing problem where vertices are numbered $1,\ldots,n$ and each customer wants an interval $[i,j]$.[2] For this problem, we give an $O(\log n)$-approximation, improving slightly over the $O(\log m+\log n)$ approximation of [15], and also give an $O(1)$-approximation for the case that all users want the *same number* of items up to a constant factor. In addition, we give a fully polynomial time approximation scheme (FPTAS) for the case that the desired subsets of different customers form a hierarchy (this is defined more precisely in Section 7).

Finally, we consider the question of what happens if we are allowed to price some items *below* their cost, and give an example in the context of graph vertex pricing in which such pricing can produce an $\Omega(\log n)$ factor more profit than possible if all items must be priced above cost. However, we do not have any good ($o(\log n)$) approximation algorithms for that setting.

Incentive-compatibility Our results described above assume the seller "understands the market": how many customers will buy different sets of items and at what prices. Thus, we are simply left with a computational problem. If we do not understand the market and are in the setting of an unlimited-supply combinatorial auction, we would instead want an algorithm that is *incentive-compatible*, meaning that it is in the bidders' self-interest to reveal their true valuations. Fortunately, a generic reduction of [3] shows that if there are sufficiently many bidders, then for problems of this type one can convert any approximation to the computational problem into a nearly-as-good approximation to the incentive-compatible auction problem. In particular, $\tilde{O}\left(hn/\varepsilon^2\right)$ bidders are sufficient for this reduction to produce only a factor $(1+\varepsilon)$ loss in approximation ratio when all valuations lie in the range $[1,h]$. Essentially, the idea of the reduction is to randomly partitions bidders into two sets S_1 and S_2, run the approximation algorithm separately on each set, and then use the prices found for S_1 on S_2 and vice-versa (making the process incentive-compatible); the results in [3] then show that $\tilde{O}\left(hn/\varepsilon^2\right)$ bidders are sufficient to ensure that the resulting profit is nearly as large as if one had used prices determined on each S_i on that set itself. Related results of [14, 13] give bounds of this form for the case of a single digital good. Thus, if one has sufficiently many bidders, one can focus attention on the computational approximation problem.

The above results assume a one-shot mechanism (sealed-bid auction) in which all bidders are present at the same time. We also consider the more demanding case that bidders arrive online, and one must present to each bidder a set of item prices that depend only on bidders seen in the past. We show how methods of [5, 6] for the online digital-good auction can be applied to our algorithms for graph (or k-hypergraph) vertex pricing to achieve good performance for these problems in the online setting as well. For the highway problem, we need a somewhat more involved argument using an algorithm of Kalai and Vempala [18].

[2] Previous work [16, 15] uses "m" to denote the number of items and "n" to denote the number of customers, viewing the *items* as edges in some network. Since we are viewing items as vertices and customers as (hyper)edges, we have reversed this notation.

Organization The rest of this paper is organized as follows. We begin with basic definitions in Section 2. We then present our 4-approximation algorithm for the graph vertex problem in Section 3 and our $O(k)$-approximation algorithm for the k-Hypergraph Vertex Pricing problem in Section 4. We discuss pricing below cost in Section 5, and present an $O(\log n)$ approximation for the highway problem in Section 6. We present a fully polynomial time approximation scheme (FPTAS) for the case that the desired subsets of different customers form a hierarchy in Section 7, and we show how our algorithms can be adapted to achieve good performance in the online setting in Section 8. We finish with a discussion of open questions in Section 9.

2 Notation and Definitions

We consider the following model introduced by Guruswami et al. [15]. We assume we have m customers (or "bidders") and n items (or "products"). We are in an *unlimited supply* setting, which means that the seller is able to sell any number of units of each item, and they each have zero marginal cost to the seller (or if they have some fixed marginal cost, we have subtracted that from all valuations and the seller may not sell any item below cost). We consider *single-minded bidders*, which means that each customer is interested in only a single bundle of items and has valuation 0 for all other bundles. Therefore, valuations can be summarized by a set of pairs (e, w_e) indicating that a customer is interested in bundle (hyperedge) e and values it at w_e. Given the hyperedges e and valuations w_e, we wish to compute a pricing of the items that maximizes the seller's profit. We assume that if the total price of the items in e is at most w_e, then the customer (e, w_e) will purchase all of the items in e, and otherwise the customer will purchase nothing. That is, we want the price vector $\mathbf{p} = (p_1, \ldots, p_n)$ with $p_i \geq 0$ for all i that maximizes

$$\text{Profit}(\mathbf{p}) = \sum_{e : w_e \geq \text{price}(e)} \text{price}(e), \qquad \text{where price}(e) = \sum_{i \in e} p_i .$$

Let $\mathbf{p}^*$ be the price vector with the maximum profit and let $\text{OPT} = \text{Profit}(\mathbf{p}^*)$.

Let us denote by E the set of customers, and V the set of items, and let $h = \max_{e \in E} w_e$. Let $G = (V, E)$ be the induced hypergraph, whose vertices represent the set of items, and whose hyperedges represent the customers. Notice that G might contain multi-edges since several customers might want the same subset of items. In the special case that all customers want at most two items, so G is a multi-graph (possibly with self-loops), we call this the *graph vertex pricing* problem. As mentioned in Section 1, this pricing problem was shown to be APX-hard in [15]. If all customers want at most k items, we call this the *k-hypergraph vertex pricing* problem. Guruswami et al. [15] present a simple $O(\log m + \log n)$ approximation algorithm for the general hypergraph vertex pricing problem.[3] Our goal is to achieve better guarantees for the graph or k-hypergraph case when $k = o(\log n)$.

3 Graph Vertex Pricing

We begin by considering the Graph Vertex Pricing problem, and show a factor 4 approximation.

[3] In fact, it has been shown recently in [4] that one can achieve an $O(\log m + \log n)$ approximation for bidders with *general valuation functions* (not only single-minded bidders).

Theorem 3.1. *There is a* 4*-approximation for the Graph Vertex Pricing problem.*

Proof. First notice that if G is *bipartite* (with self-loops allowed as well), then there is a simple 2-approximation algorithm. Specifically, consider the optimal price-vector $\mathbf{p}^*$ and let OPT_L be the amount of profit it makes from selling nodes on the left, and OPT_R be the amount it makes from selling nodes on the right (so $\mathrm{OPT} = \mathrm{OPT}_L + \mathrm{OPT}_R$). Notice that if one takes $\mathbf{p}^*$ and zeroes out all prices for nodes on the right, then this has profit at least OPT_L since all previous buyers still buy (and some new ones may too).[4] Therefore, we can algorithmically make profit at least OPT_L by setting all prices on the right to 0, and then separately fixing prices for each node on the left so as to make the most profit possible from each node. That is, for each node i, we simply order the buyers who want i by valuation $w_{e_1} \geq w_{e_2} \geq w_{e_3} \ldots$, and choose the price $p_i = w_{e_j}$ maximizing $j w_{e_j}$. (Since the graph is bipartite, the profit made from some node i on the left does not effect the optimal price for some other node i' on the left.) Similarly we can make at least OPT_R by setting prices on the left to 0 and optimizing prices of nodes on the right. So, taking the best of both options, we make

$$\max\left(\mathrm{OPT}_L, \mathrm{OPT}_R\right) \geq \frac{\mathrm{OPT}}{2}.$$

Now consider the general (non-bipartite) case. Define opt_e to be the amount of profit that OPT makes from edge e. We will think of opt_e as the *weight* of edge e, though it is unknown to our algorithm. Let E_2 be the set of edges that have two distinct endpoints, and let E_1 be the set of self-loops. Let OPT_1 be the profit made by $\mathbf{p}^*$ on edges in E_1 and let OPT_2 be the profit made by $\mathbf{p}^*$ on edges in E_2, so $\sum_{e \in E_i} opt_e = \mathrm{OPT}_i$ for $i = 1, 2$ and $\mathrm{OPT}_1 + \mathrm{OPT}_2 = \mathrm{OPT}$. Now, *randomly* partition the vertices into two sets L and R. Since each edge $e \in E_2$ has a 50% chance of having its endpoints on different sides, in expectation $\frac{1}{2}\mathrm{OPT}_2$ weight is on edges with one endpoint in L and one endpoint in R. Thus, if we simply ignore edges in E_2 whose endpoints are on the same side and run the algorithm for the bipartite case, the profit we make in expectation is at least

$$\frac{1}{2}\left[\mathrm{OPT}_1 + \frac{\mathrm{OPT}_2}{2}\right] \geq \frac{\mathrm{OPT}}{4}.$$

This proves the desired result. □

Derandomization If desired, the above algorithm can be derandomized using the fact that our analysis only needs the partitioning distribution to be pairwise-independent. In particular, pairwise-independent distributions can be realized using small (polynomial-size) sample spaces [20, 22]. Thus, given a problem instance, one can simply try each possibility in the sample space and then choose the one that produces the highest profit.

4 k-Hypergraph Vertex Pricing

We now show how to extend the algorithm in Theorem 3.1 to get an $O(k)$-approximation when each customer wants at most k items. This improves over the $O(k^2)$ bound of [7].

[4] Note that it is essential in this argument that $p_i^* \geq 0$ for all i.

Theorem 4.1. *There is an $O(k)$-approximation algorithm for the k-Hypergraph Vertex Pricing problem.*

Proof. We can use the following procedure.

Step 1 Randomly partition V into V_L and V_{rest} by placing each node into V_L with probability $\frac{1}{k}$.

Step 2 Let E' be the set of edges with *exactly* one endpoint in V_L. Ignore all edges in $E - E'$.

Step 3 Set prices in V_{rest} to 0 and set prices in V_L optimally with respect to edges in E'.

To analyze this algorithm, let $\text{OPT}_{i,e}$ denote the profit made by $\mathbf{p}^*$ selling item i to bidder e. (So $\text{OPT}_{i,e} \in \{0, p_i^*\}$ and $\text{OPT} = \sum_{i \in V, e \in E} \text{OPT}_{i,e}$.) Notice that the total profit made in Step 3 is *at least* $\sum_{i \in V_L, e \in E'} \text{OPT}_{i,e}$ because setting prices in V_{rest} to 0 can only increase the number of sales made by $\mathbf{p}^*$ to bidders in E'. Thus, we simply need to analyze the quantity $\mathbf{E}\left[\sum_{i \in V_L, e \in E'} \text{OPT}_{i,e}\right]$.

Define indicator random variable $X_{i,e} = 1$ if $i \in V_L$ and $e \in E'$, and $X_{i,e} = 0$ otherwise. We have:

$$\mathbf{E}[X_{i,e}] = \mathbf{Pr}[i \in V_L \text{ and } e \in E'] \geq \frac{1}{k}\left(1 - \frac{1}{k}\right)^{k-1}. \tag{4.1}$$

Therefore,

$$\begin{aligned}
\mathbf{E}\left[\sum_{i \in V_L, e \in E'} \text{OPT}_{i,e}\right] &= \mathbf{E}\left[\sum_{i \in V, e \in E} X_{i,e}\text{OPT}_{i,e}\right] \\
&= \sum_{i \in V, e \in E} \mathbf{E}[X_{i,e}]\,\text{OPT}_{i,e} \\
&\geq \frac{1}{k}\left(1 - \frac{1}{k}\right)^{k-1} \text{OPT} \\
&\geq \frac{\text{OPT}}{ke}.
\end{aligned}$$

□

Derandomization As with the algorithm of Theorem 3.1, the above algorithm for k-hypergraph vertex pricing can also be derandomized if desired, but in this case we need the tools of Even et al. [12]. First, note that we are only interested in the case that k is $o(\log n + \log m)$, since for larger values of k we can switch to the generic algorithm of Guruswami et al. [15]. Thus, we can allow for a blowup of $2^{O(k)}$ in our running time. Now, consider the algorithm in Theorem 4.1 and define indicator random variables $X_i = 1$ if $i \in V_L$ and $X_i = 0$ otherwise. So, each $X_i = 1$ with probability $1/k$, and notice that we need only k-wise independence among the X_i to calculate $\mathbf{E}[X_{i,e}]$ in Equation (4.1). Even et al. [12] give a construction of small sample spaces that is especially well-suited to our needs. Their construction runs in time polynomial in 2^k, n, and $1/\varepsilon$, and produces an explicit sample space with the following property: for any k-tuple $(X_{i_1}, \ldots, X_{i_k})$ of the random variables X_i and any assignment $(v_1, \ldots, v_k)$ to their values, the fraction of points in their sample space under which these variables all take on those values is within $\pm\varepsilon$ of the probability of this event under our product distribution. In particular, the k-tuples we care

about are those corresponding to edges $e \in E$, with values of the form $(1,0,\ldots,0)$ corresponding to the event that $X_{i,e} = 1$. Setting $\varepsilon = o(1/k)$, we get that under the uniform distribution over their sample space, Equation (4.1) holds up to $1 - o(1)$, which suffices for our bounds. Thus, we simply run the construction of Even al. [12] using such a value of ε, and try each partitioning in their explicit sample space, choosing the one that produces the highest profit.

5 Pricing below Cost

Following prior work [15] we have so far required solutions to satisfy $p_i \geq 0$ for all i. In fact, for the case of digital goods, it does not make sense to allow negative prices since even customers e for whom $i \notin e$ would purchase item i if $p_i < 0$ (and moreover purchase infinitely many copies). However, in the case of products of nonzero marginal cost, where we view p_i as a *profit margin* (the difference between sales price and the retailer's cost) it could make sense to allow p_i to be negative. In fact, a retailer might wish to do so in order to induce more purchases of bundles containing both those and other more expensive products. For example, consider four items A, B, C, and D, and three customers: one who values $\{A,B\}$ at \$10 above their combined cost, one who values $\{B,C\}$ at \$40 above their cost, and one who values $\{C,D\}$ at \$10 above their cost. If no item can be priced at a loss, then it is not possible to have all three customers buy at their valuations. On the other hand, by pricing A and D at \$10 below cost, and B and C at \$20 above cost, the seller could extract full profit (assuming all costs are at least \$10). More generally, we present an $\Omega(\log n)$ "positivity gap": a (bipartite) graph in which there is an $\Omega(\log n)$ gap between the optimal profit achievable without any items priced at a loss and the optimal profit if such pricing is allowed. Specifically:

Theorem 5.1. *For the graph vertex pricing problem, there exists an $\Omega(\log n)$ gap between the profit achievable when pricing below cost is allowed and the profit achievable when pricing below cost is not allowed.*

Proof. Consider a bipartite graph with vertices $\ell_1,\ldots,\ell_n$ on the left and $r_1,\ldots,r_n$ on the right, where for convenience let n be a power of 2. A set S_1 of bidders each want bundles of the form $\{\ell_i, r_{i+1}\}$ at \$1 above cost, a set S_2 of bidders each want bundles of the form $\{\ell_i, r_{i+2}\}$ at \$2 above cost, a set S_4 of bidders each want bundles of the form $\{\ell_i, r_{i+4}\}$ at \$4 above cost, and so forth, up to $S_{n/2}$. Suppose there are $(n-1)n$ bidders in S_1 (n for each value of $i \in \{1,\ldots,n-1\}$), there are $(n-2)n/2$ bidders in S_2 ($n/2$ for each value of $i \in \{1,\ldots,n-2\}$), there are $(n-4)n/4$ bidders in S_4 ($n/4$ for each value of $i \in \{1,\ldots,n-4\}$) and so on, down to $(n/2)2$ bidders in $S_{n/2}$ (2 for each $i \in \{1,\ldots,n/2\}$). In this case, if negative profit margins are allowed, then one can price each ℓ_i at profit $-i$ and each r_i at profit i and have all bidders buy at exactly their valuations, extracting full profit $\Theta(n^2 \log n)$. On the other hand, the structure of bidders is such that the set of bidders who want any given item form an "equal revenue distribution." In particular, for any pricing in which all items on one side (say on the right) are priced at zero, any pricing of items on the other can produce at most $\Theta(n)$ profit per item, or $\Theta(n^2)$ total. This in turn implies that it is not possible to get more than $\Theta(n^2)$ profit using only non-negative profit margins, since we know the optimal profit achievable using non-negative profit margins is within a factor of 2 of the profit achievable when setting all items on one side of the graph to 0. □

Notice that our approximation algorithms in Sections 3 and 4 only provide good guarantees with respect to the best profit achievable when pricing below cost is not allowed. We do not know whether it is possible to achieve a $o(\log n)$ approximation when pricing below cost is allowed, even for the case of bipartite graphs.

For the rest of the paper, we resume considering only the case that all prices must be nonnegative.

6 The Highway Problem

A particular interesting case of the hypergraph pricing problem considered in [15] is the *highway* problem. In this problem we think of the items $1, 2, \ldots, n$ as segments of a highway, and each desired subset e is an interval $[i, j]$ of the highway. A special case of this problem shown in [15] to be solvable in polynomial time is the case when all path requests share one common end-point r. For this case, Guruswami et al. [15] give an $O(m^2)$ exact dynamic programming algorithm, which we will call $\mathcal{A}$. They also give pseudo-polynomial dynamic programming algorithms for two particular cases: an $O(h^{h+2}m^{h+3})$-time exact dynamic programming algorithm for the case when all valuations are integral, and an $O(h^{k+1}m)$ time exact dynamic programming algorithm for the case that furthermore all requests have path lengths bounded by some constant k. The highway problem was recently shown to be weakly NP-hard by Briest and Krysta [7].

We present below (Theorem 6.1) an $O(\log n)$ approximation algorithm for the highway problem, improving over the $O(\log n + \log m)$ approximation guarantee of Guruswami et al. [15].

Theorem 6.1. *There is an $O(\log n)$-approximation algorithm for the highway problem.*

Proof. For convenience we assume n is a power of 2 (which we can always achieve by padding). We begin by partitioning the customers into $\log_2 n$ groups. Specifically, let S_1 be the set of all customers who want item $n/2$. Let S_2 be the set of all customers not in S_1 who want either item $n/4$ or item $3/4$. More generally, let S_i be the set of customers not in $S_1 \cup \cdots \cup S_{i-1}$ who want some item in

$$\left\{ \frac{n}{2^i}, \frac{3n}{2^i}, \ldots, \frac{(2^i - 1)n}{2^i} \right\}.$$

Now, for each set S_i we can use algorithm $\mathcal{A}$ from [15] to get a 2-approximation to the optimal profit over S_i. Specifically, for each $j \in \{1, \ldots, 2^i - 1\}$ let S_{ij} be the subset of customers in S_i who want item $jn/2^i$. Notice that by design, customers in set S_{ij} do not have any desired item in common with customers in $S_{ij'}$ for $j' \neq j$, which means we can consider each of them separately. Now, for each S_{ij} we get a 2-approximation to $\mathrm{OPT}(S_{ij})$ by running $\mathcal{A}$ twice, first zeroing out all prices for items $z < jn/2^i$ and then again zeroing out all prices for items $z > jn/2^i$ and taking the best of the two cases. Since there are only $\log_2 n$ groups S_i, we simply use the algorithm $\mathcal{A}$ from [15] to get a 2-approximation to the optimal profit over S_i, and then take the best of all options, thus obtaining a $2\log_2 n$ approximation overall. □

6.1 Special cases

Using algorithm $\mathcal{A}$ we can also get a constant-factor approximation in the special case that everyone wants exactly k items, for any (not necessarily constant) k. To see this, split the items into groups

$G_1, G_2, \ldots, G_{n/k}$ of size k, where G_1 consists of the first k items, G_2 consists of the next k items and so on, and let OPT_{even} and OPT_{odd} be the amount of money that OPT makes from the even-numbered groups and from the odd-numbered groups respectively. We can make at least $\mathrm{OPT}_{even}/2$ as follows. We first set all the prices on items in the odd groups to zero. Now notice that each customer wants items in at most one even-numbered group: let us associate that customer with that group. We can now partition the customers in each even group into two types: those who want the leftmost item in the group and those who want the rightmost item in the group (customers who want both can be assigned arbitrarily). We then run the dynamic program separately over each type, and take the best outcome. In a similar way, we can make at least $\mathrm{OPT}_{odd}/2$ by setting prices for items in the even groups to 0. So we try both and take the best, thus obtaining a factor of 4 algorithm.

Similarly we can get a factor of $4c$ approximation algorithm if for some value of k, all customers want between k/c and k elements.

7 When bidders form a hierarchy

We present a fully polynomial time approximation scheme for the case that the desired subsets of different (single-minded) customers form a hierarchy, also known as a laminar family of sets.[5] Specifically, we consider the case of a hypergraph where for any two edges e, e', we have $e \subseteq e'$ or $e \supseteq e'$ or $e \cap e' = \emptyset$. This means that the edges themselves can be viewed as forming a tree structure (actually, a forest) ordered by containment. Let T_e be the set of all bidders whose desired subset is contained in e. Note that we can assume for simplicity that we have a binary hierarchy (if the hierarchy is not binary, then we can transform it into a binary hierarchy by adding fake edges e, increasing the size of the hypergraph by at most a constant factor).

We start by presenting a pseudopolynomial algorithm for the case that the bidders have integral valuations (between 0 and h). In this case, by the integrality lemma in [15] there exists an integral optimal solution. For each $e \in E$ and nonnegative integer $s \leq h$, let us denote by n_s^e the number of bidders with desired set e whose valuations are at least s. Now, for each $e \in E$ and nonnegative integer $s \leq nh$, let $A[s,e]$ represent the maximum possible profit we get from bidders in T_e when the total sum of the prices on items in e is exactly s. Our dynamic programming algorithm for computing the quantities $A[s,e]$ can be now specified as follows.

Step 1 For each "leaf" e in the hierarchy (an edge e that does not contain any other edges e') initialize $A[s,e] = s \cdot n_s^e$.

Step 2 Consider any edge e with children e_1 and e_2 whose A-values have been computed. Compute

$$A[s,e] = \max_{s_1+s_2=s} (A[s_1,e_1] + A[s_2,e_2]) + sn_s^e.$$

Step 3 Return $\max_{s \leq nh} A[s,r]$. (Here r is the root of the hierarchy).

[5] Independently, Briest and Krysta [7] show a similar result. They also show that this problem in NP-hard.

After computing the A-values, we can then easily determine the optimal pricing vector by backtracking. Clearly, the overall procedure above runs in time polynomial in n, m and h.

If we do not want to have a polynomial dependence on h, we can instead use the above pseudopolynomial algorithm to obtain an FPTAS in a fairly standard way as follows.

Step 1 Given $\varepsilon > 0$, let $\ell = \frac{\varepsilon h}{nm}$.

Step 2 Define $w'_e = \left\lfloor \frac{w_e}{\ell} \right\rfloor$, for each hyperedge $e \in E$.

Step 3 Run the dynamic programming algorithm on the instance specified by $G = (V,E)$ and valuations w'_e, and let $\mathbf{p}'$ be the returned price vector.

Step 4 Output the price vector $\tilde{\mathbf{p}}$ defined as $\tilde{p}_i = \ell \cdot p'_i$, for $i \in V$.

Theorem 7.1. *The above algorithm is an FPTAS, achieving profit at least* $(1-\varepsilon)$OPT *in time polynomial in n, m, and* $\frac{1}{\varepsilon}$.

Proof. In the following discussion, let $\text{Profit}_{w'}(\mathbf{p})$ denote the profit made by using the price vector $\mathbf{p}$ in the rounded instance specified by $G = (V,E)$ and valuations w'_e. In order to prove that the profit we obtain by using $\tilde{\mathbf{p}}$ in the original instance (given by $G = (V,E)$ and valuations w_e) is at least $(1-\varepsilon)$OPT, we first make some observations.

Let $\mathbf{p}$ be a price vector and let W be the set of winners under the pricing scheme $\mathbf{p}$ in the original instance. If $\mathbf{p}''$ is the pricing vector defined as $p''_i = \lfloor p_i/\ell \rfloor$ for $i \in V$, then $\text{Profit}_{w'}(\mathbf{p}'') \geq (1/\ell) \cdot \text{Profit}(\mathbf{p}) - nm$. To see why this is true, notice first that $W \subseteq W''$, where W'' is the set of winners under the pricing scheme $\mathbf{p}''$ in the rounded instance (specified by $G = (V,E)$ and the valuations w'_e). This follows from the fact that $\sum_{i \in e} p_i \leq w_e$ implies $\sum_{i \in e} p''_i = \sum_{i \in e} \lfloor p_i/\ell \rfloor \leq \lfloor w_e/\ell \rfloor = w'_e$. This implies

$$\text{Profit}_{w'}(\mathbf{p}'') = \sum_{e \in W''} \sum_{i \in e} p''_i \geq \sum_{e \in W} \sum_{i \in e} \left(\frac{p_i}{\ell} - 1 \right) \geq \frac{1}{\ell} \cdot \text{Profit}(\mathbf{p}) - nm,$$

as desired.

Let $\mathbf{p}'$ be a pricing vector and let W' be the set of winners under the pricing scheme $\mathbf{p}'$ in the rounded instance. If $\tilde{\mathbf{p}}$ is the pricing vector defined as $\tilde{p}_i = \ell \cdot p'_i$ for $i \in V$, then $\text{Profit}(\tilde{\mathbf{p}}) \geq \ell \cdot \text{Profit}_{w'}(\mathbf{p}')$. To see why this is true, notice first that $W' \subseteq W$, where W is the set of winners under $\tilde{\mathbf{p}}$ in the original instance. This follows from the fact that $\sum_{i \in e} p'_i \leq w'_e$ implies $\sum_{i \in e} \tilde{p}_i = \ell \cdot \sum_{i \in e} p'_i \leq \ell \cdot w'_e = l \lfloor w_e/\ell \rfloor \leq w_e$. This implies

$$\text{Profit}(\tilde{\mathbf{p}}) = \sum_{e \in W} \sum_{i \in e} \tilde{p}_i = \sum_{e \in W} \sum_{i \in e} \ell \cdot p'_i \geq \sum_{e \in W'} \sum_{i \in e} \ell \cdot p'_i = \ell \cdot \text{Profit}_{w'}(\mathbf{p}'),$$

as desired.

We are now ready to show that $\text{Profit}(\tilde{\mathbf{p}}) \geq (1-\varepsilon)$OPT. Let $\mathbf{p}^*$ and $\mathbf{p}'$ be the price vectors with the maximum profit in the original and rounded instances respectively, and let W^* and W be the corresponding set of winners. Let $\tilde{\mathbf{p}}$ be the price vector defined as $\tilde{p}_i = \ell \cdot p'_i$ for $i \in V$ and let $\mathbf{p}''$ is the pricing vector defined as $p''_i = \lfloor p^*_i/\ell \rfloor$ for $i \in V$. According to the previous observations we have $\text{Profit}(\tilde{\mathbf{p}}) \geq \ell \cdot \text{Profit}_{w'}(\mathbf{p}')$. Since $\mathbf{p}'$ is the price vector with the maximum profit in the rounded instance we have $\text{Profit}_{w'}(\mathbf{p}') \geq \text{Profit}_{w'}(\mathbf{p}'')$. Combining these together with the fact that $\text{Profit}_{w'}(\mathbf{p}'') \geq$

$(1/\ell)\cdot \text{Profit}(\mathbf{p}^*) - nm$, we get $\text{Profit}(\tilde{\mathbf{p}}) \geq \text{Profit}(\mathbf{p}^*) - \ell nm$, which implies $\text{Profit}(\tilde{\mathbf{p}}) \geq (1-\varepsilon)\text{OPT}$, as desired.

Since $w'_e \leq nm/\varepsilon$ for all $e \in E$, we also have that our procedure runs in polynomial time in n, m, and $1/\varepsilon$, thus being an FPTAS for the hierarchy case. □

8 Online Pricing

As mentioned in Section 1, results of Balcan et al. [3] can be used to convert our algorithms into incentive-compatible mechanisms in the offline "batch" setting (i. e., a sealed-bid auction). In this section we consider a natural, more demanding *online* setting in which customers arrive one at a time, and we must assign prices to the items for customer t based only on information about customers $1,\ldots,t-1$.

8.1 The model

We assume customers arrive one at a time. Each customer will be shown a set of item prices, and will then decide whether to purchase or not at those prices. We assume customers cannot return and cannot control their time of arrival, so any take-it-or-leave-it set of prices for customer t based only on information received from customers $1,\ldots,t-1$ is incentive-compatible. In addition, we assume an *oblivious adversary* model: that is, our objective is to achieve good expected performance for any sequence of customers, but this sequence cannot depend on the outcome of any probabilistic choices made by our algorithm. As before, we use m to denote the total number of customers.

We consider two information models. In the *full information* model, we assume that after the tth customer departs, we learn his desired set e_t and valuation v_t. In the more difficult *posted-price* model, we assume we only find out whether and what the customer purchased but not his actual valuations. That is, if he purchases a subset at the current prices, we do not know if he still would have purchased at higher prices, and if he does not purchase at the current prices, we do not know if (or what) he would have purchased at lower prices. In both models, we will be interested in algorithms that perform well compared to the best fixed setting of prices for the entire sequence. Thus, we are comparing to the same notion of OPT as in the offline case.

8.2 The Online Graph and k-Hypergraph Pricing Problems

Our 4-approximation for graph vertex-pricing, and our $O(k)$-approximation for k-hypergraph vertex pricing, can be directly adapted to the online setting by using the results of [5, 6] for the online digital-good auction. In particular, for the full-information setting, Blum and Hartline [5] show the following:

Theorem 8.1 ([5]). *Consider the online pricing problem for a single item ($n = 1$) in the full-information setting. There exists an online algorithm such that for any given $\varepsilon > 0$, its expected profit on any input sequence of maximum valuation h is at least*

$$(1-\varepsilon)\text{OPT} - O\left(\frac{nh}{\varepsilon}\log\frac{1}{\varepsilon}\right).$$

We can now use this to prove the following result.

Theorem 8.2. *Consider the online hypergraph vertex pricing problem in the full-information setting. There exists an online algorithm such that for any given $\varepsilon > 0$, its expected profit on any input sequence is at least*

$$(1-\varepsilon)\text{ALG} - O\left(\frac{nh}{\varepsilon}\log\frac{1}{\varepsilon}\right),$$

where ALG *is is the profit of the* offline *approximation algorithm in Section 4 on that input.*

Proof. Note that our algorithms in Sections 3 and 4 begin by selecting a subset V_L of items to have non-zero prices, and then achieve their approximation guarantees considering only profit made from customers who want exactly one item in V_L. Thus, we can view these algorithms as effectively performing $|V_L|$ separate digital-good auctions, ignoring customers who want zero, or more than one, item from V_L. In particular, to apply these algorithms to the full-information online setting, we begin by randomly choosing the set V_L as described in the algorithms, setting prices for items in $V - V_L$ to 0. We then instantiate a separate copy of the online digital-good auction from [5] for each item $i \in V_L$. When a customer arrives, if the customer wants exactly one item i from V_L, then his valuation is given to the associated online auction algorithm. Let OPT_i denote the optimal profit achievable using a fixed price for item i from customers whose bundles contain item i but no other item in V_L. By Theorem 8.1, the expected profit of the online auction for item i will therefore be at least $(1-\varepsilon)\text{OPT}_i - O((h/\varepsilon)\log 1/\varepsilon)$. Thus, overall, we achieve profit at least

$$(1-\varepsilon)\sum_{i\in V_L}\text{OPT}_i - O\left(\frac{nh}{\varepsilon}\log\frac{1}{\varepsilon}\right),$$

where $\sum_{i\in V_L}\text{OPT}_i$ is the profit of the *offline* approximation algorithm. Note that we need the assumption of an oblivious adversary for the approximation ratios proved in Sections 3 and 4 to apply. □

In particular, so long as the offline algorithm's profit is $\Omega((nh/\varepsilon^2)\log 1/\varepsilon)$, we lose only a $(1+O(\varepsilon))$ factor in the conversion to the online setting. In the posted-price setting, we can obtain a similar result: we only need to apply the associated posted-price algorithms of [5, 6]. The only tricky issue is that a customer who chooses not to buy anything must be fed in as a non-buyer to *all* of the online algorithms, in order to ensure that the sequence of customers fed into algorithm i is a superset of the true customers for that item (so that the value of OPT for the sequence fed to the algorithm is at least as large as the true OPT for that item). In addition, the algorithms for the posted-price scenario require that the upper-bound h on the maximum valuation be known in advance.

8.3 The Online Highway Problem

For the highway problem, we cannot decompose our solution into a collection of independent digital-good auctions, so the reduction in Section 8.2 does not apply. However, we *can* convert to the online setting by placing this problem in the framework of *online geometric optimization* studied by Kalai and Vempala [18] as extended to the case of approximation algorithms by Kakade et al. [17]. In particular, [17] gives a method to convert any efficient approximation algorithm for offline optimization into an

efficient algorithm for *online* optimization with asymptotically the same approximation ratio, for any problem of the following type:

1. There is a set $S \subseteq \mathbb{R}^d$ of *feasible points*. At each time step t we must pick some point $\mathbf{p}^t \in S$; we are then given an objective function $\mathbf{v}^t \in \mathbb{R}^d$, and we obtain profit $\mathbf{p}^t \cdot \mathbf{v}^t$.

2. Our goal is to perform nearly as well as the best point $\mathbf{p} \in S$ in hindsight. That is, we want $\sum_t \mathbf{p}^t \cdot \mathbf{v}^t$ to be nearly as large as $\max_{\mathbf{p} \in S} \sum_t \mathbf{p} \cdot \mathbf{v}^t$.

3. We have an efficient α-approximation algorithm for the *offline* optimization problem: given objective function $\mathbf{v} \in \mathbb{R}^d$, find the point $\mathbf{p} \in S$ that maximizes $\mathbf{p} \cdot \mathbf{v}$.

Kalai and Vempala [18] (for the case of exact offline algorithms) and Kakade et al. [17] (for the case of approximation algorithms) give a procedure for choosing points $\mathbf{p}^t$ online for any problem of the above type such that the total profit obtained, $\sum_t \mathbf{p}^t \cdot \mathbf{v}^t$, is within a $(1-\varepsilon)/\alpha$ factor of the profit of the best $\mathbf{p} \in S$ in hindsight, minus an additive term that is polynomial in the diameter of S and the maximum L_1 magnitude of any $\mathbf{v}^t$.

We can place the highway problem into the above setting as follows. First, for simplicity of exposition, let us assume all bidders have integral valuations between 0 and h, and that we are willing to have an algorithm that runs in time polynomial in h as opposed to $\log h$ (the general case can be handled by rounding and scaling as in Section 7). Now, let S be the set of all possible item prices represented in the following way. Given a pricing of the n items, let q_{ij} denote the total cost of the bundle $[i, j]$. We represent q_{ij} as a vector of length h consisting of $q_{ij} - 1$ zeros followed by $h - q_{ij} + 1$ entries at value q_{ij}. To represent the entire pricing of the n items, we just concatenate these n^2 vectors together to create a point in $\mathbb{R}^d$ for $d = n^2 h$. A bidder who desires bundle $[i, j]$ at value w is represented as a vector of all zero entries except for a 1 in the w-th coordinate of the block corresponding to q_{ij}. By design, the dot product of this vector with a vector $\mathbf{p} \in S$ is exactly the profit that would be obtained from this bidder by the item-pricing corresponding to $\mathbf{p}$. Finally, we can use the optimization algorithm from Section 6 or from [15] as our offline optimization oracle. Since the diameter of S is polynomial in n and h, and the maximum L_1 magnitude of any $\mathbf{v}^t$ is at most 1, the total profit obtained will be within a $1 + \varepsilon$ factor of the approximation ratio guaranteed by the optimization algorithm, minus an additive loss term which is polynomial in n and h.

Note that for the posted-price version, we just need to apply known extensions of the Kalai-Vempala algorithm to the bandit setting [1, 21, 9, 17] in which only the profit $\mathbf{p}^t \cdot \mathbf{v}^t$ and not the actual vector $\mathbf{v}^t$ is revealed to the algorithm.

9 Conclusions

We present approximation and online algorithms for a number of problems of pricing items to consumers so as to maximize seller's revenue in an unlimited supply setting. We achieve an $O(k)$-approximation algorithm for the case of single-minded bidders where each consumer wants at most k items, an $O(\log n)$ approximation for the highway problem from [15], and a constant factor approximation to the highway

problem when all bidders want approximately (up to a constant factor) the same number of items. We also show how some of our approximation algorithms can be adapted to the more demanding online setting in which customers arrive one at a time, in both the full-information and posted-price settings.

9.1 Subsequent Work

Following the initial publication of our work, Krauthgamer, Mehta and Rudra [19] have provided improved approximation guarantees for the graph vertex pricing problem in the special case when the range of bidders' valuations is small. Briest and Krysta observe [8] that our algorithms in Sections 3 and 4 can be adapted to obtain similar guarantees for the case of unlimited supply *unit-demand* combinatorial auctions. Balcan et al. [2] further study how pricing certain items below their marginal cost can lead to an improvement in overall profit. In particular, they develop the *rebate* and *coupon* models for analyzing this issue and examine "profitability gaps" (to what extent can pricing below cost help to improve profit) as well as algorithms for pricing. Elbassioni et al. [11] give a quasi-polynomial time approximation scheme for the highway problem.

9.2 Open Questions

There are several natural open problems left by this work.

First, can one improve on the factor of 4 for the graph vertex-pricing problem? Any method able to reduce the factor of 2 for the bipartite case would immediately result in an improved bound. Alternatively, perhaps the reduction to the bipartite case can be improved. Second, can one achieve $o(k)$ for the hypergraph vertex-pricing problem? Even for the general case there remains a gap between the known upper bounds and the lower bound of [10].

Finally, an intriguing question related to this work is: what kind of approximation guarantees are achievable if one allows the seller to price some items below cost (i. e., to have "loss leaders")? Some progress on this direction has been made in [2]; however, we still do not know if there exists a constant-factor approximation for the graph vertex pricing problem (even the bipartite case) when negative profit margins on some items are allowed.

Acknowledgements

We would like to thank Jason Hartline for posing the graph vertex-pricing problem to us and for helpful discussions. We would also like to thank the referees for their helpful comments and suggestions.

References

[1] * B. AWERBUCH AND R. KLEINBERG: Adaptive routing with end-to-end feedback: Distributed learning and geometric approaches. In *Proc. 36th STOC*, pp. 45–53. ACM Press, 2004. [STOC:1007352.1007367]. 8.3

[2] * M.-F. BALCAN, A. BLUM, H. CHAN, AND M.T. HAJIAGHAYI: A theory of loss-leaders: Making money by pricing below cost. In *Proc. 3rd Intern. Workshop on Internet and Network Economics*, Lecture Notes in Computer Science. Springer, 2007. 9.1, 9.2

[3] * M.-F. BALCAN, A. BLUM, J. HARTLINE, AND Y. MANSOUR: Mechanism design via machine learning. In *Proc. 46th FOCS*, pp. 605–614. IEEE Computer Soc., 2005. [FOCS:10.1109/SFCS.2005.50]. 1, 8

[4] * M.-F. BALCAN, A. BLUM, AND Y. MANSOUR: Single price mechanisms for revenue maximization in unlimited supply combinatorial auctions. Technical Report CMU-CS-07-111, Carnegie Mellon University, 2007. 3

[5] * A. BLUM AND J. HARTLINE: Near-optimal online auctions. In *Proc. 16th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'05)*, pp. 1156–1163. SIAM, 2005. [SODA:1070432.1070597]. 1, 8.2, 8.1, 8.2

[6] * A. BLUM, V. KUMAR, A. RUDRA, AND F. WU: Online learning in online auctions. In *Proc. 14th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'03)*, pp. 202–204. SIAM, 2003. [SODA:644108.644143]. 1, 8.2, 8.2

[7] * P. BRIEST AND P. KRYSTA: Single-minded unlimited supply pricing on sparse instances. In *Proc. 17th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'06)*, pp. 1093–1102. SIAM, 2006. [SODA:1109557.1109678]. 1, 4, 6, 5

[8] * P. BRIEST AND P. KRYSTA: Buying cheap is expensive: Hardness of non-parametric multi-product pricing. In *Proc. 18th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'07)*, pp. 716–725. SIAM, 2007. [SODA:1283383.1283460]. 9.1

[9] * V. DANI AND T. P. HAYES: Robbing the bandit: Less regret in online geometric optimization against an adaptive adversary. In *Proc. 17th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'06)*, pp. 937–943. SIAM, 2006. [SODA:1109557.1109660]. 8.3

[10] * E. D. DEMAINE, U. FEIGE, M. HAJIAGHAYI, AND M. R. SALAVATIPOUR: Combination can be hard: Approximability of the unique coverage problem. In *Proc. 17th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'06)*, pp. 162–171. SIAM, 2006. [SODA:1109557.1109577]. 1, 9.2

[11] * K. ELBASSIONI, R. SITTERS, AND Y. ZHANG: A quasi-PTAS for profit-maximizing pricing on line graphs. In *Proc. 15th Ann. European Symp. on Algorithms (ESA'07)*, Lecture Notes in Computer Science. Springer, 2007. 9.1

[12] * G. EVEN, O. GOLDREICH, M. LUBY, N. NISAN, AND B. VELICKOVIC: Efficient approximation of product distributions. *Random Structures and Algorithms*, 13(1):1–16, 1998. [Wiley:10.1002/(SICI)1098-2418(199808)13:1¡1::AID-RSA1¿3.0.CO;2-W]. 4

[13] * A. GOLDBERG, J. HARTLINE, A. KARLIN, M. SAKS, AND A. WRIGHT: Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006. [Elsevier:10.1016/j.geb.2006.02.003]. 1

[14] * A. GOLDBERG, J. HARTLINE, AND A. WRIGHT: Competitive auctions and digital goods. In *Proc. 12th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'01)*, pp. 735–744. SIAM, 2001. [SODA:365411.365768]. 1

[15] * V. GURUSWAMI, J. HARTLINE, A. KARLIN, D. KEMPE, C. KENYON, AND F. MCSHERRY: On profit-maximizing envy-free pricing. In *Proc. 16th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'05)*, pp. 1164–1173. SIAM, 2005. [SODA:1070432.1070598]. 1, 2, 2, 4, 5, 6, 6, 7, 8.3, 9

[16] * J. HARTLINE AND V. KOLTUN: Near-optimal pricing in near-linear time. In *9th Workshop on Algorithms and Data Structures (WADS'05)*, number 3608 in Lecture Notes in Computer Science, pp. 422–431. Springer, 2005. [WADS:al2ke8gd23a44atp]. 1, 2

[17] * S. KAKADE, A. KALAI, AND K. LIGETT: Playing games with approximation algorithms. In *Proc. 39th STOC*, pp. 546–555. ACM Press, 2007. [STOC:1250790.1250870]. 8.3, 8.3

[18] * A. KALAI AND S. VEMPALA: Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005. An earlier version appears in COLT 2003. [JCSS:10.1016/j.jcss.2004.10.016, Springer:fbx1n68j5ubyv706]. 1, 8.3, 8.3

[19] * R. KRAUTHGAMER, A. MEHTA, AND A. RUDRA: Pricing commodities, or how to sell when buyers have restricted valuations. In *5th Workshop on Approximation and Online Algorithms*. Springer, 2007. 9.1

[20] * M. LUBY AND A. WIGDERSON: Pairwise independence and derandomization. Technical Report CSD-95-880, U.C. Berkeley, 1995. 3

[21] * H. B. MCMAHAN AND A. BLUM: Online geometric optimization in the bandit setting against an adaptive adversary. In *Proc. 17th Conf. on Computational Learning Theory (COLT'04)*, number 3120 in Lecture Notes in Computer Science, pp. 109–123. Springer, 2004. [Springer:38p3f4h61cgy7gg7]. 8.3

[22] * R. MOTWANI AND P. RAGHAVAN: *Randomized Algorithms*. Cambridge University Press, 1995. 3

AUTHORS

Maria-Florina Balcan [About the author]
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
ninamf@cs.cmu.edu
www.cs.cmu.edu/~ninamf/

Avrim Blum [About the author]
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
avrim@cs.cmu.edu
www.cs.cmu.edu/~avrim/

ABOUT THE AUTHORS

MARIA-FLORINA BALCAN is a Ph. D. candidate at Carnegie Mellon University under the supervision of Avrim Blum. Everyone who knows her calls her "Nina," a tradition originally started by her brother Marius when he was too young to appreciate longer names. Nina received B. S. and M. S. degrees from the University of Bucharest, Faculty of Mathematics and Informatics, Romania. Her primary research interests are Computational and Statistical Machine Learning, computational aspects of Economics and Game Theory, and Algorithms. She owes much of her interest in Computer Science to Professors Luminita State, her undergraduate mentor, and Florentina Hristea, her role model at the time, who offered Nina the opportunity of her first textbook co-authorship.

AVRIM BLUM grew up in Berkeley, California, and then went to MIT for undergraduate and graduate school. He received his Ph. D. in Computer Science under supervision of Ron Rivest, and is now Professor of Computer Science at Carnegie Mellon University. His research interests include Approximation Algorithms, Algorithmic Game Theory, and Machine Learning Theory. He has two children, Alex and Aaron, who may or may not go into the family business.