A Quasi-PTAS for Profit-Maximizing Pricing on Line Graphs

Khaled Elbassioni[1], René Sitters[2], and Yan Zhang[3]

[1] Max-Planck-Institut für Informatik, Saarbrücken, Germany;
elbassio@mpi-inf.mpg.de

[2] Department of Mathematics and Computer Science, Eindhoven University of Technology, the Netherlands; r.sitters@tue.nl

[3] Hong Kong University of Science and Technology, Hong Kong, China;
cszy@cs.ust.hk

Abstract. We consider the problem of pricing items so as to maximize the profit made from selling these items. An instance is given by a set E of n items and a set of m clients, where each client is specified by one subset of E (the bundle of items she/he wants to buy), and a budget (valuation), which is the maximum price she/he is willing to pay for that subset. We restrict our attention to the model where the subsets can be arranged such that they form intervals of a line graph. Assuming an unlimited supply of any item, this problem is known as *the highway problem* and so far only an $O(\log n)$-approximation algorithm is known. We show that a PTAS is likely to exist by presenting a quasi-polynomial time approximation scheme. We also combine our ideas with a recently developed quasi-PTAS for the unsplittable flow problem on line graphs to extend this approximation scheme to the limited supply version of the pricing problem.

1 Introduction

Suppose you have a set of items to sell and you have complete information of your possible clients, i.e., you know what each client wants to buy and how much she/he is willing to pay for this. How do you set your prices such that your total profit is maximized? If you price the items cheap, then many people will buy from you but, of course, at a low price. If you price too high, then only few people are willing to pay the price. This kind of pricing problems appear more and more in pricing mechanisms over the Internet. The seller collects statistical data and adjusts the prices online depending on the set of potential buyers. Recently, a number of papers appeared on the computational complexity of these problems [1, 2, 4–9].

Pricing problems of this type are generally very hard to solve because of their non-linear nature: each client will either buy the whole bundle of items she/he is interested in, or not buy at all; there is no solution in between (i.e. the client is not interested in buying a strict subset of the bundle). We restrict our attention to the so called *single minded* clients, i.e., each client is only interested in a single subset of the items. If there is an *unlimited* supply of any item, then the client will buy her/his preferred subset if and only if the total price of the items in the subset is no more than her/his budget (such a pricing is said to be *envy-free* since, given the pricing, no client would prefer to be assigned a different bundle). Guruswami et al. [8] show that this problem is already APX-hard if all budgets are equal to one and each client wants to buy a single pair of items. They also give an $O(\log n+\log m)$-approximation algorithm, where n and m are, respectively, the numbers of items and clients. Efficient (approximation) algorithms are known for several variants in which constraints are placed on the numbers in the input. For example, Hartline and Koltun [9] give a PTAS for the case when the number of items is constant, and Balcan and Blum [2] give an $O(k)$-approximation algorithm under the restriction that the bundle of any client contains at most k items.

In this paper, we do not impose any restriction on the sizes, prices, capacities or budgets but only restrict the system of subsets (bundles). More precisely, we assume that the subsets can be arranged such that they form intervals of a line graph. This special case is known as the *highway problem*, and recently Balcan and Blum [2] gave an $O(\log n)$-approximation algorithm improving over the $O(\log n+\log m)$ algorithm for the general problem, and an $O(1)$-approximation when all intervals are of almost the same length. NP-hardness of the highway problem was shown by Briest and Krysta [4], and Guruswami et al. [8] gave a polynomial-time exact algorithm for the case when all budgets are bounded by a constant, and a pseudo-polynomial-time algorithm for the case when all intervals are of constant length and all budgets are integral.

We consider both the *uncapacitated* and *capacitated* variants of the problem. In the capacitated variant, there is only a limited supply of any item. Hence, we have to set prices as well as to find a subset of the clients such that the supply constraints are satisfied. This problem has some similarities with the *unsplittable flow problem* on line graphs for which a quasi-PTAS was developed only recently by Bansal et al [3]. To solve the limited supply case, we consider the following generalization which may be of independent interest: Assume that there are capacities on the edges. Assume further that each buyer I wants to route a demand of $\rho(I)$ along the interval I, and has budget $B(I)$. If a buyer can route her/his demand, she/he will purchase the path, provided its price is within her/his budget. Then the question is how to price the edges such that the total routed demand on each edge, from buyers that can afford to route their demands, does not exceed the capacity of the edge and the total profit is maximized[4]. The limited supply version of the highway problem corresponds to the case with unit demands. Just as the quasi-PTAS for the unsplittable flow problem [3], the technique used here is standard for deriving polynomial time approximation schemes: first rounding, then restricting the search space further, and finally applying dynamic programming. However, for both problems the

[4] Note, however, that the resulting pricing may not be envy-free in this case, since some of the clients who can afford to purchase their paths, may not be able to do so because there might not be enough capacity on the edges.

second step is non-trivial. In the dynamic programming we do not enumerate directly over all possible (rounded) price vectors but, instead, over systems of linear inequalities on the price vectors. For each consistent set of inequalities we can easily derive the value of the solution with a feasible price vector.

2 The setting

Let $V = \{0, 1, \ldots, n\}$ and $E = \{e_1, \ldots, e_n\}$, with $e_i = \{i-1, i\}$, for $i = 1, \ldots, n$. We assume that we are given a (multi)set of intervals (bundles) $\mathcal{I} = \{I_1, \ldots, I_m\}$, defined on the set of edges (items) E, where $I_j = [s_j, t_j] \stackrel{\text{def}}{=} \{\{s_j, s_j+1\}, \{s_j+1, s_j+2\} \ldots, \{t_j-1, t_j\}\} \subseteq E$. For $I \in \mathcal{I}$, we denote by $B(I) \in \mathbb{R}_+$ the *budget* of interval I. In the *highway problem*, denoted henceforth by HP, the objective is to assign a price $p(e) \in \mathbb{R}_+$ for each edge $e \in E$, and to find a subset $\mathcal{J} \subset \mathcal{I}$, so as to maximize

$$\sum_{I \in \mathcal{J}} p(I) \tag{1}$$

subject to the budget constraints

$$p(I) \leq B(I), \quad \text{for all } I \in \mathcal{J}, \tag{2}$$

where, for $I \in \mathcal{I}$, $p(I) = \sum_{e \in I} p(e)$. When the number of copies of each edge $e \in E$, available for purchase, is limited by a given number, the problem will be called the *capacitated highway problem*. As mentioned in the introduction, this version of the problem can be cast as an instance of the following more general problem, which we call the *unsplittable flow pricing problem* and denote by UFP: Let $c(e) \in \mathbb{R}_+$, be the *available supply* or *capacity of edge* e, and $\rho(I)$ be the demand of interval $I \in \mathcal{I}$. The requirement is to assign a price $p(e) \in \mathbb{R}_+$ for each edge $e \in E$, and to find a subset $\mathcal{J} \subset \mathcal{I}$, so as to maximize (1) subject to (2) and the capacity constraints

$$\sum_{I \in \mathcal{J} : e \in I} \rho(I) \leq c(e), \quad \text{for all } e \in E. \tag{3}$$

In the following sections, we denote by $p^* : E \mapsto \mathbb{R}_+$ the optimal set of prices and by OPT $\subseteq \mathcal{I}$ the set of intervals purchased in the optimum solution. It is easy to see that specifying any one of these two sets completely determines the other, and thus any of the two is enough to completely describe the optimal solution. For a subset of intervals $\mathcal{I}' \subseteq \mathcal{I}$, and a price function $p : E \mapsto \mathbb{R}_+$, we denote by $p(\mathcal{I}') = \sum_{I \in \mathcal{I}'} p(I)$ the total price of intervals in $\mathcal{I}'$. Similar notation will be also used for the function $\rho(\cdot)$.

For an edge $e = \{u-1, u\} \in E$ and a (multi)set of intervals $\mathcal{I}$ on E, denote respectively by $\mathcal{I}_L(e)$, $\mathcal{I}_R(e)$, and $\mathcal{I}[e]$ the subsets of intervals of $\mathcal{I}$ that lie to the left of e, lie to the right of e, and span e, that is

$$\begin{aligned} \mathcal{I}_L(e) &= \{[s,t] \in \mathcal{I} : \; t \leq u-1\}, \\ \mathcal{I}_R(e) &= \{[s,t] \in \mathcal{I} : \; s \geq u\}, \\ \mathcal{I}[e] &= \{[s,t] \in \mathcal{I} : \; s \leq u-1 < u \leq t\}. \end{aligned}$$

Denote by $E_L(e)$ and $E_R(e)$ the sets of edges that lie to the left and right of edge $e \in E$, respectively:

$E_L(e) = \{\{x-1, x\} \in E : \ x \le u-1\}$ and $E_R(e) = \{\{x-1, x\} \in \mathcal{I} : \ x-1 \ge u\}$.

3 Rounding the instance

Let $\epsilon > 0$ be a given constant. We may assume without loss of generality that for every vertex $u \in V$, there is an interval $I \in \mathcal{I}$ beginning at that point (otherwise we can merge indistinguishable edges). In particular, $n \le 2m$ can be assumed. Denote by $B_{max} = \max\{B(I) : I \in \mathcal{I}\}$ the maximum budget in the instance. Clearly, $B_{max} \le p^*(\text{Opt})$ since we can assign a price of $B_{max}/|I_{max}|$ for every edge of the interval of maximum budget I_{max}, where $|I_{max}|$ is the length of I_{max}.

Consider an optimal price function p^*. We obtain a new price function p' by rounding *down* $p^*(e)$, for each $e \in E$, to the closest multiple of $\epsilon B_{max}/(nm)$. Then the total reduction in profit is $p^*(\text{Opt}) - p'(\text{Opt}) \le \epsilon B_{max} \le \epsilon p^*(\text{Opt})$, i.e. $p'(\text{Opt}) \ge (1-\epsilon)p^*(\text{Opt})$. Now we scale all budgets and prices in p' by $nm(1+\epsilon)/(B_{max}\epsilon)$ (this does not change the optimal solution), and assume, at the loss of factor of $(1-\epsilon)$ of the optimal, that all prices belong to the set $\{0, 1+\epsilon, 2(1+\epsilon), \dots, P(1+\epsilon)\}$, where $P \stackrel{\text{def}}{=} \frac{mn}{\epsilon}$. Note that we still have $p'(I) \le B(I)$ for all $I \in \text{Opt}$.

By dividing the prices further by $(1+\epsilon)$, we obtain a set of prices $\tilde{p} : E \mapsto \mathbb{R}_+$, for which can also assume that every interval $I \in \text{Opt}$ has $\tilde{p}(I) = p'(I)/(1+\epsilon) \le B(I)/(1+\epsilon)$. The total profit of such solution is $\tilde{p}(\text{Opt}) \ge \frac{1-\epsilon}{1+\epsilon}\text{Opt} \ge (1-2\epsilon)\text{Opt}$. We summarize the above facts in the following proposition.

Proposition 1. *Let p^* be an optimal solution for a given instance of* Hp*, and $\epsilon > 0$ be a given constant. Then there exists a pricing $\tilde{p} : E \mapsto \mathbb{R}_+$ for which*
(i) $\tilde{p}(e) \in \{0, 1, \dots, P\}$, for every $e \in E$, where $P = nm/\epsilon$,
(ii) $\tilde{p}(I) \le \frac{B(I)}{1+\epsilon}$, for every $I \in$ Opt, and
(iii) $\tilde{p}(\text{Opt}) \ge (1-2\epsilon)p^*(\text{Opt})$.

We shall call any pricing $\tilde{p}$, satisfying the conditions of Proposition 1, an ϵ-*optimal pricing.*

4 The highway problem

In this case, we assume that $c(e) = \infty$ for all $e \in E$. Given a price function $p : E \mapsto \mathbb{R}_+$ and an edge $e^* = \{u^*-1, u^*\} \in E$, the *accumulative price* at any edge $e = \{x-1, x\} \in E$ with respect to e^* is defined as $p([x-1, u^*])$ if e is to the left of e^* (i.e. if $x \le u^*-1$) and $p([u^*, x])$ if e is to the right of e^* (i.e. if $x-1 \ge u^*$). Obviously, starting from e^*, these accumulative prices form monotonically increasing functions to the left and right of e^*.

In this section we prove the following statement.

Theorem 1. *There is a quasi-polynomial time approximation scheme for the highway problem.*

Our QPTAS is based essentially on the same divide and conquer strategy used in [3]. It starts by picking an edge in the middle and guesses the points at which the ϵ-optimal accumulative prices increase by factors of $(1+\epsilon)$ relative to that middle edge. Having guessed such "increment points", the algorithm picks a *superset* of the optimal set of intervals containing the middle edge, then recurses independently on the two subproblems to the left and right of the middle edge. In the following, we fix $k = \lceil \log(nP)/\log(1+\epsilon)\rceil + 1$.

Definition 1. (ϵ-Relative pricings) *Let $e^* = [u^*-1, u^*] \in E$ be a given edge of E, and $1 \le k_1, k_2 \le k$ be given integers. A selection of k_1+k_2+2 points $u_1, \ldots, u_{k_1}, u_{k_1+1}, u'_1, \ldots, u'_{k_2}, u'_{k_2+1} \in V$, and k_1+k_2+2 values $p_1, \ldots, p_{k_1}, p_{k_1+1}, p'_1, \ldots, p'_{k_2}, p'_{k_2+1} \in \{0, 1, \ldots, P\}$, such that*

1. $u_{k_1+1} \le u_{k_1} < u_{k_1-1} < \cdots < u_1 \le u^* < u'_1 < u'_2 < \cdots < u'_{k_2} \le u_{k_2+1}$,
2. $p_j \ge (1+\epsilon)p_{j-1}$, *for* $j = 2, \ldots, k_1$, *and* $p'_j \ge (1+\epsilon)p'_{j-1}$, *for* $j = 2, \ldots, k_2$,
3. $p_{k_1+1} \ge p_{k_1}$ *and* $p'_{k_2+1} \ge p'_{k_2}$.

is said to be an ϵ-relative pricing *of E w.r.t. e^*, denoted by $(u^*, u_1, \ldots, u_{k_1+1}, u'_1, \ldots, u'_{k_2+1}, p_1, \ldots, p_{k_1+1}, p'_2, \ldots, p'_{k_2+1})$.*

See Figure 1 for an example. The total number of possible ϵ-relative pricings with respect to a given edge $e^* \in E$ is at most

$$L = [n(1+P)]^{2k+2} = \left(n + \frac{n^2 m}{\epsilon}\right)^{2\lceil \log \frac{n^2 m/\epsilon}{\log(1+\epsilon)}\rceil + 4}, \tag{4}$$

which is $m^{O(\log(m))}$ for every fixed $\epsilon > 0$.

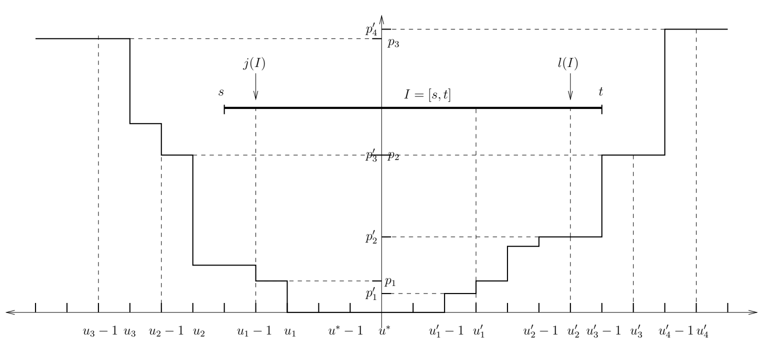

Fig. 1. ϵ-relative pricings w.r.t. to $[u^*-1, u^*]$: To the left is the plot of $p([u-1, u^*])$ for $u < u^*$, and to the right is the one for $p([u^*, u])$ for for $u \ge u^*$. The indices $j(I)$ and $l(I)$ for a given interval $I = [s,t]$ are also shown.

Let $R=(u^*,u_1,\ldots,u_{k_1+1},u'_1,\ldots,u'_{k_2+1},p_1,\ldots,p_{k_1+1},p'_2,\ldots,p'_{k_2+1})$ be an ϵ-relative pricing w.r.t. an edge e^*. Given an interval $I=[s,t]\in\mathcal{I}$, with $e^*\in I$, we associate a value $v(I,R)$ to I, defined with respect to R as follows. Let $j(I),l(I)$ be respectively the smallest and largest indices such that $e_{i_{j(I)}},e'_{i_{l(I)}}\in I$, i.e. $j(I)=\min\{i:\ u_i-1\geq s\}$ and $l(I)=\max\{i:\ u'_i\leq t\}$ (see Figure 1). Then, for $I\in\mathcal{I}$, define

$$v(I,R)=p_{j(I)}+p'_{l(I)}.$$

(If either $j(I)$ or $l(I)$ does not exist, the corresponding value $p_{j(I)}$ or $p'_{l(I)}$ is set to 0.) For a subset of intervals $\mathcal{I}'\subseteq\mathcal{I}$, we define as usual, $v(\mathcal{I}',R)=\sum_{I\in\mathcal{I}'}v(I,R)$.

Definition 2. (Consistent pricings) *Let* $R=(u^*,u_1,\ldots,u_{k_1+1},u'_1,\ldots,u'_{k_2+1},$ $p_1,\ldots,p_{k_1+1},\ p'_2,\ldots,p'_{k_2+1})$ *be an* ϵ*-relative pricing of* E *w.r.t. an edge* $e^*\in E$, *and* $p:E\mapsto\mathbb{R}_+$ *be a price function. We say that* p *and* R *are* consistent *if*

(C1) $p([u_j-1,u^*])=p_j$ *for* $j=1,\ldots,k_1+1$, *and* $p([u^*,u'_j])=p'_j$ *for* $j=1,\ldots,k_2+1$,
(C2) $p([u_j,u^*])\leq(1+\epsilon)p_{j-1}$ *for* $j=2,\ldots,k_1+1$, *and* $p([u^*,u'_{j-1}])\leq(1+\epsilon)p'_{j-1}$ *for* $j=2,\ldots,k_2+1$.

It follows that for any ϵ-relative pricing R w.r.t. an edge $e^*\in E$ and any $q:E\mapsto\mathbb{R}_+$ with which R is consistent, we have

$$v(I,R)\leq q(I)\leq(1+\epsilon)v(I,R)\text{ for all }I\in\mathcal{I}[e^*].\qquad(5)$$

This property will be used crucially in our analysis.

Lemma 1. *Let* $p:E\mapsto\mathbb{R}_+$ *be a pricing for a given instance of* HP *and* $e^*=[u^*-1,u^*]$ *be an arbitrary edge. Then there exists an* ϵ*-relative pricing of* E *w.r.t.* e^**, that is consistent with* p.

Proof. Assume that $E=\{\{0,1\},\{1,2\},\ldots,\{n-1,n\}\}$. We define a selection as follows (see Figure 1). Let $u_1=\max\{u\leq u^*:\ p([u-1,u])>0\}$, $u'_1=\min\{u>u^*:\ p([u-1,u])>0\}$, and set $p_1=p([u_1-1,u_1])$ and $p'_1=p([u'_1-1,u'_1])$. For $j=2,3,\ldots$, let $u_j=\max\{u\leq u_{j-1}:\ p([u-1,u^*])>(1+\epsilon)p_{j-1}\}$, and set $p_j=p([u_j-1,u^*])$. The highest index j for which this iteration can be done will the value of k_1. Similarly, we define k_2, and for $j=2,\ldots,k_2$, let $u'_j=\min\{u>u'_{j-1}:\ p([u^*,u])>(1+\epsilon)p'_{j-1}\}$ and set $p'_j=p([u^*,u'_j])$. Since $nP\geq p_{k_1}\geq(1+\epsilon)^{k_1-1}p_1\geq(1+\epsilon)^{k_1-1}$, we get $k_1\leq k$. Similarly, $k_2\leq k$. Finally, we let $u_{k_1+1}=1$, $u_{k_2+1}=n$, $p_{k_1+1}=p([0,u^*])$ and $p'_{k_2+1}=p([u^*,n])$. □

With every ϵ-relative pricing R, we can associate a system of linear inequalities, denoted by $S(R)$, on a set of E variables $\{p(e):\ e\in E\}$, consisting of the constraints $(C1)$ and $(C2)$ together with the non-negativity constraints $p(e)\geq0$. For two systems of inequalities S_1,S_2, we denote by $S_1\wedge S_2$ the system obtained by combining their inequalities.

The algorithm is shown in Figure 2. It is initially called with an empty $\mathcal{S}$. The procedure iterates over all ϵ-relative pricings R, consistent with $\mathcal{S}$, w.r.t.

to the middle edge e^*, then recurses on the subsets of intervals to the left and right of e^*. In line 4, we always insure that $S(R) \wedge \mathcal{S}$ is defined on variables $\{p(e) : e \in E_L(e^*)\}$. Similarly, in line 5, we always insure that $S(R) \wedge \mathcal{S}$ is defined on variables $\{p(e) : e \in E_R(e^*)\}$. This is necessary for the induction proof in Lemma 3 below to work, and can be maintained as follows. Assume at some iteration that $e^* = \{u^* - 1, u^*\}$ and $E = \{\{i, i+1\}, \{i+1, i+2\}, \ldots, \{r-1, r\}\}$. By the way we defined ϵ-relative pricings, we note that the following invariant holds throughout the algorithm:

(I) Every constraint in $\mathcal{S}$ is defined on an interval starting at $\{i, i+1\}$, or ending at $\{r-1, r\}$.

This is trivially true initially, and can be shown by induction for any iteration: (i) all the constrains of the form $(C1)$ or $(C2)$ defined on an edge of $E_L(e^*)$ are of the form $p([u, u^*]) \leq w$ (or $p([u, u^*]) = w$), where $u \in e \in E_L(e^*)$ and $w \in \mathbb{R}_+$. Note that $S(R)$ implies a constraint of the form $p(e^*) = q$ for some $q \in \mathbb{R}_+$, and thus any constraint $p([u, u^*]) \leq w$ can be reduced to the equivalent one $p([u, u^* - 1]) \leq w - q$. (ii) Any constraint in $\mathcal{S}$ of the form $p([i, u]) \leq w$, where $u > u^*$ and $w \in \mathbb{R}_+$, can be safely removed since $S(R) \wedge \mathcal{S}$ is feasible and $S(R)$ already contains the constraint $p([i, u^*]) = w'$, for some $w' \leq w$. (iii) $S(R)$ contains also the constraint $p([u^*, r]) = w$, for some $w \in \mathbb{R}_+$, and thus any constraint of $\mathcal{S}$ the form $p([u, r]) \leq w'$, where $u < u^*$ and $w' \geq w + q$ can be replaced by $p([u, u^* - 1] \leq w' - w - q$. Thus we conclude that the (new equivalent) system $\mathcal{S} \wedge S(R)$ satisfies the invariant on $E_L(e^*)$. A similar reasoning also shows the invariant is satisfied on $E_R(e^*)$.

When the procedure returns, we get two price functions $p_1 : E_L(e^*) \mapsto \mathbb{R}_+$ and $p_2 : E_R(e^*) \mapsto \mathbb{R}_+$. Let q be the value assigned to e^* by $S(R)$. We define a price function $p : E \mapsto \mathbb{R}_+$ on E as follows

$$p(e) = \begin{cases} p_1(e), & \text{if } e \in E_L(e^*) \\ p_2(e), & \text{if } e \in E_R(e^*) \\ q, & \text{If } e = e^*, \end{cases} \tag{6}$$

the procedure also returns a set of intervals which can be purchased under the returned price function p.

Lemma 2. *Algorithm* HP *runs in quasi-polynomial time in m, for any fixed $\epsilon > 0$.*

Proof. The number of possible ϵ-relative pricing is at most L, given in (4). This gives the recurrence

$$T(m) \leq \text{poly}(m) + 2L \cdot T(\frac{m}{2}).$$

for the running time. Thus $T(m) \leq L^{\log m + 1} \text{poly}(m)$ and the lemma follows. □

Lemma 3. *Algorithm* HP *returns a price function p and a set of intervals $\mathcal{J}$ such that $p(\mathcal{J}) \geq (1 - 3\epsilon) p^*(\text{OPT})$, for any $\epsilon > 0$.*

Algorithm HP$(\mathcal{I}, E, \mathcal{S})$**:**
Input: A subset of intervals $\mathcal{I}$ defined on E, and a feasible system of inequalities $\mathcal{S}$
Output: A price function $p : E \mapsto \mathbb{R}_+$ and a subset $\mathcal{J} \subseteq \mathcal{I}$ s.t. $p(I) \le B(I)\ \forall I \in \mathcal{J}$

1. **if** $|\mathcal{I}| = 0$, **then return** $(p, \emptyset)$, where p is any feasible solution of $\mathcal{S}$
2. let e^* be an edge of E such that $|\mathcal{I}_L(e^*)| \le m/2$ and $|\mathcal{I}_R(e^*)| \le m/2$
3. **for** every ϵ-relative pricing R w.r.t. e^* for which $\mathcal{S} \wedge S(R)$ is feasible **do**
4. $\quad (p_1, \mathcal{J}_1) \leftarrow \text{HP}(\mathcal{I}_L(e^*), E'_L(e^*), \mathcal{S} \wedge S(R))$
5. $\quad (p_2, \mathcal{J}_2) \leftarrow \text{HP}(\mathcal{I}_R(e^*), E'_R(e^*), \mathcal{S} \wedge S(R))$
6. $\quad$ let p be the price function defined by (6)
7. $\quad \mathcal{K} \leftarrow \{I \in \mathcal{I}[e^*] :\ v(I, R) \le B(I)/(1+\epsilon)\}$
8. $\quad \mathcal{J} \leftarrow \mathcal{K} \cup \mathcal{J}_1 \cup \mathcal{J}_2$
9. $\quad$ record $(p, \mathcal{J})$
10. **return** the recorded solution with largest $p(\mathcal{J})$ value

Fig. 2. The dynamic program for computing ϵ-approximate prices.

Proof. Fix an optimal solution OPT and an optimal price p^*, and let $\tilde{p}$ be an ϵ-optimal price. We prove by induction the following statement:

(H) Let $(\mathcal{I}, E, \mathcal{S})$ be the input to the algorithm and let $\mathcal{J}' \subseteq \mathcal{I}$ and $p' : E \mapsto \mathbb{R}_+$ be such that $p'(I) \le B(I)/(1+\epsilon)$ for all $I \in \mathcal{J}'$, and p' satisfies $\mathcal{S}$. Then the algorithm returns a pricing $p : E \mapsto \mathbb{R}_+$ and a subset $\mathcal{J} \subseteq \mathcal{I}$ such that (i) p satisfies $\mathcal{S}$, (ii) $p(I) \le B(I)$ for all $I \in \mathcal{J}$, and (iii) $p(\mathcal{J}) \ge p'(\mathcal{J}')/(1+\epsilon)$.

This can be used to prove the statement of theorem by setting $\mathcal{J}' \leftarrow$ OPT and $p' \leftarrow \tilde{p}$. Then it will follow by (H) and Proposition 1 that the algorithm returns a subset of intervals $\mathcal{J}$ and a pricing p such that all intervals in $\mathcal{J}$ are purchased, and their total price is at least $\tilde{p}(\text{OPT})/(1+\epsilon) \ge (1-3\epsilon)p^*(\text{OPT})$.

Now we prove (H). Let e^* be the middle edge. By Lemma 1, there is an ϵ-relative pricing R w.r.t e^*, consistent with p'. Note that the algorithm only considers pricings consistent with $\mathcal{S}$. One such pricing that will be eventually considered is R. Let us focus on the corresponding iteration of the loop beginning at line 3, and let $\mathcal{K}$ be the set identified in line 7 of this iteration. Since R is consistent with p', we get by (5) that $v(I, R) \le p'(I) \le B(I)/(1+\epsilon)$ for all $I \in \mathcal{J}'[e^*]$, implying that $\mathcal{J}'[e^*] \subseteq \mathcal{K}$. On the other hand, by (5) also we have $v(I, R) \ge p'(I)/(1+\epsilon)$ for all $I \in \mathcal{K} \subseteq \mathcal{I}[e^*]$. Thus

$$v(\mathcal{K}, R) \ge \frac{p'(\mathcal{K})}{1+\epsilon} \ge \frac{p'(\mathcal{J}[e^*])}{1+\epsilon}.$$

Note that the restrictions of p' to $E_L(e^*)$ and $E_R(e^*)$ together with $\mathcal{S} \wedge S(R)$ satisfy the preconditions of (H), with respect to $\mathcal{J}'_L(e^*)$ and $\mathcal{J}'_R(e^*)$, respectively. By induction, the algorithm returns two subsets of intervals $\mathcal{J}_1 \subseteq \mathcal{I}_L(e^*)$ and $\mathcal{J}_2 \subseteq \mathcal{I}_R(e^*)$, and two pricing functions $p_1 : E_L(e^*) \mapsto \mathbb{R}_+$ and $p_2 : E_R(e^*) \mapsto$

$\mathbb{R}_+$, both satisfying with $\mathcal{S} \wedge S(R)$, such that $p_1(I) \leq B(I)$ for all $I \in \mathcal{J}_1$ and $p_2(I) \leq B(I)$ for all $I \in \mathcal{J}_2$, and

$$p_1(\mathcal{J}_1) \geq p'(\mathcal{J}'_L(e^*))/(1+\epsilon) \text{ and } p_2(\mathcal{J}_2) \geq p'(\mathcal{J}'_R(e^*))/(1+\epsilon).$$

Let $\mathcal{J} = \mathcal{K} \cup \mathcal{J}_1 \cup \mathcal{J}_2$ and p be the function defined by (6). Then p satisfies $\mathcal{S} \wedge S(R)$ by construction, and $p(I) \leq B(I)$ for all $I \in \mathcal{J}_1 \cup \mathcal{J}_2$. Moreover, for any $I \in \mathcal{K}$ we have $p(I) \leq (1+\epsilon)v(I,R) \leq B(I)$, by (5) and the definition of $\mathcal{K}$. Thus p and $\mathcal{J}$ satisfy (i), (ii), and moreover,

$$\begin{aligned} p(\mathcal{J}) &= p(\mathcal{K}) + p_1(\mathcal{J}_1) + p_2(\mathcal{J}_2) \geq v(\mathcal{K},R) + p_1(\mathcal{J}_1) + p_2(\mathcal{J}_2) \\ &\geq \frac{p'(\mathcal{J}'[e^*]) + p'(\mathcal{J}'_L[e^*]) + p(\mathcal{J}'_R[e^*])}{1+\epsilon} = \frac{p'(\mathcal{J}')}{1+\epsilon}. \end{aligned} \tag{7}$$

Finally, we note that any solution $(p, \mathcal{J})$ returned by the algorithm must satisfy (i) and (ii). This is obvious for (i) since only pricings consistent with $\mathcal{S}$ are considered (see line 3). For (ii), this follows by induction for the two sets $\mathcal{J}_1$ and $\mathcal{J}_2$ computed by the algorithm, and by the definition of $\mathcal{K}$ and the fact that p is consistent with R. Thus the solution returned by the algorithm has value at least (7). □

5 The unsplittable flow pricing problem

In this case, we combine the method described in the previous section with the technique of [3]. We assume that the demands and capacities are integers bounded by $L' = 2^{\text{polylog}(m)}$. By rescaling, we can assume, at the loss of a factor of $(1-\epsilon)$ of the optimum, that $B(I) \in [1, nm/\epsilon]$ and $\rho(I) \in [1/L', 1]$, for all $I \in \mathcal{I}$. Let $Q = 1 + \lfloor \log \max_I \{B(I)/\rho(I)\} \rfloor = \text{polylog}(m)$. Given an edge e^* and an ϵ-relative pricing R w.r.t. e^*, we partition the intervals of $\mathcal{I}$ into at most Q classes according to the value of $v(I,R)/\rho(I)$: for $q \in [Q]$, $\mathcal{I}^{q,R} = \{I \in \mathcal{I} : \ 2^{q-1} \leq \frac{v(I,R)}{\rho(I)} \leq 2^q\}$.

Theorem 2. *There is a quasi-polynomial time approximation scheme for* Ufp *on line graphs, provided that the demands and capacities are integers bounded by a quasi-polynomial in the number of intervals.*

Corollary 1. *There is a quasi-polynomial time approximation scheme for the pricing problem on line graphs with limited supply.*

We recall the following definition from [3].

Definition 3. (ϵ-Restricted profiles) *Let $e = \{u-1, u\}$ be an edge of E, and $h, \epsilon \in \mathbb{R}_+$, with $1/\epsilon \in \mathbb{Z}_+$. Let $x_1, \ldots, x_{1/\epsilon}$ and $y_1, \ldots, y_{1/\epsilon}$ be start points of edges in E, such that*

$$x_1 \leq x_2 \leq \cdots \leq x_{1/\epsilon} \leq u-1 < u \leq y_{1/\epsilon} \leq \cdots \leq y_2 \leq y_1.$$

Then the vector $(\ell_1, \ldots, \ell_n)$*, where*

$$\ell_i = \begin{cases} 0, & \textit{for } i \le x_1 \textit{ and } i > y_1 \\ j\epsilon h, & \textit{for } x_j < i \le x_{j+1} \textit{ and } y_{j+1} < i \le y_j \\ h, & \textit{for } x_{1/\epsilon} < i \le y_{1/\epsilon}, \end{cases}$$

is said to be an ϵ-restricted profile *with peak* e *and* height h*, denoted by* $RP_\epsilon(e; h; x_1, \ldots, x_{1/\epsilon}; y_1, \ldots, y_{1/\epsilon})$.

The total number of ϵ-restricted profiles, with a given peak and height is at most $n^{2/\epsilon}$.

For $\mathcal{J} \subseteq \mathcal{I}$ and any edge $e \in E$, the total demands of $\mathcal{J}[e]$ define a profile $\text{prof}(\mathcal{J}[e]) \in \mathbb{R}^E$, defined by:

$$\text{prof}(\mathcal{J}[e])_{e'} = \rho(\mathcal{J}[e] \cap \mathcal{J}[e']) \stackrel{\text{def}}{=} \sum_{I \in \mathcal{J}[e] \cap \mathcal{J}[e']} \rho(I), \quad \text{for } e' \in E.$$

The following two lemmas state that any such profile can be sufficiently accurately approximated by an ϵ-restricted profile, in polynomial time (see Figure 3), provided that the demands are sufficiently small.

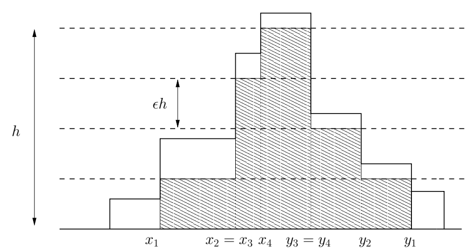

Fig. 3. A profile and its restriction.

Lemma 4 ([3]). *Let* $e^* \in E$ *be an edge,* R *be an* ϵ*-relative pricing w.r.t.* e^**, and* $\mathcal{J} \subseteq \mathcal{I}^{q,R}[e^*]$ *be a subset of class* q *intervals spanning* e^**, such that* $\rho(I) \le \delta$ *for all* $I \in \mathcal{J}$ *and some* $\delta \in \mathbb{R}_+$*. Let* h *be the largest integer multiple of* $1/2^q$ *that does not exceed* $\rho(\mathcal{J})$*. Then there exists an* ϵ*-restricted profile* π *with peak* e^* *and height* h *and a subset* $\mathcal{J}' \subseteq \mathcal{J}$*, such that*

(i) $\text{prof}(\mathcal{J}') \le \pi \le \text{prof}(\mathcal{J})$*, and*

(ii) $v(\mathcal{J}, R) - v(\mathcal{J}', R) \le 2^{q+1}(\epsilon h + \delta)$.

Lemma 5 ([3]). *Let* e^* *be an edge,* R *an* ϵ*-relative pricing w.r.t.* e^**, and* $\mathcal{J} \subseteq \mathcal{I}^{q,R}[e^*]$ *be a subset of class* q *intervals such that* $\rho(I) \le \delta$ *for all* $I \in \mathcal{J}$*. Let* π *be a restricted profile, and* $\mathcal{J}^*$ *be a subset of* $\mathcal{J}$ *such that* $\text{prof}(\mathcal{J}^*) \le \pi$ *and* $v(\mathcal{J}^*, R)$ *is maximized. Then we can find in polynomial time a subset* $\mathcal{J}' \subseteq \mathcal{J}$ *such that* $\text{prof}(\mathcal{J}') \le \pi$ *and* $v(\mathcal{J}^*, R) - v(\mathcal{J}', R) \le 2^{q+1}\delta/\epsilon$.

Remark 1. If $\rho(I) = 1$ for all $I \in \mathcal{J}$, then a subset satisfying the conditions in Lemma 5 with $v(\mathcal{J}', R) = v(\mathcal{J}^*, R)$ can be found by a greedy procedure as follows. Consider the intervals in $\mathcal{J}$ in increasing order of their $v(\cdot, R)$ values, and let $\mathcal{J}'$ be a maximal set of intervals such that $\text{prof}(\mathcal{J}') \leq \pi$. Then it is easy to see that $v(\mathcal{J}', R) \geq v(\mathcal{J}^*, R)$.

The algorithm is an extension of HP and is given in Figure 4 below. It resembles very much the algorithm in [3], with the additional incorporation of ϵ-relative pricings. Without loss of generality we can assume that $c(e) < \infty$ for all $e \in E$ by setting $c_e \leftarrow \min\{c_e, \rho(\mathcal{I}[e])\}$. For a subset of intervals $\mathcal{J}$, a restricted profile π, and a relative pricing R, we use $\text{PACK}(\mathcal{J}, \pi, R)$ to denote a procedure that returns a subset with the guarantees of Lemma 5.

The algorithm proceeds as before, by considering the middle edge e^*. It partitions the set of candidate intervals crossing e^* in a given class $q \in [Q]$ into those with large demands $\mathcal{A}_q$, and those with small demands $\mathcal{B}_q$. The number of large demands crossing e^* is small (less than $1/\epsilon^2$) and hence they can be guessed by enumerating over all of them. For each such guess of large demands, the profile of the small demands crossing e^* can be approximated by an ϵ-restricted profile using Lemmas 4 and 5. Again, among these profiles, we enumerate over the ones whose total demand requirement does not exceed the residual capacity left after routing the large demands (see [3] for more details).

Lemma 6. *Algorithm* UFP *runs in quasi-polynomial time in m, for any fixed $\epsilon > 0$.*

Lemma 7. *Algorithm* UFP *returns a solution of value at least $(1-23\epsilon)p^*(\textsc{Opt})$, for any $0 < \epsilon < 1$.*

References

1. G. Aggarwal and J. D. Hartline, *Knapsack auctions*, SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm (New York, NY, USA), ACM Press, 2006, pp. 1083–1092.
2. M. F. Balcan and A. Blum, *Approximation algorithms and online mechanisms for item pricing*, EC '06: Proceedings of the 7th ACM conference on Electronic commerce (New York, NY, USA), ACM Press, 2006, pp. 29–35.
3. N. Bansal, A. Chakrabarti, A. Epstein, and B. Schieber, *A quasi-PTAS for unsplittable flow on line graphs*, STOC '06: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing (New York, NY, USA), ACM Press, 2006, pp. 721–729.
4. P. Briest and P. Krysta, *Single-minded unlimited supply pricing on sparse instances*, SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm (New York, NY, USA), ACM Press, 2006, pp. 1093–1102.
5. ———, *Buying cheap is expensive: Hardness of non-parametric multi-product pricing*, Proc. 17th Annual ACM-SIAM Symposium on Discrete Algorithms, ACM-SIAM, 2007.

Algorithm UFP$(\mathcal{I}, E, \mathcal{S}, c)$**:**

Input: A subset of intervals $\mathcal{I}$ defined on E, a feasible system of linear inequalities $\mathcal{S}$, and a capacity vector $c = (c(e) : e \in E)$

Output: A price function $p : E \mapsto \mathbb{R}_+$ and a subset $\mathcal{J} \subseteq \mathbb{R}_+$ s.t. $p(I) \leq B(I)\ \forall I \in \mathcal{J}$ and $\rho(\mathcal{J}[e]) \leq c(e)\ \forall e \in E$

if $|\mathcal{I}| = 0$, **then return** $(p, \emptyset)$, where p is any feasible solution of $\mathcal{S}$
let e^* be an edge of E such that $|\mathcal{I}_L(e^*)| \leq m/2$ and $|\mathcal{I}_R(e^*)| \leq m/2$
for every ϵ-relative pricing R w.r.t. e^* for which $\mathcal{S} \wedge S(R)$ is feasible **do**
  foreach $q \in [Q]$, set $\mathcal{I}_q = \{I \in \mathcal{I}^{q,R}[e^*] : \ v(I, R) \leq B(I)/(1+\epsilon)\}$
  foreach Q-tuple $(\mathcal{A}_1, \ldots, \mathcal{A}_Q)$ with each $\mathcal{A}_q \subseteq \mathcal{I}_q$ and $|\mathcal{A}_q| \leq 1/\epsilon^2$ **do**
    $\mathcal{A} \leftarrow \bigcup_{q=1}^{Q} \mathcal{A}_q$
    if $\rho(\mathcal{A}) \leq c(e)\ \forall e \in E$ **then**
      foreach $e \in E$, set $c'(e) \leftarrow c(e) - \rho(\mathcal{A}[e])$
      foreach $(h_1, \ldots, h_Q) \in \mathbb{R}_+^Q$ with each h_q an integer multiple of $1/2^q$ and $\sum_{q=1}^{Q} h_q \leq c'(e)$ **do**
        foreach Q-tuple $(\pi_1, \ldots, \pi_Q)$ with each π_q an ϵ-restricted profile with peak e^* and height h_q, s.t. $\sum_{q=1}^{Q} \pi_q \leq c'$ **do**
          foreach $q \in [Q]$, set $\mathcal{B}_q \leftarrow \{I \in \mathcal{I}_q \setminus \mathcal{A}_q : \ \rho(I) \leq \epsilon^2(h_q + 1/2^q + \rho(\mathcal{A}_q))\}$
          foreach $q \in [Q]$, let $\mathcal{K}_q \leftarrow$PACK$(\mathcal{B}_q, \pi_q, R)$
          $\mathcal{K} \leftarrow \bigcup_{q=1}^{Q} \mathcal{K}_q$
          foreach $e \in E$, set $c''(e) \leftarrow c'(e) - \rho(\mathcal{B}[e])$
          $(p_1, \mathcal{J}_1) \leftarrow$ UFP$(\mathcal{I}_L(e^*), E_L(e^*), \mathcal{S} \wedge S(R), c'')$
          $(p_2, \mathcal{J}_2) \leftarrow$ UFP$(\mathcal{I}_R(e^*), E_R(e^*), \mathcal{S} \wedge S(R), c'')$
          let p be the price function defined by (6)
          $\mathcal{J} \leftarrow \mathcal{A} \cup \mathcal{K} \cup \mathcal{J}_1 \cup \mathcal{J}_2$
          record $(p, \mathcal{J})$
return the recorded solution with largest $p(\mathcal{J})$ value

Fig. 4. The dynamic program for computing ϵ-approximate prices for the capacitated version.

6. E. D. Demaine, M. T. Hajiaghayi, U. Feige, and M. R. Salavatipour, *Combination can be hard: approximability of the unique coverage problem*, SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm (New York, NY, USA), ACM Press, 2006, pp. 162–171.
7. P. W. Glynn, B. Van Roy, and P. Rusmevichientong, *A nonparametric approach to multi-product pricing*, Operations Research **54** (2006), to appear.
8. V. Guruswami, J. D. Hartline, A. R. Karlin, D. Kempe, C. Kenyon, and F. McSherry, *On profit-maximizing envy-free pricing*, SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms (Philadelphia, PA, USA), Society for Industrial and Applied Mathematics, 2005, pp. 1164–1173.
9. J. D. Hartline and V. Koltun, *Near-optimal pricing in near-linear time*, Algorithms and Data Structures - WADS 2005 (F. K. H. A. Dehne, A. López-Ortiz, and J.-R. Sack, eds.), Lecture Notes in Computer Sciences, vol. 3608, Springer, 2005, pp. 422–431.