Improved Hardness Results for Profit Maximization Pricing Problems with Unlimited Supply

Parinya Chalermsook[1⋆], Julia Chuzhoy[2⋆⋆], Sampath Kannan[3⋆⋆⋆], and Sanjeev Khanna[3†]

[1] University of Chicago, Chicago, IL and IDSIA, Lugano, Switzerland
[2] Toyota Technological Institute, Chicago, IL
[3] Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA

Abstract. We consider profit maximization pricing problems, where we are given a set of m customers and a set of n items. Each customer c is associated with a subset $S_c \subseteq [n]$ of items of interest, together with a budget B_c, and we assume that there is an unlimited supply of each item. Once the prices are fixed for all items, each customer c buys a subset of items in S_c, according to its buying rule. The goal is to set the item prices so as to maximize the total profit.

We study the unit-demand min-buying pricing ($\mathsf{UDP_{MIN}}$) and the single-minded pricing (SMP) problems. In the former problem, each customer c buys the cheapest item $i \in S_c$, if its price is no higher than the budget B_c, and buys nothing otherwise. In the latter problem, each customer c buys the whole set S_c if its total price is at most B_c, and buys nothing otherwise. Both problems are known to admit $O(\min\{\log(m+n), n\})$-approximation algorithms. We prove that they are $\log^{1-\epsilon}(m+n)$ hard to approximate for any constant ϵ, unless $\mathsf{NP} \subseteq \mathsf{DTIME}(n^{\log^{\delta} n})$, where δ is a constant depending on ϵ. Restricting our attention to approximation factors depending only on n, we show that these problems are $2^{\log^{1-\delta} n}$-hard to approximate for any $\delta > 0$ unless $\mathsf{NP} \subseteq \mathsf{ZPTIME}(n^{\log^{\delta'} n})$, where δ' is some constant depending on δ. We also prove that restricted versions of $\mathsf{UDP_{MIN}}$ and SMP, where the sizes of the sets S_c are bounded by k, are $k^{1/2-\epsilon}$-hard to approximate for any constant ϵ.

We then turn to the Tollbooth Pricing problem, a special case of SMP, where each item corresponds to an edge in the input graph, and each set

⋆ Supported in part by NSF CAREER grant CCF-0844872, Swiss National Science Foundation project 200020-122110/1, and Hasler Foundation Grant 11099. Part of this work was done while at University of Chicago

⋆⋆ Supported in part by NSF CAREER grant CCF-0844872 and Sloan Research Fellowship.

⋆⋆⋆ Supported in part by the National Science Foundation grant CCF-1137084

† Supported in part by the National Science Foundation grants CCF-1116961 and IIS-0904314

S_c is a simple path in the graph. We show that Tollbooth Pricing is at least as hard to approximate as the Unique Coverage problem, thus obtaining an $\Omega(\log^{\epsilon} n)$-hardness of approximation, assuming $\mathsf{NP} \not\subseteq \mathsf{BPTIME}(2^{n^{\delta}})$, for any constant δ, and some constant ϵ depending on δ.

1 Introduction

We study profit maximization pricing problems in the unlimited supply model. In these problems, we are given a set of m customers and a set of n items, where each customer c is associated with a budget B_c, and a subset $S_c \subseteq [n]$ of items it is interested in. Our goal is to set a price $p(i)$ for each item $i \in [n]$, so as to maximize the total revenue. Once the prices for the items are set, each customer c chooses a subset of items in S_c to buy, using its *buying rule*. We assume that we are given an unlimited supply of each item.

One of the most natural buying rules is the unit-demand min-buying rule, where each customer $c \in [m]$ buys the cheapest item $i \in S_c$ (breaking ties arbitrarily), provided that the price $p(i) \leq B_c$. We refer to the corresponding pricing problem as $\mathsf{UDP}_{\mathsf{MIN}}$. This problem was first introduced by Rusmevichientong et al. [18, 19], and subsequently Aggarwal et al. [1] have shown an $O(\log m + \log n)$-approximation algorithm for it.

The second problem that we consider is Single-Minded Pricing (SMP). Here, each customer c buys the whole set S_c of items if its total price does not exceed its budget B_c, and buys nothing otherwise. This problem was introduced by Guruswami et al. [14], who also show that the techniques of [1] can be used to obtain an $O(\log m + \log n)$-approximation algorithm for SMP. Hartline and Koltun [15] gave a $(1 + \epsilon)$-approximation algorithm for both $\mathsf{UDP}_{\mathsf{MIN}}$ and SMP when the number of items n is constant.

We remark that for pricing problems, it is natural to assume that the number of customers is much higher than the number of items, that is, $m >> n$. Even though both $\mathsf{UDP}_{\mathsf{MIN}}$ and SMP admit logarithmic approximation algorithms in terms of $(m + n)$, if we restrict ourselves to approximation factors depending only on n, nothing better than the trivial $O(n)$-approximation is known.

On the negative side, Briest [3] has shown that both $\mathsf{UDP}_{\mathsf{MIN}}$ and SMP are $\max\left\{n^{\delta}, \log^{\delta}(m + n)\right\}$-hard to approximate for some (small) $\delta > 0$, assuming that no randomized polynomial-time algorithms can approximate constant-degree Balanced Bipartite Independent Set to within arbitrarily small constant factors. He also showed similar results under an assumption that slightly strengthens Feige's Random 3SAT hypothesis [11].

In this paper, we show that both $\mathsf{UDP}_{\mathsf{MIN}}$ and SMP are $\log^{1-\epsilon}(m + n)$ hard to approximate for any constant ϵ, unless $\mathsf{NP} \subseteq \mathsf{DTIME}(n^{\log^{\epsilon'} n})$ for some constant ϵ' depending only on ϵ. If we restrict our attention to approximation factors as a function of n, then we show that both these problems are $2^{\log^{1-\delta} n}$ hard to approximate for any constant δ, under the assumption that $\mathsf{NP} \not\subseteq \mathsf{ZPTIME}(n^{\log^{\delta'} n})$, for some constant δ' depending only on δ.

We next turn to restricted versions of $\mathsf{UDP_{MIN}}$ and SMP, denoted by $\mathsf{kUDP_{MIN}}$ and kSMP respectively, where the sizes of the sets S_c are bounded by k. The kSMP problem is known to be APX-hard even for $k = 2$ [14], and Balcan and Blum [2] have shown an $O(k)$-approximation for $\mathsf{kUDP_{MIN}}$, improving on an independent work of Briest and Krysta [4], who achieved an $O(k^2)$-approximation for the problem. As for negative results, Briest [3] has proved that kSMP is k^{ϵ}-hard to approximate for some constant ϵ, assuming Feige's random 3SAT hypothesis [11], and Khandekar et al. [16] showed that the problem is $\Omega(k)$ hard to approximate for constant k, assuming the Unique Games Conjecture of Khot [17]. We show that both $\mathsf{kUDP_{MIN}}$ and kSMP are $k^{1/2-\epsilon}$-hard to approximate for any constant ϵ unless $\mathsf{P} = \mathsf{NP}$.

Finally, we consider a special case of the SMP problem called the Tollbooth Pricing problem, where we are given a graph G, and items correspond to the edges of G. The item set S_c of every customer c is some simple path in graph G, and the goal is to set the prices of the edges, so as to maximize the revenue. Since the Tollbooth Pricing problem is a special case of SMP, it admits an $O(\log m + \log n)$ approximation [14]. The problem is APX-hard [14], and from the results of Khandekar et al. [16], it is $(2 - \epsilon)$ hard to approximate even on star graphs, assuming the Unique Games Conjecture. We show that the Tollbooth Pricing problem is at least as hard to approximate as the Unique Coverage problem (to within a constant factor). In the Unique Coverage problem, we are given a collection U of n elements, and a family $\mathcal{S}$ of subsets of elements of U. The goal is to find a family $\mathcal{S}' \subseteq \mathcal{S}$ of element subsets, maximizing the number of elements that are covered by exactly one subset in $\mathcal{S}'$. The problem was introduced and studied by Demaine et al. [8], who showed that for any arbitrarily small constant δ, if $\mathsf{NP} \not\subseteq \mathsf{BPTIME}(2^{n^{\delta}})$, then Unique Coverage is hard to approximate to within a factor of $\Omega(\log^{\epsilon} n)$, where ϵ is some constant depending on δ. They also showed that the problem is hard to approximate to within $\Omega(\log^{1/3-\epsilon} n)$ for any ϵ assuming the Random 3SAT Hypothesis of Feige [11], and proved additional hardness results using a hypothesis about Balanced Bipartite Independent Set. Our reduction immediately implies similar hardness results for the Tollbooth Pricing problem.

Related Work. Briest and Krysta [4] considered a more general version of $\mathsf{UDP_{MIN}}$, where customers are allowed to have different budgets (valuations) for different items. They show an $\Omega(\log^{\epsilon} n)$-hardness for this problem for some constant ϵ, unless $\mathsf{NP} \subseteq \mathsf{DTIME}\left(n^{O(\log\log n)}\right)$, and an n^{ϵ}-hardness for some constant $\epsilon > 0$, unless $\mathsf{NP} \subseteq \mathsf{DTIME}\left(2^{O(n^{\delta})}\right)$ for all $\delta > 0$.

A special case of the Tollbooth Pricing problem, called the Highway Problem, where the input graph is restricted to be a path, has received a significant amount of attention. Elbassioni et al. [9] showed that the problem is strongly NP-hard. On the algorithmic side, Balcan and Blum [2] have shown an $O(\log n)$-approximation algorithm, and Elbassioni et al. [10] have proposed a QPTAS. Subsequently, Grandoni and Rothvoss [13] have shown a PTAS for the problem. For the special case of the Tollbooth Pricing problem where the input graph is a tree, the best known approximation ratio is $O(\log n/\log\log n)$, due to Gamzu and Segev [12]. However, when the number of leaves in the tree is bounded by a constant, the problem admits a PTAS [13].

Pricing problems with limited supply have also received a considerable amount of attention; Please refer to, e.g., [5, 7, 6] and references therein.

Our Results. We start by formally stating the pricing problems we consider. We are given a set of m customers and a set of n items, where each customer $c \in [m]$ is associated with a set $S_c \subseteq [n]$ of items and a budget B_c. Given a setting $\{p(i)\}_{i\in[n]}$ of item prices, every customer selects a subset $S'_c \subseteq S_c$ of items to buy according to its buying rule, and our goal is to maximize the total profit, $\sum_{c\in[m]}\sum_{i\in S'_c} p(i)$. In the $\mathsf{UDP_{MIN}}$ problem, the buying rule of the customers is defined as follows. Each customer $c \in [m]$ buys the cheapest item $i \in S_c$, breaking ties arbitrarily, if $p(i) \le B_c$, and buys nothing otherwise.

In the SMP problem, each customer $c \in [m]$ purchases the whole set S_c if $\sum_{i\in S_c} p(i) \le B_c$, and purchases nothing otherwise. Our main result is summarized in the following theorem.

Theorem 1. *$\mathsf{UDP_{MIN}}$ and SMP are $\log^{1-\epsilon}(m+n)$-hard to approximate for any constant $\epsilon > 0$, unless $\mathsf{NP} \subseteq \mathsf{DTIME}(n^{(\log n)^{\epsilon'}})$, where ϵ' is some constant depending only on ϵ. Moreover, assuming that $\mathsf{NP} \not\subseteq \mathsf{ZPTIME}(n^{(\log n)^{\delta'}})$, both problems are hard to approximate to within a factor of $2^{\log^{1-\delta} n}$ for any constant δ, where δ' is some constant depending only on δ.*

We next turn to special cases of both problems, denoted by $\mathsf{kUDP_{MIN}}$ and kSMP respectively, where the sizes of the sets S_c are bounded by k and prove the following theorem.

Theorem 2. *Let $\epsilon > 0$ be any constant. Then for infinitely many constants k, both $\mathsf{kUDP_{MIN}}$ and kSMP are $k^{1/2-\epsilon}$-hard to approximate unless $\mathsf{P} = \mathsf{NP}$.*

Finally we turn to the Tollbooth Pricing problem. In this problem, we are given a graph $G = (V, E)$, and a set of m simple paths $P_1, \ldots, P_m$, where each path P_c is associated with a customer c and a budget B_c. Once the price function $p : E \to \mathbb{R}$ on the edges is set, each customer c buys all edges on the path P_c if $\sum_{e\in P_c} p(e) \le B_c$, and buys nothing otherwise. The goal is to compute the edge prices $p(e)$ so as to maximize the total profit. It is clear that Tollbooth Pricing is a special case of SMP, and notice that the number of items is $n = |E(G)|$.

We perform a reduction from the Unique Coverage problem to the Tollbooth Pricing. In the Unique Coverage problem, we are given a set U of elements and a family $\mathcal{S}$ of subsets of U as input. A solution is a sub-collection $\mathcal{S}' \subseteq \mathcal{S}$ of the input sets. We say that element $u \in U$ is *satisfied* by the solution if and only if it belongs to exactly one set in $\mathcal{S}'$. Our goal is to choose $\mathcal{S}'$ so as to maximize the number of satisfied elements. Demaine et. al. [8] have shown that for any arbitrarily small constant δ, if $\mathsf{NP} \not\subseteq \mathsf{BPTIME}(2^{n^\delta})$, then Unique Coverage is hard to approximate to within a factor of $\Omega(\log^\epsilon n)$, for some constant ϵ depending on δ. They also showed that, under the assumption of Feige [11] that refuting random instances of 3SAT is hard, Unique Coverage is hard to approximate to within a factor of $\Omega(\log^{1/3-\epsilon} n)$ for any $\epsilon > 0$. We prove the following theorem:

Theorem 3. *If there is a factor α-approximation algorithm for the* Tollbooth Pricing *problem, for any approximation factor $\alpha \leq O(\log n)$, then there is a randomized $O(\alpha)$-approximation algorithm for the* Unique Coverage *problem.*

Combining this with the result of [8], we obtain the following corollary.

Corollary 1. *For any arbitrarily small constant δ, if $\mathsf{NP} \not\subseteq \mathsf{BPTIME}(2^{n^\delta})$,* Tollbooth Pricing *is hard to approximate to within a factor of $\Omega(\log^\epsilon n)$ for some constant ϵ depending on δ. Moreover, under Feige's random 3SAT assumption, this problem is hard to approximate to within a factor of $\Omega(\log^{1/3-\epsilon} n)$ for any $\epsilon > 0$.*

2 Hardness of $\mathsf{UDP}_{\mathsf{MIN}}$ and SMP

In this section we prove Theorems 1 and 2. We focus here on the $\mathsf{UDP}_{\mathsf{MIN}}$ problem only. The hardness results for SMP are obtained using similar ideas and appear in the full version of the paper.

We start with the following theorem, due to Trevisan [20]. Since we use slightly different parameters, we provide the proof in the full version.

Theorem 4. *Given an n-variable 3SAT formula φ, any sufficiently small constant $\epsilon > 0$ and any integer $\lambda > 0$, there is a randomized algorithm to construct a graph G with maximum degree at most $\Delta = 2^{\lambda \operatorname{poly}(\frac{1}{\epsilon})}$ such that w.h.p.:*

- *(*Yes-Instance*:) If φ is satisfiable, then G has an independent set of size $|V(G)|/\Delta^\epsilon$.*
- *(*No-Instance*:) If φ is not satisfiable, then G has no independent set of size $|V(G)|/\Delta^{1-\epsilon}$.*

The construction size is $|V(G)| = n^{\lambda \operatorname{poly}(\frac{1}{\epsilon})}$, and the reduction runs in time $n^{\lambda \operatorname{poly}(\frac{1}{\epsilon})}$. Moreover, the algorithm can be made deterministic with running time $2^{O(\Delta)} n^{\lambda \operatorname{poly}(\frac{1}{\epsilon})}$.

We remark that this theorem allows us to adjust parameter λ. To prove Theorem 1, we will use $\lambda = O(\log \log n)$, while we set $\lambda = O(1)$ for Theorem 2.

2.1 The Construction

Let $G = (V, E)$ be the instance of Maximum Independent Set obtained from Theorem 4, where the value of λ (and Δ) will be fixed later. We first define an intermediate instance of $\mathsf{UDP_{MIN}}$, which is then converted into a final instance.

The intermediate instance is defined as follows. The set of items contains, for each vertex $v \in V$, for each index $y \in [\Delta]$, an item $i(v, y)$. That is, the set of items is $\mathcal{I} = \{i(v, y) \mid v \in V, y \in [\Delta]\}$.

Similarly, the set of customers contains, for each vertex $v \in V$, for each index $x \in [\Delta]$, a customer $c(v, x)$. That is, the set of customers is $\mathcal{C} = \{c(v, x) \mid v \in V, x \in [\Delta]\}$.

The item set $S_{c(v,x)}$ for the customer $c(v, x)$, contains the item $i(v, x)$, and additionally, for each neighbor u of vertex v in graph G, for each index $y \in [\Delta]$, item $i(u, y)$ belongs to $S_{c(v,x)}$. Formally, $S_{c(v,x)} = \{i(u, y) \mid (u, v) \in E, y \in [\Delta]\} \cup \{i(v, x)\}$. Notice that $\left|S_{c(v,x)}\right| \le \Delta^2 + 1$ for all customers $c(v, x) \in \mathcal{C}$. Moreover for each item $i(v, y) \in \mathcal{I}$, there are at most $\Delta^2 + 1$ customers $c' \in \mathcal{C}$ such that $i(v, y) \in S_{c'}$.

We partition the set $\mathcal{C}$ of customers into Δ subsets $\mathcal{C}_1, \ldots, \mathcal{C}_\Delta$, such that for each $1 \le h \le \Delta$, set $\mathcal{C}_h$ contains customers $c(v, h)$ for all $v \in V$. Finally, for each $1 \le h \le \Delta$, each customer $c \in \mathcal{C}_h$ is assigned budget $B_c = 1/2^h$.

This finishes the definition of the intermediate instance. For convenience, we call the customers in set $\mathcal{C}$ *virtual customers*. In our final instance, we replace each virtual customer with a number of new customers.

In order to define the final instance, for each $1 \le h \le \Delta$, we replace each virtual customer $c \in \mathcal{C}_h$ with a set $\mathcal{G}(c) = \left\{c(1), \ldots, c(2^h)\right\}$ of 2^h identical new customers. Each new customer $c(h')$, for $1 \le h' \le 2^h$ has budget $B_{c(h')} = B_c$ and $S_{c(h')} = S_c$. The final set of customers is $\mathcal{C}' = \bigcup_{c \in \mathcal{C}} \mathcal{G}(c)$ and the final set of items remains unchanged, $\mathcal{I}' = \mathcal{I}$. The number of customers in the final instance is $\tilde{m} = |\mathcal{C}'| = O(2^\Delta |\mathcal{C}|) = |V| \cdot \Delta \cdot 2^{O(\Delta)} = |V| \cdot 2^{O(\Delta)}$, while the number of items is $\tilde{n} = |V| \cdot \Delta$. Moreover, for each customer $c \in \mathcal{C}'$, we have $|S_c| \le \Delta^2 + 1$. This completes the construction description.

2.2 Analysis

We analyze the construction in the following two lemmas.

Lemma 1. *In the* YES-INSTANCE*, there is a solution to the* $\mathsf{UDP_{MIN}}$ *problem instance whose value is at least* $|V|\Delta^{1-\epsilon}$.

Proof. Let $U \subseteq V$ be a maximum independent set of size $|V|/\Delta^\epsilon$ in G. We set the prices of the items $i(u,y) \in \mathcal{I}'$ as follows. If $u \notin U$, then the price of $i(u,y)$ is set to ∞. Otherwise, if $u \in U$, then we set the price of $i(u,y)$ to $1/2^y$. Notice that, since $|U| \cdot \Delta \geq |V| \cdot \Delta^{1-\epsilon}$, there are $|V| \cdot \Delta^{1-\epsilon}$ items of finite prices. We now show that this solution has value at least $|V| \cdot \Delta^{1-\epsilon}$.

Indeed, for each vertex $u \in U$ and an index $y \in [\Delta]$, consider the virtual customer $c' = c(v,y) \in \mathcal{C}_y$. Notice that $S_{c'}$ contains item $i(v,y)$ whose price is $1/2^y$, but all other items in $S_{c'}$ have price ∞. Therefore, each customer $c \in \mathcal{G}(c')$ buys the item $i(v,y)$, and pays $1/2^y$ for it. The total profit collected from customers in $\mathcal{G}(c')$ is 1, and so the total profit collected from all customers is at least $|U|\Delta \geq |V| \cdot \Delta^{1-\epsilon}$.

Lemma 2. *In the* No-Instance*, the value of the optimal solution is at most* $O(|V| \cdot \Delta^\epsilon)$.

Proof. Let p^* be an optimal solution, and let r^* be its revenue. We first argue that we can assume w.l.o.g. that for each item $i \in \mathcal{I}'$, either $p^*(i) \in \left\{1/2^h \mid 1 \leq h \leq \Delta\right\}$, or $p^*(i) = \infty$.

Indeed, suppose there is an item $i \in \mathcal{I}'$ with $p^*(i) \in (1/2^h, 1/2^{h-1})$. Then any customer who buys item i must have budget at least $1/2^{h-1}$, so increasing $p^*(i)$ to $1/2^{h-1}$ does not affect these customers, and may only increase the revenue. Therefore, from now on we assume that for each item $i \in \mathcal{I}'$, $p^*(i) \in \left\{1/2^h \mid 1 \leq h \leq \Delta\right\} \cup \{\infty\}$.

Notice that for each virtual customer $c \in \mathcal{C}$, all customers in $\mathcal{G}(c)$ contribute the same amount to the total revenue. Let k_c denote this amount. We now let $\mathcal{C}^* \subseteq \mathcal{C}$ be the set of virtual customers for which $k_c = B_c$. Equivalently,

$$\mathcal{C}^* = \left\{ c \in \mathcal{C} : \min_{i \in S_c} \{p^*(i)\} = B_c \right\}$$

Claim. The customers in $\bigcup_{c' \in \mathcal{C}^*} \mathcal{G}(c')$ contribute at least $r^*/2$ to the total revenue.

Due to space limitation, the proof of this claim appears in the full version. Notice that $|\mathcal{C}^*| \geq r^*/2$, since for each virtual customer $c \in \mathcal{C}^*$, the total budget of all customers in $\mathcal{G}(c)$ is 1.

From now on, we focus on finding an independent set U in graph G of size at least $(r^*/2 - |V|)/\Delta$ from $\mathcal{C}^*$. Since in the No-Instance, G does not contain an independent set of size more than $|V|/\Delta^{1-\epsilon}$, this implies that $(r^*/2 - |V|)/\Delta \leq |V| / \Delta^{1-\epsilon}$, and hence $r^* \leq O(|V|\, \Delta^\epsilon)$.

We construct an independent set $U \subseteq V(G)$, together with a partition $(\mathcal{C}^1, \mathcal{C}^2)$ of $\mathcal{C}^*$, as follows. Start with $U, \mathcal{C}^1, \mathcal{C}^2 = \emptyset$. We then perform Δ iterations, where in iteration y, we consider each virtual customer $c(v,y)$ in $\mathcal{C}^* \cap \mathcal{C}_y$, and do the following:

– If vertex v is already in U, we add virtual customer $c(v,y)$ into $\mathcal{C}^1$.
– If vertex v is not in U and $U \cup \{v\}$ remains an independent set, we add vertex v to set U and add $c(v,y)$ to $\mathcal{C}^1$. We say that $c(v,y)$ is *responsible* for adding vertex v into U.
– Otherwise, $v \notin U$, but there is a vertex $u \in U$ such that $(u,v) \in E(G)$. We add $c(v,y)$ to $\mathcal{C}^2$ in this case and say that vertex u prevents the algorithm from adding v into U.

In the end, when all customers in $\mathcal{C}^*$ are processed, each virtual customer in $\mathcal{C}^*$ is added to either $\mathcal{C}^1$ or $\mathcal{C}^2$, so $\mathcal{C}^* = \mathcal{C}^1 \cup \mathcal{C}^2$. Moreover, for each virtual customer $c(v,y)$ in $\mathcal{C}^1$, the corresponding vertex v belongs to U, so $|U| \geq |\mathcal{C}^1|/\Delta$. The following claim will complete the proof of the lemma.

Claim. $|\mathcal{C}^2| \leq |V|$, and so $|U| \geq |\mathcal{C}^* \setminus \mathcal{C}^2|/\Delta \geq (r^*/2 - |V|)/\Delta$.

Proof. It is sufficient to show that for each vertex $v \in V$, no two virtual customers $c(v,y), c(v,y')$ with $y \neq y'$ belong to $\mathcal{C}^2$. Assume otherwise, and let $c(v,y), c(v,y') \in \mathcal{C}^2$. By our construction, we have $c(v,y) \in \mathcal{C}_y$ and $c(v,y') \in \mathcal{C}_{y'}$. Assume w.l.o.g. that $y < y'$, so $c(v,y)$ was processed before $c(v,y')$.

Let $u \in U$ be a vertex such that $(u,v) \in E(G)$, and vertex u prevents the algorithm from adding v to set U. Let $c(u,x)$ be the customer responsible for adding u to U. Then $c(u,x)$ was processed before $c(v,y)$, and so $x \leq y < y'$.

Notice that the item $i(v,y')$ belongs to $S_{c(u,x)}$. The price of $i(v,y')$ then must be set to at least $B_{c(u,x)} = 1/2^x > 1/2^{y'} = B_{c(v,y')}$, since otherwise the customers in $\mathcal{G}(c(u,x))$ would have paid below $B_{c(u,x)}$ for item $i(v,y')$, contradicting the fact that $c(u,x) \in \mathcal{C}^*$. But then customer $c(v,y')$ must buy some item $i' \neq i(v,y')$. Assume that $i' = i(w,z)$. Then w must be a neighbor of v in G, $w \neq v$, and so $i' \in S_{c(v,y)}$ must hold. But the price of i' must be $B_{c(v,y')} = 1/2^{y'} < 1/2^y = B_{c(v,y)}$, and so the customers in $\mathcal{G}(c(v,y))$ should have paid below $B_{c(v,y)}$ for item i', contradicting the fact that $c(v,y) \in \mathcal{C}^*$.

Hardness factors: The gap between Yes-Instance and No-Instance costs is $\Delta^{1-2\epsilon}$, while the number of customers in the instance is $\tilde{m} = |V(G)| \cdot 2^{O(\Delta)}$, and the number of items is $\tilde{n} = |V(G)| \cdot \Delta$.

We first prove Theorem 1. We choose the parameter $\lambda = O(\log\log n)$ such that $\Delta = (\log n)^b$, where $b > \frac{1}{2\epsilon}$. The hardness factor then becomes $g = \Delta^{1-2\epsilon} \geq \log^{b-1} n$, while $\tilde{m} + \tilde{n} = |V(G)|2^{O(\Delta)} \leq 2^{O(\Delta \log n)} \leq 2^{\log^{b+2} n} \leq 2^{g^{1+O(\epsilon)}}$. Taking logarithm on both sides will give $g = \log^{1-O(\epsilon)}(\tilde{m}+\tilde{n})$. The deterministic reduction takes time $2^{O(\Delta)} = n^{(\log n)^{f(\epsilon)}}$ for some function f, so we have proved the first part of Theorem 1.

To prove the second part, we use the randomized version of Theorem 4, and choose $\lambda = (\log n)^b$ for some large constant b, while ϵ is set to be any small enough constant for which Theorem 4 works. In this case, we have $\Delta = 2^{O((\log n)^b)}$ and $\tilde{n} \leq |V(G)|\Delta \leq 2^{(\log n)^{b+2}}$, while $g = \Delta^{1-2\epsilon} \geq 2^{O((\log n)^b)}$. It is easy to check that $g \geq 2^{\log^{1-O(1/b)} \tilde{n}}$, as desired. Since we use the randomized reduction, the running time of the reduction is $2^{(\log n)^{O(b)}}$, and so the result holds under the assumption that $\mathsf{NP} \not\subseteq \mathsf{ZPTIME}(n^{(\log n)^{O(b)}})$.

To prove Theorem 2, we choose λ in Theorem 4 to be any sufficiently large constant. Denote by $k = \max_{c \in \mathcal{C}'} |S_c|$. Since the construction guarantees that $k \leq 2\Delta^2$, we have the hardness factor of $\Delta^{1-2\epsilon} \geq k^{1/2-\epsilon}$. In this case, the deterministic reduction only takes polynomial time, so this hardness result holds under the assumption that $\mathsf{P} \neq \mathsf{NP}$.

3 Tollbooth Pricing

In this section we prove Theorem 3. It will be useful to introduce the notion of fractional coverage and show how to convert fractional coverage to an integral one. Given an instance of Unique Coverage and a fractional solution that assigns a non-negative weight $w(S)$ to every set $S \in \mathcal{S}$, we say that an element $u \in U$ is *fractionally covered* if and only if $1/4 \leq \sum_{S: u \in S} w(S) \leq 1$. We argue that any good fractional coverage can be converted into a good integral coverage with a constant loss in the solution value. The proof of the following lemma appears in the full version of the paper.

Lemma 3. *There is an efficient randomized algorithm, that, given a fractional solution of value βn to any instance of the* Unique Coverage *problem, w.h.p. finds an integral solution of value $\Omega(\beta n)$ to the* Unique Coverage *instance.*

3.1 Construction

Let $(U, \mathcal{S})$ be an instance of Unique Coverage, where $|U| = n$ and $|\mathcal{S}| = m$. We construct an instance of Tollbooth Pricing as follows. Graph $G = (V, E)$ consists of $m+1$ vertices $v_0, \ldots, v_m$. Let $h = \lceil \log m \rceil$. For each consecutive pair (v_{i-1}, v_i) of vertices, $0 < i \leq m$, we add $h+1$ parallel edges $e^i_0, \ldots, e^i_h$. These edges are viewed as representing the set $S_i \in \mathcal{S}$. We now define the set of paths (or customers) in the graph. All paths start from v_0 and end at v_m. For each element $u \in U$, for each $j : 1 \leq j \leq h$, we have a set $\mathcal{P}(u, j)$ of 2^{h-j} paths. The budget of each path in $\mathcal{P}(u, j)$ is 2^j, the source vertex is v_0, and the sink is v_m. Each path in $\mathcal{P}(u, j)$ consists of edges $e^1_{i_1}, e^2_{i_2}, \ldots, e^m_{i_m}$, where for all $1 \leq \ell \leq m$, if $u \in S_\ell$ then $i_\ell = j$, or otherwise $i_\ell = 0$. This completes the description of the construction. Notice that the total budget is $\mathcal{B} = nh2^h$. Let $\tilde{m}$ and $\tilde{n}$ denote the number of customers (i.e. the number of paths) and items, respectively. Notice that $\tilde{m} \leq O(nm \log m)$, and $\tilde{n} \leq nh \leq O(n \log m) \leq O(n^2)$, since we can assume w.l.o.g. that $|\mathcal{S}| \leq 2^n$.

3.2 Analysis

The analysis consists of two parts. First we show that if there is a solution to Unique Coverage that satisfies a β-fraction of the elements, then there is a solution to Tollbooth Pricing of value at least $\beta \cdot \mathcal{B}$. In the second part, we show an efficient randomized algorithm, that, given any solution to Tollbooth Pricing instance G of value $\alpha \cdot \mathcal{B}$, w.h.p. finds a solution to the Unique Coverage problem that satisfies $\Omega(\alpha n)$ elements.

Lemma 4. *If there is a solution to the* Unique Coverage *instance* $(U, \mathcal{S})$ *that satisfies at least* βn*-elements, then there is a solution to the* Tollbooth Pricing *instance of value* $\beta\mathcal{B}$.

Proof. Let $\mathcal{S}' \subseteq \mathcal{S}$ be a solution to the Unique Coverage problem, and let $U' \subseteq U$ be the set of elements uniquely covered by $\mathcal{S}'$, $|U'| \geq \beta n$. For each $S_i \in \mathcal{S}'$, for each $j : 1 \leq j \leq h$, we set the price of the edge e^i_j to 2^j. The prices of all other edges (including the edges e^i_0 for all i) are set to 0. For each $u \in U'$ and $j : 1 \leq j \leq h$, we consider the revenue collected from the paths in $\mathcal{P}(u, j)$. Let S_i be the set that uniquely covers u in the solution. Then for each path in $\mathcal{P}(u, j)$, exactly one edge e^i_j on the path has a non-zero price. This price is 2^j - the same as the budget of the path, while all other edges have price 0. Therefore, each such path contributes 2^j to the solution value, and the total contribution of the paths in $\mathcal{P}(u, j)$ is 2^h. This implies the lemma.

Lemma 5. *There is an efficient randomized algorithm, that, given any solution to the* Tollbooth Pricing *instance* G *of value* $\alpha\mathcal{B}$*, w.h.p. finds a solution to the* Unique Coverage *instance* $(U, \mathcal{S})$ *that satisfies* $\Omega(\alpha n)$ *of the elements.*

Proof. Let $p^* : E \to \mathbb{R}_{\geq 0}$ be any solution of value $\alpha\mathcal{B}$ to the Tollbooth Pricing problem. Let $\mathcal{P}_1$ be the set of paths, such that each $P \in \mathcal{P}_1$ contributes at least half of its budget to the solution. Our first observation is that the profit collected from the paths in $\mathcal{P}_1$ must be at least $\alpha\mathcal{B}/2$ (otherwise, we can multiply the price of each edge by a factor of two and get a better solution). From now on, we will only focus on paths in $\mathcal{P}_1$ and we will discard all other paths. We say that a path $P \in \mathcal{P}_1$ is of type 1 if at least half the cost it pays goes to edges in set $E_0 = \left\{e^i_0 : 1 \leq i \leq m\right\}$, and it is of type 2 otherwise. Let $\mathcal{P}'$ and $\mathcal{P}''$ denote the set of paths of type 1 and 2 respectively. We distinguish between two cases.

Case 1: Paths of type 1 contribute at least $\alpha\mathcal{B}/4$ to the solution value. We claim that in this case the solution value is at most $O(\mathcal{B}/\log m)$, and therefore it is sufficient to find a solution to Unique Coverage instance that satisfies a $\Omega(1/\log m)$-fraction of the elements. We then show an algorithm to find such a solution.

Indeed, consider some element $u \in U$. Recall that, for all j, every path in the sets $\mathcal{P}(u, j)$ traverses all edges in the set $E_0(u) = \left\{e^i_0 : u \notin S_i\right\}$, and these are the only edges from E_0 traversed by these paths. Let $C_u = \sum_{e \in E_0(u)} p^*(e)$ be the total price of these edges. A path $P \in \mathcal{P}(u,j)$ can belong to $\mathcal{P}'$ only if $2^j/4 \leq C_u \leq 2^j$. This means that there are at most 3 values of $j : 1 \leq j \leq h$ for which $\mathcal{P}(u,j) \cap \mathcal{P}' \neq \emptyset$, so for each $u \in U$, the paths in set $\bigcup_{j=1}^{h} \mathcal{P}(u,j)$ only contribute at most an $O(1/h) = O(1/\log m)$-fraction of their total budget to the solution. Therefore, the solution value is at most $O(\mathcal{B}/h) = O(\mathcal{B}/\log m)$. Now we show an algorithm for the Unique Coverage problem instance that satisfies an $\Omega(1/\log m)$-fraction of the elements. From Lemma 3, it is enough to construct a fractional solution of value $\Omega(n/\log m)$. For each element $u \in U$, let $\delta(u)$ be the number of sets in $\mathcal{S}$ to which element u belongs. We partition the elements into $h = \lceil \log m \rceil$ classes $C_1, \ldots, C_h$ where class C_j contains elements u with $2^j \leq \delta(u) \leq 2^{j+1}$. Let j^* be the class containing the maximum number of elements, so $|C_{j^*}| \geq \Omega(n/\log m)$. We set the weight of every set S to be $w(S) = 1/2^{j^*+1}$. This ensures that all elements in C_{j^*} are fractionally covered. Applying Lemma 3, we obtain an integral solution of value $\Omega(n/\log m)$.

Case 2: Assume now that the paths in $\mathcal{P}''$ contribute at least $\alpha\mathcal{B}/4$ to the solution value. Let r'' denote the total revenue collected from these paths by edges in $E_1 = E \setminus E_0$. Then we have that $r'' \geq \Omega(\alpha\mathcal{B}) = \Omega(\alpha n h 2^h)$. Notice that by the definition of set $\mathcal{P}''$, each path $P \in \mathcal{P}''$ pays at least $1/4$ of its budget for the edges in set E_1 that lie on path P.

We now partition the paths in $\mathcal{P}''$ into sets $\mathcal{P}''_1, \ldots, \mathcal{P}''_h$ where set $\mathcal{P}''_j$ contains all type-2 paths whose budget is 2^j. Let j^* be the index for which the profit contributed by the paths in $\mathcal{P}''_{j^*}$ is maximized. This profit is at least $\alpha n 2^h$.

We say that element u is good if $2^{j^*}/4 \leq \sum_{i: u \in S_i} p^*(e^i_{j^*}) \leq 2^{j^*}$. From the above arguments, for each path $P \in \mathcal{P}''$, if $P \in \mathcal{P}(u, j^*)$, then the corresponding element u must be good. Moreover, if u is good, then all paths in $\mathcal{P}(u, j^*)$ belong to $\mathcal{P}''_{j^*}$. Therefore, at least $\Omega(\alpha n)$ of the elements in U must be good. We now define a fractional solution to the Unique Coverage problem, where every the weight of every set $S_i \in \mathcal{S}$ is set to $w(S_i) = p(e^i_{j^*})/2^{j^*}$. Notice that all good elements are fractionally covered, thus giving us a fractional solution where $\Omega(\alpha n)$ elements are satisfied. We finally invoke Lemma 3 to complete the proof.

Acknowledgement: The first author thanks Khaled Elbassioni for introducing him to pricing problems and for explaining the differences between various pricing models. He also thanks Danupon Nanongkai and Khaled Elbassioni for useful discussions about $\mathsf{UDP_{MIN}}$ and SMP.

References

1. Gagan Aggarwal, Tomás Feder, Rajeev Motwani, and An Zhu. Algorithms for Multi-product Pricing. In *ICALP*, pages 72–83, 2004.

2. Maria-Florina Balcan and Avrim Blum. Approximation algorithms and online mechanisms for item pricing. *Theory of Computing*, 3(1):179–195, 2007.
3. Patrick Briest. Uniform budgets and the envy-free pricing problem. In *ICALP (1)*, pages 808–819, 2008.
4. Patrick Briest and Piotr Krysta. Buying cheap is expensive: hardness of non-parametric multi-product pricing. In *SODA*, pages 716–725, 2007.
5. Ning Chen and Xiaotie Deng. Envy-free pricing in multi-item markets. In *ICALP (2)*, pages 418–429, 2010.
6. Ning Chen, Arpita Ghosh, and Sergei Vassilvitskii. Optimal envy-free pricing with metric substitutability. *SIAM J. Comput.*, 40(3):623–645, 2011.
7. Maurice Cheung and Chaitanya Swamy. Approximation algorithms for single-minded envy-free profit-maximization problems with limited supply. In *FOCS*, pages 35–44. IEEE Computer Society, 2008.
8. Erik D. Demaine, Uriel Feige, MohammadTaghi Hajiaghayi, and Mohammad R. Salavatipour. Combination can be hard: Approximability of the unique coverage problem. *SIAM J. Comput.*, 38(4):1464–1483, 2008.
9. Khaled M. Elbassioni, Rajiv Raman, Saurabh Ray, and René Sitters. On profit-maximizing pricing for the highway and tollbooth problems. In *SAGT*, pages 275–286, 2009.
10. Khaled M. Elbassioni, René Sitters, and Yan Zhang. A quasi-ptas for profit-maximizing pricing on line graphs. In *ESA*, pages 451–462, 2007.
11. Uriel Feige. Relations between average case complexity and approximation complexity. In *STOC*, pages 534–543, 2002.
12. Iftah Gamzu and Danny Segev. A sublogarithmic approximation for highway and tollbooth pricing. In *ICALP (1)*, pages 582–593, 2010.
13. Fabrizio Grandoni and Thomas Rothvoß. Pricing on paths: A ptas for the highway problem. In *SODA*, pages 675–684, 2011.
14. Venkatesan Guruswami, Jason D. Hartline, Anna R. Karlin, David Kempe, Claire Kenyon, and Frank McSherry. On profit-maximizing envy-free pricing. In *SODA*, pages 1164–1173, 2005.
15. Jason D. Hartline and Vladlen Koltun. Near-optimal pricing in near-linear time. In *In 9th Workshop on Algorithms and Data Structures*, pages 422–431, 2005.
16. Rohit Khandekar, Tracy Kimbrel, Konstantin Makarychev, and Maxim Sviridenko. On hardness of pricing items for single-minded bidders. In *APPROX*, pages 202–216, 2009.
17. Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*, STOC '02, pages 767–775, New York, NY, USA, 2002. ACM.
18. Paat Rusmevichientong. A non-parametric approach to multi-product pricing: Theory and application. *Ph. D. thesis, Stanford University*, 2003.
19. Paat Rusmevichientong, Benjamin Van Roy, and Peter W. Glynn. A nonparametric approach to multiproduct pricing. *Oper. Res.*, 54:82–98, January 2006.
20. Luca Trevisan. Non-approximability results for optimization problems on bounded degree instances. In *STOC*, pages 453–461, 2001.