

Pricing Bridges to Cross a River

Mustapha Bouhtou,¹ Alexander Grigoriev,² Stan van Hoesel,² Anton F. van der Kraaij,²
Frits C. R. Spieksma,³ Marc Uetz²

¹ France Télécom R&D, 39-40 rue du Général Leclerc, F-92131 Issy-Les-Moulineaux, France

² Maastricht University, Quantitative Economics, P.O. Box 616, NL-6200 MD Maastricht, The Netherlands

³ Katholieke Universiteit Leuven, Applied Economics, Naamsestraat 69, B-3000 Leuven, Belgium

Received 21 February 2005; revised 26 September 2006; accepted 4 November 2006

DOI 10.1002/nav.20216

Published online 2 February 2007 in Wiley InterScience (www.interscience.wiley.com).

Abstract: We consider a pricing problem in directed, uncapacitated networks. Tariffs must be defined by an operator, the leader, for a subset of m arcs, the tariff arcs. Costs of all other arcs in the network are assumed to be given. There are n clients, the followers, and after the tariffs have been determined, the clients route their demands independent of each other on paths with minimal total cost. The problem is to find tariffs that maximize the operator's revenue. Motivated by applications in telecommunication networks, we consider a restricted version of this problem, assuming that each client utilizes at most one of the operator's tariff arcs. The problem is equivalent to pricing bridges that clients can use in order to cross a river. We prove that this problem is APX-hard. Moreover, we analyze the effect of uniform pricing, proving that it yields both an m -approximation and a $(1 + \ln D)$ -approximation. Here, D is upper bounded by the total demand of all clients. In addition, we consider the problem under the additional restriction that the operator must not reject any of the clients. We prove that this problem does not admit approximation algorithms with any reasonable performance guarantee, unless $P = NP$, and we prove the existence of an n -approximation algorithm. © 2007 Wiley Periodicals, Inc. *Naval Research Logistics* 54: 411–420, 2007

Keywords: telecommunication networks; pricing; Stackelberg game; complexity; approximation

1. INTRODUCTION

The pricing problem that we study is a Stackelberg game that involves two non-cooperative groups, an operator that sets tariffs, the *leader* of the Stackelberg game, and clients that must pay these tariffs, the *followers* of the Stackelberg game. More precisely, we assume that a network is given, and a subset of the arcs, the tariff arcs, are owned by an operator. The operator can determine tariffs on these tariff arcs, while the costs for utilizing all other arcs are assumed to be given. Each client wishes to route a certain demand on a path connecting two vertices. We assume that after the tariffs have been announced, each client selfishly selects a path with minimum total cost to route her demand. Thus, before the clients select their paths, the operator must set the tariffs, which she does in order to maximize her total revenue. In order to avoid non-boundedness, we assume that clients always have the alternative of routing on a path without using any of the operators arcs.

Note that this problem is different in two aspects from the network congestion problems studied recently, e.g., by Roughgarden and Tardos [12] and Cole, Dodis, and Roughgarden [4, 5]. First, we assume that there is no congestion, hence, the clients do not influence each other. They choose minimum cost paths to route their demands, independent of each other. A game theoretic setting is only present due to the fact that there exists an operator trying to maximize the total revenue, and the clients effectively minimize the total revenue by choosing minimum cost paths. Second, due to the fact that the pricing takes place before the clients choose their paths, we are faced with a Stackelberg game. Note that the second phase of this Stackelberg game is indeed trivial, since the clients are independent of each other.

1.1. Model

In order to clarify the relation to previous work, we first formulate the general network pricing problem and then discuss the restricted version considered in this paper.

Correspondence to: M. Uetz (m.uetz@ke.unimaas.nl)

An instance of the general network pricing problem is a directed graph $G = (N, A)$, where the arc set A is partitioned into a set of m *tariff arcs* $T \subseteq A$ and a set of *fixed cost arcs* $F = A \setminus T$. There are n clients (or commodities) $k \in \{1, \dots, n\}$, and each client k has a demand d_k that must be routed from source node s_k to target node t_k .¹ Without loss of generality, we assume that all demand values d_k are scaled to be integral. The tariff for the utilization of any tariff arc $a \in T$ must be determined by the operator; it is denoted by τ_a . The tariff for the utilization of any fixed cost arc is assumed to be given for all fixed cost arcs. The clients route their demands from source to destination through a path with minimal total cost, where the total per unit cost of a path is defined as the sum of the tariffs and fixed costs on the arcs of the path. Whenever the client has a choice among multiple paths with the same total cost but with different revenues for the operator, we assume that the client takes the path that is most profitable to the operator. (This can always be achieved with arbitrary precision by reducing tariffs by some small value ε .) We assume that an (s_k, t_k) -path exists consisting only of fixed cost arcs for every client $k \in \{1, \dots, n\}$, since the problem is otherwise unbounded. Without going into further details, we mention that this problem is a classical Stackelberg game that can be modeled as a bilinear bilevel program [10].

We next describe a simple transformation of the given graph G that allows one to restrict to very specific graphs (although probably losing certain graph properties, such as planarity). When we replace all shortest paths that only consist of fixed cost arcs by direct arcs and possibly introduce additional dummy arcs with zero or infinite cost, respectively, we obtain a shortest path graph model as described by Bouhtou et al. [3]. After this transformation, we can assume that all tariff arcs are pairwise disjoint, and there exists a direct arc from the source node s_k to the tail node of any tariff arc a and a direct arc from the head node of any tariff arc a to any target node t_k . Moreover, there exists a fixed cost arc (s_k, t_k) for all clients $k = 1, \dots, n$, and the fixed cost for that arc, which we denote by u_k , represents the cheapest possible (s_k, t_k) -path (in the original graph) without using any of the tariff arcs. In other words, u_k represents the highest acceptable total per unit price for client k .

The additional assumption in the problem considered in this paper, to which we refer as the *river tariff pricing problem* (RTP), is the following: Independent of the tariffs, we assume that any client routes her demand on a path that includes at most one tariff arc. In Section 1.2, we discuss practical applications for this model, motivated by problems




Figure 1. River tariff pricing problem (RTP) with $n = 2$ clients and $m = 3$ tariff arcs.

in telecommunication networks. In the shortest path graph model, this restriction is equivalent to the deletion of any backward-arc that might exist between head nodes of tariff arcs and tail nodes of other tariff arcs. Figure 1 illustrates the shortest path graph model of an instance of the river tariff pricing problem with three tariff arcs and two clients. The tariff arcs $a_i, i \in \{1, 2, 3\}$, are the dashed arcs in the network. We may also assume without loss of generality that all fixed cost arcs incident with the target nodes t_k have zero cost, because otherwise we can just add their costs to the fixed cost arcs incident with source nodes s_k . Therefore, let us denote by c_{ka} the cost of the arc that connects customer k to tariff arc a . The value $u_k - c_{ka}$ then represents client k 's highest acceptable tariff for utilizing tariff arc a . It can as well be interpreted as client k 's *valuation* for tariff arc a . Note that the only difference to the general network pricing problem described previously is the non-existence of backward arcs in the shortest path graph model.

To summarize, the parameters that define an instance of a river tariff pricing problem are the number of tariff arcs m , the number of clients n , their demand values $d_k, k \in \{1, \dots, n\}$, and the costs for fixed cost arcs. We have c_{ka} as the cost of the fixed cost arcs that connect customers k to tariff arcs a , and u_k as the cost of arc (s_k, t_k) , the highest acceptable cost for client k . Due to the fact that any path taken by a client involves exactly one fixed cost arc with non-zero cost, we may assume without loss of generality that the costs c_{ka} of these fixed cost arcs are integral. Moreover, due to the integrality of the costs of the fixed cost arcs, it follows that any solution utilizing non-integral tariffs can be straightforwardly improved. Note that this might not be true for the general network pricing problem, where a path chosen by a client can consist of more than one tariff arc.

1.2. Applications

The present study of the river tariff pricing problem as described above is motivated by practical interest from the

¹ Note that we abuse standard notation from Graph Theory, since m denotes the number of tariff arcs, and n denotes the number of clients in the given digraph $G = (N, A)$. This is because the actual number of nodes and arcs of G are of minor interest for the present paper.




Figure 2. International interconnections market.

telecommunications industry; it was carried out within a joint research project initiated by France Télécom and Maastricht University.

An instance of the problem arises, for example, when considering the international interconnections market, where several operators offer telecom connections to a particular country. Focusing on the market for entering a particular country—France, in our example—France Télécom asks what tariffs it should use for its proprietary connections into the country such as to maximize revenue. There are several other operators competing with France Télécom by offering similar services. For a schematic illustration see Figure 2. Here, the dashed lines depict the connections offered by the operator (tariff arcs) and the solid lines depict the connections offered by the competitors (fixed cost arcs). It is a common practice that once the data of a client enter the local network of the destination country these data will be transmitted to the destination point without leaving and reentering the local network. So, given the set of clients willing to transmit their data to France and the prices of the competitors, the operator would like to determine prices for the tariff arcs such as to maximize her total revenue, therefore facing an instance of the river tariff pricing problem.

Another telecom application for the problem at hand is point-to-point markets, where an operator is offering bandwidth capacity between two points A and B . Other operators are active in this market as well. Their prices for bandwidth capacity are known. Clients can choose between different levels of Quality of Service (QoS) from each operator, and clients have a preference for the QoS-levels. We can model this problem as an instance of the river tariff pricing problem, too. Figure 3 shows a small example with two customers, represented by two commodities (s_1, t_1) and (s_2, t_2) . The operator has three QoS levels, represented by the subnetwork between

the nodes q_{is} and q_{it} , where $i \in \{1, 2, 3\}$. In this example, customer (s_1, t_1) is interested in two QoS levels, namely QoS1 and QoS2, whereas customer (s_2, t_2) is interested in QoS2 and QoS3. The preference of each customer k with regard to each QoS level is determined by the cost of the edge from the source s_k to the node q_{is} , $i \in \{1, 2, 3\}$, with smaller cost indicating a higher preference for the QoS level. The prices of other operators for the same QoS level is given by the cost on the (fixed cost) arcs $(q_{is}, q_{it}), i \in \{1, 2, 3\}$. The revenue for the operator for each QoS level $i, i \in \{1, 2, 3\}$, is then determined by setting appropriate tariffs on the tariff arcs (dashed arcs).

1.3. Related Work

A formulation of the general network pricing problem is the bilinear bilevel program described by Labbé, Marcotte, and Savard [10]. They show, among other things, that already the problem with a single client is strongly NP-hard, given that also non-negative tariffs are allowed. Roch, Marcotte, and Savard [11] prove NP-hardness for the same problem with non-negative tariffs, and propose a polynomial time $(1 + 1/2 \ln m)$ -approximation algorithm for the problem with a single client, where m is the number of tariff arcs. This result implies also a $O(n \log m)$ -approximation for the case of multiple clients. In the present paper, we consider the problem restricted to the case where each client utilizes at most one tariff arc, and we consider the case of multiple clients.

In fact, the problem at hand can equivalently be interpreted as a pricing problem for multiple products, where the tariff arcs $a \in T$ correspond to different products and each client k is interested in buying d_k units of one product. Since we consider uncapacitated networks, products are available in unlimited amount (e.g., bulk or digital goods). Whenever there is an arc between a client k and a tariff arc a in the river tariff pricing problem, the interpretation is that client k is interested in buying product a . If she decides to buy product a , she incurs a per-unit shipment cost of c_{ka} , in addition to the per-unit cost of τ_a for product a . The fixed cost u_k of the fixed cost arc (s_k, t_k) is simply interpreted as the maximum

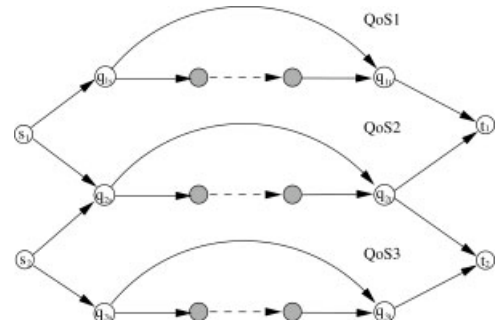


Figure 3. Point-to-point markets.

total (per unit) price a client k is willing to pay to purchase any of the products. In other words, $u_k - c_{ka}$ represents client k 's valuation for product a .

After this discussion, we can exhibit a close relation of the river tariff pricing problem considered in this paper to other papers that address multi-product pricing problems. Recently, two groups of researchers, independently of each other, reported several results for such problems. Aggarwal et al. [1], among other things, consider a multi-product pricing problem where any client k has different budgets b_{ka} for different products a , which are available in unlimited amount. The operator must determine prices for the products in order to maximize the total revenue, under the assumption that a client buys (one unit of) the cheapest product among the products she can afford. Aggarwal et al. [1] prove APX-hardness of this problem, together with a $(1 + \ln n)$ -approximation algorithm. Note that, despite the obvious similarities, the multi-product pricing problem is conceptually different from the river tariff pricing problem considered in this paper. In the river tariff pricing problem, clients choose the product with minimum total per unit cost, also taking into account the shipment costs c_{ka} , rather than the cheapest product among all affordable products.

Guruswami et al. [8] consider a profit-maximizing envy-free pricing problem. Clients have different valuations for different products, and each product is available in limited amount. The operator must determine prices for the products and allocate the products to clients such that, again, total revenue is maximized, and given the pricing, no client would prefer to be assigned a different product. Here, the clients measure their preferences in terms of the difference between their valuation and the purchase price. If the price is higher than the clients' valuation, then the client does not purchase the product. In fact, the profit-maximizing envy-free pricing problem with unlimited supply of products is equivalent to the river tariff pricing problem considered here. Guruswami et al. [8] independently prove APX-hardness of the problem and derive a $(2 \ln n)$ -approximation algorithm for the case of unit demand of clients and with limited supply of products.

1.4. Our Results

In this paper, we derive several results concerning complexity and approximability of the river tariff pricing problem. In Section 2.1, by a reduction from the MAX-2-SAT-3 problem, we show that the river tariff pricing problem is APX-hard, even if each client is connected to at most two tariff arcs. Hence, the problem does not admit a polynomial time approximation scheme, unless $P = NP$. This result coincides with the APX-hardness result of Guruswami et al. [8], obtained independently. The quality of uniform tariff pricing policies, where all arcs are priced with the same tariff, is analyzed in Section 2.2. The problem to find an optimal uniform tariff is

solvable in polynomial time by simple enumeration. We show that uniform tariff pricing is an m -approximation, and this is tight. Using a simple geometric argument, we also show that uniform tariff pricing is a $(1 + \ln D)$ -approximation, which is tight up to a constant factor. Here, D is the total demand that is served by the operator in an optimal solution, which is clearly upper bounded by the total demand. Hence, whenever the clients have unit demand, our result yields a $(1 + \ln n)$ -approximation. In Section 2.3, we empirically analyze and discuss the quality of uniform tariff pricing policies using instances of international interconnection markets provided by France Télécom.

In Section 3, we consider another variant of the problem, namely where the operator is not allowed to reject any client. We refer to that variant as *all-service* river tariff pricing problem, or all-service RTP. Note that rejecting clients might increase the total revenue, since some clients might exist that can only be served at a low price, while others would be willing to pay much more. We show, by a reduction from the INDEPENDENT SET problem, that the all-service RTP problem does not allow approximation to within a factor $O(m^{1-\epsilon})$ or $O(n^{1/2-\epsilon})$, unless $P = NP$. (Recall that m is the number of tariff arcs and n is the number of clients.) On the positive side, we can show that the problem admits an n -approximation.

2. RIVER TARIFF PRICING: COMPLEXITY AND APPROXIMATION

We first discuss the computational complexity of the river tariff pricing problem. Subsequently, we derive bounds on the quality of uniform tariff pricing policies, where all tariffs are required to be identical, and finally, we briefly discuss these results on the basis of problem instances from France Télécom.

2.1. Complexity

In a conference version of this paper [7], we proved that the river tariff pricing problem is (strongly) NP-hard by a reduction from 3-SATISFIABILITY. Here we present a modification of that reduction that yields a stronger result, namely APX-hardness of the river tariff pricing problem. Thereby, we can exclude the existence of a polynomial time approximation scheme, unless $P = NP$. Note that this result coincides with the APX-hardness of the profit maximization problem considered by Guruswami et al. [8]. Also note that Roch, Marcotte, and Savard [11] show that the general network pricing problem is strongly NP-hard, even when restricted to a single client. Their reduction works for tariff pricing problems where paths are allowed to use (and indeed, must use) several tariff arcs. In a certain sense, the problem that we consider here is "dual" to theirs, as we restrict to a single tariff arc, but allow for more than one client.

THEOREM 1 (see also Guruswami et al. [8]): The river tariff pricing problem is APX-hard, even when each client is connected to at most two tariff arcs and if the costs of the fixed cost arcs are restricted to be 0, 1, or 2.

PROOF: We use an approximation preserving reduction from MAX-2-SAT-3: Given a set of boolean variables $X = \{x_1, \dots, x_n\}$ and a collection C of clauses over X such that each clause consists of at most two variables and each variable occurs in at most 3 clauses, the question is to find a truth assignment that satisfies the maximum number of clauses. This problem is known to be APX-hard, see, e.g., [2].

For each variable x_i , $i \in \{1, \dots, n\}$ of the MAX-2-SAT-3 instance, we construct a constant-size subnetwork as shown in Figure 4. Each of these subnetworks has three clients with unit demand. For each variable x_i , let the origin–destination pairs of these clients be denoted by $\{s_{ij}, t_{ij}\}$, $j \in \{1, 2, 3\}$. Moreover, each subnetwork has two tariff arcs, one denoted by a_i representing the truth assignment $x_i = 1$ and one denoted by \bar{a}_i representing $x_i = 0$.

An upper bound on the cost of routing commodities 1 and 3 is given by fixed cost arcs (s_{i1}, t_{i1}) and (s_{i3}, t_{i3}) , each with cost 2. For commodity 2, the upper bound on the cost is given by a fixed cost arc (s_{i2}, t_{i2}) , with cost 1.

Next, for each clause C_k , $k \in \{1, \dots, m\}$, we create a *clause-commodity* k with origin destination pairs $\{s_k, t_k\}$, with unit demand. Whenever a variable x_i (\bar{x}_i , respectively) appears in clause C_k , we connect s_k to s_{i1} (s_{i3} , respectively), and t_{i1} (t_{i3} , respectively) to t_k , using arcs of zero cost. In addition, we introduce a fixed cost arc (s_k, t_k) with cost 1, defining an upper bound of 1 for the cost of routing clause-commodity k . The so-defined instance of the river tariff pricing problem has $2n$ tariff arcs, $3n + m$ commodities (or clients) and at most $7m + 11n$ fixed cost arcs, hence, the transformation is indeed polynomial.

We claim that an optimal solution of this instance of the river tariff pricing problem yields a revenue of $4n + r$, where r denotes the maximum number of satisfied clauses in the




Figure 4. Subnetwork for variable x_i , $i \in \{1, \dots, n\}$.

MAX-2-SAT-3 instance. Indeed, when given a truth assignment specifying x_i , we set the tariff of arc a_i (\bar{a}_i) to 1 and the tariff of arc \bar{a}_i (a_i) to 2 if x_i is true (if x_i is false). By observing that the maximal revenue for each subnetwork equals 4 and that we obtain 1 for each clause-commodity corresponding to a satisfied clause (by routing its demand via the tariff arc with a tariff of 1), we find a solution with value $4n + r$.

Conversely, we now argue that in an optimal solution each subnetwork contains one tariff arc with tariff 1 and the other tariff arc with tariff 2. Indeed, note that it does not make sense to use any other values for the tariffs. Further, if both tariffs equal 2, we can decrease one tariff to 1 without lowering the revenue. Moreover, if in some subnetwork both tariffs equal 1, there is one tariff arc through which at most one clause commodity is routed (this follows from the fact that each variable occurs at most three times in all clauses). When raising this tariff from 1 to 2, we are compensated for this loss of at most 1 by the corresponding commodity in the subnetwork, which brings an additional revenue of 1.

Thus, we conclude that the tariff arcs with value 1 define a valid truth assignment, and the corresponding revenue is at most $4n + r$. Hence, an inapproximability gap for MAX-2-SAT-3 translates to an inapproximability gap for the river tariff pricing problem. \square

Observe that we have delineated a borderline between easy and hard instances of the river tariff pricing problem, since if each client is connected to at most one arc the problem is trivial, while in the described reduction each client is connected to at most two tariff arcs. Guruswami et al. [8] independently presented another reduction from a restricted version of the VERTEX COVER problem, yielding exactly the same conclusions.

2.2. The Quality of Uniform Tariff Pricing

The uniform tariff pricing problem (UTP) is the same as the general tariff pricing problem, with the additional restriction that all tariffs are required to be identical. An optimal uniform tariff can be found in time $O(nm)$ by just enumerating all tariffs of the form $u_k - c_{ka}$, $k = 1, \dots, n$, $a = 1, \dots, m$. Clearly, any uniform price other than that cannot be optimal, since the tariffs could be increased by some positive amount without losing any client.

We next analyze the loss that can be experienced by adopting such a uniform tariff pricing policy for the river tariff pricing problem. Therefore, denote by Π^{UTP} the revenue for an optimal uniform tariff pricing and by Π^{RTP} the revenue for an optimal non-uniform tariff pricing. By definition, $\Pi^{\text{UTP}} \leq \Pi^{\text{RTP}}$.

LEMMA 1: If an optimal solution to the river tariff pricing problem with revenue Π^{RTP} utilizes at most r different tariffs,




Figure 5. Staircase function $f(x)$ with inscribed maximal rectangle.

then $\Pi^{\text{UTP}} \geq \Pi^{\text{RTP}}/r$, where Π^{UTP} is the optimal revenue for the case with the uniform tariff restriction.

PROOF: Consider an optimal solution to the river tariff pricing problem with non-uniform tariffs $\tau_1 \leq \dots \leq \tau_m$, and let D_i be the total demand on an arc a_i with tariff τ_i , $i \in \{1, \dots, m\}$. By $D = \sum_{k=1}^n D_k$ we denote the total demand served by the operator. Then the revenue created by this solution is the area under the following “staircase” function $f : [0, D] \rightarrow [0, \infty[$, depicted in Figure 5.

$$f(x) = \tau_i \quad \text{for all } x \text{ with } \sum_{j < i} D_j \leq x < \sum_{j \leq i} D_j, \quad i \in \{1, \dots, m\}. \quad (1)$$

Consider any of the rectangles inscribed under the graph of function $f(x)$, with area $T_i := \tau_i \cdot \sum_{j \geq i} D_j$. Then it holds that $\Pi^{\text{UTP}} \geq T_i$ for all $i \in \{1, \dots, m\}$, since the area of any such rectangle is a lower bound for the revenue yielded by the optimal uniform tariff Π^{UTP} . (Note that this does not hold for the general network pricing problem.) Hence, if only r different tariffs are utilized, we consider the r (inclusion-) maximal rectangles under function f , say T_{i_1}, \dots, T_{i_r} , and get $r \cdot \Pi^{\text{UTP}} \geq \sum_{j=1}^r T_{i_j} \geq \Pi^{\text{RTP}}$. \square

Since $r \leq m$, Lemma 1 yields the following theorem. Tightness of the result will be shown below, using Example 1.

THEOREM 2: Uniform tariff pricing is an m -approximation for the river tariff pricing problem.

We next derive another bound on the quality of uniform tariff pricing policies, developing further the same geometric argument.

THEOREM 3: Uniform tariff pricing is a $(1 + \ln D)$ -approximation for the river tariff pricing problem, where

$D \leq \sum_{k=1}^n d_k$ is the total demand that is served by the operator in an optimal solution.

PROOF: Consider an optimal non-uniform tariff pricing and recall the definition of the corresponding staircase function f in (1), as well as the inscribed rectangles, with areas $T_i = \tau_i \cdot \sum_{j \geq i} D_j$. Let ℓ be the index of the maximal area rectangle among all T_i , with area T_ℓ . Let $x_\ell := \sum_{j \geq \ell} D_j = T_\ell / \tau_\ell$. Moreover, denote by τ_{\max} the maximal tariff utilized in that optimal solution. We show

$$\Pi^{\text{UTP}} \geq \frac{\Pi^{\text{RTP}}}{1 + \ln(D\tau_{\max}/T_\ell)}. \quad (2)$$

Then, the result follows from (2) because, by definition, $T_\ell \geq \tau_{\max}$. To prove (2), let

$$g(x) := \frac{T_\ell}{D - x} \quad \text{for } x \in [0, D). \quad (3)$$

We claim that $g(x) \geq f(x)$ for $x \in [0, D)$. To see this, take any x with $\sum_{j < i} D_j \leq x < \sum_{j \leq i} D_j$, then $f(x) = \tau_i$ by definition. Furthermore, the following holds

$$g(x) = \frac{T_\ell}{D - x} \geq \frac{T_\ell}{D - \sum_{j < i} D_j} = \frac{T_\ell}{\sum_{j \geq i} D_j} = \frac{T_\ell}{T_i / \tau_i} \geq \tau_i = f(x),$$

where the first inequality follows by choice of x , and the last follows because ℓ is the index of the largest rectangle.

Hence, the area under the function $g(x)$, $0 \leq x \leq D$, is a valid upper bound for the area under the staircase function, which equals Π^{RTP} , see Figure 6. To compute the area under the function $g(x)$, $0 \leq x \leq D$, we partition it into three parts, namely the rectangle T_ℓ itself, the area under $g(x)$ on the domain $x \in [0, D - x_\ell]$, and the area to the right of $g(x)$




Figure 6. Illustration for the proof of Theorem 3.

for $g(x) \in [\tau_\ell, \tau_{\max}]$. The latter is the integral of the function $D - g^{-1}(\tau) = T_\ell/\tau$ on the domain $[\tau_\ell, \tau_{\max}]$. We thus obtain the following.

$$\begin{aligned} \Pi^{\text{RTP}} &\leq T_\ell + \int_0^{D-x_\ell} \frac{T_\ell}{D-x} dx + \int_{\tau_\ell}^{\tau_{\max}} \frac{T_\ell}{\tau} d\tau \\ &= T_\ell [1 + \ln D + \ln \tau_{\max} - \ln \tau_\ell - \ln x_\ell] \\ &= T_\ell [1 + \ln(D\tau_{\max}/T_\ell)]. \end{aligned}$$

Then, (2) follows from the above because $T_\ell \leq \Pi^{\text{UTP}}$. \square

In the case of unit demands of the clients, that is, if $d_k = 1$ for all clients $k = 1, \dots, n$, we obtain the following.

COROLLARY 1: Whenever clients have unit demands, uniform tariff pricing is a $(1 + \ln n)$ -approximation for the river tariff pricing problem.

Finally, let us show tightness of the bounds in Theorems 2 and 3.

EXAMPLE 1: Given $n = m$ clients and m tariff arcs. Every client is operating her own subnetwork with one tariff arc, thus, the entire network consists of m disjoint subnetworks and each of them contains one client and one tariff arc. Fix $b > 1$ and let the demand of client k in subnetwork k be given by $d_k = b^k - b^{k-1}$, $k \in \{1, \dots, m\}$. This way, the total demand equals $b^m - 1$. Moreover, the maximal revenue for subnetwork k is limited by a fixed cost arc (s_k, t_k) , with cost $u_k = b^{2m-k}$. Hence, the maximal tariff τ_{\max} equals b^{2m-1} . See Figure 7 for an example with $n = m = 4$. \square

In the optimal solution, the tariff for each subnetwork k is set to its maximal value, b^{2m-k} . Each subnetwork therefore contributes a revenue of $b^{2m} - b^{2m-1}$, and $\Pi^{\text{RTP}} = m(b^{2m} - b^{2m-1})$. The optimal uniform pricing consists in setting the tariff on all tariff arcs to b^m . This way, every unit of demand creates a profit of b^m , yielding a total revenue of $b^{2m} - b^m$. Other uniform tariffs would be values b^{2m-k} ,

$k \in \{1, \dots, m-1\}$. This yields a total revenue of $b^{2m} - b^{2m-k}$, which is less. Therefore, we obtain

$$\begin{aligned} \Pi^{\text{UTP}}/\Pi^{\text{RTP}} &= \frac{b^{2m} - b^m}{m(b^{2m} - b^{2m-1})} \\ &\leq \frac{b^{2m}}{m(b^{2m} - b^{2m-1})} = \frac{1}{m} \cdot \frac{b}{b-1}. \end{aligned}$$

Now, observe that in the optimal solution m different tariffs are utilized. Lemma 1 (Theorem 2, respectively) suggests that uniform pricing provides an m -approximation. Example 1 proves that this is the best possible, since b can be chosen arbitrarily large.

Moreover, Theorem 3 suggests that uniform pricing is a $(1 + \ln D)$ -approximation. In Example 1, we have $D = (b^m - 1)$ and thus $(1 + \ln D) = 1 + \ln(b^m - 1) \leq 1 + m \ln b$. Hence, Theorem 3 yields that uniform pricing is a $O(m)$ -approximation on this example. The same Example 1 shows that $O(m)$ is indeed best possible. The above discussion leads to the following result:

THEOREM 4: For uniform tariff pricing, the performance bound of Theorem 2 is the best possible, and the performance bound of Theorem 3 is the best possible up to a constant factor.

2.3. Discussion

Note that claim (2) in the proof of Theorem 3 confirms the following intuition: If the staircase function $f(x)$ approximates the straight line $x \mapsto (\tau_{\max}/D) \cdot x$, geometric intuition suggests that uniform tariff pricing yields a 2-approximation, since the size of the largest rectangle inscribed under $f(x)$ would be of exactly half of the area under $f(x)$. In that case, T_ℓ approximately equals $D\tau_{\max}/4$, and our analysis indeed yields an approximation ratio of $(1 + \ln 4) \approx 2.4$ for uniform tariff pricing, the additional 0.4 being caused by the difference between the functions $g(x)$ and $f(x)$.

Note also that the worst case example, Example 1, crucially hinges on a staircase function $f(x)$ that approximates a hyperbola; in particular it requires exponential prices. Thus, it can be conjectured that the empirical performance of uniform tariff pricing policies on practical instances, where the




Figure 7. The analysis of uniform pricing is tight.

Instance	$ N $	$ A $	m	n	Π^{RTP}	Π^{UTP}	%
RTN1	29	94	7	15	841	624	74%
RTN2	29	98	6	21	4099	3496	85%
RTN3	59	206	10	13	1118	880	79%
RTN4	59	204	10	20	2217	1512	68%
RTN5	49	120	9	21	74948	55968	74%
RTN6	33	116	15	12	28166	20328	72%

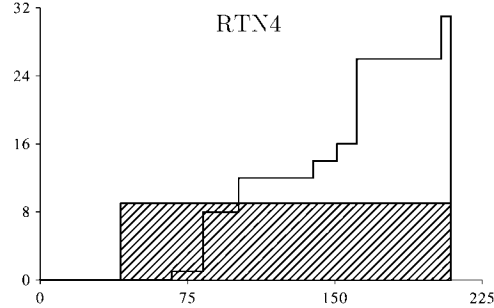


Figure 8. Uniform pricing on France Télécom instances.

price range is subexponential, outperforms the theoretical bounds we have found. Our experiments, using data from the France Télécom interconnections market, are summarized in Figure 8 and they corroborate this conjecture.

The table on the left in Figure 8 shows data for six instances that represent telecommunication networks for the international interconnections market, as described in Section 1.2. We compare the optimal solution values for uniform tariffs Π^{UTP} and non-uniform tariffs Π^{RTP} . The optimal non-uniform solution is calculated using the model and mixed integer programming formulation described by Bouhtou et al. [3]. The value of Π^{UTP} is calculated using the same formulation, requiring that all tariffs be equal. We do not compare the actual computation times here, but are only interested in effectiveness of the optimal uniform pricing. The table gives a brief description of each network, stating the number of nodes $|N|$, arcs $|A|$, tariff arcs m , and clients n . The optimal non-uniform and uniform solution values are displayed in the columns Π^{RTP} and Π^{UTP} . The final column is the approximation ratio. The graph on the right in Figure 8 shows the staircase function $f(x)$ for the optimal non-uniform solution, as well as the best uniform solution, for instance RTN4.

3. ALL-SERVICE RIVER TARIFF PRICING

In this section, we consider the following variation of the river tariff pricing problem. The operator must set tariffs in order to capture the demand of all clients; that is, tariffs must be such that no client k is forced to use the arc (s_k, t_k) . We refer to this problem as the *all-service* river tariff pricing problem.

It follows from trivial examples that the maximal revenue for the all-service problem can be an arbitrary factor away from the maximal revenue without the all-service constraint. Hence, we have an arbitrarily high “cost of regulation”, where the regulation consists of the fact that we force the operator to serve all clients. Note that in case $u_k < c_{ka}$ for some clients k and tariff arcs a , the operator might even be forced to use negative tariffs, up to the extreme case where the optimal revenue becomes negative. In such a situation, the notion of approximation algorithms is senseless. Hence, with respect to

approximability, we consider the special case where $u_k \geq c_{ka}$ for all clients k and tariff arcs a . We show that even for this restriction, the maximal revenue for the all-service river tariff pricing problem cannot be approximated within any reasonable bound.

THEOREM 5: For any $\varepsilon > 0$, the existence of a polynomial time approximation algorithm for the all-service river tariff pricing problem with n clients and m tariff arcs with worst case ratio $O(m^{1-\varepsilon})$ or $O(n^{1/2-\varepsilon})$ implies $P = NP$.

PROOF: We use an approximation preserving reduction from INDEPENDENT SET [6] to the all-service problem. The INDEPENDENT SET problem asks for finding in a graph $G = (V, E)$ a maximum cardinality subset $V' \subseteq V$ such that no two vertices in V' are connected by an edge. The transformation works as follows. For every vertex $v \in V$ we introduce a client with origin–destination pair $\{s_v, t_v\}$ and demand $d_v = |E|$, and a corresponding tariff arc a_v . We connect the source s_v to the tail of the tariff arc a_v and the head of a_v to the destination t_v , using zero cost fixed cost arcs. Moreover, there is a fixed cost arc (s_v, t_v) with cost $(|V| + 1)$ for all vertices $v \in V$. For every edge $e \in E$ we introduce a client with origin–destination pair $\{s_e, t_e\}$ and unit demand. The upper bound on the cost of routing this demand is given by the fixed cost arc (s_e, t_e) with cost 1. For all edges $e \in E$ and all vertices $v \in V$ with $v \in e$, we furthermore introduce fixed cost arcs $(s_e, \text{tail}(a_v))$ and $(\text{head}(a_v), t_e)$, with zero cost. This transformation results in an instance of the all-service problem with $|V|$ tariff arcs, and $|V| + |E|$ clients. Figure 9 gives an example of such a transformation for a graph $G = (V, E)$ with 3 nodes and 2 edges.

We claim that G has an independent set of cardinality at least k if and only if there exists a tariff policy for the all-service problem with a total revenue of $|V||E|(k + 1) + |E|$.

First, assume that G has an independent set V' of cardinality k . For all $v \in V'$, set the tariff on the corresponding tariff arc a_v to $|V| + 1$ and all other tariffs to 1. By the definition of an independent set, for any edge $e = (v, u) \in E$ at least one of the vertices, v or u , is not in V' . Therefore, the tariff of at least one of the tariff arcs, a_v or a_u , is 1. All clients




Figure 9. Reduction of INDEPENDENT SET to the all-service pricing problem.

corresponding to an edge e can thus be served, using one of the tariff arcs a_v or a_u . The clients (s_v, t_v) corresponding to the vertices $v \in V$ are also served, since the upper bound of $|V| + 1$ is not exceeded with the so-defined tariffs. Hence, all demands are served. The revenue consists of $|E|$ from all clients corresponding to the edges E of G , $|E|(|V| + 1)k$ from the clients corresponding to the independent set V' and $|E|(|V| - k)$ from the clients corresponding to $V \setminus V'$. That yields a total revenue of $|E||V|(k + 1) + |E|$.

Conversely, assume that there exists a set of tariffs that captures all demands, such that the revenue is $|E||V|(k + 1) + |E|$. We will show that this implies that the graph G has an independent set of cardinality at least k . Since all demands are captured at this tariff pricing policy, for any edge $e = (v, u) \in E$, the tariff on at least one of the arcs, a_v or a_u , is 1. Consider the set of vertices $V' := \{v \in V : t_{a_v} > 1\}$. By definition, no pair of nodes $v, u \in V'$ is connected by an edge. Hence, V' is an independent set in G . Let $k' := |V'|$. The revenue is equal to $|E| + |E|(|V| - k') + |E|(|V| + 1)k' = |E||V|(k' + 1) + |E|$, which by assumption is at least as large as $|E||V|(k + 1) + |E|$. This implies that $k' \geq k$ and thus that V' is an independent set in G of cardinality $k' \geq k$.

Now, let us assume that we have an α -approximation algorithm \mathcal{A} for the all-service problem, with $\alpha \geq 1$. Consider any instance $G = (V, E)$ of INDEPENDENT SET and the all-service problem resulting from the above reduction. We can assume that both the optimal solution and the solution produced by \mathcal{A} only utilize tariff values 1 or $|V| + 1$, because any tariff greater than 1 and not equal to $|V| + 1$ can be turned into $|V| + 1$ with a revenue gain. So $\Pi^{\text{RTP}} = |E||V|(k + 1) + |E|$ for some k , and $\Pi^{\mathcal{A}} = |E||V|(k' + 1) + |E|$ for some k' . The first part of the proof yields that the maximal independent set of G has size k , and algorithm \mathcal{A} can be used to find an independent set of size at least k' . Moreover,

$$\frac{1}{\alpha} \leq \frac{|E||V|(k' + 1) + |E|}{|E||V|(k + 1) + |E|} = \frac{1 + \frac{1}{|V|} + k'}{1 + \frac{1}{|V|} + k} \leq \frac{2 + k'}{1 + k};$$

hence, $k' \geq (k + 1)/\alpha - 2$. In other words, we have an $O(\alpha)$ -approximation algorithm for the INDEPENDENT SET problem.

It follows from Zuckerman [13], who recently improved a previous result from Håstad [9], that the INDEPENDENT SET problem cannot have a polynomial time approximation algorithm with worst case guarantee $|V|^{1-\epsilon}$ for any $\epsilon > 0$, unless $P = NP$. Since the number of tariff arcs m in our transformation equals $|V|$, the first claim of the theorem follows. Since the number of clients n in our transformation equals $|V| + |E| \in O(|V|^2)$, the second claim follows. \square

Note that this inapproximability result shows that, for the all-service RTP, we cannot even expect a performance guarantee logarithmic in the total demand D , like the one we obtained before. On the positive side, however, we can show the following.

THEOREM 6: There exists an n -approximation algorithm for the all-service river tariff pricing problem.

PROOF: In an optimal solution, at least one client contributes to the total revenue at least Π^{RTP}/n , and this contribution is achieved by utilizing a specific tariff arc at a certain tariff. The proof now works by enumeration over all $m \cdot n$ possibilities for a client using a specific arc. So assume that a tariff arc b and a client k are fixed. We claim that we can compute the maximum tariff τ_b on arc b , together with tariffs on all the other arcs, such that client k indeed utilizes arc b and all other clients are served. Taking the maximum over all $m \cdot n$ possibilities for a client using a specific arc, the revenue of this solution is obviously at least Π^{RTP}/n .

The computation of this maximum tariff τ_b on arc b , together with tariffs on all the other arcs, such that client k indeed utilized arc b and all other clients are served, can be achieved by binary search over the possible tariffs τ on arc b . Denote by c_{ka} the fixed cost for client k when utilizing arc a , and recall that u_k denotes the maximum total (per unit) cost for client k . Given that client k utilizes arc b , the maximum tariff on arc b is $u_k - c_{kb}$, which determines the interval for the binary search. Given some tariff τ on arc b , in order to

make sure that client k utilizes arc b , we just define the tariffs on all other tariff arcs a as $\tau_a = \tau + c_{kb} - c_{ka}$. It is straightforward to verify whether this yields a feasible solution with all clients served or not. \square

ACKNOWLEDGMENTS

This research was partially supported by France Télécom Research & Development. We thank Maxim Sviridenko for an insightful discussion and the suggestion for the proof of Theorem 6. Thanks also to the anonymous referees for some remarks that helped to considerably improve the paper.

REFERENCES

- [1] G. Aggarwal, T. Feder, R. Motwani, and A. Zhu, "Algorithms for multi-product pricing," Automata, Languages and Programming, volume 3142 of Lecture Notes in Computer Science, J. Diaz, J. Karhumäki, A. Lepistö, and D. Sannella (Editors), Springer, Berlin, 2004, pp. 72–83.
- [2] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Protasi, Complexity and Approximation: Combinatorial Optimization Problems and their Approximability Properties, Springer, Berlin, 1999.
- [3] M. Bouhtou, S. van Hoesel, A.F. van der Kraaij, and J.L. Lutton, Tariff optimization in networks, Meteor Research Memorandum RM03011, Maastricht University, 2003, <http://edata.ub.unimaas.nl/www-edocs/loader/file.asp?id=791>
- [4] R. Cole, Y. Dodis, and T. Roughgarden, How much can taxes help selfish routing? J Comp Syst Sci 72 (2006), 444–467.
- [5] R. Cole, Y. Dodis, and T. Roughgarden, Pricing network edges for heterogeneous selfish users, Proc. 35th Annual ACM Symposium on Theory of Computing, 2003, pp. 521–530.
- [6] M.R. Garey and D.S. Johnson, Computers and Intractability: A Guide to the Theory of NP-Completeness, W. H. Freeman, New York, 1979.
- [7] A. Grigoriev, A.F. van der Kraaij, S. van Hoesel, M. Uetz, and M. Bouhtou, "Pricing network edges to cross a river," Approximation and Online Algorithms, volume 3351 of Lecture Notes in Computer Science, G. Persiano and R. Solis-Oba (Editors), Springer, Berlin, 2005, pp. 140–153.
- [8] V. Guruswami, J.D. Hartline, A.R. Karlin, D. Kempe, C. Kenyon, and F. McSherry, On profit-maximizing envy-free pricing, Proc. 16th ACM-SIAM Symposium on Discrete Algorithms, 2005, pp. 1164–1173.
- [9] J. Håstad, Clique is hard to approximate within $n^{1-\epsilon}$, Acta Math 182 (1999), 105–142.
- [10] M. Labbé, P. Marcotte, and G. Savard, A bilevel model of taxation and its application to optimal highway pricing, Manage Sci 44 (1998), 1608–1622.
- [11] S. Roch, P. Marcotte, and G. Savard, An approximation algorithm for Stackelberg Network Pricing, <http://arxiv.org/abs/cs.GT/0409054.html>
- [12] T. Roughgarden and É. Tardos, How bad is selfish routing? J Assoc Comput Mach 49 (2002), 236–259.
- [13] D. Zuckerman, Linear Degree Extractors and the Inapproximability of MAX CLIQUE and CHROMATIC NUMBER, Proc. 38th ACM Symposium on Theory of Computing, 2006, pp. 681–690. Also available as: ECCC Report TR05-100. <http://www.eccc.uni-trier.de/eccc/>