Pricing Bridges to Cross a River *

Alexander Grigoriev [1], Stan van Hoesel [1], Anton F. van der Kraaij [1],
Frits Spieksma [2], Marc Uetz [1], and Mustapha Bouhtou [3]

[1] Maastricht University, Quantitative Economics, P.O.Box 616,
NL–6200 MD Maastricht, The Netherlands
{a.grigoriev,s.vanhoesel,a.vanderkraaij,m.uetz}@ke.unimaas.nl
[2] Katholieke Universiteit Leuven, Applied Economics,
Naamsestraat 69, B–3000 Leuven, Belgium
Frits.Spieksma@econ.kuleuven.ac.be
[3] France Télécom R&D, 39-40 rue du Général Leclerc,
F–92131 Issy-Les-Moulineaux, France
mustapha.bouhtou@francetelecom.com

Abstract. We consider a Stackelberg pricing problem in directed, uncapacitated networks. Tariffs have to be defined by an operator, the leader, for a subset of m arcs, the tariff arcs. Costs of all other arcs are assumed to be given. There are n clients, the followers, that route their demand independent of each other on paths with minimal total cost. The problem is to find tariffs that maximize the operator's revenue. Motivated by problems 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 $\mathcal{APX}$-hard. Moreover, we show that uniform pricing yields both an m–approximation, and a $(1 + \ln D)$–approximation. Here, D is upper bounded by the total demand of all clients. We furthermore discuss some polynomially solvable special cases, and present a short computational study with instances from France Télécom. In addition, we consider the problem under the additional restriction that the operator must serve all clients. We prove that this problem does not admit approximation algorithms with any reasonable performance guarantee, unless $\mathcal{NP} = \mathcal{ZPP}$, and we prove the existence of an n–approximation algorithm.

1 Introduction

The general setup for the tarification problem that we study involves two non-cooperative groups, an operator that sets tariffs, the *leader* of the Stackelberg game, and n clients that have to pay these tariffs, the *followers* of the Stackelberg game. More precisely, we assume that a network is given, and a subset of m 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. Such a path can in general involve one or several of the tariff arcs belonging to the operator, and we assume that each client selfishly selects a path with minimum total cost to route his demand. Before the clients select their paths, the operator has to set the tariffs, which he does in order to maximize 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.

Notice that this problem is different in two aspects from the network congestion problems studied recently, e.g., by Roughgarden and Tardos [13], and Cole et al. [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. The Game Theoretic setting is only introduced by the fact that

* This research was partially supported by France Télécom Research & Development. An extended abstract with parts of the results appeared at WAOA 2004 [7].

there exist an operator trying to maximize revenue using high tariffs, and the clients try to avoid high tariffs by choosing minimal cost paths. Second, the pricing takes place before the users choose their paths, so we are faced with a Stackelberg game, where the operator first sets the tariffs, and then, subject to these tariffs, the clients react selfishly.

2 Model

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

An instance of the general tarification 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 has to be routed from source node s_k to target node t_k. Because there is no congestion involved, we may assume without loss of generality 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 according to 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 modelled as a linear bilevel program [11].

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 tarification 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 4, we discuss practical applications for this model, motivated by problems 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 tarification 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. Notice that the only difference to the general tarification problem described previously is the non-existence of backward arcs in the shortest path graph model.

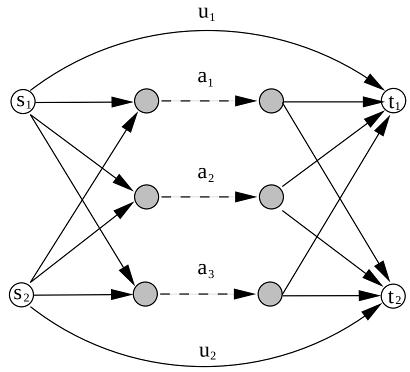

Fig. 1. River tarification problem with $n = 2$ and $m = 3$.

To summarize, the parameters that define an instance of a river tarification problem are the number of tariff arcs m, the number of clients n, their demand values d_k, $k \in \{1, \ldots, n\}$, 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 is immediate that any reasonable solution will adopt only tariffs which are integral, too. Notice that this might not be true for the general tarification problem, where a path chosen by a client can consist of more than one tariff arc.

3 Related work and results

A formulation of the general tarification problem is the linear bilevel program described by Labbé et al. [11]. They show, among other things, that already the problem with a single client is strongly $\mathcal{NP}$-hard, given that also negative tariffs are allowed. Roch et al. [12] show that the single client problem remains strongly $\mathcal{NP}$-hard, even when restricted to nonnegative tariffs. In the same paper, a polynomial time $(1 + 1/2 \ln m)$–approximation algorithm is proposed for the problem with a single client, where m is the number of tariff arcs.

In contrast to these two papers, which mainly address the general tarification problem with a single client, we consider the problem with multiple clients. However, as described above, we assume that each client utilizes at most one tariff arc, yielding an instance of the river tarification problem. In fact, this problem can equivalently be interpreted as a pricing problem for multiple products, where the tariff arcs $a \in T$ correspond to different products, and each user 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 products). Whenever there is an arc between a client k and a tariff arc a in the river tarification 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 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 tarification 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 has to 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 $\mathcal{APX}$-hardness of this problem, together with a $(1 + \ln n)$-approximation algorithm. Notice that, despite of the obvious similarities, the multi-product pricing problem is conceptually different from the tarification problem considered in this paper. In the river tarification 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.

Guruswani 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 has to 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 tarification problem considered here. Guruswani et al. [8] independently prove $\mathcal{APX}$-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.

In this paper, we derive several results concerning complexity and approximability of the river tarification problem. In Section 5.1, by a reduction from the Max-2-Sat-3 problem, we show that the river tarification problem is $\mathcal{APX}$-hard, even if each client is connected to at most two tariff arcs. This result coincides with the $\mathcal{APX}$-hardness result of [8]; obtained independently. The quality of uniform tarification policies, where all arcs are priced with the same tariff, is analyzed in Section 5.2. The problem to find an optimal uniform tariff is well-known to be solvable in polynomial time, even for the general tarification problem [14]. We show that uniform tarification is an m–approximation, and this is tight. Using a simple geometric argument, we also show that uniform tarification 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 upper bounded by the total demand. Hence, whenever the clients have unit demand, this yields a $(1 + \ln n)$–approximation. We empirically analyze the quality of uniform tarification policies in Section 5.3, using instances from France Télécom. Finally, we briefly discuss some polynomially solvable special cases of the river tarification problem in Section 5.4.

In Section 6, we consider another variant of the problem where the operator is not allowed to reject any client. Notice that this 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 this problem does not allow approximation to within a factor $\mathcal{O}(m^{1-\varepsilon})$ or $\mathcal{O}(n^{1/2-\varepsilon})$, unless $\mathcal{ZPP} = \mathcal{NP}$. (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.

4 Applications

For a first motivation, consider the internet. Whenever an autonomous system, represented by some subnetwork, has to transit data, the data may enter and exit the autonomous system at different points. Clients have to pay a price for transmitting data through the autonomous system, yielding revenue for its owner. The data flow can be modelled such that once it is routed through the autonomous system, it does not pass a second time. See Figure 2(a), where the clients have the choice between two different autonomous systems (AS1 and AS2), with two entry- and exit-points. This can be modelled as a river tarification problem, introducing a tariff arc for each entry-exit combination.

For another motivation, consider point-to-point markets, where a telecommunications operator is offering bandwidth capacity between two points A and B at a certain price. 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 tarification problem, too. Figure 2(b) 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}, $i \in \{1, 2, 3\}$, 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 in the figure).

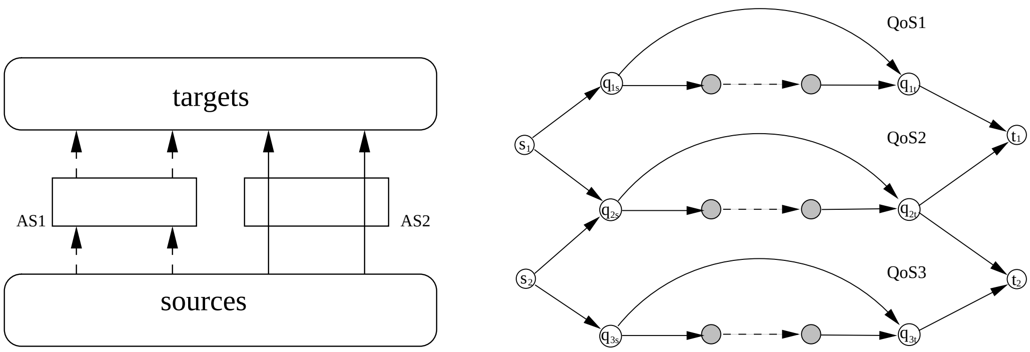

(a) Autonomous systems. (b) Point-to-point markets.

Fig. 2. Applications of the river tarification problem.

Finally, the network topology of a river tarification problem may be assumed –possibly after the simple transformation described previously– in telecommunication networks where it is known a priori that the path of each client utilizes at most one tariff arc. This occurs, e.g., in the international interconnections market, where several operators offer connections to a particular country. If we focus on the market for one particular country, we can assume that it is not profitable for any client to enter the country twice.

5 River Tarification: Complexity and Approximation

We first show that the river tarification problem is $\mathcal{APX}$-hard. Then, we derive bounds on the quality of uniform tarification policies, where all tariffs are required to be identical. We present a short computational study with instances from France Télécom, and briefly discuss some polynomially solvable special cases.

5.1 Complexity

Roch et al. [12] show that the general tarification problem is $\mathcal{NP}$-hard in the strong sense, even when restricted to a single client, using a reduction from the $\mathcal{NP}$-complete problem 3-Sat [6]. Their reduction works for tarification problems where paths are allowed to use (and indeed, must use) several tariff arcs. We consider a reduction of MAX-2-SAT-3. This is a variant of the SATISFIABILITY problem [6]. An instance of the SATISFIABILITY problem is a boolean function $f : \{0, 1\}^n \to \{0, 1\}$ on n variables $x_1, \ldots,$ x_n, in conjunctive normal form. Such a function f is the conjunction of m clauses C_k, $f = \bigwedge_{k=1}^{m} C_k$, each clause being the disjunction of one or several literals. Any literal represents either a variable x_i, or its negation $\bar{x}_i$, $i \in \{1, \ldots, n\}$. Function f is called satisfiable if there exists a truth assignment $x_1,\ldots, x_n$ such that at least one literal per clause is true. In the 2-SAT problem, each clause C_k is the disjunction of exactly 2 literals. In the 2-SAT-3 problem, in addition, each variable occurs at most 3 times in all literals. The problem to find a truth assignment that fulfills the maximum number of clauses, MAX-2-SAT-3, is known to be $\mathcal{APX}$-hard [2].

Theorem 1. *The river tarification problem is $\mathcal{APX}$-hard, even when each client is connected to at most two tariff arcs.*

Proof. We use an approximation-preserving reduction from MAX-2-SAT-3. For each variable x_i, $i \in$

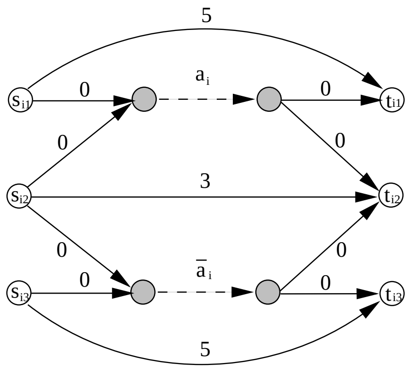

Fig. 3. Subnetwork for variable x_i, $i \in \{1, \ldots, n\}$.

$\{1, \ldots, n\}$ of the MAX-2-SAT-3 instance, we construct a constant-size subnetwork as shown in Figure 3. Each of these subnetworks has three clients with unit demand, with origin-destination pairs $\{s_{ij}, t_{ij}\}$, $j \in \{1, 2, 3\}$. Moreover, each subnetwork has two tariff arcs, a_i representing the truth assignment $x_i = 1$, and $\bar{a}_i$ 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}), both with cost 5. For commodity 2, the upper bound on the cost is given by a fixed cost arc (s_{i2}, t_{i2}), with cost 3. We call this client the *cheap client*.

Next, for each clause C_k, $k \in \{1, \ldots, m\}$, 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 as one of clause C_k's literals, 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 3, defining an upper bound of 3 for the cost of routing clause-commodity k. The so-defined instance of the river tarification problem has $2n$ tariff arcs, $3n+m$ commodities (or clients), and at most $7m+11n$ fixed cost arcs, hence the transformation is indeed polynomial.

Let an instance of MAX-2-SAT-3 be given, and denote by Π^{SAT} the maximum number of satisfiable clauses, and by Π^{RTP} denote the maximum revenue for the corresponding instance of the river tarification problem. Moreover, we denote by s any feasible solution of MAX-2-SAT-3, and we denote by S any feasible solution for the river tarification problem.

In order to exhibit an approximation-preserving reduction, we need to establish the validity of two inequalities:

(i) there exists a constant $\alpha > 0$ such that

$$\Pi^{\mathrm{RTP}} \leq \alpha \Pi^{\mathrm{SAT}} \,,$$

(ii) there exists a constant $\beta > 0$ such that for any feasible solution S of the river tarification problem, there exists a solution s for Max-2-Sat-3 with

$$\Pi^{\mathrm{RTP}} - \Pi^{\mathrm{RTP}}_S \geq \beta(\Pi^{\mathrm{SAT}} - \Pi^{\mathrm{SAT}}_s) \,.$$

First, we prove (i). The maximal revenue for each subnetwork is given by setting one of the tariffs to 3, and another one to 5, yielding a revenue of $5 + 2 \cdot 3 = 11$. In all other cases, the revenue is not more than 10. Therefore, we have that $\Pi^{\mathrm{RTP}} \leq 11n + 3m$. Since no more than two variables can be present in a clause, $m \geq n/2$. Notice that $\Pi^{\mathrm{SAT}} \geq m/2$ (just consider the best of the following two solutions: set all variables to true, and set all variables to false). It follows that $\Pi^{\mathrm{RTP}} \leq 50\Pi^{\mathrm{SAT}}$, hence (i) holds with $\alpha = 50$.

Now, we prove (ii). Given some solution S for the river tarification problem, we are going to build another solution S' for the river tarification problem such that $\Pi^{\mathrm{RTP}}_{S'} \geq \Pi^{\mathrm{RTP}}_S$. Next we associate to S' a solution s for Max-2-Sat-3. Our goal is to build a solution S' where the demand of each cheap client is routed via one of the two tariff arcs of the respective subnetwork.

First, we consider the subnetworks in which both tariffs are greater than 3. If both tariffs are greater than 3 in a subnetwork, we set arbitrarily one to 5, the other one to 3, and we route the demand of the cheap client via the arc with tariff 3. Observe that this does not decrease the subnetwork's contribution to the value of the resulting solution. If we can route a not-yet-routed clause commodity through the tariff arc with tariff 3, we do so.

Second, we consider the subnetworks in which one tariff is greater than 3, and the other one is less than or equal to 3. In that case, we set the higher tariff to 5, and the smaller tariff to 3. In case the demand of the cheap client had not yet been routed, we route it via the tariff arc with tariff 3. Again, one can verify that this does not decrease the subnetwork's contribution to the value of the resulting solution. If we can route a not-yet-routed clause commodity through the tariff arc with tariff 3, we again do so.

Finally, we consider the subnetworks in which both tariffs are less than or equal to 3. Let there be p subnetworks of this type. These p subnetworks contribute to the value of the current solution at most $9p + \ell$ where ℓ is the number of clause commodities that are routed via a tariff arc of one of the p subnetworks. For these p subnetworks, let us consider two different partial solutions. In one of these partial solutions the demand of the cheap client is routed upwards at tariff 3, the tariff of the lower arc is set to 5, and we simply disconnect all clause commodities that were previously routed via the lower arc before. In the other partial solution, the demand of the cheap client is routed downwards at tariff 3, the tariff of the upper arc is set to 5, and we simply disconnect all clause commodities that were previously routed via the upper arc. Also, if we can route a not-yet-routed clause commodity through the tariff arc with tariff 3, in both partial solutions, we do so.

We claim that at least one of these two partial solution has a contribution to the solution's value that is not less than the original contribution $9p + \ell$. Indeed, we now have $11p + \ell/2$ for at least one of these two partial solutions, since each clause commodity had to go via an upper or a lower arc. Notice that $11p + \ell/2 \geq 9p + \ell$ holds if $p \geq \ell/4$. Since we are dealing with Max-2-Sat-3 where each variable (subnetwork) is connected to at most 3 clauses, it follows that $\ell \leq 3p$, or $p \geq \ell/3$.

Now we have completed the construction of S' and we have verified that $\Pi^{\mathrm{RTP}}_{S'} \geq \Pi^{\mathrm{RTP}}_S$. Notice that in S' the demand of each cheap client has been routed via a tariff arc. This determines the truth assignment of the corresponding variables in solution s of Max-2-Sat-3: We define the truth assignment of s according to the utilized tariff arcs of the cheap clients. Thus, the number of clause commodities routed via a tariff arc in S' equals the number Π^{SAT}_s of satisfied clauses in the Max-2-Sat-3 solution s.

Let us now finish the argument for (ii). We have that $\Pi^{\mathrm{RTP}} \geq 11n + 3\Pi^{\mathrm{SAT}}$, since the optimal solution for Max-2-Sat-3 yields a feasible solution for the river tarification problem where Π^{SAT} clause commodities contribute a revenue of 3, and each subnetwork contributes a revenue of 11. Also, $\Pi_S^{\mathrm{RTP}} \leq \Pi_{S'}^{\mathrm{RTP}} = 11n + 3\Pi_s^{\mathrm{SAT}}$ by our previous construction. Thus $\Pi^{\mathrm{RTP}} - \Pi_S^{\mathrm{RTP}} \geq 11n + 3\Pi^{\mathrm{SAT}} - (11n + 3\Pi_s^{\mathrm{SAT}}) = 3(\Pi^{\mathrm{SAT}} - \Pi_s^{\mathrm{SAT}})$, which finishes the proof with $\beta = 3$. □

Observe that, in a sense, we have delineated a borderline between easy and hard instances of the river tarification 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. Hence, even for that case no polynomial time approximation scheme can exist, unless $\mathcal{P} = \mathcal{NP}$.

5.2 The quality of uniform tarification policies

The uniform tarification problem (UTP) is the same as the general tarification problem, with the additional restriction that all tariffs are required to be identical. As shown by van Hoesel et al. [14], the uniform tarification problem can be solved in polynomial time, even in the general setting where clients may use paths with several tariff arcs. The algorithm described by van Hoesel et al. [14] uses the parametric shortest path algorithm of Young et al. [15] and Karp and Orlin [10] to determine the tariff values (i.e. breakpoints) for which the shortest path tree changes for any client. Calculating the revenue for the operator at each breakpoint and maintaining the best solution yields the optimal uniform tarification policy in polynomial time.

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

Lemma 1. *If an optimal tarification for the river tarification problem with revenue Π^{RTP} utilizes at most r different tariffs, then for the optimal uniform tarification, $\Pi^{\mathrm{UTP}} \geq \Pi^{\mathrm{RTP}}/r$.*

The proof of this lemma is indeed trivial. To this end, consider an optimal non-uniform tarification with tariffs $\tau_1 \leq \cdots \leq \tau_m$, and let D_i be the total demand on an arc a_i with tariff τ_i, $i \in \{1, \ldots, m\}$. 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 $f : [0, D] \to [0, \infty[$, depicted in Figure 4.

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

Proof (of Lemma 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 $\Pi^{\mathrm{UTP}} \geq T_i$ for all $i \in \{1, \ldots, m\}$, since the area of any such rectangle is a lower bound for the revenue yielded by the optimal uniform tariff Π^{UTP}. (Notice that this does not hold for the general tarification problem.) Hence, if only r different tariffs are utilized, we consider the r (inclusion-)maximal rectangles under function f, say $T_{i_1}, \ldots, T_{i_r}$, and get $r \cdot \Pi^{\mathrm{UTP}} \geq \sum_{j=1}^{r} T_{i_j} \geq \Pi^{\mathrm{RTP}}$. □

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

Theorem 2. *Uniform tarification is an m–approximation for the river tarification problem.*

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

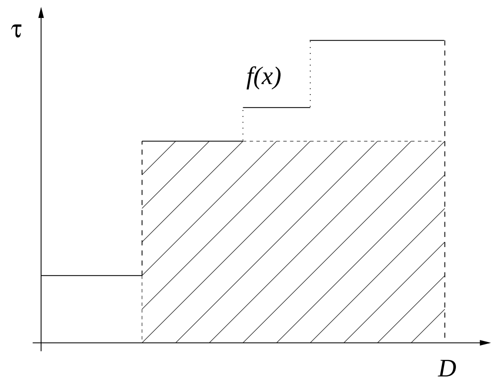

Fig. 4. Staircase function $f(x)$ with inscribed rectangle.

Theorem 3. *Uniform tarification is a* $(1+\ln D)$*–approximation for the river tarification 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. Indeed, we will even prove a slightly stronger result than claimed in Theorem 3. Consider an optimal non-uniform tarification, 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 $\tau_{\max}$ the maximal tariff utilized in that optimal solution. We show

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

Theorem 3 then follows, because $T_\ell \geq \tau_{\max}$ by definition of T_ℓ. To prove (2), we define the function

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

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. Now

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

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

Hence, the area under the staircase function, which equals Π^{RTP}, can be upper bounded in terms of the area defined by the function $g(x)$, as depicted in Figure 5. To compute this area, we partition it into three parts, namely the rectangle T_ℓ itself, the area under $g(x)$ on the domain $x \in [0, D - x_\ell]$, as well as the area to the right of $g(x)$ on the domain $\tau \in [\tau_\ell, \tau_{\max}]$. 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.

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

and since $T_\ell \leq \Pi^{\mathrm{UTP}}$, claim (2) follows. □

Notice that claim (2) confirms the following geometric intuition: The closer the staircase function $f(x)$ is to the straight line $x \mapsto (\tau_{\max}/D)\cdot x$, the closer is T_ℓ to $D\tau_{\max}/4$, which yields an approximation

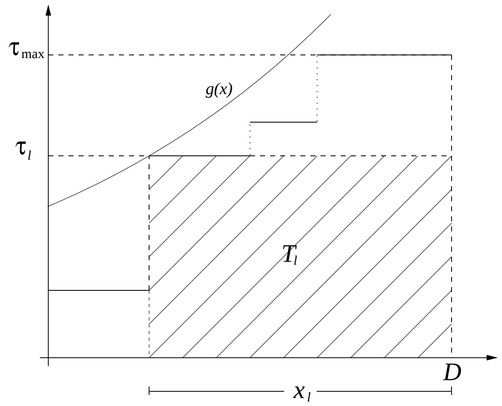

Fig. 5. Illustration for the proof of Theorem 3.

ratio of $(1 + \ln 4) \approx 2.4$ for uniform tarification. Geometric intuition indeed suggests a ratio of roughly 2, the additional 0.4 being caused be the difference between the functions $g(x)$ and $f(x)$. In Section 5.3, we compare the quality of uniform versus non-uniform tarification, based on instances obtained from France Télécom.

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 tarification is a $(1 + \ln n)$–approximation for the river tarification 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 $\tau_{\max}$ equals b^{2m-1}. See Figure 6 for an example with $n = m = 4$. □

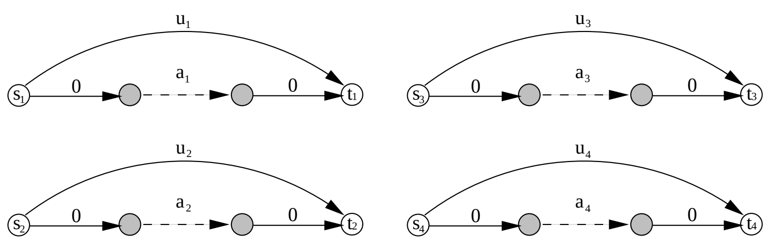

Fig. 6. The analysis of uniform tarification policies is tight.

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 tarification 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 (reasonable) uniform tariffs would be values

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

$$\Pi^{\mathrm{UTP}}/\Pi^{\mathrm{RTP}} = \frac{b^{2m} - b^m}{m(b^{2m} - b^{2m-1})} \leq \frac{b^{2m}}{m(b^{2m} - b^{2m-1})} = \frac{1}{m} \cdot \frac{b}{b-1}\,.$$

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

Moreover, Theorem 3 suggests that uniform tarification 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 tarification is a $\mathcal{O}(m)$–approximation on this example. The same Example 1 shows that $\mathcal{O}(m)$ is indeed best possible. Summarized, we thus get the following.

Theorem 4. *For uniform tarification, the performance bound of Theorem 2 is best possible, and the performance bound of Theorem 3 is best possible up to a constant factor.*

5.3 Numerical Results

As stated previously, whenever the function that describes the total revenue in an optimal non-uniform solution, i.e. the staircase function defined in (1), is close to a straight line, geometric intuition suggests a worst-case ratio for uniform tarification of approximately 2. The worst case Example 1 crucially hinges on a (staircase) function that approximates a hyperbola. Thus, it can be conjectured that the empirical performance of uniform tarification policies outperforms the theoretical bounds we have found. This is indeed confirmed in the following numerical experiments, displayed in Table 1. The study is based on instances obtained from France Télécom.

Table 1. Quality of Uniform Tarification on France Télécom instances.

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%

These instances represent telecommunication networks for the international interconnections market, as described in Section 4. We compare the optimal solutions 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 in 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 tarification policies. Table 1 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.

5.4 Polynomially solvable special cases

Several polynomially solvable special cases of the (general) tarification problem are discussed by Labbé et al. [11] and van Hoesel et al. [14]. Clearly, these results hold for the problem considered in this paper, too. In addition, we have the following polynomially solvable cases for the river tarification problem.

Theorem 5. *If the number of clients n is bounded from above by a constant, the river tarification problem is solvable in polynomial time.*

Proof. If the number of clients n is bounded from above by a constant, the number of assignments of clients to tariff arcs is bounded by m^n which is a polynomial for fixed n. Consider an assignment of clients to tariff arcs and let the tariff arc (if any) taken by a client k be denoted by a_k. The following linear program models the revenue maximization problem, given such an assignment of clients to tariff arcs.

$$\begin{aligned} \max_{\tau} \ & \sum_{k \in K} d_k \tau_{a_k} & \\ s.t. \ & \tau_{a_k} + c_{ka_k} \leq \tau_a + c_{ka}, & \forall\, k \in K, \forall\, a \in T, \\ & \tau_a \geq 0, & \forall\, a \in T. \end{aligned} \tag{4}$$

Solving m^n instances of (4), we can retrieve the optimal solution in polynomial time. □

Finally, for specific cost structures of the network, the river tarification problem is polynomially solvable, too. To this end, assume that the cost c_{ka} for connecting client k to tariff arc a is only composed of an individual cost per client c_k, independent of the tariff arc a, and of a shipment cost c_a, independent of the client k. In other words, we assume that there exist positive integers c_k, $k \in \{1, \ldots, n\}$, and c_a, $a \in T$, such that $c_{ka} = c_k + c_a$ for every client k and tariff arc a.

Theorem 6. *Assuming the cost structure described above, the river tarification problem is solvable in polynomial time.*

Proof. Consider a feasible solution of the problem and let $\tau_a, a \in T$, be the corresponding tariffs on the tariff arcs. Without loss of generality, we assume that in that solution client 1 utilizes a tariff arc a_1 at price τ_1 and client 2 utilizes a tariff arc a_2 at price τ_2. Since in any feasible solution client 1 routes her demand via the shortest path, we have that $c_{1a_1} + \tau_1 \leq c_{2a_2} + \tau_2$. From the cost structure we derive that $c_1 + c_{a_1} + \tau_1 \leq c_1 + c_{a_2} + \tau_2$ and therefore $c_{a_1} + \tau_1 \leq c_{a_2} + \tau_2$. Similarly, for client 2 we have that $c_{a_1} + \tau_1 \geq c_{a_2} + \tau_2$. Hence we obtain that $c_{a_1} + \tau_1 = c_{a_2} + \tau_2$, meaning that users 1 and 2 are indifferent between tariff arcs 1 and 2. Thus we may assume without loss of generality that they both use the same tariff arc. Proceeding iteratively this way, we may assume that all clients use the same tariff arc. For every choice of the utilized tariff arc we can now find the maximum revenue associated with this arc (assuming that the other tariff arcs are so expensive that no client can afford their utilization), using binary search on the tariff for that tariff arc. Choosing the arc that provides the highest total revenue we obtain the optimal solution for the problem. □

6 All-service river tarification problem

In this section, we consider the following variation of the river tarification 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 tarification problem. $\mathcal{NP}$-hardness of this problem follows by our previous reduction presented in Section 5.

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'. In addition, we can show that the maximal revenue for the all-service problem cannot be approximated within any reasonable bound.

Theorem 7. *For any $\varepsilon > 0$, the existence of a polynomial time approximation algorithm for the all-service river tarification problem with with n clients and m tariff arcs with worst case ratio $\mathcal{O}(m^{1-\varepsilon})$ or $\mathcal{O}(n^{1/2-\varepsilon})$ implies $\mathcal{ZPP} = \mathcal{NP}$.*

Proof. The proof uses an approximation preserving reduction from Independent Set [6] to the all-service RTP. So assume we are given a graph $G = (V, E)$, and the problem is to find a maximum cardinality subset $V' \subseteq V$ of vertices 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 $(s_e, \mathrm{tail}(a_v))$ and $(\mathrm{head}(a_v), t_e)$, with zero cost. This transformation results in an instance of the all-service RTP with $|V|$ tariff arcs, and $|V| + |E|$ clients. Figure 7 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 RTP 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 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 $|E||V|(k + 1) + |E|$.

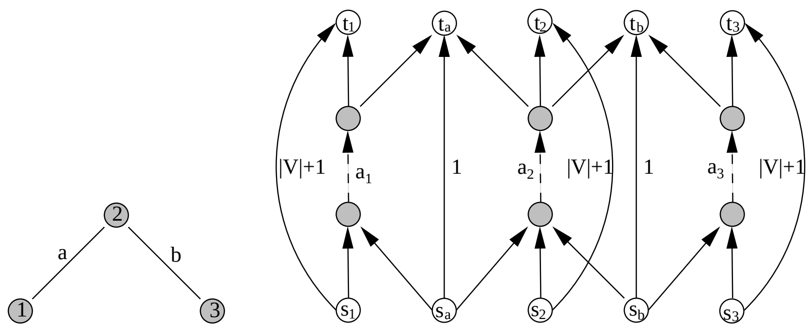

(a) Indep. Set. (b) All-service RTP.

Fig. 7. Reduction of Independent Set to all-service RTP.

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 tarification strategy, 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 RTP, with $\alpha \geq 1$. Consider any instance $G = (V, E)$ of Independent Set, and the all-service RTP 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 $\Pi^{\mathrm{RTP}} = |E||V|(k+1)+|E|$ 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 $\mathcal{O}(\alpha)$–approximation algorithm for the Independent Set problem.

It follows from Håstad [9] that the Independent Set problem cannot have a polynomial time approximation algorithm with worst case guarantee $\mathcal{O}(|V|^{1/2-\varepsilon})$ unless $\mathcal{P} = \mathcal{NP}$, and that it cannot have a polynomial time approximation algorithm with worst case guarantee $\mathcal{O}(|V|^{1-\varepsilon})$ unless $\mathcal{ZPP} = \mathcal{NP}$. 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 $|V| + |E| \in \mathcal{O}(|V|^2)$, the second claim follows. □

Notice that this inapproximability result shows that 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 8. *There exists an n-approximation algorithm for the all-service river tarification 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 affordable 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 if this yields a feasible solution with all clients served or not. □

Acknowledgements

We thank Maxim Sviridenko for an insightful discussion and the suggestion for the proof of Theorem 8, and the referees of a preliminary version of this paper [7] for some helpful remarks.

References

1. G. Aggarwal, T. Feder, R. Motwani, and A.Zhu. Algorithms for multi-product pricing. *Proceedings of the 31st International Colloquium on Automata, Languages and Programming*, Lecture Notes in Computer Science 3142:72–83, Springer, Berlin, 2004.

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.
4. R. Cole, Y. Dodis, and T. Roughgarden. How much can taxes help selfish routing? In *Proceedings of the 4th Annual ACM Conference on Electronic Commerce*, pages 98–107, 2003.
5. R. Cole, Y. Dodis, and T. Roughgarden. Pricing network edges for heterogeneous selfish users. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, pages 521–530, 2003.
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, S. van Hoesel, A.F. van der Kraaij, M. Uetz, and M. Bouhtou Pricing network edges to cross a river. In *Proceedings of the 2nd Int. Workshop on Approximation and Online Algorithms*, Lecture Notes in Computer Science 3351:140–153, Springer, Berlin, 2005.
8. V. Guruswami, J.D. Hartline, A.R. Karlin, D. Kempe, C. Kenyon, F. McSherry. On profit-maximizing envy-free pricing. *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, to appear, 2005.
9. J. Håstad. Clique is hard to approximate within $n^{1-\varepsilon}$. *Acta Mathematica*, 182:105–142, 1999.
10. R.M. Karp and J. B. Orlin. Parametric shortest path algorithms with an application to cyclic staffing. *Discrete Applied Mathematics*, 3:37–45, 1981.
11. M. Labbé, P. Marcotte, and G. Savard. A bilevel model of taxation and its application to optimal highway pricing. *Management Science*, 44:1608–1622, 1998.
12. S. Roch, G. Savard, and P. Marcotte. Design and Analysis of an approximation algorithm for Stackelberg network pricing. *Optimization Online*, 2003. `www.optimization-online.org/DB_HTML/2003/03/621.html`
13. T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the Association for Computing Machinery*, 49(2):236–259, 2002.
14. S. van Hoesel, A. F. van der Kraaij, C. Mannino, G. Oriolo, and M. Bouhtou. Polynomial cases of the tarification problem. *Meteor Resarch Memorandum* RM03053, Maastricht University, 2003.
15. N. E. Young, R. E. Tarjan, and J. B. Orlin. Faster parametric shortest path and minimum-balance algorithms. *Networks*, 21:205–221, 1991.