

Pricing Network Edges to Cross a River

Citation for published version (APA):

Grigoriev, A., van Hoesel, C. P. M., van der Kraaij, A. F., Uetz, M. J., & Bouhtou, M. (2004). Pricing Network Edges to Cross a River. In *Approximation and Online Algorithms: WAOA 2004* (Vol. 3351). Springer. Lecture Notes in Computer Science <https://doi.org/10.1007/b106130>

Document status and date:

Published: 01/01/2004

DOI:

[10.1007/b106130](https://doi.org/10.1007/b106130)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Pricing Network Edges to Cross a River^{*}

Alexander Grigoriev¹, Stan van Hoesel¹, Anton F. van der Kraaij¹,
Marc Uetz¹, and Mustapha Bouhtou²

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 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 networks. Tariffs have to be defined by an operator, the leader, for a subset of the arcs, the tariff arcs. Clients, the followers, choose paths to route their demand through the network selfishly and independently of each other, on the basis of minimal cost. Assuming there exist bounds on the costs clients are willing to bear, the problem is to find tariffs such as to maximize the operator's revenue. Except for the case of a single client, no approximation algorithm is known to date for that problem. We derive the first approximation algorithms for the case of multiple clients. Our results hold for a restricted version of the problem where each client takes at most one tariff arc to route the demand. We prove that this problem is still strongly \mathcal{NP} -hard. Moreover, we show that uniform pricing yields both an m -approximation, and a $(1 + \ln D)$ -approximation. Here, m is the number of tariff arcs, and D is upper bounded by the total demand. We furthermore derive lower and upper bounds for the approximability of the pricing problem where the operator must serve all clients, and we discuss some polynomial special cases. A computational study with instances from France Télécom suggests that uniform pricing performs better than theory would suggest.

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 set the tariffs on these arcs for renting capacity to one or several clients. Each client wishes to route a certain amount of a commodity on a path connecting two vertices. Such a path can involve one or several arcs belonging to the operator, and we assume

^{*} This research was partially supported by France Télécom Research & Development.

that each client selfishly selects a path with minimum 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.

The problem we consider here is different in two aspects from the network congestion problems studied recently, e.g., by Roughgarden and Tardos [11], and Cole et al. [2, 3]. First, we assume that there is no congestion, hence the clients do not influence each other. They choose minimum cost paths to route their commodities, independent of each other. The Game Theoretic setting is only introduced by the fact that 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 (leader) first sets the tariffs, and then, subject to these tariffs, the clients (followers) react selfishly.

A natural formulation of the problem, referred to as the (general) tariffication problem, is the bilevel linear formulation of Labbé et al. [9]. They show that already the problem with a single client is (strongly) NP-hard, given that also negative tariffs are allowed. Roch et al. [10] show that the single client problem remains (strongly) NP-hard, even when restricted to nonnegative tariffs. In the same paper, a polynomial time $\mathcal{O}(\ln m)$ -approximation algorithm for the problem with a single client is proposed, where m is the number of tariff arcs.

Our Results. We derive the first approximation results for the problem with multiple clients. However, we consider a restricted variant of the problem, since we assume that the path taken by any client utilizes at most one tariff arc. Several applications of this particular tariffication problem, to which we refer as the *river tariffication problem* (RTP) are briefly discussed in Section 2. Section 3 describes the model in detail. In Section 4, we show that the river tariffication problem is (strongly) NP-hard.

The quality of uniform tariffication policies, where all arcs are priced with the same tariff, is analyzed in Section 5. The problem to find an optimal uniform tariff is well-known to be solvable in polynomial time, even for the general tariffication problem [12]. We show that uniform tariffication is an m -approximation, and this is tight. Using a simple geometric argument, we also show that uniform tariffication 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 also consider another variant of the problem where the operator is forced to serve all clients. We show in Section 6, by a reduction from INDEPENDENT SET, that this problem is not approximable 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.

We briefly discuss some polynomially solvable special cases of the river tariffication problem in Section 7. Finally, we empirically analyze the quality of uniform tariffication policies in Section 8, using instances from France Télécom.

2 Applications

As an illustration, consider transportation networks that resemble the situation of a town that is divided by a river. Different traversal possibilities exist, and some of these are to be priced by an operator. These traversal possibilities are the tariff arcs in the network. Customers want to route certain commodities from one side of the river to the other.

Such a network topology may be assumed (after a simple transformation described below) in telecommunication networks where we know a priori that the path of each client takes 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 this particular country, we can assume that it is not profitable for any client to enter the country twice.

For another 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, thereby creating an instance of the river tariffication problem.

Finally, in point-to-point markets, a telecommunications operator is offering bandwidth capacity between two points in different qualities of service (QoS). In that setting, it is often the case that information is available concerning the prices customers are willing to pay for different levels of QoS. That pricing problem can be modelled easily as a river tariffication problem, too.

3 Model

An instance of the general tariffication 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) $k \in \{1, \dots, n\}$, where 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. We define for a commodity k the set of all possible paths from s_k to t_k by P_k . The tariff on a tariff arc $a \in T$ is denoted by τ_a , and the vector of all tariffs is given by $\tau = (\tau_a)_{a \in T}$. The cost of a fixed cost arc $a \in F$ is denoted by c_a .

The clients route their demands from source to destination according to a path with minimal total cost, where the total 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 all tariffs by some small value ε . We assume that an $\{s_k, t_k\}$ -path exists that consists 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 tariffication problem is a classical Stackelberg Game that can be modelled as (linear-linear) bilevel program [9, 1]. It follows from Jeroslow [7] that (linear-linear) bilevel programs are \mathcal{NP} -hard in general. For annotated bibliographies on bilevel programming, see Vicente and Calamai [13] or Dempe [4].

We next describe a simple transformation of the given graph G that allows us to restrict to very specific graphs (although probably losing certain graph properties, such as planarity). When replacing shortest paths using only fixed cost arcs by direct arcs, and possibly introducing some dummy arcs with zero or infinite cost, one obtains a shortest path graph model (SPGM) as defined by Bouhtou et al. [1]. In that model, all tariff arcs are disjoint, and there exists an arc from any source node s_k to the tail node of any tariff arc, and from the head node of any tariff arc to any target node t_k . Moreover, there exists a fixed cost arc (s_k, t_k) for all $k = 1, \dots, n$, and the cost c_k is the highest acceptable price for client k .

The additional assumption in the problem considered in this paper (to which we refer as the river tariffication problem) is the following: Independent of the tariffs, any client routes his demand only on a path that includes at most one tariff arc. In the shortest graph path model, that is equivalent to the deletion of any backward-arc that might exist between the head nodes of tariff arcs back to tail nodes of other tariff arcs. Figure 1 shows the shortest path graph model of an instance of the river tariffication problem with three tariff arcs and two clients. The tariff arcs $a_i, i \in \{1, 2, 3\}$ are given by 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 (by adding their costs to the fixed cost




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

arcs incident with s_k). Notice that the only difference to the general tariffication problem described above is the non-existence of backward arcs.

The essential parameters that define an instance of a (river) tariffication problem are therefore the number of tariff arcs m , the number of clients n , the demand values d_k , $k \in \{1, \dots, n\}$, and the costs c_a of the fixed cost arcs $a \in F$. 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_a of the fixed cost arcs $a \in F$ 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 hold for the general tariffication problem, where a path chosen by a client can consist of more than one tariff arc.

4 Complexity

Roch et al. [10] show that the general tariffication 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 [5]. Their reduction works for tariffication problems where paths are allowed to use (and indeed, must use) several tariff arcs. We show that the tariffication problem with multiple clients, but restricted to at most one tariff arc per path, is \mathcal{NP} -hard in the strong sense, too.

We also use a reduction from 3-SAT. Therefore, consider a boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ on n variables x_1, \dots, x_n , in conjunctive normal form. Such a function f is the conjunction of m clauses C_k ,

$$f = \bigwedge_{k=1}^m C_k, \quad (1)$$

each clause C_k being the disjunction of three literals, $C_k = (\ell_{k1} \vee \ell_{k2} \vee \ell_{k3})$. Any literal ℓ_{kj} represents either a variable x_i , or its negation \bar{x}_i , $i \in \{1, \dots, n\}$. Then f is satisfiable if there exists a truth assignment x_1, \dots, x_n such that at least one literal per clause is true.

Any function of the form (1) can be polynomially transformed to an instance of the river tariffication problem as follows. For each variable x_i , $i \in \{1, \dots, n\}$, we construct a constant-size subnetwork as shown in Figure 2. 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 $x_i = 0$, as depicted in Figure 2.

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 3. For commodity 2, the upper bound on the cost is given by a fixed cost arc (s_{i2}, t_{i2}) , with cost 2. The maximal revenue for each subnetwork is thus given by setting one of the tariffs to 2, and the other to 3, yielding a revenue of $2 \cdot 2 + 3 = 7$. In all other cases, the revenue is not more than 6.




Fig. 2. Subnetwork for variable x_i , $i \in \{1, \dots, n\}$

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 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 2, defining an upper bound of 2 for the cost of routing clause-commodity k . The so-defined instance of the river tariffication problem has $2n$ tariff arcs, $3n + m$ commodities (or clients), and $7m + 11n$ fixed cost arcs, hence the transformation is indeed polynomial.

Theorem 1. *The river tariffication problem is strongly NP-hard.*

Proof. Consider the polynomial transformation defined previously. It is now straightforward to show that a satisfying truth assignment for f exists if and only if the revenue for the river tariffication problem is equal to $2m + 7n$. \square

The reduction used for the proof of Theorem 1 shows that the river tariffication problem remains NP-hard even for unit demands, a fixed number of tariff values and when the operator is forced to use tariffs such that he serves (a given subset of) all clients.

5 The Quality of Uniform Tariffication Policies

The uniform tariffication problem (UTP) is the same as the general tariffication problem, with the additional restriction that all tariffs are required to be identical. As shown by van Hoesel et al. [12], the uniform tariffication problem can be solved in polynomial time, even in the general setting where clients may use paths with several tariff arcs. The algorithm described in van Hoesel et al. [12] uses the parametric shortest path algorithm of Young et al. [14] and Karp and Orlin [8] 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 tariffication policy in polynomial time.

We next analyze the loss that can be experienced by adopting such a uniform tariffication policy for the river tariffication problem. Apart from theoretical interest, the question is motivated by the interest in efficient tariffication strategies in a more general setting with more than one operator. In addition, although it is quite easy to think of smarter tariffication policies, so far all these policies, except uniform tariffication, resisted our attempts of a worst-case analysis.

Therefore, denote by Π^{UTP} the revenue for an optimal uniform tariffication, and by Π^{OPT} the revenue for an optimal non-uniform tariffication. By definition, $\Pi^{UTP} \leq \Pi^{OPT}$.

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

The proof of this lemma is indeed trivial. To this end, consider an optimal non-uniform tariffication with 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 $f : [0, D] \rightarrow [0, \infty[$.

$$f(x) = \tau_i \quad \text{for all } x \text{ with } \sum_{j < i} D_j \leq x < \sum_{j \leq i} D_j, \quad i \in \{1, \dots, m\}. \quad (2)$$

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 $\Pi^{UTP} \geq T_i$ 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} . (Notice that this does not hold for the general tariffication 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 $r \cdot \Pi^{UTP} \geq \sum_{j=1}^r T_{i_j} \geq \Pi^{OPT}$. \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 tariffication is an m -approximation for the river tariffication problem.*

We next derive another bound on the quality of uniform tariffication policies, using a geometric argument.

Theorem 3. *Uniform tariffication is a $(1 + \ln D)$ -approximation for the river tariffication 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 tariffication, and recall the definition of the corresponding staircase function f in (2), 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

$$\Pi^{\text{UTP}} \geq \frac{\Pi^{\text{OPT}}}{1 + \ln(D\tau_{\max}/T_\ell)}. \quad (3)$$

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

$$g(x) := \frac{T_\ell}{D-x} \text{ for } x \in [0, D). \quad (4)$$

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^{OPT} , can be upper bounded in terms of the area defined by the function $g(x)$, as depicted in Figure 3. 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.




Fig. 3. Illustration for the proof of Theorem 3

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

and since $T_\ell \leq \Pi^{\text{UTP}}$, claim (3) follows. \square

Notice that claim (3) 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 ratio of $(1 + \ln 4) \approx 2.4$ for uniform tariffication. Geometric intuition indeed suggests a ratio of roughly 2, the additional 0.4 being caused by the difference between the functions $g(x)$ and $f(x)$. In Section 8, we compare the quality of uniform versus non-uniform tariffication, 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 tariffication is a $(1 + \ln n)$ -approximation for the river tariffication problem.*

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

Example 1. Given $n=m$ commodities and m tariff arcs. Every commodity is operating its own subnetwork with one tariff arc, thus the entire network consists of m disjoint subnetworks and each of them contains one commodity and one tariff arc. Fix $b > 1$ and let the demand in subnetwork k be given by $d_k = b^k - b^{k-1}$, $k \in \{1, \dots, m\}$. This way, the total demand D equals $b^m - 1$. Moreover, the maximal revenue for subnetwork k is limited by a fixed cost arc (s_k, t_k) , with cost $c_k = b^{2m-k}$. Hence, the maximal tariff τ_{\max} equals b^{2m-1} . \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 $\Pi^{\text{OPT}} = m(b^{2m} - b^{2m-1})$. The optimal uniform tariffication 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} , $k \in \{1, \dots, m-1\}$. This yields a total revenue of $b^{2m} - b^{2m-k}$, which is less. Therefore, we obtain

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

Now, observe that in the optimal solution m different tariffs are utilized. Lemma 1 (Theorem 2, respectively) suggests that uniform tariffication 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 tariffication 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 tariffication 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 tariffication, the performance bound of Theorem 2 is best possible, and the performance bound of Theorem 3 is best possible up to a constant factor.*

6 All-Service River Tarification Problem

In this section, we consider the following variation of the river tariffication 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 tariffication problem. NP-hardness of this problem follows by our previous reduction presented in Section 4.

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 well.

Theorem 5. *For any $\varepsilon > 0$, the existence of a polynomial time approximation algorithm for the all-service river tariffication problem 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 [5] 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 $(s_e, \text{tail}(a_v))$ and $(\text{head}(a_v), t_e)$, with zero cost. This transformation results in an instance of the all-service RTP with $|V|$ tariff arcs, and $|V| + |E|$ clients. Figure 4 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 $e = (v, u) \in E$




Fig. 4. Reduction of INDEPENDENT SET to all-service RTP

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 $|E||V|(k+1) + |E|$.

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 tariffication 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 $\Pi^{\text{OPT}} = |E||V|(k + 1) + |E|$ 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 is now well known from work of Håstad [6] 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. \square

On the positive side, we can show the following.

Theorem 6. *There exists an n -approximation algorithm for the all-service river tariffication problem.*

The proof works by enumeration over all $m \cdot n$ possibilities for a maximum revenue client using a specific arc. Given that arc-client pair, we can find a corresponding optimal tariff for that arc in polynomial time using binary search, in each step solving a system of linear inequalities. We skip the details due to space limitations.

7 Polynomially Solvable Special Cases

Several polynomially solvable special cases of the (general) tariffication problem are discussed by Labbé et al. [9] and van Hoesel et al. [12]. Clearly, these results hold for the problem considered in this paper, too.

In addition, the river tariffication problem is also polynomially solvable if the number of clients n is bounded from above by a constant. In that case, the number of assignments of clients to tariff arcs is bounded by m^n which is a polynomial for fixed n . Consider therefore the following formulation, where we use notation as given next. The path taken by each client in the network is denoted by $p_k^* \in P_k$, and P_k represents the set of all possible paths taken by a client $k \in K$. The revenue associated with a path $p \in P_k$ induced by a client k with demand d_k is defined by $\pi_p(\tau, d_k) = d_k \cdot \tau_a$, where a is the (unique) tariff arc on path p . The fixed cost of a path p is given by $c_p(d_k) = d_k \sum_{a \in F \cap p} c_a$. Then $l_p(\tau, d_k) := c_p(d_k) + \pi_p(\tau, d_k)$ is the total cost of the path $p \in P_k$ for a client k .

$$\begin{aligned} \max_{\tau} \quad & \sum_{k \in K} \pi_{p_k^*}(\tau, d_k) \\ \text{s.t.} \quad & l_p(\tau, d_k) \geq l_{p_k^*}(\tau, d_k) \quad \forall k \in K, \forall p \in P_k \\ & \tau_a \geq 0 \quad \forall a \in T \end{aligned} \tag{5}$$

Since for each client, there are at most $m + 1$ paths in the network, $|P_k|$ is bounded by $m + 1$. Hence, the number of constraints is polynomial in the input data. Therefore, if we solve m^n instances of (5), we can retrieve the optimal solution in polynomial time.

8 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 (2), is close to a straight line, geometric intuition suggests a worst-case ratio for uniform tariffication 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 tariffication 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 Tariffication on France Télécom instances

Instance	$ N $	$ A $	m	n	Π^{OPT}	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 2. We compare the optimal solutions for uniform tariffs (Π^{UTP}) and non-uniform tariffs (Π^{OPT}). The optimal non-uniform solution is calculated using the model and mixed integer programming formulation described in Bouhtou et al. [1]. The value of Π^{UTP} is calculated using the same formulation, requiring that all tariffs be equal. As such, we do not compare the actual computation times, but are just interested in effectiveness of the optimal uniform tariffication policies. Table 1 gives a brief description of each network, stating the number of nodes, arcs, tariff arcs and clients. The optimal non-uniform and uniform solution values are displayed in the columns Π^{OPT} and Π^{UTP} . The final column is the approximation ratio.

Acknowledgements

We thank Maxim Sviridenko for an insightful discussion and the suggestion for the proof of Theorem 6, and the anonymous referees for some helpful remarks.

References

1. 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.

2. 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.
3. 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.
4. S. Dempe. Annotated bibliography on bilevel programming and mathematical programs with equilibrium constraints. *Optimization*, 52:333–359, 2003.
5. M.R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman, San Francisco, 1979.
6. J. Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, 182:105–142, 1999.
7. R.G. Jeroslow. The polynomial hierarchy and a simple model for competitive analysis. *Mathematical Programming*, 32:146–164, 1985.
8. R.M. Karp and J. B. Orlin. Parametric shortest path algorithms with an application to cyclic staffing. *Discrete Applied Mathematics*, 3:37–45, 1981.
9. 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.
10. S. Roch, G. Savard, and P. Marcotte. Design and Analysis of an approximation algorithm for Stackelberg network pricing. *Optimization Online*, 2003.
11. T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the Association for Computing Machinery*, 49(2):236–259, 2002.
12. S. van Hoesel, A. F. van der Kraaij, C. Mannino, G. Oriolo, and M. Bouhtou. Polynomial cases of the tariffication problem. *Meteor Research Memorandum* RM03053, Maastricht University, 2003.
13. L.N. Vicente and P. H. Calamai. Bilevel and multilevel programming: A bibliography review. *Journal of Global Optimization*, 5:291–306, 1994.
14. N. E. Young, R. E. Tarjan, and J. B. Orlin. Faster parametric shortest path and minimum-balance algorithms. *Networks*, 21:205–221, 1991.