



HAL
open science

Pricing the Internet with Multi-Bid Auctions

Patrick Maillé, Bruno Tuffin

► **To cite this version:**

Patrick Maillé, Bruno Tuffin. Pricing the Internet with Multi-Bid Auctions. [Research Report] RR-5235, INRIA. 2004, pp.31. inria-00070763

HAL Id: inria-00070763

<https://inria.hal.science/inria-00070763>

Submitted on 19 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Pricing the Internet with Multi-Bid Auctions

Patrick Maillé , Bruno Tuffin

N°5235

Juin 2004

————— Systèmes communicants —————

 ***Rapport
de recherche***

Pricing the Internet with Multi-Bid Auctions

Patrick Maillé *, Bruno Tuffin †

Systemes communicants
Projet Armor

Rapport de recherche n° 5235 — Juin 2004 — 31 pages

Abstract: Usage-based or congestion-based charging schemes have been regarded as a relevant way to control congestion and to differentiate services among users in telecommunication networks; auctioning for bandwidth appears as one of several possibilities. In a previous work, the authors designed a multi-bid auction scheme where users compete for bandwidth at a link by submitting several couples (amount of bandwidth asked, associated unit price) so that the link allocates the bandwidth and computes the charge according to the second price principle. They showed that incentive compatibility and efficiency among other properties are verified. We propose in the present paper to extend this scheme to the case of a network, by using the properties/assumptions that the backbone network is overprovisionned and the access networks have a tree structure.

Key-words: Congestion control, Economics

(Résumé : tsvp)

* pmaill@rennes.enst-bretagne.fr

† btuffin@irisa.fr

Tarification de l'Internet par enchères multiples

Résumé : Les méthodes de tarification basées sur l'usage et/ou la congestion sont considérées comme une manière pertinente de contrôler la congestion et de différencier les services entre utilisateurs dans les réseaux de communication ; les enchères pour la bande passante apparaissent comme l'une des possibilités. Dans un travail précédent, les auteurs ont construit un mécanisme d'enchères où les utilisateurs sont en compétition pour la bande passante à un lien de communication et soumettent plusieurs couples (quantité de bande passante demandée, prix unitaire associé) en fonction de quoi le lien alloue la bande passante et calcule le coût selon le principe du second prix. Ils ont démontré que la compatibilité d'incitation et l'efficacité entre autres propriétés sont vérifiées. Nous proposons dans cet article d'étendre cette méthode au cas d'un réseau, en utilisant les propriétés/hypothèses que le cœur de réseau est surdimensionné et que les réseaux d'accès ont une structure en arbre.

Mots-clé : Contrôle de congestion, Economie

1 Introduction

Many people argue that congestion pricing for communication networks is a pointless problem, since the available capacities are (or will soon be) so large that congestion will never occur. This might be true for backbone networks, that are steadily re-dimensioned using the latest available technologies and allow tremendous communication rates. However, moving data from the network backbone into houses remains a major challenge for bandwidth-demanding applications such as for instance digital-video broadcast, since wires used in local loops are provisioned for voice-grade analog service and do not allow high-speed data services. This problem is known as the “last mile bottleneck”; while major backbone routes are awash in optical fiber that they will not use for years to come, replacing all “last mile” communication links with optical fibers would be too costly (estimated at \$500-\$1500 per household [2]), and therefore is not likely to be applied.

Consequently, when several users share the same access network (like in hotel rooms, apartment units, offices or other multiunit buildings), congestion is very likely to occur. This stands especially when considering new services that are more and more bandwidth-consuming. Therefore, a fair/efficient way to share the available resources among users needs to be found. The problem lies essentially on incentives: how can we force selfish users to cooperate and share the resource efficiently? *Pricing* appears as a solution, and has become a topic of high interest (see [6, 18, 7] and references therein) thanks to its influence on users’ behaviour (see for instance the experiment that has been conducted at UC Berkeley [1]).

Currently, Internet communications are priced independently of usage, which is an incentive to overuse the network [17]. We focus in this paper on pricing schemes based on bandwidth allocation, since bandwidth is the limited resource.

In [13], we have designed the so-called multi-bid auction pricing mechanism, a one-shot auction-based scheme to allocate the capacity of a single communication link among several users. We have suggested that each user submit several bids when establishing a connection, and that the corresponding multi-bid profile be used to compute efficient allocations and prices. The multi-bid scheme, that is highly related to the Progressive Second Price mechanism (PSP) of Lazar and Semret [11], has been shown to verify several desirable properties in terms of incentives and efficiency. Moreover, due to the time of convergence and the signaling overhead needed for PSP versus the one-shot allocation of multi-bid, we have proved that multi-bid auctions are better adapted to the pricing of Internet communications, that open and close over time (the loss of efficiency of PSP in this context was studied in [12]).

However, [13] only considers the case of a single communication link, whereas PSP was also shown to apply to interconnected networks [15]. This paper aims at extending multi-bid auctions to the case of a network. Since our goal is to charge for Internet communications, we use the particular structure of the current Internet network, which involves an overprovisionned backbone network and access networks with a tree structure. Indeed, each access network can be modeled this way, as described in [3] : customers are not directly connected to high-speed backbone networks, but rather to “local” access networks, which are then connected to regional networks, themselves being connected to national backbone networks.

We therefore design the multi-bid auction scheme in that context of a tree structure, and investigate the properties of this mechanism, especially focusing on incentives and efficiency.

The paper is organized as follows. Section 2 presents the model considered in this paper by describing the network topology and users behavior. We describe the principle of multi-bid auctions and some basic definitions in Section 3. Section 4 then presents how the mechanism introduced in [13] can be adapted to the case of an access network, and the desirable theoretical results (individual rationality, incentive compatibility, efficiency) are respectively derived in Sections 5, 6, 7. Section 8 deals with the bids repartition while Section 9 studies the number of bids that should be allowed by the auctioneer. We conclude in Section 10 by giving some directions for future work.

Note that some (but not all) proofs and discussions are similar to those presented in [13] for the case of a single link. Nevertheless, we include them in this paper to make it self-contained.

2 The model

2.1 The network

We assume in this paper that congestion may only occur in access networks. Access networks (the last mile) are considered here to have a tree structure, i.e. for each user there is one and only one path to reach the congestion-free backbone network. This vision of a network is represented in Fig. 1, illustrating that several users may have to share some links in the access network to reach the backbone [3]. Moreover, we assume that users wish to access

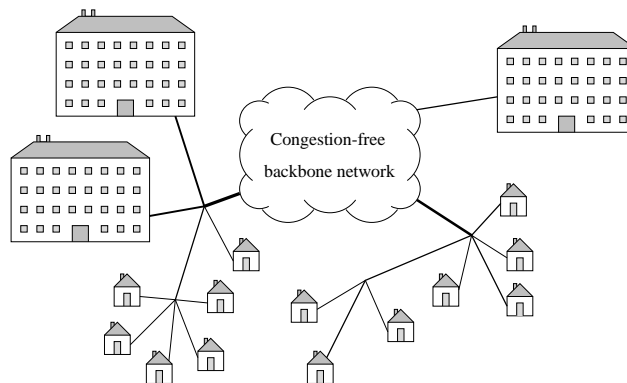


Figure 1: Network topology

services which are located in the backbone network. Therefore, a user may only create congestion in her own access network, due to the overdimensioning of the backbone. The

problem can consequently be reduced to a problem of resource allocation on independent access networks modelled by a tree structure as in Fig. 2.




Figure 2: An access network : arrows represent the users' paths to the backbone network. The thickness of the link is for its capacity.

We therefore consider a tree network with a set \mathcal{L} of links, and denote by Q^l the finite capacity of link l for $l \in \mathcal{L}$. We assume that the capacity of each link is infinitely divisible. The set of users who may use this access network is denoted by \mathcal{I} . Throughout this paper, we will use subscripts to refer to players, and superscripts to refer to links. Considering that all players in \mathcal{I} do not use all the links in \mathcal{L} , we point out those required by player i to reach the backbone network by the binary values

$$r_i^l = \begin{cases} 1 & \text{if player } i \text{ uses link } l \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

We also define the route of player i by $r_i = (r_i^l)_{l \in \mathcal{L}}$. In the example of Fig. 2, $r_1^l = 1$ for $l = l^1, l^2$ and $r_1^l = 0$ for $l = l^3, l^4, l^5, l^6$ for instance. Notice that if l^{root} denotes the link directly connected to the root of the tree, i.e. to the backbone ($l^{root} = l^1$ in Fig. 2), then $r_i^{l^{root}} = 1$ for all $i \in \mathcal{I}$.

2.2 User preferences

Since we assume in this paper that users will react to a pricing scheme so as to maximize selfishly their utility, we will study our mechanism in the framework of *game theory* [8], and then talk indifferently of users and players.

We suppose as in [13, 11, 16] that the only performance measure users are sensitive to is the bandwidth allocated to them. Of course they are sensitive to the price they are charged as well. We assume that their preferences are represented by quasi-linear utility functions.

This means that for $i \in \mathcal{I}$, user i 's utility is the difference between her valuation $\theta_i(a_i)$ of allocation a_i (her willingness-to-pay) and the price c_i that she actually pays:

$$U_i = \theta_i(a_i) - c_i. \quad (2)$$

The theoretical results derived in this paper are obtained for players with elastic demand, and smooth valuation functions such that

Assumption 1. For any $i \in \mathcal{I}$,

- θ_i is differentiable and $\theta_i(0) = 0$,
- θ'_i is positive, nonincreasing and continuous
- $\exists \gamma_i > 0, \forall z \geq 0, \theta'_i(z) > 0 \Rightarrow \forall \eta < z, \theta'_i(z) \leq \theta'_i(\eta) - \gamma_i(z - \eta)$.

Assumption 2. $\exists \kappa > 0, \forall i \in \mathcal{I}$,

- $\theta'_i(0) < +\infty$
- $\forall z, z', z > z' \geq 0, \theta'_i(z) - \theta'_i(z') > -\kappa(z - z')$.

Notice that these assumptions were first introduced in [11] for the analysis of Progressive Second Price auctions.

3 Multi-bid auctions: message process and basic definitions

In [11], Lazar and Semret introduced the Progressive Second Price auctions, suggesting that players should submit two-dimensional bids of the form $s_i = (q_i, p_i)$: q_i represents the amount of bandwidth that player i is asking and p_i the unit price that she is accepting to pay to get this quantity. The multi-bid auction scheme we are going to describe now allows players to submit simultaneously several such two-dimensional bids.

The message process is as follows. A player i entering the game (i.e. establishing a connection) submits a set of M_i two-dimensional bids $s_i = \{s_i^1, \dots, s_i^{M_i}\}$. In a multi-bid s_i , for all $m, 1 \leq m \leq M_i, s_i^m$ is a two-dimensional bid as defined in the PSP scheme: $s_i^m = (q_i^m, p_i^m) \in \mathbb{R}_+^2$. We assume without loss of generality that bids are sorted such that $p_i^1 \leq p_i^2 \leq \dots \leq p_i^{M_i}$. From the *multi-bid profile* $s = (s_i)_{i \in \mathcal{I}}$, the auctioneer computes for each player $i \in \mathcal{I}$ the allocation a_i on player i 's route and the price c_i she will be charged. To emphasize on player i 's multi-bid, we will also sometimes write $s = (s_i, s_{-i})$, where $s_{-i} = (s_j)_{j \in \mathcal{I} \setminus \{i\}}$.

In the following, S denotes the set of multi-bids that a player can submit:

$$S = \bigcup_{M \geq 0} (\mathbb{R}^+ \times \mathbb{R}^+)^M, \quad \text{with } (\mathbb{R}^+ \times \mathbb{R}^+)^0 = \emptyset.$$

Remark: If player i submits no bid ($M_i = 0$) then we write $s_i = \emptyset$.

Definition 1. We say that a player $i \in \mathcal{I}$ submits a truthful multi-bid $s_i \in S$ if $s_i = \emptyset$, or if

$$\forall m, 1 \leq m \leq M_i, p_i^m = \theta'_i(q_i^m).$$

This means that player i actually reveals her marginal value $\theta'_i(q_i)$ if she obtains quantity q_i .

We write S_i^T the set of truthful multi-bids that can be submitted by player i .

For every player $i \in \mathcal{I}$, there is a demand function associated to the valuation function θ_i . This demand function gives the quantity q of resource that player i would buy if the resource were sold at a fixed unit price p , in order to maximize her utility $\theta_i(q) - pq$.

Definition 2. Under Assumption 1, the demand function of player $i \in \mathcal{I}$ is defined as the function

$$d_i(p) = \begin{cases} (\theta'_i)^{-1}(p) & \text{if } 0 \leq p \leq \theta'_i(0) \\ 0 & \text{otherwise.} \end{cases}$$

Notice that Assumption 1 implies that the demand function is well-defined and nonincreasing.

To compute allocations and prices, the auctioneer uses two types of functions for each user i : her *pseudo-demand function* and her *pseudo-marginal valuation function*. Both functions derive from the multi-bid s_i submitted by player i in the following way :

Definition 3. Consider a player $i \in \mathcal{I}$ who has submitted a multi-bid $s_i \in S$.

• We call pseudo-demand function of i associated with s_i the function $\bar{d}_i : \mathbb{R}^+ \rightarrow \mathbb{R}^+$, defined by

$$\bar{d}_i(p) = \begin{cases} 0 & \text{if } s_i = \emptyset \text{ or } p_i^{M_i} < p \\ \max_{1 \leq m \leq M_i} \{q_i^m : p_i^m \geq p\} & \text{otherwise.} \end{cases} \quad (3)$$

• We call pseudo-marginal valuation function of i , associated with s_i , the function $\bar{\theta}'_i : \mathbb{R}^+ \rightarrow \mathbb{R}^+$, defined by

$$\bar{\theta}'_i(q) = \begin{cases} 0 & \text{if } s_i = \emptyset \text{ or } q_i^1 < q \\ \max_{1 \leq m \leq M_i} \{p_i^m : q_i^m \geq q\} & \text{otherwise.} \end{cases} \quad (4)$$

The demand, pseudo-demand, marginal valuation and pseudo-marginal valuation functions are illustrated in Fig. 3, for a truthful multi-bid.

Remark: Both pseudo-demand and pseudo-marginal valuation functions are positive, stair-step, nonincreasing and left-continuous.

For players who bid truthfully, we can compare the pseudo-demand and pseudo-marginal valuation functions to their “real” counterparts:




Figure 3: Demand and pseudo-demand functions (*left*), marginal valuation and pseudo-marginal valuation functions (*right*) for $M_i = 3$ and a truthful multi-bid.

Lemma 1. *Under Assumption 1, if player $i \in \mathcal{I}$ submits a truthful multi-bid s_i then*

$$\bar{d}_i \leq d_i \quad (5)$$

$$\bar{\theta}'_i \leq \theta'_i. \quad (6)$$

Proof. Let $x \in \mathbb{R}^+$. If $\bar{d}_i(x) = 0$ then $\bar{d}_i(x) \leq d_i(x)$ is trivial, since $d_i \geq 0$. If we assume that $\bar{d}_i(x) > 0$, then $s_i \neq \emptyset$ and

$$\begin{aligned} \bar{d}_i(x) &= \max_{1 \leq m \leq M_i} \{q_i^m : p_i^m \geq x\} \\ &= q_i^{m_0} \text{ with } p_i^{m_0} \geq x \\ &= d_i(p_i^{m_0}) \leq d_i(x) \end{aligned}$$

where the nonincreasingness of d_i is used. Relation (5) is then proved.

Relation (6) is established exactly the same way by inverting the roles of prices and quantities. \square

Fig. 3 illustrates Lemma 1.

Lemma 1 can also be used to establish some results on the composition of the pseudo-demand and pseudo-valuation functions:

Lemma 2. $\forall i \in \mathcal{I}, \forall s_i \in S,$

$$\forall x \in \mathbb{R}^+, \quad \bar{\theta}'_i(\bar{d}_i(x^+)^+) \leq x \quad (7)$$

$$s_i \neq \emptyset \Rightarrow \forall x \in [0, p_i^{M_i}], \quad \bar{\theta}'_i(\bar{d}_i(x)) \geq x, \quad (8)$$

where for $f : \mathbb{R} \rightarrow \mathbb{R}$, $f(x^+) = \lim_{z \rightarrow x, z > x} f(z)$ when the limit exists.

Lemma 2 is proved in Appendix 1, and will be used in the rest of the paper to establish our main results.

Based on the pseudo-demand functions of all players, the auctioneer can approximate the total demand function over each link, in order to determine the level of congestion:

Definition 4. Consider a set of players $i \in \mathcal{I}$, each submitting a multi-bid $s_i \in S$. We call aggregated pseudo-demand function associated with the profile $s = (s_i)_{i \in \mathcal{I}}$ the function $\bar{d} : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ defined by

$$\bar{d}(p) = \sum_{i \in \mathcal{I}} \bar{d}_i(p). \quad (9)$$

4 Multi-bid auctions for a tree network

4.1 Modified allocation rule for a single link

In [13], we introduced the multi-bid allocation rule for a single communication link. This rule implies the computation of a “pseudo-market clearing price” \bar{u} , which is defined as the highest unit price p such that aggregated pseudo-demand is strictly above the available capacity Q of the link. Such a \bar{u} always existed because the seller introduced a particular bid to ensure that the resource will not be sold at a unit price below a certain level, called the reserve price. In this paper, we choose not to introduce such a reserve price. Consequently, the definition of the pseudo-market clearing price \bar{u} and the allocation rule need to be extended to the case when the aggregated pseudo-demand is below Q (which means that there is no congestion): we therefore introduce

$$\bar{u}(s, Q) = \begin{cases} \max\{p : \bar{d}(p) > Q\} & \text{if } \bar{d}(0) > Q \\ 0 & \text{if } \bar{d}(0) \leq Q. \end{cases} \quad (10)$$

We can now define the *modified multi-bid allocation* a_i of a player i as

$$a_i(s, Q) = \begin{cases} \bar{d}_i(\bar{u}^+) + \frac{\bar{d}_i(\bar{u}) - \bar{d}_i(\bar{u}^+)}{\bar{d}(\bar{u}) - \bar{d}(\bar{u}^+)} [Q - \bar{d}(\bar{u}^+)] & \text{if } \bar{d}(0) > Q \\ \bar{d}_i(0) & \text{if } \bar{d}(0) \leq Q. \end{cases} \quad (11)$$

When the available capacity Q is not high enough to satisfy pseudo-demand, (10) and (11) ensure that each player receives the bandwidth asked at the lowest price \bar{u}^+ such that supply exceeds pseudo-demand, and that the surplus $Q - \bar{d}(\bar{u}^+)$ is proportionally shared among users who introduced a bid at price \bar{u} , with weights $\bar{d}_i(\bar{u}) - \bar{d}_i(\bar{u}^+)$.

Notice that (10) and (11) are respectively equivalent to the definition of the pseudo-market clearing price and the multi-bid allocation rule in [13] when a reserve price is introduced by the seller.

4.2 Allocation rule for the tree network

The algorithm we introduce here is inspired by the extension of PSP auctions to a tree network (see [10]). The algorithm starts from the leaves of the tree, and works toward the root, computing allocations for each link $l \in \mathcal{L}$ using the above modified multi-bid rule, among the subset of players \mathcal{I}^l whose route includes that link, i.e. $\mathcal{I}^l = \{i \in \mathcal{I} : r_i^l = 1\}$. The idea is to proceed by revising the multi-bid profile used to compute allocations after having

applied the rule (11) in order to take into account the fact that users are only sensitive to the minimum amount of bandwidth they obtain on all links toward the root. The revision of the multi-bid profile ensures that a player will never obtain more resource at an upstream link than what she gets at the current one.

Alg. 1 describes in details how allocations should be computed. It returns a vector

Alg. 1 Allocations on a tree

Input:

- the tree network defined by the set \mathcal{L} of links, and the capacities $Q^l, l \in \mathcal{L}$
 - the set of players \mathcal{I} and their routes $\{r_i, i \in \mathcal{I}\}$
 - the multi-bid profile s .
1. For all players $i \in \mathcal{I}$, define the *revised multi-bid* \underline{s}_i as $\underline{s}_i = s_i$.
 2. Pick a leaf-link $l \in \mathcal{L}$ (i.e. a link with no downstream link), and let $\underline{s}_{\mathcal{I}^l} = (\underline{s}_i)_{i \in \mathcal{I}^l}$.
 - (a) Compute $\bar{u}^l = \bar{u}(\underline{s}_{\mathcal{I}^l}, Q^l)$ and $a_i^l = a_i(\underline{s}_{\mathcal{I}^l}, Q^l)$ for all $i \in \mathcal{I}^l$, applying Equations (10) and (11) to the revised multi-bids $\underline{s}_{\mathcal{I}^l}$, i.e. using the allocation rule on link l .
 - (b) $\forall i \in \mathcal{I}^l$, modify the revised multi-bid the following way:
 - set $\underline{s}_i = \underline{s}_i \setminus \{s_i^m : q_i^m > a_i^l\}$.
 - if $a_i^l > 0$ and $\bar{u}^l > \max\{p_i^m : (q_i^m, p_i^m) \in \underline{s}_i\}$, then set $\underline{s}_i = \underline{s}_i \cup \{(a_i^l, \bar{u}^l)\}$ (we take $\max\{\emptyset\} = -\infty$).
 - (c) Set $\mathcal{L} = \mathcal{L} \setminus \{l\}$, i.e. delete link l from the tree
 3. if $\mathcal{L} \neq \emptyset$ go to 2
 else return $a = (a_1^{root}, a_2^{root}, \dots, a_{|\mathcal{I}|}^{root})$
-

$a = (a_1, \dots, a_{|\mathcal{I}|}) \in \mathbb{R}_+^{\mathcal{I}}$, where we omitted the superscript l^{root} for simplicity of notation. We then suggest that each user $i \in \mathcal{I}$ be allocated the quantity of bandwidth a_i on each link of her route to the root of the tree (i.e. to the backbone network).

Notice that the algorithm ends just after the computation of allocations for the link directly connected to the root of the tree (link l^1 in the example of Fig. 2). The aim of the algorithm is therefore to judiciously modify the revised multi-bid of all players, in order for the allocation rule (11) to have the desirable properties, i.e. making sure that demand (and then allocation) on next links cannot exceed the allocation on the current link. This

is done in Step 2b of the algorithm, and Fig. 4 illustrates that this is equivalent to revising (meaning upper-bounding) the pseudo-demand \bar{d}_i of a player $i \in \mathcal{I}^l$ into \underline{d}_i , in the following way:

$$\underline{d}_i = \min(a_i^l, \bar{d}_i). \quad (12)$$

Consequently, for each player $i \in \mathcal{I}$ we have

$$a_i = a_i^{l^{root}} = \min\{a_i^l, l \in r_i\}. \quad (13)$$

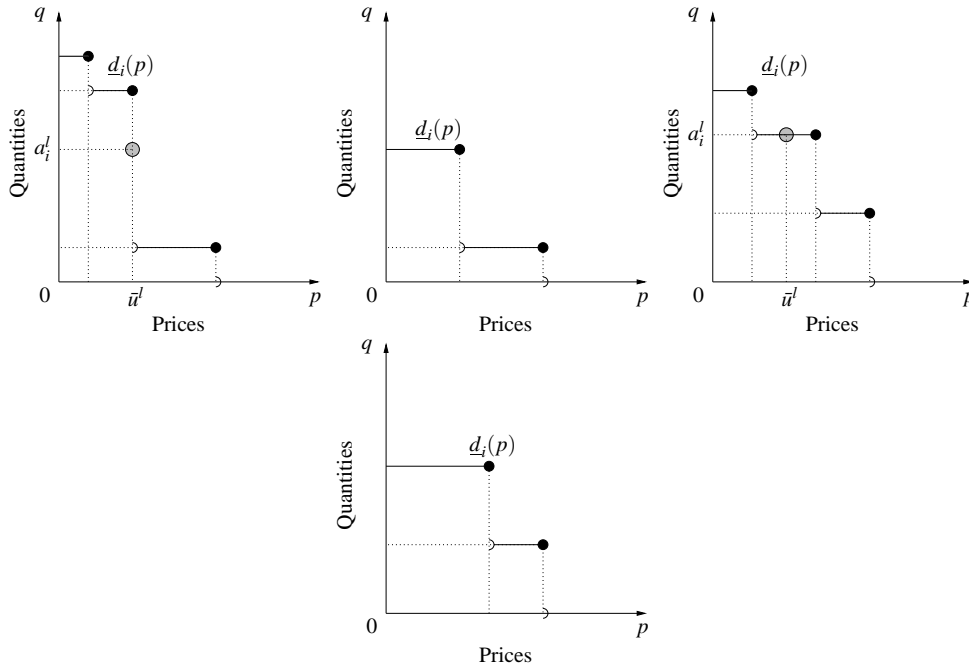


Figure 4: Effect of the revision of multi-bids (Step 2b of Alg. 1): multi-bid of player i (and associated revised pseudo-demand functions) *before* Step 2b (left) and *after* Step 2b (right), i.e. upper-bounded by the current allocation. Two cases are presented: the case when player i submitted a bid at price \bar{u}^l (top) and that when player i did not submit such a bid (bottom).

Notice that allocations on an intermediate link are computed locally, and that the only information needed is the revised multi-bids of all players using that link: each link receives those revised multi-bids from its downstream links, computes the local allocations and the new revised multi-bids, and then transmits them to the link upstream.

4.3 Example

We illustrate here the allocation rule applied to the network of Fig. 2, with link capacities $Q^{l^1} = Q^{l^5} = 10$, $Q^{l^2} = Q^{l^4} = 6$, $Q^{l^3} = Q^{l^6} = 4$. We assume that $\forall i$, player i submits $M_i = 3$ two-dimensional bids in her multi-bid. At the beginning of the algorithm (Step 1), the revised multi-bids \underline{s}_i are equal to the multi-bids s_i submitted by the players. Consider seven users as shown in Fig. 2, submitting the multi-bids of Table 1.

s_1	s_2	s_3	s_4	s_5	s_6	s_7
$\begin{pmatrix} (6,6) \\ (4,8) \\ (3,9) \end{pmatrix}$	$\begin{pmatrix} (6,5) \\ (3,7) \\ (1,8) \end{pmatrix}$	$\begin{pmatrix} (7,2) \\ (4,5) \\ (3,6) \end{pmatrix}$	$\begin{pmatrix} (6,4) \\ (3,5) \\ (2,6) \end{pmatrix}$	$\begin{pmatrix} (8,4) \\ (6,5) \\ (2,7) \end{pmatrix}$	$\begin{pmatrix} (4,1) \\ (2,2) \\ (1,4) \end{pmatrix}$	$\begin{pmatrix} (6,3) \\ (4,4) \\ (3,6) \end{pmatrix}$

Table 1: The multi-bids submitted by the users.

There are four leaf-links: l^2, l^3, l^5, l^6 . Consider link l^2 first. We have $\mathcal{I}^{l^2} = \{1, 2\}$. The computation of the pseudo-market clearing price gives $\bar{u}^{l^2} = 7$, and the corresponding allocation (Eq. (11)) yields $a_1^{l^2} = 4$ and $a_2^{l^2} = 2$. After Step 2b of the algorithm, the revised multi-bids of players in \mathcal{I}^{l^2} are respectively $\underline{s}_1 = ((4, 8); (3, 9))$ and $\underline{s}_2 = ((2, 7); (1, 8))$. We then remove link l^2 from the tree, following Step 2c of the algorithm.

Doing the same on links l^3, l^5 and l^6 respectively leads to the pseudo-market clearing prices $\bar{u}^{l^3} = 4$, $\bar{u}^{l^5} = 0$ and $\bar{u}^{l^6} = 4$, and the revised multi-bids are displayed in Table 2. Next, Step 2 of the algorithm is applied to the (remaining) network of Fig. 5 with those

\underline{s}_1	\underline{s}_2	\underline{s}_3	\underline{s}_4	\underline{s}_5	\underline{s}_6	\underline{s}_7
$\begin{pmatrix} (4,8) \\ (3,9) \end{pmatrix}$	$\begin{pmatrix} (2,7) \\ (1,8) \end{pmatrix}$	$\begin{pmatrix} (7,2) \\ (4,5) \\ (3,6) \end{pmatrix}$	$\begin{pmatrix} (4,4) \\ (3,5) \\ (2,6) \end{pmatrix}$	$\begin{pmatrix} (8,4) \\ (6,5) \\ (2,7) \end{pmatrix}$	$((0.5, 4))$	$\begin{pmatrix} (3.5, 4) \\ (3, 6) \end{pmatrix}$

Table 2: The revised multi-bids after the processing on links l^2, l^3, l^5, l^6 .

revised multi-bids. Now the allocation rule is applied to link l^4 , which has capacity $Q^{l^4} = 6$. Players who compete for capacity on this link are players 5, 6 and 7. The pseudo-market clearing price is $\bar{u}^{l^4} = 5$, and after Step 2b of the algorithm, the revised multi-bids that will be used for the computation of the allocations at link l^1 are given in Table 3.

Finally, the computation of the pseudo-market clearing price for link l^1 gives $\bar{u}^{l^1} = 6$, and the allocations are

a_1	a_2	a_3	a_4	a_5	a_6	a_7
4	2	0.75	0.5	2	0	0.75

It can be checked that the capacity constraints of all links are satisfied with this allocation vector, i.e. $\forall l \in \mathcal{L}, \sum_i r_i^l a_i \leq Q^l$.




Figure 5: The links that remain to be treated after applying the algorithm on links l^2, l^3, l^5, l^6 .

\underline{s}_1	\underline{s}_2	\underline{s}_3	\underline{s}_4	\underline{s}_5	\underline{s}_6	\underline{s}_7
$\begin{pmatrix} (4,8) \\ (3,9) \end{pmatrix}$	$\begin{pmatrix} (2,7) \\ (1,8) \end{pmatrix}$	$\begin{pmatrix} (7,2) \\ (4,5) \\ (3,6) \end{pmatrix}$	$\begin{pmatrix} (4,4) \\ (3,5) \\ (2,6) \end{pmatrix}$	$\begin{pmatrix} (3,5) \\ (2,7) \end{pmatrix}$	\emptyset	$\begin{pmatrix} (3,6) \end{pmatrix}$

Table 3: The revised multi-bids used to compute allocations on link l^1 .

4.4 Multi-bid pricing rule for a tree network

The pricing rule we choose is highly related to Vickrey-Clarke-Groves mechanisms [19, 5, 9]: the idea is that each player should pay for the declared social cost she imposes on others through her presence (this is called the *exclusion-compensation* principle, which lies behind second-price mechanisms). The expression of the price c_i charged to a player i is therefore the same as in the single-link case [13]:

$$c_i(s_i, s_{-i}) = \sum_{j \in \mathcal{I}, j \neq i} \int_{a_j(s)}^{a_j(s_{-i})} \bar{\theta}'_j(q) dq. \quad (14)$$

The computation of $(a_j(s_{-i}))_{i,j \in \mathcal{I}}$ is done during the application of Alg. 1: each link l computes $(a_j^l)_{j \in \mathcal{I}^l}$, taking into account the whole multi-bid profile s , and at the same time does the same with the multi-bid profile $(s_{-i})_{i \in \mathcal{I}^l}$.

On the example of the previous subsection, we obtain the following prices:

c_1	c_2	c_3	c_4	c_5	c_6	c_7
25	12	4.5	3	12	0	4.5

Notice that players who do not obtain some resource are not charged.

4.5 Computational considerations

This section studies the computational complexity of the multi-bid scheme. Let us consider a tree network \mathcal{L} , with a set \mathcal{I} of players, their routes $(r_i)_{i \in \mathcal{I}}$, and let us fix a multi-bid profile s . First remark that the total complexity is shared among all links in the tree, each link l applying the multi-bid allocation rule when it receives the revised multi-bid of all players in \mathcal{I}^l . We therefore focus on the complexity incurred at a link l .

- We place ourselves at a link l , and notice that the revised multi-bid for each player $i \in \mathcal{I}^l$ contains less than M_i two-dimensional bids (Step 2b of Alg. 1 can just reduce the number of two-dimensional bids in \underline{s}_i).

The computation of the aggregated pseudo-demand function needs the bids to be sorted, which can be done in time $O(\sum_{i \in \mathcal{I}^l} M_i \log(\sum_{i \in \mathcal{I}^l} M_i))$. Then the computation of the pseudo-market clearing price \bar{u}^l can be performed in time $O(\sum_{i \in \mathcal{I}^l} M_i)$. Given \bar{u}^l , the value of $\frac{Q^l - \bar{d}(\bar{u}^l)}{d(\bar{u}^l) - \bar{d}(\bar{u}^l)}$ is computed only once, therefore all allocations $(a_i^l)_{i \in \mathcal{I}^l}$ can be calculated with a total complexity $O(\sum_i M_i)$ (computing an allocation a_i^l with (11) can be done with complexity $O(M_i)$).

- To calculate charges, the computation of allocations must be done for all profiles $s_{-i}, i \in \mathcal{I}$. To do this at link l , we just need to compute $a_j(s_{-i})$ for all $i, j \in \mathcal{I}^l$, which leads to a complexity that is less than $O(|\mathcal{I}^l| \sum_{i \in \mathcal{I}^l} M_i)$, where $|\mathcal{I}^l|$ is the number of players in \mathcal{I}^l .

- Once all allocations $a_i(s_{-j})$ are calculated (i.e. at link l^{root}), a price c_i can be computed using (14) with a complexity less than $\sum_i M_i$, since we integrate a stair-step function with less than M_i discontinuity points.

Consequently, for each link l , the computational complexity involved by the multi-bid scheme for a given multi-bid profile is upper-bounded by $O(|\mathcal{I}^l| \sum_{i \in \mathcal{I}^l} M_i)$. If all players submit the same number of bids, (i.e. $\forall i \in \mathcal{I}, M_i = M$), then the complexity at each link l is less than $O(M|\mathcal{I}^l|^2)$.

Notice that PSP allocations and prices at link l can be computed with complexity $O(|\mathcal{I}^l|^2)$ (see [15]). Therefore, the computational time for both methods is of the order $|\mathcal{I}^l|^2$, this being multiplied by the number of bids for the multi-bid algorithm. However, the PSP has to compute allocations and prices several times (until the equilibrium is reached), and we believe that even if the convergence of PSP is fast (less than M iterations), the gain in signaling overhead is worth the cost in computational time.

5 General properties of the multi-bid scheme

This section establishes some basic properties of the multi-bid scheme for a tree network, as extensions of those proved in [13]. We first introduce some definitions that will play a central role in the demonstrations:

Definition 5. We denote by \bar{u}_i the highest pseudo-market price \bar{u}^l (computed in Step 2a of Alg. 1) among the links used by player i :

$$\bar{u}_i = \max\{\bar{u}^l : r_i^l = 1\}. \quad (15)$$

We also define \bar{l}_i as the highest link l in i 's route such that $\bar{u}_i^l = \bar{u}_i$, where the term "highest" is in the sense "most upstream". Link \bar{l}_i can then be seen as the "most congested link" in i 's route.

On the example of Subsection 4.3, we have $\bar{u}_i = 7$ and $\bar{l}_i = l^2$ for $i \in \{1, 2\}$, whereas $\bar{u}_i = 6$ and $\bar{l}_i = l^1$ for the other players.

We can now establish the following result, that bounds the allocations over player i 's path.

Property 3.

$$\forall i \in \mathcal{I}, \quad \bar{d}_i(\bar{u}_i^+) \leq a_i \leq \bar{d}_i(\bar{u}_i) \quad (16)$$

Property 3 is proved in Appendix 2.

The following property states that the allocation of a player i equals the allocation that was computed by Alg. 1 during Step 2a when processing the most congested link for player i , i.e. the link \bar{l}_i .

Property 4. $\forall i \in \mathcal{I}$,

$$a_i = a_i^{\bar{l}_i}, \quad (17)$$

where \bar{l}_i is given in Definition 5.

Proof. The result is trivial if \bar{l}_i is the link directly connected to the root of the tree. We now suppose that it is not the case, and establish by induction that $a_i^l = a_i^{\bar{l}_i}$ for all link l upstream from \bar{l}_i , and that after Step 2b of the algorithm at link \bar{l}_i we have

$$p \leq \bar{u}_i \Rightarrow \underline{d}_i(p) = a_i^{\bar{l}_i}, \quad (18)$$

where \underline{d}_i is the revised pseudo-demand function of player i , i.e. the pseudo-demand function derived from the revised multi-bid \underline{g}_i .

- Initialization: consider link \bar{l}_i . Step 2b of Alg. 1, and more clearly Eq. (12), implies (18).
- If we assume that (18) stands before processing a link $l \in r_i$ upstream from \bar{l}_i , then:
 - By definition of \bar{u}_i (Eq. (15)) we have $\bar{u}^l < \bar{u}_i$. Therefore the induction hypothesis implies $\underline{d}_i(\bar{u}^l) \geq \underline{d}_i(\bar{u}^{l+}) = a_i^{\bar{l}_i}$. Since (11) ensures that $\underline{d}_i(\bar{u}^l) \geq a_i^l \geq \underline{d}_i(\bar{u}^{l+})$, then $a_i^l = a_i^{\bar{l}_i}$.
 - Therefore Step 2b will not change the revised pseudo-demand function, and (18) still holds.

Property 4 is then established by applying (18) to link l^{root} . \square

We now give a lemma that will be used for establishing the main properties of our pricing scheme:

Lemma 5. *For every multi-bid profile s , if $\bar{u}_i, i \in \mathcal{I}$ are the maximum pseudo-market clearing prices defined in (15), then the multi-bid allocation $a(s)$ that Alg. 1 returns maximizes*

$$\sum_{i \in \mathcal{I}} \bar{u}_i \tilde{a}_i$$

over the set \mathcal{A} of allocations $\tilde{a} \in \mathbb{R}_+^{|\mathcal{I}|}$ satisfying the feasibility constraints

$$\begin{cases} \forall l \in \mathcal{L}, & \sum_{i \in \mathcal{I}} r_i^l \tilde{a}_i \leq Q^l \\ \forall i \in \mathcal{I}, & \tilde{a}_i \geq 0. \end{cases}$$

A proof of Lemma 5 is given in Appendix 3.

We then have the following property stating that a player will not pay more than her declared willingness-to-pay, and that submitting a truthful multi-bid always yields a non-negative utility. This is an important result, since it implies that selfish users will always participate in the auction.

Proposition 1. (individual rationality)

$$\forall i \in \mathcal{I}, \forall s \in S^{|\mathcal{I}|}, \quad c_i(s) \leq \int_0^{a_i(s)} \bar{\theta}'_i(q) dq. \quad (19)$$

Moreover, if player i submits a truthful multi-bid ($s_i \in S_i^T$), then

$$c_i(s) \leq \int_0^{a_i(s)} \theta'_i(q) dq = \theta_i(a_i(s)), \quad (20)$$

which implies that $U_i(s) \geq 0$.

A proof of Proposition 1 is given in Appendix 4.

6 Incentive compatibility

We now consider the reaction of selfish players that are faced with this pricing scheme: how can a player maximize her utility? The following proposition establishes that a player cannot do much better than revealing her true valuation, i.e. than submitting a truthful multi-bid.

Proposition 2. *If a player $i \in \mathcal{I}$ submits a truthful multi-bid $s_i \neq \emptyset$, then every other multi-bid \tilde{s}_i (truthful or not) necessarily yields an increase of utility (if any) that is less than $\int_{\bar{d}_i(\bar{u}_i^+)}^{d_i(\bar{u}_i)} (\theta'_i(q) - \bar{u}_i) dq$.*

Formally, $\forall s_i \in S_i^T, \forall \tilde{s}_i \in S, \forall s_{-i} \in S^{|\mathcal{I}|-1}$,

$$U_i(s_i, s_{-i}) \geq U_i(\tilde{s}_i, s_{-i}) - \int_{\bar{d}_i(\bar{u}_i^+)}^{d_i(\bar{u}_i)} (\theta'_i(q) - \bar{u}_i) dq. \quad (21)$$

A proof of Proposition 2 is provided in Appendix 5.

This result is illustrated in Fig. 6 where the shaded area corresponds to the maximum utility gain player i could expect by submitting a different multi-bid.

Since the pseudo-market clearing price is necessarily higher than 0, we can therefore also note in Fig. 6 that the quantity

$$\int_{\bar{d}_i(\bar{u}_i^+)}^{d_i(\bar{u}_i)} (\theta'_i(q) - \bar{u}_i) dq$$

is always less than

$$\max_{0 \leq m \leq M_i} \left\{ \int_{d_i(p_i^{m+1})}^{d_i(p_i^m)} (\theta'_i(q) - p_i^m) dq \right\},$$

where $p_i^{M_i+1} = \theta'_i(0)$ and $p_i^0 = 0$.

This last quantity is the largest shaded area in Fig. 7. The following proposition is then straightforward:

Proposition 3. Under Assumption 1 we have $\forall i \in \mathcal{I}, \forall s_i \in S_i^T \setminus \emptyset, \forall \tilde{s}_i \in S, \forall s_{-i} \in S^{|\mathcal{I}|-1}$,

$$U_i(s_i, s_{-i}) \geq U_i(\tilde{s}_i, s_{-i}) - C_i, \quad (22)$$

where

$$C_i = \max_{0 \leq m \leq M_i} \left\{ \int_{d_i(p_i^{m+1})}^{d_i(p_i^m)} (\theta'_i(q) - p_i^m) dq \right\} \quad (23)$$

with $p_i^{M_i+1} = \theta'_i(0)$ and $p_i^0 = 0$.

Note that C_i can also be written

$$C_i = \max_{0 \leq m \leq M_i} \left\{ \int_{q_i^{m+1}}^{q_i^m} (\theta'_i(q) - \theta'_i(q_i^m)) dq \right\}$$

with $q_i^{M_i+1} = 0$ and $q_i^0 = d_i(0)$.

Proposition 3 implies that a player can give a truthful multi-bid that brings him the best utility possible, up to a value C_i that can be controlled through the choice of the bids s_i^m on the demand curve. One important point is that this value does not depend on the number of other players, nor on the multi-bid they submit.




Figure 6: The multi-bid $s_i = (s_i^1, s_i^2, s_i^3, s_i^4)$ is optimal for player i up to the value $\int_{\bar{d}_i(\bar{u}_i^+)}^{d_i(\bar{u}_i)} (\theta'_i(q) - \bar{u}_i) dq$ of shaded surface




Figure 7: The multi-bid $s_i = (s_i^1, s_i^2, s_i^3)$ is optimal for player i up to a constant C_i , whatever the multi-bids submitted by others s_{-i} be. C_i is the surface of the darkest shaded area.

7 Efficiency

We prove now that our mechanism provides efficient allocations. The criterion we consider is *social welfare*, i.e. the total utility of all participants in the game, including the seller (player 0). Since the seller's utility is the sum of all prices $U_0(s) = \sum_{i \in \mathcal{I}} c_i(s)$, this is equivalent to maximizing the total valuation of the users $\sum_{i \in \mathcal{I}} \theta_i(a_i)$:

$$\sum_{i \in \mathcal{I} \cup \{0\}} U_i(s) = \sum_{i \in \mathcal{I}} c_i(s) + \sum_{i \in \mathcal{I}} \theta_i(a_i(s)) - c_i(s) = \sum_{i \in \mathcal{I}} \theta_i(a_i(s)).$$

The following proposition states that when players bid truthfully, this quantity is maximized, which means that the resource effectively goes to players who value it most.

Proposition 4. *If Assumptions 1 and 2 hold, then for every truthful multi-bid profile s ,*

$$\max_{\mathcal{A}} \sum_i \theta_i(\tilde{a}_i) - \sum_i \theta_i(a_i(s)) \leq Q^{l^{\text{root}}} \sqrt{2\kappa \max_{i \in \mathcal{I}} C_i},$$

where

$$\mathcal{A} = \left\{ \tilde{a} \in \mathbb{R}_+^{|\mathcal{I}|} : \forall l \in \mathcal{L}, \sum_i r_i^l \tilde{a}_i \leq Q^l \right\}$$

and C_i is defined in Eq. (23).

A proof of Proposition 4 is provided in Appendix 6.

8 Which bids should be submitted?

This section follows the same principles as for the case of a single link. We assume that a user i intends to ensure a utility as close as possible to the maximum. According to the user's knowledge of the bandwidth demand, there are two different possibilities:

- players may have beliefs (meaning *a priori* probability distributions, as in [4]) on the number of users in the game and on their preferences. From this knowledge, a probability distribution of the pseudo-market price \bar{u} can be computed. Note that this way, the auction is a game with population uncertainty [14]. The best strategy for user i is then to use Proposition 2 to choose her bids so as to minimize $\mathbb{E} \left[\int_{\bar{d}_i(\bar{u}_i^+)}^{d_i(\bar{u}_i)} (\theta'_i(q) - \bar{u}_i) dq \right]$. However, estimating such a distribution on the pseudo-market price may imply high cost for information-gathering and market appraisal [19].
- The alternative possibility, that we will adopt in the rest of the paper, assumes that players have no prediction of the pseudo-market price. Then, in order to be as close as possible to the optimum, independently of the multi-bid profile, a natural goal is to minimize the quantity C_i of Proposition 3. We further assume here that the number of bids M_i is fixed (as will be discussed in Section 9): indeed, if player i is allowed to submit as many bids as she wants in her multi-bid, then submitting a number M_i of bids as large as possible will make C_i close to zero, but this will increase the required message process, which the auctioneer will prevent. Assuming M_i fixed, the multi-bid $(s_i^1, \dots, s_i^{M_i})$ that minimizes C_i is such that $\forall m, n, 0 \leq m, n \leq M_i$,

$$\int_{d_i(p_i^{m+1})}^{d_i(p_i^m)} (\theta'_i(q) - p_i^m) dq = \int_{d_i(p_i^{n+1})}^{d_i(p_i^n)} (\theta'_i(q) - p_i^n) dq$$

with $p_i^{M_i+1} = \theta'_i(0)$ and $p_i^0 = 0$,

i.e., considering equal values of $\int_{d_i(p_i^{m+1})}^{d_i(p_i^m)} (\theta'_i(q) - p_i^m) dq \forall m$ will minimize the corresponding value of C_i . In the following, we will call *quantile uniform* this bid repartition. An example of quantile uniform repartition of bids is presented in Fig. 8.

Example: For parabolic valuation functions, i.e. of the form

$$\theta_i(q) = \alpha_i \left[-(q \wedge \bar{q}_i)^2 / 2 + \bar{q}_i (q \wedge \bar{q}_i) \right]$$

with parameters α_i and \bar{q}_i , the marginal valuation function is linear:

$$\theta'_i(q) = \alpha_i [\bar{q}_i - q]^+.$$

The quantile uniform repartition of bids is then easy to compute: prices $p_i^m, 1 \leq m \leq M_i$ are such that

$$p_i^m = m \frac{\theta'_i(0)}{M_i + 1} = m \frac{\alpha_i \bar{q}_i}{M_i + 1}.$$




Figure 8: *Quantile uniform* repartition of bids for $M_i = 4$: the five shaded zones have the same surface.

9 Determination of the number of bids admitted by the auctioneer

According to what we said in previous sections, we assume that the auctioneer imposes the number of bids for all players $i \in \mathcal{I}$ to be fixed. We further assume that it is the same value M for all players and that the players choose their multi-bid according to the quantile uniform distribution previously described.

In order to help the auctioneer in fixing the value of M , we introduce a *cost function* $C(M, \mathcal{I})$ that models the negative effects that are the signaling overhead, the memory storage and the complexity of all underlying allocation and price computations for a corresponding value of M . The auctioneer benefit can then be computed as $B(M, \mathcal{I})$:

$$B(M, \mathcal{I}) = \sum_{i \in \mathcal{I}} c_i - C(M, \mathcal{I}),$$

where allocations and prices correspond to the situation when each user submits exactly M bids.

Since the auctioneer has no *a priori* knowledge of the exact set of users competing for bandwidth, we assume that those users come from a set \mathcal{T} of possible player types, characterized by their valuation function (in other words, a type- t player has valuation function $\theta_{(t)}$). The auctioneer is assumed to know the distribution of number of players of each type by $\mathbb{P}_{\mathcal{T}}$ on $\mathbb{N}^{\mathcal{T}}$. Therefore, the expected revenue $\mathbb{E}_{\mathcal{T}}[R_M]$ can be computed when

players submit M bids by

$$\begin{aligned}\mathbb{E}_{\mathcal{I}}[R_M] &= \mathbb{E}_{\mathcal{I}} \left[\sum_{i \in \mathcal{I}} c_i \middle| M \right] \\ &= \int_{\mathcal{I} \in \mathbb{N}^{\mathcal{T}}} \left(\sum_{i \in \mathcal{I}} c_i \middle| M \right) d\mathbb{P}_{\mathcal{I}}(\mathcal{I}).\end{aligned}$$

The following proposition shows that there exists an optimal value of M :

Proposition 5. *If the marginal valuation functions $(\theta'_{(t)})_{t \in \mathcal{T}}$ are uniformly bounded by a value p_{\max} (meaning that $\forall t \in \mathcal{T}, \theta'_{(t)}(0) \leq p_{\max}$), and if additionally the expected cost $\mathbb{E}_{\mathcal{I}}[C](M) = \int_{\mathcal{I} \in \mathbb{N}^{\mathcal{T}}} C(M, \mathcal{I}) d\mathbb{P}_{\mathcal{I}}(\mathcal{I})$ is nondecreasing, and tends to infinity when M tends to infinity:*

$$\lim_{M \rightarrow +\infty} \mathbb{E}_{\mathcal{I}} C(M) = +\infty,$$

then there exists a finite M that maximizes the expected net benefit of the seller, i.e. that maximizes

$$\mathbb{E}_{\mathcal{I}} \left[\sum_{i \in \mathcal{I}} c_i \middle| M \right] - \mathbb{E}_{\mathcal{I}} [C(M, \mathcal{I})].$$

Proof. Applying Proposition 1, we have $\forall \mathcal{I}, \forall M$,

$$\begin{aligned}\sum_{i \in \mathcal{I}} c_i &\leq \sum_{i \in \mathcal{I} \cup \{0\}} \theta_i(a_i) \leq \sum_{i \in \mathcal{I} \cup \{0\}} a_i \theta'_i(0) \\ &\leq p_{\max} \sum_{i \in \mathcal{I} \cup \{0\}} a_i = p_{\max} Q.\end{aligned}$$

Consequently $\mathbb{E}_{\mathcal{I}} [\sum_{i \in \mathcal{I}} c_i \middle| M] \leq p_{\max} Q$ for all $M \in \mathbb{N}$. Therefore

$$\lim_{M \rightarrow +\infty} \mathbb{E}_{\mathcal{I}} \left[\sum_{i \in \mathcal{I}} c_i - C(M, \mathcal{I}) \middle| M \right] = -\infty,$$

which ensures us that there exists a finite M that maximizes the expected net benefit $\mathbb{E}_{\mathcal{I}} [\sum_{i \in \mathcal{I}} c_i \middle| M] - \mathbb{E}_{\mathcal{I}} [C(M, \mathcal{I})]$. \square

The assumption that $\lim_{M \rightarrow +\infty} \mathbb{E}_{\mathcal{I}} C(M) = +\infty$ seems intuitive: if we account for memory costs, $C(M, \mathcal{I}) = M|\mathcal{I}|$, so the assumption is verified as soon as $\mathbb{E}_{\mathcal{I}}[|\mathcal{I}|] > 0$; if we account for computation costs as in Subsection 4.5 or signaling costs, it is verified as well.

Finally, note that if the expected net benefit is nonpositive for all $M \geq 1$, organizing the auction is too expensive for the seller of the resource, so that $M = 0$, i.e. she will prefer not to sell the resource.

10 Conclusions and perspectives

This paper describes how multi-bid auctions can help in pricing and allocating bandwidth among Internet users. It assumes that the core network is over-provisioned, and then uncongested, so that congestion occurs only in access networks, which are assumed to have a tree structure. The resulting pricing scheme is proved to be individually rational, incentive compatible and efficient (by optimizing the overall utility). The results extend those in [13] for the case of a single link.

As direction for future research, we would like to apply multi-bid auctions to a general network topology, but this requires further attention if we wish to apply it in a distributed manner.

Appendix

1 Proof of Lemma 2

Proof. We first notice that $\bar{d}_i(\cdot)^+$ and $\bar{\theta}'_i(\cdot)^+$ can respectively be written the following way: $\forall i \in \mathcal{I}, \forall s_i \in S, \forall x, y \in \mathbb{R}^+$,

$$\bar{d}_i(x^+) = \begin{cases} 0 & \text{if } s_i = \emptyset \text{ or } p_i^{M_i} \leq x, \\ \max_{1 \leq m \leq M_i} \{q_i^m : p_i^m > x\} & \text{otherwise.} \end{cases}$$

$$\bar{\theta}'_i(y^+) = \begin{cases} 0 & \text{if } s_i = \emptyset \text{ or } q_i^1 \leq y, \\ \max_{1 \leq m \leq M_i} \{p_i^m : q_i^m > y\} & \text{otherwise.} \end{cases}$$

We focus here on (7):

- if $s_i = \emptyset$ then $\bar{\theta}'_i(\cdot) = 0 \leq x$, and (7) is verified.
- If $s_i \neq \emptyset$ and $p_i^{M_i} \leq x$, (7) comes from $\bar{\theta}'_i(\cdot) \leq p_i^{M_i}$.
- If $s_i \neq \emptyset$ and $p_i^{M_i} > x$ then

$$\bar{d}_i(x^+) = \max_{1 \leq m \leq M_i} \{q_i^m : p_i^m > x\}. \quad (24)$$

Now let us assume that $\bar{\theta}'_i(\bar{d}_i(x^+)^+) > x$ and see that we arrive at a contradiction. The fact that $\bar{\theta}'_i(\bar{d}_i(x^+)^+) > 0$ implies that we are in the case when

$$\bar{\theta}'_i(\bar{d}_i(x^+)^+) = \max_{1 \leq m \leq M_i} \{p_i^m : q_i^m > \bar{d}_i(x^+)\}.$$

Since $\bar{\theta}'_i(\bar{d}_i(x^+)^+) > x$,

$$\exists m_1, 1 \leq m_1 \leq M_i : \begin{cases} s_i^{m_1} = (q_i^{m_1}, p_i^{m_1}) \in s_i \\ p_i^{m_1} = \bar{\theta}'_i(\bar{d}_i(x^+)^+) > x \\ q_i^{m_1} > \bar{d}_i(x^+). \end{cases}$$

We now remark that this contradicts the definition of $\bar{d}_i(x^+)$ (Eq. (24)). Therefore (7) is verified.

Now we establish (8). By definition, $\bar{d}_i(x) = \max_{1 \leq m \leq M_i} \{q_i^m : p_i^m \geq x\}$. This means that there exists $m_0 \leq M_i$ such that $\bar{d}_i(x) = q_i^{m_0}$ and $p_i^{m_0} \geq x$.

Consequently we have

$$\bar{\theta}'_i(\bar{d}_i(x)) = \max_{1 \leq m \leq M_i} \{p_i^m : q_i^m \geq q_i^{m_0}\} \geq p_i^{m_0} \geq x,$$

which gives (8). \square

2 Proof of Property 3

Proof. Equations (10) and (11) imply that for all link l such that $r_i^l = 1$ we have

$$\underline{d}_i^l(\bar{u}^{l+}) \leq a_i^l \leq \underline{d}_i^l(\bar{u}^l), \quad (25)$$

where \underline{d}_i^l is the revised pseudo-demand function of player i for link l just before Step 2a of the algorithm, i.e. the pseudo-demand function associated to the multi-bid \underline{s}_i .

Applying (13) and the fact that revised pseudo-demand functions are lower than the original pseudo-demand function (from (12)), we have for each link l such that $r_i^l = 1$,

$$a_i \leq \underline{d}_i^l(\bar{u}^l) \leq \bar{d}_i(\bar{u}^l).$$

This inequality holds for all $l \in r_i$, therefore $a_i \leq \bar{d}_i(\bar{u}_i)$.

Now let us establish the right-hand side of (16). Equations (13) and (25) yield

$$\begin{aligned} a_i = \min_{l \in r_i} a_i^l &\geq \min_{l \in r_i} \underline{d}_i^l(u_i^{l+}) \\ &\geq \min_{l \in r_i} \underline{d}_i^l(u_i^+), \end{aligned} \quad (26)$$

where we have used the nonincreasingness of revised pseudo-demand functions.

We now prove by induction that

$$p > \bar{u}_i \Rightarrow \forall l \in r_i, \underline{d}_i^l(p) = \bar{d}_i(p) : \quad (27)$$

- (27) holds at the beginning of the algorithm since $\underline{s}_i = s_i$ (the revised pseudo-demand equals the pseudo-demand).
- Assume (27) holds just before Alg. 1 processes the allocation on link l . The definition of \bar{u}_i implies that Step 2a will determine a local pseudo-clearing price $\bar{u}^l \leq \bar{u}_i$. Then after Step 2b we have for $p > \bar{u}^l$:

$$\underline{d}_i(p) = \min(a_i^l, \bar{d}_i(p)). \quad (28)$$

The induction hypothesis and (25) imply that $a_i^l \geq \underline{d}_i^l(\bar{u}^{l+}) \geq \underline{d}_i^l(\bar{u}_i^+) \geq \underline{d}_i^l(p) = \bar{d}_i(p)$, since the revised pseudo-demand functions are nonincreasing.

Relation (27) is then established, and gives in particular $\underline{d}_i^l(u_i^+) = \bar{d}_i(u_i^+)$. This last result together with (26) concludes the proof of the property. \square

3 Proof of Lemma 5

Proof. We first define for $i \in \mathcal{I}$ and $x \leq \bar{u}_i$,

$$\bar{l}_i(x) \triangleq \text{the highest link } l \text{ in } i\text{'s route such that } \bar{u}^l \geq x \quad (29)$$

For every user $j \in \mathcal{I}$ whose route includes link $\bar{l}_i(x)$, we know from the definition of \bar{u}_j that $\bar{u}_j \geq \bar{u}^{\bar{l}_i(x)} \geq x$, meaning that $\bar{l}_i(x)$ is above $\bar{l}_j(\bar{u}_j)$ in the tree. Note that ‘‘above’’ means that either $\bar{l}_i(x) = \bar{l}_j(\bar{u}_j)$ or $\bar{l}_i(x)$ is a link upstream from $\bar{l}_j(\bar{u}_j)$.

>From Property 4, and the fact that the allocation cannot increase along the path, $a_j^{\bar{l}_i(x)} = a_j^{\bar{l}_j} = a_j$. Thus if $x > 0$ we have

$$\sum_{j \in \mathcal{I}} r_j^{\bar{l}_i(x)} a_j = \sum_{j \in \mathcal{I}} r_j^{\bar{l}_i(x)} a_j^{\bar{l}_i(x)} = Q^{\bar{l}_i(x)}, \quad (30)$$

as a direct consequence of (11) when the pseudo-market clearing price is strictly positive. This means that the capacity of the links in $\{\bar{l}_i(x) : i \in \mathcal{I}, x \in (0, \bar{u}_i]\}$ is completely allocated.

Now consider two players $i, j \in \mathcal{I}$ and $x \in (0, \bar{u}_i]$:

- if $r_j^{\bar{l}_i(x)} = 1$, then we know that every link l above $\bar{l}_i(x)$ verifies $\bar{u}^l < x$, thanks to the definition of $\bar{l}_i(x)$, leading to $\bar{l}_j(x) = \bar{l}_i(x)$.
- if $\bar{l}_j(x) = \bar{l}_i(x)$ then $r_j^{\bar{l}_i(x)} = 1$ since $\bar{l}_j(x)$ is by definition on player j 's route toward the root of the tree.

Consequently we can write

$$\forall i, j \in \mathcal{I}, \forall x \in (0, \bar{u}_i], r_j^{\bar{l}_i(x)} = 1 \Leftrightarrow \bar{l}_j(x) = \bar{l}_i(x). \quad (31)$$

For any allocation $\tilde{a} \in \mathcal{A}$ and $x > 0$, we have

$$\begin{aligned} \sum_{i: \bar{u}_i \geq x} (a_i - \tilde{a}_i) &= \sum_{l: \exists i \in \mathcal{I}, \bar{l}_i(x) = l} \sum_{j: \bar{l}_j(x) = l} (a_j - \tilde{a}_j) \\ &= \sum_{l: \exists i \in \mathcal{I}, \bar{l}_i(x) = l} \underbrace{\sum_{j \in \mathcal{I}} r_j^l (a_j - \tilde{a}_j)}_{\geq 0 \text{ from (30) and } \tilde{a} \in \mathcal{A}} \end{aligned}$$

where the second line comes from (31). Therefore

$$\forall x > 0, \forall \tilde{a} \in \mathcal{A}, \quad \sum_{i: \bar{u}_i \geq x} (a_i - \tilde{a}_i) \geq 0. \quad (32)$$

To complete the proof of the lemma, we sort the maximum pseudo-market clearing prices $\{\bar{u}_i, i \in \mathcal{I}\}$ in an ascending order: $\bar{u}_{(1)} > \bar{u}_{(2)} > \dots > \bar{u}_{(K)}$, and define $D_{(k)} =$

$\sum_{i: \bar{u}_i \geq \bar{u}_{(k)}} (a_i - \tilde{a}_i)$. Equation (32) implies that $D_{(k)} \geq 0$ for all k such that $\bar{u}_{(k)} > 0$, and consequently for all $k \leq K - 1$ since $\bar{u}_{(K)} \geq 0$.

If we introduce $D_{(0)} = 0$, we have

$$\begin{aligned} \sum_{i \in \mathcal{I}} \bar{u}_i (a_i - \tilde{a}_i) &= \sum_{k=1}^K \bar{u}_{(k)} \sum_{i: \bar{u}_i = \bar{u}_{(k)}} (a_i - \tilde{a}_i) \\ &= \sum_{k=1}^K \bar{u}_{(k)} (D_{(k)} - D_{(k-1)}) \\ &= \bar{u}_{(K)} D_{(K)} + \sum_{k=1}^{K-1} (\bar{u}_{(k)} - \bar{u}_{(k+1)}) D_{(k)} \\ &\geq 0 \end{aligned}$$

since all the terms in the sum are nonnegative, and

- if $\bar{u}_{(K)} = 0$ then $\bar{u}_{(K)} D_{(K)} = 0$
- if $\bar{u}_{(K)} > 0$ then we have $D_{(K)} \geq 0$ and consequently $\bar{u}_{(K)} D_{(K)} \geq 0$.

Lemma 5 is then proved. \square

4 Proof of Proposition 1

Proof. We first establish that $\forall \mathcal{I}, \forall s \in S^{|\mathcal{I}|}, \forall j \in \mathcal{I}, \forall y \in \mathbb{R}^+$,

$$\int_y^{a_j(s)} \bar{\theta}'_j \geq \bar{u}_j (a_j(s) - y), \quad (33)$$

where \bar{u}_j is the highest pseudo-market clearing price on user j 's route, as defined in (15).

This comes from the fact that $\bar{\theta}'_j$ is left-continuous and nonincreasing, and from (16) and Lemma 2:

- If $y \leq a_j(s)$, then

$$\begin{aligned} \int_y^{a_j(s)} \bar{\theta}'_j &\geq \underbrace{\bar{\theta}'_j(a_j(s))}_{\geq \bar{\theta}'_j(\bar{d}_j(\bar{u}_j))} \underbrace{(a_j(s) - y)}_{\geq 0} \\ &\geq \bar{u}_j (a_j(s) - y). \end{aligned}$$

- If $y > a_j(s)$, then

$$\begin{aligned} \int_y^{a_j(s)} \bar{\theta}'_j &\leq \underbrace{\bar{\theta}'_j(a_j(s)^+)}_{\leq \bar{\theta}'_j(\bar{d}_j(\bar{u}_j^+))} \underbrace{(y - a_j(s))}_{\geq 0} \\ &\leq \bar{u}_j (y - a_j(s)). \end{aligned}$$

To prove Proposition 1, we apply (33) with $y = a_j(s_{-i})$ and get

$$\begin{aligned} c_i(s) &= \sum_{j \in \mathcal{I} \setminus \{i\}} \int_{a_j(s)}^{a_j(\emptyset, s_{-i})} \bar{\theta}'_j \\ &\leq \sum_{j \in \mathcal{I} \setminus \{i\}} \bar{u}_j (a_j(\emptyset, s_{-i}) - a_j(s)). \end{aligned} \quad (34)$$

Furthermore, from Lemma 5 we have

$$\sum_{j \in \mathcal{I}} \bar{u}_j (a_j(\emptyset, s_{-i}) - a_j(s)) \leq 0,$$

which implies that $\sum_{j \in \mathcal{I} \setminus \{i\}} \bar{u}_j (a_j(\emptyset, s_{-i}) - a_j(s)) \leq \bar{u}_i a_i$, and therefore

$$c_i(s) \leq \bar{u}_i a_i.$$

The end of the proof works like the proof of the individual rationality in the single-link case [13]:

- if $a_i(s) = 0$ then $c_i(s) \leq 0$ and (19) is established.
- if $a_i(s) > 0$ then necessarily $\bar{d}_i(\bar{u}_i) > 0$ (Eq. (16)), so we have $s_i \neq \emptyset$ and $p_i^{M_i} \geq \bar{u}_i$. Lemma 2 thus gives $\bar{\theta}'_i(\bar{d}_i(\bar{u}_i)) \geq \bar{u}_i$. The nonincreasingness of $\bar{\theta}'_i$ and the inequality $a_i(s) \leq \bar{d}_i(\bar{u}_i)$ imply that

$$\int_0^{a_i(s)} \bar{\theta}'_i \geq a_i(s) \bar{\theta}'_i(a_i(s)) \geq a_i(s) \bar{\theta}'_i(\bar{d}_i(\bar{u}_i)) \geq a_i(s) \bar{u}_i,$$

giving (19).

Relation (20) is a straightforward consequence of (19) and Lemma 1. \square

5 Proof of Proposition 2

Proof. Following the lines of [13], we consider two multi-bids: a truthful multi-bid s_i and another one \tilde{s}_i , not necessarily truthful, for user i . The difference of charges is

$$\begin{aligned} c_i(\tilde{s}_i, s_{-i}) - c_i(s) &= \sum_{j \in \mathcal{I} \setminus \{i\}} \int_{a_j(\tilde{s}_i, s_{-i})}^{a_j(s)} \bar{\theta}'_j \\ &\geq \sum_{j \in \mathcal{I} \setminus \{i\}} \bar{u}_j (a_j(s) - a_j(\tilde{s}_i, s_{-i})) \end{aligned} \quad (35)$$

where we used (33).

On the other hand, consider the difference of valuations $D_{\theta_i} := \theta_i(a_i(s)) - \theta_i(a_i(\tilde{s}_i, s_{-i}))$. We distinguish several cases:

- if $a_i(s) > a_i(\tilde{s}_i, s_{-i})$, then

$$\begin{aligned} D_{\theta_i} &= \int_{a_i(\tilde{s}_i, s_{-i})}^{a_i(s)} \theta' \geq \int_{a_i(\tilde{s}_i, s_{-i})}^{a_i(s)} \bar{\theta}'_i \\ &\geq \bar{u}_i(a_i(s) - a_i(\tilde{s}_i, s_{-i})) \end{aligned}$$

from inequalities (6) and (33).

- If $a_i(s) \leq a_i(\tilde{s}_i, s_{-i})$ and $\bar{u}_i \geq \theta'_i(0)$, then

$$\begin{aligned} D_{\theta_i} &= \int_{a_i(\tilde{s}_i, s_{-i})}^{a_i(s)} \theta' \geq \theta'_i(0)(a_i(s) - a_i(\tilde{s}_i, s_{-i})) \\ &\geq \bar{u}_i(a_i(s) - a_i(\tilde{s}_i, s_{-i})). \end{aligned}$$

- If $a_i(s) \leq a_i(\tilde{s}_i, s_{-i})$ and $\bar{u}_i < \theta'_i(0)$, then $\theta'_i(d_i(\bar{u}_i)) = \bar{u}_i$ and

$$\begin{aligned} D_{\theta_i} &= \theta_i(a_i(s)) - \theta_i(d_i(\bar{u}_i)) + \theta_i(d_i(\bar{u}_i)) - \theta_i(a_i(\tilde{s}_i, s_{-i})) \\ &\geq \int_{d_i(\bar{u}_i)}^{a_i(s)} (\theta'_i(q) - \bar{u}_i) dq + \bar{u}_i(a_i(s) - d_i(\bar{u}_i)) + \bar{u}_i(d_i(\bar{u}_i) - a_i(\tilde{s}_i, s_{-i})) \\ &\geq \int_{d_i(\bar{u}_i)}^{\bar{d}_i(\bar{u}_i^+)} (\theta'_i(q) - \bar{u}_i) dq + \bar{u}_i(a_i(s) - a_i(\tilde{s}_i, s_{-i})) \end{aligned}$$

where the last line comes from (16) and from the fact that $\theta'_i(q) - \bar{u}_i \geq 0$ for all $q \leq d_i(\bar{u}_i)$.

Finally we always have

$$D_{\theta_i} \geq \bar{u}_i(a_i(s) - a_i(\tilde{s}_i, s_{-i})) - \int_{\bar{d}_i(\bar{u}_i^+)}^{d_i(\bar{u}_i)} (\theta'_i(q) - \bar{u}_i) dq. \quad (36)$$

To conclude the proof, from (35) and (36), we get

$$\begin{aligned} U_i(s) - U_i(\tilde{s}_i, s_{-i}) &= D_{\theta_i} + c_i(\tilde{s}_i, s_{-i}) - c_i(s) \\ &\geq - \int_{\bar{d}_i(\bar{u}_i^+)}^{d_i(\bar{u}_i)} (\theta'_i(q) - \bar{u}_i) dq + \sum_{j \in \mathcal{I}} \bar{u}_j(a_j(s) - a_j(\tilde{s}_i, s_{-i})) \\ &\geq - \int_{\bar{d}_i(\bar{u}_i^+)}^{d_i(\bar{u}_i)} (\theta'_i(q) - \bar{u}_i) dq, \end{aligned}$$

where we applied Lemma 5. □

6 Proof of Proposition 4

Proof. First notice that if Assumptions 1 and 2 hold, then $\forall i \in \mathcal{I}, \forall e, f : 0 < e \leq f \leq \theta'_i(0)$,

$$d_i(e) - d_i(f) \geq \frac{f - e}{\kappa}. \quad (37)$$

Consider a player $i \in \mathcal{I}$ such that $a_i(s) > 0$. Since $a_i(s) \leq \bar{d}_i(\bar{u}_i)$, we have $d_i(\bar{u}_i) \geq \bar{d}_i(\bar{u}_i) > 0$. Thus $\theta'_i(d_i(\bar{u}_i)) = \bar{u}_i$, and

$$\theta'_i(a_i(s)) \geq \theta'_i(\bar{d}_i(\bar{u}_i)) \geq \theta'_i(d_i(\bar{u}_i)) = \bar{u}_i. \quad (38)$$

On the other hand, we have

$$\theta'_i(a_i(s)) \leq \theta'_i(\bar{d}_i(\bar{u}_i^+)). \quad (39)$$

- If $\theta'_i(0) \leq \bar{u}_i$ then $\theta'_i(a_i) \leq \bar{u}_i$.
- If $\theta'_i(0) > \bar{u}_i$ then

$$\begin{aligned} \theta'_i(\bar{d}_i(\bar{u}_i^+)) &= \min_{1 \leq m \leq M_i+1} \{p_i^m : p_i^m > \bar{u}_i\} \\ &\leq \bar{u}_i + \max_{0 \leq m \leq M_i} (p_i^{m+1} - p_i^m). \end{aligned} \quad (40)$$

with $p_i^{M_i+1} = \theta'_i(0)$ and $p_i^0 = 0$.

Now we remark that for all $m, 0 \leq m \leq M_i$

$$\begin{aligned} \int_{d_i(p_i^{m+1})}^{d_i(p_i^m)} (\theta'_i(q) - p_i^m) dq &= \int_{p_i^m}^{p_i^{m+1}} (d_i(p) - d_i(p_i^{m+1})) dp \\ &\geq \int_{p_i^m}^{p_i^{m+1}} \frac{p_i^{m+1} - p}{\kappa} dp \\ &\geq \frac{(p_i^{m+1} - p_i^m)^2}{2\kappa} \end{aligned}$$

where the second line comes from (37).

Finally, (23) implies that $\forall m, 0 \leq m \leq M_i, p_i^{m+1} - p_i^m \leq \sqrt{2\kappa C_i}$. Therefore (39) and (40) give

$$\theta'_i(a_i(s)) \leq \bar{u}_i + \sqrt{2\kappa C_i}.$$

Define $\mathcal{A} = \{\tilde{a} \in \mathbb{R}_+^{|\mathcal{I}|} : \sum_i \tilde{a}_i \leq Q\}$, and take any $\tilde{a} \in \mathcal{A}$. Let $\mathcal{I}_+ = \{k : \tilde{a}_k \geq a_k(s)\}$ and $\mathcal{I}_- = \{k : \tilde{a}_k < a_k(s)\}$. For $i \in \mathcal{I}_-$, we have $a_k(s) > \tilde{a}_k \geq 0$, and therefore (38) implies $\theta'_i(a_i(s)) \geq \bar{u}_i$. Applying (40), we then have

$$\begin{aligned}
\sum_i \theta_i(\tilde{a}_i) - \theta_i(a_i(s)) & \\
& \leq \sum_{\mathcal{I}_+} \theta'_i(a_i(s))(\tilde{a}_i - a_i(s)) - \sum_{\mathcal{I}_-} \theta'_i(a_i(s))(a_i(s) - \tilde{a}_i) \\
& \leq \sum_i \bar{u}_i(\tilde{a}_i - a_i(s)) + \sum_{\mathcal{I}_+} \sqrt{2\kappa C_i}(\tilde{a}_i - a_i(s)) \\
& \leq Q^{root} \sqrt{2\kappa \max_{i \in \mathcal{I}} C_i},
\end{aligned}$$

where we used Lemma 5. The proposition is then established. \square

References

- [1] J. Altmann, B. Rupp, and P. Varaiya. Internet demand under different pricing schemes. In *Proc. of ACM conference on Electronic Commerce, (EC'99)*, Denver, Colorado, USA, Nov 1999.
- [2] L. Bernstein. Managing the last mile. *IEEE Communications Magazine*, 35(10):72–76, Oct 1997.
- [3] C. Bisdikian, K. Maruyama, D. I. Seidman, and D. N. Serpanos. Cable access beyond the hype: On residential broadband data services over HFC networks. *IEEE Communications Magazine*, 34(11):128–135, Nov 1996.
- [4] H. Chen and Y. Li. Intelligent flow control under game theoretic framework. In D. W. Come, G. D. Smith, and M. J. Oats, editors, *Telecommunications Optimization: Heuristic and Adaptive Techniques*. Wiley, John & Sons, Jan 2000.
- [5] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
- [6] C. Courcoubetis and R. Weber. *Pricing Communication Networks: Economics, Technology and Modelling*. Wiley & Sons, Inc., 2003.
- [7] L. A. DaSilva. Pricing for QoS-enabled networks: A survey. *IEEE Communications Surveys*, 3(2):2–8, 2000.
- [8] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, Cambridge, Massachusetts, 1991.
- [9] T. Groves. Incentives in teams. *Econometrica*, 41(3):617–631, Jul 1973.
- [10] A. A. Lazar and N. Semret. Auctions for network resource sharing. Technical Report 468-97-02, Columbia University, Feb 1997.

-
- [11] A. A. Lazar and N. Semret. Design and analysis of the progressive second price auction for network bandwidth sharing. *Telecommunication Systems - Special issue on Network Economics*, 1999.
 - [12] P. Maillé and B. Tuffin. The progressive second price mechanism in a stochastic environment. *Netnomics*, 5(2):119–147, Nov 2003.
 - [13] P. Maillé and B. Tuffin. Multi-bid auctions for bandwidth allocation in communication networks. In *Proc. of IEEE INFOCOM*, Mar 2004.
 - [14] R. B. Myerson. Population uncertainty and Poisson games. *International Journal of Game Theory*, 27(3):375–392, 1998.
 - [15] N. Semret. *Market Mechanisms for Network Resource Sharing*. PhD thesis, Columbia University, 1999.
 - [16] N. Semret, R. R.-F. Liao, A. T. Campbell, and A. A. Lazar. Pricing, provisioning and peering: Dynamic markets for differentiated internet services and implications for network interconnections. *IEEE Journal on Selected Areas in Communications*, 18(12):2499–2513, Dec 2000.
 - [17] J. Shu and P. Varaiya. Pricing network services. In *Proc. of IEEE INFOCOM*, 2003.
 - [18] B. Tuffin. Charging the internet without bandwidth reservation: an overview and bibliography of mathematical approaches. *Journal of Information Science and Engineering*, 19(5):765–786, Sept 2003.
 - [19] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, Mar 1961.



Unité de recherche INRIA Lorraine, Technopôle de Nancy-Brabois, Campus scientifique,
615 rue du Jardin Botanique, BP 101, 54600 VILLERS LÈS NANCY
Unité de recherche INRIA Rennes, Irista, Campus universitaire de Beaulieu, 35042 RENNES Cedex
Unité de recherche INRIA Rhône-Alpes, 655, avenue de l'Europe, 38330 MONTBONNOT ST MARTIN
Unité de recherche INRIA Rocquencourt, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex
Unité de recherche INRIA Sophia-Antipolis, 2004 route des Lucioles, BP 93, 06902 SOPHIA-ANTIPOLIS Cedex

Éditeur
INRIA, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399