

An Auction-Based Pricing Scheme for Bandwidth Sharing with History-Dependent Utility Functions

Patrick Maillé
GET/ENST-Bretagne
2, rue de la Châtaigneraie CS 17607
35576 Cesson-Sévigné Cedex, France
patrick.maillé@enst-bretagne.fr

Bruno Tuffin
IRISA/INRIA
Campus de Beaulieu
35000 Rennes Cedex, France
btuffin@irisa.fr

Abstract

Pricing telecommunication networks has become a topic of high interest in order to deal with the increasing number of subscribers as well as more and more demanding applications. Users' behavior (or preferences) is usually represented by means of the so-called utility function, but in most cases this function expresses the instantaneous level of satisfaction for the quality of service provided. In this paper, we aim at extending a previous work on auctions for bandwidth to the case where users (or applications) are sensitive to the history of their previous allocations. We introduce a mechanism which takes bids in the form of a finite set of three-(or more)-dimensional points, indicating the bidders willingness to pay for a given quantity, contingent on a given history of allocations or prices to the bidder. Bids are partial representations of continuous utility surfaces. The mechanism computes prices and allocations that are approximately (within specified bounds) and myopically incentive compatible and approximately (within specified bounds) and myopically efficient. The results are approximate because we use piecewise linear approximations of the unknown true continuous utility functions. This extension also includes a refinement of the scheme previously published by providing a closer approximation of real user demand functions.

1 Introduction

During the last decade, the Internet has suffered from congestion due to an exponentially increasing number of subscribers, while capacity did not increase in the same way. At the same time, applications have become increasingly bandwidth and quality of service (QoS) demanding. Many solutions have been proposed to cope with congestion, but a natural one is to introduce pricing schemes to

control demand: if prices increase, demand decreases, and conversely. Although many people argue that pricing is not likely to be implemented thanks to the use of optical fiber so that capacity will still be far away from demand, this does not seem to be the case for wireless networks, where available frequencies are limited, as well as for many access networks, where switching from traditional copper lines to optical fiber would be very expensive (this is traditionally called the last mile problem [2]).

Pricing is thus the subject of a large literature; the reader is advised to look at [7, 9, 12, 22] and the references therein for overviews on the range of methods that have been developed so far. We deal here with auction schemes for bandwidth. Auctions have been first used in the *smart market* proposal [14] at the packet level in such a way that packets with the highest bid are served as far as capacity is not exceeded, and the remaining ones are discarded. The per-packet price almost follows the second-price principle: each admitted packets is charged the lowest bid among all admitted packets. In order to alleviate the per-packet management, progressive second price (PSP) auctions have been developed [13]. In that scheme, users submit two-dimensional bids composed of the amount of bandwidth asked, and the unit price for it. Users with the highest unit price are allocated the desired quantity. The scheme is studied using the framework of *game theory* [10]: the game is played until a Nash equilibrium is reached, meaning that no user (or player) has an incentive to deviate from his current allocation. Here again, charges obey the second price principle. Incentive compatibility (users' interest is to declare their real valuation of bandwidth), individual rationality (users will always gain by entering the game) and efficiency (the social welfare of the resulting allocation is maximized) are proved to be verified. Nevertheless, that scheme (still) presents the drawbacks of requiring a convergence phase (meaning a loss of steady state efficiency, especially if players leave the game or new ones enter [15]), and it also

requires that the bid profile (the list of users' bids) be broadcasted to all players at each step, resulting in signalization overhead. Moreover, the properties of the scheme in terms of incentives and efficiency rely on the strong assumption that users are short-sighted, i.e. they do not take into account the convergence phase of the auction game [16]. To cope with those problems, the authors have developed the so-called multi-bid scheme where players submit a set of two-dimensional bids for once when they enter the game [17], meaning that they do not *need* to send new bids afterwards. Allocations and charges are then immediately computed. It is thus a one-shot scheme where players do not need to know the bid profile before submitting. Again, incentive compatibility, individual rationality and efficiency are proved to be verified. Note that other auction schemes exist in the literature, but we do not describe them here for sake of conciseness [1, 3, 5, 6, 19, 20, 21].

In all those pricing schemes, the user behavior is modeled by the so-called utility function, representing users' preferences. This function may depend on various QoS measures, but the obtained throughput or allocated bandwidth is often the measure of interest. In all the previously presented pricing schemes, the utility function depends on the instantaneous allocation, and is therefore independent of the allocations a player obtained previously. However, we can easily imagine applications for which the utility of getting a given amount of bandwidth depends on what was previously obtained: for instance, it is the case of applications targeting a given average bandwidth.

Our goal is thus to extend the work on multi-bid auctions to the case of history-dependent valuation functions. We additionally provide a refinement of the results in [17] by using a better approximation of real valuation functions, resulting in improved error bounds. This work has been inspired in part by ATHENA [8] where a history-dependent utility function has been developed. Nevertheless, ATHENA considers only users with strict requirements, whereas we consider here elastic users, and it does not include the analytical results in terms of incentives or efficiency for instance that are provided here.

This paper is organized as follows. In Section 2 we present the modeling of history-dependent utility functions and discuss its practical validity. Section 3 presents the multi-bid scheme applied to this model. This can be seen as a generalization of the work on multi-bids in [17] where it is applied only to history-independent utility functions, with tighter bounds thanks to a modification of the allocation procedure. Section 4 then describes the properties verified by the scheme, that are individual rationality, (conditional) incentive compatibility and (conditional) efficiency. Note that those inter-temporal results are independent of future allocations. This is typically justified when connection durations are random and unknown, and may end at any

time. The consequences of the three aforementioned properties in terms of user behavior and efficiency/complexity trade-off are investigated in Section 5. Section 7 gives our conclusions and directions for future work.

2 A history-dependent model to represent users' preferences

We describe here the mathematical representation we will consider to take into account the fact that the utility of a user at a given time is a function not only of her current allocation and charge, but also of what happened from her arrival in the network.

Time is divided into slots: allocations and prices cannot change within a time slot, which implies that slots are short enough to prevent a user entering the game from waiting too long before her request be treated, and long enough to allow synchronization among the different parts of the auction game and control computational complexity.

We assume that utility functions are quasi-linear, which means that for a user i whose resource allocation and charge at a given time slot t are respectively $a_{i,t}$ and $c_{i,t}$, the utility $U_{i,t}$ is the difference between what user i thinks the resource $a_{i,t}$ is worth to her (her willingness-to-pay) and the price $c_{i,t}$ she is charged.

In this paper, we assume that a user willingness-to-pay for the resource at a given time slot may depend on her previous allocation and/or charges. Formally, the utility of a user i at time t is

$$U_{i,t} = \theta_i(a_{i,t}, f(X_{i,t-1})) - c_{i,t}, \quad (1)$$

where θ_i is user i 's valuation (or willingness-to-pay) function, $X_{i,t-1}$ is the history of user i 's allocations and costs starting from her arrival until time slot $t - 1$, and f is a function that represents what criteria user i is sensitive to. Following are examples of such criteria.

- $f(X_{i,t-1}) = \emptyset$: this is the case when the user is only sensitive to her current allocation and price. This (simple) model was considered in [13, 17].
- $f(X_{i,t-1}) = a_{i,t-1}$: the bandwidth allocated in the previous slots. This case may correspond to a user that is sensitive to the "continuity" of her allocation. For example, consider a user whose target is to experience a certain average throughput. Then this user will value more the resource at time t if she did not obtain enough at the previous time slot: θ_i in this case is non-increasing in its second argument.
- $f(X_{i,t-1}) = (\sum_{k=t_i^0}^{t-1} a_{i,k}, \sum_{k=t_i^0}^{t-1} c_{i,k})$, where t_i^0 is the time slot when user i entered the game: here the

user is sensitive to the cumulated resource she has obtained and the total charge she has paid until the current slot. This model can apply for example to users downloading a file: $\sum_{k=t_0}^{t-1} a_{i,k}$ is proportional to the amount of data that user i has downloaded from the beginning of her connection, and she may value more or less the resource depending on the amount of data that remains to be transferred, and on the amount of money she has spent.

In this paper, we will treat the case when $f(X_{i,t-1})$ is one-dimensional, because it is more intuitive and enables graphical interpretations. However, all next definitions and properties also hold when $f(X_{i,t-1})$ has several dimensions. Consequently, the utility experienced by a user i during time slot t depends on three parameters, that are her current allocation $a_{i,t}$, the price she is charged $c_{i,t}$, and the value of $f(X_{i,t-1})$ that we call the *history-relevant criterion* for user i . To simplify the notations, we will note ξ_i the value of the history-relevant criterion for a user i at the current time slot, and Ξ_i the set of possible values of ξ_i .

We assume that users have elastic demand, that is to say their valuation function satisfies some regularity assumptions. Moreover, we assume that the dependency of θ_i on the history-relevant criterion is monotone¹. Those properties are summarized in Assumption A.

Assumption A $\forall i \in \mathcal{I}$,

- $\forall \xi_i \in \Xi_i$, $\theta_i(0, \xi_i) = 0$ and $\theta_i(\cdot, \xi_i)$ is non-decreasing,
- $\forall \xi_i \in \Xi_i$, $\theta_i(\cdot, \xi_i)$ is concave.
- $\forall q \in \mathbb{R}^+$, $\theta_i(q, \cdot)$ is a monotone function over Ξ_i .

Such valuation functions are displayed in Figure 1 in the case when $f(X_{i,t-1}) = a_{i,t-1}$. In the following, $\theta'_{i,\xi}$ will denote the partial derivative $\frac{\partial \theta_i(q, \xi)}{\partial q}$ with respect to the first argument q , i.e. the marginal valuation conditionally on the past.

3 Multi-bid auctions to compute allocations at each time slot

In this section, we describe how the definition of the multi-bid auction scheme introduced in [17] can be extended to take into account the time dependency of users' valuation in the history. The mechanism we suggest here implies that at her arrival in the game, each user i submits a certain number M_i of 3-dimensional bids² of the form

$$s_i \triangleq \{(q_i^m, \xi_i^m, \rho_i^m), 1 \leq m \leq M_i\}, \quad (2)$$

¹when $f(X_{i,t-1})$ is a K -dimensional vector, then we need θ_i to be monotone in each of its K components.

²In the general case where $f(X_{i,t-1})$ is a K -dimensional vector, then each s_i^m should be of dimension $K + 2$.

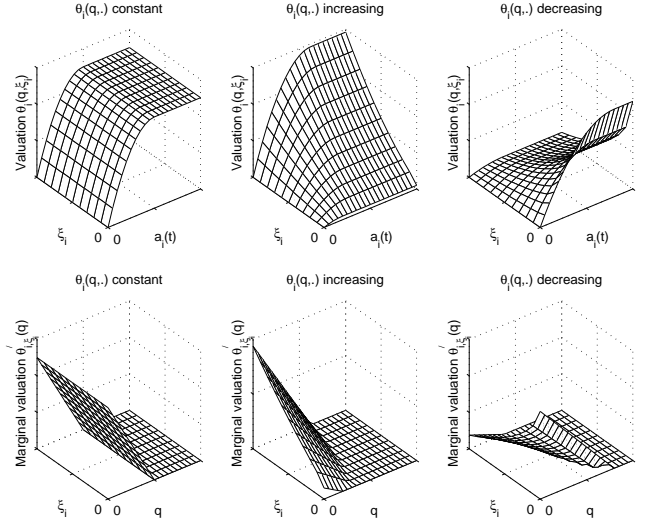


Figure 1. Valuation (top) and marginal valuation (bottom) functions for three types of users: indifferent to the history-relevant criterion (left), increasing (center) and decreasing (right) in that criterion.

that are interpreted as follows: a user submitting such a bid declares she is willing to pay $\rho_i^m \geq 0$ to obtain an allocation $q_i^m > 0$ at the current time slot, if the value of her history-relevant criterion is ξ_i^m . The set $s_i = (s_i^m)_{1 \leq m \leq M_i}$ is called the *multi-bid* submitted by user i . Remark that the scheme we define here is a one-shot scheme, in the sense that a user does not modify her bid once entered the auction game: the multi-bid s_i will be taken into account by the mechanism until the departure of user i whatever the modifications in the network conditions. A user is also asked the sense of variation of her valuation function in the history-relevant criterion; when $f(X_{i,t-1})$ is one-dimensional, one bit added to the multi-bid s_i is sufficient to specify whether $\theta_i(q, \cdot)$ is increasing or decreasing for all q .

Definition 1 We say that user i bids truthfully, or submits a truthful multi-bid, if

- all points $(q_i^m, \xi_i^m, \rho_i^m)$, $1 \leq m \leq M_i$ are on her valuation function curve, i.e. $\rho_i^m = \theta_i(q_i^m, \xi_i^m)$ for all m ,
- user i reveals her true sense of variation in the history-relevant criterion.

At time slot t , we denote \mathcal{I}_t the set of users that are in the game: this set may change over time, as users may enter and leave the network (like connections starting and ending).

At a given time slot t , the auctioneer will then compute an allocation $a_{i,t}$ and a price to pay $c_{i,t}$ for each user i present in the game, based on the submitted multi-bids, and taking into account the history of the game through the computation of the history-relevant criterion $f(X_{i,t-1})$ of each user.

3.1 Reserve price

We assume that the seller sets a reserve unit price p_0 under which she prefers not to sell the resource. For this to be taken into account by the mechanism, the seller (who will be denoted as player 0) may be seen as a player submitting a bid $s_0 = (q_0, q_0 p_0)$ (with $M_0 = 1$), where $q_0 > Q$. Remark that if we want that bid to have the same form as the multi-bids submitted by the users, then every multi-bid of the form $(q_0, \xi_0, q_0 p_0)$, where the seller declares she is indifferent to the history-relevant criterion, is appropriate. In this paper, we will denote $\mathcal{I}_t^0 \triangleq \mathcal{I}_t \cup \{0\}$ the set of all users present at time slot t (including the seller), and $s_t \triangleq (s_i)_{i \in \mathcal{I}_t}$ the set of all competing multi-bids, that we will call the *multi-bid profile* at time t . Throughout the paper, for $i \in \mathcal{I}_t^0$, we denote $s_{t,-i} \triangleq (s_j)_{j \in \mathcal{I}_t^0 \setminus \{i\}}$ the multi-bid profile without user i .

3.2 Conditional multi-bid depending on the history

When users' valuation functions are independent of the history (which is the case treated in [17]), the multi-bid of each user i is a set of two-dimensional bids from which the functions used to compute allocations and prices are defined. We describe here how the mechanism extracts a set of 2-dimensional points from the multi-bid s_i and a given value ξ_i of the history-relevant criterion $f(X_{i,t-1})$. This set of 2-dimensional points, that we call the *conditional multi-bid* and denote s_{i,ξ_i} , is computed as follows from the 3-dimensional bids:

- if user i declared that $\theta_i(q, \cdot)$ is an **increasing** function for all given q , the conditional multi-bid corresponding to the history-relevant criterion ξ_i is

$$s_{i,\xi_i} \triangleq \{(q_i^m, \rho_i^m) : \xi_i^m \leq \xi_i\}; \quad (3)$$

- if user i declared that $\theta_i(q, \cdot)$ is **decreasing** for all given q , the conditional multi-bid corresponding to the history-relevant criterion ξ_i is

$$s_{i,\xi_i} \triangleq \{(q_i^m, \rho_i^m) : \xi_i^m \geq \xi_i\}. \quad (4)$$

Therefore only bids with history criterion under/above ξ_i , depending on the form of monotonicity, are considered by

just skipping the ξ_i^m values. The conditional multi-bid s_{i,ξ_i} is designed in a way that each pair $(q, \rho) \in s_{i,\xi_i}$ is interpreted as meaning that at the current time slot, user i is willing to pay less than ρ to obtain q units of resource because her history-relevant criterion has value ξ_i . Figure 2 illustrates the conditional multi-bid s_{i,ξ_i} for a truthful bidder whose valuation function $\theta_i(q, \cdot)$ is decreasing for all fixed q .

If player i declared she is indifferent to the history-relevant criterion, then the mechanism simply ignores the ξ_i^m in the bids, and takes $s_{i,\xi_i} \triangleq \{(q_i^m, \rho_i^m) : 1 \leq m \leq M_i\}$ as in [17].

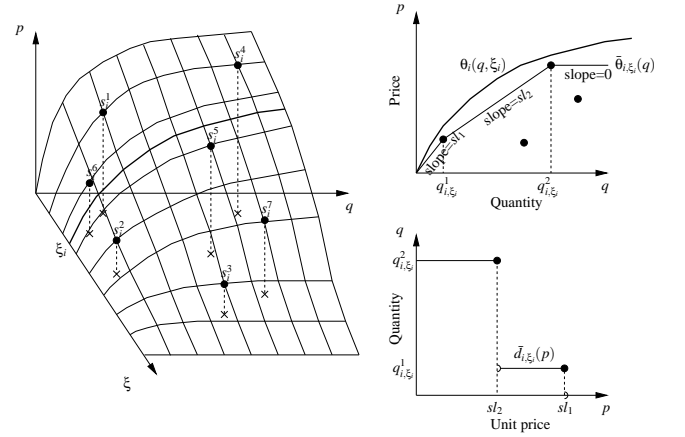


Figure 2. Valuation function and a truthful multi-bid submitted by a user (left), conditional multi-bid s_{i,ξ_i} for a given value ξ_i of the history-relevant criterion. The bids from s_i that appear in s_{i,ξ_i} are s_i^2, s_i^3, s_i^5 and s_i^7 . On the right are the conditional pseudo-valuation function $\bar{\theta}_{i,\xi_i}$ (top-right), and associated conditional pseudo-demand function $\bar{d}_{i,\xi_i}(p)$ (bottom-right) as described in Section 3.3.

3.3 Allocation rule

We now describe how allocations are computed at each time slot t . The mechanism works as follows:

- the value ξ_i of the history-relevant criterion is calculated for each player i , and the conditional multi-bid s_{i,ξ_i} is derived from s_i and ξ_i , as described in subsection 3.2.
- The auctioneer then computes a *conditional pseudo-valuation function* for each player i :

Definition 2 *At a given time slot, the conditional pseudo-valuation function for a user i , who submitted*

the multi-bid s_i and whose history-relevant criterion equals ξ_i is the function $\bar{\theta}_{i,\xi_i} : \mathbb{R}^+ \rightarrow \mathbb{R}^+$, defined as the **lowest positive and concave function** such that $\bar{\theta}_{i,\xi_i}(q) \geq \rho$ for all pair $(q, \rho) \in s_{i,\xi_i}$.

The conditional pseudo-valuation function is displayed in Figure 2 for a given value of $\xi_i(t)$. The goal is to obtain an approximation of the true valuation function from the multi-bid points. The valuation function θ_i being concave in its first argument, the conditional pseudo-valuation function $\bar{\theta}_{i,\xi_i}$ associated with a truthful multi-bid is such that

$$\forall \xi_i, \forall q \geq 0 \quad \bar{\theta}_{i,\xi_i}(q) \leq \theta_i(q, \xi_i). \quad (5)$$

- c) From the conditional pseudo-valuation function, which is left-differentiable, we compute the *conditional pseudo-marginal valuation function*, denoted by $\bar{\theta}'_{i,\xi_i}$, as the left derivative function of $\bar{\theta}_{i,\xi_i}$. Since $\bar{\theta}'_{i,\xi_i}(q)$ is not defined at $q = 0$, we define $\bar{\theta}'_{i,\xi_i}(0)$ as the right limit of $\bar{\theta}'_{i,\xi_i}$ at 0 (which exists since $\bar{\theta}_{i,\xi_i}$ is piecewise linear), so that $\bar{\theta}'_{i,\xi_i}$ is continuous at 0.
- d) Then a *conditional pseudo-demand function* \bar{d}_{i,ξ_i} is computed for each user:

$$\forall p \in \mathbb{R}^+ \quad \bar{d}_{i,\xi_i}(p) \triangleq \sup\{q : \bar{\theta}'_{i,\xi_i}(q) \geq p\}, \quad (6)$$

with the convention $\sup \emptyset \triangleq 0$. \bar{d}_{i,ξ_i} may also be defined as the largest quantity q of resource that maximizes $\bar{\theta}_{i,\xi_i}(q) - pq$, i.e. the quantity that a user with valuation function $\bar{\theta}_{i,\xi_i}$ would buy to optimize her utility if the resource were sold at a fixed unit price p . Like the conditional pseudo-marginal valuation, the conditional pseudo-demand function \bar{d}_{i,ξ_i} is stair-step, left-continuous, positive and non-increasing.

Those functions are different from the ones initially published in [17] where users were asked to declare their marginal valuation in their multi-bid. This new choice allows to get a closer approximation of the valuation and demand functions, while keeping the properties proved in [17], yielding a reduced gap with respect to the optimal values. Indeed, the new choice computes a concave function for the pseudo-valuation function that better approaches the actual one than the stair-step function in [17].

- e) Based on all conditional pseudo-demand functions of the users present in the game $\bar{\theta}_{i,\xi_i}$, $i \in \mathcal{I}_t$, allocations are determined (using the same rule as in [17]): the aggregated conditional pseudo-demand function \bar{d}_ξ is computed:

$$\bar{d}_\xi \triangleq \sum_{i \in \mathcal{I}_t^0} \bar{d}_{i,\xi_i}. \quad (7)$$

Then the pseudo-market clearing price corresponding to \bar{d}_ξ is defined as

$$\bar{u}_t \triangleq \sup\{p : \bar{d}_\xi(p) > Q\}, \quad (8)$$

and the total capacity Q is shared among flows according to their conditional pseudo-demand functions: the allocation $a_{i,t}(s_t)$ for a user $i \in \mathcal{I}_t^0$ can be written

$$a_{i,t}(s_t) \triangleq \bar{d}_{i,\xi_i}(\bar{u}_t^+) + \frac{\bar{d}_{i,\xi_i}(\bar{u}_t) - \bar{d}_{i,\xi_i}(\bar{u}_t^+)}{\bar{d}_\xi(\bar{u}_t) - \bar{d}_\xi(\bar{u}_t^+)}(Q - \bar{d}_\xi(\bar{u}_t^+)), \quad (9)$$

where, for every function f , $f(x^+)$ denotes the right limit at x (which exists here since conditional pseudo-demand functions are stair-step). The first term of the allocation (9) corresponds to the quantity player i asks at the lowest price \bar{u}_t^+ for which supply exceeds pseudo-demand. The second term is strictly positive if all the resource is not allocated at \bar{u}_t^+ , the surplus $Q - \bar{d}_\xi(\bar{u}_t^+)$ being shared among players who submitted a bid at price \bar{u} , with weights proportional to the ‘‘hops’’ of the pseudo-demand functions $\bar{d}_{i,\xi_i}(\bar{u}_t) - \bar{d}_{i,\xi_i}(\bar{u}_t^+)$.

Let us now introduce some remarks that will be helpful in the proofs of properties. Since $\bar{\theta}'_{i,\xi_i}$ is left-continuous, the sup in (6) is a max when $\{q : \bar{\theta}'_{i,\xi_i}(q) \geq p\} \neq \emptyset$, i.e. when $p \leq \bar{\theta}'_{i,\xi_i}(0)$. It implies that

$$p \leq \bar{\theta}'_{i,\xi_i}(0) \Rightarrow \bar{\theta}'_{i,\xi_i}(\bar{d}_{i,\xi_i}(p)) \geq p. \quad (10)$$

Moreover,

$$\forall p \in \mathbb{R}^+ \quad \bar{d}_{i,\xi_i}(p^+) = \sup\{q : \bar{\theta}'_{i,\xi_i}(q) > p\}, \quad (11)$$

(still using $\sup \emptyset = 0$). Therefore

$$\forall p \in \mathbb{R}^+, \quad \bar{\theta}'_{i,\xi_i}(\bar{d}_{i,\xi_i}(p^+)^+) \leq p, \quad (12)$$

since if it were not the case then there would exist a $q > \bar{d}_{i,\xi_i}(p^+)$ such that $\bar{\theta}'_{i,\xi_i}(q) > p$, which would contradict (11). Equation (12) also holds for $p \geq \bar{\theta}'_{i,\xi_i}(0)$.

3.4 Pricing rule

The price $c_{i,t}$ each user $i \in \mathcal{I}_t$ is charged at time slot t is computed according to the multi-bid pricing rule defined in [17], based on the conditional pseudo-valuation functions:

$$c_{i,t}(s_t) \triangleq \sum_{j \neq i} \bar{\theta}_{j,\xi_j}(a_{j,t}(s_{t,-i})) - \bar{\theta}_{j,\xi_j}(a_{j,t}(s_t)), \quad (13)$$

where $a_{i,t}(s_{t,-j})$ is the allocation that the mechanism would have given to user i at time slot t if player j had just leaved the game, i.e. if the bid profile at time t had been

$s_{t,-i}$. This pricing rule, inspired by the Vickrey-Clarke-Groves pricing rule [4, 11, 23], can be interpreted as an “exclusion-compensation” principle: at time slot t , each user pays for the loss of valuation she imposed on the others by staying in the game.

4 Properties of the scheme

We prove in this section that the auction scheme described above verifies three important properties (at least up to pre-determined constants): (conditional) incentive compatibility which states that users’ best interest is to truthfully declare their valuations, individual rationality which states that entering the game will never yields a negative utility, and (conditional) efficiency, stating that the resulting allocation is the one providing the highest social welfare (that is the sum of valuations of users, auctioneer included).

Before proving those three properties, let us prove a lemma stating that the auction mechanism allocation maximizes the declared “pseudo-social welfare”, i.e. the sum of pseudo-valuations of all users, at each time slot t , given the history of allocations. Define $a_t(s_t) \triangleq (a_{i,t}(s_t))_{i \in \mathcal{I}_t}$ as the vector of allocations at time slot t .

Lemma 1 *The multi-bid allocation $a_t(s_t)$ maximizes the sum of the pseudo-valuations of all users (including the seller) under the capacity constraint:*

$$\forall s_t, \quad a_t(s_t) \in \arg \max_{\tilde{a} \in \mathcal{A}_t} \sum_{i \in \mathcal{I}_t} \bar{\theta}_{i,\xi_i}(\tilde{a}_i) \quad (14)$$

where $\mathcal{A}_t \triangleq \{\tilde{a} = (\tilde{a}_i)_{i \in \mathcal{I}_t^0} \in [0, Q]^{|\mathcal{I}_t^0|} : \sum_{i \in \mathcal{I}_t^0} \tilde{a}_i \leq Q\}$.

Proof: We start by showing that for every multi-bid profile s_t , $i \in \mathcal{I}_t$ and $y \in \mathbb{R}^+$, we have

$$\bar{\theta}_{i,\xi_i}(a_{i,t}(s_t)) - \bar{\theta}_{i,\xi_i}(y) \geq \bar{u}_t(a_{i,t}(s_t) - y), \quad (15)$$

where \bar{u}_t is the pseudo-market clearing price computed in (8).

- if $y < a_{i,t}(s_t)$, then

$$\begin{aligned} & \bar{\theta}_{i,\xi_i}(a_{i,t}(s_t)) - \bar{\theta}_{i,\xi_i}(y) \\ & \geq \underbrace{\bar{\theta}'_{i,\xi_i}(a_{i,t}(s_t))}_{\geq \bar{\theta}'_{i,\xi_i}(\bar{d}_{i,\xi_i}(\bar{u}_t))} \underbrace{(a_{i,t}(s_t) - y)}_{\geq 0} \\ & \geq \bar{u}_t(a_{i,t}(s_t) - y), \end{aligned}$$

where (10) is applied, holding because $0 \leq y < a_{i,t}(s_t)$ by hypothesis (thus $\bar{d}_{i,\xi_i}(\bar{u}_t) \geq a_{i,t}(s_t) > 0$ and consequently $\bar{u}_t \leq \bar{\theta}'_{i,\xi_i}(0)$).

- If $y \geq a_{i,t}(s_t)$, then

$$\begin{aligned} & \bar{\theta}_{i,\xi_i}(y) - \bar{\theta}_{i,\xi_i}(a_{i,t}(s_t)) \\ & = \int_{a_{i,t}(s_t)}^y \bar{\theta}'_{i,\xi_i} \\ & \leq \underbrace{\bar{\theta}'_{i,\xi_i}(a_{i,t}(s_t)^+)}_{\leq \bar{\theta}'_{i,\xi_i}(\bar{d}_{i,\xi_i}(\bar{u}_t)^+)} \underbrace{(y - a_{i,t}(s_t))}_{\geq 0} \\ & \leq \bar{u}_t(y - a_{i,t}(s_t)), \end{aligned}$$

by using (12).

Therefore (15) holds, and Lemma 1 immediately follows: if we denote by \tilde{a} another allocation vector satisfying the capacity constraint, we have

$$\begin{aligned} & \sum_{i \in \mathcal{I}_t^0} \bar{\theta}_{i,\xi_i}(a_{i,t}(s_t)) - \sum_{i \in \mathcal{I}_t^0} \bar{\theta}_{i,\xi_i}(\tilde{a}_{i,t}) \\ & \geq \bar{u}_t \left(\underbrace{\sum_{i \in \mathcal{I}_t^0} a_{i,t}(s_t)}_{=Q} - \underbrace{\sum_{i \in \mathcal{I}_t^0} \tilde{a}_{i,t}}_{\leq Q} \right) \\ & \geq 0. \end{aligned}$$

4.1 Incentive compatibility

This subsection aims at studying an important notion called incentive compatibility. It states that a selfish user’s best interest reacting to the auction scheme we have defined in the previous section, in order to optimize her utility, is to play truthfully by declaring her valuation of bandwidth, whatever the bids submitted by the other players be. Indeed, we are going to prove that at each time slot t , in the worst case, a truthful user is ensured that the gap between the utility brought by the multi-bid she submitted and the maximum utility that she could have obtained by bidding differently is less than $\max_{q \in [0, Q]} \theta_i(q, \xi_i) - \bar{\theta}_{i,\xi_i}(q)$, where $\xi_i = f(X_{i,t-1})$.

Proposition 2 (incentive compatibility) $\forall t, \forall i \in \mathcal{I}_t, \forall s_{t,-i}, \forall s_i, \tilde{s}_i, \forall \xi_i = f(X_{i,t-1})$,

s_i truthful \Rightarrow

$$U_{i,t}((s_i, s_{t,-i}), \xi_i) \geq U_{i,t}((\tilde{s}_i, s_{t,-i}), \xi_i) - C_{i,\xi_i}, \quad (16)$$

with

$$C_{i,\xi_i} \triangleq \max_{q \in [0, Q]} \theta_{i,t}(q, \xi_i) - \bar{\theta}_{i,\xi_i}(q). \quad (17)$$

Proof: Let $s_t = (s_i, s_{t,-i})$. Note that the conditional pseudo-valuation functions $\theta_{j,\xi_j}, j \neq i$, are the same if the multi-bid profile is s_t or $(\tilde{s}_i, s_{t,-i})$ and that, from the definition of C_{i,ξ_i} , $\bar{\theta}_{i,\xi_i}(q) \leq \theta_{i,t}(q, \xi_i) \leq \bar{\theta}_{i,\xi_i}(q) + C_{i,\xi_i}, \forall q \in [0, Q]$. We therefore have, using (13):

$$\begin{aligned}
& U_{i,t}(s_t, \xi_i) - U_{i,t}(\tilde{s}_i, s_{t,-i}, \xi_i) \\
& \geq \bar{\theta}_{i,\xi_i}(a_{i,t}(s_t)) + \sum_{j \neq i} \bar{\theta}_{j,\xi_j}(a_{j,t}(s_t)) - C_{i,\xi_i} \\
& \quad - \bar{\theta}_{i,\xi_i}(a_{i,t}(\tilde{s}_i, s_{t,-i})) - \sum_{j \neq i} \bar{\theta}_{j,\xi_j}(a_{j,t}(\tilde{s}_i, s_{t,-i})) \\
& \geq -C_{i,\xi_i} + \underbrace{\sum_{j \in \mathcal{I}_t^0} (\bar{\theta}_{j,\xi_j}(a_{j,t}(s_t)) - \bar{\theta}_{j,\xi_j}(a_{j,t}(\tilde{s}_i, s_{t,-i})))}_{\geq 0 \text{ (Lemma 1)}},
\end{aligned}$$

which establishes Proposition 2.

Bidding truthfully therefore ensures that, at each time slot t , player i maximizes her utility up to a constant that is the largest difference between her conditional valuation and pseudo-valuation functions. Notice that incentive compatibility is proved only when reasoning per time slot, conditionally on the past; there may exist untruthful strategies that lead in some situations to a larger overall utility during the presence period of a player, given the fact that your current allocation affects your future utility. However, playing truthfully may anyway be the best strategy for users who do not have an idea of their (random) connection duration, so that it could stop at any time, or of the evolution of network conditions in the future (with other users randomly entering/leaving the game).

4.2 Individual rationality

A second interesting property is individual rationality, stating that a user who submitted a truthful multi-bid is ensured to obtain a non-negative utility, i.e. to be charged a price equal to or below her valuation for her allocation. This property implies that a rational user will always enter the game and continue to play at each time slot.

Proposition 3 $\forall t, \forall i \in \mathcal{I}_t, \forall s_{t,-i}, \forall \xi_i = f(X_{i,t-1})$, if user i submitted a truthful multi-bid s_i then

$$U_{i,t}(s_i, s_{t,-i}) = \theta_i(a_{i,t}(s_t), \xi_i) - c_{i,t}(s_t) \geq 0. \quad (18)$$

Proof: Equation (5) implies

$$\begin{aligned}
U_{i,t}(s_t, \xi_i) &= \theta_{i,t}(a_{i,t}(s_t)) - c_{i,t}(s_t) \\
&= \theta_{i,t}(a_{i,t}(s_t)) + \sum_{j \neq i} [\bar{\theta}_{j,\xi_j}(a_{j,t}(s_t)) - \bar{\theta}_{j,\xi_j}(a_{j,t}(s_{t,-i}))] \\
&\geq \sum_{j \in \mathcal{I}_t} \bar{\theta}_{j,\xi_j}(a_{j,t}(s_t)) - \sum_{j \neq i} \bar{\theta}_{j,\xi_j}(a_{j,t}(s_{t,-i})) \\
&\geq \sum_{j \in \mathcal{I}_t} \bar{\theta}_{j,\xi_j}(a_{j,t}(s_t)) - \bar{\theta}_{j,\xi_j}(a_{j,t}(s_{t,-i})) \geq 0,
\end{aligned}$$

where the last line comes from Lemma 1.

4.3 Efficiency

The last important property we wish to show is (conditional) efficiency. We prove that, at each time slot t , conditionally to the past and independently of the future, the auction scheme allocates efficiently the available resource among users. The efficiency measure that we consider here is *social welfare*, that is the total valuation of users (including the seller) for the allocation: $\sum_{i \in \mathcal{I}_t^0} \theta_i(a_{i,t}, \xi_i)$. This quantity is also the sum of utilities of all users, if we consider that the utility of the seller is her valuation for her allocation plus her total revenue, $U_{0,t} = \theta_0(a_{0,t}) + \sum_{i \in \mathcal{I}_t} c_{i,t}$ for each t , where $\theta_0(q) = p_0 q$.

Like for the incentive compatibility property, efficiency is instantaneous (i.e. at each time slot t), which means that allocations are not necessarily efficient if we consider periods of several time slots (then in general a discount factor has to be introduced to compare valuations at different times). However, once again, such a property indicates that the mechanism behaves in a good fashion, and may be chosen if the network conditions of the future (number of users and submitted multi-bids) cannot be predicted.

To establish this property, we add some regularity assumptions on valuation functions:

Assumption B $\exists \kappa > 0 : \forall i \in \mathcal{I}, \forall \xi_i \in \Xi_i$,

- $\theta_i(\cdot, \xi_i)$ is differentiable in its first argument and θ'_{i,ξ_i} (the derivative in the first argument) is continuous,
- $\forall z, z', z > z' \geq 0, \theta'_{i,\xi_i}(z) - \theta'_{i,\xi_i}(z') > -\kappa(z - z')$,
- $\theta'_{i,\xi_i}(Q) = 0$.

The first and second points of Assumption B were introduced by Lazar and Semret in [13] for valuation functions independent of the history to prove the efficiency of Progressive Second Price auctions. The last point states that the available resource of the link is sufficient to fully satisfy a user if she is the only one using that link.

For such valuation functions, we have the following efficiency result:

Proposition 4 Under Assumptions A and B, the auction allocation at each time slot is close to the social welfare optimum, conditionally on the past, when players bid truthfully: $\forall \mathcal{I}_t, \forall (\xi_i)_{i \in \mathcal{I}_t}$ and truthful multi-bid profile s_t ,

$$\sum_{i \in \mathcal{I}_t^0} \theta_i(a_{i,t}(s_t), \xi_i) \geq \sup_{\tilde{a} \in \mathcal{A}_t} \left(\sum_{i \in \mathcal{I}_t^0} \theta_i(\tilde{a}_i, \xi_i) \right) - Q \sqrt{8\kappa \max_{i \in \mathcal{I}_t} C_{i,\xi_i}}$$

with $\mathcal{A} \triangleq \{\tilde{a} = (\tilde{a}_i)_{i \in \mathcal{I}_t^0} \in [0, Q]^{|\mathcal{I}_t^0|} : \sum_{i \in \mathcal{I}_t^0} \tilde{a}_i \leq Q\}$.

Proof: Relations (5) and (17) lead to

$$\forall q, \tilde{q} \in [0, Q], \forall \xi_i, \\ \left| \theta_i(q, \xi_i) - \bar{\theta}_{i, \xi_i}(q) - (\theta_i(\tilde{q}, \xi_i) - \bar{\theta}_{i, \xi_i}(\tilde{q})) \right| \leq C_{i, \xi_i},$$

or equivalently

$$\left| \int_{\tilde{q}}^q [\theta'_{i, \xi_i}(x) - \bar{\theta}'_{i, \xi_i}(x)] dx \right| \leq C_{i, \xi_i}. \quad (19)$$

Assume now that there exist $K > 0$, ξ_i and $q \geq 0$ such that

$$\theta'_{i, \xi_i}(q) > \bar{\theta}'_{i, \xi_i}(q) + K. \quad (20)$$

Then from Assumption B, $\theta'_{i, \xi_i}(q + \frac{K}{\kappa}) \geq \theta'_{i, \xi_i}(q) - \kappa \frac{K}{\kappa} > 0$, so that $q + \frac{K}{\kappa} \in [0, Q]$. Therefore (19) gives

$$C_{i, \xi_i} \geq \int_q^{q + \frac{K}{\kappa}} (\theta'_{i, \xi_i}(x) - \bar{\theta}'_{i, \xi_i}(x)) dx \\ \geq \int_q^{q + \frac{K}{\kappa}} (\theta'_{i, \xi_i}(q) + \kappa(q - x) - \theta'_{i, \xi_i}(q) + K) dx = \frac{K^2}{2\kappa},$$

thus from (20) we have

$$\forall \xi_i, \forall q \in [0, Q], \quad \theta'_{i, \xi_i}(q) - \bar{\theta}'_{i, \xi_i}(q) \leq \sqrt{2\kappa C_{i, \xi_i}}. \quad (21)$$

On the other hand, consider $K > 0$, ξ_i and $q > 0$ such that $\theta'_{i, \xi_i}(q) < \bar{\theta}'_{i, \xi_i}(q) - K$. Then from Assumption B, $q - \frac{K}{\kappa} \geq 0$, unless we would have $\theta'_{i, \xi_i}(0) < \bar{\theta}'_{i, \xi_i}(q)$, which would contradict (5). Applying again (19), we get

$$-C_{i, \xi_i} \leq \int_{q - \frac{K}{\kappa}}^q (\theta'_{i, \xi_i}(x) - \bar{\theta}'_{i, \xi_i}(x)) dx \\ \leq \int_{q - \frac{K}{\kappa}}^q (\theta'_{i, \xi_i}(q) + \kappa(q - x) - \theta'_{i, \xi_i}(q) - K) dx = \frac{-K^2}{2\kappa},$$

therefore

$$\forall \xi_i, \forall q \in [0, Q], \quad \theta'_{i, \xi_i}(q) - \bar{\theta}'_{i, \xi_i}(q) \geq -\sqrt{2\kappa C_{i, \xi_i}}. \quad (22)$$

We now establish Proposition 4: from (21), (22) and the immediate relation $\bar{d}_{i, \xi_i}(\bar{u}_t^+) \leq a_{i, t}(s_t) \leq \bar{d}_{i, \xi_i}(\bar{u}_t)$ (that is a consequence of (9)), we easily get, applying (10) and (12),

$$\left\{ \begin{array}{l} \theta'_{i, \xi_i}(a_{i, t}(s_t)) \leq \bar{u}_t + \sqrt{2\kappa C_{i, \xi_i}} \\ a_{i, t}(s_t) > 0 \Rightarrow \theta'_{i, \xi_i}(a_{i, t}(s_t)) \geq \bar{u}_t - \sqrt{2\kappa C_{i, \xi_i}}. \end{array} \right.$$

Let $\mathcal{I}_{t,+} = \{i \in \mathcal{I}_t^0 : \bar{a}_i \geq a_{i, t}(s_t)\}$ and $\mathcal{I}_{t,-} = \{i \in \mathcal{I}_t^0 :$

$\bar{a}_i < a_{i, t}(s_t)\}$. We have

$$\sum_{i \in \mathcal{I}_t^0} \theta_i(a_{i, t}, \xi_i) - \theta_i(\bar{a}_i, \xi_i) \\ \geq \sum_{i \in \mathcal{I}_{t,-}} \theta'_{i, \xi_i}(a_{i, t}(s_t))(a_{i, t}(s_t) - \bar{a}_i) - \\ - \sum_{i \in \mathcal{I}_{t,+}} \theta'_{i, \xi_i}(a_{i, t}(s_t))(\bar{a}_i - a_{i, t}(s_t)) \\ \geq \sum_{i \in \mathcal{I}_{t,-}} (\bar{u}_t - \sqrt{2\kappa C_{i, \xi_i}})(a_{i, t}(s_t) - \bar{a}_i) - \\ - \sum_{i \in \mathcal{I}_{t,+}} (\bar{u}_t + \sqrt{2\kappa C_{i, \xi_i}})(\bar{a}_i - a_{i, t}(s_t)) \\ \geq -\sqrt{2\kappa} \sum_{i \in \mathcal{I}_t} \sqrt{C_{i, \xi_i}} |a_{i, t}(s_t) - \bar{a}_i| \\ \geq -2Q \sqrt{2\kappa \max_i C_{i, \xi_i}},$$

which gives the proposition.

5 Player and auctioneer behavior

5.1 Multi-bid choice for a user

As pointed out in the previous section, bidding truthfully at the arrival into the game ensures a user i to obtain at each time slot t a utility that is close to the maximum possible, up to $C_{i, f(X_{i, t-1})}$. If user i has an a priori probability distribution of the allocations she will obtain during her connection period, then she may try to choose her multi-bid so as to minimize the expected value of $C_{i, f(X_{i, t-1})}$. If she does not have such an a priori distribution, then a natural way to choose her bid s_i is to try to be as close as possible to the maximum utility at each time slot *in the worst case*, that is

$$s_i \in \arg \min \left\{ \max_{\xi_i \in \Xi_i} C_{i, \xi_i} : s_i = (s_i^m)_{1 \leq m \leq M} \text{ truthful} \right\}. \quad (23)$$

For general valuation functions θ_i , finding an analytical expression of a multi-bid satisfying (23) is beyond the scope of this paper, but could otherwise be obtained by numerical optimization procedures (meta-heuristics for instance).

5.2 Trade-off between economical efficiency and computational complexity

Proposition 4 states that the allocation at a given time slot t is close to the socially optimal one, up to a value that decreases (if users are rational) when the number of bids M_i allowed in a multi-bid increases. Moreover, we proved in [17] that the complexity of computing allocations and prices given the conditional multi-bids $(s_{i, \xi_i})_{i \in \mathcal{I}_t}$ are of the order $O(|\mathcal{I}_t| \times \sum_{i \in \mathcal{I}_t} M_{i, \xi_i})$, where M_{i, ξ_i} is the number

of two-dimensional bids in s_{i,ξ_i} . Notice that at each time slot, selecting the conditional multi-bids s_{i,ξ_i} for all players in \mathcal{I}_t needs around $\sum_{i \in \mathcal{I}_t} M_i$ operations.

Therefore, increasing M_i for all users in \mathcal{I}_t ensures that the conditional social welfare be close to the optimal one, but also yields a higher complexity. Depending on the computational capacities of the auctioneer and the relative importance of efficiency *versus* complexity, the auctioneer may admit at most M bids of the form (q_i^m, ξ_i^m, p_i^m) in a multi-bid, and choose M to balance the trade-off. A study of this trade-off, as well as a comparison with the PSP scheme in terms of both complexity and efficiency, has been realized in [18] in the case of atemporal utility functions, illustrating the advantages of the multi-bids scheme.

6 Numerical illustration

In this section, we present an example of scenario with different types of users entering and leaving the game over time. The valuation functions that are used are those described in Figure 1. The history-relevant criterion considered here for each user is his allocation at the previous slot. We illustrate here the behavior of our scheme for when the following scenario is applied (the corresponding allocations and pseudo-market clearing price are plotted in Figure 3 for $Q = 10$, $M = 200$ and truthful bids taken with equally spaced values of q_i^m and ξ_i^m):

- the game begins at slot $t = 0$ with only one player, of type 2 (valuation increasing in the history-relevant criterion), who leaves the game at time $t = 28$,
- another type-2 player bids for bandwidth between $t = 7$ and $t = 16$,
- a player of type 3 (valuation decreasing in the history-relevant criterion) enters the game at $t = 23$, and another one arrives at $t = 35$.

Between time slots $t = 7$ and $t = 16$, two identical players of type 2 are competing for the resource. However, player 1 was already present upon arrival of player 2, and had obtained some resource at time $t = 6$. Therefore due to his type, he values the resource more than player 2 at time $t = 7$, which explains that he obtains more resource at that time. However he obtains less than in the previous slot, hence his conditional valuation for slot $t = 8$ will decrease, whereas for the same reason the conditional valuation of player 2 increases. For the following slots, the allocations of both players therefore get closer.

The remaining type-2 player then faces a type-3 player in the interval $[23, 28]$: when that player enters the game, he obtains a large amount of resource (due to an important valuation), which will decrease his conditional valuation for

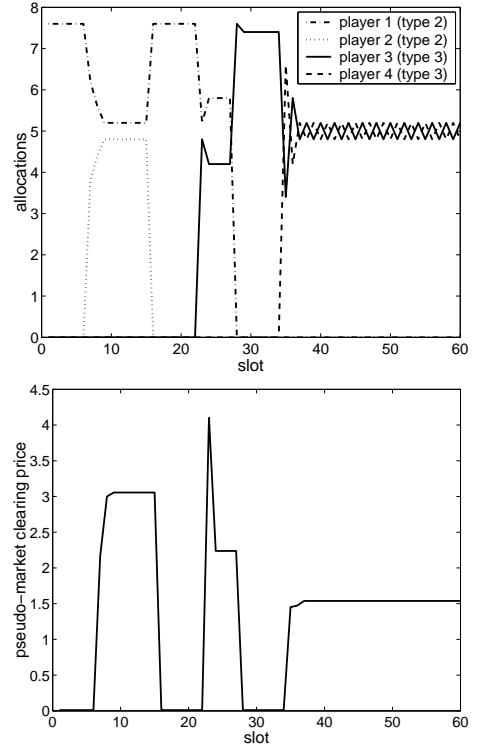


Figure 3. Allocations $a_{i,t}$ (top) and pseudo-market clearing price \bar{u}_t (illustrating the congestion level) (bottom) for the scenario described in Section 6.

the next slot. That is the reason why the corresponding allocation of player 3 is lower at slot 24 than at slot 23. The symmetric reasoning applies to the type-2 player.

After slot 35, we have a situation where two identical players of type 3 compete for bandwidth. Since player 3 had obtained some resource at slot 34, his conditional valuation is lower than that of player 4, consequently at slot 35 player 4 obtains more resource than player 3. At each time slot, the player who obtains more resource is the one who obtained less at the previous slot, which explains the oscillations that occur in Figure 3.

7 Conclusions

In this paper, we have designed an auction scheme to allocate the available bandwidth of a communication link among several users *over time*. Based on the previously defined multi-bid auction mechanism, we adapted the scheme to the case when the valuation of users for the resource depends on their history since their arrival in the game, including also tighter bounds due to a better approximation of the true valuation functions. The signaling overhead is low, since users submit a number M of 3-dimensional bid once only for all their connection duration, and do not need to receive any information, unlike in some other schemes. Moreover, the computational complexity can be controlled, and the scheme has been proved to satisfy some properties in terms of incentives and social welfare maximization.

Some more work can be done to take into account history-dependent valuations: in this paper we assumed the dependency to be monotone in the history-relevant criterion, in order to define properly the conditional multi-bid and the conditional valuation function. Adapting the scheme described here to the case when the dependency is not monotone would deserve some attention.

References

- [1] C. Ambühl, A. E. F. Clementi, P. Penna, G. Rossi, and R. Silvestri. Energy consumption in radio networks: Selfish agents and rewarding mechanisms. In *Proc. of WAOA 2003, LNCS 2909*, pages 248–251, 2003.
- [2] L. Bernstein. Managing the last mile. *IEEE Communications Magazine*, 35(10):72–76, Oct 1997.
- [3] M. Bitsaki, G. Stamoulis, and C. Courcoubetis. An efficient auction mechanism for hierarchically structured bandwidth markets. In *Proc. of ICQT'04, LNCS 3266*, Barcelona, Spain, Oct 2004. Springer.
- [4] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
- [5] C. Courcoubetis, M. P. Dramitinos, and G. D. Stamoulis. An auction mechanism for bandwidth allocation over paths. In *Proc. of the 17th International Teletraffic Congress (ITC)*, Dec 2001.
- [6] J. Crowcroft, R. Gibbens, F. Kelly, and S. Östring. Modelling incentives for collaboration in mobile ad hoc networks. In *Proc. of WiOpt'03*, Mar 2003.
- [7] L. A. DaSilva. Pricing for QoS-enabled networks: A survey. *IEEE Communications Surveys*, 3(2):2–8, 2000.
- [8] M. P. Dramitinos, C. Courcoubetis, and G. D. Stamoulis. Auction-based resource reservation in 2.5/3G networks. In *Proc. of WiOpt'03: Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, INRIA, Sophia-Antipolis, France, Mar 2003.
- [9] M. Falkner, M. Devetsikiotis, and I. Lambadaris. An overview of pricing concepts for broadband IP networks. *IEEE Communications Surveys & Tutorials*, 3(2), 2000.
- [10] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, Cambridge, Massachusetts, 1991.
- [11] T. Groves. Incentives in teams. *Econometrica*, 41(3):617–631, Jul 1973.
- [12] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan. Rate control in communication networks: Shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49:237–252, 1998.
- [13] 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.
- [14] J. K. MacKie-Mason and H. R. Varian. Pricing the internet. In B. Kahin and J. Keller, editors, *Public Access to the Internet*, pages 269–314. MIT Press, 1995.
- [15] P. Maillé and B. Tuffin. The progressive second price mechanism in a stochastic environment. *Netnomics*, 5(2):119–147, Nov 2003.
- [16] P. Maillé and B. Tuffin. A progressive second price mechanism with a sanction bid from excluded players. In *Proc. of ICQT'03, LNCS 2816*, pages 332–341, Munich, Germany, Sept 2003. Springer.
- [17] P. Maillé and B. Tuffin. Multi-bid auctions for bandwidth allocation in communication networks. In *Proc. of IEEE INFOCOM*, Mar 2004.
- [18] P. Maillé and B. Tuffin. Multi-bid versus progressive second price auctions in a stochastic environment. In *Proc. of ICQT'04, LNCS 3266*, pages 318–327, Barcelona, Spain, Oct 2004. Springer.
- [19] P. Marbach. Pricing differentiated services networks: Bursty traffic. In *Proc. of IEEE INFOCOM*, 2001.
- [20] P. Reichl, G. Fankhauser, and B. Stiller. Auction models for multiprovider internet connections. In *Proc. of Messung, Modellierung und Bewertung (MMB'99)*, Trier, Germany, Sept 1999.
- [21] J. Shu and P. Varaiya. Pricing network services. In *Proc. of IEEE INFOCOM*, 2003.
- [22] 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.
- [23] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, Mar 1961.