

An Auction Mechanism for Bandwidth Allocation over Paths*

Costas Courcoubetis^{1,2}, Manos P. Dramitinos¹, and George D. Stamoulis^{1,2}

¹ Department of Informatics, Athens University of Economics and Business
76 Patision Str., Athens, GR 10434, Greece

² Institute of Computer Science, Foundation for Research and Technology Hellas
Tel: +30 1 8203549, Fax: +30 1 8828655

E-mail: {courcou, mdramit, gstamoul}@aueb.gr

Abstract: The demand for bandwidth contracts over networks has recently been growing rapidly. Moreover, due to increased competition among providers, customized short-term bandwidth contracts are now preferred by users to static, long-term contracts of the past. To this end, auctions appear to be a proper trading mechanism. When purchasing bandwidth over a path of a network, it is only meaningful for a user to reserve the same quantity at all the constituent links. We have developed a simple, yet efficient auction for allocating bandwidth on a network basis to users who wish to utilize it for the same time period. This mechanism (referred to as MIDAS) consists of a set of simultaneous multi-unit Dutch (i.e. descending-price) auctions, one per link of the network. In order to win bandwidth over a certain path, it suffices for a user to simultaneously bid for the quantity desired at all relevant auctions. Thus, instant allocation of bandwidth is attained. An important feature of our approach is that prices at the various links are reduced at different rates, so that prices reflect the demand exhibited so far for each link. We have evaluated experimentally two price reduction policies, in terms of the social welfare associated with the resulting allocation, and argue (both theoretically and by means of experimental results) that it is indeed efficient to introduce such rules rather than reduce all prices at the same rate. We have also briefly addressed the issue of incentive compatible pricing.

Keywords: auctions, bandwidth markets, efficiency, resource allocation, electronic commerce.

1. INTRODUCTION

The rapid growth of Internet has resulted in increased need for bandwidth. At the same time, the introduction of innovative technologies, such as optical fibres, wireless and satellite links, has increased the supply of bandwidth, which is now

* This work is supported in part by EU's IST project "Market Managed Multiservice Internet" - M3I (RTD No IST-1999-11429).

offered at lower prices. Moreover, due to increased competition among providers, customized short-term bandwidth contracts are now preferred by users to static, long-term contracts of the past. These facts have created an overwhelming interest in dynamic bandwidth brokering applications.

Auctions offer the advantages of: i) simplicity in determining market-based prices, and ii) efficiency, since, if the auction is properly designed, goods are acquired by those that value them most. Furthermore, auctions may lead to higher revenues for the providers compared to traditional methods of selling goods. There are many possible ways to classify auctions, such as open outcry or sealed bid, ascending or descending, single or multi-unit, etc; see [1] for an overview. Besides the well-known simple auction formats, (i.e. English, Dutch, First-price sealed, and Vickrey), there are numerous variations thereof ([1], [2]). Though such auction mechanisms can be adopted for allocating bandwidth on a single link, they cannot be applied to the case of networks.

Thus, we are lacking of an auction-based self-regulating, efficient and fair mechanism for bandwidth allocation in a network. Such a mechanism could be used by bandwidth markets to enable their customers to dynamically build end-to-end paths, or multicast trees, or complex VPNs of arbitrary topology. It could also be used as an internal bandwidth allocation mechanism for networks of academic and research institutes. We have developed a simple, yet efficient auction mechanism for allocating the bandwidth of all links of a network to users who wish to utilize it for the same time period. This mechanism consists of a set of simultaneous multi-unit Dutch (i.e. descending-price) auctions, one per link of the network. In order for a user to win bandwidth over a certain path, it suffices to simultaneously bid for the amount desired at all relevant auctions. Note that it is only meaningful for a user to reserve the *same* quantity of bandwidth at all links of the path of interest, which complicates our problem. This property can be guaranteed by our mechanism, due to the instant allocation of bandwidth. User strategies can be based on the price per unit of bandwidth and the spare capacity of the various links, which are sent as feedback to users, as well as on the pricing rule. An important feature of our approach is that prices at the various links are reduced at different rates, following rules specified so that prices reflect the demand exhibited so far for each link. Thus, auctions at different links are coupled due to the demand for paths; e.g. more popular links are in general more expensive, which is fair from an economic point of view. We have evaluated experimentally two price reduction policies, in terms of the seller's (provider's) revenue and the social welfare associated with the resulting allocation. We argue (both theoretically and by means of experimental results) that it is efficient to introduce such rules rather than reduce all prices at the same rate. We also address the issue of payment rule and its implication on bidders' behavior.

The remainder of this paper is organized as follows: In Section 2, we define the problem addressed. In Section 3, we present the proposed auction mechanism, while in Section 4 we present experimental results assessing our auction mechanism w.r.t. social welfare. In Section 5, we briefly deal with the payment rule. Finally, in Section 6, we discuss some other properties of our auction mechanism and give some concluding remarks.

2. BUILDING PATHS THROUGH SINGLE-LINK AUCTIONS

We assume that there are N links whose bandwidth is auctioned simultaneously and I users bidding for bandwidth over paths (sets of consecutive links) of the network. Bandwidth will be utilized by all users for the same period, though this assumption can be relaxed (see [6]). When bandwidth is auctioned, it is desirable that users reveal their true needs in terms of bandwidth quantity and willingness-to-pay. In this paper, we develop such an auction mechanism. Our main objective is that each user should be able to buy the *same* quantity of bandwidth at all the links of his path; see Figure 1, for an example. This goal is very important, since buying a pipe with different bandwidth at its constituent links is inefficient. We assume that each user is associated with a utility function $U_i(\cdot)$ that expresses his valuation (and hence his willingness-to-pay) for quantities of bandwidth over a particular path. (Note that the terms utility and valuation will henceforth be used interchangeably.) For simplicity, we take that each user is interested in a single path.




Figure 1: User's objective

Although our mechanism is applicable to a variety of information models, we implicitly assume henceforth that the model of independent private valuations applies. That is, each user judges how much a quantity of bandwidth in a path is worth to him, independently of the bids of other users. Indeed, this is the case when bandwidth is to be allocated to end-users (who will use it themselves), or to large independent entities, such as the divisions of an institution, or ISPs who resell it to markets of limited overlap.

2.1 The Progressive Second Price auction (PSP)

Auctioning bandwidth over a network has also been examined by Semret and Lazar in [4], where they present the PSP auction. This is first presented for the case of a single link, for which an allocation rule is applied repeatedly until the bids converge. At every round users place bids (i.e., pairs of quantity and price) and gradually “raise” their bids until demand matches supply; the auction ends at this point. The amount each user pays covers the “social opportunity cost” due to the bandwidth allocated to him. This mechanism is then extended to the case of a network; it is suggested that a bidder splits equally his bid among the links that constitute his path. However, in our opinion, this is not the best approach (see [5] and [6]).

3. THE MIDAS MECHANISM

The solution we propose to the problem of selling bandwidth over a network is a mechanism that consists of simultaneous multi-unit Dutch auctions, one for each link. This mechanism (henceforth referred to as MIDAS) is presented in detail below. First, we briefly motivate the design of MIDAS. In particular, we opted for employing *simultaneous auctions*, one per link, rather than allow users to place

combinatorial bids for the set of links for which they are interested in, in order to avoid the well-known problems of combinatorial auctions. That is, the *threshold* and *free-rider* problems (see [2]) and, its major disadvantage, the high computational complexity; winner determination is in general an NP-complete problem [1]. Also, we opted for an *open* rather than a *sealed* auction, since, in the former case, bidders have access to information concerning the bids of others, while in the latter case each bidder places his bid on the basis of his own valuation only. Since combinatorial bids are not allowed, in a sealed auction a path bidder would have to place a sealed bid for each link of the path. Thus, he must decide without any information on the market demand on how to split the total value for the path among the various links. Hence, there is a big chance that this user loses due to unsuccessful splitting of the budget among the path's links. Finally, we opted for *descending* auctions rather than *ascending* ones, because we have proved (see [6]) that when employing ascending auctions, it is impossible to synchronize the auctions of the various links so that all of them terminate at the same time. This is due to the difference in the demand per link.

MIDAS mechanism comprises simultaneous multi-unit Dutch auctions. In particular, for each link i , the total capacity auctioned C_i is announced together with the initial unit price $p_i(0)$. This initial unit price should be high, for example three times a reasonable market price. At each link, the unit price is reduced as time elapses. Users place their bids and are *instantly* allocated bandwidth over links. Although not enforced by rule, it is implicitly assumed that users interested in a path reserve bandwidth at *all constituent links simultaneously*. Why this is a meaningful assumption on user behavior is discussed below. Bidders are allowed to bid several times for the same link(s), thus accumulating bandwidth. The amount to be paid by a user depends on the payment rule. One such option is the "pay-your-bid" rule; i.e. a user bidding at time t for x units of bandwidth at links 1 and 2, at prices $p_1(t)$ and $p_2(t)$ respectively, pays $x[p_1(t) + p_2(t)]$. It should be noted however that the payment rule greatly influences the strategy of the customer. In particular, "pay-your-bid" may result in *bid shading*, that is a bidder bids lower than his valuation in order to achieve a discount. Nevertheless, for the time we assume that customers are truthful in order to assess our mechanism w.r.t. the social welfare. The issue of payment rule is revisited in Section 5.

As already explained, the main merit of employing Dutch auctions in our context lies in instant allocation. However, meeting the objective of efficiency is not straightforward at all. For this purpose, prices should be reduced in a smart way, so that more popular links are more expensive, because it is unfair for bandwidth of the most popular links to be offered at the same price with links having less demand. By charging the links with low and high demand at the same prices, network resources are not allocated to those who value them most. Furthermore, the aggregation of bandwidth over paths may become unreasonably costly for users interested in paths. Consider for example the case depicted in Figure 2. Although path formation may be efficient from an economic point of view, it is hard in practice due to symmetric pricing. Indeed, if link 1 users start placing bids, then the path user has to pay in total (per unit of bandwidth) double the price of the popular link 1. (Otherwise, he risks due to the upcoming exhaustion of capacity of link 1.) This may prevent such a user

to bid. Link 2 should have been priced much lower, since it has much lower demand.




Figure 2: Different demand for two links

In the subsections to follow, we present certain price reduction policies. The way prices are computed under these policies enables the provider to take competition into consideration, and thus balance the various links' spare capacities. This facilitates the purchase of bandwidth over the constituent links of a path, and helps users to develop meaningful strategies. Moreover, it is efficient for users to reserve bandwidth at different links at the same time, which would not be the case under symmetric pricing. This way, users do not risk being allocated different bandwidth quantities at the various links of their path. Of course, it is not obvious whether each of the price reduction policies to be presented deals successfully with the objective of optimal (or at least near-optimal) social welfare; this will be made apparent in Section 4.

4. PRICE REDUCTION POLICIES AND EXPERIMENTAL ASSESSMENT

We have specified and assessed two price reduction policies whereby prices are reduced at different rates, reflecting competition. These policies are described below and then compared to symmetric pricing with respect to social welfare. To keep the presentation simple, we describe the price reduction policies in discrete (rather than) continuous time. It is also important to note that our policies can be implemented in a distributed fashion, employing only *local* information at each link, which ensures scalability for large numbers of links and users.

4.1 Variable Reduction Rates (VRR)

This policy involves a decrement rate of the per unit price of bandwidth per link. At every time t , the price at link l is given by the equation:

$$p_l(t) = p_l(t-1) - \max\{[C_{spare}(t;l) / C_{init}(l)] * MaxDrop, 1\}.$$

That is, the decrement rate of link l at time t is proportional to the fraction of the current spare capacity $C_{spare}(t;l)$ divided by its initial value $C_{init}(l)$. The price at each link is reduced at every step at least by 1 and at most by $MaxDrop$. Thus, prices reflect the demand already exhibited.

4.2 Price Freezing policy(PF)

Under this policy, prices are constantly reduced at a fixed rate r , which is expressed in monetary units per time unit. When an allocation of bandwidth takes place in a link, its price freezes for some time that is proportional to the quantity x of the bandwidth just allocated in this link; that is, the freezing period equals $f \cdot x$, where f is a constant expressed in time units per bandwidth unit. If additional allocations occur during the period of freezing then the price is kept frozen for more time accordingly. It is easily seen that, when the price in a link l is not frozen, then:

$$p_l(t) = p_l(0) - r \cdot [t - f \cdot x_l(t)],$$

where: $p_l(t)$ is the price per unit of bandwidth in link l at time t , $p_l(0)$ is the corresponding initial price, and $x_l(t)$ is the total quantity of bandwidth allocated at link l by time t . Clearly, the values of r and f influence the pace of the auction. This explicit relation between the price and the spare capacity of each link is the property motivating this policy. Indeed, by selecting appropriately the initial price, this relation leads to the following: For two or more links whenever their prices are not frozen, the link with the least spare capacity has the highest price.

4.3 Comparing price reduction policies with symmetric pricing

As already mentioned, the above price reduction policies have been designed in a way that prices tend to reflect the difference in the demand for the various links of the network. This does not apply when all prices are symmetric and hence are reduced at the same rate. Below, we show by means of an example that our price reduction policies result in improved efficiency. We consider a network consisting of two links, namely links 1 and 2. Each link's initial capacity is C_0 , which is assumed to be integer. The initial price in both links is p_0 , which is large. There are four groups of users, namely groups A , B , C , and D . Figure 3 depicts for each group the total quantity of bandwidth demanded, the corresponding link(s), and the valuation per unit of bandwidth; note that the parameters ε and ε' are taken to be positive and very small. For simplicity, we assume that users are truth-telling. Notice that users of groups A , B and D demand bandwidth over one link, while group C comprises path users. Since all of them compete for bandwidth within the same network, we can compare their willingness-to-pay only w.r.t. the mean valuation per unit of bandwidth per hop. For single link users this equals their valuation, while for path users this equals their valuation divided by 2; indeed, under symmetric pricing, a path user has to pay double the price at each link in order to acquire a unit of bandwidth. Below, we derive the different outcomes of the auction, for the price reduction policies, and compare the social welfare attained.

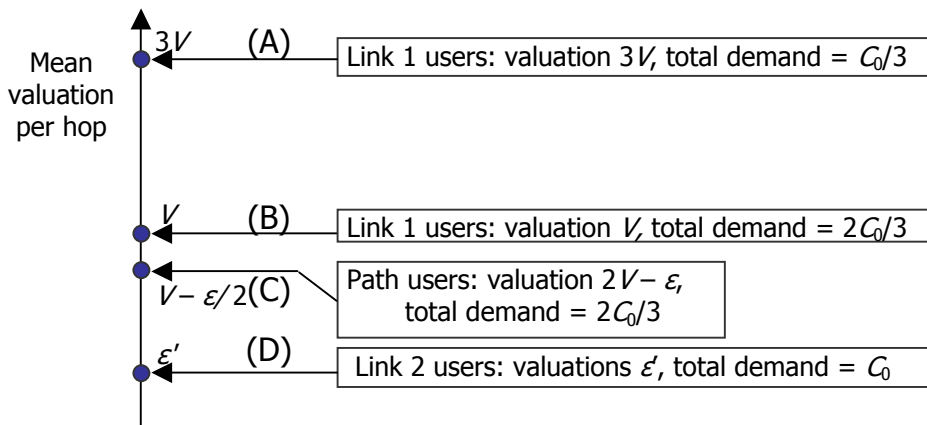


Figure 3: Demand and users' valuations distribution for a simple network

If prices drop symmetrically, then the outcome of this auction is the following: First, group A of users reserves $C_0/3$ units of bandwidth at link 1 when $p_l(t) = 3V$. At some

point in time the price at each link equals V . Then, group B of users reserves $2C_0/3$ units of bandwidth at link 1, exhausting its capacity and preventing path users (users of group C) to reserve any bandwidth. Thus, the entire capacity of link 2 is allocated to group D of users at a very low price, namely ε' . Hence, the social welfare attained under the symmetric price policy is $(C_0/3) \cdot (3V) + (2C_0/3) \cdot V + C_0 \cdot \varepsilon'$.

If either the PF or the VRR policy is applied, then again group A of users reserves $C_0/3$ units of bandwidth at link 1 first. This results in a greater price at link 1 compared to that at link 2. Therefore, group C of users reserve $2C_0/3$ units of bandwidth at each of links 1 and 2, while group B of users still have to wait. However, the spare capacity of link 1 is exhausted, which implies that group B of users will not reserve any bandwidth at all. On the other hand, link 2 has still a spare capacity of $C_0/3$ units of bandwidth, that is reserved by group D users when $p_2(t) = \varepsilon'$. Thus, the social welfare attained under the price freezing policy equals $(C_0/3) \cdot (3V) + (2C_0/3) \cdot (2V - \varepsilon) + (C_0/3) \cdot \varepsilon'$, which exceeds the social welfare attained under symmetric pricing by $(2C_0/3) \cdot (V - \varepsilon - \varepsilon')$. Therefore, it is beneficial to apply our price reduction policies instead of symmetric pricing.

4.4 Experimental Results

Besides the theoretical, an experimental assessment of our auction mechanism has been carried out in a Java testbed for comparing auction types and estimating social welfare.

The experiments carried out regard two network topologies, linear and hierarchical (see Figure 4 and 5 respectively). A number of emulated users participate at each experiment. Each of them is fully specified by means of a utility function and a strategy. The user utility functions implemented are:

- guaranteed, pertaining to users demanding a specific amount of bandwidth;
- linear, pertaining to users with constant marginal utility; such users can purchase any amount of bandwidth but can only afford prices below this threshold;
- elastic, pertaining to users with diminishing marginal utility. (see Figure 6)




Figure 4: A linear network




Figure 5: A hierarchical network




Figure 6: User utility as a function of bandwidth quantity

A variety of user strategies (truth-telling, bid-shading by a predefined factor, and bid-shading adaptively to competition) have been implemented to evaluate the mechanism’s performance in terms of social welfare under several payment rules. Presenting these strategies is beyond the scope of this paper (see [6] for details). Below, we restrict attention to the evaluation of social welfare under the assumption that users are truth-telling. Whether this is a reasonable assumption under the various payment rules is briefly discussed in Section 6.

Below, we present some experimental results for truth-telling users with linear or elastic utility functions. In order to facilitate comparison, the same users participate in the same network auction under both price reduction policies. No immediate comparison however can be made among different experiments. Table 1 contains experimental data for two networks. The first experiment concerns a network of 3 links with capacity 100 Mbps. There are 50 truth-telling bidders whose path size is selected to be 1, 2, or 3 links with probability 1/3; each bidder demands at most 10 Mbps. The second experiment concerns a hierarchical network of 15 links. The central link’s capacity is 500 Mbps while all the others’ is 100 Mbps. There are 350 linear bidders whose path length is 1 or 2, or 3, or 4 with probability 1/4; each bidder demands at most 10 Mbps. For each case, we provide the social welfare attained under the Price Freezing (PF) and the Variable Reduction Rates (VRR) policies.

	Linear network of 3 links		Hier/cal network of 15 links	
	PF	VRR	PF	VRR
Social Welfare	1827	1827	43557	40082

Table 1: Experimental results for two networks

Table 2 contains experimental data regarding two linear networks. The first experiment concerns a network of 3 links of 10 Mbps. There are 5 truth-telling bidders; two of them try to reserve 5Mbps at links 1,2 and 2,3 respectively, are of the guaranteed type and have the highest valuations among all bidders. The other three truth-telling bidders have linear utility functions, and try to reserve any amount from 1 to 10 Mbps at one link each. Thus, at every link a single-link bidder competes against at least one path bidder (at Link 2 there are two path bidders). The second experiment concerns a linear network of 2 links of 10 Mbps. The demand for the second link is the highest. There are 30 bidders of the linear type and 20 bidders of the guaranteed type. The bidders’ path size is either 1 or 2; all demand 1 Mbps.

	Linear network of 3 links		Linear network of 2 links	
	PF	VRR	PF	VRR
Optimal Social welfare	1572 (computed exhaustively)		1556 (computed by a special algorithm)	
Social Welfare	1572 (100% of optimal)	1506 (96% of optimal)	1529 (98% of optimal)	1330 (85.48% of optimal)

Table 2: Experimental results for two linear networks

The experiments above, together with several others carried out, lead to the following conclusion: No matter if the demand for bandwidth at different links is the same or varies considerably, our auction mechanism performs very well in terms of efficiency. Losses are up to 2% of the optimal social welfare for PF and typically around 10% for VRR. Also, the PF policy is typically the most efficient. Finally, it is worth noting that the optimal social welfare is computed exhaustively except for the case of a linear network of two links for which we have developed a special algorithm (see [6] for more).

We have also compared our price reduction policies with symmetric pricing w.r.t. social welfare. It appears that symmetric pricing is significantly less efficient (typically by 15%) than our price reduction policies. Table 3 contains typical experimental results.

Hierarchical network of 15 links			
Social Welfare	Symmetric	PF	VRR
	40052	44209	43943

Table 3: Price reduction policies versus symmetric pricing

Next, we focus on the time for an auction to terminate, which is a very important parameter. Both policies seem to be rather fast. The VRR policy is the fastest one, and thus the “market price” is reached quickly. What is important however is that for both price reduction policies there exists an upper bound on the duration (number of steps) of the auction (see [6] for more).

We finally provide two figures that display the auction’s evolution for a network of two links. As shown in Figures 7a and 7b both policies ultimately balance the links’ spare capacities.

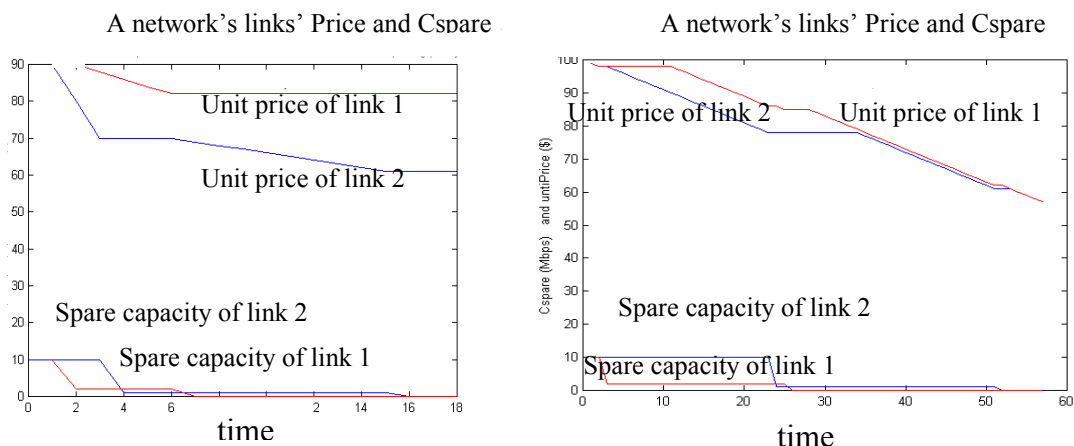


Figure 7a: A network auction under VRR. Figure 7b: A network auction under PF.

5. PAYMENT RULE

So far, we have studied the mechanism of the auction and the allocation rule. The payment rule is also very important, because it affects users' behavior and strategies, thus determining the auction outcome. Our proposal has been described so far as a set of simultaneous multi-unit Dutch (i.e. descending-price) auctions that aim to achieve efficiency. The payment rule of the simple Dutch auction specifies that the winner pays his bid. In practice, such a rule would result in bid shading both for the simple and the multi-unit Dutch auctions. For the former case, it is proved in [1] that bid shading may result in inefficiency (except for the case of symmetric users). In general, efficiency cannot be achieved in the simultaneous multi-unit Dutch auctions when users pay their own bids, because users do not have the incentive for honest revelation of their valuations. Therefore, an *incentive compatible* payment rule is needed; that is, a rule that provides users with the incentives to bid their true valuations. In [6] we address this issue and propose and analyse two payment rules other than "pay-your-bid". We briefly discuss these rules below.

The first payment rule prescribes that each user pays for each unit of bandwidth he reserves at a link the corresponding stop-out price, i.e., the price at which the last unit of bandwidth is reserved at this link. This rule should be applied both to single-link users and path users. It can be seen as a modification of the well-known VCG mechanism. In our case however, there are no losing bids; that is why the stop-out price is employed. Under this payment rule, each user of the guaranteed type has the incentive to be truth-telling, except when being the last winner in one or more links, where he does pay his bid. This is not the case for elastic users; in this case, *demand reduction* [3] is bidders' dominant strategy, thus ruining efficiency.

An efficient approach that produces near-optimal revenue is proposed and analyzed in [6]. The main idea is to create losing bids by not announcing the spare capacity throughout the evolution of MIDAS and charge for the bandwidth allocated according to the VCG rule.

6. CONCLUDING REMARKS

An auction mechanism for reserving bandwidth over paths in a network has been proposed and proven to be a promising approach to a hard problem. This mechanism is also suitable for multicast trees and VPNs. Experiments show that our mechanism is efficient, i.e. the social welfare attained is close to the optimal one. Furthermore, the mechanism has small computational complexity, is not susceptible to collusions and is scalable for large number of links and users. Finally, no assumptions regarding users' utility functions are made, which along with the above features, makes this mechanism easily applicable to real-world networks that serve many different types of customers. Detailed presentation of i) the various payment rules and corresponding bidding strategies, ii) the seller's revenue issue, iii) more experimental results, iv) the extensions of the mechanism for users purchasing bandwidth for different time scales and v) an overall assessment and comparison with related work, is given in [6].

REFERENCES

1. E. Wolfstetter, "Topics in Microeconomics", Cambridge University Press, 1999.
2. P. Milgrom, "Putting Auction Theory to Work: The Simultaneous Ascending Auction", Stanford University, May 1998, <http://www.stanford.edu/~milgrom/>
3. L. Ausubel, "An efficient ascending-bid auction for multiple objects", Univ. of Maryland, June 1997, <http://www.ausubel.com/larry/auction-papers.htm>.
4. A. Lazar and N. Semret, Design, "Analysis and Simulation of the Progressive Second Price Auction for Network Bandwidth Sharing", Columbia University, 1998.
5. Discussion between Huimin Chen and N. Semret on PSP Auction Model, <http://jerry.au.tsinghua.edu.cn/ZLJ/chenhm/res3.htm>
6. C. Courcoubetis, M. Dramitinos, and G. D. Stamoulis, "An auction mechanism for bandwidth allocation over paths", June 2001, <http://www.csi.forth.gr/~mdramit/fullPaper.html>.