Available online at www.sciencedirect.com
ScienceDirect
Computer Networks 51 (2007) 4979–4996
Computer Networks
www.elsevier.com/locate/comnet

An auction mechanism for allocating the bandwidth of networks to their users ☆

Manos Dramitinos *, George D. Stamoulis, Costas Courcoubetis

Department of Informatics, Athens University of Economics and Business, 76 Patision Street, Athens, GR 10434, Greece

Received 20 September 2006; received in revised form 7 February 2007; accepted 23 August 2007
Available online 6 September 2007

Responsible Editor: N. Akar

Abstract

We present a mechanism for auctioning bandwidth on a network-wide basis to end users or ISPs that will utilize it for the same time period. This mechanism consists of a set of simultaneous multi-unit descending-price (i.e. Dutch) auctions, one per link of the network. The per unit prices of bandwidth at the various links are asymmetric, thus reflecting the asymmetry of demand for these links. A user can be instantly allocated bandwidth over a certain path, by simultaneously bidding for the quantity desired at all relevant auctions. This winner determination rule is complemented by a payment rule of the VCG (Vickrey–Clarke–Groves) type, which provides users with the incentive to bid truthfully, thus simplifying bidding. Also, the mechanism enables the auctioneer to use his prior information on market demand anticipated and its spreading among the various links in order to set effectively the auction's parameters. We argue that our mechanism attains nearly efficient allocation of the network's bandwidth (i.e. the resulting social welfare is close to the respective maximum for the quantity decided to be sold by the auctioneer), while it is simple, scalable and applicable to real networks, even for auctioning the capacity of links owned by multiple providers and then splitting the revenue among them. Alternatively, the mechanism offers the provider the opportunity to optimize his revenue, rather than the social welfare. Since our mechanism's computational complexity is low it can serve as a fast, practical, and near-optimal solution to a generally NP-hard optimization problem.

Keywords: Auctions; Networks; Bandwidth markets; Efficiency; Resource allocation

1. Introduction

The rapid growth of the Internet has resulted in an increased need for bandwidth. Link capacities are now very high. Nevertheless, various parts of the Internet remain congested and overprovisioning is considered an ‘‘economically prohibitive luxury’’ [12]. Thus, efficient exploitation of the available

☆ A preliminary version of part of this work was presented in [9]. The authors would like to thank Bjorn Hansen, Frank Kelly, Robin Mason and John Tsitsiklis, for useful discussions on the subject of the paper.
* Corresponding author. Tel.: +30 210 8203693; fax: +30 210 8203686.
E-mail addresses: mdramit@aueb.gr (M. Dramitinos), gstamoul@aueb.gr (G.D. Stamoulis), courcou@aueb.gr (C. Courcoubetis).

1389-1286/$ - see front matter
doi:10.1016/j.comnet.2007.08.007

capacities is of high importance for providers. Moreover, the competition among providers is high, while the market for Virtual Private Networks (VPNs) is expanding rapidly. Thus, the static, long-term bandwidth contracts that were in place in the past are being replaced by short-term customized ones that are of higher value to customers. The market for bandwidth becomes more liquid over time [16] and emerging technologies, such as VoIP and Grid, further intensify this trend.

Within this competitive economic context, pricing can serve as a control mechanism for allocating the network resources. A popular pricing scheme is to advertise posted prices for the consumption of units of network resources. The main merit of this approach is that it is simple for the users, which is an attractive feature [24]. However, this scheme cannot adapt satisfactorily to varying network conditions and prevent congestion [10]. Though there have been proposed schemes that theoretically can be asymptotically efficient, computing the prices that can achieve this, and adapting them to varying traffic mixes, is a computationally intractable problem [25]. This limits such schemes' applicability. On the contrary, dynamic pricing is superior to static pricing in economic terms. This has been both proven in theory [10] and observed in practice [2,11]. To this end, a popular method of allocating goods is the use of auctions. Auctions offer the advantage of transparency and simplicity in determining market-based prices and economic efficiency (i.e. social welfare maximization), since certain auctions can guarantee that goods are acquired by those that value them the most. Furthermore, auctions may lead to higher revenues for the providers compared to traditional trading methods, due to the competition arising. There are numerous auction mechanisms, related software and applications, newsletters and specialized search engines; see [10] and references therein.

Nowadays, special markets called bandwidth markets (or exchanges) offer a large number of point to point circuits of certain capacity [6]. Other bandwidth exchanges, such as Arbinet [3], eSwitch [14] and Min-X [23] focus on real-time aggregation and matching of supply and demand for call-minutes in spot markets. Liquid bandwidth exchanges offer real-time trading of bandwidth for time scales as small as 5 min and allow the buyers to purchase bandwidth from any available seller, without binding contracts [15]. An extensive list and description of bandwidth markets can be found in [26]. Furthermore, it is stated in [6] that the site of such an exchange, namely that of Bandwidth Market Ltd., contains over 300,000 offers and bids for circuits. It is also stated that "*commodity traders think bandwidth will be a major new commodity*". Moreover, vendors of network equipment sell a wide variety of routers that enable the multiplexing of thousands of simultaneous VPN or private circuits sessions [8]. Note also that this functionality can be further enhanced by advanced Traffic Engineering software that implements the virtual router (VR) concept [13]. Thus, it is feasible for the network providers nowadays to implement in their network advanced resource allocation policies of the finest granularity. Also, utilities – such as electricity companies – that have telecommunications networks and services are expected to offer part of their networks' capacity for sale to bandwidth markets. Therefore, it is theoretically interesting and practically important to design an auction mechanism for selling bandwidth on a network-wide basis to possibly many bidders.

However, so far very few innovative auction mechanisms have been devised in order to allocate the bandwidth of a *network* of arbitrary topology (rather than that of just one link). This is mainly due to the high complexity of the problem, because of the large size of the network, the large population of competing users and most importantly the inherent characteristics of users' demand in this context: Users demand the reservation of *multiple* units of bandwidth across paths; it must be ensured that the *same* quantity of bandwidth shall be allocated at *all* the links constituting each user's path. This large set of user-imposed constraints combined with the large number of resources whose demand is interdependent in the market, render the problem of optimal resource allocation (even in the case of full market demand) NP-hard.

In this paper, we propose and evaluate an innovative auction mechanism for the allocation of a network's bandwidth to its users. This mechanism consists of multi-unit Dutch auctions. The auctions' clocks are reduced according to a price dropping policy already defined and assessed experimentally under the assumption of truthful bidding in [9]. In this paper, a new payment rule is introduced that provides indeed the incentives to users to be truthtelling. This simplifies the determination of the optimal biding strategy. Furthermore, it is shown how the proposed auction can utilize the provider's prior information so that the auction's performance is improved and the network's bandwidth is allocated (nearly) efficiently. The proposed mechanism is decentralized, applicable in practice, performs well with respect to the social welfare attained and does not suffer from a set of problems that appear in most related work (as explained in Section 5).

The remainder of this paper is organized as follows: In Section 2, we provide some background on auctions. In Section 3, we define the problem addressed, highlight its most important aspects and assess certain approaches regarding the design of an appropriate auction. In Section 4, we present the proposed auction mechanism, experimentally assess our mechanism's performance and introduce an incentive-compatible payment rule. In Section 5, we compare our work with related research on network bandwidth allocation, while in Section 6, we present an extension of our mechanism for auctioning bandwidth over a longer time horizon. Finally, in Section 7, we provide some concluding remarks.

2. Background on auctions

In this section, we present some fundamental definitions, theorems and results from auction theory; further details are presented in several references; e.g. see [10,17].

An auction is a mechanism based on a pair of rules, namely the *allocation rule* that defines which good is allocated to whom and the *payment rule* that defines the charge of the auction winner(s). A participant of an auction is called *bidder*, while the entity conducting the auction *auctioneer*. Auctions are referred to as *simple* or *single-unit* if only one good is auctioned and *multi-unit* if multiple units of a good (e.g. integral units of a link's bandwidth) are to be traded. Moreover, depending on whether bids are made in public or submitted as sealed envelopes, the auction is referred to as *open* or *sealed* respectively. Auctions maximizing seller's revenue are referred to as *optimal* and those maximizing social welfare are referred to as *efficient*. If the auction is conducted in rounds, then it is called *progressive*.

A *bid* in the context of simple auctions is the amount of money offered by a bidder for the item auctioned. The best-known mechanism is by far the *English* auction, where the seller starts with a minimum price that is gradually incremented until there is only one person claiming the item, whom the item is awarded to. The *Dutch* auction corresponds to the opposite mechanism. The price is initially high and is gradually decremented until a bidder claims the object. The item is awarded to him for a charge equal to the current price. The sealed bid auctions (1st price and 2nd *price* or Vickrey) consist of two phases: (a) the first one, where bidders submit sealed envelopes with their bids and (b) the second one, where these envelopes are opened. The item is then awarded to the bidder who submitted the highest bid. The winner pays his bid at the 1st price auction and the highest losing bid – i.e. the 2nd highest bid – at the Vickrey auction. It has been proved that under the Vickrey auction it is best for each bidder to honestly bid his true value for the item being awarded. This property is referred to as *incentive compatibility*.

A *bid* in the context of multi-unit auctions is defined to be the pair (p,q) of the per unit expressed willingness to pay p for a quantity q of units. All simple auctions can be generalized to multi-unit auctions. Incentive compatibility holds only for the generalizations of the Vickrey auction. The rules of the Generalized Vickrey Auction (GVA) prescribe that: (i) each user reports his valuation for a subset or for all points of his demand function for units of the good auctioned, (ii) units are allocated to the highest bids until demand exhausts supply, (iii) each user is charged according to the social opportunity cost that his presence entails. Hence, each user pays for the units he is awarded the losing bids that would have become winning if all his own bids were set to zero. Hence, winners pay less than their respective bids.

3. Auctioning bandwidth

In this section we present the various requirements that should be met by our design and motivate our choice of a Dutch type auction for allocating bandwidth. We also demonstrate why some other well known approaches are not applicable in practice.

3.1. Assumptions and objectives

We assume that there are N links whose bandwidth is being auctioned *simultaneously*. In particular, the bandwidth of each link is sold in integral "units". To keep the presentation general, we refrain from mapping the term "unit" to a specific network technology or bit rate. This is also in accordance with the fact that, as a result of the links' large capacities, providers sell virtual circuits of various bit rates. The granularity of the latter, i.e. the "unit", is determined by the underlying technology and cannot be arbitrary. However, note that in practice the demand will also take this into account. Technically, it is natural to assume that since users' access and transport rates are not arbitrary, the same applies to their demand (and quantities of bandwidth asked in their bids). Nevertheless, for users of guaranteed applications, this impact is expected to be minimal, because the rates at which these applications are offered is in line with the granularity of bandwidth offered in networks.

Our objectives regarding the mechanism sought are:

1. Bandwidth should be awarded efficiently, i.e. to those users that value it the most.
2. Each user should be able to reserve the *same* quantity of bandwidth at all the links he is bidding for.
3. The mechanism should be practically applicable. This implies that the mechanism should be (i) *scalable* with respect to both the size of the network and the number of users and (ii) *reliable* and *not susceptible to dishonest dealing*.
4. The mechanism should enable providers not to disclose to users any unnecessary information regarding the size of their network, without affecting efficiency.
5. The mechanism should enable providers to use their prior private information on market demand in order to set the auction's parameters so that its performance be optimized.
6. The mechanism should be able to simultaneously auction links that belong to different networks, allowing users to build paths by reserving bandwidth on them. In this case, it is required that the "sharing" of the revenue of paths to the respective owners of the links comprising those paths in a both fair and acceptable way be feasible.

Note that the aforementioned assumptions are common among all research works on auction mechanisms addressing the same problem (and discussed in Section 5). In particular, without simultaneous auctions, it is impossible to guarantee that users will be able to reserve the desired quantity of bandwidth at all the links of their interest under any auction or any other trading mechanism and subsequently evaluate the mechanism's performance. This is crucial, since a path with different bandwidth at its constituent links contains some portions of bandwidth that are useless to the user, thus introducing losses to both the user and the social welfare. For instance, a user reserving at the two links comprising his path 2 Mbps and 4 Mbps, would be better off by reserving 2 Mbps everywhere since the latter reservation is cheaper than the former for the same pipe bit rate. On the contrary, sequential auctioning or trading in general cannot guarantee that such costly and inefficient aggregations of resources will not happen. In the context of auctions, this is referred to as the *exposure problem*: Bidding individually for complementary items exposes bidders that seek synergistic combinations to aggregation risk. This risk limits competition in the market and often leads to inefficient allocations. Therefore, simultaneous auctions are more preferable than sequential ones [10]. Finally, note that in our mechanism, the simultaneous auctioning of all links, combined with the incentive compatibility property, simplifies the bidding strategy of our mechanism: Users bid for a unit when it becomes profitable to do so, based on the link prices offered at the auction. This would not be the case if sequential auctioning of the links had been adopted.

3.2. Business models

Our mechanism can be applied in bandwidth markets, in order to enable their client networks to dynamically build point-to-point interconnection, multicast trees or complex VPNs of arbitrary topology (as depicted in Fig. 1). It is implicitly assumed that the bandwidth of the various links is auctioned simultaneously for a predefined time period, which has been publicly announced. It is possible that a bandwidth market sells the capacity of links that belong to several network providers on their behalf, by aggregating them and auctioning them simultaneously. As a result of the properties of our mechanism, *revenue sharing* among different providers is feasible, which would not necessarily be the case with other mechanism designs, e.g. with a combinatorial approach (as explained later in this paper).

The business relationship model and the network conceptual model are depicted in Fig. 1. An important related case is that of Internet Backbone Providers (IBPs) who offer connectivity through their backbone to their customers, namely the Internet Service Providers (ISPs). In this context, it is clear that revenue is an important objective. However, customer satisfaction and thus economic efficiency is also important. These two objectives can be in-

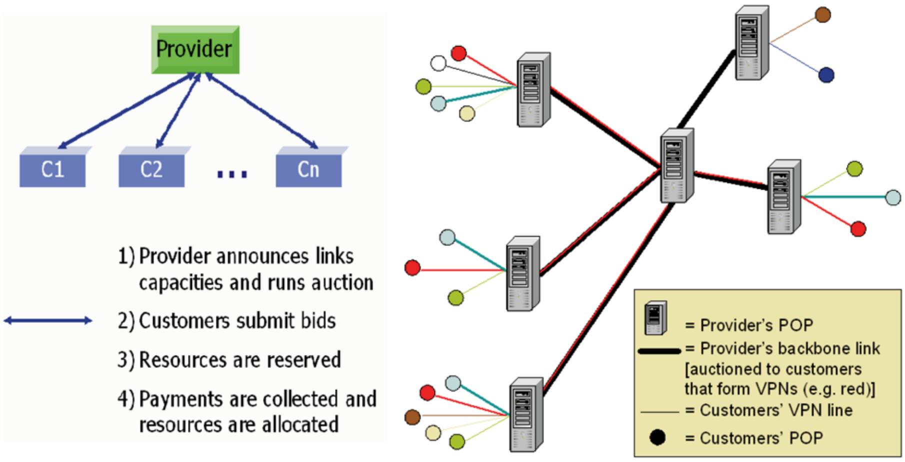

Fig. 1. The business relationship model and the network's links and customers.

line, at least to a certain extent, since by awarding more bandwidth to the high value customers, the provider is well off in terms of revenue as well.

Another field of application of our mechanism is that of a network interconnecting academic or research institutions. These institutions consist of a large number of departments and laboratories that are geographically scattered and need bandwidth for interconnection and Internet access purposes in partly overlapping paths. Our auction can be used as an internal bandwidth allocation mechanism, since efficiency in the exploitation of the bandwidth available is the primary objective in such an environment.

3.3. Why combinatorial approaches are inapplicable

The requirement that each user should be able to reserve the *same* quantity of bandwidth at all the links he is bidding for, seems to motivate the choice of a combinatorial mechanism. In a combinatorial auction of M items-in our context each unit of bandwidth of every link is such an item-each bidder may bid $b_i(S) > 0$ for any combination (bundle) of items $S \subseteq M$. However, winner determination in this case is in general NP-hard [17]. Therefore, this approach is inapplicable for real world networks. Hence, it is important that the auction's complexity be low, even for large networks and populations of users. Clearly, this can be achieved only if the auction is implemented in a distributed, non-combinatorial fashion, which is the case with our mechanism.

Moreover, *revenue sharing* among different providers would not be feasible with a combinatorial approach: Since bidders declare their total willingness to pay for a path consisting of links belonging to different network providers, it is impossible for the auctioneer to decide on the "sharing" of the revenue of paths to the respective owners of the links comprising those paths in a both fair and acceptable way.

3.4. Why ascending auctions are inapplicable

The main problem when employing ascending auctions for bandwidth in multiple links is as follows: Because of the difference in the demand per link, it is impossible to synchronize the auctions of the various links so that all of them terminate at the same time. For simplicity, we assume that there are three users participating in the auctions: user 1, who is interested in purchasing bandwidth in the first link; user 2, who is interested in the second link; and, user 3, who is a path user, i.e. he is interested in purchasing bandwidth in the path. Each of the users has a utility function $u_i(\cdot)$ that describes his valuation for various quantities of bandwidth. The auction outcome consists of the quantities of bandwidth x_1, x_2 and x_3 purchased by the three users, and by the corresponding payments. This outcome is efficient if the allocation of bandwidth maximizes social welfare, subject to capacity constraints:

$$\max \quad \{u_1(x_1) + u_2(x_2) + u_3(x_3)\}, \tag{1}$$

$$\text{subject to} \quad x_1 + x_3 \leqslant C, \tag{2}$$

$$x_2 + x_3 \leqslant C. \tag{3}$$

This problem may be solved by means of Lagrange multipliers (see pp. 122–123 of [10]), which correspond to the per unit prices of bandwidth p_1^*, p_2^* at each link. At the efficient outcome, there should hold:

$$u_1'(x_1^*) = p_1^*, \tag{4}$$

$$u_2'(x_2^*) = p_2^*, \tag{5}$$

$$u_3'(x_3^*) = p_1^* + p_2^*, \tag{6}$$

$$p_1^*(C - x_1^* - x_3^*) = 0, \tag{7}$$

$$p_2^*(C - x_2^* - x_3^*) = 0. \tag{8}$$

The above set of equation defines the unknown pair of prices, that when offered simultaneously at the two links, lead each player to individually select such a quantity of bandwidth that the overall social welfare is maximized. Assume now that the auctioneer tries to "discover" these prices, by means of two simultaneous ascending auctions, one per link. That is, the unit price in each link increases with time t, and users decrease their demand accordingly.

In order for such an approach to be successful, it must be ensured that the optimal pair of prices will always be offered simultaneously at some point in time in the two links. That is, it should always be feasible to discover dynamically a monotonic price path that passes through the optimal point. Unfortunately, this is not always possible, even if the price path adapts to the demand expressed so far. To see this, first note that we are essentially looking for a monotonic price path that both lies in the interior of the box shown in Fig. 2 and crosses the optimal point (p_1^*, p_2^*). Note that if prices are within the box of the Fig. 2, the following set of inequalities hold:

$$x_1(t) \geqslant x_1^*, \tag{9}$$

$$x_2(t) \geqslant x_2^*, \tag{10}$$

$$x_3(t) \geqslant x_3^*. \tag{11}$$

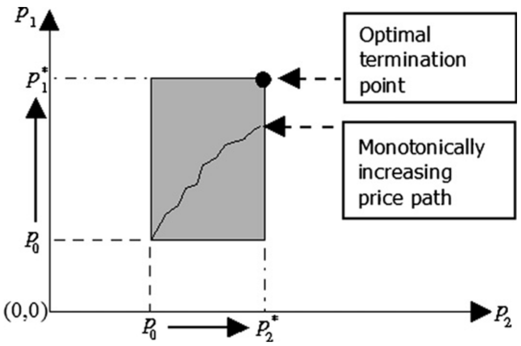

Fig. 2. Attempt of ascending auctions to reach the optimal pair of prices.

The objective of the auctioneer is to attain all equalities at the same time, otherwise the auctions will not terminate simultaneously. This is impossible since any price increasing policy can lead to a price vector $\langle p_1', p_2' \rangle$ for which exactly one of the prices exceeds its respective optimal value (i.e. either $p_1' > p_1^*$ or $p_2' > p_2^*$), while the excess demand for each of the links is positive both for the price vector $\langle p_1', p_2' \rangle$ as well as for all vectors of the trajectory that led to $\langle p_1', p_2' \rangle$. In such a case, the price vector $\langle p_1', p_2' \rangle$ is outside the shaded box of Fig. 2, either to its left or above it, while, due to monotonicity, the optimal price vector cannot be attained in the future. For instance, consider a price increasing policy that led to the price vector $\langle p_0 + \epsilon, p_2^* + \epsilon \rangle$ so that $p_0 + \epsilon < p_1^*$; this is depicted as the black square of Fig. 2. Because of the low per unit price of bandwidth at the first link, the cost of the bandwidth for the path is significantly less than $p_1^* + p_2^*$; this renders the excess demand at both links positive, while the per unit price of bandwidth at the second link has exceeded the optimal price p_2^*. Therefore, the network's bandwidth cannot be allocated efficiently. Moreover, in such cases auctions will not terminate simultaneously. Hence, *it is impossible for the path bidders to reserve the same quantity of bandwidth at their path's constituent links.*

Concluding, ascending auctions can meet neither the objective of efficient bandwidth allocation nor the requirement that each path bidder is allocated the same quantity of bandwidth at all the constituent links of his path. It is worth noting that this would not be the case if Dutch auctions were employed. This stems from the fact that in the latter case bids are *instantly* mapped to allocations, as opposed to ascending auctions where resource allocation is performed when a certain price is reached where demand equals supply. However, appropriate price synchronization is also necessary with Dutch auctions in order to attain efficiency. This is because the prices offered simultaneously at the various links should enable the allocation of bandwidth to all path- and single-link bidders that value it the most (see 4.2).

4. The MIDAS mechanism

4.1. The mechanism in brief

The proposed mechanism consists of simultaneous multi-unit Dutch auctions one for each link (MIDAS). Although these auctions run in an independent fashion, they need to be coupled by pre-specifying different offered prices at the various link auctions, as explained later. For link l, the total capacity auctioned C_l is announced together with the initial unit price $p_l(0)$. This initial unit price should be high, for example a few times a reasonable market price. At each link, the price is reduced as time elapses, until it becomes zero. Users place their bids and are instantly allocated bandwidth over links or paths. Although not enforced by rule, it is implicitly assumed that users interested in a path would try to reserve bandwidth at *all constituent links simultaneously*. Bidders are allowed to bid several times even for the same link(s), thus accumulating bandwidth if desirable. The feedback of the auction to the bidders is limited to the prices of the links. Without loss of generality, we henceforth assume that each user is interested in reserving bandwidth over a pre-decided network path. Note that the auctioneer may set a reserve price at each link. Unless otherwise specified, this is taken equal to 0. In the remainder of this section, we illustrate the fundamental properties ensuring appropriateness of our mechanism, under the assumption of truth-telling bidders. We show how the performance of the proposed mechanism can be optimized by employing information on the demand that the auctioneer may have available. We also present a payment rule of the VCG type, which complements our mechanism and enforces incentive compatibility and thus truthful bidding. Finally, we provide a comprehensive example of a MIDAS auction execution.

4.2. Coupling of demand and auctions among links

As already explained, the main merit of employing descending auctions in our context lies in instant allocation. In order for MIDAS to meet the objectives of efficiency and high revenue, we argue that link prices should be reduced in a smart way. In particular, the decrements of the link prices should not be all the same. This should be contrasted with the policy where the prices in all links are the same at any time; i.e., all initial prices are the same, and all decrements are the same too. This policy will be henceforth referred to as *symmetric-prices*, and suffers from the following drawback: The bandwidth of the most popular links is offered at the same price with the bandwidth of the links having less demand. This is clearly inefficient – since competition may not be the same everywhere – and may lead to high charges for paths. Also, by charging the links with low and high demand at the same prices, the network resources are not allocated to those who value them most.

Consider for example the case depicted in Fig. 3. Although allocation of bandwidth to a path bidder may be efficient from an economic point of view, it is not likely to happen in the auction if the approach of symmetric prices is adopted. Indeed, if link 1 users start placing their bids, then the path user is forced to pay in total (per unit of bandwidth) double the price of the congested link 1, in order to avoid exhaustion of capacity of link 1. This may prevent such a user to bid; hence the path may not be built. Link 2 should have been priced much lower, because it has actually lower demand than link 1, thus helping path-bidders reserve bandwidth at all constituent links simultaneously. In order for the auction to attain improved efficiency it is necessary that prices at different links be reduced asymmetrically, in accordance to their respective expressed demand.

This shows that since demand for bandwidth at the various links is "coupled", so are the optimal market-clearing prices at these links. Indeed, bottleneck links greatly affect the final outcome with respect to the users' bandwidth reservations and charge. This is to be expected since users' paths are mainly constrained in terms of both quantity and budget by the popular links, where the competition is fierce. Thus, it is important that the prices of the congested links throughout the auction remain higher than those of the non-popular links.

Prior to proceeding with the details of the price reduction policy that we propose for MIDAS, we describe how the mechanism works: Consider a user who wishes to purchase bandwidth over links 1 and 2. He observes the trajectory of the corresponding prices $p_1(t)$ and $p_2(t)$ and based on his utility function $u(x)$ submits a bid for a quantity y of bandwidth. The value of y depends on the auction's payment rule and thus the bidder's strategy; this issue is thoroughly presented in Section 4.4. As already mentioned, bidding simultaneously for both

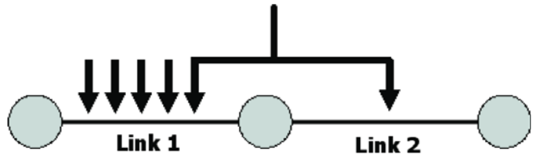

Fig. 3. Different demand for two links.

links of interest ensures that the same quantity of bandwidth is purchased over the links that constitute his path. As soon as the bid is submitted for each of these links, the auctioneer instantly allocates y units of bandwidth to this user, and updates the prices of these links, according to the MIDAS price reduction policy presented in the following subsections.

4.3. MIDAS price reduction policy

We have designed and evaluated experimentally a certain price reduction policy that is an integral part of MIDAS, namely the Price Freezing (PF) Policy: An initial high price is set at each link. Subsequently, link 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_l of the bandwidth just allocated in this link; that is, the freezing period equals $s \cdot x_l$. If additional allocations occur during the period of freezing then the link price is kept frozen for more time accordingly. A comprehensive example is provided at Section 4.8. It is easily seen that, when the price of a link l is not frozen, the following equation holds:

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

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 the parameters r and s influence the pace of the auction. The motivation for this policy is that it leads to the above explicit relation between the price and the spare capacity of each link.

Proposition 1. *There exists a choice of initial prices for the various links such that the following property holds: At any time instant, for every pair of links for none of which prices are frozen, the link with the lowest spare capacity has the highest price.*

Proof. Assume that at time t, link l has less spare capacity than link j, i.e. $C_l - x_l(t) < C_j - x_j(t)$. In order to obtain the right ordering of prices, namely $p_l(t) > p_j(t)$, we should have: $p_l(0) - r \cdot t + r \cdot s \cdot x_l(t) > p_j(0) - r \cdot t + r \cdot s \cdot x_j(t) \iff p_l(0) + r \cdot s \cdot [x_l(t) + C_l - C_l] > p_j(0) + r \cdot s \cdot [x_j(t) + C_j - C_j] \iff [p_l(0) + r \cdot s \cdot C_l] - r \cdot s \cdot [C_l - x_l(t)] > [p_j(0) + r \cdot s \cdot C_j] - r \cdot s \cdot [C_j - x_j(t)]$ Since $C_l - x_l(t) < C_j - x_j(t)$ and $r \cdot s > 0$, there holds $-r \cdot s \cdot [C_l - x_l(t)] > -r \cdot s \cdot [C_j - x_j(t)]$. Therefore, in order for the inequality desired to apply, it suffices to have $p_l(0) + r \cdot s \cdot C_l = \tilde{p}\ \forall l$, where $\tilde{p} > 0$, i.e. set the initial price of every link to $p_l(0) = \tilde{p} - r \cdot s \cdot C_l$. □

Therefore, when the auctioneer has no information on the demand, he should configure the auction so that the price of bandwidth at links of less capacity (i.e. supply) is offered at higher prices than that of links where bandwidth is less scarce. However, the auction's performance can be further improved if the initial prices are set in a way that price is kept higher at links where *competition* for bandwidth is more intense. Indeed, it is competition that defines the optimal cut-off price per link. We revisit the problem of configuring MIDAS when such information is indeed available so that the outcome is nearly efficient at the following subsections of the paper.

Proposition 2. *The maximum MIDAS execution time is proportional to the link's initial price and capacity.*

Proof. The auction of each link will have surely terminated at time t_l^0 when price drops to 0. Setting $0 = p_l(t_l^0) = p_l(0) - r \cdot t_l^0 + r \cdot s \cdot x_l(t_l^0)$, we obtain $t_l^0 = \frac{p_l(0)}{r} + s \cdot C_l$.

If we further assume that $p_l(0) = \tilde{p} - r \cdot s \cdot C_l$ we obtain: $t_l^0 = \frac{\tilde{p}}{r}$. □

Proposition 2 indicates that there is a maximum MIDAS execution time, that can be defined *a priori* by the provider, by selecting proper values for $p_l(0)$, r and s. This maximum duration is also independent of the number of the bidders participating, as opposed to other mechanisms that address the same problem (see Section 5). Also, each link's price policy depends only on local information on the demand exhibited so far, and thus MIDAS is implemented in a distributed fashion. Therefore, MIDAS is scalable with respect to the number and capacities of links and the number of users. Note also that in practice if there is enough competition among bidders, the auction can be terminated much earlier than the $t_l(0)$ of Proposition 2, i.e. bandwidth will be exhausted at prices higher than 0. This is also depicted in the example of Section 4.8.

4.4. Assessment of MIDAS

As already mentioned, the MIDAS price reduction policy has been specified in such a way that prices tend to reflect the difference in the demand for the various links of the network. We have performed extensive experimental assessment of our auction mechanism under the assumption of *truth-telling* users. A more detailed exposition of these experiments and their results is given in [9]. The experiments carried out regard two network topologies, that is linear and dumbbell (with up to 20 links), which have been provided as Figs. 4 and 5. The linear topology was chosen because of its simplicity, while the dumbbell was chosen because of the hierarchical form of the Internet.

For all the experiments provided below, as well as for those presented in [9], the initial link prices are set according to Proposition 1. Hence, it is implicitly assumed that the provider has no prior information on market demand, which corresponds to the least favorable conditions of the assessment. These experiments assess our price reduction policies with respect to the social welfare attained. In particular, no matter if the demand for bandwidth at different links is the same or varies a lot, our auction mechanism performs well in terms of efficiency. Losses are up to 2% of the maximum social welfare for PF. In order to further assess the effectiveness of the PF policy, a comparison of it has been made with symmetric-prices. This policy is significantly less efficient than our price reduction policy. The social welfare attained under this policy is typically 85% of that attained under PF. Moreover, whenever the optimal social welfare can be analytically computed, the proximity of the social welfare attained under MIDAS with the optimal has been studied. Tables 1 and 2 contain some typical experimental results.

Next, we provide an intuitive explanation of the success of the MIDAS PF policy which has been illustrated in this subsection by means of its experimental assessment. In order for MIDAS to maximize the social welfare attained, it suffices to offer simultaneously at all links the vector of optimal prices (denoted as $\langle p_l^* \rangle$). Obviously, the p_l^* of popular links are expected to be higher than those of less popular links. The MIDAS PF policy guarantees that the relation of prices offered in the auction is the same with that of the unknown $\langle p_l^* \rangle$, by dynamically adapting to the demand exhibited so far: At a given price, more units are purchased in a popular link than at a less popular one, due to the higher demand. This implies that under the MIDAS PF policy, the price of a popular link freezes for more time than that of a less popular one. Therefore, the price of the former link is prevented from decreasing as fast as the one of the latter, and therefore is in general higher. Thus, the ordering of the link prices $\langle p_l(t) \rangle$ gradually becomes the same with that of the $\langle p_l^* \rangle$. This also explains the reduced performance of symmetric pricing. Indeed, this policy prescribes that $\langle p_l(t) \rangle$ is the same for all links at any time t of the auction, thus leading to an incorrect relation of the link prices and typically low social welfare. In fact, the higher the size of the network and the less uniform the spread of demand across links, the worse this policy performs.

Fig. 4. A linear network of three links where the middle link is congested.

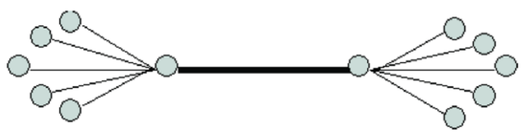

Fig. 5. A dumbbell network where the middle link is congested.

Table 1
Price reduction policies versus symmetric reduction of prices

	Dumbbell network of 15 links		Linear network of 15 links	
	Symmetric pricing	PF policy	Symmetric pricing	PF policy
Social welfare	40,052	44,209	36,280	42,371

However, the fact that MIDAS PF orders the link prices properly – according to demand – does not suffice to ensure that the optimal price vector $\langle p_l^* \rangle$ is indeed offered in the auction. Therefore, it is possible that units of bandwidth may be misallocated, thus resulting in loss of social welfare. This loss depends on many factors, such as the distance of the MIDAS price vectors from the optimal, the elasticity of the users' demand for prices around $\langle p_l^* \rangle$, the topology of the network and the spread of the demand across links and paths. Extensive experimental assessment carried out has revealed that MIDAS PF policy is in general successful, approaching the efficient allocation. Nevertheless, no general bound of its performance can be obtained for the case where the provider is totally unaware of the market demand. However, if such knowledge is available, the mechanism can be

Table 2
Experimental results for two linear networks

	Linear network of three links (10 Mbps each), five users	Linear network of two links (10 Mbps each), 50 users
Optimal	1572	1556
Social welfare	(computed exhaustively)	(computed by means of a special algorithm [9])
MIDAS social welfare	1572	1529
(Price freezing)	(100% of optimal)	(98.26% of optimal)

"fine-tuned" in order to mitigate this loss. This is discussed in the following subsections, where it is shown how the maximum social welfare can be achieved or approximated by exploiting such information.

4.5. MIDAS attains efficiency under full information

We briefly present here how MIDAS can be ex-ante "configured" so that by taking into account the auctioneer's full knowledge about market demand, efficiency is achieved. Indeed, a network provider has good prior knowledge of market demand at his various links, so it is important for him to be able to use this knowledge for his benefit. This is done by setting at the various links such initial prices – and then applying the PF policy – so that after some time elapses all the link prices will simultaneously reach the optimal set of prices (denoted as p_l^* for each link l). Let us consider a network of arbitrary topology. We assume that the auctioneer has full information about market demand. Had the link prices been fixed and announced by the provider, the demand would have given rise to the efficient allocation. This assumption will be relaxed later. By applying the Lagrangian methodology, the provider may derive the optimal link prices. Assuming that the provider would still be interested in running the MIDAS auction (e.g. for transparency reasons), could he configure the PF policy so that bandwidth is allocated efficiently after running MIDAS for some fixed total time t?

Proposition 3. *MIDAS PF policy allocates bandwidth efficiently when the optimal prices of the links are known a priori.*

Proof. The price of a link l under PF will be $p_l(t) = p_l(0) - r \cdot [t - s \cdot x_l(t)]$. The optimal link price p_l^* and respective allocation x_l^* can be computed by means of Lagrange multipliers. Therefore, in order to have for some t, $p_l(t) = p_l^*$ and $x_l(t) = x_l^*$ $\forall l$, it suffices to set each initial link price to be

$$p_l(0) = p_l^* + r \cdot t - r \cdot s \cdot x_l^* \quad \text{or equivalently}$$
$$p_l(0) = \bar{p} + p_l^* - r \cdot s \cdot x_l^*,$$

where $\bar{p} > 0$ determines the highest initial link price that the auctioneer may wish to set. □

Note that assuming competition for a certain link, i.e. $x_l^* = C_l$, we obtain from Proposition 3 that $p_l(0) = \bar{p} + p_l^* - r \cdot s \cdot C_l$. This is similar to the result of Proposition 1, in the case that the auctioneer has no information regarding the market demand, where $p_l(0) = \tilde{p} - r \cdot s \cdot C_l$. Therefore, in order to optimize the auction's performance, the auctioneer uses his information regarding the market demand, by replacing the "common" $\tilde{p}$ of the first formula with the "differentiated" $\bar{p} + p_l^*$ of the second formula.

It is worth emphasizing that not all auction mechanisms could take good advantage of full information. For example, the Progressive Second Price auction of [18] (discussed in 5) does not involve any parameters or rules which the auctioneer could fine-tune using such information.

4.6. Using the available information

One might argue that this property of MIDAS stated in Proposition 3, though attractive in theory, is useless in practice, since the full information assumption is too restrictive. However, in practice, the provider has in general some information about the market demand. He is aware to a certain extent of his customers' needs with respect to the desired paths and the quantities of bandwidth they are interested in purchasing. Thus, the various links can be classified according to their popularity; the number and the paths of the users that are interested in purchasing bandwidth at the various links can also be approximated. In order to quantify the demand for bandwidth on a link, we assume a specific utility function u of a "small" user. Large customers can be seen as users of the same type yet having a total scaled demand function. Equivalently, a large customer can be seen as a group of small users. Hence, the higher the number of small users crossing a network link, the higher the demand for the bandwidth of this specific link. This information can be exploited in order to ensure that the initial per unit prices at the various links are ordered according to their respective popularity, thus improving the mechanism's performance.

4.6.1. The case of linear demand

Next, we assume that the demand of bandwidth per link can be approximated by a linear function having a slope that is proportional to the number of users crossing the link; all users are taken as identical in terms of demand. We show how MIDAS PF initial link prices can be set to allocate the network bandwidth efficiently.

Indeed, let the demand per link be expressed by the linear demand function $d_l(p_l) = D_l^{\max} - \frac{g}{n_l} \cdot p_l$, where g is a normalization factor and n_l denotes the number of unary users whose paths cross the link.

Assuming that there is competition at all links, we can derive the optimal price p_l^* for link l as follows the fact: $d_l(p_l^*) = C_l$, which implies that $D_l^{\max} - \frac{g}{n_l} \cdot p_l^* = C_l \; \forall l$. Therefore, $p_l^* = \frac{D_l^{\max} - C_l}{g} \cdot n_l$. Thus, it suffices to set initial link prices $p_l(0) = p_l^* + \bar{p} - r \cdot s \cdot C_l = \frac{D_l^{\max} - C_l}{g} \cdot n_l + \bar{p} - r \cdot s \cdot C_l$. Also note that we can further assume that all unary users crossing link l are interested in buying one unit of bandwidth at most, i.e. at a price equal to 0. Then, we obtain $D_l^{\max} = n_l$, which implies that $p_l(0) = \frac{n_l - C_l}{g} \cdot n_l + \bar{p} - r \cdot s \cdot C_l$.

This is actually a special case of that of Proposition 3 in case the estimate of $\langle n_l \rangle$ is accurate for all links. This result shows how the auctioneer should set the initial prices at the various links in the general network case. The assumption of linear demand functions is common in the literature [4] and though in our case the demand at the various links is interdependent—because of the path bidders—this is a necessary simplification that provides an intuitive explanation why the MIDAS PF policy is successful. Obviously, the more accurate the estimate regarding $\langle n_l \rangle$, the better the mechanism's performance. In practice, the auctioneer may obtain an accurate estimate of $\langle n_l \rangle$ after running MIDAS for a few times: The auctioneer *does* acquire full information about the market demand *after* the auction is conducted (ex post). This information can be used to order the links according to their popularity and therefore set the respective initial link prices accordingly. Moreover, since the auction is repeated frequently, it is feasible for the auctioneer to keep statistics of the price fluctuation patterns of his links and obtain a fairly accurate approximation of the market demand $\langle n_l \rangle$ of the next MIDAS auction.

4.6.2. Various cases of network topologies

We proceed to show know how the aforementioned general principle regarding the setting of the initial prices of the links, can be applied to some simple network topologies for which there is substantial information regarding the market demand.

Let us consider the linear network of Fig. 4. (The methodology presented here can be applied to more complex network topologies as well and is based on the aforementioned notation of "small" users.) Assume that the network provider has some information regarding customer demand for paths on the aforementioned network. In our example, assume that there are path bidders interested in building a path consisting of links 1 and 2, and some others whose path consists of links 2 and 3. Also assume that there are some single-link bidders interested solely in link 2. Hence, the middle link is the most popular. The respective (aggregate) utility functions, denoted as $u_{12}(x_{12})$, $u_{23}(x_{23})$ and $u_2(x_2)$ respectively, are:

$$u_{12}(x_{12}) = N_{12} \cdot u(x_{12}/N_{12}), \tag{12}$$

$$u_{23}(x_{23}) = N_{23} \cdot u(x_{23}/N_{23}), \quad \text{and} \tag{13}$$

$$u_2(x_2) = N_2 \cdot u(x_2/N_2), \tag{14}$$

where $u_{12}(x_{12})$, $u_{23}(x_{23})$, $u_2(x_2)$ denotes the utility of an individual user of each category of users.

In order for the provider to compute the optimal allocations of bandwidth to each group of users and the optimal prices of the links 1, 2 and 3, it suffices to solve the following optimization problem:

$$\max \quad \left\{ N_{12} \cdot u\left(\frac{x_{12}}{N_{12}}\right) + N_{23} \cdot u\left(\frac{x_{23}}{N_{23}}\right) + N_2 \cdot u\left(\frac{x_2}{N_2}\right) \right\}, \tag{15}$$

$$\text{subject to} \quad x_{12} \leqslant C_1, \tag{16}$$

$$x_{12} + x_2 + x_{23} \leqslant C_2, \tag{17}$$

$$x_{23} \leqslant C_3. \tag{18}$$

This optimization problem is solved by means of Lagrange multipliers, which correspond to the per unit market-clearing prices of bandwidth p_1^*, p_2^*, p_3^* at each link:

$$\max\left\{N_{12}\cdot u\left(\frac{x_{12}}{N_{12}}\right)+N_{23}\cdot u\left(\frac{x_{23}}{N_{23}}\right)+N_2\cdot u\left(\frac{x_2}{N_2}\right)-\lambda_1\cdot(C_1-x_{12})-\lambda_2\cdot(C_2-x_{12}-x_2-x_{23})-\lambda_3\cdot(C_3-x_{23})\right\}. \quad (19)$$

By partial differentiation of (19) we obtain:

$$u'\left(\frac{x_{12}}{N_{12}}\right)+\lambda_1+\lambda_2=0, \quad (20)$$

$$u'\left(\frac{x_2}{N_2}\right)+\lambda_2=0, \quad (21)$$

$$u'\left(\frac{x_{23}}{N_{23}}\right)+\lambda_2+\lambda_3=0. \quad (22)$$

Hence, the optimal bandwidth allocations per user – denoted y_{12}, y_2 and y_{23} – and link prices can be derived by the following equations:

$$N_{12}\cdot u'(y_{12})+\lambda_1+\lambda_2=0, \quad (23)$$

$$N_2\cdot u'(y_2)+\lambda_2=0, \quad (24)$$

$$N_{23}\cdot u'(y_{23})+\lambda_2+\lambda_3=0, \quad (25)$$

$$N_{12}\cdot y_{12}=C_1 \text{ or } \lambda_1=0, \quad (26)$$

$$N_{23}\cdot y_{23}=C_3 \text{ or } \lambda_3=0, \quad (27)$$

$$N_{12}\cdot y_{12}+N_2\cdot y_2+N_{23}\cdot y_{23}=C_2 \text{ or } \lambda_2=0. \quad (28)$$

It is reasonable to assume that the provider knows which links of his network are congested. That is, he knows at which links the capacity is to be exhausted and at which links the Lagrange multipliers equal zero. In our example, the middle link is the congested link of the network. Hence, we obtain:

$$p_1^*=p_3^*=0, \quad (29)$$

$$p_2^*=-\lambda_2=N_{12}\cdot u'(y_{12})=N_2\cdot u'(y_2)=N_{23}\cdot u'(y_{23}), \quad (30)$$

$$N_{12}\cdot y_{12}+N_2\cdot y_2+N_{23}\cdot y_{23}=C_2. \quad (31)$$

Thus, the optimal link prices can be derived from the above set of equations and the provider may fine-tune MIDAS as explained in Section 4.5.

Let us now consider the dumbbell network of Fig. 5. Assume that the network provider is aware of the fact that a large number N of customers demand paths that start from the left part of the network and end at the right part and that the capacity of all edge links is C except for the middle link whose capacity C_m is $C < C_m < N_c$. Therefore, it can be approximated that each user chooses randomly his path's origin and destination nodes from a uniform distribution. In this case, it is expected that on the average the middle link's demand will be N times that of the edge links. Therefore, to simplify the mechanism, the auctioneer may set a low fixed price for the bandwidth of the edge links and run the auction just for the middle link, which is the bottleneck. This example also depicts that though users adapt to the auction prices in the short auction time-scale, the network providers also adapt to the demand expressed in the auction in longer time-scale. Providers can make use of the *demand revelation* attained by means of the auction in order to dimension their network. Hence, in the aforementioned dumbbell network case, in case the provider upgrades the middle link's capacity so that it becomes $N\cdot C$, then this link would no longer be the bottleneck. In this case, congestion may occur at the access links, where MIDAS can be run with the initial link prices set accordingly. This case can be generalized for more complex hierarchical topologies, such as the one depicted in Fig. 6.

The aforementioned example also illustrates that though in this paper for the sake of the generality of the presentation, we propose MIDAS for auctioning the bandwidth of all the links of a network, MIDAS could be used in practice for allocating the bandwidth of the *popular* links of the network. In that case, the price to pay for bandwidth of

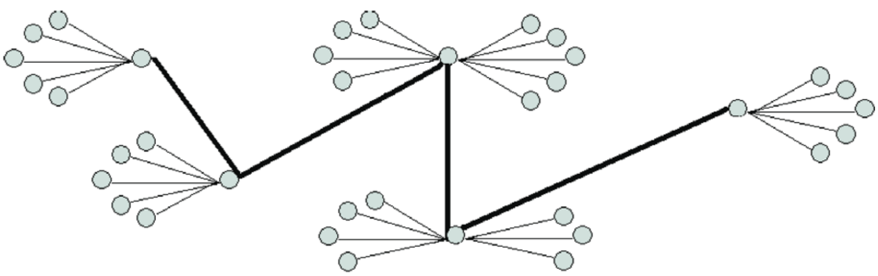

Fig. 6. A typical topology of an ISP's network.

non-congested links could be a predefined low price. In fact, this essentially happens under MIDAS as well, as a result of the price dropping policy adopted.

4.7. The payment rule

We now proceed to show how MIDAS bidders can be indeed provided the incentive to be truth-telling, by combining MIDAS with an incentive-compatible payment rule. The payment rule associated with MIDAS is also important, because it affects users' strategies, thus determining the auction outcome and whether it is successful w.r.t. efficiency (which is our primary objective) and provider's revenue. It is possible to attain social welfare that is close to the maximum by means of MIDAS, if the seller is risk averse and thus it is acceptable if a quantity of bandwidth is not sold. Indeed, the auction should involve uncertainty about the spare capacity of the links (and the corresponding termination point of the auction). Moreover, this feature should be combined with the use of reserve prices. Ausubel and Cramton discuss the issue of optimality and generalize the Vickrey auction to allow for reserve pricing in a multi-unit auction [5]. It is known that the seller by withholding quantity can improve revenues and mitigate collusion, while rendering truthful bidding the users' *dominant strategy*. The authors of [5] design a generalized Vickrey auction with reserve prices that maximizes both social welfare and seller's revenues.

Thus, the introduction of reserve prices possibly along with some uncertainty w.r.t. the auction termination and limited user feedback results in efficiency with MIDAS. The underlying idea thereof is to ''create'' losing bids and then apply the VCG (Vickrey–Clarke–Groves) payment rule, i.e. charge users with the social opportunity cost their presence entails. In particular, MIDAS consists of two successive phases:

Phase 1: Users submit their bids expressing their demand, receiving feedback regarding only the link prices until all prices become zero, thus revealing their demand.

Phase 2: The seller decides on the vector of prices (among the ones already offered) in which he decides that the auction should have ended. By doing so, he also determines the quantity of bandwidth to be sold for each link. The seller can then decide on the *optimal termination point*. Bidders are then informed about the quantity of bandwidth reserved and their corresponding charge derived by means of the VCG payment rule.

The criterion for deciding on the optimal termination point can be one of the following:

1. *Revenue maximization*, as in [5]. In particular, for every vector of prices offered, the auctioneer computes the revenue that would be attained if the auction were to terminate at that point. The vector of link prices that maximizes the revenue is selected as the optimal termination point.
2. *Social welfare maximization.* The price vector that maximizes the social welfare is the latest feasible vector of link prices: That is, the vector of link prices last offered for which all the bids submitted up to then can be satisfied with the capacities available.

By applying the aforementioned rules, the auctioneer runs the MIDAS auction until prices drop to 0. Subsequently, he applies one of the aforementioned criteria and decides on the auction's termination point. A comprehensive example is presented in Subsection 4.8. It is worth noting that winner determination is *immediate*, since winning bids are the bids submitted before the ''optimal termination point'', while path bidders' constraints are also met. Note that these properties would not apply if we had opted to compute an optimal termination quantity per link, since it would be impossible to satisfy all path bidders' quantity constraints. Indeed, in such a case it is possible that some path bidders are allocated different quantities in the various links of their respective paths. Also, the computation of the social opportunity cost can be done online, i.e. as the auction progresses, and on a per round and link basis.

If MIDAS is employed for a single link, then it is equivalent to the Generalized Vickrey auction of [5]. Due to the VCG payment rule, bidders gradually (by reacting to current link prices) submit their demand curve and the complete demand of the market is revealed to the seller. Therefore, if the first criterion is selected, both the provider's revenue and the social welfare *for the quantity allocated* are maximized. If the second criterion is applied, then the link's capacity is allocated efficiently.

However, in the case of a general network, there might be loss in efficiency, since the optimal prices may not be offered simultaneously at the auction. Hence, if the optimal vector of prices is not part of the auction's price path, some units of bandwidth will not be allocated to the users who value them the most (see Subsection 4.2). However, the closer the final price vector of the auction to the optimal price vector, the smaller the loss of efficiency is. Indeed, the experimental assessment of MIDAS of Subsection 4.5 indicates that this loss is in general limited because of the effectiveness of the price freezing policy.

Concluding, MIDAS provides a practical, fast, near-optimal solution to a generally NP-hard optimization problem. Also, the fact that MIDAS reveals limited information, is a desirable feature for the network providers, who want to reveal as little information as possible regarding their networks to their customers (and competitors). Also, another attractive feature of the mechanism is the fact that revenue sharing among multiple providers is feasible. Concluding, MIDAS is scalable, nearly efficient and is applicable for practical use (e.g. in bandwidth markets).

4.8. MIDAS execution example

A comprehensive example of a MIDAS auction execution is presented in Table 3. To keep the presentation simple, assume a linear network of three links; the capacity of all links is taken to be $C = 4$, while the parameters r, s are taken to be 1. Also, assuming that the provider has no information regarding demand, he sets the initial link prices to be all equal to 50. Also, the bidders participating are: (i) A path bidder interested (and in fact succeeding in the auction) in reserving two units of bandwidth at links 1 and 2 for \$144. (ii) A multitude of both path and single-link bidders interested in reserving one unit of bandwidth; their respective valuations are omitted for brevity reasons and are taken to be equal to the link prices offered in the auction when they bid. Therefore, the charge for the first bidder equals the sum of the two highest losing bids of links 1 and 2, while the rest of the win-

Table 3
An example of MIDAS execution for a linear network of three links

t	Link 1 (p_1, SC_1)	Link 2 (p_2, SC_2)	Link 3 (p_3, SC_3)	Brief explanation
0	(50, 4)	(50, 4)	(50, 4)	Auction begins
...	...	...	...	
14	(36, 4 → 2)	(36, 4 → 2)	(36, 4)	Two units reserved at links 1 and 2
15	(36, 2)	(36, 2)	(35, 4)	Price freezes at links 1 and 2
16	(36, 2)	(36, 2)	(34, 4)	Price still frozen at links 1and 2
17	(35, 2)	(35, 2)	(33, 4)	
18	(34, 2)	(34, 2 → 1)	(32, 4 → 3)	One unit reserved at links 2 and 3
19	(33, 2)	(34, 1)	(32, 3)	Price freezes at links 2 and 3
20	(32, 2)	(33, 1)	(31, 3 → 2)	One unit reserved at link 3
21	(31, 2)	(32, 1)	(31, 2)	Price freezes at link 3
22	(30, 2)	(31, 1)	(30, 2 → 1)	One unit reserved at link 3
23	(29, 2 → 1)	(30, 1 → 0)	(30, 1)	Price freezes at link 3 One unit reserved at links 1 and 2
24	(29, 1)	(30, 0)	(29, 1)	Price freezes at links 1 and 2
25	(28, 1)	(29, 0)	(28, 1)	
24	(27, 1)	(28, 0)	(27, 1)	
25	(26, 1 → 0)	(27, 0)	(26, 1)	One unit reserved at link 1
26	(26, 0)	(26, 0)	(25, 1)	
27	(25, 0)	(25, 0)	(24, 1)	
28	(24, 0)	(24, 0)	(23, 1 → 0)	One unit reserved at link 3
29	(23, 0)	(23, 0)	(23, 0)	Price freezes at link 3
...	...	...	...	
33	(19, 0 → −1)	(19, 0 → −1)	(19, 0)	One unit reserved at links 1 and 2
34	(19, −1)	(19, −1)	(18, 0)	Price freezes at links 1 and 2
...	...	...	...	
41	(12, −1 → −2)	(12, −1 → −2)	(11, 0)	One unit reserved at links 1 and 2
42	(12, −2)	(12, −2)	(10, 0)	Price freezes at links 1 and 2
43	(11, −2)	(11, −2)	(9, 0 → −1)	No need to continue the auction

ners are charged with the highest losing bid of their respective links.

Each row of Table 3 corresponds to an auction round, and comprises the link spare capacities $\langle SC_l(t)\rangle$, as well as the respective link prices $\langle p_l(t)\rangle$. For brevity reasons, we omit the rounds after $t = 43$, because they do not affect the auction outcome. Note that since users are unaware of the link spare capacities, they continue bidding even after the available bandwidth has been exhausted; in this example this occurs at $t = 33$. Finally, we remind the reader that as a result of the mechanism's incentive compatibility property, each user bids for an additional unit of bandwidth at a path, whenever the sum of auction prices of the links comprising this path is less than or equal to the user's respective marginal utility.

In case the auctioneer aims for revenue maximization, he must examine the revenue attained at every round of the auction in order to decide whether this round will be selected as the auction's termination point. It is worth noting that the sooner the auction terminates, the higher the highest losing bids that determine the charge of the winners for the units allocated. However, the quantity of the units allocated increases with the number of rounds. For instance, if the auction were terminated at $t = 14$, then there is a winner reserving two units of bandwidth at links 1 and 2. According to the payment rule, his charge is the sum of the two highest losing bids at links 1 and 2. Therefore, the attained revenue is $(29 + 26) + (34 + 30) = \$119$. Table 4 depicts the revenue attained for every possible termination round. Clearly, the optimal termination point is $t = 23$, corresponding to the vector or prices $\langle 29, 30, 30\rangle$, for which 3, 4 and 3 units of bandwidth are allocated at links 1, 2 and 3 respectively. The respective revenue is \$209. On the contrary, if the auctioneer opts for social welfare maximization, the optimal termination point is $t = 28$, corresponding to the vector or prices $\langle 24, 24, 23\rangle$. The respective revenue is \$174, which is considerably lower than the revenue attained under the first criterion. The reverse ordering applies for the values of the resulting social welfare, which are \$330 and \$379 respectively.

5. Comparison with related research

This section provides a comparison of MIDAS with other auction mechanisms that have been proposed for bandwidth allocation in communication networks. It is worth noting that we refrain from referring to auctions for a single either wired or wireless link (see [1] and references therein for an overview). This is mostly because single-link auctions are essentially multi-unit auctions; these have been studied thoroughly by economists [17]. On the contrary, as already explained in Section 1, the problem of bandwidth allocation of a network is far more complex and cannot be addressed by standard multi-unit auction mechanisms.

MacKie-Mason and Varian [19] propose to run at each router packet-level Vickrey auctions in short time-scales. Despite its simplicity, this mechanism cannot provide end-to-end guarantees for users' traffic. Also, users' willingness to pay can be unambiguously defined only for end-to-end flows, rather than local per-router packet forwarding. Therefore, the performance of the "smart market" approach of [19] with respect to efficiency is unclear.

Auctioning bandwidth over a network has been examined by Lazar and Semret in [18], where they present the Progressive Second Price (PSP) auction. In PSP, at every round, users submit their bids (i.e. pairs of quantity and price) and get as feedback the bandwidth allocations they would have got if this were the final round. Bidders gradually "raise" their

Table 4
MIDAS revenue per round for the example of Table 3

Termination	Attained revenue
$0 \leqslant t \leqslant 13$	0
$14 \leqslant t \leqslant 17$	$(29 + 26) + (34 + 30) = 119$
$18 \leqslant t \leqslant 19$	$(29 + 26) + (30 + 19 + 30) + 31 = 165$
$20 \leqslant t \leqslant 21$	$(29 + 26) + (30 + 19 + 30) + (30 + 30) = 194$
$t = 22$	$(29 + 26) + (30 + 19 + 30) + (23 + 23 + 23) = 203$
$23 \leqslant t \leqslant 24$	$(26 + 19 + 26) + (19 + 12 + 19 + 19) + (23 + 23 + 23) = 209$
$25 \leqslant t \leqslant 27$	$(19 + 12 + 19 + 19) + (19 + 12 + 19 + 19) + (23 + 23 + 23) = 207$
$28 \leqslant t \leqslant 32$	$(19 + 12 + 19 + 19) + (19 + 12 + 19 + 19) + (9 + 9 + 9 + 9) = 174$
$t \geqslant 33$	0 (There are not any feasible allocations)

bids (raising the price and reducing the quantity) by replying to rival bids until no bids are raised; the auction ends at this point and the quantity each user is allocated, is the maximum possible given his bid. Each user pays so as to exactly cover the social opportunity cost.

In our opinion, PSP has certain drawbacks. PSP requires a long, computationally expensive iterative process until convergence is achieved, that results in significant message overhead for the network. Moreover, the PSP auction, even for one link, has more than one equilibria, which can lead to an unfair result, while its "truthful equilibrium" (PSP's main merit) is unlikely to appear in practice. Indeed, it is proved in [20], that the first bidder has no incentive to give his true valuation of the bandwidth; it is more beneficial for him to declare a high willingness to pay for the link's capacity. By doing so, other bidders are prevented from bidding and the first bidder is awarded the entire link for a small price-the link's reserve price. Though a solution to this problem is proposed in [20] by means of enabling bidders to place truthful "sanction" losing bids, users have no incentive for submitting such truthful bids. In our opinion, the above facts justify revisiting the problem of bandwidth allocation on a network basis, looking for alternative solutions.

Indeed, in [21] the multi-bid auction is proposed. This is an elegant sealed-bid "discretization" of PSP for one link. This mechanism has also been extended for networks of *tree* topology, such as access networks [22]. Indeed, in [22], the authors run an instance of the multi-bid auction of [21] at every link of the tree, while a centralized algorithm ensures that the users will finally reserve the same quantity of bandwidth throughout their path. The auction's computational complexity per link is a quadratic function of the number of bidders at this link. The fact that a centralized algorithm is used, combined with the fact that the multi-bid auction's complexity depends on the number of bidders, limit its scalability. Most importantly, there is no generalization of the multi-bid auction for the general network case. On the contrary, MIDAS is both scalable and applicable to networks of arbitrary topology.

Also note that contrary to MIDAS, a network provider employing PSP, cannot utilize any information regarding the market demand, as already explained in Section 4.5. Last but most importantly, under our mechanism and assuming private values (as PSP does), it is *dominant strategy* for the users to be truth-telling. This is not the case with PSP where truthful bidding is just an equilibrium strategy. Indeed, it has been proved in [7] that other strategies than the one suggested by the authors of [18] may be more profitable for users in the network case.

Finally, it could also be argued that the VCG mechanism with combinatorial bids could be applied. This way path bidders would express accurately their demand and bandwidth should be allocated efficiently. This property is indeed attractive in theory. However, the winner determination is an NP-hard problem and so is the computation of the respective (VCG) charge for every winner. This renders this mechanism inapplicable in practice. A secondary problem of this mechanism is that *revenue sharing* among different providers is not feasible, as already explained earlier in the paper.

6. Auctioning bandwidth over a longer time

So far we have assumed that our auction mechanism is ran periodically and at each time the allocation of bandwidth applies for a certain period of time. That is, the bidders competing for bandwidth are considered "synchronized" in a certain period, which can be taken as a slot of unit duration without loss of generality. Below, we explain how this restriction can be relaxed. In particular, we extend our mechanism for a number k of time periods, where k is an integer greater than 1.

So far we had the objective that a user wishing to reserve bandwidth in two neighboring links (for the same time slot) should reserve the same quantity of bandwidth. Clearly, the same applies if a user is interested in using the same link for two consecutive slots. Thus, we can extend our auction mechanism by noting that two links may be "neighboring" either spatially (i.e., two links sharing a common node) or temporally (i.e., the same physical link used in two consecutive slots). Thus, the requirement for reserving the same quantity of bandwidth in two (or more) neighboring links applies for both temporal and spatial neighbors. The generalization of the mechanism in order to accommodate non-synchronized bidders is now straightforward. In order to auction the bandwidth of N links for T time slots $N \cdot T$ multi-unit Dutch auctions should be conducted. Bidders create paths over time the same way that they create paths in space. Thus, the paths become two-dimensional over the space and time axes, accommodating bidders that are not synchro-

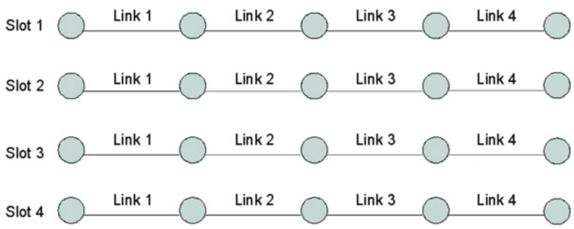

Fig. 7. Bidding w.r.t. spatial and temporal locality.

nized and each of which can reserve bandwidth for the time slots desired.

A related example, for $T = 4$ slots is depicted in Fig. 7; a bidder wishing to reserve bandwidth at links 2, 3, 4 for time slots 2 and 3 should bid in the respective 6 simultaneous auctions, also depicted as "gray links" in Fig. 7.

7. Conclusions

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 auctioning multicast trees and VPNs. The social welfare attained under our mechanism is in general close to the optimal one, while under full information the respective maximum is reached. This is achieved in polynomial time by means of both a successful reduction policy of the auction's link prices (Price Freezing policy) and an incentive-compatible payment rule. Furthermore, the mechanism has low computational complexity and is scalable for large numbers of links and users. MIDAS reveals limited information for provider's networks and serves as a fast, near-optimal solution to a generally NP-hard optimization problem.

References

[1] E. Altman, T. Boulogne, R. El-Azouzi, T. Jimenez, L. Wynter, A survey on networking games in telecommunications, Computers and Operations Research 33 (2) (2006) 286–311.

[2] J. Altmann, B. Rupp, P. Varaiya, Internet demand under different pricing schemes, in: Proceedings of ACM conference on Electronic Commerce (EC99), Denver, Colorado, USA, Nov 1999.

[3] Arbinet. <http://www.arbinet.com>.

[4] L. Ausubel, P. Cramton, Demand Reduction and Inefficiency in multi-unit auctions, Working paper, No 96-07, Univ. of Maryland, 1996, <http://www.ausubel.com/larry/auction-papers.htm>.

[5] L. Ausubel, P. Cramton, Vickrey auctions with reserve pricing, Economic Theory 23 (April) (2004) 493–505.

[6] Bandwidth Market Ltd. site, Bandwidth Markets and Brokering: Wave of the Future, December 2004. <http://www.bandwidthmarket.com/>.

[7] M. Bitsaki, G.D. Stamoulis, C. Courcoubetis, A new strategy for bidding in the network-wide progressive second price auction for bandwidth, in: Proceedings of CoNEXT'05, 2005, pp. 146–155.

[8] Cisco VPN IPsec architecture. <http://www.cisco.com/>.

[9] C. Courcoubetis, M. Dramitinos, G.D. Stamoulis, An auction mechanism for bandwidth allocation over paths, in: Proceedings of Seventeenth International Teletraffic Congress (ITC), Teletraffic Engineering in the Internet Era, Brazil, 2001, pp. 1163–1174.

[10] C. Courcoubetis, R. Weber, Pricing Communication Networks: Economics, Technology and Modelling, John Wiley and Sons, 2003, ISBN 0-470-85130-9.

[11] R. Edell, P. Varaiya, Providing internet access: what we learn from INDEX, IEEE Network 13 (1999) 18–25.

[12] P. Gevros, J. Crowcroft, P. Kirstein, S. Bhatti, Congestion control mechanisms and the best effort service model, IEEE Network 15 (2001) 6–26.

[13] Internet Draft, Network based IP VPN Architecture Using Virtual Routers. <http://www.ietf.org/internet-drafts/draft-ietf-l3vpn-vpn-vr-03.txt>.

[14] Internet news. <http://www.internetnews.com/bus-news/article.php/376491>.

[15] Invisible hand report. <http://web.invisiblehand.net/ihn-bmr/>.

[16] C. Kenyon, G. Cheliotis, Network Market Dynamics, Telecoms Capacity 1 (2001/2) 38–39. <www.zurich.ibm.com/pdf/bandwidth/network_market_dynamics.pdf>.

[17] V. Krishna, Auction Theory, Academic Press, 2002, ISBN 0-12426297-X.

[18] A. Lazar, N. Semret, Design and analysis of the progressive second price auction for network bandwidth sharing, Telecommunications Systems (1999), Special Issue on Network Economics.

[19] J.K. MacKie-Mason, H.R. Varian, Pricing the internet, in: B. Kahin, J. Keller (Eds.), Public Access to the Internet, MIT Press, 1995, pp. 269–314.

[20] P. Maillé, B. Tuffin, A progressive second price mechanism with a sanction bid from excluded players, in: Proceedings of ICQT03, LNCS 2816, Munich, 2003, pp. 332–341.

[21] P. Maillé, B. Tuffin, Multi-bid auctions for bandwidth allocation in communication networks, in: Proceedings of IEEE Infocom 2004, Hong Kong.

[22] P. Maillé, B. Tuffin, Pricing the internet with multi-bid auctions, IEEE/ACM Transactions on Networking 14 (5) (2006) 992–1104.

[23] Min-X. <http://www.min-X.com>.

[24] A. Odlyzko, The current state and likely evolution of the Internet, in: Proceedings of Globecom '99, 1999, pp. 1869–1875.

[25] I. Ch. Paschalidis, Y. Liu, Pricing in multiservice loss networks: static pricing, asymptotic optimality, and demand substitution effects, IEEE/ACM Transactions on Networking 10 (3) (2002) 425–438.

[26] <http://dir.yahoo.com/Business_and_Economy/Business_to_Business/Communications_and_Networking/Telecommunications/Bandwidth_Auctions_ and_Exchanges/>.

Manos Dramitinos received his B.Sc. and Masters Degree in Computer Science from the University of Crete in 1998 and 2000 respectively. He received his Ph.D from the Informatics Department of Athens University of Economics and Business (AUEB) in 2006. He specializes in game theory and auction mechanism design. He has designed innovative mechanisms for the dynamic allocation of computer network resources. He received several awards for being admitted first in his class in 1994 and the Ericsson Award of Excellence in Telecommunications in 1998 for his diploma thesis. He was also awarded a scholarship from the Greek National Scholarship Foundation (IKY) for completing his postgraduate studies first among his class in 2000. He also received the best student paper award of the ICQT'06 conference. His research interests include auctions, mechanism design, network services and economics. He can be reached at mdramit@aueb.gr. Homepage: http://nes.aueb.gr/~mdramit/.

George D. Stamoulis received the Diploma in Electrical Engineering (1987, with highest honours) from the National Technical University of Athens, Greece, and the M.S. (1988) and Ph.D. (1991) degrees in Electrical Engineering from the Massachusetts Institute of Technology, Cambridge, Massachusetts, USA. From 1991 to 1993 he was a Research Associate with the Communication Networks group at the NTUA, participating in RACE projects, while serving in the Hellenic Navy, as a lecturer in the Hellenic Naval Academy. From 1993 to 1995 he was with the Development Programmes Department of INTRACOM, as a coordinator of RACE projects. From 1995 to 2000 he was an Assistant Professor at the Computer Science Department of the University of Crete, Greece, and a member of the Telecommunications and Networks Division of the Institute of Computer Science, Foundation for Research and Technology Hellas (ICS-FORTH). Since 2001 he is a faculty member, currently Professor, in the Department of Informatics of Athens University of Economics and Business (AUEB). His research interests are in network economics, auction mechanisms for bandwidth, charging for telecommunications services, reputation mechanisms for electronic environments, and telecommunications regulation. He has published about 55 papers in scientific journals such as IEEE/ACM Transactions on Networks, Telecommunications Systems, Computer Communications, IEEE Transactions on Communications, Journal of the ACM, and IEEE Transactions on Information Theory, as well as in conferences such as INFOCOM, ITC, ACM SIGMETRICS, GLOBECOM and GP2PC. He has also collaborated with the Greek Regulatory Authorities for Communications and Power on auction design as well as with Greek telecom companies on service pricing.

Costas Courcoubetis was born in Athens, Greece. He received his Diploma (1977) from the National Technical University of Athens, Greece, in Electrical and Mechanical Engineering, his M.S. (1980) and Ph.D. (1982) from the University of California, Berkeley, in Electrical Engineering and Computer Science. From 1982 until 1990 he was Member of the Technical Staff (MTS) in the Mathematical Sciences Research Center, Bell Laboratories, Murray Hill, NJ, and from 1990 until 1999 he was Professor in the Computer Science Department at the University of Crete in Heraklion, Greece, and headed the Telecommunications and Networks Group at the Institute of Computer Science, FORTH. Since Fall 1999 he is a Professor in the Computer Science Department at the Athens University of Economics and Business.

His current research interests are economics of networks with emphasis in the development of pricing schemes that reduce congestion and enhance stability and robustness, quality of service and management of integrated services, performance and traffic analysis of large systems, applied probability models. Other interests include the combination of e-commerce technologies with telecommunications, and formal methods for software verification.

He has published over 80 papers in scientific journals such as Operations Research, Mathematics of Operations Research, Journal on Applied Probability, IEEE Transactions in Communications, IEEE Transactions in Automatic Control, IEEE JSAC, SIAM Journal on Computing, Stochastic Processes, Probability in Engineering and Information Sciences, Queuing Systems, ACM TOPLAS, JACM, Formal Methods in System Design, Telecommunications Systems, Information and Computation, Theoretical Computer Science, and in conferences such as FOCS, STOC, LICS, INFOCOM. GLOBCOM, ITC, ACM SIGMETRICS, etc. His work has over 430 citations in scientific journals, and over 1200 citations in Internet published documents according to the NECI Scientific Literature Digital Library.