

Design and Analysis of the Progressive Second Price Auction for Network Bandwidth Sharing*

Aurel A. Lazar

Dept. of Electrical Engineering, Columbia University

New York, NY, 10027-6699, USA

`aurel@comet.columbia.edu`

and

Nemo Semret[†]

Dept. of Electrical Engineering, University of California, Los Angeles

Los Angeles, CA, 90095-1594, USA

`nemo@ee.ucla.edu`

April 1, 1998

Revised: September 30, 1999

Abstract

We present the Progressive Second Price auction, a new decentralized mechanism for allocation of variable-size shares of a resource among multiple users. Unlike most mechanisms in the economics literature, PSP is designed with a very small message space, making it suitable for real-time market pricing of communication bandwidth. Under elastic demand, the PSP auction is incentive compatible and stable, in that it has a “truthful” ϵ -Nash equilibrium where all players bid at prices equal to their marginal valuation of the resource. PSP is economically efficient in that the equilibrium allocation maximizes total user value. With simulations using a prototype implementation of the auction game on the Internet, we investigate how convergence times scale with the number of bidders, as well as the trade-off between engineering and economic efficiency. We also provide a rate-distortion

*Parts of this work were presented at the 8th *International Symposium on Dynamic Games and Applications*, Maastricht, The Netherlands, July 1998, and at the *DIMACS Workshop on Economics, Game Theory, and the Internet*, Rutgers, NJ, April 1997.

[†]Corresponding author.

theoretic basis for valuation of bandwidth, which leads naturally to the elastic demand model that is assumed in the analysis of the mechanism.

Keywords: resource allocation, auctions, game theory, mechanism design, network pricing.

1 Introduction

Communication networks are characterized by what economists call externalities. The value a user gets from the network depends on the other users. The positive externalities are that a communication network is more valuable if more people are connected. The negative externalities are that resources are shared by users who – because of distance, population size, or selfishness – cannot or will not coordinate their actions sufficiently to achieve the most desirable allocation of resources. The recognition of this reality in many aspects of networks and distributed computations has led in recent years to the emergence of game theoretic approaches in their analysis and design [23, 9, 25, 33, 15, 16].

Prices, whether they relate to “real money” in a public network or “funny money” (based on quotas) in a private system, play a key role as allocation control signals. In the former case, this role is of course intimately tied to another, which is to allow a network provider to remain in business [7].

The telephone system and the current Internet represent two extremes of the relationship between resource allocation and pricing. The resources allocated to a telephone call are fixed, and usage prices are based on the predictability of the total demand at any given time. On the Internet, the current practice of pricing by the maximum capacity of the user’s connection (flat-rate pricing) decouples the allocation (actual use) of resources from the prices.

In the emerging multiservice networks (ATM, Next-Generation Internet), neither of these approaches are viable. The former because of the wide and rapidly evolving range of applications (including some which adapt to resource availability) will make demand more difficult to predict. And the latter because, once the flat fee is paid, there are no incentives to limit usage since increasing consumption benefits the user individually, whereas limiting it to sustainable levels brings benefits which are shared by all. This makes it vulnerable to the well-known “tragedy of the commons”. With flat pricing alone, the tendency is toward increasing congestion which chases away high-value users, or increasing prices which exclude low-value users [6], in both cases leading to decreased network revenue.

Thus there is a need to develop new approaches to pricing of network resources. Among the requirements are: sensitivity to the range of resource requirements (either through a sufficiently rich range of traffic classes which are priced differently, or by allowing users to explicitly quantify resource requirements); prices must be dynamically responsive to unpredictable demand (market based system); perhaps most importantly the pricing architecture should constrain as little as possible the efficiency trade-offs of the policies.

Indeed, the fundamental issue in designing pricing policies is the trade-off between engineering efficiency and economic efficiency. This trade-off, which is more or less constrained by the underlying network technology, has many dimensions, including:

- how much measurement (from usage to capacity pricing),
- the granularity of differently priced service offerings (e.g. number of traffic classes),
- the level of resource aggregation – both in time and in space – at which pricing is done (per packet/cell or per connection, at the edge of the network or at each hop), and
- the information requirement (how much *a priori* knowledge of user behavior and preferences is required/assumed by the network in computing prices).

An approach which achieves economic efficiency is the smart-market approach of [19], wherein each packet contains a bid, and if it is served, pays a clearing price given by the highest bid among packets which are denied service (dropped). This approach is incentive compatible in that the optimal strategy for a (selfish) user is to set the bid price in each packet equal to the true valuation. Each node in the network becomes an efficient market, but the engineering cost (sorting packets by bid price, as well as per-packet and per-hop accounting) could be significant if line speeds are high relative to the processing power in the router. In [14], users are charged according to a combination of declared and measured characteristics of traffic. By taking an equivalent bandwidth model of resource utilization, and assuming appropriate traffic models, a menu of pricing plans indexed by the declared traffic can be offered which encourages users to make truthful declarations (e.g. of the mean rate), and also encourages the users' characterization efforts to be directed where they are most relevant to the network resource

allocation. As the pricing is relative, [14] does not aim to address the problem of determining the actual monetary values of the market price (that users would be willing to pay). Another pricing scheme which incorporates multiplexing gain is formulated in [12]. These and a number of other schemes are summarized in [11], in a comprehensive view of the connection establishment process, which identifies the user-network negotiation as the key “missing link” in network engineering/economic research. In terms of our taxonomy of the previous paragraph, this is part of the information requirement trade-off. Indeed, in the absence of formal mechanisms to deal with the information problem, complex and (at least intuitively) undesirable things happen. For example, some providers offer expensive “front of the book” rates to uninformed customers, and lower “back of the book” rates to informed customers who may be about to defect to another carrier (see [7] and also the recent wars between AT&T and MCI in consumer long-distance service in the United States). In [34], it is argued that architectural considerations such as where charges are assessed should take precedence over the pursuit of optimal efficiency, and edge pricing (spatial aggregation in terms of our taxonomy) is proposed as a useful paradigm.

In this paper, we propose a new auction mechanism which accommodates various dimensions of the engineering-economics trade-off. The mechanism applies to a generic arbitrarily divisible and additive resource model (which may be equivalent bandwidth, peak rate, contract regions, etc., at any level of aggregation.) It does not assume any specific mapping of resource allocation to quality of service. Rather, users are defined as having an explicit monetary valuation of quantities of resource, which the network doesn’t or can’t know a priori. Thus, in terms of our trade-off taxonomy, this mechanism aims for unlimited granularity, flexibility in the level of aggregation and minimal information requirement.

In the most likely auction scenario, users would be aggregates of many flows data flows for which bulk capacity is being purchased for e.g. Virtual Paths, Virtual Private Networks, or edge capacity [2, 31].

We begin in Section 2 by formally presenting the design of our Progressive Second Price auction mechanism for sharing a single arbitrarily divisible resource, and relating it to classical mechanism design from the economics literature. In Section 3, after describing our model of user preferences and the elastic demand assumption, we prove that PSP has the desired properties of incentive compatibility, stability, and efficiency. The section concludes with simulation results on the convergence properties, and the efficiency trade-offs. Appendix A describes an information theoretic basis

for valuations of the type that are assumed in the analysis of Sections 3, as one possible justification.

2 Design of an Auction for a Divisible Resource

2.1 Message Process

Following [36], it is useful to expose the design in terms of its two aspects: realization, where a message process that enables a certain allocation objective is defined; and Nash implementation, where allocation rules are designed with incentives which drive the players to an equilibrium where the (designer's) desired allocation is achieved.

In this section we define the message process. Here we make the fundamental choice which will constrain the subsequent aspects of the design. Our first concern here is with engineering. For the sake of scalability in a network setting, we shall aim for a process where a) the exchanged messages are as small as possible, while still conveying enough information to allow resource allocation and pricing to be performed without any a-priori knowledge of demand (market research, etc.); and b) the amount of computation at the center is minimized.

Given a quantity Q of a resource, and a set of players $\mathcal{I} = \{1, \dots, I\}$, an auction is a mechanism consisting of: 1) players submitting bids, i.e. declaring their desired share of the total resource and a price they are willing to pay for it, and 2) the auctioneer allocating shares of the resource to the players based on their bids.

Player i 's bid is $s_i = (q_i, p_i) \in \mathcal{S}_i = [0, Q] \times [0, \infty)$, meaning he would like a quantity q_i at a *unit* price p_i . A bid profile is $s = (s_1, \dots, s_I)$. Following standard game theoretic notation, let $s_{-i} \equiv (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_I)$, i.e. the bid profile of player i 's opponents, obtained from s by deleting s_i . When we wish to emphasize a dependence on a particular player's bid s_i , we will write the profile s as $(s_i; s_{-i})$.

The allocation is done by an **allocation rule** A ,

$$A : \quad \mathcal{S} \longrightarrow \mathcal{S} \\ s = (q, p) \longmapsto A(s) = (a(s), c(s)),$$

where $\mathcal{S} = \prod_{i \in \mathcal{I}} \mathcal{S}_i$.

The i -th row of $A(s)$, $A_i(s) = (a_i(s), c_i(s))$, is the allocation to player i : she gets a quantity $a_i(s)$ for which she is charged $c_i(s)$. Note that p is a price per unit and c is a total cost.

An allocation rule A is **feasible** if $\forall s$,

$$\sum_{i \in \mathcal{I}} a_i(s) \leq Q$$

and $\forall i \in \mathcal{I}$,

$$\begin{aligned} a_i(s) &\leq q_i, \\ c_i(s) &\leq p_i q_i. \end{aligned}$$

Remark a: The above formulation is a generalization of what is usually meant by an auction. The latter is the special case where $a_w(s) = Q$ for some winner $w \in \mathcal{I}$ and $a_i(s) = 0, \forall i \neq w$, i.e. the sale of a single indivisible object to one buyer, for which the theory is well developed [22, 24]. In our approach, allocations are for arbitrary shares of the total available quantity of resource. Equivalently, one could slice the resource into many small units, each of which is auctioned as an indivisible object. But in a practical implementation of auctions for sharing a resource, a process of bidding for each individual unit would result in a tremendous signaling overhead. More importantly, since the users would be bidding on a discrete grid of quantities, analytical predictions of outcomes could be misleading since they could be sensitive to the particular choice of grid¹.

Remark b: Most of the mechanism design literature in Economics makes use of the following ‘‘Revelation Principle’’:

Given any feasible auction mechanism, there exists an equivalent² feasible direct revelation mechanism which gives to the seller and all bidders the same expected utilities as the given mechanism. ([24], Lemma 1)

In this sharing context, a direct revelation mechanism would be one where each user message consists of the user’s type, which is the valuation³ of the resource over the whole range of their possible demands, i.e. a function $\theta_i : [0, Q] \rightarrow [0, \infty)$, and the budget (see Section 3.1). A consequence of revelation principle is that the mechanism designer can restrict her attention to direct revelation mechanisms, find the best mechanism in terms of

¹For a more detailed discussion of this point, see [4] p. 34, and references therein.

²By equivalent, in [24] it is meant that, at some equilibrium, all players get the same utility. There may be other, possibly ill-behaved, equilibria.

³The valuation of a given amount of resource is how much the user is willing to pay for that quantity. The inverse of the valuation is the user’s demand function, giving a desired quantity for each price.

her (economic) efficiency objectives, and then – if necessary – transform it into an equivalent mechanism in the desired message space. This is convenient because one can exclude the infinitely many mechanisms with larger message spaces, without fear of missing any better designs. The design process is usually the solution of an optimization (mathematical programming) problem. For this reason, in the literature, mechanism design problems are mostly solved for cases where the space of users’ types is one dimensional, or at most finite dimensional [21], using message spaces that are of the same dimension.

In our sharing problem, the conventional approach is unsatisfactory in two ways:

- First, a user’s type is infinite-dimensional, as we do not restrict the valuation functions beyond some very general assumptions (see Section 3.1), and so the conventional “programming” approach of deriving the mechanism from the revelation principle would lead to an intractable problem.
- Second, the conventional (direct revelation) approach, even if it was tractable, implies that a single message (bid) can theoretically be infinitely long, because it has to contain a description of the function θ_i . Clearly, this is not desirable in a communication network, where signaling load is a key consideration. For engineering reasons, we choose a message space that is 2-dimensional. Therefore, a given message can come from many possible types, so there is no single way to do the transformation from the direct revelation mechanism to the desired one.

Thus, unlike most of the mechanism design literature, we will take a direct approach, where we posit an allocation rule for our desired message space, and then show that it has an equilibrium, and that the design objective is met at equilibrium⁴. This is equivalent to guessing the right direct-revelation-to-desired-mechanism transformation and building it into the allocation rule from the start.

⁴Our aim is to show that we *can* use this smaller message space and still achieve our objective.




Figure 1: Exclusion-compensation principle: the intuition behind the PSP rule

2.2 Allocation Rule

Define, for $y \geq 0$

$$\underline{Q}_i(y, s_{-i}) = \left[Q - \sum_{p_k \geq y, k \neq i} q_k \right]^+ . \quad (1)$$

and

$$Q_i(y, s_{-i}) = \lim_{\eta \searrow y} \underline{Q}_i(\eta, s_{-i}) = \left[Q - \sum_{p_k > y, k \neq i} q_k \right]^+ .$$

The “progressive second price” (PSP) allocation rule is defined as follows:

$$a_i(s) = q_i \wedge \underline{Q}_i(p_i, s_{-i}), \quad (2)$$

$$c_i(s) = \sum_{j \neq i} p_j [a_j(0; s_{-i}) - a_j(s_i; s_{-i})], \quad (3)$$

where \wedge means taking the minimum.

Remark a: For a fixed opponent profile s_{-i} , $Q_i(p_i, s_{-i})$ represents the maximum available quantity at a bid price of p_i . **The intuition behind PSP** is an exclusion-compensation principle: player i pays for his allocation

so as to exactly cover the “social opportunity cost” which is given by the declared willingness to pay (bids) of the users who are excluded by i ’s presence (see Figure 1), and thus also compensates the seller for the maximum lost potential revenue. Note that this amounts to implicitly assuming that the bid price accurately reflects the marginal valuation θ'_i on the range $[a_i, q_i]$. In other words, by this rule the auctioneer is saying to the player: “if you bid (q_i, p_i) , I take it to mean that in the vicinity of q_i , θ_i can be approximated by a line of slope p_i .” This is the (built-in) transformation from the direct-revelation mechanism to the desired message process discussed in the second remark at the end of Section 2.1.

The charge c_i increases with a_i in a manner similar to the income tax in a progressive tax system. For a fixed opponent profile s_{-i} , imagine player i is increasing q_i , starting from 0. The first few units that player i gets will be taken away from the lowest clearing opponent (i.e. $m = \arg \min_j \{p_j : a_j > 0\}$), and player i will pay a price (marginal cost) p_m per unit. When a_m reaches 0, the subsequent units that player i gets will cost him $p_{m'} > p_m$, where m' is the new lowest clearing player, the one just above m . The PSP rule is the natural generalization of second-price auctions (or Vickrey auctions). In a Vickrey auction of a *single non-divisible object*, each player submits a sealed bid, and the object is sold to the highest bidder at the bid price of the second highest bidder, which is what happens here if $q_i = Q, \forall i$. This is widely known to have many desirable properties [35, 24, 4], the most important of which is that it has an equilibrium profile where all players bid their true valuation. As we will presently show, this property is preserved by the PSP rule in the more general case of sharing an arbitrarily divisible resource, and this leads to stability (Nash equilibrium). The PSP rule is analogous to Clarke-Groves mechanisms [3, 8, 20] in the direct-revelation case.

Remark b: When two players bid at exactly the same price, and they are asking for more than is available at that price, (2) punishes both of them. For example, if $Q = 100$ and $s_1 = (4, 60)$ and $s_2 = (4, 70)$, the allocations would be $a_1 = 60 \wedge (100 - 70) = 30$, and $a_2 = 70 \wedge (100 - 60) = 40$. Since the bid prices are equal, there is no “right” way to decide who to give the remaining capacity to. One could divide it equally, or proportionally to their requests, etc. For the subsequent analysis, it turns out it is simpler to not give it to either one (of course, it will be allocated to the lower bidders if there are any). This is just a technicality since by deciding this, we ensure that it will never happen (at equilibrium), since the users will always prefer to change their prices and/or reduce their quantity.

Considering the computational complexity of PSP, a straightforward implementation would at worst, sort the bids in time $I \log I$, perform (2) in linear time, and (3) can be done in time I^2 . Thus, the complexity of computing the allocations is $O(I^2)$.

3 Analysis of the Progressive Second Price Auction

3.1 User Preferences

Since the allocation rule A is given by design, the only analytical assumptions we make is on the form of the players' preferences.

Player i 's preferences are given by his utility function

$$\begin{aligned} u_i : \mathcal{S} &\longrightarrow (-\infty, \infty) \\ s &\longmapsto u_i(s). \end{aligned}$$

Player i has a **valuation** of the resource $\theta_i(a_i(s)) \geq 0$, which is the total value to her of her allocation. Thus, for a bid profile of s , under allocation rule A , player i getting an allocation $A_i(s)$ has the quasi-linear utility

$$u_i(s) = \theta_i(a_i(s)) - c_i(s) \tag{4}$$

which is simply the value of what she gets minus the cost.

In addition, the player can be constrained by a **budget** $b_i \in [0, \infty]$, so the bid s_i must lie in the set

$$S_i(s_{-i}) = \{s_i \in \mathcal{S}_i : c_i(s_i; s_{-i}) \leq b_i\}. \tag{5}$$

In the proofs of the following section, we will assume that users have elastic demand, that is:

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

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

The last item says that as long as the valuation is strictly increasing, it must also be strictly concave (with minimum curvature γ_i). However, it is allowed to “flatten” beyond a certain amount of resource.

Functions of this (concave) form have wide applicability as models of resource valuation, and can be justified from the economic standpoint (diminishing returns) as well as from information theoretic standpoint – see Appendix A. For examples of valuations satisfying Assumption 1, see Section 3.4 and Appendix A.

3.2 Equilibrium of PSP

The auction game is given by (Q, u_1, \dots, u_I, A) , that is, by specifying the resource, the players, and a feasible allocation rule. We analyze it as a strategic game of complete information [4].

Define the set of best replies to a profile s_{-i} of opponents bids: $S_i^*(s_{-i}) = \{s_i \in S_i(s_{-i}) : u_i(s_i; s_{-i}) \geq u_i(s'_i; s_{-i}), \forall s'_i \in S_i(s_{-i})\}$. Let $S^*(s) = \prod_i S_i^*(s_{-i})$. A Nash equilibrium is a fixed point of the point-to-set mapping S^* , i.e. a profile $s \in S^*(s)$. In other words, it is a point from which no player will want to unilaterally deviate. Such a point is what is most accepted as a consistent prediction of the actual outcome of a game, and has been repeatedly confirmed by experiments, as well as a wide range of theoretical approaches. Indeed, in a dynamic game, where players recompute the best response to the current strategy profile of their opponents, this iteration can only converge to a Nash equilibrium (if it converges at all). In addition, an important trend in modern game theory is the development of learning models, and there too, it has been shown that Nash equilibria result also from rational learning through repeated play among the same players [13].

A more general (and hence weaker) notion of stability is the existence of an ϵ -Nash equilibrium. Let the ϵ -best replies be $S_i^\epsilon(s_{-i}) = \{s_i \in S_i(s_{-i}) : u_i(s_i; s_{-i}) \geq u_i(s'_i; s_{-i}) - \epsilon, \forall s'_i \in S_i(s_{-i})\}$. An ϵ -Nash equilibrium is a fixed point of S^ϵ .

In a dynamic auction game, $\epsilon > 0$ can be interpreted as a *bid fee* paid by a bidder each time they submit a bid. Thus, the user will send a best reply bid as long as it improves her current utility by ϵ , and the game can only end at an ϵ -Nash equilibrium.

Define

$$P_i(z, s_{-i}) = \inf \{y \geq 0 : Q_i(y, s_{-i}) \geq z\}. \quad (6)$$

Thus, for fixed s_{-i} , $\forall y, z \geq 0$,

$$z \leq Q_i(y, s_{-i}) \Rightarrow y \geq P_i(z, s_{-i}) \quad (7)$$

and⁵

$$y > P_i(z, s_{-i}) \Rightarrow z \leq Q_i(y, s_{-i}). \quad (8)$$

The graph of $P_i(\cdot, s_{-i})$ is the “staircase” shown in Figure 1, and that of $Q_i(\cdot, s_{-i})$ is obtained by flipping it 90 degrees.

It is readily apparent that

$$c_i(s) = \int_0^{a_i(s)} P_i(z, s_{-i}) dz. \quad (9)$$

The key property of PSP is that, for a given opponent profile, a player cannot do much better than simply tell the truth, which in this setting means bidding at a price equal to the marginal valuation, i.e. set $p_i = \theta'_i(q_i)$. By doing so, she can always get within $\epsilon > 0$ of the best utility.

Let $\mathcal{T}_i = \{s_i \in \mathcal{S}_i : p_i = \theta'_i(q_i)\}$, the (unconstrained) set of player i 's truthful bids, and $\mathcal{T} = \prod_i \mathcal{T}_i$.

Proposition 1 (Incentive compatibility) *Under Assumption 1, $\forall i \in \mathcal{I}$, $\forall s_{-i} \in \mathcal{S}_{-i}$, such that $Q_i(0, s_{-i}) = 0$, for any $\epsilon > 0$, there exists a truthful ϵ -best reply $t_i(s_{-i}) \in \mathcal{T}_i \cap S_i^\epsilon(s_{-i})$.*

In particular, let

$$G_i(s_{-i}) = \left\{ z \in [0, Q] : z \leq Q_i(\theta'_i(z), s_{-i}) \text{ and } \int_0^z P_i(\eta, s_{-i}) d\eta \leq b_i \right\}.$$

Then with $v_i = [\sup G_i(s_{-i}) - \epsilon/\theta'_i(0)]^+$ and $w_i = \theta'_i(v_i)$, $t_i = (v_i, w_i) \in \mathcal{T}_i \cap S_i^\epsilon(s_{-i})$.

The truthful best reply can be found in a straightforward manner, as illustrated in Figure 2.

Proof: Fix $s_{-i} \in \mathcal{S}_{-i}$. Let $z_i = \sup G_i(s_{-i})$ and $y_i = \theta'_i(z_i)$.

By definition of z_i , $\exists \{z(n)\} \subset G_i(s_{-i})$ such that $\lim_n z(n) = z_i$. Hence $b_i \geq \lim_n \int_0^{z(n)} P_i(\eta, s_{-i}) d\eta = \int_0^{z_i} P_i(\eta, s_{-i}) d\eta \geq c_i(t_i; s_{-i})$, where the equality comes from the boundedness of P_i and the Lebesgue dominated convergence theorem, and the second inequality from (9) and (2). Thus $t_i \in \mathcal{T} \cap S_i(s_{-i})$.

⁵Actually, since $Q_i(\cdot, s_{-i})$ is upper-semi-continuous (jumps up), we have $z \leq Q_i(y, s_{-i}) \Leftrightarrow y \geq P_i(z, s_{-i})$




Figure 2: Truthful ϵ -best reply

Next we show that $t_i \in S_i^\epsilon(s_{-i})$. First, $z_i = \lim_n z(n) \leq \lim_n Q_i(\theta'_i(z(n)), s_{-i}) \leq Q_i(\lim_n \theta'_i(z(n)), s_{-i})$, where the inequalities follow respectively from $z(n) \in G_i(s_{-i})$, and the upper semi-continuity of $Q_i(\cdot, s_{-i})$. Now by the continuity of θ'_i , $Q_i(\lim_n \theta'_i(z(n)), s_{-i}) = Q_i(\theta'_i(z_i), s_{-i}) = Q_i(y_i, s_{-i})$, hence

$$z_i \leq Q_i(y_i, s_{-i}). \quad (10)$$

Now, we claim that $a_i(t_i; s_{-i}) = v_i$. Indeed, if $z_i = 0$ then $v_i = 0$ and $a_i(t_i; s_{-i}) = 0$. If $z_i > 0$, then by (10), $Q_i(y_i, s_{-i}) > 0$ and since by hypothesis $Q_i(0, s_{-i}) = 0$, we have $\theta'_i(z_i) = y_i > 0$. Also, $z_i > 0$ implies $v_i < z_i$. Therefore, by Assumption 1, we have $w_i = \theta'_i(v_i) > \theta'_i(z_i) = y_i$. Hence, since $\underline{Q}_i(\cdot, s_{-i})$ is non-decreasing, $\underline{Q}_i(w_i, s_{-i}) \geq \lim_{\eta \searrow y} \underline{Q}_i(\eta, s_{-i}) = Q_i(y_i, s_{-i}) \geq z_i > v_i$. Thus, by (2),

$$a_i(t_i; s_{-i}) = v_i \quad (11)$$

Now $\forall s_i \in S_i(s_{-i})$,

$$\begin{aligned}
& u_i(t_i; s_{-i}) - u_i(s) \\
&= \theta_i(a_i(t_i; s_{-i})) - \theta_i(a_i(s)) - c_i(t_i; s_{-i}) + c_i(s) \\
&= \int_{a_i(t_i; s_{-i})}^{a_i(s)} [P_i(z, s_{-i}) - \theta'_i(z)] dz \\
&= \int_{z_i}^{a_i(s)} [P_i(z, s_{-i}) - \theta'_i(z)] dz + \int_{v_i}^{z_i} [P_i(z, s_{-i}) - \theta'_i(z)] dz
\end{aligned}$$

$$\geq \int_{z_i}^{a_i(s)} [P_i(z, s_{-i}) - \theta'_i(z)] dz - \epsilon \quad (12)$$

where the inequality follows from $(z_i - v_i) \leq \epsilon/\theta'_i(0)$ and the fact that θ'_i is non-increasing. Thus, it suffices to show that the integral is ≥ 0 .

If $z_i < a_i(s)$, take any $z \in (z_i, a_i(s)]$. By the definition of z_i , $z \notin G_i(s_{-i})$. Now $s_i \in S_i(s_{-i})$ implies $b_i \geq c_i(s) = \int_0^{a_i(s)} P_i(\eta, s_{-i}) d\eta \geq \int_0^z P_i(\eta, s_{-i}) d\eta$. Therefore, we must have $z > Q_i(\theta'_i(z))$, which by (8), implies $\theta'_i(z) \leq P_i(z)$ and the integrand in (12) is ≥ 0 as desired.

Suppose $z_i \geq a_i(s_i)$. Since θ'_i is non-increasing, $Q_i(\cdot, s_{-i})$ is non-decreasing and $P_i(\cdot, s_{-i}) \geq 0$, any point to the left of z_i is in the set $G_i(s_{-i})$, $\forall z < z_i, z \in G_i(s_{-i})$, hence $z \leq Q_i(\theta'_i(z), s_{-i})$ which by (7), implies $\theta'_i(z) \geq P_i(z, s_{-i})$, so the integrand in (12) is ≤ 0 as desired. \square

Figure 3 shows the utility function of player 4, $u_4(s_4)$, in a PSP auction with $I = 5$ players, with s_{-4} fixed, and a valuation $\theta_4(q) = 10q$. The plateaus correspond to the points where $q_4 \geq Q_4(p_4, s) = \left[Q - \sum_{\{j:p_j > p_4\}} a_j(s)\right]^+$, and $a_4(s)$ can no longer be increased at that bid price – see (2). At bid prices $p_4 > p_5$, the utility decreases when $a_4 > Q - q_5$, because after that point, each additional unit of resource is taken away from player 5, and thus costs p_5 , which is more than θ'_4 its value to player i . Thus, each additional unit starts bringing negative utility. This is what discourages users from bidding above their valuation. Proposition 1 is illustrated by the fact that for any given quantity q_4 , the utility u_4 is maximized on the plane $p_4 = \theta'_4 \equiv 10$.

Remark: When the players have linear valuations and no budget constraint ($b_i = \infty$), PSP becomes identical to a second-price auction for a non-divisible object. Then the existence of a Nash equilibrium follows directly from incentive compatibility.

Note that in PSP, the incentive compatibility (optimality of truth-telling) is in the price dimension, for a given quantity. With the message space we have designed, there is no single “true” quantity to declare, the optimal quantity depends on opponent bid prices. Were the message process such that players declared a price and a budget (rather than desired quantity), it may have been possible to design an allocation rule A such that they are inclined to reveal their true budget, thus obtaining incentive compatibility in both dimensions, and hence equilibrium. But such a rule A would likely not have a simple closed form like (2) and (3). In essence, the computational load of translating budgets into shares would be centralized at the auctioneer, thus making the system less scalable to large numbers of users. On the




Figure 3: Utility $u_4(s_4)$ for $s_1 = (100, 1)$, $s_2 = (10, 2)$, $s_3 = (20, 4)$, $s_5 = (20, 7)$, $s_6 = (30, 12)$

other hand, decentralization has a cost too, which is the signaling overhead resulting from players possibly adjusting bids based on opponent bids in the iterated game. Our design is based on the premise that the latter approach is the more scalable of the two (indeed that was the reason for choosing a small message space).

The next property is that the truthful best reply is continuous in opponent profiles (this can be seen in Figure 2: as the “staircase” is varied smoothly, the point of intersection with θ'_i moves smoothly, provided θ'_i is not flat – which is given by the last time item in Assumption 1). To prove that, we will need the following:

Lemma 1 $\forall s, s' \in \mathcal{S}, \forall y, z \geq 0, \forall \delta > 0$, if $\|s_{-i} - s'_{-i}\| < \delta$ then

$$Q_i(y + \delta, s_{-i}) + \delta\sqrt{I} \geq Q_i(y, s'_{-i}) \geq Q_i(y - \delta, s_{-i}) - \delta\sqrt{I}, \quad (13)$$

and

$$P_i(z + \delta\sqrt{I}, s_{-i}) + \delta \geq P_i(z, s'_{-i}) \geq P_i(z - \delta\sqrt{I}, s_{-i}) - \delta. \quad (14)$$

Proof: First, $\|s_{-i} - s'_{-i}\| < \delta$ implies $\sum_k |q'_k - q_k| < \delta\sqrt{I}$, and $p_k + \delta > p'_k > p_k - \delta$. Thus, $\sum_k q_k 1_{\{p_k + \delta > y\}} + \delta\sqrt{I} \geq \sum_k q'_k 1_{\{p'_k > y\}} \geq \sum_k q_k 1_{\{p_k - \delta > y\}} - \delta\sqrt{I}$. Then, using (1) and the identity $(a + b)^+ \leq (a)^+ + (b)^+$, the first result follows.

For any $y < P_i(z, s'_{-i})$, by (7), we have $z > Q_i(y, s'_{-i}) \geq Q_i(y - \delta, s_{-i}) - \delta\sqrt{I}$, which by (8), implies $y - \delta \leq P_i(z + \delta\sqrt{I}, s_{-i})$. Letting $y \nearrow P_i(z, s'_{-i})$, we get $P_i(z, s'_{-i}) \leq P_i(z + \delta\sqrt{I}, s_{-i}) + \delta$.

For any $y > P_i(z, s'_{-i})$, by (8), we have $z \leq Q_i(y, s'_{-i}) \leq Q_i(y + \delta, s_{-i}) + \delta\sqrt{I}$, which by (7) implies $y + \delta \geq P_i(z - \delta\sqrt{I})$. Letting $y \searrow P_i(z, s'_{-i})$, we get $P_i(z, s'_{-i}) \geq P_i(z - \delta\sqrt{I}) - \delta$. \square

Lemma 2 (*Continuity of best reply*) *Under Assumption 1, $\forall i \in \mathcal{I}$, the ϵ -best reply t_i given in Proposition 1 is continuous in s_{-i} on any subset $V_i(\underline{P}, \overline{P}) = \{s_{-i} \in \mathcal{S}_i : \forall z > 0, \overline{P} \geq P_i(z, s_{-i}) \geq \underline{P}\}$, with $\infty > \overline{P} \geq \underline{P} > 0$.*

Proof: Let $z_i = \sup G_i(s_{-i})$. We will show z_i is continuous, and the continuity of $v_i = [z_i - \epsilon/\theta'_i(0)]^+$ and $w_i = \theta'_i(v_i)$ follow immediately (recall that by Assumption 1, θ'_i is continuous).

Suppose there is a discontinuity at some s_{-i} . Then, $\exists \epsilon_1 > 0$, such that $\forall \delta > 0$, $\exists s'_{-i} \in V_i(\underline{P}, \overline{P})$ with $\|s_{-i} - s'_{-i}\| < \delta$ and $\|z_i - z'_i\| \geq \epsilon_1$, where $z'_i = \sup G_i(s'_{-i})$.

Suppose $z_i + \epsilon_1 \leq z'_i = \sup G_i(s'_{-i})$ (the case $z'_i + \epsilon_1 \leq z_i$ is handled identically, with s_{-i} and s'_{-i} interchanged.). Consider the definition of G_i ; since θ'_i is decreasing and $Q_i(\cdot, s_{-i})$ is non-decreasing and $P_i(\cdot, s_{-i}) \geq 0$, any point to the left of z'_i is in the set $G_i(s'_{-i})$, therefore

$$z_i + \epsilon_1 \in G_i(s'_{-i}). \quad (15)$$

Thus, $z_i + \epsilon_1 \leq Q_i(\theta'_i(z_i + \epsilon_1), s'_{-i}) = \left[Q - \sum_k q'_k 1_{\{P'_k > \theta'_i(z_i + \epsilon_1)\}}\right]^+$. Therefore,

$$z_i + \epsilon_1 \leq Q_i(\theta'_i(z_i + \epsilon_1) + \delta, s_{-i}) + \delta\sqrt{I},$$

using Lemma 1.

Also, by (7) $z_i + \epsilon_1 \leq Q_i(\theta'_i(z_i + \epsilon_1), s'_{-i}) \Rightarrow \theta'_i(z_i + \epsilon_1) > P(z_i + \epsilon_1, s'_{-i})$. Now since $s'_{-i} \in V_i(\underline{P}, \overline{P})$, this last expression is $\geq \underline{P} > 0$, hence $\theta'_i(z_i + \epsilon_1) > 0$. Then using Assumption 1, $\theta'_i(z_i + \epsilon_2) \geq \theta'_i(z_i + \epsilon_1) + \gamma_i(\epsilon_1 - \epsilon_2) > \theta'_i(z_i + \epsilon_1) + \delta$, for $\delta < \delta_1 = \frac{\epsilon_1}{\sqrt{I}} \wedge \epsilon_1 \gamma_i$, and $0 < \epsilon_2 < (\epsilon_1 - \delta_1 \sqrt{I}) \wedge (\epsilon_1 - \delta_1 / \gamma_i)$. Therefore, since $Q_i(\cdot, s_{-i})$ is non-decreasing,

$$\begin{aligned} z_i + \epsilon_2 &\leq Q_i(\theta'_i(z_i + \epsilon_2), s_{-i}) + \delta\sqrt{I} - \epsilon_1 + \epsilon_2 \\ &< Q_i(\theta'_i(z_i + \epsilon_2), s_{-i}). \end{aligned} \quad (16)$$

Now (15) also implies that

$$\begin{aligned} b_i &\geq \int_0^{z_i + \epsilon_1} P_i(\eta, s'_{-i}) d\eta \\ &\geq \int_0^{z_i + \epsilon_3} P_i(\eta, s_{-i}) d\eta + \int_0^{z_i + \epsilon_3} [P_i(\eta, s'_{-i}) - P_i(\eta, s_{-i})] d\eta + (\epsilon_1 - \epsilon_3)\underline{P}, \end{aligned}$$

and this holds $\forall \epsilon_3 < \epsilon_1$. Now, using Lemma 1,

$$\begin{aligned}
& \int_0^{z_i + \epsilon_3} [P_i(\eta, s'_{-i}) - P_i(\eta, s_{-i})] d\eta \\
& \geq -\delta Q + \int_0^{z_i + \epsilon_3} P_i(\eta - \delta\sqrt{T}, s_{-i}) d\eta - \int_0^{z_i + \epsilon_3} P_i(\eta, s_{-i}) d\eta \\
& \geq -\delta Q - \int_{z_i + \epsilon_3 - \delta\sqrt{T}}^{z_i + \epsilon_3} P_i(\eta, s_{-i}) d\eta \\
& \geq -(Q + \bar{P})\delta\sqrt{T}.
\end{aligned}$$

Let $\delta_2 = \frac{\epsilon_1 \bar{P}}{(Q + \bar{P})\sqrt{T}}$, and ϵ_3 such that $0 < \epsilon_3 < \epsilon_1 - (Q + \bar{P})\delta_2\sqrt{T}/\underline{P}$. Then

$$b_i \geq \int_0^{z_i + \epsilon_3} P_i(\eta, s_{-i}) d\eta, \quad (17)$$

for $\delta < \delta_2$.

Now choosing $\delta < \delta_1 \wedge \delta_2$, (16) and (17) imply that $G_i(s_{-i}) \ni (z_i + \epsilon_3) \wedge (z_i + \epsilon_2) > z_i = \sup G_i(s_{-i})$, a contradiction. \square

We introduce one additional player, player 0, whose valuation is $\theta_0(z) = p_0 z$, and whose bid can therefore be fixed at $s_0 = (q_0, p_0) = (Q, p_0)$. Player 0 can be viewed as the auctioneer, and $p_0 > 0$ as a “reserve price” at which the seller is willing to “buy” all of the resource from himself. From (1), the presence of the bid $s_0 = (Q, p_0)$ implies $\forall i \in \mathcal{I}, Q_i(y, s_{-i}) = 0, \forall y < p_0$. In particular, setting $y = 0$, the condition of Proposition 1 holds. Thus, we can restrict our attention to truthful strategies only, and still have feasible best replies. This forms a “truthful” game embedded within the larger auction game, where the strategy space is $\mathcal{T} \subset \mathcal{S}$, the feasible sets are $\mathcal{T}_i \cap S_i(s_{-i})$, and the best replies are $R_i^\epsilon(s) = \mathcal{T}_i \cap S_i^\epsilon(s)$. A fixed point of R^ϵ in \mathcal{T} is a fixed point of S^ϵ in \mathcal{S} . Thus an equilibrium of the embedded game is an equilibrium of the whole game.

Proposition 2 (*Nash equilibrium*) *In the auction game with the PSP rule given by (2) and (3) and a reserve price $p_0 > 0$, and players described by (4) and (5), if Assumption 1 holds, then for any $\epsilon > 0$, there exists a truthful ϵ -Nash equilibrium $s^* \in \mathcal{T}$.*

Proof: $\forall s \in \mathcal{T}, \forall i \in \mathcal{I}, \forall z > 0$, we have $z > 0 = Q_i(p_0/2, s_{-i})$, which by (8) implies $P_i(z, s_{-i}) \geq p_0/2 = \underline{P}$. Let $\bar{P} = \max_{k \in \mathcal{I} \cup \{0\}} \theta'_k(0)$. Then, the conditions of Lemma 2 are satisfied and $t = (v, w)$ is **continuous in s on \mathcal{T}** . By Assumption 1, θ'_i is continuous therefore $v(q, p) = v(q, \theta'(q))$ (as defined in Proposition 1),

can be viewed as a continuous mapping of $[0, Q]^I$ onto itself. By Brouwer's fixed-point theorem (see for example [10]), any continuous mapping of a convex compact set into itself has at least one fixed point, i.e. $\exists q^* = v(q^*) \in [0, Q]^I$. Now with $s^* = (q^*, \theta'(q^*))$, we have $s^* = t(s^*) \in \mathcal{T}$. \square

3.3 Efficiency

The objective in designing the auction is that, at equilibrium, resources always go to those who value them most. Indeed, the PSP mechanism does have that property. This can be loosely argued as follows: for each player, the marginal valuation is never greater than the bid price of any opponent who is getting a non-zero allocation. Thus, whenever there is a player j whose marginal valuation is less than player i 's and j is getting a non-zero allocation, i can take some away from j , paying a price less than i 's marginal valuation, i.e. increasing u_i , but also increasing the total value, since i 's marginal value is greater. Thus at equilibrium, i.e. when no one can unilaterally increase their utility, the total value is maximized. Formally, consider $a \in \arg \max_{\mathcal{A}} \sum_i \theta_i(a_i)$. The Karush-Kuhn-Tucker [17] optimality conditions are that there exists a Lagrange multiplier λ such that $\theta'_i(a_i) = \lambda$, if $a_i > 0$, and $\theta'_i(0) \leq \lambda$, if $a_i = 0$.

Assumption 2 For any $i \in \mathcal{I}$, $b_i = \infty$, and θ'_i satisfies⁶

$$\theta'_i(z) - \theta'_i(z') > -\kappa(z - z'),$$

whenever $z > z' \geq 0$.

Given any $\epsilon > 0$, for any ϵ -Nash-equilibrium $s^* \in \mathcal{T}$, let $a^* \equiv a(s^*)$, and let $\underline{a}^* \equiv \min_{i \in \mathcal{I} \cup \{0\}, a_i > 0} a_i^*$, the smallest non-zero allocation. The following is the “ ϵ version” of the Karush-Kuhn-Tucker conditions.

Lemma 3 Suppose Assumptions 1 and 2 hold. If for some j , $a_j^* > \sqrt{\epsilon/\kappa}$, then $\forall i \in \mathcal{I} \cup \{0\}$,

$$\theta'_i(a_i^*) < \theta'_j(a_j^*) + 2\sqrt{\epsilon\kappa}.$$

An immediate corollary is that if $\underline{a}^* > \sqrt{\epsilon/\kappa}$ then

$$\lambda^* - 2\sqrt{\epsilon\kappa} < \theta'_i(a_i^*) < \lambda^* + 2\sqrt{\epsilon\kappa}$$

⁶If θ'_i is differentiable, the condition is $0 \geq \theta''_i > \Leftrightarrow \kappa$.

if $a_i^* > \sqrt{\epsilon/\kappa}$ and

$$\theta'_i(a_i^*) < \lambda^* + 2\sqrt{\epsilon\kappa},$$

if $a_i^* = 0$, for some $\lambda^* \geq 0$.

Proof: Suppose $\theta'_i(a_i^*) \geq \theta'_j(a_j^*) + 2\sqrt{\epsilon\kappa}$. Since $\theta'_j(a_j^*) \geq \theta'_j(q_j^*) \geq p_j^*$, we have $\theta'_i(a_i^*) \geq p_j^* + 2\sqrt{\epsilon\kappa}$.

Since $a_j^* > 0$, if player i bids at a price above p_j^* , he can take all of player j 's allocation, without losing anything of his own, i.e. $a_i^* + a_j^* \leq Q_i(p_j^*, s_{-i}^*)$. By (7), this implies

$$p_j^* \geq P_i(a_i^* + a_j^*, s_{-i}^*).$$

Let $q_i = (a_i^* + \sqrt{\epsilon/\kappa})$ and $s_i = (q_i, \theta'_i(q_i))$. Then

$$\begin{aligned} u_i(s_i; s_{-i}^*) - u_i(s^*) &= \int_{a_i^*}^{a_i^* + \sqrt{\epsilon/\kappa}} \theta'_i(z) - P_i(z, s_{-i}^*) dz \\ &\geq \left[\theta'_i(a_i^* + \sqrt{\epsilon/\kappa}) - p_j^* \right] \sqrt{\epsilon/\kappa} \\ &\geq \left[\theta'_i(a_i^*) - \kappa\sqrt{\epsilon/\kappa} - p_j^* \right] \sqrt{\epsilon/\kappa} \\ &\geq \left[2\sqrt{\epsilon\kappa} - \kappa\sqrt{\epsilon/\kappa} \right] \sqrt{\epsilon/\kappa} \\ &= \epsilon \end{aligned}$$

which contradicts the fact that s^* is an ϵ -Nash equilibrium. \square

Proposition 3 (Efficiency) Suppose Assumptions 1 and 2 hold. If $\underline{a}^* > \sqrt{\epsilon/\kappa}$, then

$$\max_{\mathcal{A}} \sum_i \theta_i(a_i) - \sum_i \theta_i(a_i^*) = O(\sqrt{\epsilon\kappa}),$$

where $\mathcal{A} = \{a \in [0, Q]^{I+1} : \sum_i a_i \leq Q\}$.

Proof: (of Proposition 3) Let $\mathcal{I}^+ = \{k : a_k > a_k^*\}$ and $\mathcal{I}^- = \{k : a_k < a_k^*\}$. For $i \in \mathcal{I}^+$, we have $\theta'_i(a_i^*) \leq \lambda^* + 2\sqrt{\epsilon\kappa}$. For $i \in \mathcal{I}^-$, we have $a_i^* > a_i \geq 0$, therefore by the lemma, $\theta'_i(a_i^*) > \lambda^* - 2\sqrt{\epsilon\kappa}$. Therefore,

$$\begin{aligned} \sum_{\mathcal{I}} \theta_i(a_i) - \theta_i(a_i^*) &\leq \sum_{\mathcal{I}^+} \theta'_i(a_i^*)(a_i - a_i^*) - \sum_{\mathcal{I}^-} \theta'_i(a_i^*)(a_i^* - a_i) \\ &\leq (\lambda^* + 2\sqrt{\epsilon\kappa})\Delta - (\lambda^* - 2\sqrt{\epsilon\kappa})\Delta, \end{aligned}$$

where $\Delta = \sum_{\mathcal{I}^+} (a_i - a_i^*) = \sum_{\mathcal{I}^-} (a_i^* - a_i)$. Since $\Delta \leq Q$ the result follows, with the bound $4Q\sqrt{\epsilon\kappa}$. \square

Remark a: The condition $b_i = \infty$ is sufficient, but not necessary to achieve efficient outcomes. In fact with any budget profile, efficiency can be achieved if the users cooperate. For example, if they all choose a bid quantity close to what they can actually obtain (which they do if they use the strategy given by Proposition 1), then the price paid would be p_0 per unit for all the allocations, and if p_0 or the shares a_i^* are not too large, then budget constraints are irrelevant and a^* is efficient. More generally, efficiency is attained if the budgets are not too far out of line with the valuations, i.e. there are no players with very high demand and very low budget.

Remark b: (*Welfare and Efficiency*) A more common measure of efficiency is the social welfare, which is the sum of all the players' utility $\sum_i u_i$, including the seller $i = 0$. The natural definition of the seller's utility is the value of the leftover capacity $a_0 = Q - \sum_{i \neq 0} a_i$ plus the revenue, i.e.

$$u_0 = \theta_0(a_0) + \sum_{i \neq 0} c_i.$$

Then, $\sum_i u_i = \sum_{i \neq 0} (\theta_i - c_i) + u_0 = \sum_{i \neq 0} \theta_i + \theta_0(a_0) = \sum_i \theta_i$. Thus, $\sum_i u_i$ is equivalent to the efficiency measure used above, which is $\sum_i \theta_i$. Another measure is the seller's revenue. Even though PSP is not, in general, revenue-maximizing, it tends to the revenue maximizing allocations and prices as demand increases [18].

Remark c: Proposition 3 provides a key to understanding the basic trade-off between engineering and economic efficiency. The smaller ϵ , the closer we get to the value-optimal allocations. But in a dynamic game, where players iteratively adjust their bids to the opponent profile, a player will bid as long as he can gain at least ϵ utility (since that is the cost of the bid), thus a smaller ϵ makes the iteration take longer to converge, i.e. entails more signaling.

3.4 Convergence

An issue of obvious concern is whether the game converges under dynamic play: it turns out that it does, when users behave rationally (see Proposition 4 in Chapter 2 of [29]). Moreover, irrational or malicious behavior – like intentionally trying to prevent convergence by making unnecessary bids – can always be controlled by setting the bid fee ϵ high enough to make such behavior prohibitively costly for the culprit.




Figure 4: Parabolic valuation with $\kappa = 0.5$ and $\bar{q} = 70$

Another issue is how the convergence time scales with the number of bidders. We now consider this experimentally using the software described in Appendix B.

In all our simulations we let $Q = 100$. For each user, the valuation is strictly increasing and concave up to a maximum corresponding to a physical line capacity, and flat beyond that. Since, as shown by Proposition 3, only the second derivative of the valuation is needed to measure the efficiency of the PSP auction, a second order (parabolic) model is deemed sufficient. Thus we use valuations of the form:

$$\theta_i(z) = -\kappa_i(z \wedge \bar{q}_i)^2/2 + \kappa_i\bar{q}_i(z \wedge \bar{q}_i),$$

where \bar{q}_i is the line rate, and $\kappa_i > 0$ (see Figure 4).

We generate our user population with independent random variables $\{\theta'_i(0)\}_{\mathcal{I}}$ (corresponding to the maximum unit price the user would pay) uniformly distributed on $[10, 20]$, and $\kappa_i = \theta'_i(0)/\bar{q}_i$, and \bar{q}_i uniformly distributed on $[50, 100]$. All players have a budget $b_i = 100$. The bid fee is fixed at $\epsilon = 5$. Each user has a bidding agent which can submit at most one bid per second (see Algorithm 1 in Appendix B).

With this set-up, the results are shown in Figures 5-6. Simulations were run for 11 population sizes ranging from 2 to 96 players. Each point is




Figure 5: Mean (+/- std. dev.) number of bids – solid line. The dashed line is $I + I^2/10$.

simulated 10 times with new random valuations for all players. The overall mean is 11.9 bids per player. From Figure 5, the number of bids seems to grow as the square of the number of players.

The actual time to converge, shown in Figure 6, grows more slowly, since the computation of bids is done in parallel by all the players. In fact, for small numbers of players, the time decreases. This can best be explained as follows. Suppose there are only two players, with similar valuations. They will both start by asking for their maximum quantity, at their marginal valuation (which at their maximum quantity is near zero). Then as each sees the other's bid, each will reduce the quantity and increase the price a little bit. And they go on taking turns, gradually raising the market price until they reach an equilibrium. However if there are 10 players, in between two bids by the same player, the 9 others will already have bid up the price, so he will jump to higher price than if there was only one opponent. Thus the equilibrium market price will be reached more quickly. For large populations, this effect becomes small compared to the sheer volume of bids, and the convergence time starts to grow.

The trade-off between signaling and economic efficiency discussed in light of Proposition 3 is illustrated by Figures 7-8. Increasing the bid fee speeds




Figure 6: Mean (+/- std. dev.) convergence time in seconds (for a 1 second bid interval).

up convergence, at a cost of lost efficiency. A resource manager should select a bid fee which optimally balances the two for the particular context.

Figure 8 also illustrates the validity of the lower bound given by Proposition 3.

4 Conclusion

Auctions are one of oldest surviving classes of economic institutions [...] As impressive as the historical longevity is the remarkable range of situations in which they are currently used. [22]

We proposed the progressive second price auction, a new auction which generalizes key properties of traditional single non-divisible object auctions to the case where an arbitrarily divisible resource is to be shared. We have shown that our auction rule, assuming an elastic-demand model of user preferences, constitutes a stable and efficient allocation and pricing mechanism in a network context. Even though we are motivated by problems of bandwidth and buffer space reservation in a communication network, the auction was formulated in a manner which is generic enough for use in a wide range




Figure 7: Number of bids to converge vs ϵ




Figure 8: $\sum_i \theta_i(a_i^*)$ – solid line, and $\max_a \sum_i \theta_i(a_i) - 4(\epsilon\kappa)^{-1/2}$ – dashed line, vs ϵ .

of situations. In the sequel to this work, we show that the key results – namely incentive compatibility, equilibrium, and efficiency – generalize to a setting where multiple networked resources are auctioned, with users bidding on arbitrary but fixed routes and topologies [29, 31]. In related work, we consider the case of stochastically arriving players bidding for advance reservations (i.e. resources for a given period of time) [29, 30].

An interesting direction of future work is learning strategies, and evolutionary behavior which can emerge from repeated inter-action between the same players.

5 Acknowledgements

The authors would like to thank Wang Ke for his valuable help in the simulations.

A Information-theoretic basis for the valuation

In general, valuations are simply assumed to be given as external factors. Indeed, the fundamental assumption in any market theory is that buyers know what the goods are worth to them. The “elastic demand” or “diminishing returns” nature of Assumption 1 is fully justified from a purely economic standpoint for virtually all resources in everyday life.

In the case of variable bandwidth, we can go even further by better quantifying what the goods are. For a user sending video, say, how much value is lost when the channel capacity goes from 1.5 to 1.2 Mbps? Ultimately, the value lies not in the amount of raw bandwidth but in the information that is successfully sent. Our goal in this section is to give a brief description of how Information Theory can be used for a bottom-up construction of bandwidth valuations – based on the fundamental thing the user cares about which is communication of information – and that such valuations will generally be of the type in Assumption 1.

Any information source has a function $D(\cdot)$, such that when compressed to a rate R , the signal has a distortion of at least $D(R)$ [1]. The distortion is the least possible expected “distance” between the original and compressed signals, where the minimization is over all possible coding/decoding schemes. In this context, we make the distance measure the monetary cost of the error. This cost can be chosen, for example, to be proportional to some common measures like the mean squared error, the Hamming distance (probability of error), the maximum error, etc., or heuristic measures based on experiments with human perception. Given that modern source-coding techniques can, given a distortion measure, achieve distortions close to the theoretical lower bound[5], it is not unreasonable to use the rate-distortion curve as an indication of the value of the bandwidth share.

Let $D_i(\cdot)$ be the distortion-rate function of $\{X_i(t)\}$, a stochastic process modeling the source of information associated with user i . X_i is encoded as Y_i which has a rate of R bits per second.

Shannon’s channel coding theorem[32] states that Y_i can be received without errors if and only if the channel has a capacity $C > R$. In our auction context, user i has capacity (bandwidth allocation) $C = a_i$, and thus has to suffer a distortion of at least $D_i(a_i)$. The value of the bandwidth is then

$$\theta_i(a_i) = \bar{\theta}_i - D_i(a_i), \tag{18}$$

where $\bar{\theta}_i$ is the value of the full information.

The relevant properties of the distortion-rate functions are:

- when the rate is greater than the entropy of the source, the distortion is zero, and
- for many common source models and cost functions, the distortion-rate function is convex, and has a continuous derivative.

It is easy to see that, with these properties, (18) satisfies Assumption 1.

Example 1: Let $\{X\}$ be a Bernoulli source, taking two values with probabilities p and $1 - p$. Without loss of generality, let $p \in [0, 1/2]$. In this case, since the source is i.i.d, one can define the distortion on a per symbol basis. Using a Hamming cost function

$$d(X, Y) = 1_{\{X \neq Y\}},$$

i.e. assuming it costs one unit of money every time one bit is wrong, we have the distortion $D = Ed(X, Y) = P(X \neq Y)$. the rate-distortion function is

$$R(D) = [H(p) - H(D)]^+,$$

where $H(x) = -x \log(x) - (1 - x) \log(1 - x)$, and the distortion-rate function is the inverse function. It can be easily seen that $D(R)$ is strictly convex and decreasing for $0 \leq R < H(p)$, and $D(R) = 0$ for $R \geq H(p)$. The continuity of D' on $0 \leq R < H(p)$ and $R > H(p)$ is obvious. At the critical point ($R = H(p), D = 0$), we have

$$\begin{aligned} \lim_{R \nearrow H(p)} D'(R) &= \lim_{D \searrow 0} 1/R'(D) \\ &= \lim_{D \searrow 0} 1/\log(D/1 - D) \\ &= 0 \\ &= \lim_{R \searrow H(p)} D'(R). \end{aligned}$$

Thus continuity of D' holds throughout, and Assumption 1 is valid for the valuation of the form (18) for this source – see Figure 9.

Example 2: Let $\{X\}$ be a Gaussian source with Markovian time-dependency, i.e a covariance matrix $\Phi = (\sigma^2 r^{|i-j|})_{i,j}$, $r \in [0, 1)$. Suppose we use the squared error cost, i.e. it costs one unit of money for one unit of energy in the error signal. Then, we have for low distortions $D \leq (1-r)/(1+r)$, $R(D) = \frac{1}{2} \log \frac{1-r^2}{D}$, or

$$D(R) = (1 - r^2) \sigma^2 2^{-2R},$$




Figure 9: Distortion-rate based valuation for a Bernoulli $p = 1/2$ source

and Assumption 1 clearly holds for (18). In the i.i.d. case ($r = 0$), the formula holds for all R .

As the source models get more complex, it rapidly becomes impossible to give closed-form expressions for either $R(D)$ or $D(R)$. Often parametric forms are available, and the functions can be evaluated numerically. Fortunately, the convexity property extends to a wide class of models, including for example auto-regressive sources, even when the generating sequence is non-Gaussian (see [1] for a full treatment of $R(D)$, including the above cases).

It can happen, e.g. for some video source models, that the R-D curve, which gives the best (R,D) pairs achievable by any coder/decoder, is not convex. But, for tractability, practical codecs are usually optimized on a convex hull of the space of possible (R,D) pairs[26, 27]. Thus, even when the theoretical $D(R)$ curve is not convex, the actual distortion achieved in real-life systems almost always varies in a convex manner with the available bandwidth.

B Simulation software and bidding algorithm

A prototype software agent based implementation of the auction game, called TREX, has been developed and extensively used since December 1995. Much of the intuition behind the mechanism design and the analysis in this work came from experiments done on this inter-active distributed auction game on the World Wide Web, using the Java programming language. The game can be played in real-time by any number of players from anywhere on the Internet [28].

Each user plays in the dynamic auction game using the following:

- Algorithm 1**
- 1 Let $s_i = 0$, and $\hat{s}_{-i} = \emptyset$. Start an independent thread which receives updates of \hat{s}_{-i} .
 - 2 Compute the truthful ϵ -best-reply of Proposition 1, $t_i \in \mathcal{T}_i \cap S_i^\epsilon(\hat{s}_{-i})$.
 - 3 If $u_i(t_i; \hat{s}_{-i}) > u_i(s_i; \hat{s}_{-i}) + \epsilon$, then send the bid $s_i = t_i$.
 - 4 Sleep for 1 second.
 - 5 Go to 2.

No assumption is made on the order of the turns. Players join the game at different times, and depending on the execution context of the client program, the sleep time of 1 second is more or less approximate. This, along with communication delays which make the times at which bids arrive at the server and updates at the clients essentially random times, make the distributed game completely asynchronous.

Algorithm 1 can be described as selfish and short-sighted. Selfish because it will submit a new bid if and only if it can improve its own utility (by more than the fee for the bid). Thus, the game can only converge⁷ to an ϵ -Nash equilibrium. And short-sighted because it does not take the extensive form of the game into account, i.e. does not use strategies which may result in a temporary loss but a better utility in the long run.

References

- [1] T. Berger. *Rate Distortion Theory: A Mathematical Basis for Data Compression*. Prentice-Hall, 1971.

⁷The fact that it does converge is established by Proposition 4 in Chapter 2 of [29].

- [2] D. Clark. Internet cost allocation and pricing. In L. W. McKnight and J. P. Bailey, editors, *Internet Economics*. MIT Press, 1997.
- [3] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 8:17–33, 1971.
- [4] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991.
- [5] A. Gersho and R. M Gray. *Vector Quantization and Signal Compression*. Kluwer Academic Publishers, 1992.
- [6] J. Gong and S. Marble. Pricing common resources under stochastic demand. preprint – Bellcore, April 1997.
- [7] J. Gong and P. Sriganesh. An economic analysis of network architectures. *IEEE Network*, pages 18–21, March/April 1996.
- [8] T. Groves. Incentives in teams. *Econometrica*, 41(3):617–631, July 1973.
- [9] M. T. T. Hsiao and A. A. Lazar. Optimal decentralized flow control of markovian queueing networks with multiple controllers. *Performance Evaluation*, 13(3), Dec. 1991.
- [10] Hurewicz and Wallman. *Dimension Theory*. Princeton University Press, 1948.
- [11] H. Jiang and S. Jordan. Connection establishment in high speed networks. *IEEE J. Select. Areas Commun.*, 13(7):1150–1161, 1995.
- [12] H. Jiang and S. Jordan. The role of price in the connection establishment process. *European Trans. Telecommunications*, 6(4):421–429, 1995.
- [13] E. Kalai and E. Lehrer. Rational learning leads to Nash equilibrium. *Econometrica*, 61(5):1019–1045, September 1993.
- [14] F. P. Kelly. Charging and accounting for bursty connections. In L. W. McKnight and J. P. Bailey, editors, *Internet Economics*. MIT Press, 1997.
- [15] Y. A. Korilis and A. A. Lazar. On the existence of equilibria in noncooperative optimal flow control. *J. ACM*, 42(3), May 1995.

- [16] Y. A. Korilis, A. A. Lazar, and A. Orda. Architecting noncooperative networks. *IEEE J. Select. Areas Commun.*, 13(7), September 1995.
- [17] H. W. Kuhn and A. W. Tucker. Non-linear programming. In *Proc. 2nd Berkeley Symp. on Mathematical Statistics and Probability*, pages 481–492. Univ. Calif. Press, 1961.
- [18] A. A. Lazar and N. Semret. Auctions for network resource sharing. Technical Report CU/CTR/TR 468-97-02, Columbia University, 1997. <http://comet.columbia.edu/publications>.
- [19] J. K. MacKie-Mason and H. R. Varian. Pricing the internet. In B. Kahin and J. Keller, editors, *Public Access to the Internet*. Prentice Hall, 1994. <ftp://ftp.econ.lsa.umich.edu/pub/Papers>.
- [20] A. Mas-Colell and M. D. Whinston. *Microeconomic Theory*. Oxford University Press, 1995.
- [21] R. P. McAfee and J. McMillan. Multidimensional incentive compatibility and mechanism design. *Journal of Economic Theory*, 46(2):335–354, 1988.
- [22] P. R. Milgrom. *Advances in Economic Theory: Fifth World Congress*, chapter Auction Theory, pages 1–31. Number 12 in Econometric Society Monographs. Cambridge University Press, 1987.
- [23] M. S. Miller and K. E. Drexler. Markets and computation: Agoric open systems. In Bernardo Huberman, editor, *The Ecology of Computation*. Elsevier Science Publishers/North-Holland, 1988.
- [24] R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, February 1981.
- [25] A. Orda, R. Rom, and N. Shimkin. Competitive routing in multiuser communication networks. *IEEE/ACM Trans. Networking*, 1(5), October 1993.
- [26] G. Pacifici, G. Karlsson, M. Garrett, and N. Ohta, editors. *Real-Time Video Services in Multimedia Networks*, volume 15 of *IEEE J. Select. Areas Commun.*, Aug. 1997.
- [27] G. M. Schuster and A. K. Katsaggelos. *Rate-Distortion Based Video Compression*. Kluwer Academic Publishers, 1997.

- [28] N. Semret. TREX-The Resource EXchange, 1995-96. <http://comet.columbia.edu/~nemo/Trex>.
- [29] N. Semret. *Market Mechanisms for Network Resource Sharing*. PhD thesis, Columbia University, 1999.
- [30] N. Semret and A. A. Lazar. Spot and derivative markets in admission control. In *Proc. IEE Int'l Teletraffic Congress*, Edinburgh, UK, June 1999. Elsevier Science Publishers B.V. <http://comet.columbia.edu/publications/conference.html>.
- [31] N. Semret, R. R.-F. Liao, A. T. Campbell, and A. A. Lazar. Market pricing of differentiated internet services. In *IEEE/IFIP 7th Int. Workshop on Quality of Service*, 1999.
- [32] C. E. Shannon. A mathematical theory of communication. *Bell Syst. Tech. J.*, 27, 1948.
- [33] S. Shenker. Making greed work in networks: A game-theoretic analysis of switch service disciplines. In *Proc. ACM SIGCOMM*, 1994.
- [34] S. Shenker, D. Clark, D. Estrin, and S. Herzog. Pricing in computer networks: Reshaping the research agenda. *ACM Comput. Commun. Review*, 26(2):19–43, 1996.
- [35] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, 16, 1961.
- [36] S. R. Williams. Realization and Nash implementation: Two aspects of mechanism design. *Econometrica*, 54(1):139–151, January 1986.