CTR Technical Report: CU/CTR/TR 468-97-02

Auctions for Network Resource Sharing

Aurel A. Lazar and Nemo Semret
Center for Telecommunications Research,
Columbia University
{aurel, nemo}@ctr.columbia.edu

February 11, 1997

Abstract

We consider the use of an auction as a decentralized mechanism for efficiently and fairly sharing resources in a network. We present a new auction rule called Progressive Second Price, and prove that it has a "truthful" Nash equilibrium. We give necessary conditions for the resulting allocations to be fair (envy-free), and demonstrate that, with a large population of small users, the allocations are increasingly efficient when total demand exceeds supply. We also present an algorithm extending the auction to a network of resources, and show that in the case of a tree topology, the efficiency of the allocations is preserved. This work is applicable, for example, in joint admission control and pricing on a network where users may request different, arbitrary levels of quality of service.

1 Introduction

Resource allocation is among the most keenly debated and least understood issues in networking today. While it is widely believed that pricing plays a determining role in achieving desirable uses of resources, in reality commercial Internet service providers, on-line services, and even telephone companies are struggling to develop pricing policies that satisfy their customers' changing needs and demands, and result in efficient use of resources. The centralized policies of the single service world of traditional telephony are inadequate in packet switched multiservice integrated networks where the demands range over wide quantitative and qualitative ranges. Here, "given that applications have very different sensitivities to service quality, it seems preferable to place the bulk of the variability where it can be done in the most informed way." [13]

Rather than trying to centrally and explicitly compute prices which achieve some system wide objective, auctions decentralize the decision making, in that prices emerge from the users' valuations of (and willingness to pay for) resources. The motivation for auctions is that one can obtain better allocations because, a careful design of the auction rules can leverage the collective "intelligence" of the users into a result greater than that of a centralized system. A centralized[1] pricing policy in reality leads to somewhat arbitrarily set prices. This is because there are a number of fundamental barriers to an optimal centralized approach. First, the policy maker would have to know the preferences of the different users in terms of resources. The relationship between preceived quality of service and allocated resources can be very different for different users, so the policy necessarily assumes that users are cooperative in announcing them truthfuly or that they are somehow known. Second, it assumes that it is possible and computationally scalable to *objectively* compare and optimally trade-off users' utilities in a centralized way. These are at best naive, if not invalid, assumptions, in a multiservice, multimedia network. The recognition of similar realities in many aspects of networks and distributed computations has lead in recent years to the emergence of game theoretic approaches in the analysis and design of these systems [8, 12, 4].

[1]By centralized, we mean that the relationship between demand and price is decided by an a-priori formula, even if the actual price levels may vary according to demand – e.g. time-of-day pricing in the telephone network. In an auction, the policy is decentralized in the sense that a given demand may lead to different prices depending on the users valuation of the resources and their bidding strategies.

In this paper, we introduce and analyse an auction rule for resource sharing, whose design is guided by the basic principles of automated agent mechanism design: stability, simplicity, efficiency, and fairness [10]. Our rule is a generalization of Vickrey auctions [14] (which applies to the sale of a single non-divisible object) to the case of shared (divisible) resources.

Auctions have been used to allocate computation and communication resources in [1, 15]. A "smart market" mechanism similar to an auction is suggested in [6] for pricing Internet service. In their approach, because the bids are in one dimension only (price per packet), the market clearing price has to be centrally set to equal a marginal congestion cost, which is computed with an explicitly assumed utility function for the users. Thus, theirs is essentially a centralized pricing policy, and is thus subject to the scalability and "knowability" drawbacks described above. This appears unavoidable in a pure datagram network with no notion of flows for which a specific amount of resources can be reserved. By contrast, here, the possibility of per-flow (or connection) resource reservation is taken as a given, so we can make the bids two dimensional (price and quantity). Thus, a clearing price arises directly from the bids only, and the users' preferences are not part of the mechanism itself.[2]

After formaly presenting the problem in Section 2, in Section 3, we present a new auction rule which we call Progressive Second Price, and show that it has the desired properties of simplicity and stability (in the sense of Nash equilibrium). With some additional assumptions on players, fairness and efficiency are achieved. Then in Section 4, we look at combinations of resources, e.g. paths or circuits in a network. We present an algorithm which gives an allocation rule for multiple resources, which works by applying a single-resource rule recursively. In the case where the set of resources (the network) has a tree structure, if the single-resource rule is efficient, then so is the tree rule. In addition to making fair and efficient allocations, the mechanism is computationally efficient. For a number of users I, and a number of resources L, the allocations are computed in time $O(I^2L)$.

The proofs of all the results are given in Appendix A.

[2] In our mechanism, no knowledge about the users goes into the computation of the allocations. However, we do of course assume certains forms of the players' utilities in the *analysis* of their behaviour.

2 Formulation of an Auction for a Divisible Resource

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

2.1 Allocation Rule

Player i's bid is $s_i \stackrel{def}{=} (q_i, p_i) \in \mathcal{S}_i \stackrel{def}{=} [0, Q] \times [0, \infty)$, meaning he would like a quantity q_i at a *unit* price p_i. A bid profile is $s = (s_1^T, \ldots, s_I^T)^T$. The i-th row of s is player i's bid. Let q denote the operator which extracts the first column of its operand, i.e. $qs \equiv s(1,0)^T = (q_1, \ldots, q_I)^T$, and ps is defined similarly. Note that $qs_i = q_i$ and $ps_i = p_i$.

Following standard game theoretic notation, let $s_{-i} \equiv (s_1^T, \ldots, s_{i-1}^T, s_{i+1}^T, \ldots, s_I^T)^T$, i.e. the bid profile of player i's opponents, obtained from s by deleting the row 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,

$$\begin{array}{rrcl} A: & \mathcal{S} & \longrightarrow & \mathcal{S} \\ & s = (qs, ps) & \longmapsto & A(s) = (qA(s), pA(s)), \end{array}$$

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

The i-th row of $A(s)$, $A_i(s) = (qA_i(s), pA_i(s))$, is the allocation to player i: she gets a quantity $qA_i(s)$ at a unit price $pA_i(s)$.

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

$$\begin{array}{rcl} \sum_{i \in \mathcal{I}} qA_i(s) & \leq & Q \\ A(s) & \leq & s \end{array}$$

where $\leq$ for matrices is taken element by element.

Player i's preferences are given by his utility function

$$\begin{array}{rrcl} u_i: & \mathcal{S} & \longrightarrow & [0, \infty) \\ & s & \longmapsto & u_i(s). \end{array}$$

An auction game is given by $(Q, u_1, \ldots, u_I, A)$, that is by completely specifying the resource, the players, and a feasible allocation rule.

The above formulation is a generalization of what is usually meant by auction. If the condition $qA_w(s) = Q$ for some winner $w \in \mathcal{I}$ and $qA_i(s) = 0, \forall i \neq w$, then we get the traditonal winner-take-all type of auction, i.e. as the sale of a single indivisible object to one buyer, for which the theory is well developed [7, 9]. 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. However, that would assume either that the value of the resource to a user is the same for each unit of resource (a considerable loss in flexibility since one might have, for example, a user whose valuation is decreasing for each additional unit of resource beyond a minimum quantity), or that the user has potentially as many utility functions as there are units of resource (which obviously is not tractable). Furthermore, in a practical implementation of auctions for sharing a resource, a process of bidding for each individual unit would result in a tremendous "signalling" overhead. More importantly, since the users would be bidding on a discrete grid of quantities, the analysis would be susceptible to being very sensitive to the choice of a particular grid and can thus easily give misleading predictions of outcome[3].

2.2 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.

We assume player i has a **valuation** of the resource $\theta_i \geq 0$ of each unit of resource that she gets. Thus, the total value to her of her allocation is $\theta_i qA_i(s)$.[4] Thus, for a bid profile of s, under allocation rule A, player i getting an allocation $A_i(s)$ has utility

$$u_i(s) = \theta_i qA_i(s) - qA_i(s)pA_i(s), \tag{1}$$

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

In addition, we assume that the player is constrained by a **budget** $b_i \geq 0$, so the bid s_i must lie in the set

$$S_i(s_{-i}) = \{s_i \in \mathcal{S}_i : qA_i(s_i; s_{-i})pA_i(s_i; s_{-i}) \leq b_i\}. \tag{2}$$

[3]For a more detailed discussion of this point, see [2] p. 34, and references therein.

[4]For full generality, the value of the allocation should be $\Theta_i(qA_i(s))$, where Θ_i is any positive, non-decreasing function.

We assume there always exists a player 0, whose bid is fixed at $s_0 = (qs_0, ps_0) \equiv (Q, \theta_0)$. Player 0 can be viewed as the auctioner himself. It is natural to set $\theta_0 = 0$, since the auctioner is always willing to "buy" all of the resource from himself at a price of zero, in other words he keeps whatever he can't sell. However, we will assume that $\theta_0 > 0$, i.e., that the auctioner imposes a minimum unit price. This "reservation price" can be made arbitrarily small, and can be used to model a cost-recovery requirement on behalf of the resource controller.

3 Progressive Second Price Rule

We present a specific auction rule A which we will show achieves our design goals. The "progressive second price" (PSP) allocation rule is defined as follows.

$$qA_i(s) = qs_i \wedge \left[Q - \sum_{\{j: ps_j > ps_i\}} qA_j(s) \right]^+ \tag{3}$$

$$pA_i(s) = \begin{cases} \frac{\sum_{j \neq i} ps_j [qA_j(0; s_{-i}) - qA_j(s_i; s_{-i})]}{\sum_{j \neq i} [qA_j(0; s_{-i}) - qA_j(s_i; s_{-i})]}, & \text{if } qA_i(s) > 0 \\ 0, & \text{if } qA_i(s) = 0 \end{cases} \tag{4}$$

where $\wedge$ means taking the minimum.

Because of player 0's fixed bid of (Q, ps_0), with $ps_0 > 0$,

$$\sum_{i=0}^{I} qA_i(s) = Q, \tag{5}$$

and therefore, the denominator of $pA_i(s)$ is always positive whenever $qA_i(s) > 0$.

The intuition behind PSP is an exclusion-compensation principle: you pay a price per unit which is the average of all other players' bid prices, each weighted by how much the allocation of that player is decreased by your bid. Equivalently, for each infinitesimal share of the resource, the player who is getting it pays the maximum amount that the player who is being denied it by him would have been willing to pay for it. The unit price paid by player i, pA_i, increases with qA_i in a manner similar to the income tax rate in a progressive tax system. For a fixed opponent profile s_{-i}, imagine player i is increasing qs_i, starting from 0. The first few units that player i gets will be taken away from the lowest clearing opponent (i.e. $m = \arg\min_j\{ps_j : qA_j > 0\}$), and player i will pay ps_m. When qA_m reaches 0, the subsequent units that player i gets will cost him $ps_{m'} > ps_m$, 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. This is widely known to have many the desirable properties [14, 9, 2], 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) and fairness.

The computational complexity of PSP is very low. A straightforward implementation would at worst, sort the bids in time $I \log I$, perform (3) in linear time, and (4) can be done in time I^2. Thus, the complexity of computing the allocations is $O(I^2)$.

3.1 Equilibrium of PSP

With players as defined in Section 2.2, the PSP rule has a number of nice properties, ultimately leading up to the existence of a fair Nash equilibrium.

Define the set of best replies to a profile s_{-i} of opponents bids: $S_i^*(s_{-i}) \stackrel{def}{=} \{s_i \in S_i(s_{-i}) : u_i(s_i; s_{-i}) \geq u_i(s_i'; s_{-i}), \forall s_i' \in S_i(s_{-i})\}$. Let $S^*(s) \stackrel{def}{=} \prod_i S_i^*(s_{-i})$. A Nash equilibrium is a fixpoint of S^*, i.e. a profile s^* such that $s_i^* \in S_i^*(s_{-i}^*), \forall i$. 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, the dominant 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 [3]. Thus, we take the existence of a Nash equilibrium as the definition stability.

Define the *partial* ordering $\geq$ on $\mathcal{S}$ to be the usual component-wise relation: $x = (x_1, x_2) \geq y = (y_1, y_2)$ if $x_1 \geq x_2$ and $y_1 \geq y_2$. Let $x \wedge y = (x_1 \wedge y_1, x_2 \wedge y_2)$.

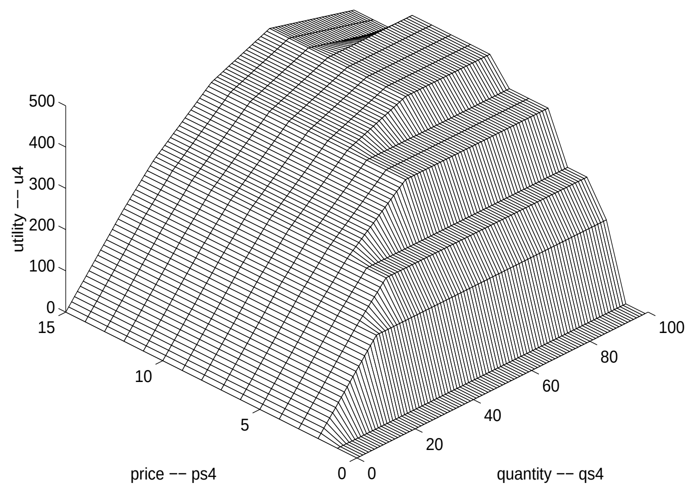

Figure 1: Utility $u_4(s_4)$ for $s_0 = (100, 1)$, $s_1 = (10, 2)$, $s_2 = (20, 4)$, $s_3 = (20, 7)$, $s_5 = (30, 12)$

The key property of PSP, the importance of which cannot be overstated, is that a player cannot do better than simply tell the truth, i.e. set $ps_i = \theta_i$. *Bidding his valuation is a dominant strategy.* This is known in the economics litterature as incentive compatibility [9].

Lemma 1 *(*Incentive compatibility*) For each player $i \in \mathcal{I}$, $\forall (s_i; s_{-i}) \in \mathcal{S}$,*

$$u_i((qs_i, \theta_i); s_{-i}) \geq u_i(s_i; s_{-i}). \tag{6}$$

Figure 1 shows the utility function of player 4, $u_4(s_4)$, in an PSP auction with $I = 5$ players, with s_{-4} fixed, and $\theta_4 = 10$. The plateaus correspond to the points where $qs_4 \geq \left[Q - \sum_{\{j:ps_j > ps_4\}} qA_j(s)\right]^+$, and $qA_4(s)$ can no longer be increased at that bid price – see (3). At bid prices $ps_4 > \theta_5$, the utility decreases when $qA_4 > Q - qs_5$, because after that point, each additional unit of resource is taken away from player 5, and thus costs $ps_5 = \theta_5$, which is more than θ_4 it's value to player i. Thus, each additional unit starts bringing negative utility. This is what discourages users from bidding above their valuation. Lemma 1 is illustrated by the fact that for any given quantity qs_4, the utility u_4 is maximized on the plane $ps_4 = \theta_4$.

The next result says that not only is the truthful strategy dominant, it is also always feasible, and thus there is always a "truthful best-reply". Let

$\mathcal{T}_i \stackrel{def}{=} \{s_i \in \mathcal{S}_i : ps_i = \theta_i\}$, the (unconstrained) set of player i's truthful bids, and $\mathcal{T} = \prod_i \mathcal{T}_i$. Define $T_i(s_{-i}) \stackrel{def}{=} \mathcal{T}_i \cap S_i(s_{-i})$, and $T(s) = \prod_i T_i(s_{-i})$. Let $R_i(s_{-i}) \stackrel{def}{=} \mathcal{T}_i \cap S_i^*(s_{-i})$ be the set of "truthful" best replies.

Lemma 2 (Existence of truthful best reply) *For each player $i \in \mathcal{I}$, $\forall s_{-i} \in \prod_{j \neq i} \mathcal{S}_j$,*

$$R_i(s_{-i}) \neq \emptyset.$$

This means that, thanks to the incentive compatibility property, we can restrict our attention to truthful strategies only, and still have feasible best replies. This forms a proper "truthful" subgame, where the strategy space is $\mathcal{T} \subset \mathcal{S}$, the feasible sets are $T_i(s_{-i}) \subset S_i(s_{-i})$, and the best replies are $R(s) \subset S^*(s)$. A fixpoint of R in $\mathcal{T}$ is a fixpoint of S^* in $\mathcal{S}$. Thus an equilibrium of the subgame is an equilibrium of the whole game.

Though u_i is neither continuous nor concave on $\mathcal{S}$ (see Figure 1), its restriction to $\mathcal{T}$ is continuous in s and concave in s_i. This makes the subgame a *convex game* for which the methodology for establishing equilibria is well-known [2].

Proposition 1 *(Nash equilibrium) In the auction game with the PSP rule given by (3) and (4), and players with utilities of the form (1), there exists a Nash equilibrium $s^* \in \mathcal{T}$.*

Even if one or any group of players are non-reactive to the others, the remaining players still have a Nash equilibirum. This is not surprising, since, in establishing Proposition 1, no assumption was made on the number of players or the amount Q of resources. This robustness characteristic of the stability of the mechanism is important for practical implementations of our auction in distributed systems where variable latencies may cause some players to react more slowly than others.

Corollary 1 *Suppose a player $m \in \mathcal{I}$ fixes his bid at $\sigma_m \in \mathcal{S}_m$. Then, for the remaining players, the game has a Nash equilibrium $\sigma^*_{-m} \in \mathcal{T}_{-m}$.*

3.2 Equilibrium Properties: Fairness and Efficiency

Fairness can be defined in a number of ways, the most intuitively appealing of which is the notion of an "envy-free" allocation [12]. An allocation $A(s)$ is said to be envy-free if $\forall i, j \in \mathcal{I}$,

$$u_i(A_i(s)) \geq u_i(A_j(s)). \quad (7)$$

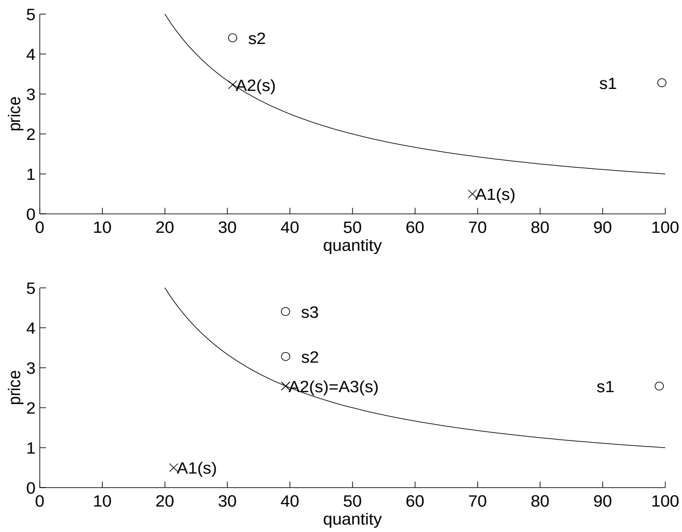

Figure 2: Examples of unfair (top) and fair (bottom) equilibrium

The next result provides a simple fairness test. It gives sufficient conditions, based only on the players' characterstics (b and θ) and the total amount of resource Q, which are all known *a-priori*.

To simplify notation, assume without loss of generality that $\theta_0 \leq \ldots \leq \theta_i \leq \theta_{i+1} \leq \ldots \leq \theta_I$. Thus, the equilibrium bid prices satisfy $ps_0^* \leq \ldots \leq ps_i^* \leq ps_{i+1}^* \leq \ldots \leq ps_I^*$.

Proposition 2 *(Fairness) If $\exists m \in \mathcal{I}$ such that*

$$\sum_{j>m} \frac{b_j}{\theta_m} + \frac{b_m}{\theta_{m-1}} \geq Q > \sum_{j>m} \frac{b_j}{\theta_m} \tag{8}$$

and, $\forall i > m$,

$$\sum_{\{j:j>m\}} \frac{b_j}{\theta_m} + \frac{\theta_i - \theta_m}{\theta_i - \theta_{m-1}} \frac{b_i}{\theta_m} \geq Q, \tag{9}$$

then there exists an equilibrium $s^* \in \mathcal{T}$ such that the allocation $A(s^*)$ is fair.

Figure 2 shows the equilibrium outcomes from the auction game being played on a distributed inter-active implementation on the World Wide Web in the Java programming language[5]. The outcomes illustrate Proposition 2,

[5]This paper has a companion web page [11], where a Java applet acting as a bidding agent enables anyone to play the auction against others interactively over the Internet.

for $Q = 100$, and $\theta_0 = 0.5$. In the first scenario (top), there are two players with valuations of $\theta_1 = 3.28$ and $\theta_2 = 4.4$, and budgets of $b_1 = b_2 = 100$, and $Q = 100$. Neither (8) nor (9) is satisfied, and the allocation turns out to be unfair: clearly player 2 will envy player 1, since 1 pays a higher price $pA_1(s) > pA_2(s)$, and gets a smaller quantity $qA_1(s) < qA_2(s)$. In the second scenario, a third player is added so that $\theta_1 = 2.54$, $\theta_2 = 3.28$ and $\theta_3 = 4.4$, and $b_1 = b_2 = b_3 = 100$. Both (8) and (9) are satisfied, and the allocation is fair.

The first condition (8) is increasingly likely to be satified, as the "demand" of an individual player becomes smaller relative to the total amount of resource Q. Indeed, if $b_i/\theta_j \ll Q$, $\forall i,j$, then (8) is approximately the same as $\sum_{j \geq m} \frac{b_j}{\theta_m} \geq Q > \sum_{j>m} \frac{b_j}{\theta_m}$, which is necessarily true for some m.

An efficiency measure for an allocation rule is a function of the form

$$m(A(s)) = \sum_{i \in \mathcal{I}} m_i(A_i(s)). \tag{10}$$

An allocation rule A is m-efficient at s if $m(A(s)) \geq m(A'(s))$ for all feasible A'. For example, as a user-centric measure of efficiency, one may take the sum of the utilities of the players $\sum u_i(A_i(s))$, in which case m-efficiency is equivalent to Pareto optimality. In the case of auctions, the most common measure of efficiency is the seller's revenue, which here is $\sum qA_i(s)pA_i(s)$. Even if the objective is not to raise money, it is natural for the designer to seek to get the most "value" from the given amount of resource Q. Thus we will focus on efficiency as measured by revenue.

Although ideally we would like A to be efficient at all s, the efficiency will primarily be measured at Nash equilibrium profiles. Different rules will in general lead to different equilibria. An allocation game with rule A and equilibrium s^* is efficient if $m(A(s^*)) \geq m(A'(s'^*))$ for any A' with NE s'^*, i.e. we compare the efficiency of each game's rule *at its own equilibirium*[6].

The highest revenue one can generate from any allocation would occur by getting a) all the players to pay their valuation, $pA_i(s) = \theta_i$, and b) getting the most out of player I (i.e. b_I, which is obtained when $qA_I(s) = b_I/\theta_I$), then player $I-1$, etc., until all of the resource has been sold. This would

[6] Note the difference between efficiency of the *game*, which is the equilibrium efficiency, and the efficiency of the *rule*, which is at any profile s. Thus a game with rule A may be efficient because at its NE it is better than all other games at their NE, but the allocation rule A itself may be inefficient (less efficient than others) at non-equilibrium points. Conversely a rule which is efficient at all s may lead to an equilibrium which is less efficient than the equilibria reached by other rules.

lead to a total revenue of

$$\max_m \sum_{i>m} b_i + \theta_m \left(Q - \sum_{i>m} b_i/\theta_i \right), \tag{11}$$

which will occur for $m = \max_n\{n : Q \leq \sum_{i\geq n} b_i/\theta_i\}$. Clearly, the PSP allocation rule will not achieve this, since the price paid is always strictly less than the bid, and it is not clear that there is any rule which will achieve it at equilibrium[7]. However, it is useful as a benchmark. If (8) is satisfied, the revenue generated at the equilibrium of Proposition 2 is

$$\sum_{i>m} b_i + \theta_{m-1} \left(Q - \sum_{i>m} b_i/\theta_m \right), \tag{12}$$

(this is true even if (9) does not hold and the equilibirium is not fair – see proof of Proposition 2 in Appendix A). But, as noted above, (8) can be assumed to hold when the game involves large numbers of "small" players. In thoses cases, the ratio ρ which we define as (12) over (11) is a good indicator of the relative efficiency. Figure 3 plots the efficiency ρ versus the offered load factor $\lambda = \frac{1}{Q}\sum_i b_i/\theta_i$. For each point on the curve, a new set of $I = 200$ players (large enough for the "small individual player" approximations to be reasonable) is generated with random θ_i uniformly distributed in the range $(0.5, 9.5)$, $b_i = 100$, $\forall i$, and $Q = \sum_i b_i/(\theta_i \lambda)$.

As can be seen from Figure 3, the PSP rule is reasonably efficient even for small demand, and tends to maximal efficiency as the offered load factor (or ratio of demand to supply) gets large.

4 Auctioning on a Network

4.1 Formulation

Suppose there is a set of resources $\mathcal{L} = \{1, \ldots, L\}$, of which the quantities are $Q = \{Q^1, \ldots, Q^L\}$, and as before, a set of players $\mathcal{I} = \{1, \ldots, I\}$. Bids are of the form $s_i = (q_i, p_i, r_i)$, where $r_i \in \{0, 1\}^L$.

The interpretation is that the resources $\mathcal{L}$ are links in a network, and s_i is a bid is for a "circuit" of bandwidth q_i which goes on a route[8] given by

[7] In the traditional auction of an indivisble object, even first price auctions, where you pay the price you bid, in general cannot, *at equilibrium*, raise higher revenues than a second price auction [7].

[8] We assume initially that the allocation rule is not concerned about wether r_i forms a continuous path or not – route means any arbitray set of links.

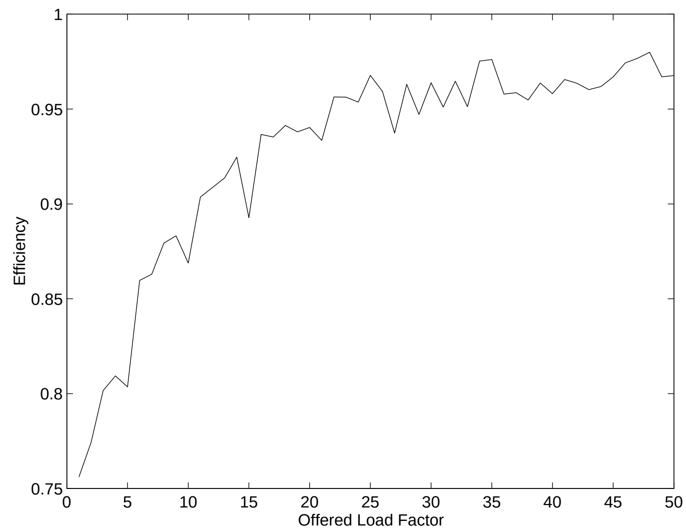

Figure 3: ρ vs. "offered load" λ.

r_i as follows: link l is in player i's route if and only if $r_{i,l} = 1$.

Let $rs = (r_1^T, \ldots, r_L^T)^T$. The i-th row of rs is user i's route r_i . Then an allocation rule F,

$$
\begin{array}{rccc}
F: & [0,Q]^I \times [0,\infty)^I \times \{0,1\}^{L\times I} & \longrightarrow & [0,Q]^I \times [0,\infty)^I \times \{0,1\}^{L\times I} \\
 & s = (qs, ps, rs) & \longmapsto & F(s) = (qF(s), pF(s), rF(s)).
\end{array}
$$

is feasible if $\forall s$,

$$
\begin{aligned}
qF(s)rs &\leq Q && (13)\\
(qF(s), pF(s)) &\leq (qs, ps) && (14)\\
rF(s) &= rs. && (15)
\end{aligned}
$$

When bids are allowed for any "route" r_i, an efficient feasible auction rule is computationally hard, unlike the case of a single resource, where (3) involves only simple arithmetic operations. The difficulty arises from the need to avoid situations where a bidder using a single link could block a bidder using many links. An efficient rule would be such that $\forall s$, computing $F(s)$ is equivalent to solving

$$\max_{F(s)} m(F(s)) \qquad (16)$$

subject to (13)-(15).

This means solving a large optimization problem[9] just to compute an allocation from a bid profile, which would make it impractical for, say, making pricing and admission control decisions within a realistic call setup time. Also, there is no guarantee that such an "optimal" rule would even be stable, i.e. lead to a Nash equilibrium.

4.2 Sharing on a tree

If the set of resources $\mathcal{L}$ has some "structure", and we restrict bids to certain specific subsets of resources (i.e. routes), we can devise simple rules that are efficient. In other words, we are still solving (16), but the matrix rs is restricted to having some special structure that makes it easy.

Consider the case when $\mathcal{L}$ is the set of edges of a tree, and bids are only allowed for routes which begin at the root of the tree and form a continuous path to some other vertex (not necessarily a leaf). Then we can use a simple single-resource allocation rule at each individual edge and combine them by a dynamic programing type of algorithm to obtain an allocation rule F which is efficient for the whole tree, and computationally simple.

Let A be some allocation rule for a single resource as in Section 2. To distinguish which resource we are applying the rule to, we will write $A(Q^l, .)$ instead of simply $A(.)$ when the rule is being applied on resource l of which there is a quantity Q^l. Also, for $s = (qs, pq, rs)$, by abuse of notation we will write simply $A(Q^l, s)$ although, since the routes rs are not required for the single-resource allocation, we should write $A(Q^l, (qs, ps))$.

For any $\mathcal{J} \subset \mathcal{I}$, let $1_{\mathcal{J}}$ be the operator which deletes all rows i such that $i \notin \mathcal{J}$, thus $1_{\mathcal{J}}(s)$ is the profile of the subset of players $\mathcal{J}$.

The allocation rule F for the tree is given by the algorithm in Table 1.

The algorithm starts from the leaves, and works back toward the root, doing an allocation for each edge using the single resource rule F, among the subset of players whose route includes that edge. Note that, as the algorithm backtracks along a player's route, at step 5, *his "bid" for the next edge is automatically reduced to his actual allocation on the current edge*. Thus the allocations are decreasing, i.e. his "pipe" is "thining" as it gets closer to the source. Since the useable "thickness" of a player's "pipe" is only the thickness of his thinnest link allocation, it may seem like the downstream allocations are wasteful of resources. However, it is

[9] With the additional restricition that $qF_i(s) = q_i$ or 0, this optimization is in fact N-P complete – see [5].

- `input:`
  `The vertices/nodes` $\mathcal{V}$`, the edges/links` $\mathcal{L}$`,` $Q = (Q^1, \ldots, Q^L)$`,` $s = \{(q_i, p_i, r_i)\}_{i=1}^I$.
- `1.`
  - `Pick a leaf` $v \in \mathcal{V}$`; let` $l_v \in \mathcal{L}$ `be the (unique) edge going into` v`.`
  - `Let` $\mathcal{I}(l_v) = \{i \in \mathcal{I} : r_{i,l_v} = 1\}$`, the subset of players which use link` l_v`.`
- `2. Compute` $\underline{s} = F(Q^{l_v}, 1_{\mathcal{I}(l_v)}(s))$`, i.e. perform the allocation on resource` l_v`.`
- `3. For each` $i \in \mathcal{I}(l_v)$`, set` $s_i = \underline{s}_i$`.`
- `4. Set` $\mathcal{V} = \mathcal{V} - \{v\}$`, and` $\mathcal{L} = \mathcal{L} - \{l_v\}$`, i.e. delete that leaf and that edge from the tree.`
- `5. If` $\mathcal{L} \neq \emptyset$ `goto 1, else set` $F(s) = s$ `and we're done.`

Table 1: **Algorithm of Tree Allocation Rule**

not so, because as you move closer to the root of the tree along a given path, the set of players is increasing, so the resources that appear wasted downstream are not wasted because there would have been no one else to allocate them to. More precisely, if the single resource rule A is efficient, then this algorithm consitutes an allocation rule F which is efficient for the tree. This is expressed formally by the following result.

Proposition 3 *Let A be a single-resource allocation rule, and F the corresponding Tree Allocation Rule. Suppose A is monotone with respect to the efficiency measure m, i.e. $\forall s, s'$*

$$m(s) \geq m(s') \Rightarrow m(A(s)) \geq m(A(s')). \tag{17}$$

Then, $\forall s$,

A m-efficient at $s \Rightarrow F$ m-efficient at s.

It can be shown that, with revenue as the measure of efficiency m, the

PSP rule given in Section 3 satisfies 17, and thus is suitable for extension to multiple resources using the tree algorithm.

4.3 Practical Considerations

Auctions, as formulated here, are applicable in a setting where the resources are arbitrarily divisible, and these divisions enforceable. Thus, this approach would fit particularly well in a system where arbitrary amounts of resources can be requested and obtained, rather than systems with a small fixed menu of service classes to choose from. In the jargon of network engineering, this translates to user controlled quality of service, plus packet scheduling algorithms which can assign and guarantee an arbitrary share of, say, the bandwidth to each individual flow.

It is natural to assume that bidding would be on a per flow basis, or per unit time on an appropriate time scale. Not only is active bidding on the time scale of packets probably not feasible, but the relationship between the resource allocation of a single packet and the user-perceived quality of service is a daunting, if not impossible, thing to estimate [13]. On the flow or call level though, this is certainly possible, particularly if the "user" is a software bidding agent, possibly embedded into the applications, which develops over time and with human feedback, an accurate relationship Θ translating resources into value or perceived QoS.

5 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 currenlty used. [7]

Though the most natural application of an auction mechanism would be for pricing, it is not necessarily the only one. Auctions can also be used in an environment where there is no "real money" involved (such as a non-commercial network, or a corporate intranet) with fictitious budgets of "funny money" assigned to many intelligent agents according to the importance of their tasks, and using the auction mechanism as a means for these tasks to share resources without central control.

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 a flexible and intuitive model of user preferences, constitutes a stable resource allocation mechanism. We presented some sufficient conditions for the equilibrium allocations to be rigorously fair, and illustrated how the allocations tend to optimal efficiency with increasing load. We have presented an algorithm for applying an auction mechanism on a network of resources, and in the particular case of a tree topology, showed that the efficiency is perserved. Even thoug 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 of situations,

Directions for future analytical work on the basic mechanism include finding weaker conditions for fairness, as well as considering off equilibrium play, learning strategies, and evolutionary behaviour which can emerge from repeated inter-action between the same players. Much of the intuition behind the mechanism design and the analysis in this work came from experiments done on an inter-active distributed implementation of this auction game on the World Wide Web, using the Java programming language [11].

References

[1] Agorics Technical Report ADd004P. Real-time video delivery with market-based resource allocation. Technical report, Sun Microsystems Labs, SunConnect, Agorics, Inc., 1994.

[2] D. Fudenberg and J. Tirole. *Game Theory.* The MIT Press, 1991.

[3] E. Kalai and E. Lehrer. Rational learning leads to nash equilibrium. *Econometrica*, 61(5):1019–1045, September 1993.

[4] Y. A. Korilis, A. A. Lazar, and A. Orda. Architecting noncooperative networks. *IEEE J. Select. Areas Commun.*, 13(7), September 1995.

[5] R. M. Harstad M. H. Rothkopf, A. Pekec. Computationally manageable combinatorial auctions. Technical Report 95-09, DIMACS, Rutgers University, 1995.

[6] J. K. MacKie-Mason and H. R. Varian. Pricing the internet. In *Public Access to the Internet*, 1993. ftp://gopher.econ.lsa.umich.edu/pub/Papers/Pricing_the_Internet.ps.Z.

[7] 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.

[8] 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.

[9] R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, February 1981.

[10] J. F. Rosenschein and G. Zlotkin. *Rules of Encounter*. The MIT Press, 1994.

[11] N. Semret. TREX-The Resource EXchange, 1995-96. http://argo.ctr.columbia.edu:1024/~nemo/Trex/trex.html.

[12] S. Shenker. Making greed work in networks: A game-theoretic analysis of switch service disciplines. In *Proc. of the ACM SIGCOMM*, 1994.

[13] 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.

[14] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, 16, 1961.

[15] C. A. Waldspurger, T. Hogg, B. A. Huberman, J. O. Kephart, and W. S. Stornetta. Spawn: A distributed computational economy. *IEEE Trans. on Software Engineering*, 18(2):103–117, February 1992.

A Proofs

We begin by stating some very basic properties of the the PSP rule given by (3), (4).

Lemma 3 *For each player i, $\forall s, s' \in \mathcal{S}$,*

$$\begin{aligned} s_i \le s_i' \quad &\Rightarrow \quad qA_i(s_i; s_{-i}) \le qA_i(s_i'; s_{-i}) \\ &\qquad \textit{and } pA_i(s_i; s_{-i}) \le pA_i(s_i'; s_{-i}), \qquad (18) \\ s_{-i} \le s_{-i}' \quad &\Rightarrow \quad qA_i(s_i; s_{-i}) \ge qA_i(s_i; s_{-i}') \\ &\qquad \textit{and } pA_i(s_i; s_{-i}) \le pA_i(s_i; s_{-i}'). \qquad (19) \end{aligned}$$

Also,

$$s_i \le s_i' \in S_i(s_{-i}) \quad \Rightarrow \quad s_i \in S_i(s_{-i}), \qquad (20)$$

$$s_{-i} \le s_{-i}' \quad \Rightarrow \quad S_i(s_{-i}) \supseteq S_i(s_{-i}'). \qquad (21)$$

Proof: All the assertions can be derived by straightforward manipulations from the definitions (3), (4) and (2). □

Now, we give the formula for the "derivative" of a player's utility with respect to his own bid (for a fixed opponent profile), which we will use in all subsequent results.

Lemma 4 *(*Derivative*) For all profiles $s_{-i} \in \mathcal{S}_{-i}$, $\forall s_i' \in \mathcal{S}_i$,*

$$u_i(s_i'; s_{-i}) - u_i(s_i; s_{-i}) = \sum_{j \ne i} (\theta_i - ps_j) \left[qA_j(s_i; s_{-i}) - qA_j(s_i'; s_{-i})\right]. \qquad (22)$$

Proof: It can be easily seen from (5) that

$$qA_i(s) = \sum_{j \ne i} \left[qA_j(0; s_{-i}) - qA_j(s_i; s_{-i})\right], \qquad (23)$$

which is also intuitively clear[10]. Thus, substituting (4) and (23) into (1), we get

$$u_i(s_i; s_{-i}) = \theta_i \sum_{j \ne i} \left[qA_j(0; s_{-i}) - qA_j(s_i; s_{-i})\right] - \sum_{j \ne i} ps_j \left[qA_j(0; s_{-i}) - qA_j(s_i; s_{-i})\right].$$

[10] Whatever player i gets is "taken away" from some other players who would would have gotten it if player i was not there. Note that if there is no player 0, who is always willing to take up to Q, (23) has an extra term on the right hand side, $(Q - \sum_{j \ne i} q_j)^+ \wedge qs_i$, to account for the possible "leftover" that player i could get without taking anything away from the others.

Subtracting the analogous expression for $u_i(s'_i; s_{-i})$ from this last equation gives the desired result. $\square$

Next is the incentive compatibility property of Lemma 1, which we repeat here for convenience.

Lemma 5 *(*Incentive compatibility*) For each player $i \in \mathcal{I}$, $\forall (s_i; s_{-i}) \in \mathcal{S}$,*

$$u_i((qs_i, \theta_i); s_{-i}) \geq u_i(s_i; s_{-i}). \tag{24}$$

Proof: From Lemma 4,

$$u_i((qs_i, \theta_i); s_{-i}) = u_i(s) + \sum_{j \neq i} (\theta_i - ps_j) \left[qA_j(s_i; s_{-i}) - qA_j((qs_i, \theta_i); s_{-i}) \right].$$

If $ps_i \leq \theta_i$, then from (19) it follows that $qA_j(s_i; s_{-i}) - qA_j((qs_i, \theta_i); s_{-i}) \geq 0$ for all j. It is clear from (3) that the quantity qA_j allocated to the opponents whose bids remain above or below the bid price of player i remain unchanged. In particular, $qA_j(s_i; s_{-i}) - qA_j((qs_i, \theta_i); s_{-i}) = 0$ for all j such that $ps_j > \theta_i$. Thus, the terms in the summation are all ≥ 0.

Similarly, if $ps_i > \theta_i$, then $qA_j(s_i; s_{-i}) - qA_j((qs_i, \theta_i); s_{-i}) \leq 0$ for all j, and in particular, $= 0$ for j such that $ps_j < \theta_i$. Thus, the terms in the summation are all ≥ 0. $\square$

We now repeat and prove Lemma 2.

Lemma 6 *(*Existence of truthful best reply*) For each player $i \in \mathcal{I}$, $\forall s_{-i} \in \prod_{j \neq i} \mathcal{S}_j$,*

$$R_i(s_{-i}) \neq \emptyset.$$

Proof: Fix $s_{-i} \in \prod_{j \neq i} \mathcal{S}_|$. Now $(0, p_i) \in S_i(s_{-i})$ for any $p_i \geq 0$, therefore $S_i(s_{-i}) \neq \emptyset$, hence $S_i^*(s_{-i}) \neq \emptyset$. Pick any $s_i \in S_i^*(s_{-i})$.

As shorthand notation, let $u_i(.) \equiv u_i(.; s_{-i})$, $qA_i(.) \equiv qA_i(.; s_{-i})$, and $pA_i(.) \equiv pA_i(.; s_{-i})$.

- If $\theta_i \leq ps_i$, then by (18), $0 \leq qA_i((qs_i, \theta_i)) \leq qA_i(s_i)$ and $0 \leq pA_i((qs_i, \theta_i)) \leq pA_i(s_i)$, hence

  $$qA_i((qs_i, \theta_i))pA_i((qs_i, \theta_i)) \leq qA_i(s_i)pA_i(s_i) \leq b_i.$$

  Therefore, $(qs_i, \theta_i) \in S_i(s_{-i})$, and by Lemma 1, $(qs_i, \theta_i) \in S_i^*(s_{-i})$. Now trivially, $(qs_i, \theta_i) \in \mathcal{T}_i$, therefore $(qs_i, \theta_i) \in R_i(s_{-i})$.

- If $\theta_i > ps_i$, let $q'_i = qA_i(s_i)$. From (18), $q'_i \geq qA_i((q'_i, \theta_i)) \geq qA_i((q'_i, ps_i)) = q'_i$. Therefore, we have equality throughout, and $qA_i((q'_i, \theta_i)) = qA_i(s_i)$. Thus

$u_i((q_i', ps_i)) = u_i(s_i)$. Using Lemma 1, we get $u_i((q_i', \theta_i)) \geq u_i((q_i', ps_i)) = u_i(s_i)$, that is,

$$\theta_i qA_i((q_i', \theta_i)) - qA_i((q_i', \theta_i))pA_i((q_i', \theta_i)) \geq \theta_i qA_i(s_i) - qA_i(s_i)pA_i(s_i),$$

which implies,

$$qA_i((q_i', \theta_i))pA_i((q_i', \theta_i)) \leq qA_i(s_i)pA_i(s_i) \leq b_i,$$

where the last inequality is from the fact that $s_i \in S_i(s_{-i})$. Therefore, $(q_i', \theta_i) \in S_i(s_{-i})$, and by Lemma 1, $(q_i', \theta_i) \in S_i^*(s_{-i})$. Now trivially, $(q_i', \theta_i) \in \mathcal{T}_i$, therefore $(q_i', \theta_i) \in R_i(s_{-i})$.

Thus, $R_i(s_{-i}) \neq \emptyset$, which completes the proof. □

Before proceeding to the main result, we need to verify that key rather technical conditions are satisfied by the game we have defined. The reader will note in the proof of Lemma 7 that we will assume that player 0, the auctioner, bids a strictly positive (but possibly very small) price $ps_0 > 0$ for the whole quantity Q. This in effect amounts to imposing a small minimum price (or "reservation fee") on each unit of resource. Lemma 7 and Proposition 1 show that this is sufficient to ensure the existence of an equilibrium. Though we have not yet proved this, from some preliminary studies, we conjecture that this "reservation fee" is a necessary condition for an equilibriumto exist.

Lemma 7 *(Continuity properties) In the subgame with strategy space $\mathcal{T}$, for each player i, the utility function u_i is concave in s_i, continuous in s, and has a continuous maximum, i.e.*

$$\begin{array}{rccc} u_i^* : & \mathcal{T}_{-i} & \longrightarrow & [0, \infty) \\ & s_{-i} & \longmapsto & u_i^*(s_{-i}) = \max_{r_i \in \mathcal{T}_i(s_{-i})} u_i(r_i; s_{-i}) \end{array}$$

is continuous.

Proof: From Lemma 4,

$$u_i(s) = \sum_{j \neq i} (\theta_i - ps_j) \left[qA_j(0; s_{-i}) - qA_j(s_i; s_{-i}) \right].$$

From (3), with ps fixed, qA_j is continuous in qs. Therefore u **is continuous in** s **on** $\mathcal{T}$.

Now we fix $s_{-i} \in \prod_{j \neq i} \mathcal{T}_j$, we will show that u_i is concave in s_i on $\mathcal{T}_i$. Recall that $s_i \in \mathcal{T}_i \Rightarrow ps_i = \theta_i$, therefore we need only show that u_i is concave in qs_i. Let

$$Q_j = Q - \sum_{ps_k > ps_j, k \neq i} qs_k,$$

the "leftovers" after those who bid higher than ps_j (except player i) have been served. For convenience, define the labels $l(0) = 0$, and

$$l(n) = \arg \min_{k \notin \{l(1), l(2), \ldots, l(n-1)\}} ps_k,$$

for $0 < n \leq I$. This is simply a relabeling of the players so that: $ps_{l(n)} \leq ps_{l(n+1)}$, and $Q_{l(n)} \leq Q_{l(n+1)}$, $0 \leq n < I$. In other words, $l(n) = j$ means ps_j is the n-th lowest bid price. Let $m = \min\{n \neq i : Q_{l(n)} > 0\}$, the lowest player[11] who can get a non-zero allocation. Let $M = l^{-1}(i) - 1$, the highest player just below player i. Now (4) can be written as:

$$pA_i(s) = \begin{cases} ps_{l(m)} & \text{if } 0 < qs_i \leq Q_{l(m)} \\ \frac{1}{qs_i}\left[ps_{l(m)}Q_{l(m)} + ps_{l(m+1)}(qs_i - Q_{l(m)})\right] & \text{if } Q_{l(m)} < qs_i \leq Q_{l(m+1)} \\ \vdots & \\ \begin{array}{l} \frac{1}{qs_i}\left[ps_{l(m)}Q_{l(m)} + ps_{l(m+1)}Q_{l(m+1)} + \ldots \right. \\ \qquad \left. + ps_{l(M)}(qs_i - Q_{l(M-1)})\right] \end{array} & \text{if } Q_{l(M-1)} < qs_i \leq Q_{l(M)} \\ \begin{array}{l} \frac{1}{Q_{l(M)}}\left[ps_{l(m)}Q_{l(m)} + ps_{l(m+1)}Q_{l(m+1)} + \ldots \right. \\ \qquad \left. + ps_{l(M)}(Q_{l(M)} - Q_{l(M-1)})\right] \end{array} & \text{if } Q_{l(M)} < qs_i, \end{cases}$$

and (3) as

$$qA_i(s) = \begin{cases} qs_i & \text{if } 0 < qs_i \leq Q_{l(M)} \\ Q_{l(M)} & \text{if } Q_{l(M)} < qs_i. \end{cases}$$

Substituting into the definition of the utility (1) and differentiating with respect to qs_i, we get

$$\frac{du_i}{d(qs_i)} = \begin{cases} \theta_i - ps_{l(m)} & \text{if } 0 < qs_i \leq Q_{l(m)} \\ \theta_i - ps_{l(m+1)} & \text{if } Q_{l(m)} < qs_i \leq Q_{l(m+1)} \\ \vdots & \\ \theta_i - ps_{l(M)} & \text{if } Q_{l(M-1)} < qs_i \leq Q_{l(M)} \\ 0 & \text{if } Q_{l(M)} < qs_i. \end{cases}$$

Since $ps_{l(n)} \leq ps_{l(n+1)} \leq ps_i = \theta_i$, and $Q_{l(n)} \leq Q_{l(n+1)}$, $0 \leq n < I$, $du_i/d(qs_i)$ is decreasing but non-negative, i.e. u_i **is concave and non-decreasing in** qs_i on $\mathcal{T}_i$. In particular, we see that u_i will reach it's maximum for all points $qs_i \geq Q_{l(M)}$.

Now consider the constraint set: $T_i(s_{-i}) = \{s_i \in \mathcal{T}_i : qA_i(s_i; s_{-i})pA_i(s_i; s_{-i}) \leq b_i\}$. Let $B(s_{-i}) = qA_i((Q_{l(M)}, \theta_i), s_{-i})pA_i((Q_{l(M)}, \theta_i), s_{-i})$.

- If $B(s_{-i}) \leq b_i$, then $(Q_{l(M)}, \theta_i) \in T_i(s_{-i})$ and

$$u_i^*(s_{-i}) = \max_{r_i \in T_i(s_{-i})} u_i(r_i; s_{-i}) = u_i((Q_{l(M)}, \theta_i); s_{-i}). \tag{25}$$

  Since $Q_{l(M)}$ is a continuous function of $s_{-i} \in \mathcal{T}_{-i}$ (as are indeed all Q_j), $u_i^*(s_{-i})$ is continuous on $\{s_{-i} \in \mathcal{T}_{-i} : B(s_{-i}) \leq b_i\}$.

[11] By "lowest player", we mean the player with the lowest bid price.

- If $B(s_{-i}) > b_i$, then $(Q_{l(M)}, \theta_i) \notin T_i(s_{-i})$, and since by (18) qA_i and pA_i are both non-decreasing in s_i, we have for any $s_i \in \mathcal{T}_i$, $qs_i \geq Q_{l(M)} \Rightarrow s_i \notin T_i(s_{-i})$, i.e. $T_i(s_{-i}) \subset [0, Q_{l(M)}] \times \{\theta_i\}$. But on that set, $qA_i(s) = qs_i$ is strictly increasing, pA_i is positive non-decreasing, and both are continuous, therefore the product $qA_ipA_i(.; s_{-i})$ is strictly increasing and continuous. In fact, it's derivative is

$$\frac{d(qA_ipA_i)}{d(qs_i)} = \begin{cases} ps_{l(m)} & \text{if } 0 < qs_i \leq Q_{l(m)} \\ ps_{l(m+1)} & \text{if } Q_{l(m)} < qs_i \leq Q_{l(m+1)} \\ \vdots & \\ ps_{l(M)} & \text{if } Q_{l(M-1)} < qs_i \leq Q_{l(M)}, \end{cases}$$

which is always $> ps_0 > 0$. Thus the inverse function $(qA_ipA_i(s_{-i}))^{-1}(.)$ is well-defined, continuous and increasing as a function on $[0, B(s_{-i})]$ to $[0, Q_{l(M)}]$, and $T_i(s_{-i}) = (qA_ipA_i(s_{-i}))^{-1}([0, b_i])$.

Now define $q_i^{max}(s_{-i}) = \sup T_i(s_{-i})$. Since $T_i(s_{-i})$ is compact, $(q_i^{max}(s_{-i}), \theta_i) = \max T_i(s_{-i}) = (qA_ipA_i(s_{-i}))^{-1}(b_i)$. We claim $q_i^{max}(.)$ is continuous at all s_{-i} in its domain $\{s_{-i} \in \mathcal{T}_{-i} : B(s_{-i}) > b_i\}$. Suppose the contrary, i.e. $\exists \epsilon > 0, \forall \delta > 0, \exists s'_{-i}$ such that $\|s_{-i} - s'_{-i}\| < \delta$ and $|q_i^{max}(s_i) - q_i^{max}(s'_i)| \geq \epsilon$. Then, by the fundamental theorem of calculus,

$$|qA_ipA_i(q_i^{max}(s_{-i}), s'_{-i}) - qA_ipA_i(q_i^{max}(s'_{-i}), s'_{-i})| \geq \epsilon ps_0.$$

Now since $qA_ipA_i(q_i^{max}(s'_{-i}), s'_{-i}) = b_i = qA_ipA_i(q_i^{max}(s_{-i}), s_{-i})$, we have

$$|qA_ipA_i(q_i^{max}(s_{-i}), s'_{-i}) - qA_ipA_i(q_i^{max}(s_{-i}), s_{-i})| \geq \epsilon ps_0.$$

Since qA_ipA_i is continuous in it's second argument, we can pick $\delta > 0$ small enough that $|qA_ipA_i(q_i^{max}(s_{-i}), s'_{-i}) - qA_ipA_i(q_i^{max}(s_{-i}), s_{-i})| < \epsilon ps_0$. Thus we have the desired contradiction, which means $q_i^{max}(.)$ is continuous.

Since u_i is increasing in qs_i, it's maximum in $T(s_{-i})$ will be achieved at $q_i^{max}(s_{-i})$. Thus

$$u_i^*(s_{-i}) = \max_{r_i \in T_i(s_{-i})} u_i(r_i; s_{-i}) = u_i((q_i^{max}(s_{-i}), \theta_i); s_{-i}), \qquad (26)$$

is continuous on $\{s_{-i} \in \mathcal{T}_{-i} : B(s_{-i}) > b_i\}$.

Now, we "stitch" (25) and (26) together, and verify that continuity holds across the boundary $\{s_{-i} \in \mathcal{T}_{-i} : B(s_{-i}) = b_i\}$, i.e. we want

$$x = lim_{B(s_i) \searrow b_i} q_i^{max}(s_{-i}) = Q_{l(M)}$$

Since $(q_i^{max}(s_{-i}), \theta_i) = (qA_ipA_i(s_{-i}))^{-1}(b_i)$ is continuous on $\{s_{-i} \in \mathcal{T}_{-i} : B(s_{-i}) > b_i\}$, we have

$$\begin{aligned} (x, \theta_i) &= lim_{B(s_i) \searrow b_i} (qA_ipA_i(s_{-i}))^{-1}(B(s_i), \theta_i) \\ &= (Q_{l(M)}, \theta_i). \end{aligned}$$

where the last line follows from the definition $B(s_{-i}) = qA_i pA_i((Q_{l(M)}, \theta_i), s_{-i})$. Thus u_i^* is continuous at all $s_{-i} \in \mathcal{T}_i$. $\square$

The main result of Section 3 is Proposition 1, which is the following.

Proposition 4 *(Nash equilibrium) In the auction game with the PSP rule given by (3) and (4), and players with utilities of the form (1), there exists a Nash equilibrium $s^* \in \mathcal{T}$.*

Proof: First we verify that $R(.)$ **has a closed graph**: if $(s(n), t(n)) \to (s, t)$ with $s(n) \in \mathcal{T}$ and $t(n) \in R(s(n))$, $\forall n$, then $t \in R(s)$. Suppose the contrary. Then, for some i, t_i is a sub-optimal reply to s_{-i}, i.e.

$$\begin{aligned} u_i^*(s_{-i}) &= \max_{r_i \in T_i(s_{-i})} u_i(r_i; s_{-i}) \\ &> u_i(t_i; s_{-i}) = \lim_{n\to\infty} u_i(t_i(n); s_{-i}(n)) \\ &= \lim_{n\to\infty} \max_{r_i \in T_i(s_{-i}(n))} u_i(t_i(n); s_{-i}(n)) \\ &= \lim_{n\to\infty} u_i^*(s_{-i}(n)), \end{aligned}$$

which contradicts the continuity of u_i^*.

From Lemma 2, $R(s)$ **is non-empty** for all s.

Now we show that, for all s, $R(s)$ **is a convex set**: if $r, r' \in R(s)$, then given any $\lambda \in [0,1]$, $\lambda r + (1-\lambda) r = \overline{r} \in R(s)$. By Lemma 7, for each i, $u_i(., s_{-i})$ is concave in it's first argument, therefore

$$\begin{aligned} u_i(\overline{r}_i; s_{-i}) &\geq \lambda u_i(r_i; s_{-i}) + (1-\lambda) u_i(r_i'; s_{-i}) \\ &= \lambda u_i^*(s_{-i}) + (1-\lambda) u_i^*(s_{-i}) = u_i^*(s_{-i}) \end{aligned}$$

which implies $\overline{r}_i \in R_i(s_{-i})$. Since this is true for all i, we have $\overline{r} \in R(s)$.

Clearly, since $\mathcal{T} = \prod_i \{[0, Q] \times \{\theta_i\}\}$, $\mathcal{T}$ **is compact, convex and non-empty**.

Now, by Kakutani's fixed-point theorem (see [2]), the four statements in bold font together imply that there exists a fixed-point $s^* \in R(s^*)$. $\square$

We now repeat and prove Corollary 1.

Corollary 2 *Suppose a player $m \in \mathcal{I}$ fixes his bid at any $\sigma_m \in \mathcal{S}_m$. Then, for the remaining players, the game has a Nash equilibrium $\sigma_{-m}^* \in \mathcal{T}_{-m}$.*

Proof: We simply verify that each step of the proof of Proposition 1 holds for the subset of players $\mathcal{I} - \{m\}$, when s_m is fixed.

With s_m fixed, u_i^* is simply restricted to a subset of it's domain, therefore it is still continuous in s_{-i}, i.e. u_i has a continuous maximum.

The proof of Lemma 2 does not rely on any particular choice of s_{-i}, thus, the set of best replies is non-empty.

Finally, the concavity of u_i as a function of $s_i \in \mathcal{T}_i$ does not rely on any particular choice of s_{-i}, so it still holds for all $i \in \mathcal{I} - \{m\}$.

Thus the best reply correspondence R for the subset of players $\mathcal{I} - \{m\}$ has a fixed-point. $\square$

To simplify notation, assume without loss of generality that $\theta_0 \leq \ldots \leq \theta_i \leq \theta_{i+1} \leq \ldots \leq \theta_I$. Thus, the equilibrium bid prices satisfy $ps_0^* \leq \ldots \leq ps_i^* \leq ps_{i+1}^* \leq \ldots \leq ps_I^*$.

Proposition 5 *(Fairness) If $\exists m \in \mathcal{I}$ such that*

$$\sum_{j>m} \frac{b_j}{\theta_m} + \frac{b_m}{\theta_{m-1}} \geq Q > \sum_{j>m} \frac{b_j}{\theta_m} \tag{27}$$

and, $\forall i > m$,

$$\sum_{\{j:j>m\}} b_j + \frac{\theta_i - \theta_m}{\theta_i - \theta_{m-1}} b_i \geq \theta_m Q, \tag{28}$$

then there exists an equilibrium $s^* \in \mathcal{T}$ such that the allocation $A(s^*)$ is fair.

Proof: Set $s_m = (Q, \theta_m)$. By Corollary 1, the other players have an equilibrium $s_{-m}^* \in \mathcal{T}_{-m}$. By Lemma 1, s_m is always a best reply provided it is feasible.

It is easy to see from (3) and (4) that since $qs_m = Q$, $pA_i(s_m; s_{-m}^*) \geq \theta_m$, therefore $qA_i(s_m; s_{-m}^*) \leq b_i/\theta_m$, $\forall i > m$. Hence, $qA_m(s_m; s_{-m}^*) \geq Q - \sum_{i>m} b_i/\theta_m > 0$, by the second inequality of (27). But that implies $pA_i(s_m; s_{-m}^*) \leq \theta_m$. Thus, $pA_i(s_m; s_{-m}^*) = \theta_m$, and $qA_m(s_m; s_{-m}^*) = Q - \sum_{i>m} b_i/\theta_m$. Also, since $s_{-m}^* \in \mathcal{T}_{-m}$, we have $pA_m(s_m; s_{-m}^*) \leq \theta_{m-1}$. Therefore,

$$pA_m(s_m; s_{-m}^*) qA_m(s_m; s_{-m}^*) \leq \theta_{m-1} \left(Q - \sum_{i>m} b_i/\theta_m \right) \leq b_m,$$

using the first inequality of (27). Thus, s_m is feasible, hence a best reply to s_{-m}^*, which means

$$s^* = (s_m; s_{-m}^*) = \begin{pmatrix} b_1/\theta_m & \theta_1 \\ b_2/\theta_m & \theta_2 \\ \vdots & \\ Q & \theta_m \\ \vdots & \\ Q & \theta_0 \end{pmatrix} \tag{29}$$

is a Nash equilibrium, yielding an allocation

$$A(s^*) = \begin{pmatrix} b_1/\theta_m & \theta_m \\ b_2/\theta_m & \theta_m \\ \vdots & \\ Q - \sum_{i>m} b_i/\theta_m & \theta_{m-1} \\ 0 & 0 \\ \vdots & \\ 0 & 0 \end{pmatrix}. \tag{30}$$

Now, we show that $A(s^*)$ is fair (envy-free). The players $j < m$ cannot be envied by any other player, because $A_j(s^*) = (0,0)$.

Moreover, $j < m$ cannot envy any other player $i \geq m$ because $pA_i(s^*) \geq \theta_m \geq \theta_j \Rightarrow u_j(A_i(s^*)) \leq 0$.

Any two players $i, i' > m$ cannot envy one another, because they pay the same unit price θ_m and they are all getting the maximum qA_i for their budget.

Can m envy a player $i > m$? No, because $pA_i(s^*) = \theta_m \Rightarrow u_m(A_i(s^*)) = 0$. Can $i > m$ envy m? No because, from (28),

$$\begin{aligned} \sum_{\{j: j>m\}} \frac{b_j}{\theta_m} + \frac{\theta_i - \theta_m}{\theta_i - \theta_{m-1}} \frac{b_i}{\theta_m} &\geq Q \\ \Leftrightarrow (\theta_i - \theta_m)\frac{b_i}{\theta_m} &\geq \left(Q - \sum_{\{j: j>m\}} \frac{b_j}{\theta_m} \right)(\theta_i - \theta_{m-1}) \\ &\Leftrightarrow u_i(A_i(s^*)) \geq u_i(A_m(s^*)). \end{aligned}$$

□

Proposition 6 *Let A be a single-resource allocation rule, and F the corresponding Tree Allocation Rule. Suppose A is monotone with respect to the efficiency measure m, i.e. $\forall s, s'$*

$$m(s) \geq m(s') \Rightarrow m(A(s)) \geq m(A(s')). \tag{31}$$

Then, $\forall s$,

A m-efficient at $s \Rightarrow F$ m-efficient at s.

Proof: Fix s. We proceed by induction on the depth of the tree d. For $d = 1$, we have simply L resources with disjoint sets of players. The constraint $qF(s)r(s) \leq Q$ has decoupled columns, i.e. it can be written as L independent constraints $\sum_{i \in \mathcal{I}(l)} qF_i(s) \leq Q^l$. Thus

$$m(F(s)) = \sum_{i \in \mathcal{I}} m_i(F_i(s)) = \sum_{l \in \mathcal{L}} \sum_{i \in \mathcal{I}(l)} m_i(F_i(s)),$$

is maximized iff each of the L terms $\sum_{i \in \mathcal{I}(l)} m_i(F_i(s)))$ is independently maximized. Now since a player i will go through one allocation only, $F_i(s) = A_i(Q^l, 1_{\mathcal{I}(l)}(s))$, and

$$\sum_{i \in \mathcal{I}(l)} m_i(F_i(s)) = \sum_{i \in \mathcal{I}(l)} m_i(A_i(Q^l, 1_{\mathcal{I}(l)}(s))) = m(A(Q^l, 1_{\mathcal{I}(l)}(s))).$$

Since A is m-efficient at s, the above expression is maximal for each l, therefore $m(F(s))$ is maximized, i.e. F is m-efficient.

Now suppose the result holds for all trees of depth $< d$. Let F' be any feasible allocation rule on the tree of depth d. We will show that that $m(F(s)) \geq m(F'(s))$. Let $\mathcal{L}^1 = \{$ depth 1 edges $\}$, i.e. all the edges going out of the root. We can decouple the problem into $|\mathcal{L}^1|$ independent subtrees. Each player only uses one of the links in $\mathcal{L}^1$, so

$$m(F(s)) = \sum_{l \in \mathcal{L}^1} \sum_{i \in \mathcal{I}(l)} m_i(F_i(s))$$

is maximized iff each of the $|\mathcal{L}^1|$ terms of the outer sum is independently maximized. Pick $l \in \mathcal{L}^1$. Let v be the vertex at the head of l. Let $F(v, .)$ denote the tree allocation rule $F(.)$ applied to the subtree rooted at v. From the algorithm, it is easy to see that $F_i(s) = A_i(Q^l, F(v, 1_{\mathcal{I}(l)}(s)))$. Therefore

$$m(F(s)) = \sum_{l \in \mathcal{L}^1} m(A(Q^l, F(v, 1_{\mathcal{I}(l)}(s)))). \tag{32}$$

Now

$$m(F'(s)) = \sum_{l \in \mathcal{L}^1} \sum_{i \in \mathcal{I}(l)} m_i(F'_i(s)) = \sum_{l \in \mathcal{L}^1} m(1_{\mathcal{I}(l)}(F'(s))). \tag{33}$$

Since by the induction hypothesis, F is m-efficient on the tree of depth $< d$ rooted at v, $m(F(v, 1_{\mathcal{I}(l)}(s))) \geq m(1_{\mathcal{I}(l)} F'(s))$, i.e. the whole-tree allocation rule F' cannot beat the subtree-efficient rule $F(v, .)$ on the subtree. Therefore, by the monotonicity property (31), we have

$$\begin{aligned} m(A(Q^l, F(v, 1_{\mathcal{I}(l)}(s)))) &\geq m(A(Q^l, 1_{\mathcal{I}(l)} F'(s))) & (34) \\ &\geq m(1_{\mathcal{I}(l)}(F'(s))), & (35) \end{aligned}$$

where (35) follows from the fact that F being m-efficient it beats in particular the identity allocation $Id(s) = s$, which is feasible in this case. Now, combining (32), (33) and (35), we get $m(F(s)) \geq m(F'(s))$. $\square$