COMMUNICATIONS IN INFORMATION AND SYSTEMS
Vol. 3, No. 4, pp. 305-324, September 2004

004

COMBINATORIAL EXCHANGE MECHANISMS FOR EFFICIENT BANDWIDTH ALLOCATION*

RAHUL JAIN† AND PRAVIN P. VARAIYA‡

Dedicated to Sanjoy Mitter on the occasion of his 70th birthday.

Abstract. The size, scale and multiple ownership of communication network resources makes it important to consider an economic framework wherein we can investigate the efficiency of network operation taking agents' incentives into account. Such a framework has been considered in the design and analysis of pricing mechanisms to regulate congestion and share bandwidth over short time scales. We consider time scales of a few months over which owners of communication links lease bandwidth to network service providers. As is well-known, economic efficiency is related to how close an allocation is to a competitive equilibrium. We first show that achieving economic efficiency through a market mechanism depends on network topology. We then show that in finite networks a competitive equilibrium may not exist. But a competitive equilibrium does exist in an idealized continuum model, in which all agents are infinitesimal compared with the size of the network. This suggests that approximate competitive equilibria with good performance may be attainable in real networks. We finally introduce a market mechanism called the combinatorial seller's bid double auction whose outcome, in the continuum model, is a competitive equilibrium.

1. Introduction. Communication networks have increased in scale and heterogeneity. There are multiple network owners and operators, each with their own heterogeneous endowments and privately known cost and revenue models. So, an allocation of network resources efficient from an engineering perspective need not be realized in the market. Hence, it is useful to find economic mechanisms that result in efficient resource allocations among selfish agents.

The efficiency of bandwidth allocation among competing users in a communication network has been studied within an economic framework before. The previous research literature focuses attention on congestion that results when aggregate user demand exceeds capacity, and proposes usage-dependent pricing as a method for controlling congestion [23, 24]. However, several difficulties must be faced in coming up with a good price-adjustment mechanism. First, congestion occurs over a very short time scale: Aggregate traffic is bursty and exceeds nominal capacity for periods of a few seconds [10]. So the pricing mechanism must react very quickly to changes in aggregate demand. Second, the 'correct' congestion price depends on how much users are willing to pay for the service, which is not known. This has prompted invention of distributed price iteration algorithms that converge to the 'correct' price [16, 22]. These algorithms are not practical but permit a re-interpretation of TCP-like proto-

*Accepted for publication on June 21, 2004. Research supported by DARPA contract F 30602-01-2-0548.

†EECS department at the University of California, Berkeley, E-mail: {rjain, varaiya} @eecs.berkeley.edu

cols as rate-control algorithms that embody user responses to congestion notification signals that serve as price surrogates [18, 21].

In contrast with this literature on congestion control, this paper is concerned with the long-term economic efficiency of a communication network. Multiple agents are involved in a large scale communication network such as the Internet. There are owners of capacities on communication links like AT&T and Sprint and network service providers like AOL and Earthlink. In the current Internet, there are many such entities all over the world. Currently, the service providers enter into long-term bilateral contracts with capacity owners. However, such bilateral exchanges do not result in efficient allocation and we must seek other ways of promoting it.

In the late 90s, Enron Broadband joined by Williams Communications proposed the setting of an exchange where transmission capacity (bandwidth) is traded like a 'commodity'. However, the market failed to establish itself. There were too few market players, each big enough to influence the market. The lack of sufficient number of pooling points made efficient trading of bandwidth difficult. Excess capacity also compounded the problem. This situation however may change in the future [8, 20].

We investigate the following questions: When is economic efficiency achievable in a communication network through a market mechanism? How does economic efficiency depend on network topology? What market mechanisms are available to achieve economic efficiency?

We study such markets within mathematical models that emphasize certain distinctive features. First, capacity is traded in *indivisible* units: One can buy or sell multiple trunks on a link, each with a fixed bandwidth, say OC-3, but not a fraction of a trunk. Second, demand for bandwidth is *combinatorial*: A buyer must purchase bandwidth in several links to create an end-to-end route; a seller owning two links may wish to sell at least one of them. (Indivisibility and combinatorial demand are also characteristic of spectrum and logistics auctions.)

We formulate the exchange between network players as a double auction, which allows room for a study of agents' strategic behavior. However, our emphasis is on *competitive equilibria*, that is, allocations in which aggregate demand equals aggregate supply, and prices that support such allocations. It is well-known that economic efficiency i.e., maximization of social welfare is related to competitive equilibrium [26]. We thus study the existence of competitive equilibrium and market mechanisms whose outcome are a competitive equilibria.

The set-up is as follows: Consider a network $G = (N, L)$ with a finite set of nodes N, and links L. A link l's capacity C_l is an integer number of trunks, each trunk with the same bandwidth. There are M agents, each with an initial endowment of money and link bandwidth. We conduct a double auction between buyers and sellers: each buyer specifies the bundle of links (comprising a route), the bandwidth (number of trunks) on each link, and the maximum price it is willing to pay for the bundle;

each seller specifies a similar bundle and the minimum price it is willing to accept. We assume that each agent's preferences are monotonic over the bundle (they prefer larger bundles to strictly smaller ones) and continuous in money. Moreover, we assume that buyers insist on getting the same bandwidth on all links in their bundles. The framework is quite general and can be extended to the case where the network consists of several autonomous systems and their owners are trying to negotiate service level agreements (SLAs) about capacity, access and QoS issues.

We first show that when agents have quasi-linear utility functions, existence of competitive equilibrium, and hence of economically efficient market mechanisms depends on network topology. We show an example of a finite network with a finite number of agents, for which no competitive equilibrium exists. We then model a perfect competition economy as one with a continuum of agents, each with negligible influence on the final allocation and prices [3]. Such idealized models are used frequently in economics. They are helpful in characterizing and finding approximate equilibria that are nearly efficient for finite settings. We show that a competitive equilibrium exists in a continuum model of a network. This is accomplished using the Debreu-Gale-Nikaido lemma, a useful corollary of Kakutani's fixed point theorem.

In the second part of the paper, we propose and analyze a combinatorial double auction mechanism that maximizes the social surplus. While the auction mechanism is general enough for any combinatorial auction, we keep the bandwidth trading market foremost in our mind.

There is much recent work on combinatorial auctions [19, 35], and interest in the interplay between economic, game-theoretic and computational issues leading to new areas such as algorithmic mechanism design [30]. There are a number of mechanisms for combinatorial auctions that are iterative, ascending price and achieve efficiencies close to the Vickrey auction. Some of them rely on approximation of the integer program [27] while others rely on linear programming [6]. The focus of this paper is on combinatorial double auction mechanism design, which is relatively unexplored. A modified Vickrey double auction, related to the k-double auction, is presented in [38], while [9] considers truthful double auction mechanisms and obtain upper bounds on the profit of any such auction. But the setting for both papers is non-combinatorial and these results for single-item auctions and/or single-sided auctions do not extend to combinatorial double auctions.

We consider multi-item combinatorial double auctions. We assume that sellers offer "loose" bundles, each with just one link (or one type of item). For example, if a seller has 5 units of link A and 5 units of link B, he makes two OR offers, one with 5 units of link A and another with 5 units of link B, but then within each bundle only a fraction of the units may get sold, say 3 out of 5 units. The buyer's bundles on the other hand are of "all-or-none" kind. If a buyer bids for 5 units of *both* link A and link B, and if this bid is accepted, the buyer must receive all 5 units on each of the two links. As mentioned earlier, this requirement is motivated by realistic situations where buyers want to acquire routes on communication networks. The assumption of "loose" bundles for sellers allows us to define uniform prices on links.

The allocation reached by the auction is obtained by solving a linear integer program that maximizes the trade surplus. Based on the allocation, the mechanism declares the highest matched seller's ask-bid as the price on the link. The mechanism thus specifies the 'rules of a game' among buyers and sellers. It is known that when there is only one type of good, the Nash equilibrium strategies of the associated game converge to truth-telling [11, 36]. Moreover, the rate of convergence is $O(1/n)$, where n is the number of players on each side of the market [32]. When all players bid truthfully, the allocation determined by the mechanism is efficient.

Our proposed mechanism called *c-SeBiDA* (combinatorial sellers' bid double auction) is combinatorial and in a framework that allows us to define uniform and anonymous (competitive) prices on the links. Such prices are highly desirable from an economic perspective as they yield socially efficient and Pareto-optimal outcomes. This is achieved by few auction mechanisms.

The analysis of combinatorial auctions is usually very difficult, and even more so for combinatorial double auctions. We thus consider the continuum model and show that the auction outcome is a competitive equilibrium. While the model is an idealization of the scenario where there are a large number of agents such that no single agent can affect the auction outcome by himself, it suggests that the auction outcome is likely an approximate competitive equilibrium, and hence close to efficient. The methodology used in the proof is novel in that it casts the mechanism in an optimal control framework and appeals to the maximum principle to conclude that the outcome is indeed a competitive equilibrium. The c-SeBiDA mechanism has been implemented in a web-based software test-bed [2] and has been used in experiments to validate the mechanism.

The paper is organized as follows: In section 2, we present some examples of finite networks, and show that if bandwidth is indivisible, a competitive equilibrium may not exist. Section 3 presents existence results for the continuum model of a network. Section 4 introduces the combinatorial sellers' bid double auction algorithm. Section 5 presents analysis of the proposed mechanism. Section 6 presents conclusions. The proofs of the main theorems are technical and presented in the appendices.

2. Network Topology and Competitive Equilibrium. We first prove that a competitive equilibrium exists if the routes that buyers want form a tree and all agents (buyers and sellers) have utilities that are linear in bandwidth and money. Examples are given to show that a competitive equilibrium may not exist if the routes do not form a tree or if utilities are nonlinear.

Links are indexed $j = 1, 2, \cdots$; link j provides C_j trunks of bandwidth (C_j an integer). Its owner, j, can lease $y_j \leq C_j$ trunks and has a per trunk reservation price or cost a_j. Buyer i, $i = 1, 2, \cdots$, wishes to lease x_i trunks on each link j in route R_i. The value to buyer j of one trunk along route R_i is b_i. Let $A = \{A_{ij}\}$ be the edge-route incidence matrix, i.e. $A_{ij} = 1(0)$, if link $j \in (\notin) R_i$.

With this notation, the allocation (x^*, y^*) with the maximum surplus solves the following integer program:

$$\max_{x,y} \quad \sum_i b_i x_i - \sum_j a_j y_j \tag{1}$$

$$s.t. \ \sum_i A_{ij} x_i \leq y_j \leq C_j, \ \forall j \tag{2}$$

$$x_i, y_j \in \{0, 1, 2, \cdots\}, \ \forall i, j. \tag{3}$$

The allocation (x^*, y^*) together with a link price vector $p^* = \{p_j^*\}$ is a *competitive equilibrium* if every buyer i maximizes his surplus at x_i^*,

$$\max_{x_i = 0, 1, \cdots} (b_i - \sum_{j \in R_i} p_j^*) x_i,$$

and every seller j maximizes his profit at y_j^*,

$$\max_{y_j = 0, 1, \cdots, C_j} (p_j^* - a_j) y_j.$$

A matrix is *totally unimodular* (TU) if the determinant of every square submatrix is 0, 1 or -1 [33]. If the routes that buyers want in a network form a tree, its edge-route incidence matrix is TU.

THEOREM 1. *If A is TU, in particular if the routes form a tree, there is a competitive equilibrium.*

Proof. Consider the relaxed LP version of problem (1) in which the integer constraint (3) is dropped. Because A is TU, the convex set of allocations (x, y) that satisfy constraint (2) has integer-valued vertices. Hence there is an optimal solution (x^*, y^*) to the LP problem which is integer-valued. The Lagrange multipliers $\{p_j^*\}$ associated with the constraint (2), together with (x^*, y^*), form a competitive equilibrium, as can be verified from the Duality Theorem of LP. □

The proposition has a partial converse: If $(p^*, (x^*, y^*))$ is a competitive equilibrium, (x^*, y^*) is a solution to the relaxed LP problem.

It is well known that a competitive equilibrium exists if every buyer i (seller j) has a utility (cost) function $u_i(x_i)(v_j(y_j))$ that is concave (convex), monotone and continuous (along with some boundary conditions) [1] and *fractional* trunks can be traded. This fact is exploited in [16, 18] to infer existence of competitive equilibrium prices for bandwidth on each link.

Examples 1,2 are non-TU networks that do not have a competitive equilibrium.

EXAMPLE 1. *Consider the cyclic network in figure 1 with buyers $1, \cdots, 4$, who want routes $\{e1, e2\}$, $\{e2, e3\}$, $\{e3, e1\}$, and $\{e_3\}$, respectively. Buyers $1, 2, 3$ receive*

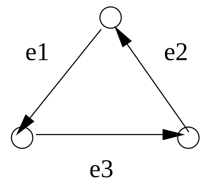

FIG. 1. *A cyclic network that is not TU.*

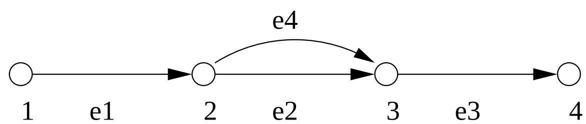

FIG. 2. *An acyclic network that is not TU.*

benefit $b_i = 1$ *per trunk; buyer 4 receives* $b_4 = \alpha(< 0.5)$. *Sellers own one trunk on each link, and their reservation price* $a_j = 0$ *for all links. The network is not TU, as can be easily checked. Surplus maximization allocates route* $\{e1, e2\}$ *to user 1 and* $\{e3\}$ *to user 4. If prices* p_1, p_2, p_3 *were to support this allocation, they must satisfy the conditions,* $1 = p_1 + p_2 \leq \min(p_2 + p_3, p_1 + p_3)$ *and* $0.5 > \alpha \geq p_3$, *which is impossible. So there is no competitive equilibrium.* □

EXAMPLE 2. *Consider the acyclic network of figure 2 again with buyers* $1, \cdots, 4$, *desired routes* $\{e1, e2\}, \{e2, e3\}, \{e1, e4, e3\}$, *and* $\{e3\}$ *and benefits* b_i *as before. Each link supports one trunk, and the sellers are as before. Surplus maximization again allocates route* $\{e1, e2\}$ *to user 1 and* $\{e3\}$ *to user 4. Competitive prices supporting this allocation must satisfy* $1 = p_1 + p_2 \leq \min(p_2 + p_3, p_1 + p_4 + p_3)$, $0.5 > \alpha \geq p_3$, *and* $p_4 = 0$, *which is impossible.* □

Next we see a TU network with nonlinear concave utilities for which there is no competitive equilibrium.

EXAMPLE 3. *Consider a network with two links, each with two trunks of capacity. There are two buyers. Buyer 1 wants a route through both links with bandwidth* x_1 *and has concave utility function:*
$u_1(x; \{l_1, l_2\}) = 1.1x, 0 \leq x \leq 1; = x + 0.1, 1 \leq x \leq 2.$
Buyer 2 demands bandwidth x_2 *only on link 2 and has concave utility function:*
$u_2(x; l_2) = 1.5x, 0 \leq x \leq 1; 1.1(x-1) + 1.5, 1 < x \leq 1 + \epsilon; (x - 1 - \epsilon) + 1.6, 1 + \epsilon < x \leq 2,$ *where* $\epsilon = 0.1/1.1$.
The sellers have reservation price of 0 on each link. It is easy to check that if fractional trunks can be traded there is a competitive equilibrium with allocation $x^* = (0.4, 1.6)$ *and prices* $p^* = (0, 1.1)$. *However, if trades must be in integral trunks, there is no competitive equilibrium.* □

A competitive equilibrium is efficient, because it maximizes total surplus $\sum_i b_i x_i - \sum_j a_j y_j$. To find an equilibrium, one normally proposes an iterative mechanism (often called 'Walrasian') involving an 'auctioneer' who in the nth round proposes link prices $\{p_j^n\}$, to which agents respond: buyer i places demand x_i^n, seller j offers to supply $y_j \leq C_j$ trunks. The auctioneer calculates the 'excess demand' on link j, $\xi_j^n = \sum A_{ij} x_i^n - y_j^n$, and begins round $n+1$ with price p_j^{n+1} higher or lower than p_j^n, accordingly as ξ_j^n is positive or negative. The equilibrium is reached when $\xi_j \leq 0$ for all j.

Two questions arise: Will the iterations converge? And are such mechanisms practically implementable? Of course, if there is no competitive equilibrium, caused perhaps by indivisibilities, the price-adjustment algorithms will not converge and no practical mechanisms can exist. Thus, in the next section we study the existence of competitive equilibrium in an ideal model.

3. Economic Efficiency in an ideal model. Consider a competitive communication network with indivisible trunks. We assume that there is perfect competition in which no single agent can influence the outcome, by considering a continuum of agents (buyers and sellers). The continuum network economy first introduced by Aummann [3] is an idealized model of perfect competition in which no agent has significant 'market power' to be able to alter the outcome. From a practical perspective, the existence of a competitive model in the ideal model can be used to establish existence of approximate competitive equilibria (which are approximately efficient) in finite network economies.

With money. Consider a network $\mathcal{G}$ with L links (goods $1, \cdots, L$), each with capacity C_l (which is integral), and bandwidth on each link is available in integral (discrete) units. Also assume that there is money (denoted as good 0) which however is divisible. There is a continuum of agents indexed $t \in X = [0, M]$, with a given *non-atomic* measure space $(X, \mathcal{B}(X), \mu)$. Suppose there are M routes and each agent t demands some route R_i. For example, all agents $t \in (m, m+1]$ demand route R_{m+1}, for $m+1 \leq M$. Agents' preferences $\succeq_t$ are monotonic, and continuous in money. (Monotonicity simply means that if $A \subseteq B$, then $B \succeq_t A$. Continuity means that if $A_n \to A$ and $B \succeq_t A_n$, then $B \succeq_t A$.) As a result preferences are continuous. A particular example of such preferences is when utility functions are quasi-linear in money, i.e., linear in money. Agent t has an initial endowment ω_t, which is a $L+1$-tuple. Though the following discussion and the results are for any general initial endowment, a particular example is of an auction setting when agent 0 is an auctioneer, endowed with $\omega_0 = (0, C_1, \cdots, C_L)$ (the whole network) and any other agent $t (> 0)$ has $\omega_t = (m_t, 0, \cdots, 0)$, in which m_t is t's money endowment. Similarly, the price vector $p = (p_0, p_1, \cdots, p_L)$ is a $L+1$-tuple; p_0 is the price of money and p_l is the price of unit bandwidth on link l. We will call a system as described above a *network exchange economy* $\mathcal{E}$.

We begin with a few definitions: Let $p \in \Theta = \mathbb{R}_+^{L+1}$ be a price vector. By $p > 0$, we shall mean that all components are non-negative with $p \neq 0$, and by $p \gg 0$, we shall mean that all components are strictly positive.

Commodity space: $\Omega = \mathbb{R}_+ \times \mathbb{Z}_+^L$. Thus, $\omega = (\omega_0, \cdots, \omega_L) \in \Omega$ denotes ω_0 units of money and ω_l units of good l, for $l = 1, \cdots, L$. Note that ω_l for $l > 0$, must be an integer, indicating indivisibility.

Unit price simplex: $\Delta = \{p \in \Theta : \Sigma_0^L p_l = 1\}$. Prices lie in the unit simplex. Later, we can normalize prices so that the price of money, $p_0 = 1$, and we get the prices of other goods in terms of money.

Budget set: $B_t(p) = \{z \in \Omega : p.z \leq p.\omega_t\}$, which gives the allocations that agent t can afford based on its initial endowment at given prices p.

Preference level sets: $\mathcal{P}_t(z) = \{z' \in \Omega : z' \succeq_t z\}$, is the set of allocations preferred by agent t to the allocation z.

Individual demand correspondence: $\psi_t(p) = \{z \in B_t(p) : z \succeq_t z', \forall z' \in B_t(p)\}$. At given prices p, $\psi_t(p)$ is the set of t's most-preferred allocations. There may be more than one most preferred allocation, so ψ_t is a demand correspondence rather than a demand function.

Aggregate excess demand correspondence:

$$\Phi(p) = \int_X \psi_t(p) d\mu - \int_X \omega_t d\mu.$$

We denote the first integral by $\Psi(p)$—the aggregate demand correspondence, and the second integral by $\bar{\omega}$, the total endowment of all agents, or the aggregate supply.

Definition 1 (Competitive Equilibrium). *A pair (x^*, p^*) with $p^* \in \Delta$ and $x^* \in \Omega$ is a competitive equilibrium if $x_t^* \in \psi_t(p^*)$ and $0 \in \Phi(p^*)$.*

A competitive equilibrium comprises an allocation and a set of prices such that the prices support an allocation for which aggregate demand equals aggregate supply, or in other words, the aggregate excess demand is zero. Moreover, the allocation to each agent is what it demands at those prices.

We make the following assumptions:

Assumptions.

1. $\bar{\omega} \gg 0$ (component-wise positive), and $\omega_t > 0, \forall t$ (component-wise non-negative with some component positive).
2. $\Phi(p)$ is homogeneous in p.
3. Boundary condition: Suppose $p^\nu \to p^*$, and $p_l^* = 0$ for some l. Then, $z_l^\nu \to \infty$, $\forall z^\nu \in \Phi(p^\nu)$.
4. Walras' Law holds: $p.z = 0, \forall z \in \Phi(p), \forall p \in \Delta^0$, the relative interior of Δ.

The first assumption simply states that there is a strictly positive endowment of each good (capacity on each link), and moreover each agent has a strictly positive endowment of some good (either money, or capacity, or both). The second assumption ensures that scaling of prices does not alter the competitive allocation if it exists. The third assumption is a boundary condition that holds in the absence of undesirable goods. The fourth assumption, Walras' law, can be shown to hold for the economy under consideration. But we shall assume it without proof. Essentially, it means that if there is positive excess demand for a good at given prices, its price can be reduced still further towards zero.

Now we can show the following:

THEOREM 2 (Existence). *Under assumptions (1)-(4), a competitive equilibrium exists in the network exchange economy $\mathcal{E}$.*

The proof is long and technical, and relegated to the appendix. It relies on lemma 1 which is a corollary of Kakutani fixed point theorem and the Lyapunov-Richter theorem which states that the integral of a correspondence with respect to a non-atomic measure is closed and convex-valued [4]. We can set the price of money $p_0 = 1$, and we get the other prices in units of money. Since we only require that preference be monotonic (and not strictly monotonic), the next result follows.

COROLLARY 1. *If demands are for source-destination pairs (and agents are indifferent between different routes for same source-destination pair), then theorem 2 still holds. And in that case, prices for various alternative routes (given by the sum of link prices along the routes) for a given source-destination pair are same.*

Without money. The role of money is crucial in the above result. The following example shows that in the absence of money, a competitive equilibrium may not exist even in a continuum network exchange economy.

EXAMPLE 4. *Consider the networks of figure 1, with demands as discussed before in example 3. Now instead of one user of each type demanding a particular route, we have a continuum of users. Let $X = [0, M]$ and let all users in $[0, 1]$, where M is the total number of routes, demand the same route and have identical preferences. We make the same assumption for the other M disjoint intervals of unit length. This reduces the continuum case to the same as example 3, for which a competitive equilibrium does not exist.*

4. Combinatorial Sellers' Bid Double Auction (c-SeBiDA). We now propose a combinatorial double auction mechanism for bandwidth trading. A buyer places buy bids for bandwidth along his route. A buyer's bid is *combinatorial* or 'all or nothing': he must receive the same bandwidth on every link along the route or nothing at all. On the other hand, sellers make separate ask bids on *individual* links.

The mechanism collects all announced bids, matches a subset of these to maximize the 'surplus' ((4), below) and declares a settlement price for bandwidth on each link at which the winning buy and ask bids—which we call the matched bids—are transacted. This constitutes the payment rule. The announced prices are uniform (i.e., same for all items of the same type) and anonymous (i.e., same for all players independent of their bids and true valuations). As will be seen, each matched buyer's buy bid is larger, and each matched seller's ask bid is smaller, than the settlement price, so the outcome respects *individual rationality.*

There is an asymmetry: buyers make multi-link combinatorial bids, but sellers make single-link bids. This permits the mechanism to select uniform and anonymous settlement prices for each link.

A player can make multiple bids. The mechanism treats these as XOR bids, so at most one bid per player is a winning bid. Therefore the outcome is the same as if a matched player only makes (one) winning bid. Thus, in the formal description of the *combinatorial Sellers' Bid Double Auction* (c-SeBiDA), each player places only one bid. c-SeBiDA is a 'double auction' because both buyers and sellers bid; it is a 'sellers' bid' auction because the settlement price depends only on the matched sellers' bids, as we will see.

Formal mechanism. Consider a network with links $1, \cdots, L$. Buyer i has (true) reservation value $v_i \in [0, V]$ per trunk along route R_i, and submits a buy bid of b_i per trunk and demands $\delta_i \in [0, D]$ trunks along the route. Seller j has (true) per trunk cost $c_j \in [0, C]$ and offers to sell up to $\sigma_j \in [0, S]$ trunks in link l_j at a unit price of a_j. Denote $L_j = \{l_j\}$.

The mechanism receives all these bids, and matches some buy and sell bids. The matches are described by 0-1 variables x_i, y_j: x_i is 1, if buyer i's buy bid is accepted, 0 otherwise; similarly, y_j is 1 if seller j's ask bid is accepted, 0 otherwise.

The mechanism determines the allocation (x^*, y^*) as the solution of the surplus maximization problem **MIP**:

$$\begin{aligned} \max_{x,y} \quad & \sum_i x_i \delta_i b_i - \sum_j y_j \sigma_j a_j \\ s.t. & \sum_j y_j \sigma_j \mathbb{1}(l \in L_j) - \sum_i x_i \delta_i \mathbb{1}(l \in R_i) \geq 0, \forall l \in [1 : L], \\ & x_i \in \{0, 1\}, \forall i, \quad y_j \in [0, 1], \forall j. \end{aligned} \tag{4}$$

MIP is a mixed integer program: Buyer i's bid is either fully matched, or not at all ($x_i^* \delta_i \in \{0, \delta_i\}$); but seller j's bid may be partially matched ($y_j^* \sigma_j \in [0, \sigma_j]$). Nevertheless, $\sum_{j=1}^{M} y_j^* \sigma_j \mathbb{1}(l \in L_j)$ is integral for all l.

The settlement price is the highest ask bid among matched sellers,

$$\hat{p}_l = \max\{a_j : y_j^* > 0, l \in L_j\}. \tag{5}$$

The payments are determined by these prices. Matched buyers pay the sum of the link prices along their route; matched sellers receive a payment equal to the number of trunks sold times the price for the link. Unmatched buyers and sellers do not participate. This completes the mechanism description.

For use later, we will also define the following quantity:

$$\check{p}_l := \min\{a_j : y_j = 0, l \in L_j\}.$$

If i is a matched buyer ($x_i^* = 1$), his bid $b_i \geq \sum_{l \in R_i} \hat{p}_l$; for otherwise, the surplus (4) can be increased by eliminating this bid. Similarly, if j is a matched seller ($y_j^* > 0$), and $l \in L_j$, his bid $a_j \leq \hat{p}_l$, for otherwise the surplus can be increased by eliminating his bid. Thus the outcome of the auction respects individual rationality.

It is easy to understand how the mechanism picks matched sellers. For each link j, a seller with lower ask bid will be matched before one with a higher bid. So sellers with bid $a_j < \hat{p}_l$ sell all their bandwidth ($y_j^* = 1$). At most one seller with ask bid $a_j = \hat{p}_l$ sells a fraction of his bandwidth ($y_j^* < 1$). On the other hand, because their bids are combinatorial, the matched buyers are selected only after solving MIP.

Our focus is on price determination, so we will assume that the quantities in players' bids (namely, $\{\delta_i, \sigma_j\}$) are fixed. A strategy for seller i is a buy bid b_i, a strategy for buyer j is an ask- bid a_j. Let θ denote a collective strategy. Given θ, the mechanism determines the allocation (x^*, y^*) and the prices $\{\hat{p}_l\}$. So the payoff of buyer i and seller j are, respectively,

$$u_i^b(\theta) = x_i^* \delta_i (v_i - \sum_{l \in R_i} \hat{p}_l), \tag{6}$$

$$u_j^s(\theta) = y_j^* \sigma_j (\sum_{l} \hat{p}_l \mathbb{I}(l \in L_j) - c_j). \tag{7}$$

The bids b_i, a_j may be different from the true valuations v_i, c_j, which however figure in the payoff calculations above.

5. Analysis of c-SeBiDA. We would like to know the behavior of the outcome of this auction as the number of players is large enough such that no single player by himself can affect the outcome. As before, we consider a continuum of agents. Since each agent cannot by himself alter the outcome, we will assume that the bids are truthful and analyze if the auction outcome is a competitive equilibrium, and hence efficient.

Assume the continuum of buyers indexed by $t \in [0, 1]$, and continuum of sellers indexed by $\tau \in [0, 1]$. There are N types of buyers and M types of sellers. Let $B_1, \cdots, B_N$ and $S_1, \cdots, S_M$ partition $[0, 1]$ so that all buyers in B_i demand the same set of items R_i (corresponding to a route), and all sellers in S_j offer the same item l_j, $L_j = \{l_j\}$. We assume that the partitions B_i's and S_j's are subintervals.

A buyer $t \in B_i$ has true value $v(t)$, bids $p(t)$ per unit for the set R_i, and demands $\delta(t) \in [0, D]$ units. Suppose $v(t), p(t) \in [0, V]$. A seller $\tau \in S_j$ has true cost $c(\tau)$ and asks $q(\tau)$ for the item(s) L_j with supply $\sigma(\tau) \in [0, S]$ units, with $c(\tau), q(\tau) \in [0, C]$. Let $x(t)$ and $y(\tau)$ be the decision variables, i.e. buyer t's $x(t)$ is 1, if bid is accepted, 0 otherwise. And similarly seller τ's $y(\tau)$ is 1 if offer is accepted, 0 otherwise. Without loss of generality, we assume that within each partition B_i, the buyers' bid function $b(t)$ is non-increasing, and within each partition S_j, the sellers' bid function $q(\tau)$ is nondecreasing.

Denote the indicator function by $\mathbb{1}(\cdot)$ and as before, consider the surplus maximization problem **cLP**:

$$
(8)\quad \sup_{x,y} \int_0^1 \sum_{i=1}^{N} x(t)\delta(t)p(t)\,\mathbb{1}(t \in B_i)dt - \int_0^1 \sum_{j=1}^{M} y(\tau)\sigma(\tau)q(\tau)\,\mathbb{1}(\tau \in S_j)d\tau,
$$

$$
\text{s.t.} \int_0^1 \sum_{j=1}^{M} y(\tau)\sigma(\tau)\,\mathbb{1}(l \in L_j, \tau \in S_j)d\tau - \int_0^1 \sum_{i=1}^{N} x(t)\delta(t)\,\mathbb{1}(l \in R_i, t \in B_i)dt \geq 0,\ \forall\, l,
$$

$$
x(t), y(\tau) \in \{0,1\},\ \forall\, t, \tau \in [0,1].
$$

The mechanism determines $((x^*, y^*), \hat{p})$ where (x^*, y^*) is the solution of the above continuous linear program and for each $l \in [1 : L]$,

$$
\hat{p}_l = \sup\{q(\tau) : y(\tau) > 0, \tau \in S_l, l \in L_j\},
$$

and

$$
\check{p}_l = \inf\{q(\tau) : y(\tau) = 0, \tau \in S_l, l \in L_j\}.
$$

The mechanism announces prices $\hat{p} = (\hat{p}_1, \cdots, \hat{p}_L)$; the matched buyers (those for which $x^*(t) = 1$) pay the sum of the prices in their bundle while the matched sellers (those for which $y^*(\tau) = 1$) get a payment equal to the number of their items sold times the price of the item.

When buyers and sellers bid truthfully, the following result holds:

THEOREM 3. *Suppose that the bid function of the sellers* $q : [0,1] \to [0, C]$ *is piecewise continuous and nondecreasing (i.e. continuous and nondecreasing in each partition* S_j *of* $[0,1]$*), then* (x^*, y^*) *is a competitive allocation and* $\hat{p}$ *is a competitive price.*

The proof is technical and is relegated to the appendix. It casts the surplus maximization problem cLP in the optimal control framework to conclude the existence of (x^*, y^*) and the dual variables $\lambda^* = (\lambda_1^*, \cdots, \lambda_L^*)$ using the maximum principle [28]. The assumptions are then used to conclude that $\hat{p} = \lambda^*$, and hence are competitive prices.

In section 3, the existence result we established was in a much more general framework though we considered money as another good in the economy. In this section, the existence result is for the special framework of a combinatorial double auction, and we do not consider money as another good.

We now show that the assumption that the sellers' bid function is piecewise continuous and nondecreasing is necessary for the c-SeBiDA's prices to be competitive prices.

EXAMPLE 5. *Suppose there is only one item. Buyers* $t \in [0, 0.5]$ *have valuation 1 while buyers* $t \in (0.5, 1]$ *have valuation 4. Sellers* $t \in [0, 0.5]$ *have valuation 3 while sellers* $t \in (0.5, 1]$ *have valuation 2. The mechanism will match buyers in* $(0.5, 1]$ *and* *sellers* $(0.5, 1]$ *with surplus* $0.5 \times 2 = 1$. *Thus,* $\hat{p} = 2$ *which is not equal to* $\check{p} = 3$. *As can be easily checked, the competitive price in this case is* $\lambda^* = 3$ *different from* $\hat{p}$.

Remarks. The outcome of the designed auction mechanism is a competitive equilibrium if players are truthful. But it is not clear how it relates to the Nash equilibria of the game. This is a very difficult question that has drawn the attention of mathematical economists for a long time.

The problem of relating Nash equilibria of games to competitive equilibrium was first investigated in [29], in which it is shown that when consumers announce their preferences, then as the economy is replicated, the gain that any player can achieve by misrepresenting his demand goes to zero. In [12] it is shown that under certain regularity conditions, the economy replicated enough has an allocation which is incentive-compatible, individually-rational and ex-post ϵ-efficient. In [13] it is shown that the demand functions that an agent might consider based on strategic considerations converge to the competitive demand functions. Further, under certain conditions on beliefs of individual agents, not only do the strategic behaviors of individual agents converge to the competitive behavior but the Nash equilibrium allocations also converge to the competitive equilibrium allocation [14].

A buyer's bid double auction with single type of item that maximizes surplus is studied in [37]. It is shown that with Bayesian Nash strategies, the mechanism is asymptotically incentive efficient in the sense that a player's maximum expected payoff using any strategy is the same as his competitive payoff when the number of players is sufficiently large. The rate of convergence of the Nash equilibria to the competitive equilibria for single-item buyer's bid double auction is investigated in [11, 31, 32]. Implementation and mechanism design in the continuum setting are discussed in [25].

6. Conclusions. We studied combinatorial bandwidth trading markets. We showed that for finite networks, prices that yield socially efficient allocation may not exist. We then used a model of perfect competition with a continuum of agents, and showed that with money, it is possible to support the socially efficient allocation with a certain price vector. The key here is the Lyapunov-Richter theorem that enables a convexification of the economy. However, such a result does not hold for countable economies. The main reason is that defining the average of a sequence of correspondences is trickier as the limit may not exist.

The continuum model is useful in showing the existence of enforceable approximate equilibria (when we require that supply exceeds demand) in finite networks. Such approximate equilibria were presented in [15]. It is well-known that the set of the competitive allocations is contained in the *core* (the set of Pareto-optimal allocations). However, it is unknown if the two sets are equal. This is an interesting question and part of future work.

We presented the combinatorial seller's bid auction mechanism. Since uniform and anonymous prices are very desirable, we designed the auction for the case when sellers do not submit bundle bids. We showed that under the continuum model, the prices the mechanism announces are competitive prices and the allocation is a competitive allocation.

Apendix. Proof of theorem 2.

Proof. Consider any non-empty, closed convex subset S of Δ. We will first make some claims about the properties of the aggregate excess demand correspondence [1].

CLAIM 1. *Φ is non-empty and convex-valued on S.*

From assumption 1, Φ is non-empty. Fix $p \in S$. By Lyapunov-Richter's theorem [4] with μ a non-atomic measure on X, and $\psi_t(p)$ a correspondence for each p, $\int_X \psi_t(p) d\mu(t)$ is convex. Hence, Φ is convex.

CLAIM 2. *Φ is compact-valued, hence bounded on S.*

Note that S is compact and for each $p \in S$, $p \gg 0$. Write

$$\psi_t(p) = \bigcap_{z \in B_t(p)} [B_t(p) \cap \mathcal{P}_t(z)].$$

Then, $\mathcal{P}_t(z)$ is closed by continuity of preferences. $B_t(p)$ is closed and bounded for $p \gg 0$. Thus, their intersection is closed. And so is the outer intersection. It is bounded as well. Thus, $\psi_t(p)$ is compact for each $p \gg 0$.

CLAIM 3. *$p \cdot z \leq 0$, $\forall p \in \Delta^0$, $z \in \Phi(p)$.*

Fix $p \in \Delta^0$. By definition,

$$p \cdot z \leq p \cdot \omega_t, \forall z \in \psi_t(p), \forall t \in X.$$

Or, with an abuse of notation:

$$\int_X p \cdot \psi_t(p) d\mu \leq \int_X p \cdot \omega_t d\mu,$$

$$p \cdot \Psi(p) \leq p \cdot \bar{\omega},$$

$$p \cdot \Phi(p) \leq 0,$$

CLAIM 4. *ψ_t is closed and upper semi-continuous (u.s.c.) in S $\forall t \in X$. Hence, Φ is closed and u.s.c. in S.*

Fix $t \in X$. To show ψ_t is closed, we have to show that for any sequences, $\{p^\nu\}, \{z^\nu\}$, $[p^\nu \to p^0, z^\nu \to z^0, z^\nu \in \psi_t(p^\nu)] \implies z^0 \in \psi_t(p^0)$. From the definition of demand correspondence, $p^\nu \cdot z^\nu \leq p^\nu \cdot \omega_t$. Taking limit as $\nu \to \infty$, we get $p^0 \cdot z^0 \leq p^0 \cdot \omega_t$, i.e. $z^0 \in B_t(p^0)$. It remains to show: $z^0 \succeq_t z, \forall z \in B_t(p^0)$.

[1] A good discussion of properties of correspondences can be found in [5].

Consider any $z \in B_t(p^0)$. Then

Case 1: $p^0 \cdot z < p^0 \cdot \omega_t$.

Then, for large enough ν, $p^\nu \cdot z < p^\nu \cdot \omega_t$. This implies that $z \in B_t(p^\nu)$. Now, $z^\nu \in \psi_t(p^\nu)$. Hence, $z^\nu \succeq_t z$. And by continuity of preferences, we get $z^0 \succeq_t z$.

Case 2: $p^0 \cdot z = p^0 \cdot \omega_t$.

Define $z'^\nu := ((1-1/\nu)z_0, z_1, \cdots, z_L) \in \Omega$, by divisibility of money. So, $p^0 \cdot z'^\nu < p^0 \cdot \omega_t$. Then, by the same argument as above: $z^0 \succeq_t z'^\nu$. And by continuity of preferences, we get $z^0 \succeq_t z$.

This implies $z^0 \in \psi_t(p)$, i.e. it is closed. Now, to show it is u.s.c., we have to show by proposition 11.11 in [7], that for any sequence $p^\nu \to p^0$, and any $z^\nu \in \psi_t(p^\nu)$, there exists a convergent subsequence $\{z^{\nu_k}\}$ whose limit belongs to $\psi_t(p^0)$.

Now, $p^\nu \to p^0 \gg 0$. Hence, $\exists \nu_0$ s.t. $p^\nu \gg 0, \forall \nu > \nu_0$. Define

$$\pi := \inf\{p_l^\nu : \nu > \nu_0, l = 0, \cdots, L\}.$$

Then, $p^\nu \cdot z^\nu \leq p^\nu \cdot \omega_t$ implies for all $\nu > \nu_0$,

$$0 < z^\nu \leq \frac{p^\nu \cdot \omega_t}{\pi},$$

i.e. the sequence $\{z^\nu\}$ is bounded. By the Bolzano-Weierstrass theorem, there exists a convergent subsequence $\{z^{\nu_k}\}$ converging to say, z^0. Since ψ_t is closed in S, $z^0 \in \psi_t(p^0)$. Thus, it is upper semi-continuous in S.

We now show that Φ is u.s.c (hence closed) as well. Let $p^\nu \to p^0$ in S. Consider $\xi^\nu \in \Psi(p^\nu) = \int_X \psi_t(p^\nu) d\mu$. Then, $\exists z_t^\nu$ s.t. $\xi^\nu = \int_X z_t^\nu d\mu$. Now, ψ_t is compact-valued and u.s.c. in S. Thus, by proposition 11.11 in [7], the sequence $\{z_t^\nu\}$ has a convergent subsequence $\{z_t^{\nu_k}\}$ s.t. $z_t^{\nu_k} \to z_t^0 \in \psi_t(p^0)$. Define $\xi^0 := \int_X z_t^0 d\mu$. Thus,

$$\xi^0 \in \int_I \psi_t(p^0) d\mu = \Psi(p^0).$$

As argued before, Ψ is compact-valued. Hence, by reapplication of the same theorem, it is u.s.c. in S. And so is Φ.

We now need the following lemma.

Lemma 1 (Debreu-Gale-Nikaido). [2] *Let S be a non-empty closed convex subset in the unit simplex $\Delta \subset \mathbb{R}^n$. Suppose the correspondence $\Phi : \Delta \to \mathcal{P}(\mathbb{R}^n)$ satisfies the following:*

(i) Φ is non-empty, convex-valued $\forall p \in S$,

(ii) Φ is closed,

(iii) $p \cdot z \leq 0, \forall p \in S, z \in \Phi(p)$,

(iv) $\Phi(p)$ is bounded $\forall p \in S$.

Then, $\exists$ $p^* \in S$ and $z^* \in \Phi(p)$ s.t. $p \cdot z^* \leq 0$, $\forall p \in S$.

[2] This can be found in [1, 7].

Then, using the above lemma, we get the following proposition.

PROPOSITION 1. *For any non-empty, closed convex subset S of Δ^0, $\exists p^0 \in S, z^0 \in \Phi(p^0)$ s.t. $p \cdot z^0 \leq 0, \forall p \in S$.*

Consider a increasing sequence of sets $S^\nu \uparrow \Delta$. Let p^ν, z^ν be those given by the above proposition. Then, $p^\nu \in S^\nu \subset \Delta$, which is compact. Thus, $\exists$ a convergent subsequence $p^{\nu_k} \to p^* \in \Delta$.

Without loss of generality, consider this subsequence as the sequence. Consider any $z^\nu \in \Phi(p^\nu)$. We have the following lower bound on the sequence

$$z^\nu \geq -\bar{\omega}, \forall \nu. \tag{9}$$

To get an upper bound, take any $\tilde{p} \gg 0 \in S^\nu$. It exists because $S^\nu \uparrow \Delta$. Using the proposition above, we get

$$\tilde{p} \cdot z^\nu \leq 0, \tag{10}$$

for large enough ν. Equations (9) and (10) imply $\{z^\nu\}(\subset \Omega)$ is bounded. Thus, there exists a convergent subsequence with limit say, z^*.

By assumption 1, $\bar{\omega} \gg 0$. Also, $p^* \in \Delta$. Hence, $p^* \cdot \bar{\omega} > 0$. Further, $p^* \gg 0$ since if $p_l^* = 0$ for some l, $z_l^{\nu_k} \to \infty$, by the boundary condition, which then contradicts the boundedness of the subsequence above. Further, since Φ is closed, $z^* \in \Phi(p^*)$.

This establishes the following lemma.

LEMMA 2. *$\exists p^* \gg 0 \in \Delta, z^* \in \Phi(p^*)$ s.t. $p^\nu \to p^*, z^\nu \to z^*$, and $p \cdot z^* \leq 0, \forall p \in \Delta$.*

We are now ready to prove the theorem: Walras' law implies $p \cdot z = 0$, $\forall z \in \Phi(p)$, and $\forall p \in \Delta^0$. This implies $p^* \cdot z^* = 0$. From lemma above, $p^* \cdot z^* \leq 0$, and $p^* \gg 0$. This yields $z^* = 0$. ☐

Proof of theorem 3.

Proof. We first show the existence of (x^*, y^*) and $\lambda^* = (\lambda_1^*, \cdots, \lambda_L^*)$, the dual variables corresponding to the demand less than equal to supply constraints. We do this by casting cLP above as an optimal control problem and then appeal to Pontryagin's maximum principle [34]. Define

$$\begin{aligned}
\dot{\zeta}(t) &:= \Sigma_{i=1}^N x(t)\delta(t)p(t)\,\mathbb{I}(t \in B_i) - \Sigma_{j=1}^M y(t)\sigma(t)q(t)\,\mathbb{I}(t \in S_j),\\
\dot{\xi}_l(t) &:= \sum_{j=1}^M y(t)\sigma(t)\,\mathbb{I}(l \in L_j, t \in S_j) - \sum_{i=1}^N x(t)\delta(t)\,\mathbb{I}(l \in R_i, t \in B_i),\\
\theta(t) &:= (\xi_1(t), \cdots, \xi_L(t), \zeta(t))',
\end{aligned}$$

where θ is the state of the system, x and y are controls, and $\zeta(t)$ and $\xi(t)$ describe the state evolution as a function of the controls. The objective is to find the optimal control (x^*, y^*) which maximizes $\zeta(1)$. Let

$$\Sigma(t) := \{\dot{\theta}(t) : x_l(0) = 0, \forall l \text{ and } x(t), y(t) \in \{0, 1\}, \forall t \in [0, 1]\}.$$

Observe that $\Sigma(t)$ has cardinality at most 2^{L+1} in $\mathbb{R}^{L+1}$. $\int_0^1 \Sigma(\tau)d\tau$ is the set of reachable states under all allowed control functions, namely, all measurable functions x and y with $x(\tau), y(\tau) \in \{0,1\}$. Note that $\zeta(1)$ is the total surplus; i.e., buyer surplus minus seller surplus, and $\xi_l(1)$ is the excess supply for item l; i.e., total supply minus total demand for item l. Define

$$\Gamma := \{\theta(1) \in \mathbb{R}^{L+1} : \theta(1) \in \int_0^1 \Sigma(\tau)d\tau, \xi_l(1) \geq 0, \forall l\},$$

the set of final reachable states under all control functions such that state evolution happens according to the equations above, and excess supply is non-negative.

LEMMA 3. *Γ is a compact, convex set.*

Proof. By assumption, $\delta(t), p(t), \sigma(t)$, and $q(t)$ are bounded. By Lyapunov-Richter theorem [4], $\int_0^1 \Sigma(\tau)d\tau$ is a closed and convex set. Since x and y are bounded functions, the integral is bounded as well. Thus, it is also compact. Moreover, $\xi_l(1)$ is a hyperplane, and $\xi(1) \geq 0$ defines a closed subset of $\mathbb{R}^L$. Therefore, $\{\theta(1) : \theta(1) \in \int_0^1 \Sigma(\tau)d\tau\} \bigcap \{\theta(1) : \xi_l(1) \geq 0, l = 1, \cdots, L\}$ is a compact, convex set. $\square$

Our optimal control problem is: $\sup_{\theta(1)\in\Gamma} \zeta(1)$. But observe that one component of $\theta(1)$ is $\zeta(1)$. Since Γ is compact and convex, the supremum is achieved and so an optimal control exists. By the maximum principle [28], there exist adjoint functions $p_0^*(t)$ and $p_l^*(t)$, $l = 1, \cdots, L$ such that $\dot{p}_0^*(t) = 0$, and $\dot{p}_l^*(t) = 0$, (i.e., $p_l^*(t) = \lambda_l^*$, a constant) for $l = 0, \cdots, L$.

Defining the Lagrangian over the objective function and the third (demand less than equal to supply) constraint

$$L(x, y; \lambda) = \zeta(1) + \sum_{l=1}^{L} \lambda_l \xi_l(1),$$

we get from the saddle-point theorem [34],

$$L(x, y; \lambda^*) \leq L(x^*, y^*; \lambda^*) \leq L(x^*, y^*; \lambda).$$

This leads to the next lemma.

LEMMA 4. *If $((x^*, y^*), \lambda^*)$ is a saddle point satisfying the inequality above, the λ^* are competitive equilibrium prices. Moreover, $\hat{p}_l \leq \lambda_l^* \leq \check{p}_l, \forall l = 1, \cdots, L$.*

Proof. Rewrite the Lagrangian as

$$L(x, y; \lambda) := \Sigma_{i=1}^{N} \int_{B_i} \delta(t)x(t)(p(t) - \Sigma_{l\in R_t}\lambda_l)dt + \Sigma_{j=1}^{M} \int_{S_j} \sigma(\tau)y(\tau)(\lambda_{l(\tau)} - q(\tau))d\tau.$$

where $l(\tau)$ is the item offered by seller τ. Using the first saddle-point inequality, we get $x^*(t) = \mathbb{1}(p(t) > \Sigma_{l\in R_t}\lambda_l^*)$ and $y^*(\tau) = \mathbb{1}(q(\tau) < \lambda_{l(\tau)}^*)$, which imply that the Lagrange multipliers are competitive equilibrium prices. To prove the second part, note that by definition, for a given τ, $y(\tau) > 0$ implies $q(\tau) \leq \lambda_l^*$ for $\tau \in S_l$, which gives the first inequality. Again from the definition, $y(\tau) = 0$ implies that $q(\tau) \geq \lambda_l^*$ for $\tau \in S_l$, which gives the second inequality. $\square$

To conclude the proof of the theorem, we observe that if q is continuous and non-decreasing in each interval S_j of $[0, 1]$, then $\hat{p}_l = \check{p}_l$ for each l, which then equals λ_l^* by lemma 4. $\square$

REFERENCES

[1] K. J. Arrow and F. H. Hahn, *General Competitive Analysis*, Holden-Day, 1971.

[2] *Auctions Theory Group*, http://auctions.eecs.berkeley.edu, EECS Department, University of California, Berkeley, October 2003.

[3] R. J. Aumann, *Markets with a continuum of traders*, Economterica, 32:1(1964), pp. 39–50.

[4] R. J. Aumann, *Integrals of set-valued functions*, J. of Mathematical Analysis and Applications, 12(1965)), pp. 1–12.

[5] C. Berge, *Topological Spaces*, Dover Books, 1997.

[6] S. Bikhchandani and J. Ostroy, *Ascending price Vickrey auctions*, to appear in: Games and Economic Behavior, 2004.

[7] K.C. Border, *Fixed Point Theorems with Applications to Economics and Game Theory*, Cambridge University Press, 1985.

[8] S. Borthick, *Cheap bandwidth: How low can it go*, Business Communications Review, pp. 14-16, August 2001.

[9] K. Deshmukh, A.Y. Goldberg, J.D. Hartline and A.R. Karlin, *Truthful and competitive double auctions*, Proc. European Symp. on Algorithms, 2002.

[10] R.J. Edell, N. McKeown and P. Varaiya *Billing users and pricing for TCP*, IEEE J. on Selected Areas in Communications, 13:7(1995), pp. 1162–1175.

[11] T. Gresik and M. Satterthwaite, *The rate at which market converges to efficiency as the number of traders increases: An asymptotic result for optimal trading mechanism*, J. of Economic Theory, 48(1989), pp. 304–332.

[12] F. Gul and A. Postlewaite, *Asymptotic efficiency in large exchange economies with assymetric information*, Econometrica, 60:6(1992), pp. 1273–1292.

[13] M. Jackson, *Incentive compatibility and competitive allocations*, Economic Letters, 40(1992), pp. 299–302.

[14] M. Jackson and A. Manelli, *Approximately competitive equilibria in large finite economies*, J. of Economic Theory, 77:2(1997), pp. 354–376.

[15] R. Jain, A. Dimakis and P. Varaiya, *On the existence of competitive equilibria in bandwidth markets*, Proc. Allerton Conf., 2002.

[16] F. P. Kelly, *Charging and rate control for elastic traffic*, European Trans. on Telecommunications, 8:1(1996), pp. 33–37.

[17] F. P. Kelly, *Mathematical modelling of the Internet*, In: *Mathematics Unlimited: 2001 and beyond*, Springer-Verlag, 2001.

[18] F. P. Kelly, A. K. Maullo, and D. K. H. Tan, *Rate control in communication networks: Shadow prices, proportional fairness and stability*, J. of the Operational Research Soc., 49:1(1998), pp. 237–252.

[19] P. Klemperer, *Auction theory: A guide to the literature*, J. of Economic Surveys, 13:3(1999), pp. 227–286.

[20] R. Kuckro, *Bandwidth trading: Down but not out*, Business Communications Review, pp. 30–34, April, 2002.

[21] S. H. Low, *A duality model of TCP and queue management algorithms*, IEEE/ACM Trans. on Networking, 11:4(2003), pp. 525–536.

[22] S. H. Low and D. E. Lapsley, *Optimization flow control I: Basic algorithm and convergence*, IEEE/ACM Trans. on Networking, 7:6(1999), pp. 861–75.

[23] S. H. Low and P. Varaiya, *A new approach to service provisioning in ATM networks*, IEEE/ACM Trans. on Networking, 1:5(1993), pp. 547–553.

[24] J. K. MacKie-Mason and H. Varian, *Pricing congestible network resources*, IEEE J. on Selected Areas in Communication, 13:7(1995), pp. 1141–1149.

[25] A. Mas-Colell and X. Vives, *Implementation in economies with a continuum of agents*, Rev. Economic Studies, 60:3(1993), pp. 613–629.

[26] A. Mas-Colell, M. Whinston, and J. Green, *Microeconomic theory*, Oxford University Press, 1995.

[27] D. Parkes, *Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency*, PhD Thesis, University of Pennsylvania, 2001.

[28] L. Pontryagin, V. Boltyanskii, R. Gamkrelidze, and E. Mischenko, *The Mathematical Theory of Optimal Processes*, John Wiley, New York, 1962.

[29] D. Roberts and A. Postlewaite, *The incentives for price-taking behavior in large economies*, *Econometrica*, 44:1(1976), pp. 115–127.

[30] A. Ronen, *Solving Optimization Problems among Selfish Agents*, PhD Thesis, Hebrew University, 2000.

[31] A. Rustichini, M. Satterthwaite, and S. Williams, *Convergence to efficiency in a simple market with incomplete information*, Econometrica, 62:5(1994), pp. 1041–1063.

[32] M. Satterthwaite and S. Williams, *The rate of convergence to efficiency in the buyer's bid double auction as the market becomes large*, Rev. Economic Studies, 56(1989), pp. 477–498.

[33] A. Schrijver, *Theory of Linear and Integer Programming*, Wiley-Interscience, 1986.

[34] P. Varaiya, *Lecture Notes on Optimization*, Nostrand-Holland, 1971.

[35] S. de Vries and R. Vohra, *Combinatorial auctions: A survey*, INFORMS J. of Computing, 15:3(2003), pp. 284–309.

[36] S. Williams, *Existence and convergence to equilibria in the buyer's bid double auction*, Rev. Economic Studies, 58(1991), pp. 351–374.

[37] R. Wilson, *Incentive efficiency of double auctions*, Econometrica, 53:5(1985), pp. 1101–1115.

[38] K. Yoon, *The modified Vickrey double auction*, J. of Economic Theory, 101(2001), pp. 572–584.