A MARKET-BASED APPROACH TO OPTIMAL RESOURCE ALLOCATION IN INTEGRATED-SERVICES CONNECTION-ORIENTED NETWORKS

PANAGIOTIS THOMAS

Advanced Network Technology, Motorola, Inc., Arlington Heights, Illinois 60004, panos.thomas@motorola.com

DEMOSTHENIS TENEKETZIS

Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, Michigan 48109, teneketzis@eecs.umich.edu

JEFFREY K. MACKIE-MASON

Department of Economics and School of Information, University of Michigan, Ann Arbor, Michigan 48109, jmm@umich.edu

(Received August 1999; revisions received June 2000, May 2001; accepted May 2001)

We present an approach to the admission control and resource allocation problem in connection-oriented networks that offer multiple services to users. Users' preferences are summarized by means of their utility functions, and each user is allowed to request more than one type of service. Multiple types of resources are allocated at each link along the path of a connection. We assume that the relation between Quality of Service (QoS) and resource allocation is given, and we incorporate it as a constraint into a static optimization problem. The objective of the optimization problem is to determine the amount of and required resources for each type of service to maximize the sum of the users' utilities. We prove the existence of a solution of the optimization problem and describe a competitive market economy that implements the solution and satisfies the informational constraints imposed by the nature of the decentralized resource allocation problem. The economy consists of four different types of agents: resource providers, service providers, users, and an auctioneer that regulates the prices based on the observed aggregate excess demand. The goods that are sold are: (i) the resources at each link of the network, and (ii) services constructed from these resources and then delivered to users. We specify an iterative procedure that is used by the auctioneer to update the prices, and we show that it leads to an allocation that is arbitrarily close to a solution of the optimization problem in a finite number of iterations.

1. INTRODUCTION

The idea of a single shared physical network that will support multiple heterogeneous applications, that is, applications with different traffic characteristics and different Quality of Service (QoS) requirements, is widely regarded as the way to meet the telecommunication challenges of the future. Packet-switched, connection-oriented networks have been proposed to offer the QoS guarantees in integrated-services networks, because in connectionless networks individual packets may exhibit a significant variation in network service quality.

The challenge in integrated-services connection-oriented networks is to determine general admission and resource allocation schemes that have the following desirable features: (i) they meet the QoS requirements of the services provided by the network; (ii) they efficiently allocate resources by appropriately distributing the QoS among the various resources at each link along the path of a connection; (iii) they satisfy the user's preferences and simultaneously are social-welfare maximizing; and (iv) they satisfy the informational constraints imposed by the network; specifically they take into account the fact that the network is an informationally decentralized system where the number of users is unknown, and each user's preferences and input traffic characteristics are private information. Consequently, the development of efficient admission and resource allocation schemes requires: (a) a quantification of the interaction between resource allocation and QoS requirements; such a quantification will be described by expressions that are function of the allocation of multiple resources at each link, as well as the allocation of resources along multilink paths; and (b) the discovery of mechanisms that delegate resource allocation decisions to individual users, yet they lead to social welfare maximizing solutions.

The above considerations have guided most of the research on admission control and resource allocation in integrated-services networks. The same considerations, in particular (iv) and (b), have led to the use of microeconomic methods in the investigation of the aforementioned problems. Two basic microeconomic approaches have been used for the development of efficient decentralized resource allocation schemes in integrated-services networks: *resource-directed* and *price-directed* (Hurwicz

Subject classifications: Networks: pricing schemes for resource allocation. Programming: nonlinear, algorithms. Economics: resource allocation in integrated-services networks.
Area of review: TELECOMMUNICATIONS

1973). In the resource-directed approach, each user (agent) computes the *marginal* values for his current resources, and communicates them to the rest of the users. The allocation is then changed so that users with an above average marginal utility receive more of this resource and users with a below average marginal utility receive less. This approach has been used in Kurose and Simha (1989) to develop decentralized algorithms for optimally allocating a single resource to a set of interconnected computing agents. In the price-directed approach, an initial allocation of resources is made and an arbitrary set of systemwide initial resource prices is chosen. Prices then are iteratively changed to accommodate the "demands" for resources until the total demand for a resource exactly equals the total amount available. Most of the results on decentralized resource allocation currently available in the literature are based on the price-directed approach (Cocchi et al. 1993, MacKie-Mason and Varian 1995b, Jordan and Jiang 1995, Jiang and Jordan 1995, Wang et al. 1997, Murphy and Murphy 1994, Murphy et al. 1994, Parris et al. 1992, Parris and Ferrari 1992, Kelly 1994, Kelly et al. 1998, Gupta et al. 1997, Courcoubetis et al. 1997, MacKie-Mason and Varian 1995a, Low and Varaiya 1993, de Veciana and Baldick 1998, Thomas and Teneketzis 1997). Because we follow the price-directed approach to resource allocation in this paper, we critically review the results reported in (Cocchi et al. 1993, MacKie-Mason and Varian 1995b, Jordan and Jiang 1995, Jiang and Jordan 1995, Wang et al. 1997, Murphy and Murphy 1994, Murphy et al. 1994, Parris et al. 1992, Parris and Ferrari 1992, Kelly 1994, Kelly et al. 1998, Gupta et al. 1997, Courcoubetis et al. 1997, MacKie-Mason and Varian 1995a, Low and Varaiya 1993, de Veciana and Baldick 1998, Thomas and Teneketzis 1997) so that we can point out to the contributions of our work.

The work currently available on decentralized resource allocation by price-directed methods has addressed, either by analysis (Jiang and Jordan 1995, Wang et al. 1997, Murphy et al. 1994, Kelly 1994, Kelly et al. 1998, de Veciana and Baldick 1998, Gupta et al. 1997, Courcoubetis et al. 1997, MacKie-Mason and Varian 1995a, Low and Varaiya 1993, Thomas and Teneketzis 1997), or simulation and analysis (Cocchi et al. 1993, Murphy and Murphy 1994, Parris et al. 1992, Parris and Ferrari 1992, Gupta et al. 1997), a subset of the issues outlined in the second paragraph of this section. A significant part of this work has dealt with single link networks (Wang et al. 1997, Murphy and Murphy 1994, Parris et al. 1992, de Veciana and Baldick 1998), or with the allocation of a single resource per connection (Wang et al. 1997, Murphy and Murphy 1994, Parris et al. 1992, de Veciana and Baldick 1998, Jiang and Jordan 1995, Murphy et al. 1994, Kelly et al. 1998, Gupta et al. 1997, MacKie-Mason and Varian 1995a).

In several papers (Jiang and Jordan 1995, Wang et al. 1997, Murphy et al. 1994, Kelly et al. 1998, Gupta et al. 1997, Low and Varaiya 1993, de Veciana and Baldick 1998, Thomas and Teneketzis 1997) the following general philosophy to resource allocation by price-directed methods has been adopted: (1) formulate a centralized constrained optimization problem where the objective is the maximization of a social welfare function and the constraints are imposed by the QoS requirements and the availability of network resources; and (2) use pricing methods to devise a decentralized scheme that realizes the solution of the centralized problem and satisfies the informational constraints imposed by the network. The existence of a solution to the centralized problem is shown, and market methods are used to structure and develop the solution. The existence of a set of prices that induce users to request the "optimal" allocation is established and, in some cases (Wang et al. 1997, Gupta et al. 1997, Low and Varaiya 1993, de Veciana and Baldick 1998, Thomas and Teneketzis 1997), an iterative scheme for adjusting the prices based on the users' requests is described. However, none of the papers specify a mechanism to force the successive prices to converge to the "optimal" set of prices.

In this paper, we follow the price-directed approach and the philosophy presented in the previous paragraph, to address the admission control and resource allocation problem in integrated-services networks. Our formulation of the admission control and resource allocation problem captures the issues and considerations discussed in the second paragraph of this section. In particular, we consider a connection-oriented network that offers multiple services to users. Users' preferences are summarized by means of their utility functions, and each user is allowed to request more than one type of service. No assumption is made on the functional form of the utility functions, although some mild regularity conditions are imposed. We assume that the relation between QoS and resource allocation is given (see discussion in §2), and we incorporate it as a constraint into a static optimization problem. The objective of the optimization problem is to determine the amount and, required resources for each type of service to maximize the sum of the users' utilities. We prove the existence of a solution of the optimization problem, and describe a competitive market economy that implements the solution and satisfies the informational constraints imposed by the nature of the decentralized resource allocation problem. The economy consists of four different types of agents: resource providers, service providers, users, and an auctioneer that regulates the prices based on the observed aggregate excess demand. The goods that are sold are: (i) the resources at each link of the network; and (ii) services constructed from these resources and then delivered to users. We specify an iterative procedure that is used by the auctioneer to update the prices, and we prove its convergence. In particular, we show that it leads to an allocation that is arbitrarily close to a solution of the optimization problem in a finite number of iterations.

The contributions of this paper are: (1) the formulation of a general optimization problem that has the following features: (a) incorporation of the QoS requirements along

paths; (b) allocation of resources along paths with multiple links; (c) allocation of multiple type of resources at each link; (d) independent self-interest users that request multiple types of service; (2) the proof of existence of a welfare maximizing solution; and (3) the construction of an algorithm for finding a welfare maximizing solution with arbitrary precision in a finite number of steps.

The remainder of the paper is organized as follows: In §2 we formulate the optimization problem. In §3 we prove the existence of a solution, and in §4 we describe (§4.1) and analyze (§4.2) the competitive market economy that leads to an optimum allocation. In §5 we summarize the results, discuss their implications, and present some open problems that arise from this approach.

2. PROBLEM FORMULATION

Consider a communication network that consists of a set $\mathbf{L} = \{1, 2, \ldots, L\}$ of links. The network offers one-way connections to a set $\mathbf{N} = \{1, 2, \ldots, N\}$ of users. There is also a set $\mathbf{M}_i = \{1, 2, \ldots, M_i\}$ of different *types* of connections associated with each user i. Each type is characterized by the origin, the destination, and the quality level of the connection. By quality level we mean specific quality of service (QoS) guarantees (for example guarantees on maximum packet loss and delay), that the network is obliged to provide once the connection is accepted. Users can request a multiple number of connections over the network.

We represent user i's request by the vector $x^i = (x_1^i, x_2^i, \ldots, x_{M_i}^i)$, that takes values in $\mathbb{R}_+^{M_i}$. In particular, x_j^i is the number of type j connections requested by user i. User's i preferences on $\mathbb{R}_+^{M_i}$ are summarized by means of a quasilinear utility function $x_0^i + u_i(x^i)$, where $x_0^i > 0$ is the numeraire commodity (Mas-Colell et al. 1995). The choice of representing users' preferences by quasilinear objective functions imposes the constraint that there are no income effects on network service demand; that is, changes in income or budget available to the users does not change the amount of network services they wish to purchase. This is a typical simplifying assumption in the economic literature when the budget share of the services of interest is small, e.g., when network services are only a relatively small amount of the users' total expenditures.

Users announce their requests to the network during the call setup phase. The network then, assigns virtual circuits (routes) to each origin-destination pair. We denote by $\mathbf{V}^{i,j}$ the set of links that belong to the route of a type j connection requested by user i. Once the routes are established the network needs to determine the number of connections that will be accepted as well as the amount of resources (e.g., bandwidth and buffers) that need to be allocated along the path of each connection in order to guarantee its requested quality level. The allocation of resources must also maximize a social welfare function $\sum_{i \in \mathbf{N}} u_i(x^i)$, which is the total utility of the users.[1]

We assume that there is a set $\mathbf{K} = \{1, 2, \ldots, K\}$ of different *types* of resources that the network allocates at each link in order to meet the QoS requirements of a connection that is accepted. For example, if bandwidth and buffers are the two types of resources that the network allocates in order to establish connections, then $\mathbf{K} = \{1, 2\}$. In this case, $k = 1$ refers to bandwidth and $k = 2$ refers to buffers. We denote by $r_{l,k}^{i,j}$ the amount of a type k resource that is reserved at link l for a connection of type j requested by user i. Also, the maximum amount of a type k resource that can be allocated at link l is denoted by $R_{l,k} > 0$. We denote by $\mathbf{F}^{i,j}$, the set of all resource allocations $r^{i,j} \triangleq \{r_{l,k}^{i,j}\}_{l \in \mathbf{L}, k \in \mathbf{K}}$, that guarantee the quality level specified by the type j connection requested by user i.[2] This set results from relations that describe the interaction between resource allocation and QoS requirements along routes; such relations can be found in Thomas and Teneketzis (1997). In this paper we assume that the sets $\mathbf{F}^{i,j}$ of all resource allocations that guarantee the end-to-end QoS requirements are known and given.

The following assumptions are made on the utility functions $u_i(\cdot)$ and the sets $\mathbf{F}^{i,j}$:

ASSUMPTION (A1). *The function $u_i(\cdot)$ is continuous, differentiable, locally nonsatiated,*[3] *and strictly concave in* $\mathbf{X}$, *for all* $i \in \mathbf{N}$.

ASSUMPTION (A2). *The sets* $\mathbf{F}^{i,j}$, $i \in \mathbf{N}$, $j \in \mathbf{M}_i$, *are compact and strictly convex.*[4]

Assumption (A1) implies that the underlying preference relations are rational, continuous, strictly convex, and locally nonsatiated on $\mathbf{X}$ (Mas-Colell et al. 1995). Assumption (A2) is technical and is needed to establish the existence of optimal allocation strategies that can be obtained as the result of a tatônment process. In general there may not exist a convex relation between QoS and resource allocation along paths within the network; see Thomas and Teneketzis (1997). In such a case, we take the sets $\mathbf{F}^{i,j}$, $i \in \mathbf{N}$, $j \in \mathbf{M}_i$ to be the maximum strictly convex subsets of the sets specified by the relationship between QoS and resource allocation. A discussion of this issue is presented in §5.

Based on the above, we formulate the following optimization problem:

$$\max_{x,r} \sum_{i \in \mathbf{N}} u_i(x^i) \qquad \textbf{(MAX 1)}$$

subject to

$$x^i \in \mathbf{X}^i, \quad i \in \mathbf{N}, \tag{1}$$

$$r^{i,j} \in \mathbf{F}^{i,j}, \quad i \in \mathbf{N}, j \in \mathbf{M}_i, \tag{2}$$

$$\sum_{i \in \mathbf{N}} \sum_{j \in \mathbf{M}_i} x_j^i r_{l,k}^{i,j} \leqslant R_{l,k}, \quad l \in \mathbf{L}, k \in \mathbf{K}, \tag{3}$$

where $(x, r) \triangleq \{x_j^i, r^{i,j}\}_{i \in \mathbf{N}, j \in \mathbf{M}_i}$,

$$\mathbf{X}^i \triangleq \{x^i \in \mathbb{R}^{M_i} : 0 \leqslant x_j^i \leqslant B < \infty, \ j \in \mathbf{M}_i\}, \tag{4}$$

and B is a sufficiently large real number. The restriction of x^i to the compact space $\mathbf{X}^i$ is made for technical reasons that appear in the proof of existence of a solution to

(**MAX 1**). This assumption does not affect the optimality condition, since an infinite value for x_j^i would violate the constraint on the available resources—see (3), and notice that the amount of resources $r_{l,k}^{i,j}$, $k \in \mathbf{K}$ needed at each link $l \in \mathbf{V}^{i,j}$ to establish the type j connection for user i, cannot all be zero.

In summary, our problem formulation incorporates the following features: (1) the QoS requirements are included into the service provisioning problem, via constraints (2); (2) multiple types of resources $r_{l,k}^{i,j}$, $k \in \mathbf{K}$ are allocated at each link l of the network; (3) a virtual path $\mathbf{V}^{i,j}$ consists of multiple links; and (4) a user can request multiple types of services based on his own preferences. A critique and discussion of the assumptions and features of the problem formulation are presented in §5.

In the next section we establish the existence of a solution to (**MAX 1**).

3. EXISTENCE OF A WELFARE-MAXIMIZING SOLUTION

The main result of this section is summarized by the following theorem:

THEOREM 3.1. *There exists a solution to problem* (**MAX 1**).

PROOF. We first fix r to an allocation that satisfies (2). Then, Weierstrass's Theorem (Simon and Blume 1994, p. 823) guarantees that the problem (**MAX 2**) defined below, has a solution.

$$\max_x \sum_{i \in \mathbf{N}} u_i(x^i) \qquad \textbf{(MAX 2)}$$

subject to

$$x \in \mathbf{X}, \tag{5}$$

$$\sum_{i \in \mathbf{N}} \sum_{j \in \mathbf{M}_i} x_j^i r_{l,k}^{i,j} \leqslant R_{l,k}, \quad l \in \mathbf{L}, k \in \mathbf{K}, \tag{6}$$

where (5) is an abbreviation for (1).

If we show that the maximum value function $U(r) \triangleq \sum_{i \in \mathbf{N}} u^i(x^i(r))$, i.e. the objective function of (**MAX 2**) evaluated at a maximizer $x(r)$, is a continuous function of r, then the continuity of $U(r)$ together with Assumption (A2) and the Weierstrass' theorem will guarantee the existence of a solution to the following problem:

$$\max_r U(r) \qquad \textbf{(MAX 3)}$$

subject to

$$r^{i,j} \in \mathbf{F}^{i,j}, \quad i \in \mathbf{N}, j \in \mathbf{M}_i. \tag{7}$$

The existence of a solution to the original problem (**MAX 1**) is then established, since a solution to (**MAX 3**) is also a solution to (**MAX 1**).

We now show that $U(r)$ is continuous. For that matter we use Proposition 3.1, that is given below. The proof of the proposition is presented in Appendix A.

PROPOSITION 3.1. *Consider the correspondence $\Gamma : \mathbf{R} \mapsto \mathbf{X}$ defined by $\Gamma r = \{x \in \mathbf{X} : g(x, r) \leqslant 0\}$, where $\mathbf{X}$ is compact and g is a continuous function in $\mathbf{X} \times \mathbf{R}$. Assume that for all $r \in \mathbf{R}$ there exists $x \in \mathbf{X}$ with $g(x, r) < 0$. Then Γ is continuous in $\mathbf{R}$.*

Define the correspondence

$$\Gamma : \mathbf{R} \mapsto \mathbf{X}, \quad \Gamma r = \{x \in \mathbf{X} : g(x, r) \leqslant 0\},$$

where $\mathbf{R} \triangleq \mathbb{R}_+^{D \cdot L}$, $\mathbf{X}$ is the same as in (5), and $g : \mathbf{X} \times \mathbf{R} \to \mathbb{R}^{K \cdot L}$ is defined by $g_{l,k}(x, r) = \sum_{i \in \mathbf{N}} \sum_{j \in \mathbf{M}_i} x_j^i r_{l,k}^{i,j} - R_{l,k}$, $l \in \mathbf{L}$, $k \in \mathbf{K}$. The set $\mathbf{X}$ is compact. Furthermore, g is continuous on $\mathbf{X} \times \mathbf{R}$ and $g_{l,k}(0, r) < 0$ for all $l \in \mathbf{L}$ and $k \in \mathbf{K}$. Therefore the conditions of Proposition 3.1 are satisfied and Γ is a continuous correspondence in $\mathbf{R}$. By Assumption (A1) the objective function of (**MAX 2**) is continuous on $\mathbf{X}$. The Maximum Theorem (Border 1989, Theorem 12.1, p. 64) then implies that the correspondence

$$\Phi : \mathbf{R} \mapsto \mathbf{X},$$

$$\Phi r = \left\{ x \in \Gamma r : x \text{ maximizes } \sum_{i \in \mathbf{N}} u_i(x^i) \text{ on } \Gamma r \right\},$$

is closed and upper hemicontinuous and that the function

$$U(r) = \sum_{i \in \mathbf{N}} u_i(x^i) \quad \text{for} \quad x \in \Phi r,$$

is continuous on $\mathbf{R}$. This completes the proof of Theorem 3.1. □

4. A COMPETITIVE MARKET INTERPRETATION

In the previous section we established the existence of a (welfare-maximizing) solution to problem (**MAX 1**). We are now interested in the existence of an algorithm that can lead to a solution of (**MAX 1**), and satisfies the informational constraints imposed by the nature of the network. To establish the existence of such an algorithm we proceed as follows: First, we describe a competitive market economy that consists of three types of agents and an auctioneer that regulates the prices through an iterative procedure that is based on the aggregate demand and supply. Then, we prove that the procedure used by the auctioneer leads to an allocation that is arbitrarily close to a solution of (**MAX 1**), in a finite number of steps. This result is given by Theorem 4.1, which is presented later in the section. We now proceed with the description of the market.

4.1. Description of the Market

The economy consists of four different types of agents: resource providers, service providers, users, and an auctioneer that was mentioned before. We assume that the resource providers, the service providers, and the users are price takers. That is, they behave as if their behavior has, and can have no effect on the equilibrium prices that are reached by the market allocation process. The raw materials that are

sold are the resources at each link of the network. The price for a type k resource at link l is denoted by $\lambda_{l,k}$. These resources are treated as if they were owned by the resource providers, who then sell them to the service providers. The service providers, in turn, buy resources at each link in order to produce flow services, which they sell to end users. For example, a service provider might want to produce a "bounded-delay guaranteed" service from node A to node B. To do so, it will purchase the use of buffers and bandwidth at each link along some path between A and B in amounts sufficient to provide the promised service. Users purchase services from service providers; they do not purchase the primitive resources, nor do they need to know how many buffers or how much bandwidth is being dedicated to their service, nor the path of links that has been established for their flows.

Although we construct the problem as one in which there are two markets—for network resources and for services—our algorithm works by solving only for a vector of resource prices. We are able to do this because the assumption of price taking that we impose on our agents enables us to directly derive the service prices as a deterministic function of the resource prices.[5] Nonetheless, the characterization of the problem in terms of the two separate markets for separate resources and services is significant, because the information hiding conditions described above are maintained: The users do not need to know the resource prices, nor the resources being purchased and combined to provide services on their behalf. In the sequel we describe each of the above agents separately.

Resource Providers. Resource providers own the resources (bandwidth, buffers, etc.) at each link of the network. We assume that they are price takers and that there is no cost associated with the supply of their resources to the market. The aggregate supply $y(\lambda)$ at prices $\lambda \triangleq \{\lambda_{l,k}\}_{l\in\mathbf{L},k\in\mathbf{K}}$ is given by

$$y(\lambda) \in \arg\max_{y\in Y} \sum_{l\in\mathbf{L}} \sum_{k\in\mathbf{K}} \lambda_{l,k} y_{l,k}, \tag{8}$$

where $Y \triangleq \{y \in \mathbb{R}^{KL} : 0 \leqslant y_{l,k} \leqslant R_{l,k}, l \in \mathbf{L}, k \in \mathbf{K}\}$, and K, L are the cardinal numbers of the sets $\mathbf{K}$, $\mathbf{L}$ respectively. Since there is no cost associated with the production, (8) is trivially satisfied for every $\lambda \geqslant 0$, when $y_{l,k} = R_{l,k}$, for all l and k. Therefore,

$$y(\lambda) = \{R_{l,k}\}_{l\in\mathbf{L},k\in\mathbf{K}} \quad \text{for every } \lambda \geqslant 0. \tag{9}$$

Service Providers. Service providers accept requests for connections from users. Each request is characterized by the triplet: origin, destination, and quality level, and is indexed by the pair i, j, where $i \in \mathbf{N}$ refers to the user and $j \in \mathbf{M}_i$ refers to the type of connection. For each user request i, j, the service provider establishes a one-way route $V^{i,j}$ between the origin and the destination and allocates resources $r^{i,j}(\lambda)$ that minimize the cost of the connection and guarantees the requested quality level. We assume that the service providers are price takers and that they can freely enter or exit the market. An implication of this is that their profit is zero. Based on the above, service providers allocate resources to each type of connection by determining

$$r^{i,j}(\lambda) \in \arg\min_{r^{i,j}\in\mathbf{R}^{KL}} \sum_{l\in V^{i,j}} \sum_{k\in\mathbf{K}} \lambda_{l,k} r^{i,j}_{l,k}, \tag{10}$$

for each $i \in \mathbf{N}$ and $j \in \mathbf{M}_i$. Once resources have been allocated, service providers announce to users the following prices:

$$p^i_j(\lambda) \triangleq \sum_{l\in V^{i,j}} \sum_{k\in\mathbf{K}} \lambda_{l,k} r^{i,j}_{l,k}(\lambda), \quad i \in \mathbf{N}, j \in \mathbf{M}_i. \tag{11}$$

Users. Users request one way connections from the service providers. As mentioned before, each request is characterized by the triplet: origin, destination, and quality level. Service providers announce prices $p^i_j(\lambda)$, $i \in \mathbf{N}$, $j \in \mathbf{M}_i$, and users respond by requesting $x(\lambda)$ number of connections. Users' demands satisfy

$$x^i(\lambda) \in \arg\max_{x^i\in\mathbf{X}^i} \left[u_i(x^i) - \sum_{j\in\mathbf{M}_i} x^i_j p^i_j(\lambda) \right] \quad \text{for all } i \in \mathbf{N}. \tag{12}$$

For each $i \in \mathbf{N}$, the optimization problem in (12) is equivalent to the standard utility maximization problem (Mas-Colell et al. 1995).

Auctioneer. The auctioneer regulates the market prices based on the aggregate excess demand vector $z(\lambda)$

$$z_{l,k}(\lambda) \triangleq \sum_{i\in\mathbf{N}} \sum_{j\in\mathbf{M}_i} (x^i_j(\lambda) r^{i,j}_{l,k}(\lambda)) - y_{l,k}(\lambda), \quad l \in \mathbf{L}, k \in \mathbf{K}. \tag{13}$$

Later, in the proof of Theorem 4.1, we show that $z(\lambda)$ is a continuous function of λ (see Proposition 4.1).

During the setup phase users announce to service providers the types of connections (origin, destination, quality level) that they are willing to buy. Then, the service providers establish the routes $V^{i,j}$ for each type of connection i, j, $i \in \mathbf{N}$, $j \in \mathbf{M}_i$. After the routes have been established, an iterative procedure begins, during which the auctioneer announces prices λ, and then uses the corresponding aggregate excess demand function $z(\lambda)$ to compute the next set of prices λ. Every time a new price vector λ is announced by the auctioneer, the service providers determine the amount of resources $r(\lambda)$ that according to (10) guarantee the requested quality level and minimize the cost per connection. Then, the prices (11) are announced to the users. Based on these prices users respond by requesting $x(\lambda)$ number of connections, that satisfy (12). The response of the resource providers to any price vector announced by the auctioneer is given by (9).

A detailed description of the algorithm used by the auctioneer to update market prices is presented in Appendix B. In Figure 1, we present in block-diagram form the auctioneer's algorithm and provide an explanation about the variables that appear in the description of the algorithm.

Figure 1. Auctioneer's algorithm.

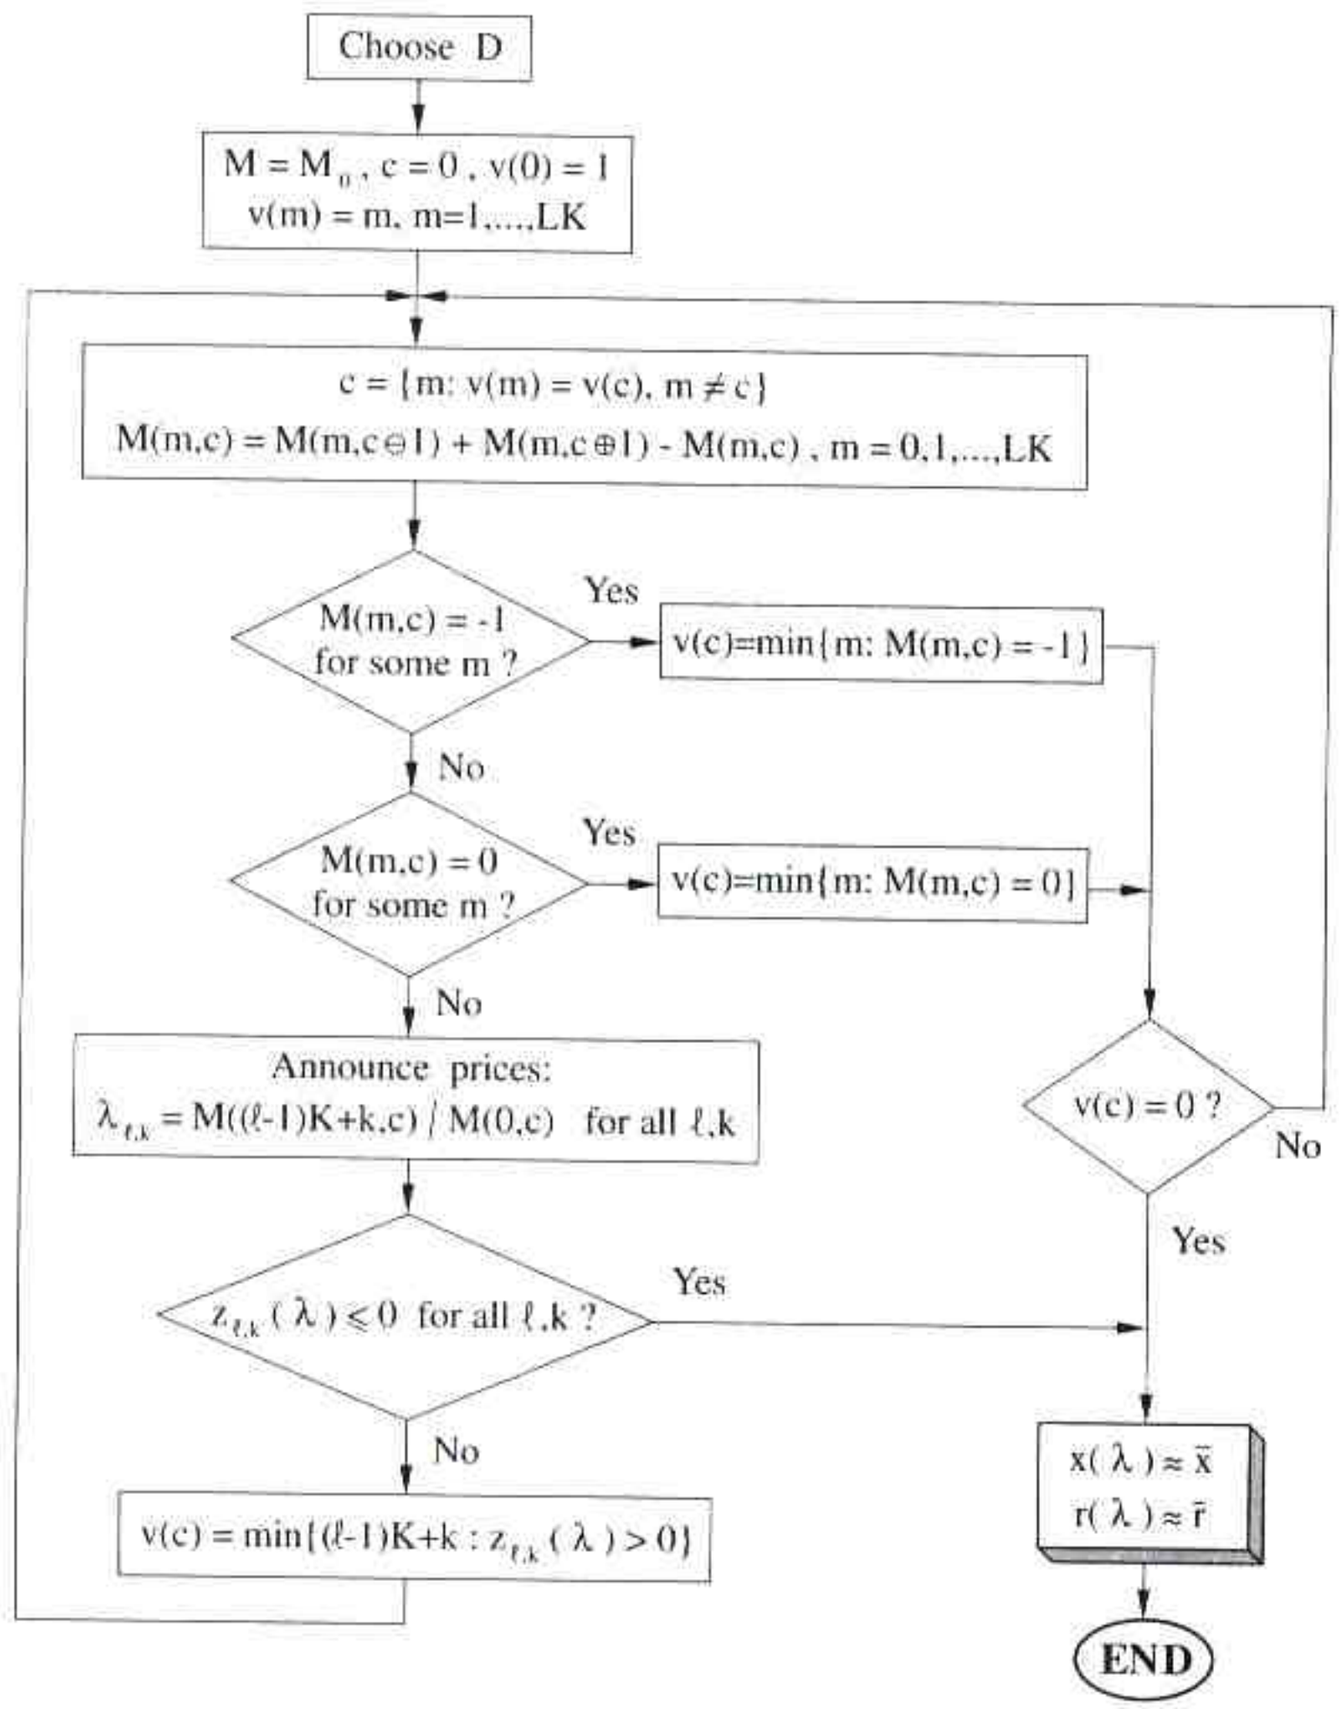

The algorithm starts by choosing a positive integer value for the parameter D. As shown in §4.2, where we prove the convergence of the algorithm, the value of D determines how close is the allocation that results from the algorithm to a solution $(\bar{x}, \bar{r})$ of problem (**MAX 1**). In particular, we show that by increasing the value of D, we can come arbitrarily close to $(\bar{x}, \bar{r})$. However, by increasing D we also increase the number of iterations needed for the algorithm to terminate. Two other elements used by the algorithm are the $(LK+1) \times (LK+1)$ matrix M and the $(LK+1)$ column ν. The following notation is used for the entries of these matrices: $M(i, j)$ is the element in row i and column j of matrix M, and $\nu(i)$ is the element in row i of column ν. Initially, matrix M is assigned the value

$$M_0 = \begin{bmatrix} D & D+1 & \cdots & D+1 \\ 0 & -1 & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \ldots & -1 \end{bmatrix}. \tag{14}$$

Finally, the symbols $\ominus, \oplus$ denote modulo LK subtraction and addition, respectively. The proof of convergence of the auctioneer's algorithm is given by Theorem 4.1, which is presented in the section that follows.

4.2. Analysis of the Market

4.2.1. Preliminaries. In this section we show that the recursive process, used by the auctioneer and described in §4.1 and Appendix B, leads in a finite number of iterations, to an allocation that is arbitrarily close to a solution of problem (**MAX 1**).

We proceed as follows. First, we present briefly some known results that are crucial for the analysis of the market. Then, we prove Theorem 4.1, which is the main result of this section. The known results used in the analysis of the market can be found in Scarf (1973, Chapters 2–4, 6, Appendix 1). Here we present them in a way that suits the needs of our analysis.

Define the simplex S:

$$S \triangleq \left\{ q \in \mathbb{R}_+^{LK+1} : \sum_{m=0}^{LK} q_m = 1 \right\},$$

where L and K are, as before, the cardinal numbers of the sets $\mathbf{L}$ and $\mathbf{K}$ respectively, and denote by s^m, $m = 0, 1, \ldots, LK$, its sides, that is,

$$s^m \triangleq \{q \in S : q_m = 0\}, \quad m = 0, 1, \ldots, LK.$$

Consider the regular grid of vectors on S, that is the list P of all vectors of the form $(n_0/D, n_1/D, \ldots, n_{LK}/D)$, with n_i representing nonnegative integers summing to D. We introduce a special type of subsimplex in S, called the *primitive set*, that is a key concept for the proof of Theorem 4.1.

DEFINITION 4.1. The $LK + 1 - n$ vectors $q^{j_0}, \ldots, q^{j_{LK-n}}$ of P along with the n sides $s^{i_1}, \ldots, s^{i_n}$ of S, form a *primitive set* in S if no vector $q \in P$ is interior to the simplex defined by $x_{i_1} \geqslant 0, \ldots, x_{i_n} \geqslant 0$ and $q_m \geqslant \min\{q_m^{j_0}, \ldots, q_m^{j_{LK-n}}\}$, for $m \neq i_1, \ldots, i_n$.

We now introduce the family $\mathcal{M}$ of $(LK+1) \times (LK+1)$ matrices

$$\begin{bmatrix} n_{0,0} & \cdots & n_{0,LK} \\ n_{1,0} & \cdots & n_{1,LK} \\ \vdots & & \vdots \\ n_{LK,0} & \cdots & n_{LK,LK} \end{bmatrix}, \tag{15}$$

that have the following properties: The entries are integers between -1 and $D+1$ and the column sums are equal to D. For the purpose of our analysis we associate to each column of $M \in \mathcal{M}$, a vector or a side from the simplex S and also a label from the set $\{0, \ldots, LK\}$.

Association of a Vector or Side from the Simplex with a Column. The following rule is used to associate to a column of $M \in \mathcal{M}$, either a vector from the regular grid P, or a side of the simplex S: If column c has at least one negative entry then it is associated with side s^i, where $i = \min\{m : M(m, c) = -1\}$; if column c has nonnegative entries, it is associated with the vector $(M(0, c)/D, \ldots, M(LK, c)/D)$.

Association of a Column with a Label. To specify the rule that associates each column of $M \in \mathcal{M}$ with a label from the set $\{0, \ldots, LK\}$, we first need to make the following definitions:

For each $q \in S$ with $q_0 > 0$ define the price vector λ':

$$\lambda' = \{\lambda'_{l,k}\}_{l \in \mathbf{L}, k \in \mathbf{K}} \triangleq \left\{\frac{q_1}{q_0}, \frac{q_2}{q_0}, \ldots, \frac{q_{LK}}{q_0}\right\}. \tag{16}$$

Define also the following subsets of S:

$$C_0 \triangleq \{q \in S : q_0 = 0 \text{ or } z_{l,k}(\lambda') \leqslant 0, \text{ for all } l \in \mathbf{L}, k \in \mathbf{K}\}, \tag{17}$$

$$C_{l,k} \triangleq \{q \in S : q_0 > 0 \text{ and } q_{(l-1)K+k} = 0 \text{ or } z_{l,k}(\lambda') \geqslant 0\}, \quad l \in \mathbf{L}, k \in \mathbf{K}, \tag{18}$$

where

$$z_{l,k}(\lambda) = \sum_{i \in \mathbf{N}} \sum_{j \in \mathbf{M}_i} x_j^i(\lambda) r_{l,k}^{i,j}(\lambda) - R_{l,k}, \quad l \in \mathbf{L}, k \in \mathbf{K}, \tag{19}$$

and $r(\lambda)$, $x(\lambda)$ are solutions of the problems:

$$\min_{r^{i,j} \in \mathbf{F}^{i,j}} \sum_{l \in \mathbf{L}} \sum_{k \in \mathbf{K}} \lambda_{l,k} r_{l,k}^{i,j} \text{ for all } i \in \mathbf{N}, j \in \mathbf{M}_i, \tag{20}$$

$$\max_{x^i \in \mathbf{X}^i} \left[u_i(x^i) - \sum_{j \in \mathbf{M}_i} x_j^i \sum_{l \in \mathbf{L}} \sum_{k \in \mathbf{K}} \lambda_{l,k} r_{l,k}^{i,j}(\lambda) \right] \text{ for all } i \in \mathbf{N}. \tag{21}$$

The definition of the sets C_0 and $C_{l,k}$, $l \in \mathbf{L}, k \in \mathbf{K}$ assumes that the sign of the aggregate excess demand $z_{l,k}(\lambda)$ is well defined for all $l \in \mathbf{L}, k \in \mathbf{K}$. Later in the section we show that $z(\lambda)$ is a continuous function of λ (Proposition 4.1). This makes the above definitions feasible, because $z(\cdot)$ is a function of λ and not a correspondence.

We assign to each column c of M the following label $\nu(c)$: If column c is associated with side s^i then $\nu(c) = i$; if c is associated with vector q then, if $q \in C_0$ set $\nu(c) = 0$, otherwise set $\nu(c) = \min\{(l-1)K + k : q \in C_{l,k}\}$.

The main result of this section is based on a combinatorial theorem due to Scarf (1973, Theorem 2.5.1). Scarf's theorem and the main idea of the proof, that are both crucial for the analysis of our algorithm, are presented in the next paragraph.

Consider a labeling process that assigns to each side s^m $(m = 0, \ldots, LK)$ of the simplex S the integer m, and to each vector of the regular grid P an arbitrary integer from the set $\{0, \ldots, LK\}$. Then according to Scarf (1973, Theorem 2.5.1), there exists a primitive set each of whose vectors or sides has a different label. The proof of this theorem is based on the following iterative procedure, (**IP1**):

(**IP1**) 1. Define a unique replacement operation for a single side or vector of a primitive set in S, that leads to another primitive set in S.

2. Start with a primitive set all of whose vectors and sides have distinct (nonzero) labels, except for a single pair whose labels are identical.

3. Replace one of the elements (vector or side), in the pair with the identical labels, using the operation defined in Step 1.

4. If the new vector or side has a zero label, terminate the procedure. Otherwise, replace the element in the primitive set that has the same label with the element that was just brought in, and repeat Step 4.

The above procedure will terminate only with a primitive set of the required type. It is shown in Scarf (1973, pages 47–48) that the algorithm never returns to the same position and that every required replacement can be carried out. Therefore, since there are a finite number of possible primitive sets that can be composed from the sides of the simplex and the vectors of the regular grid, the algorithm will terminate in a finite number of steps.

The concepts and results presented above provide the background necessary to proceed with the analysis of the market described in §4.1.

4.2.2. Convergence of the Auctioneer's Algorithm. The main result of §4.2 is given by the following theorem.

THEOREM 4.1. *The algorithm described in §4.1 and depicted in Figure 1 leads, in a finite number of steps, to a price vector λ that induces an allocation $x(\lambda)$, $r(\lambda)$ which is arbitrary close to a solution $\bar{x}$, $\bar{r}$ of problem* (**MAX 1**).

PROOF. The proof is based on the combinatorial result (Scarf 1973, Theorem 2.5.1) presented above, and proceeds in three steps:

(i) We first define a labeling process for the sides of S and the vectors of P, that satisfies the conditions of Theorem 2.5.1 in Scarf (1973). This establishes the existence of a primitive set T in S with vectors and sides that have distinct labels.

(ii) We show that the auctioneer's algorithm, described in §4.1, leads in a finite number of steps to the set T (as it implements the steps of the iterative procedure (**IP1**)).

(iii) We prove that for any $q \in T$ with $q_0 > 0$, the price vector given by (16), induces an allocation that is arbitrarily close to a solution of (**MAX 1**).

We now proceed to prove the statements made in Steps (i)–(iii).

(i) Consider the following labeling process for the sides of S and the vectors of P: Side s^m is labeled with integer m; and $q \in P$ is assigned the label of its associated column of $M \in \mathcal{M}$ (see the two association rules defined above). This labeling process satisfies the conditions of Theorem 2.5.1 in Scarf (1973), thus, it establishes the existence of a primitive set $T \subset S$ such that all of its sides and vectors have distinct labels.

(ii) We show that the auctioneer's algorithm implements the steps of (**IP1**).

By construction, the auctioneer's algorithm is an iterative process with steps (structure) identical to those of (**IP1**). Therefore, we only need to prove that the auctioneer's algorithm satisfies the conditions of Steps 1 and 2 of (**IP1**).

Let $M \in \mathcal{M}$ be a matrix associated with a primitive set in S,[6] and consider the following rule for replacing an element (side or vector) of the primitive set: Find the column c of M associated with the element of the primitive set that is replaced; replace c by the sum of columns $c \ominus 1$ and $c \oplus 1$ minus the column c itself, as in the auctioneer's algorithm; find the vector or side associated with that column and use it as a replacement for the original element of the primitive set. Then, according to Scarf (1973, Chapter 2 and Appendix 1), this replacement operation leads to another primitive set in S. The replacement is also unique. Consequently, the replacement operation in the auctioneer's algorithm satisfies the conditions of Step 1 of (**IP1**).

Let $A_0 \subset S$ be the subsimplex associated with the matrix M_0, introduced in the description of the auctioneer's algorithm (see Equation (14)). Matrix M_0 satisfies the conditions of Theorems 6.2.1 and 6.2.9 in Scarf (1973, note 6), therefore A_0 is a primitive set in S. Furthermore, according to our association rule, all the vectors and sides of A_0 have distinct labels from the set $\{1, \ldots, LK\}$, except for a single pair whose labels are identical. Therefore, A_0, the initially chosen primitive set in the auctioneer's algorithm, meets the conditions of Step 2 of the iterative process (**IP1**).

Consequently, the auctioneer's algorithm leads, in a finite number of steps, to the primitive set T in S with vectors and sides that have distinct labels. The set T contains at least one vector from the regular grid P which is in the interior of the simplex S (i.e., a vector the components of which are all positive).[7]

(iii) We prove that for any $q \in T$ with $q_0 > 0$, the price vector given by (16), induces an allocation that is arbitrarily close to a solution of (**MAX 1**). According to the labeling process, T has a nonzero intersection with each one of the sets C_0, $C_{l,k}$, $l \in \mathbf{L}$, $k \in \mathbf{K}$, defined by (17), (18), respectively. Define

$$C \triangleq C_0 \cap \left(\bigcap_{l \in \mathbf{L},\, k \in \mathbf{K}} C_{l,k} \right).$$

In Proposition 4.2 we prove that C is nonempty, and in Proposition 4.3 we show that for every $q \in C$, the allocation $x(\lambda')$, $r(\lambda')$ (induced by the price vector λ' defined by (16)) solves (**MAX 1**). Consequently, if we make the regular grid sufficiently fine (by increasing the parameter D), we can guarantee that the primitive set T has a nonzero intersection with C. This result, along with the continuity of the functions $x(\lambda')$, $r(\lambda')$ with respect to λ' (shown in Proposition 4.1), prove the statement of Step (iii).

Propositions, 4.1, 4.2, and 4.3, that we referred to above, are presented below. Their proofs are given in Appendices C, D, and E, respectively.

PROPOSITION 4.1. *The allocation $x(\lambda)$, $r(\lambda)$ and the aggregate excess demand $z(\lambda)$ are continuous functions of λ.*

PROPOSITION 4.2. *The set C is nonempty.*

PROPOSITION 4.3. *For every $q \in C$, $x(\lambda')$, $r(\lambda')$ solves* (**MAX 1**).

This concludes the proof of Theorem 4.1. □

5. CONCLUSIONS—REFLECTIONS

We have presented an approach for optimal admission and resource allocation control in multiservice connection-oriented networks.

The main contribution of our work in this paper is the specification of a convergent and decentralized iterative procedure that leads to a solution of a fairly general admission and resource allocation problem.

The main features of our approach are the following:

(1) The objective of the resource allocation process is to maximize the total value of the network to its users.

(2) The agents are price takers in the markets in which they participate.

(3) There is no cost associated with the supply of network resources.

(4) The sets $\mathbf{F}^{i,j}$ are strictly convex, for all i, j.

(5) Resource allocation decisions are based on the solution of a constrained static optimization problem.

We now discuss and critique each one of the above features.

(1) In this analysis we assumed that the objective function of interest was to maximize the sum of individual network users' utility functions. It may not be obvious why this is a reasonable objective to consider.

It is important to realize that our point of view is primarily *normative*, not *descriptive*. That is, we have taken a particular objective function—one which we believe is often reasonable, but see below—and studied whether a network resource pricing scheme exists that can achieve an optimum for that particular function, and how one might *implement that allocation with a market-based algorithm*. Thus, we have demonstrated the feasibility of using pricing to achieve a particular performance goal. We are not claiming that this goal describes any particular actual network environment. Nor are we making the stronger normative claim that this objective function *should* be adopted in any particular setting.

What we have shown is that economic market methods can be harnessed to implement network resource allocations that fulfill a particular objective function. We believe that this work provides a template for developing market *methods for implementing other objective functions* as well. For example, some might think it appropriate for short-run allocation problems, when network resource capacities are relatively fixed and there is little opportunity for entry by competing resource providers, to devise automated market mechanisms that maximize the sum of user utilities plus resource provider profits (possibly with some particular weighting on how this total surplus is shared between the users and the providers). To study this problem we would augment the (**MAX1**) objective function by adding the sum of resource provider profits, and then proceed to analyze the existence and implementation of an optimizing price vector.

We do, in fact, believe that maximizing the sum of user utilities is a reasonable description for a wide variety of network allocation problems. Suppose we are considering a corporate intranet. If the corporation's overall objective is to maximize its profits (in present value), then the appropriate interpretation of our problem is to define each user's "utility" as that user's contribution to corporate profit as a function of the network services it consumes. In other words, the corporation is not (directly) interested in how personally happy an employee is with the network, but on how much value the network enhances the employee's productivity. Then the sum of user utilities will be the contribution of network services to corporate profits, which is precisely the firm's objective function for this part of the *overall management problem. Although it may seem difficult* to come up with a reasonable representation of the effect of network services on each user's contribution to corporate profits, at some level this is precisely the problem corporations need to solve for allocating equipment, office space, subordinates and so forth to each employee—it is well beyond the scope of our research to worry about how the corporation specifically formulates these valuation functions.

In a market context, it might be in the interest of a network resource provider to design a marketing system with the goal of maximizing its profits, rather than maximizing the realized utility of its customers. However, if the network provider is competing with other providers and networks are interconnected, then in long-run equilibrium it will tend to be the case that the successful allocation schemes are those that have the net effect of maximizing users' utilities.

Thus, although our method of using prices to allocate network resources cannot be directly applied to every allo*cation problem with any reasonable objective function*, we believe that it has broad applicability to many existing situations. In any case, when our objective function *is* the desired goal, we have carefully analyzed the existence and implementability of a pricing scheme to support that objective.

(2) In §4.1 we imposed the price-taking assumption to the agents of the market economy. In some network allocation settings price-taking will not be a reasonable assumption, for at least some of the agents. For example, suppose that we are concerned with a single physical net*work owned by a single resource provider. That provider*, then, would be a monopolist. If that provider had the goal of profit-maximization, say, then it would likely know or quickly learn that it could earn more profits by directly manipulating price. In terms of Equation (9), the resource provider might offer $y_{l,k}$ less than $R_{l,k}$ for some prices λ, in order to drive up the price and increase profits (selling fewer units at a sufficiently higher price). Likewise, there might be only a single service provider. Indeed, in a corporate intranet, the information services division might be both the resource provider and the service provider to the users (corporate employees in general).

How useful is the price-taking assumption? It is not essential for a proof that an algorithm exists that will clear the markets and reach some equilibrium allocation of network resources. However, in general, that allocation will not be a solution of our original optimization problem, (**MAX1**). As a general matter we could show that equilibrium allocations based on behavior other than price-taking will lead to less efficient allocations, that is, allocations that do not maximize the sum of user utilities subject to the technology constraints (1)–(3). Therefore, for the purpose of this work we do not study markets in which agents exhibit different types of strategic behavior, but limit ourselves to the price-taking behavior that we can show can be harnessed to yield a solution to (**MAX1**).

Restricting attention to the price-taking case may not in practice be as restrictive as it seems. Consider the example of a corporate intranet with a single monopoly provider of resources and services. If the management instructs the resource and service provider to behave "as if" it is a price taker (and provides compensation incentives that make it in the provider's best interests to do so) then the desired outcome can be achieved. Essentially, this requires compensating the provider based on the value of the allocation to the company as a whole, rather than based on the provider's

own local "profits." If the network is to be managed with an agent-based control system, the agents should be programmed to act as price takers, whether or not other programmable strategies might seem more desirable from the local viewpoint of the agents.

In a more open, conventionally market-based system, such as a commercial market for virtual data circuits, it is also possible that at a given moment some participants might have some market power, which is to say that they are cognizant of an opportunity to improve their position by acting strategically with respect to price-setting, rather than as a price taker. In such a setting, it might not be possible to directly control behavior to make those participants behave "as if" they are price takers. However, if there are no artificial barriers to entry by other providers—for example, if it is possible for another competing firm to build an interconnected network of links with buffers and bandwidth—then it will tend to be the case that in a long-run equilibrium surviving agents will be those who behave as price takers (competition will drive others out of the market). Therefore, we believe there are many circumstances under which the conditions will exist, or can be imposed, that are necessary for our method to provide an equilibrium that is a solution to (**MAX1**).

(3) When we introduced the resource providers' problem in §4.1, we assumed that there is no cost in supplying network resources (bandwidth, buffers, etc.) to the market. This cost can be incorporated into our model if we subtract it from the objective function of the optimization problem. We believe that the new problem will have the same qualitative properties with the problem presented in this paper, thus it may lead to a similar type of result.

(4) In §2 we made the assumption that the sets $\mathbf{F}^{i,j}$ are strictly convex. This condition is necessary for the convergence of the market-based algorithm.[8] If this condition is not satisfied for some pair i, j, then, in order to be able to use the algorithm, we need to replace $\mathbf{F}^{i,j}$ with its maximum strictly convex subset. An important future problem therefore, is to specify an algorithm for finding that relaxed set, and also to be able to characterize how inefficient is the resulting allocation compared to the optimum solution, if we knew how to find it.

(5) Our approach to resource allocation is based on the solution of a static constrained optimization problem. Resources are not reserved in anticipation of future requests, and response to users' requests may be delayed until the next allocation epoch. If we wanted to consider a formulation where at allocation instances resources will be reserved in anticipation of future requests, we would have to take into account statistics of future arrivals and statistics of service times among other things. These statistics may not be readily available; even if they were, we believe they would lead to an intractable decentralized dynamic stochastic optimization problem. We leave the study of such a problem as a future challenge.

APPENDIX A

Proof of Proposition 3.1

We first show that Γ is upper hemicontinuous. The graph of Γ, $\mathrm{Gr}\,\Gamma \triangleq \{(x, r) \in \mathbf{X} \times \mathbf{R} : g(x, r) \leqslant 0\}$ is a closed set. This follows from the fact that g is continuous in $\mathbf{X} \times \mathbf{R}$ and therefore it is lower-semicontinuous in $\mathbf{X} \times \mathbf{R}$ (Border 1989, p. 15). Therefore Γ is a closed correspondence of $\mathbf{R}$ into $\mathbf{X}$. Since $\mathbf{X}$ is compact, it follows from Border (1989, Proposition 11.9(b)), that Γ is upper hemicontinuous in $\mathbf{R}$.

Next we show that Γ is lower hemicontinuous. Thus, we must prove that for every $r' \in \mathbf{R}$ and each open set G in $\mathbf{X}$ meeting $\Gamma r'$ there is a neighborhood $\Omega(r')$ such that

$$r \in \Omega(r') \quad \Rightarrow \quad \Gamma r \cap G \neq \varnothing. \tag{A.1}$$

We proceed to determine for every $r' \in \mathbf{R}$ and every open set G in $\mathbf{X}$ meeting $\Gamma r'$, a neighborhood $\Omega(r')$ such that (A.1) is satisfied. For that matter, we define the correspondence, $\overline{\Gamma} : \mathbf{R} \mapsto \mathbf{X}$, by

$$\overline{\Gamma} r \triangleq \{x \in \mathbf{X} : g(x, r) < 0\}, \tag{A.2}$$

and show that the lower inverse of G under $\overline{\Gamma}$, defined by

$$\overline{\Gamma}^{-}[G] \triangleq \{r \in \mathbf{R} : \overline{\Gamma} r \cap G \neq \varnothing\}, \tag{A.3}$$

is a neighborhood of r' that satisfies (A.1).

We first show that $\overline{\Gamma}^{-}[G]$ is an open set. By assumption, $\overline{\Gamma} r \neq \varnothing$. Furthermore, since g is upper semicontinuous (as it is continuous) on $\mathbf{X} \times \mathbf{R}$, the graph of $\overline{\Gamma}$, defined by

$$\mathrm{Gr}\,\overline{\Gamma} = \{(x, r) \in \mathbf{X} \times \mathbf{R} : g(x, r) < 0\},$$

is open. Consequently, $\overline{\Gamma}$ is an open correspondence, hence it is lower hemicontinuous (Border 1989, Proposition 11.9(c)). Since $\overline{\Gamma}$ is lower hemicontinuous and G is an open set, $\overline{\Gamma}^{-}[G]$ is open (Border 1989, Definition 11.3).

We proceed now to establish (A.1) with $\overline{\Gamma}^{-}[G]$ playing the role of $\Omega(r')$. Let $r' \in \mathbf{R}$ and consider an open set G in $\mathbf{X}$ such that

$$\Gamma r' \cap G \neq \varnothing. \tag{A.4}$$

Then from (A.4), the fact that $\mathrm{cl}\,(\overline{\Gamma} r') = \Gamma r'$,[9] and the identity, $\mathrm{cl}\,(\overline{\Gamma} r') \cap G \subseteq \mathrm{cl}\,(\overline{\Gamma} r' \cap G)$, we conclude that $\mathrm{cl}\,(\overline{\Gamma} r' \cap G) \neq \varnothing$. Consequently, $\overline{\Gamma} r' \cap G \neq \varnothing$. Therefore by (A.3), $r' \in \overline{\Gamma}^{-}[G]$ and $\overline{\Gamma}^{-}[G]$ is an open neighborhood of r'. Furthermore, by (A.3),

$$r \in \overline{\Gamma}^{-}[G] \quad \Rightarrow \quad \overline{\Gamma} r \cap G \neq \varnothing. \tag{A.5}$$

Since $\overline{\Gamma} r \subseteq \Gamma r$, (A.5) gives that for every $r \in \overline{\Gamma}^{-}[G]$, $\Gamma r \cap G \neq \varnothing$. Thus, (A.1) is satisfied, as (A.4) holds and $\overline{\Gamma}^{-}[G]$ is an open neighborhood of r'. The lower hemicontinuity of Γ is thus established.

Therefore, Γ is continuous in $\mathbf{R}$, since it is both upper hemicontinuous and lower hemicontinuous in $\mathbf{R}$. This completes the proof of Proposition 3.1. □

APPENDIX B

Description of Auctioneer's Algorithm

The algorithm is based on the construction of a specific sequence of $(LK+1)\times(LK+1)$ matrices, each of which is obtained from the previous one by replacing a single column. The replacement for column c equals the sum of columns $c\ominus 1$ and $c\oplus 1$ minus column c itself. The sequence begins with the matrix M_0, given in (14). As shown in the proof of Theorem 4.1, replacing a column of M_0, other than the first, results in a matrix that has the same properties as M_0; that is, its entries are between -1 and $D+1$ and the column sums are equal to D.

The rule determining the column that is replaced at each step is the following:

Initial Step: The first column of M_0 $(m=0)$ is labeled with the integer $\nu(0)=1$ and each column m $(m=1,2,\ldots,LK)$ is labeled with the integer $\nu(m)=m$. Thus, all columns of M_0 have distinct labels except for the first two, whose labels are identical. The algorithm begins by replacing the second column (according to the procedure described above).

Iteration Step: An integer $\nu(c)\in\{0,1,\ldots,LK\}$ is assigned to a new column c, that is brought into the matrix as a replacement, according to the following rule: If the column has a negative entry, then $\nu(c)$ is the smallest subscript for which this is true. If the column has a zero entry, then $\nu(c)$ is the smallest subscript for which this is true. If none of the above is true, then the label of column c is *determined by observing the aggregate excess demand* $z(\lambda)$ that results from the announcement of the price vector

$$\lambda_{l,k}=M((l-1)K+k,c)/M(0,c),\quad l\in\mathbf{L},k\in\mathbf{K},$$

to the market:[10] if $z_{l,k}(\lambda)\leqslant 0$ for all $l\in\mathbf{L},k\in\mathbf{K}$ then $\nu(c)=0$; otherwise, $\nu(c)=\min\{(l-1)K+k: z_{l,k}(\lambda)>0\}$.

The matrix resulting after an integer $\nu(c)$ is assigned to a new column c has one of the following two features:

(F1) 1. None of the columns is associated with the label 0.

2. All of the columns have distinct labels, except for a single pair whose labels are identical.

3. One member of the pair of columns with identical labels has just been brought into the matrix.

(F2) All columns of the matrix have distinct labels that span the set $\{0,1,\ldots,LK\}$ (in this case the column whose label is 0 has just been brought into the matrix).

If the matrix constructed by the algorithm has feature **(F1)**, the algorithm proceeds by eliminating from the matrix that column in the pair with identical labels that has *not just been brought into the matrix. If the matrix constructed* by the algorithm has feature **(F2)**, the algorithm terminates. As shown in §4.2, any matrix that is constructed by the algorithm and has feature **(F2)** contains at least one column with all positive entries. This implies that the auctioneer announces at least one set of prices before the algorithm terminates. The demand $x(\lambda)$ and the resource allocation $r(\lambda)$ corresponding to the auctioneer's latest prices are arbitrarily close to a solution of **(MAX 1)**.

APPENDIX C

Proof of Proposition 4.1

We first use the strict convexity of the sets $\mathbf{F}^{i,j}$ to show *that the solution $r(\lambda)$ to (20) is unique. Then we use the* Maximum Theorem, to show that $r(\lambda)$ is continuous. The solution $x(p(\lambda))$ to problems (21) is then shown to be continuous in $p(\lambda)$. This, along with the continuity of $p(\lambda)$ shows that $x(\lambda)$ is also a continuous function of λ. The continuity of both $x(\lambda)$ and $r(\lambda)$ validates the statement of the proposition.

We show that for each λ the solution $r(\lambda)$ to (20) is unique. Assume that, for given $i\in\mathbf{N}$, $j\in\mathbf{M}_i$ the solution $r^{i,j}(\lambda)$ to problem

$$\min_{r^{i,j}\in\mathbf{F}^{i,j}}\ \sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}}\lambda_{l,k}\ r^{i,j}_{l,k}, \tag{A.6}$$

is not unique. Therefore, there are at least two distinct vectors, r and $\dot r$ that both solve (A.6). By the strict convexity of the set $\mathbf{F}^{i,j}$ we have that, $\ddot r\triangleq\alpha r+(1-\alpha)\dot r$, $0<a<1$ belongs to the interior of $\mathbf{F}^{i,j}$. This implies that we can find a vector $\bar r\in\mathbf{F}^{i,j}$, whose elements are strictly less than the corresponding elements of $\ddot r$. Thus,

$$\sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}}\lambda_{l,k}\bar r^{i,j}_{l,k}<\sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}}\lambda_{l,k}\ddot r^{i,j}_{l,k}. \tag{A.7}$$

But,

$$\begin{aligned}\sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}}\lambda_{l,k}\ \ddot r^{i,j}_{l,k}&=\alpha\sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}}\lambda_{l,k}\ r^{i,j}_{l,k}+(1-\alpha)\sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}}\lambda_{l,k}\ \dot r^{i,j}_{l,k}\\&=\sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}}\lambda_{l,k}r^{i,j}_{l,k}(\lambda).\end{aligned} \tag{A.8}$$

Equations (A.7) and (A.8) lead to a contradiction, thus $r(\lambda)$ is unique.

We proceed by showing that $r(\lambda)$ is a continuous function of λ. Let $i\in\mathbf{N}$, $j\in\mathbf{M}_i$, and consider the correspondence $\Phi^{i,j}:\mathbb{R}^{LK}_+\mapsto\mathbb{R}^{LK}_+$ defined by $\Phi^{i,j}(\lambda)=\{r^{i,j}\in\mathbb{R}^{LK}_+: r^{i,j}\in\mathbf{F}^{i,j}\}=\mathbf{F}^{i,j}$. Then, $\Phi^{i,j}$ is constant and therefore continuous in $\mathbb{R}^{LK}_+$. It is also compact-valued since $\mathbf{F}^{i,j}$ is a compact set. The function $f^{i,j}:\mathbb{R}^{LK}_+\to\mathbb{R}$, defined by $f^{i,j}(r^{i,j})=-\sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}}\lambda_{l,k}\ r^{i,j}_{l,k}$ is continuous on $\mathbb{R}^{LK}_+$. Therefore, the conditions of the Maximum Theorem (see Border 1989, Theorem 12.1, p. 64) are satisfied and the solution $r^{i,j}(\lambda)$ to problem (A.6) is an upper semicontinuous correspondence at λ. This result, along with the fact that $r^{i,j}(\lambda)$ is unique, establishes the continuity of $r(\lambda)$ at every λ.

In the sequel we show that $x(\lambda)$ is also a continuous function of λ. Let $i\in\mathbf{N}$, and consider the following problem:

$$\max_{x^i\in\mathbf{X}^i}\left[u_i(x^i)-\sum_{j\in\mathbf{M}_i}p^i_jx^{i,j}\right], \tag{A.9}$$

where $p^{i,j}\triangleq\sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}}\lambda_{l,k}\ r^{i,j}_{l,k}(\lambda)$.

The function $u_i(\cdot)$ is continuous and strictly convex on $\mathbb{R}^M_+$. Therefore, according to Mas-Colell et al. (1995, Proposition 3.D.2 (iii) and Proposition 3AA.1) the solution $x^i(p^i)$ to (A.9) is continuous at all vectors $p^i \triangleq \{p^i_j\}_{j\in\mathbf{M}_i}$ that satisfy $p^i_j > 0$, for all $j \in \mathbf{M}_i$. This result, along with the continuity of p^i with respect to λ, implies that $x(\lambda)$ is a continuous function of λ. The continuity of both $x(\lambda)$ and $r(\lambda)$, establish the fact that the aggregate excess demand function given by (19) is a continuous function of λ. This completes the proof of Proposition 4.1. □

APPENDIX D

Proof of Proposition 4.2

We show that C is nonempty by constructing a vector $\bar{q}$ that belongs in C.

Fix r to a feasible allocation and consider problem (**MAX 2**) that is presented in §3. The compactness of $\mathbf{X}$, the linearity of constraint (6), and the continuity of $u_i(\cdot)$ for all $i \in \mathbf{N}$, guarantee, by the Weierstrass Theorem (Simon and Blume 1994, p. 823), the existence of a finite solution $x(r)$ to (**MAX 2**). Notice also that $\mathbf{X}$ is convex, $u_i(\cdot)$ is concave for all $i \in \mathbf{N}$ (by Assumption (A1)), and that the following constraint qualification holds true: $\hat{x} = 0$ belongs in X and is such that $\sum_{i\in\mathbf{N}}\sum_{j\in\mathbf{M}_i} \hat{x}^i_j r^{i,j}_{l,k} < R_{l,k}$, for all $l \in \mathbf{L}, k \in \mathbf{K}$. Therefore, according to Bazaraa et al. (1993, Corollary, p. 210), there exists a vector $\lambda(r) \in \mathbb{R}^{LK}_+$ that satisfies

$$\sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}} \lambda_{l,k}(r)\left(\sum_{i\in\mathbf{N}}\sum_{j\in\mathbf{M}_i} x^i_j(r) r^{i,j}_{l,k} - R_{l,k}\right) = 0, \tag{A.10}$$

and is such that

$$U(r) = \max_{x\in\mathbf{X}} \min_{\lambda\geqslant 0} \Lambda(x, r, \lambda), \tag{A.11}$$

where

$$\Lambda(x, r, \lambda) \triangleq \sum_{i\in\mathbf{N}} u_i(x^i) - \sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}} \lambda_{l,k}\left(\sum_{i\in\mathbf{N}}\sum_{j\in\mathbf{M}_i} x^i_j\, r^{i,j}_{l,k} - R_{l,k}\right) \tag{A.12}$$

is the Lagrangian, $x(r)$ is a solution to (**MAX 2**), and $U(r)$ is its maximum value function.

Based on the above, if $\bar{r}$ is a solution to problem (**MAX 3**) (also presented in §3), then $(x(\bar{r}), \bar{r})$ solves (**MAX 1**). The maximum value function of (**MAX 1**) is then given by

$$\max_{r\in\mathbf{F}} \max_{x\in\mathbf{X}} \min_{\lambda\geqslant 0} \Lambda(x, r, \lambda) = \Lambda(\bar{x}, \bar{r}, \bar{\lambda}), \tag{A.13}$$

where $\bar{x} \triangleq x(\bar{r})$, $\bar{\lambda} \triangleq \lambda(\bar{r})$, and $r \in \mathbf{F}$ is an abbreviation for (7).

Consider the vector $\bar{q}$, defined by

$$\bar{q}_0 \triangleq \left(1 + \sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}} \bar{\lambda}_{l,k}\right)^{-1}, \tag{A.14}$$

$$\bar{q}_{(l-1)K+k} \triangleq \bar{\lambda}_{l,k}\left(1 + \sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}} \bar{\lambda}_{l,k}\right)^{-1}, \quad l \in \mathbf{L}, k \in \mathbf{K}. \tag{A.15}$$

We establish that $\bar{q} \in C$. Note that, by construction, $\bar{q} \in S$. Also, since $(\bar{x}, \bar{r})$ is feasible to (**MAX 1**),

$$\sum_{i\in\mathbf{N}}\sum_{j\in\mathbf{M}_i} \bar{x}^i_j\, \bar{r}^{i,j}_{l,k} - R_{l,k} \leqslant 0 \quad \text{for all} \quad l \in \mathbf{L}, k \in \mathbf{K}. \tag{A.16}$$

Thus, according to (17) and (19), $\bar{q} \in C_0$. We proceed by showing that $\bar{q}$ belongs to the sets $C_{l,k}$, for all $l \in \mathbf{L}$ and $k \in \mathbf{K}$.

If $\bar{q}_{(l-1)K+k} = 0$ for some $l \in \mathbf{L}, k \in \mathbf{K}$ then $\bar{q} \in C_{l,k}$ (see (18)). In the sequel we consider the case where $\bar{q}_{(l-1)K+k} > 0$. In this case we proceed as follows: We first show that $\bar{r}$ and $\bar{x}$ solve (20) and (21) respectively, for $\lambda = \bar{\lambda}$. Then we show that

$$\sum_{i\in\mathbf{N}}\sum_{j\in\mathbf{M}_i} \bar{x}^i_j\, \bar{r}^{i,j}_{l,k} - R_{l,k} = 0. \tag{A.17}$$

These two results establish, according to (16)–(21), that $\bar{q} \in C_{l,k}$.

To show that $\bar{r}$, $\bar{x}$ solve (20) and (21), respectively, for $\lambda = \bar{\lambda}$, we denote by $r(\bar{\lambda})$ and $x(\bar{\lambda})$ the solutions to problems (20) and (21) and note that, since $\bar{x}, \bar{r}$ solves (**MAX 1**), $\bar{r}$ is a feasible solution to (20) and $\bar{x}$ is a feasible solution to (21). Therefore the following inequalities are true:

$$\sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}} \bar{\lambda}_{l,k}\, \bar{r}^{i,j}_{l,k} \geqslant \sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}} \bar{\lambda}_{l,k}\, r^{i,j}_{l,k}(\bar{\lambda}), \quad i \in \mathbf{N}, j \in \mathbf{M}_i, \tag{A.18}$$

$$u_i(x^i(\bar{\lambda})) - \sum_{j\in\mathbf{M}_i} x^i_j(\bar{\lambda}) \sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}} \bar{\lambda}_{l,k}\, r^{i,j}_{l,k}(\bar{\lambda})$$
$$\geqslant u_i(\bar{x}^i) - \sum_{j\in\mathbf{M}_i} \bar{x}^i_j \sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}} \bar{\lambda}_{l,k}\, r^{i,j}_{l,k}(\bar{\lambda}), \quad i \in \mathbf{N}. \tag{A.19}$$

From the above it follows that

$$\Lambda(x(\bar{\lambda}), r(\bar{\lambda}), \bar{\lambda}) \geqslant \Lambda(\bar{x}, r(\bar{\lambda}), \bar{\lambda}) \geqslant \Lambda(\bar{x}, \bar{r}, \bar{\lambda}), \tag{A.20}$$

where, the first inequality is obtained from (A.19) by summing over all i and adding the term $\sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}} \bar{\lambda}_{l,k}\, R_{l,k}$, and the second from (A.18) by multiplying by $-\bar{x}^i_j$, summing over all i, j, and adding the term $\sum_{i\in\mathbf{N}} u_i(\bar{x}^i) + \sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}} \bar{\lambda}_{l,k}\, R_{l,k}$. Assume that $\bar{r}, \bar{x}$ do not solve (20) and (21) for $\lambda = \bar{\lambda}$. This implies that at least one of the inequalities in (A.18), (A.19) is satisfied with strict inequality. Therefore (A.20) gives

$$\Lambda(x(\bar{\lambda}), r(\bar{\lambda}), \bar{\lambda}) > \Lambda(\bar{x}, \bar{r}, \bar{\lambda}). \tag{A.21}$$

From (A.13) we obtain that

$$\max_{r\in\mathbf{F}} \max_{x\in\mathbf{X}} \min_{\lambda=\bar{\lambda}} \Lambda(x, r, \lambda) = \Lambda(\bar{x}, \bar{r}, \bar{\lambda}), \tag{A.22}$$

and this leads to a contradiction, since (A.21) violates (A.22). Therefore, $\bar{r}, \bar{x}$ solve (20) and (21) respectively, for $\lambda = \bar{\lambda}$.

We next show that

$$\sum_{i\in\mathbf{N}}\sum_{j\in\mathbf{M}_i} \bar{x}^i_j\, \bar{r}^{i,j}_{l,k} - R_{l,k} = 0. \tag{A.23}$$

Equation (A.10) for $r = \bar{r}$ gives

$$\sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}} \bar{\lambda}_{l,k}\left(\sum_{i\in\mathbf{N}}\sum_{j\in\mathbf{M}_i} \bar{x}_j^i\ \bar{r}_{l,k}^{i,j} - R_{l,k}\right) = 0. \tag{A.24}$$

This combined with (A.16) and the fact that $\bar{\lambda}_{l,k} \geqslant 0$ for all $l \in \mathbf{L}, k \in \mathbf{K}$, implies that

$$\bar{\lambda}_{l,k}\left(\sum_{i\in\mathbf{N}}\sum_{j\in\mathbf{M}_i} \bar{x}_j^i\ \bar{r}_{l,k}^{i,j} - R_{l,k}\right) = 0 \quad \text{for all} \quad l \in \mathbf{L}, k \in \mathbf{K}. \tag{A.25}$$

By assumption, $\bar{q}_{(l-1)K+k} > 0$. Therefore $\bar{\lambda}_{l,k} > 0$, which combined with (A.25), leads to (A.23).

We have thus established the fact that $\bar{q} \in C_{l,k}$, for all $l \in \mathbf{L}, k \in \mathbf{K}$. This completes the proof of Proposition 4.2. □

APPENDIX E

Proof of Proposition 4.3

Let $q \in C$. Then $q_0 > 0$ because $q \in C_{l,k}$. Furthermore, since $q \in C_0$ and $q_0 > 0$, we have that

$$\sum_{i\in\mathbf{N}}\sum_{j\in\mathbf{M}_i} x_j^i(\lambda')\ r_{l,k}^{i,j}(\lambda') \leqslant R_{l,k} \quad \text{for all} \quad l \in \mathbf{L},\ k \in \mathbf{K}. \tag{A.26}$$

We also have from (20) and (21) that $x(\lambda') \in \mathbf{X}$ and $r^{i,j}(\lambda') \in \mathbf{F}^{i,j}$, for all i, j. Therefore $x(\lambda'), r(\lambda')$ is a feasible solution to problem (**MAX 1**).

Let x, r be any other feasible solution to (**MAX 1**). From (20) and (21) we get

$$\sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}} \lambda_{l,k}'\ r_{l,k}^{i,j} \geqslant \sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}} \lambda_{l,k}' r_{l,k}^{i,j}(\lambda'), \quad i \in \mathbf{N},\ j \in \mathbf{M}_i, \tag{A.27}$$

and

$$\begin{aligned} u_i(x^i(\lambda')) &- \sum_{j\in\mathbf{M}_i} x_j^i(\lambda') \sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}} \lambda_{l,k}'\ r_{l,k}^{i,j}(\lambda') \\ &\geqslant u_i(x^i) - \sum_{j\in\mathbf{M}_i} x_j^i \sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}} \lambda_{l,k}' r_{l,k}^{i,j}(\lambda'), \quad i \in \mathbf{N}, \end{aligned} \tag{A.28}$$

respectively. From the above it follows that

$$\Lambda(x(\lambda'), r(\lambda'), \lambda') \geqslant \Lambda(x, r(\lambda'), \lambda') \geqslant \Lambda(x, r, \lambda'), \tag{A.29}$$

where the first inequality is obtained from (A.28) by summing over all i and adding the term $\sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}} \lambda_{l,k}'\ R_{l,k}$, and the second from (A.27) by multiplying by $-x_j^i$, summing over all i, j, and adding the term $\sum_{i\in\mathbf{N}} u_i(x^i) + \sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}} \lambda_{l,k}'\ R_{l,k}$.

From (A.29), we have

$$\begin{aligned} \sum_{i\in\mathbf{N}} u_i(x^i(\lambda')) &- \sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}} \lambda_{l,k}'\left(\sum_{i\in\mathbf{N}}\sum_{j\in\mathbf{M}_i} x_j^i(\lambda') r_{l,k}^{i,j}(\lambda') - R_{l,k}\right) \\ &\geqslant \sum_{i\in\mathbf{N}} u_i(x^i) - \sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}} \lambda_{l,k}'\left(\sum_{i\in\mathbf{N}}\sum_{j\in\mathbf{M}_i} x_j^i\ r_{l,k}^{i,j} - R_{l,k}\right). \end{aligned} \tag{A.30}$$

Note that $q \in C_{l,k}$ implies that $q_{(l-1)K+k} = 0$ (i.e. $\lambda_{l,k}' = 0$), or that

$$\sum_{i\in\mathbf{N}}\sum_{j\in\mathbf{M}_i} x_j^i(\lambda')\ r_{l,k}^{i,j}(\lambda') \geqslant R_{l,k}. \tag{A.31}$$

If $q_{(l-1)K+k} > 0$, then from (A.26) and (A.31) it follows that

$$\sum_{i\in\mathbf{N}}\sum_{j\in\mathbf{M}_i} x_j^i(\lambda')\ r_{l,k}^{i,j}(\lambda') = R_{l,k}. \tag{A.32}$$

Therefore,

$$\sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}} \lambda_{l,k}'\left(\sum_{i\in\mathbf{N}}\sum_{j\in\mathbf{M}_i} x_j^i(\lambda')\ r_{l,k}^{i,j}(\lambda') - R_{l,k}\right) = 0, \tag{A.33}$$

because for all $l \in \mathbf{L}, k \in \mathbf{K}$ either $\lambda_{l,k}' = 0$ or $\sum_{i\in\mathbf{N}}\sum_{j\in\mathbf{M}_i} x_j^i(\lambda') r_{l,k}^{i,j}(\lambda') - R_{l,k} = 0$.

Since x, r is a feasible solution to (**MAX 1**) we have that

$$\sum_{i\in\mathbf{N}}\sum_{j\in\mathbf{M}_i} x_j^i\ r_{l,k}^{i,j} \leqslant R_{l,k} \quad \text{for all} \quad l \in \mathbf{L}, k \in \mathbf{K}. \tag{A.34}$$

Multiplying (A.34) by $\lambda_{l,k}' \geqslant 0$ and summing over all l, k, gives

$$\sum_{l\in\mathbf{L}}\sum_{k\in\mathbf{K}} \lambda_{l,k}'\left(\sum_{i\in\mathbf{N}}\sum_{j\in\mathbf{M}_i} x_j^i\ r_{l,k}^{i,j} - R_{l,k}\right) \leqslant 0. \tag{A.35}$$

Inequalities (A.30), (A.33), and (A.35) give

$$\sum_{i\in\mathbf{N}} u_i(x^i(\lambda')) \geqslant \sum_{i\in\mathbf{N}} u_i(x^i), \tag{A.36}$$

which shows that $x(\lambda'), r(\lambda')$ solves (**MAX 1**). This concludes the proof of Proposition 4.3. □

ENDNOTES

1. The numeraire quantities x_0^i, $i \in \mathbf{N}$, are omitted without loss of generality from the objective function, because they do not affect the marginal conditions for optimization.
2. We assume that $r_{l,k}^{i,j} = 0$ if link l doesn't belong to the path $V^{i,j}$ assigned to that connection.
3. That is, for every $x \in \mathbf{X}$ and every $\epsilon > 0$, there is $x^0 \in \mathbf{X}$ such that $\|x^0 - x\| \leqslant \epsilon$ and $u_i(x^0) > u_i(x)$, where $\|x^0 - x\|$ is the Euclidean distance between points x^0 and x.
4. A set is strictly convex if every proper convex combination of two points of the set belongs to the interior of the set.
5. In the concluding §5, we discuss the role of the price-taking assumption in our market-based implementation.
6. Scarf (1973, Theorems 6.2.1 and 6.2.9) derives conditions for a matrix $M \in \mathbf{M}$, in order for the vectors or sides of S that are associated with the columns of M to form a primitive set in S.
7. Otherwise T would be formed by sides and/or vectors of P that belong to the sides of S, thus, it would coincide with S which is not a primitive set (as it contains vectors of P in its interior).
8. We have seen in §3, that compactness is the necessary condition for the existence of an optimal allocation.
9. cl(A) denotes the closure of a set A.

10. $M(p, q)$ denotes the element in the p row and q column of matrix M.

ACKNOWLEDGMENTS

This research was supported in part by NSF Grant NCR-9204419 and ARO Grant DAAH04-96-1-0377.

REFERENCES

Bazaraa, M., H. Sherali, C. Shetty. 1993. *Nonlinear Programming Theory and Algorithms*. John Wiley, New York.

Border, K. 1989. *Fixed Point Theorems with Applications to Economics and Game Theory*. Cambridge University Press, New York.

Cocchi, R., D. Estrin, S. Shenker, L. Zhang. 1993. Pricing in computer networks: Motivation, formulation, and example. *IEEE/ACM Trans. on Networking* **1**(6) 614–627.

Courcoubetis, C., F. P. Kelly, R. R. Weber. 1997. Measurement-based usage charges in communications networks. Technical Report 19, Statistical Laboratory, University of Cambridge, Cambridge, U.K.

De Veciana, G., R. Baldick. 1998. Resource allocation in multiservice networks via pricing. *Computer Networks ISDN Systems* **30**(9–10) 951–962.

Gupta, A., D. Stahl, A. Whinston. 1997. A stochastic equilibrium model of Internet pricing. *J. Econom. Dynam. Control* **21**(4–5) 697–722.

Hurwicz, L. 1973. The design of mechanisms for resource allocation. *Amer. Econom. Rev.* **63**(2) 1–30.

Jiang, H., S. Jordan. 1995. The role of price in the connection establishment process. *Eur. Trans. Telecomm.* **6**(4) 421–429.

Jordan, S., H. Jiang. 1995. Connection establishment in high speed networks. *IEEE Selected Areas in Comm.* **13**(7) 1150–1161.

Kelly, F. 1994. On tariffs, policing and admission control for multiservice networks. *Oper. Res. Letters* **15**(1) 1–9.

——, A. Maulloo, D. Tan. 1998. Rate control for communication networks: Shadow prices, proportional fairness and stability. *J. Oper. Res. Soc.* **49** 237–252.

Kurose, J., R. Simha. 1989. A microeconomic approach to optimal resource allocation in distributed systems. *IEEE Trans. Comput.* **38**(5) 705–717.

Low, S., P. Varaiya. 1993. A new approach to service provisioning in ATM Networks. *IEEE/ACM Trans. Networking*, **1**(5) 547–553. See also Low and Varaiya 1994 for corrections.

——, ——. 1994. Corrections to: A new approach to service provisioning in ATM Networks. *IEEE/ACM Trans. Networking* **2**(3) 312.

MacKie-Mason, J., H. Varian. 1995a. Pricing congestible network resources. *J. Selected Areas in Comm.* **13**(7) 1141–1149.

——, ——. 1995b. Some FAQs about usage-based pricing. *Comput. Networks and ISDN Systems* **28**(1–2) 257–265.

Mas-Colell, A., M. D. Whinston, J. R. Green. 1995. *Microeconomic Theory*. Oxford University Press, New York.

Murphy, L., J. Murphy. 1994. Bandwidth allocation by pricing in ATM networks. *Broadband Communications, II, Proceedings of the IFIP TC6 Second International Conference on Broadband Communications* **C-24**. Elsevier Science Publishers B.V. (North-Holland), Paris, France.

——, ——, E. Posner. 1994. Distributed pricing for embedded ATM networks. *Proc. Internat. Teletraffic Congress ITC-14*. Elsevier Science Publishers B.V. (North-Holland), Antibes, France.

Parris, C., D. Ferrari. 1992. A resource based pricing policy for real-time channels in a packet-switching network. Technical report. International Computer Science Institute, Berkeley, CA.

——, S. Keshav, D. Ferrari. 1992. A framework for the study of pricing in integrated networks. Technical report TR-92-016. International Computer Science Institute, Berkeley, CA.

Scarf, H. 1973. *The Computation of Economic Equilibria*. Yale University Press, New Haven, CT, and London, U.K.

Simon, C., L. Blume. 1994. *Mathematics for Economists*. W. W. Norton, New York.

Thomas, P., D. Teneketzis. 1997. An approach to service provisioning with quality of service requirements in ATM networks. *J. High Speed Networks* **6**(4) 263–291.

Wang, Q., J. M. Peha, M. A. Sirbu. 1997. Optimal pricing for integrated-services networks. L. W. McKnight, J. P. Bailey, eds. *Internet Economics*, 3rd ed. MIT Press, Cambridge, MA, 353–376.

Warren B. Powell ("Implementing Real-Time Optimization Models: A Case Application from the Motor Carrier Industry") is Director of CASTLE Laboratory and a Professor in the Department of Operations Research and Financial Engineering at Princeton University. His research has focused on models and algorithms for stochastic, dynamic problems arising in freight transportation. After several projects to implement load matching systems in the truckload industry, this project grew out of a desire to study both the benefits of these systems and the challenges with implementing them successfully.

Kumar Rajaram ("Product Cycling with Uncertain Yields: Analysis and Application to the Process Industry") is an Assistant Professor of Operations and Technology Management at the Anderson School at UCLA. A significant portion of his research deals with developing techniques to improve the productivity of process operations. For his earlier work in this area, in 1998 he was a finalist for the Franz Edelman Award for Achievement in Operations Research and the Management Sciences, an award of INFORMS. His other research interests include developing models to better match supply with uncertain demand for products with short life cycles and highly unpredictable demand.

Richard E. Rosenthal ("Optimizing Military Airlift") is a Professor of OR at the Naval Postgraduate School. In 2001, he was a Visiting Senior Scientist in the Decision Technologies Department of Hewlett-Packard Labs. He is interested in the application of optimization models to industry, government, and the military.

José L. Ruiz ("A Statistical Test for Nested Radial DEA Models") has the degree of M.Sc. in mathematics from the University of Granada and a Ph.D. from the University Miguel Hernández of Elche. He is Associate Professor at the Department of Statistics and Applied Mathematics of the University Miguel Hernández of Elche and is a member of the Operational Research Center of the University. His research interest is primarily the analysis of efficiency and productivity.

Trilochan Sastry ("Facets and Reformulations for Solving Production Planning with Changeover Costs") is a Professor in the Production and Quantitative Methods Area at the Indian Institute of Management, Ahmedabad. His research interests include operations research, supply-chain management, and operations strategy. This paper is part of a dissertation written under the supervision of Professor Thomas L. Magnanti. It applies cutting plane techniques to NP-hard scheduling problems.

Hanif D. Sherali ("On Tightening the Relaxations of Miller-Tucker-Zemlin Formulations for Asymmetric Traveling Salesman Problems") is the W. Thomas Rice Endowed Chaired Professor of Engineering in the Industrial and Systems Engineering Department at Virginia Polytechnic Institute and State University. His area of research interest is in discrete and continuous optimization, with applications to location, transportation, and engineering design problems. He has coauthored several papers in operations research journals, has written four books in this area, and serves on the editorial boards of eight journals. He is a member of the U.S. National Academy of Engineering.

Inmaculada Sirvent ("A Statistical Test for Nested Radial DEA Models") has the degree of M.Sc. in mathematics from the University of Valencia and a Ph.D. from the University Miguel Hernández of Elche. She is Associate Professor at the Department of Statistics and Applied Mathematics of the University Miguel Hernández of Elche and is a member of the Operational Research Center of the university. Her research interest is primarily the analysis of efficiency and productivity.

Joel S. Sokol ("Spare-Capacity Assignment for Line Restoration Using a Single-Facility Type") is an Assistant Professor in the School of Industrial and Systems Engineering at the Georgia Institute of Technology. His current research includes the use of OR models to answer questions in telecommunications, logistics, biology, and sports.

Matthew Stafford ("Note: Online Searching/On the Optimality of the Geometric Sequences for the m Ray Search") is a Principal Member of Technical Staff, Cingular Wireless Inc. Dr Stafford's current research interests are in data networking, particularly data services in 3rd Generation wireless networks. The paper "Online Searching" is part of his Ph.D. dissertation at The University of Texas at Austin.

Demosthenis Teneketzis ("A Market-Based Approach to Optimal Resource Allocation in Integrated-Services Connection-Oriented Networks") received a B.S. degree in electrical engineering from the University of Patras, Greece, and M.S., E.E., and Ph.D. degrees in electrical engineering from MIT. He is a Professor of Electrical Engineering and Computer Science at the University of Michigan, Ann Arbor. His current research interests include stochastic control, decentralized systems, queueing and communication networks, stochastic scheduling and resource allocation problems, mathematical economics, and discrete-event systems.

Panagiotis Thomas ("A Market-Based Approach to Optimal Resource Allocation in Integrated-Services Connection-Oriented Networks") received a B.E. degree in electrical engineering from the Aristotle University of Thessaloniki, Greece, and M.S. and Ph.D. degrees in electrical engineering systems from the University Michigan, Ann Arbor. He works for Motorola Inc., Arlington Heights, Illinois, in the Advanced Network Technology department. His research interests include authorization, accounting, and quality of service provisioning in next generation wireless networks.

Yi Wang ("Spare-Capacity Assignment for Line Restoration Using a Single-Facility Type") is a consulting manager at i2 Technologies, Inc., a leading provider of supply chain planning solutions to enterprises. Her current vocational