Architecting Noncooperative Networks

Yannis A. Korilis, *Student Member, IEEE*, Aurel A. Lazar, *Fellow, IEEE*, and Ariel Orda, *Member, IEEE*

Abstract— **In noncooperative networks users make control decisions that optimize their individual performance measure. Focusing on routing, two methodologies for architecting noncooperative networks are devised, that improve the overall network performance. These methodologies are motivated by problem settings arising in the provisioning and the run time phases of the network. For either phase, Nash equilibria characterize the operating point of the network. The goal in the provisioning phase is to allocate link capacities that lead to systemwide efficient Nash equilibria. The solution of such design problems is, in general, counterintuitive, since adding link capacity might lead to degradation of user performance. For systems of parallel links, it is shown that such paradoxes cannot occur and that the optimal solution coincides with the solution in the single-user case. Extensions to general network topologies are derived. During the run time phase, a manager controls the routing of part of the network flow. The manager is aware of the noncooperative behavior of the users and makes its routing decisions based on this information while aiming at improving the overall system performance. We obtain necessary and sufficient conditions for enforcing an equilibrium that coincides with the global network optimum, and indicate that these conditions are met in many cases of interest.**

I. Introduction

Control decisions in large-scale networks are often made by each user independently, according to its own individual performance objectives.[1] Such networks are henceforth called *noncooperative*, and game theory [1], [2] provides the systematic framework to study and understand their behavior. Game theoretic models have been employed in the context of flow control [3]–[7], routing [8]–[10], virtual path bandwidth allocation [11], and pricing [12] in modern networking. These studies mainly investigate the structure of the network operating points, i.e., the *Nash equilibria* of the respective games. Such equilibria are inherently inefficient [13] and, in general, exhibit suboptimal network performance.

The goal of this paper is to demonstrate that, while users make noncooperative decisions, there is still room for improving network performance. Improvements can be achieved both during the provisioning phase, i.e., when the network parameters are sized, and during the run time phase, i.e., during the operation of the network. Focusing on routing, we give a uniform methodology for achieving such improvements. This methodology is based on architecting the network equilibria. The related analysis involves comparisons of equilibria of different games. Such comparisons are scarcely attempted in the game theoretic literature, mainly due to the complex structure—or lack thereof—of the underlying game. One exception is [14], which addresses the problem of designing the service discipline of a switch shared by users performing flow control.

In the *provisioning phase*, the designer allocates link capacities, i.e., architects the *capacity configuration* of the network, so that the resulting equilibrium is systemwide "efficient" or "optimal." We consider several efficiency criteria for the designer, such as the "price" (marginal cost) as seen by each user, the total cost of each user, or some combination of the above. The designer has to decide how much capacity should be allocated to each link, while satisfying lower bounds specified per link and an upper bound on the total capacity. The designer seeks an allocation of capacities that achieves the best performance, according to the chosen efficiency criterion. The immediate question that arises is whether the designer should attempt to employ all the available resources. Surprisingly, in general, the answer is no! To illustrate this counterintuitive behavior of noncooperative networks, we adapt the Braess paradox [15], [16] to our setting and show that addition of resources may result in degradation of user performance.

Example: Consider the network depicted in Fig. 1. There are I users, each with an average throughput demand r, sending flow from node 1 to node 4. Links $(1,2)$ and $(3,4)$ have each capacity c_1. Link $(1,3)$ represents a path of n tandem links, each with capacity c_2.[2] Similarly, links $(2,4)$ and $(2,3)$ are paths of n links, each with capacities c_2 and c_3, respectively. Each user routes its demand r over the available paths, so as to minimize its total cost defined as the sum of its delays over all links. The delay per unit of flow on each link is given by the M/M/1 delay formula. Prices (marginal costs) represent derivatives of the cost with respect to user flows. For this system there exists a unique and symmetrical Nash equilibrium [10], i.e., the equilibrium flows (and thus, the costs and prices) of the users are equal. Figs. 2 and 3 show, correspondingly, the user price and cost as functions of c_3, for $c_1 = 2.7$, $c_2 = 27$, $n = 54$, $I = 10$ and $r = 0.2$. The figures indicate that, for any $c_3 > 0$, both the price and the cost of each user are higher than for $c_3 = 0$, i.e., eliminating the path $(2,3)$ leads to an improvement of performance for

Manuscript received September 30, 1994; revised April 1, 1995. This work was supported by the Office of Naval Research under Contract N00014-90-J-1289. This paper was presented in part at IEEE INFOCOM'95, Boston, MA, April 4–6, 1995.

Y. A. Korilis and A. A. Lazar are with the Department of Electrical Engineering and the Center for Telecommunications Research, Columbia University, New York, NY 10027 USA.

A. Orda is with the Department of Electrical Engineering, Technion, Haifa 32000, Israel. This work was in part carried out while he was visiting with the Center for Telecommunications Research at Columbia University.

IEEE Log Number 9413102.

[1] The term "user" is purposely left ambiguous. It may refer to a network user itself or, in case that the user's traffic consists of multiple connections, to individual connections that are controlled independently.

[2] For $c_2 \gg Ir$, each of the paths $(1,3)$ and $(2,4)$ approximates a link with nonnegligible delay that has low sensitivity to flow changes; such constructions are required in order to reproduce the classical Braess paradox in a queueing setting [16].

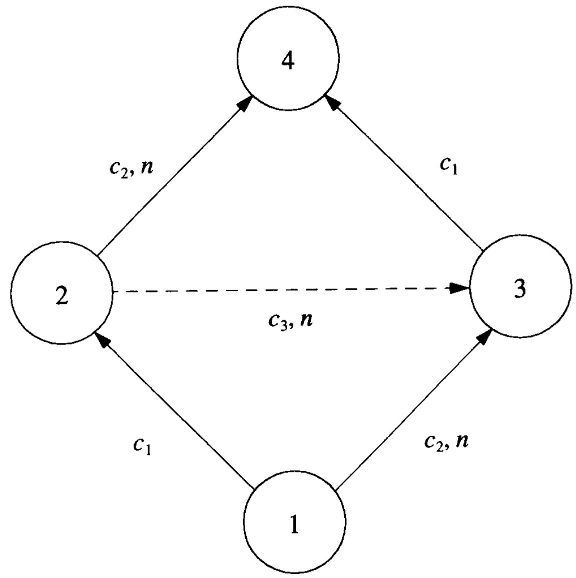

Fig. 1. Network paradox.

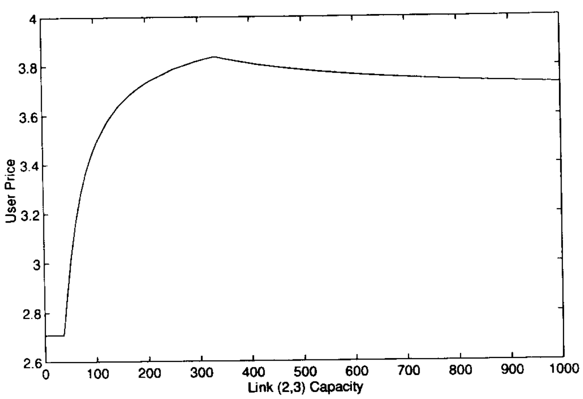

Fig. 2. User price as a function of the link capacity c_3.

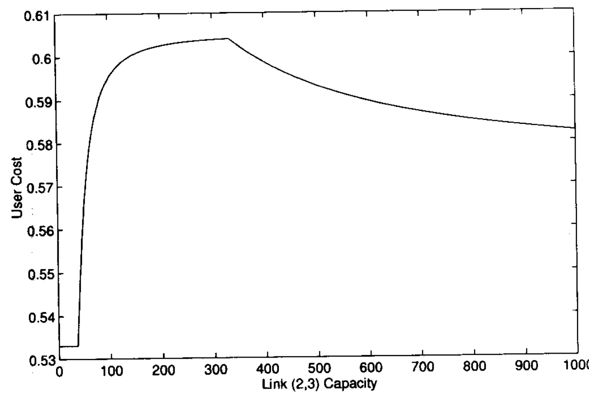

Fig. 3. User cost as a function of the link capacity c_3.

all users. More surprisingly, this paradoxical behavior persists even if $c_3 = \infty$, i.e., when nodes 2 and 3 are merged into a single node.

For a system of parallel links we show that the Braess paradox cannot occur, that is, addition of capacity improves the network performance. We then consider the problem of allocating such additional capacity to links in an optimal way. We show that the best design strategy is to allot the additional capacity exclusively to the link with the originally highest capacity. This solution coincides with the optimal capacity allocation in a network where routing is centrally controlled. We extend some of these results to general network topologies.

In the *run time phase*, we assume that, apart from the noncooperative users, there is also a manager, that attempts to optimize the system performance, by deciding upon the routing of an additional, network-controlled flow. The manager is aware of the noncooperative behavior of the users, and thus it can predict their reaction to any routing strategy that it chooses. This information enables the manager to implement a routing strategy that drives the users to the "best" Nash equilibrium in terms of system performance, architecting, this way, the *flow configuration* of the network. This is the typical scenario of a Stackelberg game [1], in which the manager acts as a leader and imposes its strategy on the users which behave as followers. Stackelberg strategies have been investigated in the context of flow control in [17]. In that reference, however, the leader is a selfish user concerned about its own, rather than the system's, performance.

For the parallel links model, we derive necessary and sufficient conditions that guarantee that the manager can enforce an equilibrium that coincides with the network optimum (the optimal solution of the routing problem when all the flow in the network is centrally controlled), and indicate that these conditions are met in many cases of practical interest. In other words, the manager is often able to obtain, through limited control, the same system performance as in the case of centralized control. Moreover, when these conditions are satisfied, we show that there exists a unique strategy of the manager that drives the system to the network optimum, and specify its structure explicitly.

We note that systems of parallel links, albeit simple, represent an appropriate model for seemingly unrelated networking problems. Consider, for example, a network in which resources are pre-allocated to various routing paths that do not interfere. Such scenaria are common in modern networking. In broadband networks bandwidth is separated among different virtual paths, resulting effectively in a system of parallel and noninterfering "links" between source/destination pairs. Another example is that of internetworking, in which each "link" models a different sub-network.

The outline of the paper is the following. In Section II, we present the noncooperative parallel links model. The design issues arising in the provisioning phase are investigated in Section III. Specifically, after formulating the problem in Section III-A, we outline the main results in Section III-B. In Section III-C, we explore the structure of the underlying Nash equilibria. In Section III-D, we establish that addition of capacity to a network of parallel links cannot degrade performance. With this result at hand, we investigate, in

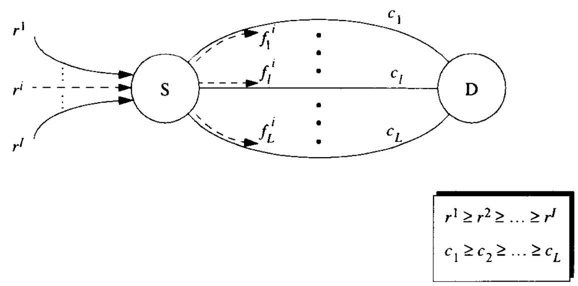

Fig. 4. The system of parallel links.

Section III-E, the optimal strategy for adding capacity to networks of parallel links. In Section III-F, we extend some of these results to general topologies. The management issues arising in the run time phase are considered in Section IV. Finally, Section V summarizes the main results and delineates their practical implications. Due to size constraints, most of the formal proofs are omitted; for these proofs the reader is referred to [18], [19].

II. The Model

We consider a set $\mathcal{I} = \{1, \cdots, I\}$ of users, that share a set $\mathcal{L} = \{1, \cdots, L\}$ of communication links interconnecting a common source to a common destination node (Fig. 4). Let c_l be the capacity of link l, and $C = \sum_{l \in \mathcal{L}} c_l$ be the total capacity of the system. Each user i has a throughput demand that is some process with average rate $r^i > 0$. We assume that $r^1 \geq r^2 \geq \cdots \geq r^I$. Let $R = \sum_{i \in \mathcal{I}} r^i$ denote the total demand of the users. We only consider capacity configurations $\mathbf{c} = (c_1, \cdots, c_L)$ that can accommodate the total user demand, i.e., configurations with $C > R$.

User i ships its flow by splitting its demand r^i over the set of parallel links, according to some individual performance objective. Let f_l^i denote the expected flow that user i sends on link l. The user flow configuration $\mathbf{f}^i = (f_1^i, \cdots, f_L^i)$ is called a routing *strategy* of user i and the set $F^i = \{\mathbf{f}^i \in I\!R^L : 0 \leq f_l^i \leq c_l,\ l \in \mathcal{L};\ \sum_{l \in \mathcal{L}} f_l^i = r^i\}$ of strategies that satisfy the user's demand is called the strategy space of user i. The system flow configuration $\mathbf{f} = (\mathbf{f}^1, \cdots, \mathbf{f}^I)$ is called a routing *strategy profile* and takes values in the product strategy space $F = \otimes_{i \in \mathcal{I}} F^i$.

The performance objective of user i is quantified by means of a cost function $J^i(\mathbf{f})$. The user aims to find a strategy $\mathbf{f}^i \in F^i$ that minimizes its cost. This optimization problem depends on the routing decisions of the other users, described by the strategy profile $\mathbf{f}^{-i} = (\mathbf{f}^1, \cdots, \mathbf{f}^{i-1}, \mathbf{f}^{i+1}, \cdots, \mathbf{f}^I)$, since J^i is a function of the system flow configuration $\mathbf{f}$. A *Nash equilibrium* of the routing game is a strategy profile from which no user finds it beneficial to unilaterally deviate. Hence, $\mathbf{f} \in F$ is a Nash equilibrium if

$$\mathbf{f}^i \in \arg\min_{\mathbf{g}^i \in F^i} J^i(\mathbf{g}^i, \mathbf{f}^{-i}), \qquad i \in \mathcal{I}. \tag{1}$$

The problem of existence and uniqueness of equilibria has been investigated in [10] for certain classes of cost functions. Here, we consider cost functions that are the sum of link cost functions

$$J^i(\mathbf{f}) = \sum\nolimits_{l \in \mathcal{L}} J_l^i(\mathbf{f}_l), \qquad J_l^i(\mathbf{f}_l) = f_l^i T_l(f_l), \qquad l \in \mathcal{L} \tag{2}$$

where $\mathbf{f}_l = (f_l^1, \cdots, f_l^I)$, and $T_l(f_l)$ is the average delay per unit of flow on link l that depends only on the total flow $f_l = \sum_{i \in \mathcal{I}} f_l^i$ on that link. In particular, we concentrate on the M/M/1 delay function

$$T_l(f_l) = \begin{cases} \frac{1}{c_l - f_l}, & f_l < c_l \\ \infty, & f_l \geq c_l. \end{cases} \tag{3}$$

Equations (2) and (3) imply that $J^i(\mathbf{f})/r^i$ is the average time-delay that the flow of user i experiences under strategy profile $\mathbf{f}$. Note that the stability constraint $f_l < c_l$ of link l is manifested through the definition of T_l. In particular, since the total user demand R does not exceed the total capacity C of the network, (1) and (3) guarantee that at any Nash equilibrium $f_l < c_l$ for all $l \in \mathcal{L}$.

Given a strategy profile $\mathbf{f}^{-i}$ of the other users, the cost of user i, as defined by (2) and (3), is a convex function of its strategy $\mathbf{f}^i$. Hence, the minimization problem in (1) has a unique solution. The Kuhn-Tucker optimality conditions [20], then, imply that $\mathbf{f}^i \in F^i$ is the optimal response of user i to $\mathbf{f}^{-i}$ if and only if there exists a (Lagrange multiplier) λ^i, such that

$$\lambda^i = \frac{\partial J^i}{\partial f_l^i}(\mathbf{f}), \qquad \text{if } f_l^i > 0, \qquad l \in \mathcal{L} \tag{4}$$

$$\lambda^i \leq \frac{\partial J^i}{\partial f_l^i}(\mathbf{f}), \qquad \text{if } f_l^i = 0, \qquad l \in \mathcal{L}. \tag{5}$$

Thus, a strategy profile $\mathbf{f} \in F$ is a Nash equilibrium, if and only if there exist λ^i, $i \in \mathcal{I}$, such that the optimality conditions (4)–(5) are satisfied for all $i \in \mathcal{I}$. The above conditions imply that the Lagrange multiplier λ^i is, in fact, the marginal cost of user i at the optimality point. In accordance with the economics terminology, λ^i will be referred to as the *price* of user i [21].

For the cost function $J^i(\mathbf{f})$ given by (2) and (3), we have

$$\frac{\partial J^i}{\partial f_l^i}(\mathbf{f}) = f_l^i T_l'(f_l) + T_l(f_l) = \frac{c_l - f_l^{-i}}{(c_l - f_l)^2} \tag{6}$$

where T_l' is the derivative of T_l with respect to f_l, and $f_l^{-i} = \sum_{j \neq i} f_l^j$ is the total flow that all users except the ith send on link l. Note that $T_l' = T_l^2$.

In [10] it has been shown that the routing game described above has a *unique* Nash equilibrium.

At times we will concentrate on special types of users, defined in the sequel.

Definition 1: Users are called *identical* if their demands are equal, i.e., $r^i = r^j$ for all $i, j \in \mathcal{I}$.

The Nash equilibrium of identical users is symmetrical, i.e., $f_l^i = f_l^j = f_l/I$ for all $i, j \in \mathcal{I}$ [10].

Definition 2: A user is is said to be *simple* if all of its flows are routed through links (or paths) of minimal delay.

Users often route their flows according to the "simple" scheme due to practical considerations. Many typical routing algorithms send flows through shortest paths, without accounting for derivatives (T_l') and thus bifurcating flows. The Nash equilibrium of simple users in a system of parallel links is unique with respect to the *total* link flows [10], and the corresponding necessary and sufficient conditions require the existence of some λ, such that

$$\lambda = T_l, \text{ if } f_l > 0, \qquad l \in \mathcal{L} \tag{7}$$

$$\lambda \leq T_l, \text{ if } f_l = 0, \qquad l \in \mathcal{L}. \tag{8}$$

We shall refer to the value of λ as the price of the simple users. From (7)–(8), it is easy to see that users that route according to the optimality conditions (4)–(5) become simple as their population grows to infinity and their individual demands become infinitesimally small, while their total demand remains R. This is the typical scenario in a transportation network.

Definition 3: Users are said to be *consistent* (for a given capacity configuration) if, at the Nash equilibrium, they all use the same set of links.

Due to the structure of their Nash equilibrium, identical users are consistent. It is easy to verify that simple users are also consistent. Finally, consistent users are typical of systems with heavy traffic, i.e., when R approaches C, in which case each user sends flow on all links in the network.

III. Architecting the Capacity Configuration in the Provisioning Phase

A. Problem Formulation

Consider a network of parallel links with initial capacity configuration $\mathbf{c}^0$ and total capacity $C^0 > R$. Assume that $c_1^0 \geq \cdots \geq c_L^0 > 0$. Suppose that there exists some additional capacity allowance of at most Δ, which the network designer can distribute among the network links. The aim of the designer is to implement a capacity configuration $\mathbf{c}$, with $c_l \geq c_l^0$ for all links $l \in \mathcal{L}$, that results in a network with a total capacity of at most $C^0 + \Delta$, that is "efficient" at the corresponding Nash equilibrium. Without loss of generality, we concentrate on capacity configurations $\mathbf{c}$ that preserve the initial link order, i.e., configurations with $c_1 \geq \cdots \geq c_L$.[3] The set of all capacity configurations that can be implemented by the designer is $\mathcal{C}_\Delta = \{\mathbf{c} \in I\!R_+^L : c_1 \geq \cdots \geq c_L;\ c_l \geq c_l^0,\ l \in \mathcal{L};\ \sum_{l\in\mathcal{L}}(c_l - c_l^0) \leq \Delta\}$. Each capacity configuration in $\mathcal{C}_\Delta$ induces a routing game that has a unique Nash equilibrium. Therefore, we can define a function $\mathcal{N} : \mathcal{C}_\Delta \to F$, that assigns to each $\mathbf{c} \in \mathcal{C}_\Delta$ the Nash equilibrium $\mathcal{N}(\mathbf{c})$ of its respective routing game. $\mathcal{N}$ will be referred to as the *Nash mapping*. The set $\mathcal{C}_\Delta$ will be called the *space of routing games*.

The designer may have different measures to characterize the efficiency of a capacity configuration. We shall concentrate on measures that are expressed by means of either the user prices or costs. Although the user's cost is a direct measure of its level of satisfaction, prices may be a more important measure from the system's point of view, since they account for the level of congestion as seen by users and are the direct indication of how each user could accommodate fluctuations in the system's state. The designer can consider various ways of combining either the prices or the costs of the users. We shall concentrate on *user* optimization, i.e., trying to reduce the price or cost of each and every user, and *overall* optimization, i.e., trying to reduce the sum of all prices or costs. The various performance measures of the designer are formally stated in the following definitions:

Definition 4: Consider two capacity configurations $\mathbf{c}$ and $\hat{\mathbf{c}}$ and let λ^i and $\hat{\lambda}^i$ (J^i and $\hat{J}^i$) be the price (cost) of user i at the respective equilibrium. Then:

1) Configuration $\hat{\mathbf{c}}$ is said to be *user price (cost) efficient* relative to configuration $\mathbf{c}$, if $\hat{\lambda}^i \leq \lambda^i$ ($\hat{J}^i \leq J^i$), for all $i \in \mathcal{I}$.
2) Configuration $\hat{\mathbf{c}}$ is said to be *overall price (cost) efficient* relative to configuration $\mathbf{c}$, if $\sum_{i\in\mathcal{I}} \hat{\lambda}^i \leq \sum_{i\in\mathcal{I}} \lambda^i$ ($\sum_{i\in\mathcal{I}} \hat{J}^i \leq \sum_{i\in\mathcal{I}} J^i$).

Definition 5: Given a set of capacity configurations $\mathcal{C}$, a capacity configuration $\mathbf{c}^* \in \mathcal{C}$ is called:

1) *user price (cost) optimal* in $\mathcal{C}$, if it is user price (cost) efficient relative to any $\mathbf{c} \in \mathcal{C}$,
2) *overall price (cost) optimal* in $\mathcal{C}$, if it is overall price (cost) efficient relative to any $\mathbf{c} \in \mathcal{C}$.

Obviously, user efficiency (optimality) implies overall efficiency (optimality). Price and cost efficiency (optimality), however, do not imply each other in either direction. Note also that, in general, existence of user optima cannot be guaranteed even if overall optima do exist.

The optimal capacity allocation problem corresponding to the various performance measures is described as follows:

> Given a system of parallel links $\mathcal{L}$ with users $\mathcal{I}$, an (initial) capacity configuration $\mathbf{c}^0$ and an additional capacity allowance Δ, find a capacity configuration $\mathbf{c}^*$ that is user/overall price/cost optimal in $\mathcal{C}_\Delta$.

Although the problem is formulated as allocating additional capacity to an existing network, this formulation is equivalent to the typical capacity allocation problem, where the capacity of each link has to be higher than a lower bound, e.g., due to reliability considerations.

Solving the optimal capacity allocation problem in a network shared by noncooperative users amounts to comparing the Nash equilibria of the routing games induced by different capacity configurations in $\mathcal{C}_\Delta$. Comparing the outcomes of different games is, in general, a highly complex task and is feasible only if an explicit characterization of the respective equilibria is available. The structure of the unique Nash equilibrium of the routing game is investigated in Section III-C. Before we proceed, let us first summarize the main results of this section.

B. Outline of Results

Following is an informal summary of the main results on the design problem.

1) Addition of capacity to any link results in a configuration that is user price efficient.
2) Addition of capacity to the link with the initially highest capacity results in an overall cost efficient configuration.

[3] The properties of the Nash equilibrium in a system of parallel links with capacity configuration $\mathbf{c}$ depends on the actual link capacities and not on the link "labels," that are determined by the initial configuration $\mathbf{c}^0$. Hence, renaming the links, so that $c_1 \geq \cdots \geq c_L$, does not affect the characteristics of the resulting equilibrium.

3) For consistent users (thus, in particular, for identical or simple users), addition of capacity to any link results with a user cost efficient configuration.
4) The capacity configuration that results from allocating the entire additional capacity allowance to the link with the initially highest capacity is user price optimal in $\mathcal{C}_\Delta$.
5) The user price optimal capacity configuration is also user cost optimal, if the lower bounds on the link capacities are the same for all links.
6) User cost optimality of the above configuration is also established when users are consistent (thus, in particular, when they are identical or simple), and also in the case of two users.
7) Considering general topologies (i.e., not necessarily systems of parallel links), we obtain methods for adding capacity to links so the Braess paradox does not occur.

C. Structure of the Nash Equilibrium

In this subsection we study the structure of the Nash equilibrium of the routing game in a network of parallel links with capacity configuration $\mathbf{c}$. A set of intuitive monotonicity properties of the Nash equilibrium have been established in [10] and are summarized in the following:

Lemma 1: Let $\mathbf{f}$ be the unique Nash equilibrium of the routing game in a network of parallel links with capacity configuration $\mathbf{c}$. Then:

1) The expected flow of any user $i \in \mathcal{I}$ decreases in the link number, i.e., $f_1^i \geq f_2^i \geq \cdots \geq f_L^i$.
2) For any link $l \in \mathcal{L}$, the flows decrease in the user number, i.e., $f_l^1 \geq f_l^2 \geq \cdots \geq f_l^I$.
3) The residual capacity is decreasing in the link number, i.e., $c_1 - f_1 \geq c_2 - f_2 \geq \cdots \geq c_L - f_L$.
4) For every user $i \in \mathcal{I}$, the residual capacity $c_l^i = c_l - f_l^{-i}$ seen by the user on link l is decreasing in the link number, i.e., $c_1^i \geq c_2^i \geq \cdots \geq c_L^i$.

Let $\mathcal{L}^i$ denote the set of links that receive some flow from user i, and $\mathcal{I}_l$ denote the set of users that send flow over link l. The first statement in Lemma 1 implies that for every user i, there exists some link L^i, such that $f_l^i > 0$ for all $l \leq L^i$, and $f_l^i = 0$ for $l > L^i$, that is, $\mathcal{L}^i = \{1, 2, \cdots, L^i\}$. Similarly, the second statement in the Lemma implies that for every link l, there exists some user I_l, such that $f_l^i > 0$ for all $i \leq I_l$, and $f_l^i = 0$ for $i > I_l$, that is, $\mathcal{I}_l = \{1, 2, \cdots, I_l\}$.

Consider now the best reply $\mathbf{f}^i$ of user i to a fixed strategy profile $\mathbf{f}^{-i}$ of the other users. This is the unique solution to the optimal routing problem in a network of parallel links with capacities c_l^i, $l \in \mathcal{L}$. In the sequel, we give an explicit characterization of the structure of the user's equilibrium strategy $\mathbf{f}^i$, as a function of $\mathbf{c}^i = (c_1^i, \cdots, c_L^i)$, which depends on the capacity configuration $\mathbf{c}$ and the strategy profile $\mathbf{f}^{-i}$ of the other users. To this end, let us define

$$G_l^i = \sum_{m=1}^{l-1} c_m^i - \sqrt{c_l^i} \sum_{m=1}^{l-1} \sqrt{c_m^i}, l = 2, \cdots, L \quad (9)$$
$$G_1^i = 0, \quad G_{L+1}^i = \sum_{m=1}^{L} c_m^i = C - R^{-i}$$

for all $i \in \mathcal{I}$, where $R^{-i} = \sum_{j \neq i} r^j$ is the total demand of all users except the ith. Then, $c_l^i \geq c_{l+1}^i$ (see Lemma 1) implies that $G_l^i \leq G_{l+1}^i$, for all $l \in \mathcal{L}$. We are now ready to state the following:

Proposition 1: The Nash equilibrium $\mathbf{f}$ of the routing game in a system of parallel links with capacity configuration $\mathbf{c}$ is described by

$$f_l^i = \begin{cases} c_l^i - (\sum_{m=1}^{L^i} c_m^i - r^i) \frac{\sqrt{c_l^i}}{\sum_{m=1}^{L^i} \sqrt{c_m^i}}, & 1 \leq l \leq L^i \\ 0, & L^i < l \leq L \end{cases} \quad (10)$$

for every user $i \in \mathcal{I}$, where the threshold L^i is determined by $G_{L^i}^i < r^i \leq G_{L^i+1}^i$. The equilibrium price and the equilibrium cost for user i are, respectively,

$$\lambda^i = \left\{ \frac{\sum_{l=1}^{L^i} \sqrt{c_l^i}}{\sum_{l=1}^{L^i} c_l^i - r^i} \right\}^2, \qquad J^i = \frac{\left\{\sum_{l=1}^{L^i} \sqrt{c_l^i}\right\}^2}{\sum_{l=1}^{L^i} c_l^i - r^i} - L^i. \quad (11)$$

Remark: The proposition implies that the information user i needs to determine its best reply $\mathbf{f}^i$ to any strategy profile $\mathbf{f}^{-i}$ of the other users is the residual capacity c_l^i seen by the user on every link $l \in \mathcal{L}$ (see (10) and (9)), and not a detailed description of $\mathbf{f}^{-i}$. In practice, information about the residual capacities c_l^i, $l \in \mathcal{L}$, can be acquired by means of an appropriate estimation technique.

From Proposition 1, and especially the expressions for the equilibrium prices and costs, it is clear that the set of links over which each user sends its flow has a prominent role in the properties of the Nash equilibrium. To investigate the capacity allocation problem, we need to compare the equilibria of games that are induced by different capacity configurations in $\mathcal{C}_\Delta$. If the resulting equilibria are such that the sets of links over which each user sends its flow do not coincide at both equilibria, such comparisons are extremely complex, if possible at all. In [18], we exploit the structure of the Nash equilibrium, to show that the Nash mapping $\mathcal{N}$ is *continuous*. In the same reference, we show that this fundamental property allows us to investigate the general capacity allocation problem based solely on comparisons between capacity configurations, that are such that each user sends its flow over the same links under both configurations.

D. Efficiency of Capacity Addition

In this subsection we investigate the addition of capacity to systems of parallel links, and show that, under various conditions, the Braess paradox does not occur in this setting.

A capacity configuration $\hat{\mathbf{c}}$ is called an *augmentation* of configuration $\mathbf{c}$, if $\hat{c}_l \geq c_l$ for all l and $\sum_l \hat{c}_l > \sum_l c_l$. Throughout this subsection we shall compare the Nash equilibrium of a capacity configuration $\mathbf{c}$ to that of some augmentation $\hat{\mathbf{c}}$. "Hat" values will refer to configuration $\hat{\mathbf{c}}$, while "nonhat" values to $\mathbf{c}$.

The first lemma shows that addition of capacity is always efficient as with respect to prices.

Lemma 2: If a capacity configuration $\hat{\mathbf{c}}$ is an augmentation of configuration $\mathbf{c}$, then $\hat{\mathbf{c}}$ is user price efficient relative to $\mathbf{c}$, i.e., $\hat{\lambda}^i \leq \lambda^i$, for all $i \in \mathcal{I}$. Moreover, the equilibrium delay of each link l is lower (not higher) under configuration $\hat{\mathbf{c}}$, i.e., $\hat{T}_l \leq T_l$, for all $l \in \mathcal{L}$.

The following lemma shows that if capacity is added solely to the link with the initially highest capacity (i.e., to link 1), the resulting configuration is overall cost efficient.

Lemma 3: Let $\mathbf{c}$ and $\hat{\mathbf{c}}$ be two capacity configurations such that $\hat{c}_l = c_l$ for all $l > 1$ and $\hat{c}_1 > c_1$. Then $\hat{\mathbf{c}}$ is overall cost efficient relative to $\mathbf{c}$.

The following two lemmata establish user cost efficiency of capacity addition in some special cases of interest.

Lemma 4: Let $\mathbf{c}$ and $\hat{\mathbf{c}}$ be two capacity configurations such that $\hat{\mathbf{c}}$ is an augmentation of $\mathbf{c}$. Assume that users are consistent under both $\hat{\mathbf{c}}$ and $\mathbf{c}$. Then $\hat{\mathbf{c}}$ is user cost efficient relative to $\mathbf{c}$, that is, $\hat{J}^i \leq J^i$, for all $i \in \mathcal{I}$.

The above result applies, in particular, both to identical users and to simple users, since they belong to the class of consistent users, under all capacity configurations.

Lemma 5: Let $\mathbf{c}$ and $\hat{\mathbf{c}}$ be two capacity configurations such that $\hat{c}_l = c_l$ for all $l > 1$ and $\hat{c}_1 > c_1$. Then, for $I = 2$, $\hat{\mathbf{c}}$ is user cost efficient relative to $\mathbf{c}$.

E. Optimal Capacity Allocation

We now proceed to investigate the optimal capacity allocation problem, according to the various performance measures defined in Section III-A. The main results of this section, namely Theorems 1, 2, and 3, assert that the capacity configuration $\mathbf{c}^* = \mathbf{c}^0 + \Delta\mathbf{e}_1$,[4] that results from allocating the entire additional capacity to the link with the initially highest capacity is (i) user price optimal in $\mathcal{C}_\Delta$ and (ii) user cost optimal in $\mathcal{C}_\Delta$ if the lower bounds on the link capacities are equal for all links. Furthermore, $\mathbf{c}^*$ will be shown to be user cost optimal for a number of special cases of interest. While this is a simple and intuitive result, its proof requires systematic analysis that establishes some "order" in the complex structure of the underlying game. Although most of the formal proofs are excluded from the main text, the lemmata presented in this section delineate the methodology through which that task has been achieved.

We start by considering two capacity configurations $\mathbf{c}$ and $\hat{\mathbf{c}}$ in $\mathcal{C}_\Delta$, such that $\hat{\mathbf{c}} = \mathbf{c} + \Delta_q(\mathbf{e}_1 - \mathbf{e}_q)$ is derived from $\mathbf{c}$ by a transfer of capacity Δ_q from some link $q > 1$, with $c_1 > c_q > c_q^0$, to link 1, and show that $\hat{\mathbf{c}}$ is user price efficient with respect to $\mathbf{c}$. Hence, if an additional capacity of exactly Δ is to be allocated to the system, then $\mathbf{c}^*$ is the optimal capacity configuration. By virtue of the price efficiency of capacity addition (Lemma 2), $\mathbf{c}^*$ is also user price optimal in the entire space of games $\mathcal{C}_\Delta$, that allows for addition of capacity not necessarily equal, but also less than Δ. As before, "hat" values will refer to configuration $\hat{\mathbf{c}}$, while "nonhat" values to the initial configuration $\mathbf{c}$.

The comparison of capacity configurations $\hat{\mathbf{c}}$ and $\mathbf{c}$ is carried out in a series of lemmata. The first lemma shows that the transfer of capacity from link q to link 1 decreases the equilibrium delay on link 1, while it increases the delay on link q.

Lemma 6: Consider two capacity configurations $\mathbf{c}, \hat{\mathbf{c}} \in \mathcal{C}_\Delta$ with $\hat{\mathbf{c}} = \mathbf{c} + \Delta_q(\mathbf{e}_1 - \mathbf{e}_q)$. Then, $\hat{T}_1 \leq T_1$ and $\hat{T}_q > T_q$.

[4] $\mathbf{e}_l$ is the vector in $\mathbb{R}^L$ with the lth component equal to 1 and all other components equal to 0.

In the sequel, we present two lemmata that will play a key role in the proof of price efficiency of $\hat{\mathbf{c}}$ relative to $\mathbf{c}$. Both refer to the case where the transfer of capacity Δ_q from link q to link 1 is such that each user sends its flow over the same set of links under $\mathbf{c}$ and $\hat{\mathbf{c}}$, i.e., $\hat{\mathcal{L}}^i = \mathcal{L}^i$ for all $i \in \mathcal{I}$, and $\hat{\mathcal{I}}_l = \mathcal{I}_l$ for all $l \in \mathcal{L}$. The first lemma asserts that the capacity transfer affects the prices of all users that send flow to link q in the same way, that is, either $\hat{\lambda}^i > \lambda^i$ for all $i \in \mathcal{I}_q$, or $\hat{\lambda}^i \leq \lambda^i$ for all $i \in \mathcal{I}_q$. Similarly, either all links with capacity lower than link q increase their equilibrium delays with $\hat{\mathbf{c}}$, or else all of them decrease their equilibrium delays.

Lemma 7: Consider two capacity configurations $\mathbf{c}, \hat{\mathbf{c}} \in \mathcal{C}_\Delta$ with $\hat{\mathbf{c}} = \mathbf{c} + \Delta_q(\mathbf{e}_1 - \mathbf{e}_q)$, where Δ_q is such that $\hat{\mathcal{L}}^i = \mathcal{L}^i$ for all $i \in \mathcal{I}$. Suppose that $\mathcal{I}_q \neq \emptyset$. Then, either:

1) $\hat{\lambda}^i > \lambda^i$ for all $i \in \mathcal{I}_q$, and $\hat{T}_l > T_l$ for all links l with $q < l \leq L^1$, or
2) $\hat{\lambda}^i \leq \lambda^i$ for all $i \in \mathcal{I}_q$, and $\hat{T}_l \leq T_l$, for all links l with $q < l \leq L^1$.

The following lemma shows that if the delay of some link l in $\{2, \cdots, q-1\}$ is higher under configuration $\hat{\mathbf{c}}$, then the same is true for all links in $\{l+1, \cdots, q-1\}$.

Lemma 8: Consider two capacity configurations $\mathbf{c}, \hat{\mathbf{c}} \in \mathcal{C}_\Delta$ with $\hat{\mathbf{c}} = \mathbf{c} + \Delta_q(\mathbf{e}_1 - \mathbf{e}_q)$, where $\Delta_q > 0$ is such that $\hat{\mathcal{L}}^i = \mathcal{L}^i$, for all $i \in \mathcal{I}$. For any link $l < q-1$, if $\hat{T}_l > T_l$ then $\hat{T}_n > T_n$ for all links $n \in \{l+1, \cdots, q-1\}$.

We are now ready to prove that $\hat{\mathbf{c}}$ is user price efficient compared to $\mathbf{c}$. The proof is given in the following theorem, which asserts also that the equilibrium delays on all links except link q are lower under configuration $\hat{\mathbf{c}}$.

Theorem 1: Consider two capacity configurations $\mathbf{c}, \hat{\mathbf{c}} \in \mathcal{C}_\Delta$ with $\hat{\mathbf{c}} = \mathbf{c} + \Delta_q(\mathbf{e}_1 - \mathbf{e}_q)$, $0 < \Delta_q \leq c_q - c_q^0$. Then:

1) Configuration $\hat{\mathbf{c}}$ is user price efficient compared to $\mathbf{c}$, i.e., $\hat{\lambda}^i \leq \lambda^i$ for all users $i \in \mathcal{I}$.
2) $\hat{T}_l \leq T_l$ for all $l \in \mathcal{L} \backslash \{q\}$ and $\hat{T}_q > T_q$.

Proof: Note that $\hat{T}_1 \leq T_1$ and $\hat{T}_q > T_q$ have been established in Lemma 6, thus we only have to prove the remaining statements in the theorem. We will establish these claims under the assumption that each user sends its flow over the same set of links under $\mathbf{c}$ and $\hat{\mathbf{c}}$, i.e., that $\hat{\mathcal{L}}^i = \mathcal{L}^i$ for all $i \in \mathcal{I}$. As explained in [18], the results readily generalize to the case where $\hat{\mathcal{L}}^i \neq \mathcal{L}^i$ for some user i.

If no user sends flow to link q, i.e., if $L^1 < q$, transferring capacity from link q to link 1 is, in fact, equivalent to adding capacity to the system of parallel links $\mathcal{L}' = \{1, \cdots, L^1\}$, and the result is immediate from Lemma 2. Thus, we have to consider only the case $\mathcal{I}_q \neq \emptyset$, i.e., $L^1 \geq q$. Without loss of generality, we will assume that user 1 sends flow on all links in the network, i.e., that $L^1 = L$.

Let us first show that $\hat{\lambda}^1 \leq \lambda^1$. Assume by contradiction that $\hat{\lambda}^1 > \lambda^1$. Then, by Lemma 7, $\hat{T}_l > T_l$, for all links $l \in \{q+1, \cdots, L\}$. An immediate consequence of Lemma 8 is that there exists some link l_0, $1 \leq l_0 < q$, such that $\hat{T}_l \leq T_l$ for all $l \in \{1, \cdots, l_0\}$, and $\hat{T}_l > T_l$ for all $l \in \{l_0+1, \cdots, q\}$. Let us now define: $y_l = |(\hat{c}_l - \hat{f}_l) - (c_l - f_l)|$, $l \in \mathcal{L}$. Note that

$$\sum_{l=1}^{l_0} y_l = \sum_{l=l_0+1}^{L} y_l \tag{12}$$

since $\sum_{l\in\mathcal{L}}(\hat{c}_l - \hat{f}_l) = \sum_{l\in\mathcal{L}}(c_l - f_l) = C - R$. Recalling that $c_l - f_l \geq c_{l+1} - f_{l+1}$ $(1 \leq l < L)$ we have

$$\begin{aligned}\sum_{l=1}^{L}(\hat{c}_l - \hat{f}_l)^2 - \sum_{l=1}^{L}(c_l - f_l)^2 &= \sum_{l=1}^{L} y_l^2 + 2\left\{\sum_{l=1}^{l_0}(c_l - f_l)y_l - \sum_{l=l_0+1}^{L}(c_l - f_l)y_l\right\}\\ &\geq \sum_{l=1}^{L} y_l^2 + 2\left\{(c_{l_0} - f_{l_0})\sum_{l=1}^{l_0} y_l - (c_{l_0+1} - f_{l_0+1})\sum_{l=l_0+1}^{L} y_l\right\}\\ &= \sum_{l=1}^{L} y_l^2 + 2\{(c_{l_0} - f_{l_0}) - (c_{l_0+1} - f_{l_0+1})\}\sum_{l=1}^{l_0} y_l > 0 \quad (13)\end{aligned}$$

where the last equality is obtained using (12). From (4), (6), and (13), we have[5]

$$\hat{\lambda}^1 = \frac{\sum_{l\in\mathcal{L}^1} c_l - R^{-1}}{\sum_{l\in\mathcal{L}^1}(\hat{c}_l - \hat{f}_l)^2} < \frac{\sum_{l\in\mathcal{L}^1} c_l - R^{-1}}{\sum_{l\in\mathcal{L}^1}(c_l - f_l)^2} = \lambda^1 \quad (14)$$

since $\mathcal{L}^1 = \mathcal{L}$. But this contradicts the assumption $\hat{\lambda}^1 > \lambda^1$. Therefore, $\hat{\lambda}^1 \leq \lambda^1$. Lemma 7, then, implies that $\hat{T}_l \leq T_l$ for all $l > q$, and $\hat{\lambda}^i \leq \lambda^i$ for all $i \in \mathcal{I}_q$. Thus, $\sum_{i\in\mathcal{I}_q}\hat{\lambda}^i \leq \sum_{i\in\mathcal{I}_q}\lambda^i$. As explained in [18], this implies that $\hat{T}_{q-1} \leq T_{q-1}$. Applying Lemma 8 inductively for $l = q-2, \cdots, 2$, it follows that $\hat{T}_l \leq T_l$, for every link l in $\{2, \cdots, q-1\}$. This concludes the proof of the second statement in the theorem.

It remains to be shown that $\hat{\lambda}^i \leq \lambda^i$, for all users $i \in \mathcal{I}$. Assume by contradiction that there exists some user j, such that $\hat{\lambda}^j > \lambda^j$. Then, $j \in \mathcal{I}\backslash\mathcal{I}_q$, since $\hat{\lambda}^i \leq \lambda^i$ for all $i \in \mathcal{I}_q$. Therefore, $\hat{T}_l \leq T_l$, for all $l \in \mathcal{L}^j$. Since $f_l^j T_l' + T_l = \lambda^j < \hat{\lambda}^j = \hat{f}_l^j \hat{T}_l' + \hat{T}_l$, this implies that $\hat{f}_l^j > f_l^j$ for all $l \in \mathcal{L}^j$. Thus, $r^j = \sum_{l\in\mathcal{L}^j}\hat{f}_l^j > \sum_{l\in\mathcal{L}^j} f_l^j = r^j$, which is a contradiction. Therefore, for all $i \in \mathcal{I}$, we have $\hat{\lambda}^i \leq \lambda^i$, and this completes the proof. ■

We are now ready to prove the main result of this section, namely that the capacity configuration that is obtained by allocating the entire additional capacity Δ to link 1 is user price optimal in $\mathcal{C}_\Delta$.

Theorem 2: Consider a system of parallel links with initial capacity configuration $\mathbf{c}^0$, shared by I noncooperative users, and an additional capacity allowance Δ. The capacity configuration $\mathbf{c}^* = \mathbf{c}^0 + \Delta\mathbf{e}_1$, that results from allocating the entire additional capacity to the link with the initially highest capacity, is user price optimal in $\mathcal{C}_\Delta$.

Proof: Let $\mathcal{D}_\delta$ denote the subspace of routing games that is generated by allocating an additional capacity of exactly δ, $0 \leq \delta \leq \Delta$, to a system of parallel links with initial capacity configuration $\mathbf{c}^0$. Then $\mathcal{C}_\Delta = \cup_{0\leq\delta\leq\Delta}\mathcal{D}_\delta$. For every δ, define $\mathbf{c}^*(\delta) = \mathbf{c}^0 + \delta\mathbf{e}_1$. Theorem 1 , then, implies that $\mathbf{c}^*(\delta)$ is user price optimal in $\mathcal{D}_\delta$. To see this consider any $\mathbf{c} \in \mathcal{D}_\delta$. From Theorem 1, the capacity configuration $\mathbf{c}+(c_L - c_L^0)(\mathbf{e}_1 - \mathbf{e}_L)$ is user price efficient compared to $\mathbf{c}$. Proceeding inductively, for every $m > 1$, the configuration $\mathbf{c}+\sum_{l=m}^{L}(c_l - c_l^0)(\mathbf{e}_1 - \mathbf{e}_l)$ is user price efficient compared to $\mathbf{c}+\sum_{l=m+1}^{L}(c_l - c_l^0)(\mathbf{e}_1 - \mathbf{e}_l)$. Hence, $\mathbf{c}^0 + \delta\mathbf{e}_1 = \mathbf{c} + \sum_{l=2}^{L}(c_l - c_l^0)(\mathbf{e}_1 - \mathbf{e}_l)$ is user price efficient with respect to $\mathbf{c}$, that is, $\mathbf{c}^*(\delta)$ is user price optimal in $\mathcal{D}_\delta$. From Lemma 2, $\mathbf{c}^* = \mathbf{c}^*(\Delta)$ is user price efficient with respect to any $\mathbf{c}^*(\delta)$ with $0 \leq \delta < \Delta$. Therefore, $\mathbf{c}^*$ is user price optimal in $\mathcal{C}_\Delta$. ■

[5] Equations (4) and (6) give $\lambda^1(c_l - f_l)^2 = c_l - f_l + f_l^1$, and summing over $l \in \mathcal{L}^1$, the equalities in (14) follow.

The following theorem shows that $\mathbf{c}^*$ is user cost optimal, if the lower bounds on the link capacities are equal for all links.

Theorem 3: Consider a system of parallel links with initial capacity configuration $\mathbf{c}^0$, shared by I users, and an additional capacity allowance Δ. If $c_l^0 = c_m^0$, for all $l, m \in \mathcal{L}$, then the capacity configuration $\mathbf{c}^* = \mathbf{c}^0 + \Delta\mathbf{e}_1$, that results from allocating the entire additional capacity to link 1, is user cost optimal in $\mathcal{C}_\Delta$.

The following two propositions establish that the user price optimal capacity configuration $\mathbf{c}^*$ is also user cost optimal in some special cases of interest.

Proposition 2: Consider a system of parallel links shared by I users, consistent at all capacity configurations in $\mathcal{C}_\Delta$. The capacity configuration $\mathbf{c}^* = \mathbf{c}^0 + \Delta\mathbf{e}_1$ is user cost optimal in $\mathcal{C}_\Delta$.

Note that the above result applies to the special cases of simple and of identical users.

Proposition 3: Consider a system of parallel shared by two users. The capacity configuration $\mathbf{c}^* = \mathbf{c}^0 + \Delta\mathbf{e}_1$ is user cost optimal in $\mathcal{C}_\Delta$.

F. General Topologies

The example presented in the Introduction shows that adding capacity to a network, even in infinite amounts, may result in an increase of both the price and the cost of each and every user. This indicates that an upgrade of a general network, in terms of capacity and link addition, should be carried out in a cautious way. In this subsection we devise methods for upgrading a general network, so that the Braess paradox does not occur. The terminology introduced in previous subsections for the parallel links model, readily extends to the general case. Due to space limits the details are omitted; the reader is referred to [18].

We consider now a network $(\mathcal{V}, \mathcal{L})$, where $\mathcal{V}$ is a finite set of nodes and $\mathcal{L} \subseteq \mathcal{V} \times \mathcal{V}$ is a set of directed links. A set $\mathcal{I}$ of users share the network and ship flow from a common source s to a common destination d. User i has a throughput demand that is some process with average rate r^i, and ships its flow by splitting this demand through the various paths connecting the source to the destination, according to its performance objective. The terms of user flow f_l^i, user routing strategy $\mathbf{f}^i$, user strategy space F^i and system flow configuration $\mathbf{f}$, originally defined in the context of parallel links, readily extend to general topologies, except that now the strategy space F^i of user i should account for the conservation of flow at the nodes [18]. The cost function J^i of user i is the sum of link cost functions J_l^i, taken over all network links $l \in \mathcal{L}$. The concepts of Nash equilibria, Nash mapping and optimality conditions are derived similarly as for parallel links. The various versions of the design problem apply also for the case of general topologies.

The class of problems investigated in this paper is well defined if the Nash equilibrium, under any capacity configuration, is unique. Whether this property holds in general topologies is an open question. Thus, we shall concentrate on cases for which uniqueness has been established, such as those of identical users and simple users [10].

The following proposition shows that the potential danger of degradation in performance can be avoided by upgrading the network "uniformly," i.e., by multiplying the capacity of each link by some constant factor $\alpha > 1$.

Proposition 4: In a general topology, consider two capacity configurations $\hat{\mathbf{c}}$ and $\mathbf{c}$, such that $\hat{c}_l = \alpha c_l$ for all $l \in \mathcal{L}$, $\alpha > 1$. Then:

1) If the users are simple, then $\hat{\mathbf{c}}$ is user price and cost efficient relative to $\mathbf{c}$.
2) If the users are identical, then $\hat{\mathbf{c}}$ is user price efficient relative to $\mathbf{c}$; moreover, for $\alpha > I$, $\hat{\mathbf{c}}$ is also user cost efficient.

Consider now an upgrade achieved by adding capacity to a direct link between the source s and the destination d (and, as a special case, adding a new (s,d) link). Denote by $\mathbf{c}$ and $\hat{\mathbf{c}}$, respectively, the capacity configurations before and after this addition. We say that $\hat{\mathbf{c}}$ is a *direct augmentation* of $\mathbf{c}$. We then have:

Proposition 5: In a general topology, consider two capacity configurations $\hat{\mathbf{c}}$ and $\mathbf{c}$, such that $\hat{\mathbf{c}}$ is a direct augmentation of $\mathbf{c}$. Then:

1) If the users are simple, then $\hat{\mathbf{c}}$ is user price and cost efficient relative to $\mathbf{c}$.
2) If the users are identical, then $\hat{\mathbf{c}}$ is user price efficient relative to $\mathbf{c}$.

This result suggests that yet another way to avoid the paradox is to upgrade the network through direct connections between source and destination.

IV. Architecting the Flow Configuration in the Run-Time Phase

Improvements of the overall performance of a noncooperative network can be achieved not only in the provisioning phase, but also during the actual operation of the network. In this section we demonstrate this approach based on the noncooperative routing model described in Section II. We assume that, apart from the flow generated by the self-optimizing users, there is also some flow whose routing is controlled by a central network entity, that will be referred to as the "manager." Typical examples of such flows are the traffic generated by signaling and/or control mechanisms, as well as traffic of users that belong to virtual networks. The manager has the following goals and capabilities: (1) it aims at optimizing the system performance, i.e., the average delay of *all* flow in the network, and (2) it is cognizant of the noncooperative structure of user routing. The first property makes the manager just another user, whose cost function corresponds to the system's (rather than its own) performance. The second property, however, enables the manager to predict the response of the users to any strategy that it chooses, and hence to determine a strategy of its own flow that would pilot them to a Nash equilibrium that minimizes the system's cost. Therefore, instead of *reacting* to the routing strategies of the users, the manager *fixes* this strategy and lets them converge to their respective equilibrium.

This is the typical scenario of a Stackelberg game [1], in which the manager plays the role of the "leader," and the noncooperative users play the role of the "followers."[6] The presence of sophisticated users that can acquire information about the demands and the cost functions of the other users and become Stackelberg leaders in order to optimize their own performance is in general undesirable [14]. In the problem considered here, however, the cost function of the manager is that of the system, and therefore it plays a social rather than a selfish role.

In this section we investigate the optimal strategy of the leader. In particular, we address the following question: is it possible for the leader to impose a strategy that drives the system into the network optimum, i.e., to the point that corresponds to the solution of a routing problem, in which the leader has full control over the entire flow? Intuitively, one would expect that the leader cannot enforce the network optimum, since it controls only part of the flow, while the rest is controlled by noncooperative users. Rather surprisingly, the results reported in the sequel show that in most cases the leader does have such capability. Due to space limits, we confine ourselves to a general and brief overview of the results; details can be found in [19]. We begin with an informal statement of the results:

1) In the special case of a single follower, the manager can always enforce the network optimum.
2) In the general case of any (finite) number of followers, the manager can enforce the network optimum if and only if its demand exceeds some threshold $\underline{r}^0$.
3) The threshold $\underline{r}^0$ is feasible, in the sense that the total demand of the users plus $\underline{r}^0$ is lower than the total capacity of the network.
4) In heavily loaded networks it is "easy" for the manager to enforce the network optimum (i.e., the threshold $\underline{r}^0$ is small).
5) As the number of users increases, it becomes harder for the manager to enforce the network optimum (i.e., the threshold $\underline{r}^0$ increases).
6) The higher the difference in the throughput demand of any two users, the easier it becomes for the manager to enforce the network optimum.

We proceed with a more detailed description of these results. Consider a system of parallel links $\mathcal{L} = \{1, \cdots, L\}$ shared by a set $\mathcal{I} = \{1, \cdots, I\}$ of noncooperative users (the followers), and the manager (the leader) that is labeled as user 0. Denote $\mathcal{I}_0 = \mathcal{I} \cup \{0\}$, and extend the notation of Section II in order to indicate the presence of the additional user 0. In particular, let $f_l = \sum_{i \in \mathcal{I}_0} f_l^i$ be the total flow on link $l \in \mathcal{L}$, and $\mathbf{f} = (\mathbf{f}^0, \mathbf{f}^1, \cdots, \mathbf{f}^I)$ the system flow configuration. Let $r = \sum_{i \in \mathcal{I}} r^i$ denote the total demand of the followers, and $R = r^0 + r$ the total demand offered to the network. We assume that $R < C$. Each user $i \in \mathcal{I}$ tries to minimize its individual cost function given by (2), while the manager aims

[6] The terms "manager" and "leader," as well as "users" and "followers," will be used interchangeably.

at minimizing the total cost of the system

$$J(\mathbf{f}) = \sum_{i=0}^{I} J^i(\mathbf{f}) = \sum_{l=1}^{L} \frac{f_l}{c_l - f_l} \tag{15}$$

that is proportional to the average time-delay experienced by the total flow offered to the network.

Each strategy $\mathbf{f}^0$ of the manager induces a unique Nash equilibrium, denoted by $\mathcal{N}^0(\mathbf{f}^0)$, of the noncooperative users, that can be determined from Proposition 1 by replacing $\mathbf{c}$ with $\mathbf{c} - \mathbf{f}^0$. The manager has knowledge of the noncooperative behavior of the users, and makes its routing decisions based on this information. In particular, the manager seeks a strategy $\mathbf{f}^0 \in F^0$ that minimizes $J(\mathbf{f}^0, \mathcal{N}^0(\mathbf{f}^0))$. It is worth mentioning that this optimization problem is similar to the optimal capacity allocation problem studied in the previous sections. Indeed, the two problems are similar, in the sense that the manager modifies the link capacities that are available to the users. They are different, in the sense that its routing decisions incur a cost for the manager's flow that has to be accounted for in these decisions.

Let $(f_1^*, \cdots, f_L^*)$ denote the unique solution [19] to the problem of optimally routing the total demand R over the set of parallel links, i.e., the link flow configuration that minimizes the total cost of the system. $(f_1^*, \cdots, f_L^*)$ will be referred to as the *network optimum*, and it is determined from Proposition 1, by replacing f_l^i with f_l^*, and c_l^i with c_l, $l \in \mathcal{L}$. In the sequel, we consider the problem of finding a strategy of the manager that drives the system to the network optimum, i.e., a strategy $\mathbf{f}^0 \in F^0$ such that if $\mathbf{f}^{-0} = \mathcal{N}^0(\mathbf{f}^0)$, then $\sum_{i \in \mathcal{I}_0} f_l^i = f_l^*$ for all $l \in \mathcal{L}$. Any such strategy of the manager achieves the minimal cost of the system and, therefore, leads to the most efficient utilization of network resources. Accordingly, let us introduce the following:

Definition 6: Let $\mathbf{f}^0 \in F^0$ be a strategy of the manager and $\mathbf{f}^{-0} = \mathcal{N}^0(\mathbf{f}^0)$. Strategy $\mathbf{f}^0$ is called *maximally efficient* if it achieves the network optimum, i.e., if $\sum_{i \in \mathcal{I}_0} f_l^i = f_l^*$ for all $l \in \mathcal{L}$.

Note that, although an optimal strategy of the manager always exists [19], existence of a maximally efficient strategy cannot be guaranteed, in general. Evidently, if a maximally efficient strategy exists, then it is an optimal strategy of the manager.

In the sequel, we present necessary and sufficient conditions that guarantee existence of a maximally efficient strategy of the manager. Moreover, provided that these conditions are met, we show that the maximally efficient strategy of the manager is unique and we specify its structure explicitly. To that end, define

$$H_l = \sum_{n=1}^{l-1} f_n^* - \frac{f_l^*}{c_l} \sum_{n=1}^{l-1} c_n, \qquad l = 2, \cdots, L \tag{16}$$

$$H_0 = 0, \qquad H_{L+1} = \sum_{n=1}^{L} f_n^* = R.$$

Then, as shown in [19], we have $H_l \le H_{l+1}$, for all $l \in \mathcal{L}$.

Consider first the case of a single follower. Except for being the simplest case of the general Stackelberg routing game, this case is of interest since it represents practical situations, in which different types of traffic (say, "system" and "data") are routed by different entities, one of which is cognizant of the operation of the other, hence the leader-follower setting.

Theorem 4: In the single-follower Stackelberg routing game, there exists a unique maximally efficient strategy $\mathbf{f}^0$ of the leader that is given by

$$f_l^0 = \begin{cases} c_l \frac{\sum_{n=1}^{L^1} f_n^* - r^1}{\sum_{n=1}^{L^1} c_n}, & l = 1, \cdots, L^1 \\ f_l^*, & l = L^1 + 1, \cdots, L \end{cases} \tag{17}$$

where L^1 is determined by $H_{L^1} < r^1 \le H_{L^1+1}$.

The theorem indicates that, the leader can enforce the network optimum, independently of the relative sizes, in terms of demands, of the leader and the follower. In other words, it is enough to have control on just a nonzero portion of flow in order to "tame" a single selfish user.

We now proceed to the general case of any (finite) number of users. The following lemma describes the maximally efficient strategy of the leader, provided that such a strategy exists.

Lemma 9: In a multifollower Stackelberg routing game, if there exists a maximally efficient strategy $\mathbf{f}^0$ of the leader, then it is unique and is given by

$$f_l^0 = c_l \sum_{i \in \mathcal{I}_l} \frac{\sum_{n=1}^{L^i} f_n^* - r^i}{\sum_{n=1}^{L^i} c_n} - (I_l - 1) f_l^*, \qquad l \in \mathcal{L} \tag{18}$$

where, for every $i \in \mathcal{I}$, L^i is determined by $H_{L^i} < r^i \le H_{L^i+1}$, and for every $l \in \mathcal{L}$, $\mathcal{I}_l = \{i \in \mathcal{I} : l \le L^i\}$ and $I_l = |\mathcal{I}_l|$. In that case, the equilibrium strategy $\mathbf{f}^i$ of user $i \in \mathcal{I}$ is described by

$$f_l^i = \begin{cases} f_l^* - c_l \frac{\sum_{n=1}^{L^i} f_n^* - r^i}{\sum_{n=1}^{L^i} c_n}, & l = 1, \cdots, L^i \\ 0, & l = L^i + 1, \cdots, L. \end{cases} \tag{19}$$

Conversely, if $\mathbf{f}^0$ described by (18) is an admissible strategy of the leader, i.e., if $\mathbf{f}^0 \in F^0$, then it is its maximally efficient strategy.

Note that if the leader employs strategy $\mathbf{f}^0$, then (19) implies that the set of links used by follower i is precisely $\{1, \cdots, L^i\}$, thus $\mathcal{I}_l$ is the set of followers that send flow on link l. In general, $\mathbf{f}^0$ might fail to be an admissible strategy of the leader. In [19], we show that $\mathbf{f}^0$ is admissible, if and only if the demand of the leader is higher than a threshold $\underline{r}^0$.[7] Therefore, we have the following:

Theorem 5: There exists some $\underline{r}^0$, with $0 \le \underline{r}^0 < C - r$, such that the leader in a multifollower Stackelberg routing game can enforce the network optimum, if and only if its throughput demand r^0 satisfies $\underline{r}^0 \le r^0 < C - r$. Then, the maximally efficient strategy of the leader is given by (18).

From the theorem, it follows that, for any set of followers for which $r < C$, there is a (feasible) leader, with $\underline{r}^0 \le r^0 < C - r$, that can enforce the network optimum. Moreover, when $r \to C$, we have $\underline{r}^0 \to 0$, meaning that in heavily loaded networks it suffices to control just a small portion of the flow in order to drive the system into the network optimum. This result is quite encouraging, because it is in heavily loaded networks where the presence of a manager is particularly important.

[7] The expression for determining $\underline{r}^0$ can be found in [19].

If the leader can enforce the network optimum, it can determine its maximally efficient strategy given the throughput demand r^i of every follower $i \in \mathcal{I}$ and the network optimum $(f_1^*, \cdots, f_L^*)$, as indicated by Lemma 9. The network optimum can be readily computed from Proposition 1 given the total load R offered to the network. Hence, the leader needs information only about the throughput demand of every follower. Since user flows are accepted by means of some admission control mechanism, this information is available to the manager. Each time a user arrives to or departs from the network, the manager can readily adjust its strategy to the maximally efficient one, using the information about the throughput demand of that user. In that sense, the proposed mechanism of enforcing the network optimum by means of the manager's routing strategy is *scalable*.

The minimum throughput demand $\underline{r}^0$ that guarantees that the leader can enforce the network optimum depends on the number and the throughput demands of the followers. This dependence is summarized in the following two propositions. The first gives the dependence of $\underline{r}^0$ on the number of followers when their total throughput demand r is fixed. To simplify the formulation of the problem, we concentrate on the case of identical users. The proposition shows that as the number of users increases, the harder it becomes for the leader to enforce the network optimum.

Proposition 6: Suppose that the followers are identical and their total throughput demand r is fixed. Then, the minimum throughput demand $\underline{r}^0$ that enables the leader to enforce the network optimum is nondecreasing with the number of followers.

Let us now concentrate on the dependence of $\underline{r}^0$ on differences of the demands of the followers, when their total throughput demand r is fixed. The following proposition shows that the higher the difference in the throughput demand of any two followers, the easier it becomes for the leader to enforce the network optimum.

Proposition 7: Suppose that the total throughput demand r of the followers is fixed. Then, for any two followers j and k, the minimum throughput demand $\underline{r}^0$ that enables the leader to enforce the network optimum is nonincreasing with $|r^j - r^k|$. Therefore, $\underline{r}^0$ attains its maximum value when all followers are identical.

Let us now demonstrate the properties of $\underline{r}^0$ by means of a numerical example. We consider a system of parallel links with capacity configuration $\mathbf{c} = (12, 7, 5, 3, 2, 1)$, shared by I identical followers with total demand r. The threshold $\underline{r}^0$ of the leader is depicted in Fig. 5 as a function of r, for various values of I. We concentrate on total follower demands that exceed half the total capacity of the network. In the same figure, we also show the saturation line "$\underline{r}^0 + r = C$". From the figure, one can see that $\underline{r}^0$ always lies below the saturation line, in accordance with Theorem 5. Furthermore, $\underline{r}^0$ increases with the number of followers. An important observation from the figure is that $\underline{r}^0$ decreases as the total demand of the followers increases, not only in the heavy load region, but also for moderate loads.

Finally, in [19], we also consider the case of an infinite number of followers, i.e., the case of simple followers. In particular, we explain that the leader cannot enforce, in general, the network optimum. For that case, we specify the structure of an optimal strategy of the leader and we provide a simple algorithm to compute it.

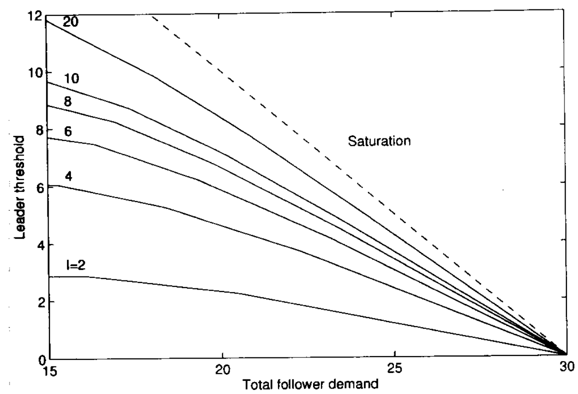

Fig. 5. Leader threshold as a function of total follower demand.

V. CONCLUSIONS

Design and management strategies for improving the performance of noncooperative networks were considered. A practical implication of this work is that design rules for noncooperative networks may follow the same simple patterns that apply to centrally controlled networks, and limited controllability can be as powerful as full controllability.

The first strategy called for devising proper design rules during the provisioning phase of the network. The problem was formulated as one of allocating additional capacity to an existing noncooperative network. Not only is this problem prohibitively complex and hard to analyze, but also it exhibits paradoxical behavior, according to which added resources might degrade user performance. For a system of parallel links we established that addition of capacity guarantees improved performance for all users. Given this result, we showed that the capacity allocation problem has a simple and intuitive solution: the optimal allocation assigns the additional capacity exclusively to the link with the initially highest capacity. It is worth noting that, although the noncooperative setting makes the analysis tedious, this solution coincides with the optimal capacity allocation when routing is centrally controlled.

The second strategy called for improving the performance of the network during its actual operation. This can be achieved by a management entity, that has control on only part of the network flow, and is cognizant of the presence of noncooperative users. Specifically, we considered a network manager that acts as a Stackelberg leader. Considering a system of parallel links, we showed that, in a wide range of cases, by controlling just a small portion of the total flow, the network operating point can be driven into the network optimum. This result suggests that, even with limited controllability, proper run time actions can diminish considerably, or even avoid altogether, the inefficiency implicated by the noncooperative behavior of the users.

Methodologies for upgrading general networks while avoiding the Braess paradox were also investigated. The related results indicate that capacity should be added across the network, rather than on a local (e.g., single link) scale. This fits well with engineering practice, where common folklore suggests that local improvement may result in transferring the problem somewhere else in the system. Another indication is that upgrades should be aimed at direct connections between the source and the destination. This is yet a further indication of the potential benefit of decoupling complex structures in a network, so that the controllers are presented with simple choices.

ACKNOWLEDGMENT

The authors would like to thank Raphael Rom and Nahum Shimkin for helpful discussions.

REFERENCES

[1] R. B. Myerson, *Game Theory: Analysis of Conflict*. Cambridge, MA: Harvard University Press, 1991.

[2] D. Fudenberg and J. Tirole, *Game Theory*. Cambridge, MA: M.I.T. Press, 1992.

[3] M.-T. T. Hsiao and A. A. Lazar, "Optimal decentralized flow control of Markovian queueing networks with multiple controllers," *Perform. Eval.*, vol. 13, no. 3, pp. 181–204, 1991.

[4] C. Douligeris and R. Mazumdar, "A game theoretic approach to flow control in an integrated environment," *J. Franklin Inst.*, vol. 329, pp. 383–402, Mar. 1992.

[5] Z. Zhang and C. Douligeris, "Convergence of synchronous and asynchronous greedy algorithms in a multiclass telecommunications environment," *IEEE Trans. Commun.*, vol. 40, pp. 1277–1281, Aug. 1992.

[6] E. Altman, "Flow control using the theory of zero-sum Markov games," *IEEE Trans. Automat. Contr.*, vol. 39, pp. 814–818, Apr. 1994.

[7] Y. A. Korilis and A. A. Lazar, "On the existence of equilibria in noncooperative optimal flow control," *J. Assoc. Comput. Mach.*, vol. 42, May 1995.

[8] A. A. Economides and J. A. Silvester, "Multi-objective routing in integrated services networks: A game theory approach," in *Proc. INFOCOM'91*, 1991, pp. 1220–1225.

[9] E. Altman and N. Shimkin, "Worst-case and Nash routing policies in parallel queues with uncertain service allocations," IMA Preprint Series 1120, Inst. Math., Applicat., Univ. Minn., Minneapolis, 1993.

[10] A. Orda, R. Rom, and N. Shimkin, "Competitive routing in multiuser communication networks," *IEEE/ACM Trans. Networking*, vol. 1, pp. 510–521, Oct. 1993.

[11] A. A. Lazar, A. Orda, and D. E. Pendarakis, "Virtual path bandwidth allocation in multi-user networks," in *Proc. IEEE INFOCOM'95*, Boston, MA, Apr. 1995, pp. 312–320.

[12] R. Cocchi, S. Shenker, D. Estrin, and L. Zhang, "Pricing in computer networks: Motivation, formulation and example," *IEEE/ACM Trans. Networking*, vol. 1, pp. 614–627, Dec. 1993.

[13] P. Dubey, "Inefficiency of Nash equilibria," *Math. Op. Res.*, vol. 11, pp. 1–8, Feb. 1986.

[14] S. Shenker, "Making greed work in networks: A game-theoretic analysis of switch service disciplines," in *Proc. ACM SIGCOMM'94*, Sept. 1994, pp. 47–57.

[15] W. I. Zangwill and C. B. Garcia, *Pathways to Solutions, Fixed Points and Equilibria*. Englewood Cliffs, NJ: Prentice-Hall, 1981.

[16] J. E. Cohen and F. P. Kelly, "A paradox of congestion in a queuing network," *J. Appl. Probab.*, vol. 27, pp. 730–734, 1990.

[17] C. Douligeris and R. Mazumdar, "Multilevel flow control of queues," in *Proc. Johns Hopkins Conf. Inform. Sci.*, Baltimore, MD, Mar. 1989.

[18] Y. A. Korilis, A. A. Lazar, and A. Orda, "Capacity allocation under noncooperative routing," Center for Telecommun. Re., Columbia Univ., New York, NY, CTR Tech. Rep. 372-94-19, 1994.

[19] Y. A. Korilis, A. A. Lazar, and A. Orda, "Achieving network optima using Stackelberg routing strategies," Center for Telecommun. Re., Columbia Univ., New York, NY, CTR Tech. Rep. 384-94-31, 1994.

[20] D. G. Luenberger, *Linear and Nonlinear Programming*, 2nd ed. Reading, MA: Addison-Wesley, 1984.

[21] W. Nicholson, *Intermediate Microeconomics and its Applications*, 6th ed. New York: Dryden Press, 1994.

Yannis A. Korilis (S'89) was born in Thessaloniki, Greece, in 1966. He received the Diploma from the National Technical University of Athens, Greece, in 1989, and the M.S., M.Phil., and Ph.D. degrees from Columbia University, New York, NY, in 1990, 1992, and 1995, respectively, all in electrical engineering.

During the summer of 1993, he worked as a Senior Staff Technologist at Bellcore, Morristown, NJ. His research interests include network control, network and service management, modeling, analysis and design of broadband networks, network pricing, economic and game-theoretic models for networking, and distributed algorithms and complexity.

Dr. Korilis has received various academic excellence awards from the National Scholarship Foundation and the Technical Chamber of Greece. He has also been a winner in the Interschool Fellowship Competition at Columbia University for four consecutive years. In May 1995, he received the E. I. Jury award for excellence in Systems, Communications or Signal Processing. He has organized the session "Networking Games" in the 23rd TIMS International Conference, Singapore, June 1995. He is a member of the ACM and INFORMS.

Aurel A. Lazar (S'77–M'80–SM'90–F'93) has been a Professor of Electrical Engineering at Columbia University, New York, NY, since 1988. His research interests span both theoretical and experimental studies of telecommunication networks and multimedia systems. The theoretical research he conducted during the 1980's pertains to the modeling, analysis, and control of broadband networks. He formulated optimal flow and admission control problems and, by building on the theory of point processes, derived control laws for Markovian queueing network models in single control as well as game theoretic settings. He was the Chief Architect of two experimental networks generically called MAGNET. This work introduced traffic classes with explicit quality-of-service constraints to broadband switching and led to the concepts of schedulable, admissible load, and contract regions in real-time control of broadband networks. He is currently leading the COMET project of the Center for Telecommunications Research at Columbia University on the foundations of the control and management architecture of future giant gigabit networks. His management and control research pioneered the application of virtual reality to the management of ATM-based broadband networks. He is also leading investigations into networking architectures supporting interoperable exchange mechanisms for interactive and on-demand multimedia applications with quality-of-service requirements. The main focus of this work is on building an open signalling infrastructure that enables the rapid creation, deployment, and management of multimedia services.

Ariel Orda (S'89–M'92) was born in Argentina in 1961. He received the B.Sc. (summa cum laude), M.Sc., and D.Sc. degrees in electrical engineering from the Technion–Israel Institute of Technology, Haifa, in 1983, 1985, and 1991, respectively.

From 1983 to 1985, he was a Teaching Assistant at the Technion. From 1985 to 1990, he served in the IDF as a Research Engineer. From 1990 to 1991, he was a Teaching Instructor at the Technion. In 1991, he joined the Faculty of Electrical Engineering at the Technion as a Lecturer. During the academic year 1993–1994 he was a Visiting Scientist with the Center for Telecommunications Research, Columbia University, New York, NY. Since 1994, he has been a Senior Lecturer at the Technion. During 1991–1994, he held consulting positions in AT&T Bell Laboratories and in the Israeli industry. During the summer of 1995, he was an Academic Visitor at IBM T. J. Watson Research Center. His current research interests include the control and management of broadband networks, the application of game theory to computer networking, and on-line network algorithms.

In 1991, Dr. Orda received the award of the Chief Scientist in the Ministry of Communication in Israel and a Gutwirth Award for Outstanding Distinction. In 1993, he received the research award of the Association of Computer and Electronic Industries in Israel.