Virtual Path Bandwidth Allocation in Multiuser Networks

Aurel A. Lazar, *Fellow, IEEE*, Ariel Orda, *Senior Member, IEEE*, and Dimitrios E. Pendarakis, *Member, IEEE*

***Abstract*— We consider a multiuser network that is shared by noncooperative users. Each user sets up virtual paths that optimize its own selfish performance measure. This measure accounts for the guaranteed call level quality of service, as well as for the cost incurred for reserving the resource. The interaction among the user strategies is formalized as a noncooperative game. We show that the game has a unique Nash equilibrium and that it possesses a certain fairness property. We investigate the dynamics of this game and prove convergence to the Nash equilibrium of both a Gauss–Seidel scheme and a Jacobi scheme. We extend our study to various general network topologies. Finally, the formal results and some extensions thereof are tested by emulating the schemes on an experimental network.**

***Index Terms*— Bandwidth allocation, distributed algorithms, game theory, Nash equilibrium, network control, virtual path.**

I. Introduction

A TRADEOFF encountered in communication networks is between bandwidth utilization and the complexity of call setup. Consider, for example, virtual circuit networks in which bandwidth is reserved upon call setup and released immediately upon call completion. This results in an optimal bandwidth utilization but at the expense of a complex per-call resource management. This complexity might prove prohibitive in large networks, such as emerging broadband asynchronous transfer mode (ATM)-based networks. In such systems, users would experience large call setup delays. In such cases, the user will be willing to sacrifice bandwidth utilization in order to reduce the setup complexity. A common way to obtain that is offered by the virtual path approach [15], [18]. A virtual path is a set of preestablished virtual circuits, dedicated to a specific source-destination pair. These resources are reserved for a time span longer than the duration of a specific call.

A major question arising from the virtual path approach is how the available bandwidth is to be distributed among users. By the term "user" we refer to some typically large network customer, such as a corporation setting up a virtual private network. According to the traditional approach, such a task is handled by a single administrative entity. In other words, users are essentially passive and would quite often reduce their own performance for the benefit of the entire network. Recently, it has been recognized that system-wide (single) administration is an impractical paradigm for the control of giant broadband networks. Moreover, the practical difficulties of coordinating among controllers in large-scale networks often require that each controller would effectively attempt to optimize its own performance. According to this formulation, the network is a common resource shared and competed by "selfish" users. This is a typical scenario of a "game" [13]. The application of game-theoretic tools to networking problems has been gaining increasing interest within the past few years [2], [3], [5], [7], [8], [11], [16], [20]. Competition on network resources can be observed, for example, in AT&T's customer-specified routing [14] and in networks implementing the intelligent network (IN) concept [4], such as the Deutsche Bundespost Telekom [19]. A natural competition among users arises when they attempt to reserve network bandwidth in order to establish their virtual paths.

An obvious implication of the virtual path scheme is that, typically, part of the reserved bandwidth will be unused (hence, the sacrifice in resource utilization). Yet another implication is that at times a user will experience lack of bandwidth to accommodate all incoming calls, thus resulting in call blocking. In a noncooperative ("game") scenario, as described above, each user would attempt to reserve the maximum possible amount of bandwidth for its own virtual path. Coupling this greedy behavior of users with the limited amount of bandwidth would lead to undesirable effects such as instability and unfairness.

As is typical of games [13], what is needed is to impose the interaction among the users on their individual strategies. This can be done by way of a cost function that depends on the availability of bandwidth in the network. It is important to note that, while users are interlinked via these costs, the specific parameters and control process of each user are localized within the user, in accordance with the framework described above. Noncooperative games often exhibit complex and even pathological behavior, e.g., lack or instability of equilibria. Thus, while the distributed noncooperative approach seems to be an appropriate one for large-scale multiuser networks, it calls for careful design and analysis.

Manuscript received October 25, 1995; revised October 1996 and June 1997; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor D. Mitra. This work was supported by the National Science Foundation under Grant CDR-88-11111. An earlier version of this paper was presented at the IEEE Infocom'95 Conference, 1995.

A. A. Lazar is with the Department of Electrical Engineering, Columbia University, New York, NY 10027 USA (e-mail: aurel@ctr.columbia.edu).

A. Orda was with the Center for Telecommunications Research, Columbia University, New York, NY 10027 USA. He is now with the Department of Electrical Engineering, Technion—Israel Institute of Technology, Haifa 3200, Israel (e-mail: ariel@ee.technion.ac.il).

D. E. Pendarakis was with Department of Electrical Engineering, Columbia University, New York, NY 10027 USA. He is now with the IBM Thomas J. Watson Research Center, Yorktown Heights, NY 10598 USA (e-mail: dimitris@watson.ibm.com).

Publisher Item Identifier S 1063-6692(97)09090-0.

After formulating the model, we investigate the equilibrium properties and the dynamics of the system. We show that a Nash equilibrium exists and that it is unique. We demonstrate that this unique equilibrium exhibits certain desirable properties, such as fairness. We then prove convergence to the unique equilibrium under the Gauss–Seidel and the Jacobi schemes [1]. In general, users may be sharing more than one resource, such that the underlying network of resources constitutes a graph of general topology. We extend several of our results to apply to these cases. Finally, the formal results and some extensions thereof are tested by emulating the schemes on an experimental network.

The paper is organized as follows. The model is formulated in Section II. In Section III we investigate the Nash equilibrium and prove its existence and uniqueness. The dynamics of the system and convergence to the equilibrium point are investigated in Section IV. Extensions to more general configurations are presented in Section V. The experimental results are presented and discussed in Section VI. Concluding remarks appear in Section VII. Due to space limitations, most technical proofs are omitted and they can be found in [9].

II. Problem Formulation

We consider a set of N users, denoted by $\mathcal{I} = \{1, 2, \cdots, N\}$, that share a resource of total capacity (bandwidth) of B units. We refer to this resource as a *(single) link*. The common resource may stand for more than a single link, as long as, from the users' standpoint, the capacity available in the various links can be "condensed." However, if different reservation rules apply at each link or if users prefer one link to the other (e.g., due to their qualities), then such a system cannot be treated as a single link and will be investigated in the more general context of Section V.

Each user reserves some of the resource capacity in order to establish a virtual path for its incoming calls. Calls that are blocked at the virtual path level, i.e., that find the corresponding virtual path full upon arrival, can be assumed to either be lost or else to be accommodated through an alternative virtual circuit scheme. In the latter case, the user faces the need to consume processing resources and waste time on call setup. Thus, in either case, blocking at the virtual path level leads to performance degradation.

Users are noncooperative, meaning that each seeks to optimize its own performance. To that end, the user minimizes a cost function. This function should account for the following tradeoff. On the one hand, each user should try to minimize the blocking probability of its incoming calls at the virtual path level, which is a decreasing function of the reserved capacity of the user's virtual paths. On the other hand, reserving capacity becomes more difficult (and thus costlier) as the system's resources are less available. This tradeoff is formally quantified as follows. Denote by C_i the *strategy* of user i, i.e., the amount of capacity that user i reserves. We can assume, without loss of generality, that a user always obtains the amount of capacity it requests (see [9] for a related discussion). The *strategy space* of user i is [0, B], i.e., $C_i \in [0, B]$. Hence, $\mathbf{C} = (C_1, C_2, \cdots, C_N)$ is the *game strategy vector*, and the *game strategy space* is $\mathcal{C} = \{\mathbf{C} \,|\, \forall i \, 0 \le C_i \le B\}$. We also denote by C the total amount of reserved capacity, i.e., $C = \sum_i C_i$. The cost function for user i, denoted by J_i, is of the following form:

$$J_i(\mathbf{C}) = J_i(C_i, C) = F_i(C_i, C) + G_i(C_i) \qquad (1)$$

where the function F_i accounts for the availability of resources as perceived by the ith user, whereas the function G_i accounts for the effect that the amount of reserved capacity has on the performance of that user. Following the above discussion, these functions are assumed to have the following properties:

- **F1** $F_i(\cdot, \cdot)$ monotonically increases in each of its two arguments.
- **F2** $F_i(C_i, C)$ is continuously differentiable with respect to C_i.
- **F3** $F_i(C_i, C)$ is strictly convex in C_i.
- **F4** $\partial F_i(C_i, C)/\partial C_i$ is nondecreasing with respect to C (and, due to [F3], it is also strictly increasing with respect to C_i).
- **F5** $\lim_{C_i \to B} F_i(C_i, C) = \infty$.
- **G1** $G_i(C_i)$ is continuously differentiable.
- **G2** $G_i(C_i)$ is strictly decreasing.
- **G3** $G_i(C_i)$ is convex.
- **G4** $\lim_{C_i \to 0} G_i(C_i) = \infty$.

We shall denote $F_i'(C_i, C) \triangleq \partial F_i(C_i, C)/\partial C_i$ and $G_i'(C_i) \triangleq dG_i(C_i)/dC_i$. We note that the F_i function increases with both the reservation demand of user i and with the total amount of reservations (assumption F1). Assumption F5 imposes the restriction that a user is not allowed to exhaust the resource on its own. This assumption will be further strengthened by assumption F6, in the sequel. The G_i function decreases with C_i, thus reflecting the improvement in user i's performance. Assumption G4 states that any user needs *some* amount of capacity. The reader may note that the above assumptions do not prevent users from reserving capacities in a total amount that exceeds the available capacity B. This issue will be addressed in the sequel. We make the following additional assumption concerning the entire model:

- **M1** For every strategy vector $\mathbf{C} = (C_1, C_2, \cdots, C_N)$ of reserved capacities, either all user costs $J_i(\mathbf{C}) = J_i(C_i, C)$ are finite, or else there is at least one user j with infinite cost $(J_j(C_j, C) = \infty)$ that can change its own strategy into, say, $\hat{C}_j$, such that its cost becomes finite, i.e., for $\hat{\mathbf{C}} = (C_1, \cdots, C_{j-1}, \hat{C}_j, C_{j+1}, \cdots, C_N)$, $J_i(\hat{\mathbf{C}}) < \infty$.

Since the cost function of each user depends on the strategies of all the other users, we are faced with a *noncooperative game* [13]. We are interested in the Nash equilibrium solution of the game, i.e., we seek a game strategy vector such that no user finds it beneficial to change its strategy. Formally, a game strategy vector $\mathbf{C}^*$ in the game strategy space $\mathcal{C}$ is a Nash equilibrium point (NEP) if, for all $i = 1, 2, \cdots, N$, the following holds:

$$\begin{aligned} J_i(\mathbf{C}^*) &= J_i(C_1^*, \cdots, C_{i-1}^*, C_i^*, C_{i+1}^*, \cdots, C_N^*) \\ &= \min_{0 \le C_i \le B} J_i(C_1^*, \cdots, C_{i-1}^*, C_i, C_{i+1}^*, \cdots, C_N^*). \;\Box \end{aligned} \qquad (2)$$

Assumption M1 immediately implies that, at any NEP, the costs of all users must be finite. Moreover, from assumptions F5, G4, and M1, we can deduce that the boundaries of the strategy space of any user are not reached under an optimal response of the user (to the strategies of the other users) and, in particular, at any NEP.

As already mentioned, the assumptions made thus far do not prevent users from reaching an NEP at which the total reserved capacity C is more than the available amount B. It is tempting to add an assumption of the form

$$\lim_{C \to B} F_i(C_i, C) = \infty$$

(and $F_i(C_i, C) = \infty$ for $C \geq B$). However, such an assumption would not be consistent with assumption M1. Indeed, consider a system of three users ($N = 3$) where each reserves a capacity of $B/2$. If we assume that $F_i(C_i, C) = \infty$ for $C \geq B$, then no user will be able to make its cost finite just by an action of its own. This is prohibitive in terms of fairness. Moreover, as will be observed in the sequel, assumption M1 is essential in our analysis for establishing the existence of a (finite) NEP. Thus, we must elaborate on the additional required assumption in a more careful way. Denote by B_i the capacity that is left available to user i, i.e., $B_i = B - \sum_{j \neq i} C_j$. Also, to each user i we assign a value $\epsilon_i > 0$. Then, the additional assumption is defined as follows:

$$\textbf{F6} \quad \lim_{C_i \to \max(B_i, \epsilon_i)} F_i(C_i, C) = \infty.$$

This assumption guarantees that the cost of a user cannot go to infinity solely because of the strategies of the other users. In other words, the term $\max(B_i, \epsilon_i)$ imposes a kind of "fairness" constraint: If the unreserved capacity available to user i is less than ϵ_i, user i will be "charged" *as if* the available capacity were ϵ_i. This ensures that user i can always reserve an amount $C_i < \epsilon_i$ while encountering a *finite* cost, regardless of the actions of other users. In view of this interpretation, the following assumptions are made on ϵ_i:

- **E1** $\sum_i \epsilon_i < B$, i.e., there is enough capacity to satisfy the minimal guaranteed amounts.
- **E2** There exists a $C_{i,\min}$, such that $C_{i,\min} \in (0, \epsilon_i)$ and $G_i(C_i) < \infty$ for all $C_i \in (C_{i,\min}, \epsilon_i)$, i.e., the amount guaranteed to each user complies with its required call-level quality of service (QoS).

Assumption E1 simply means that there is enough capacity to satisfy the minimal guaranteed amounts. Assumption E2 is needed when the function G_i goes to infinity before C_i goes to zero, i.e., when a user needs a strictly positive minimal amount of capacity in order to satisfy its QoS.[1] In such cases, assumption E2 states that ϵ_i is higher than that minimal amount.

Remark: We note that the above formulation accommodates the case in which a user has calls of different classes. In such a case, a user should find an appropriate scheme for sharing the capacity by the various types of calls, i.e., one that satisfies the requirements of each type. This amounts to computing the *schedulable region* [6]. The interpretation of the schedulable region in our context can be done in two alternative ways. According to the first, each user reserves a capacity C_i and then computes the schedulable region that corresponds to that C_i. Alternatively, a global schedulable region can be computed for the whole system and then subdivided among the users according to the values of C_i, i.e., user i would get a portion C_i/B of the schedulable region. We note that the second approach is applicable only in cases in which all users have the same type of calls.

Example: Motivated by [15] and [18], let the function F_i be of the following form:

$$F_i(C_i, C) = \begin{cases} \dfrac{C_i}{\max(B_i, \epsilon_i) - C_i}, & \text{if } C_i < \max(B_i, \epsilon_i) \\ \infty, & \text{otherwise.} \end{cases} \tag{3}$$

In other words, for each unit of reserved capacity the cost to user i is $1/(\max(B_i, \epsilon_i) - C_i)$. For G_i, we employ the Erlang-B loss function. Assume that the call arrival process of each user i is Poisson with rate α_i. Thus, the Erlang-B loss function corresponding to user i is a function of α_i and C_i, and is denoted by $E(\alpha_i, C_i)$. Let $\kappa_i \leq 1$ be an upper bound on user i's call blocking probability, as determined by its call-level QoS requirements. We then define G_i to be

$$G_i(C_i) = \begin{cases} \dfrac{1}{\kappa_i - E(\alpha_i, C_i)}, & \text{if } \kappa_i > E(\alpha_i, C_i) \\ \infty, & \text{otherwise.} \end{cases} \tag{4}$$

We now show that, for the above example and for parameters ϵ_i that satisfy E1 and E2,[2], our assumptions F1–F6 and G1–G4 hold. Consider first the F_i functions. Assumptions F1–F4 immediately follow by observing that $\max(B_i, \epsilon_i)$ is independent of C_i, and that $\max(B_i, \epsilon_i) - C_i = \max(B - C, \epsilon_i - C_i)$. Since $B_i \leq B$ (by definition) and $\epsilon_i < B$ (by assumption E1), we have that $\max(B_i, \epsilon_i) < B$, thus assumption F5 holds. Assumption F6 is immediate. Turning now to the G_i functions, we note that assumptions G1–G3 follow from the properties of the Erlang-B loss function shown in [12], whereas G4 follows from the fact that $\lim_{C_i \to 0} E(\alpha_i, C_i) = 1$ while $\kappa_i \leq 1$. Finally, assumption M1 holds because a user i can always choose a C_i such that $C_{i,\min} < C_i < \epsilon_i$, which makes $G_i(C_i) < G_i(C_{i,\min}) < \infty$ (by assumption E2) and $F_i(C_i, C) \leq \infty$.

III. Existence and Uniqueness of the Nash Equilibrium

We begin by showing that, for the model set in the previous section, an NEP exists. By its definition, the game strategy space $\mathcal{C}$ is a convex, closed, and bounded set $\mathcal{C} \subset R^N$. Furthermore, the cost function of each user $J_i(C_i, C)$ is, wherever finite, continuous in C_i and convex in C_i for every fixed value of C. By assumption M1, all cost functions must be finite at any NEP (if it exists). Therefore, as shown in [16], [17, Th. 1] can be accommodated to guarantee that an NEP exists for this game. We note that, in order to establish the existence of the NEP, it is enough to assume that the strategy spaces are $C_i \in [0, B]$, that the functions F_i and G_i are continuous (a

[1] We note that, in fact, this is the typical case. However, for generality, we assumed the more relaxing assumption G4.

[2] Note that E2 trivially holds for $\kappa_i \equiv 1$.

milder assumption than F2 together with G1), and that F3, G3, and M1 hold. In other words, existence per se is guaranteed with a more relaxed set of assumptions than the one that we imposed.

We proceed by proving that, at any NEP, the total amount of reserved capacity C is less than the resource capacity B.

Lemma 1: At any NEP $\mathbf{C} = (C_1, C_2, \cdots, C_N)$, $C = \sum_{i=i}^{N} C_i < B$.

Proof: Let $\mathbf{C} = (C_1, C_2, \cdots, C_N)$ be an NEP. If $C_i \le \epsilon_i$ for all i, $i = 1, 2, \cdots, N$, then $\sum_i C_i \le \sum_i \epsilon_i < B$, where the last inequality is due to assumption E1.

Assume, then, that there exists some i for which $C_i > \epsilon_i$. Since cost functions are finite at any NEP (due to assumption M1), it follows from assumption F6 that $C_i < B_i$, i.e., $C_i < B - \sum_{j \ne i} C_j$, which means that $C < B$. □

We observe that the cost function of each user is convex against any set of strategies of the other users. The Kuhn–Tucker optimality conditions [10] imply that a strategy vector $\mathbf{C} = (C_1, C_2, \cdots, C_N)$ is an NEP if and only if, for all users i, $1 \le i \le N$

$$F_i'(C_i, C) = -G_i'(C_i). \tag{5}$$

Note that the constraints $0 \le C_i$ and $C_i \le B$ are not manifested in the above conditions since assumptions G4 and F5 guarantee that the boundaries are not met at a (finite) optimality point.

The following theorem states the uniqueness of the NEP. We note that the result follows immediately from the more general uniqueness result obtained in the sequel for systems of parallel links; see Theorem 5.

Theorem 1: In a system of N noncooperative users sharing a resource of capacity B, as formalized in Section II, the NEP is unique.

Next, we derive a fairness property of the (unique) NEP. To that end, we make the further assumption that all cost functions F_i are *identical for all users*, i.e., for all $i \in \mathcal{I}$, $F_i(C_i, C) = F(C_i, C)$. The fairness property will guarantee that a user that is in greater "need" for capacity will indeed get more capacity at the NEP. In the following, all references to capacity values are to those at the NEP.

Definition 1: A user i is said to have a (strictly) higher loss-sensitivity than a user j if $G_i'(x) \le G_j'(x)$ (correspondingly, $G_i'(x) < G_j'(x)$) for all (positive) values of x.

Note that, since the derivative of the G function is negative, the second condition implies that i is more sensitive than j to call loss.

Definition 2: An NEP $\mathbf{C} = (C_1, C_2, \cdots, C_N)$ is said to be *fair* if, for each pair of users i and j, such that i has a (strictly) higher loss-sensitivity than j, we have $C_i \ge C_j$ (correspondingly, $C_i > C_j$).

We then have

Theorem 2: In a system of N noncooperative users sharing a resource of capacity B, as formalized above, in which all functions F_i are identical, the (unique) NEP is fair.

Proof: This is a particular case of Theorem 6, presented in Section V. The proof is thus omitted. □

The last theorem implies that, if all users have the same functions F_i and G_i, then at the NEP we have $C_i = C/N$ for all $i \in \mathcal{I}$.

IV. CONVERGENCE TO THE NEP

In this section we will explore the system dynamics and show convergence to the NEP under two different iterative schemes, namely Gauss–Seidel and Jacobi.

A. The Gauss–Seidel Scheme

In this subsection we consider a dynamic scheme in which users operate iteratively in an asynchronous fashion. At each iteration, each user recomputes its reserved capacity so as to optimize its cost function with respect to the current state of the system. More precisely, we assume that values of reserved capacities are recomputed in a sequence of steps. At each step only one user updates the value of its reserved capacity based on information on *the total capacity reserved by all other users* up to the previous step. No particular order in which users update their values is assumed; we will, however, require that the number of steps between any two subsequent updates by the same user is *finite*. Following common terminology, we will call this mode of operation the *Gauss–Seidel scheme*. Throughout this section we will denote by C^* and C_i^* the respective values at the (unique) NEP.

In what follows we will show that, under the Gauss–Seidel scheme, the vector of capacities reserved by all users at step n, $\mathbf{C}(n) = (C_1(n), C_2(n), \cdots, C_N(n))$, converges to the unique NEP $(C_1^*, C_2^*, \cdots, C_N^*)$, as $n \to \infty$. If the time interval between any two subsequent steps is *finite*, then $\mathbf{C}(n)$ converges to the unique NEP as $t \to \infty$.

Let $C(n) = \sum_{i=1}^{N} C_i(n)$ be the total capacity reserved by all users after step n, and define the following metric:

$$S(\mathbf{C}(n)) = \sum_{i=1}^{N} |C_i(n) - C_i^*| + |C(n) - C^*|. \tag{6}$$

It will be shown that $S(\mathbf{C}(n))$ converges to 0. We begin with the following lemma:

Lemma 2: Assume that user i performs the $(n+1)$th step, then

$$C(n+1) \ge C^* \Rightarrow C_i(n+1) \le C_i^*$$

(and, conversely)

$$C(n+1) \le C^* \Rightarrow C_i(n+1) \ge C_i^*.$$

Furthermore, the above holds also when all inequalities are strict.

Proof: See [9]. □

We next prove that the sequence $S(\mathbf{C}(n))$, $n \in \mathbb{N}$, is nonincreasing.

Lemma 3: Under the Gauss–Seidel scheme

$$S(\mathbf{C}(n+1)) \le S(\mathbf{C}(n))$$

for all $n \in \mathbb{N}$.

Proof: See Appendix. □

We now show that $S(\mathbf{C}(n))$ strictly decreases after a finite number of steps.

Lemma 4: For all n there exists a (finite) $k > n$, such that, if $(C_1(n), \cdots, C_N(n)) \neq (C_1^*, \cdots, C_N^*)$, then $S(\mathbf{C}(k)) < S(\mathbf{C}(n))$.

Proof: See [9]. □

We are now ready to prove the convergence of the Gauss–Seidel scheme to the unique NEP.

Theorem 3: Under the Gauss–Seidel scheme, the vector of capacities reserved by all users at step n, $\mathbf{C}(n) = (C_1(n), C_2(n), \cdots, C_N(n))$, converges to the unique NEP $(C_1^*, C_2^*, \cdots, C_N^*)$ as $n \to \infty$, that is, $\lim_{n\to\infty} \mathbf{C}(n) = \mathbf{C}^*$.

Proof: Follows from Lemma 4 and continuity arguments. For the details see [9]. □

B. The Jacobi Scheme

At each step of the Jacobi scheme, the values of reserved capacities of all users are updated simultaneously. We note that the assumption on simultaneous updating should not be interpreted as a strict timing synchronization requirement on the users, but rather as the more relaxed assumption that, at each step, the network provider will consider capacity reservations placed by all users simultaneously (however, users can still place their requests at any instance in time).

We observe that, in general, this scheme can lead to unstable system behavior. For example, at some iteration each user might reserve an amount of $2B/N$, thus ending with a total reserved capacity of $2B$. In the next iteration, each user i would reserve an amount that is less than ϵ_i, thus resulting in a new total reserved capacity that is (possibly much) lower than B, thus enabling each user to reserve, once again, an amount in the order of $2B/N$, etc. Therefore, in general, we cannot expect such a scheme to converge. However, we will show that a modified version of the Jacobi scheme does converge to the unique NEP. Denote by $\hat{C}_i(n+1)$ the capacity that user i would reserve at the nth step of the (regular) Jacobi scheme. Then, rather than reserving this value, user i reserves

$$C_i(n+1) = C_i(n) + \frac{1}{N}(\hat{C}_i(n+1) - C_i(n)) \qquad (7)$$

where $1/N$ is a "damping" constant and N is the total number of users in the network.

As with the Gauss–Seidel scheme, we define a metric $S(\mathbf{C}(n))$, show that it is decreasing in n, and conclude that the modified Jacobi scheme converges to the unique NEP. The metric to be used is

$$S(\mathbf{C}(n)) = \sum_{i=1}^{N} |C_i(n) - C_i^*| + |C(n) - C^*|. \qquad (8)$$

We note that $S(\mathbf{C}(n))$ has the same form as the metric defined for the Gauss–Seidel scheme. However, note that the capacity values are not those derived from the exact optimization condition, but from their damped counterparts, calculated according to (7). Also recall that, at step n, *all* users update their capacity. In [9], the following result is established:

Lemma 5: If $\mathbf{C}(n) \neq \mathbf{C}^*$, then $S(\mathbf{C}(n+1)) < S(\mathbf{C}(n))$.

We then obtain the following convergence property:

Theorem 4: Under the modified Jacobi scheme, $\lim_{n\to\infty} \mathbf{C}(n) = C^*$.

Proof: Follows from Lemma 5 in similar lines to the proof of Theorem 3. □

V. EXTENSIONS

As explained in the previous section, the resources available to users cannot always be modeled by a single link. For example, capacity might be available in various links of different quality, thus a user would have preferences among the links. Moreover, such links may not directly connect the source node to the destination node, meaning that resources may take the form of a graph of general topology. Such extensions to our model are investigated in this section and in [9].

A. Parallel Links

Consider the following extension of the model. Instead of a single link, users share a system of L parallel links, denoted as $\mathcal{L} = \{1, 2, \cdots, L\}$, each with a total capacity B^l. This corresponds to the case in which there are various distinguishable resources that *directly* interconnect a source to a destination. User i reserves an amount of capacity C_i^l on a link l, subject to the following constraints:

- $C_i^l \geq 0$
- $C_i^l \leq B^l$.

We denote by $C^l = \sum_i C_i^l$ the total capacity reserved on a link l. Also, we denote by $C_i = \sum_l C_i^l$ the total capacity reserved (on all links) by user i. The cost function of a user is composed of the sum of two types of functions, namely F_i^l and G_i.

As before, F_i^l is a function of two arguments, namely the capacity reserved by i on link l (i.e., C_i^l) and the total capacity reserved there by all users (C^l). In other words, F_i^l accounts for the cost of reserving capacity for a user on each link, as perceived by that user. Note that, according to this formulation, a user may pay different prices for reserving the same total amount of capacity, but distributing it differently among the links. Moreover, the F_i^l functions can reflect user preferences among the links, for example, if some link l has some undesirable property, such as high propagation delay, a user i which is sensitive to end-to-end delay could add to the respective F_i^l function a large additive constant or multiply it by a large constant factor, etc. We make the same assumptions on the F-functions as before, that is:

PF1 $F_i^l(\cdot, \cdot)$ monotonically increases in each of its two arguments.

PF2 $F_i^l(C_i^l, C^l)$ is continuously differentiable with respect to C_i^l.

PF3 $F_i^l(C_i^l, C^l)$ is convex in C_i^l.

PF4 $\partial F_i^l(C_i^l, C^l)/\partial C_i$ is nondecreasing with respect to C^l (and, due to [PF3], it is also strictly increasing with respect to C_i^l).

PF5 $\lim_{C_i^l \to B^l} F_i^l(C_i^l, C^l) = \infty$.

A user can assign an incoming call to any link (or combination of links) on which the amount of capacity, which has been reserved by the user but was not used yet, can accommodate that call. This means that the loss process of a user depends only on the total amount of capacity reserved by that user and not on the precise distribution of that capacity among the links. As a consequence, the cost function G_i (which accounts for the loss process), takes as its argument the total capacity C_i reserved by user i. The assumptions made on G_i are the same as before.

We still need an assumption on the F_i^l functions that is similar to F6, i.e., that guarantees that the resource B^l at each link is not exhausted. To that end, and following the previous terminology, denote by B_i^l the capacity that is left available to user i at link l, i.e., $B_i^l = B^l - \sum_{j \neq i} C_j^l$. Again, to each user i and for each link l we assign a value ϵ_i^l; however, in the case of parallel links it is enough to assume that $\epsilon_i^n > 0$ for *at least one* link n (rather than for all links), while at other links l the value may be $\epsilon_i^l = 0$. We also denote $\epsilon^i = \sum_l \epsilon_i^l$. The additional assumption is then defined as follows:

PF6 $\lim_{C_i^l \to \max(B_i^l, \epsilon_i^l)} F_i^l(C_i^l, C^l) = \infty$.

As before, the following assumptions are made on ϵ_i^l:

PE1 $\sum_i \epsilon_i^l < B^l$, i.e., at each link, there is enough capacity to satisfy the minimal guaranteed amounts.

PE2 There exists a $C_{i,\min}$ with $C_{i,\min} \in (0, \epsilon_i)$, such that $G_i(C_i) < \infty$ for all $C_i \in (C_{i,\min}, \epsilon_i)$.

From the above discussion it follows that, as far as call routing is concerned, all links behave as one logical link whose total capacity is $B = \sum_{l \in \mathcal{L}} B^l$. The distinction among links is in the different costs for reserving capacities, which is manifested by the F_i^l-functions.

To conclude, the cost function J_i of a user i is defined as e12

$$J_i(\mathbf{C}) = \sum_l F_i^l(C_i^l, C^l) + G_i(C_i).$$

We shall denote $F_i^{l'}(C_i^l, C^l) \triangleq \partial F_i^l(C_i^l, C^l)/\partial C_i^l$ and, as before, $G_i'(C_i) \triangleq dG_i(C_i)/dC_i$.

Here too we assume that, at any configuration of capacities, each user can change its own capacity to make its cost finite (i.e., assumption M1).

The strategy space of each user i on each link l is $C_i^l \in [0, B^l]$. The existence of the NEP follows from [17, Th. 1] in an identical way as for the single link case. Also, the proof that, at any NEP, the total amount of reserved capacity C^l at any link l does not exceed the resource capacity B^l follows as in Lemma 1.

Note that, due to the assumptions made on the cost function, the boundary B^l is not reached by a user i on any link l at any optimal response of that user and, in particular, at the NEP. Thus, the Kuhn–Tucker conditions ignore the upper-bound constraint on C_i^l (but not the lower-bound constraint) and they can be stated in the following form:

$$F_i^{l'}(C_i^l, C^l) \geq \lambda_i; \quad F_i^{l'}(C_i^l, C^l) = \lambda_i, \qquad \text{if } C_i^l > 0 \quad \forall i, l \tag{9}$$

$$G_i'(C_i) = -\lambda_i \tag{10}$$

where λ_i are nonnegative Lagrange multipliers [10].

The following result establishes the uniqueness of the NEP for a network of parallel links.

Theorem 5: In a system of N noncooperative users sharing L parallel links as formalized above, the NEP is unique.

Proof: Part of the proof follows similar lines to the proof of [16, Th. 5.1]. For details, see [9]. □

We borrow the definitions of *loss sensitivity* and *fairness*, as stated in the previous section, only that now the definition of fairness applies to the total capacity C_i reserved by each user. We derive a fairness property for the (unique) NEP, similar to the one derived for a case of a single link.

Theorem 6: In a system of N noncooperative users sharing L parallel links as formalized above, for which the functions F_i^l are identical, the (unique) NEP is fair.

Proof: See [9]. □

As was the case with a single link, the last theorem implies that, if all users have the same functions F_i^l and G_i, then the NEP is symmetrical, i.e., for all $l \in \mathcal{L}$ and for all $i \in \mathcal{I}$ we have $C_i^l = C^l/N$.

Suppose that users are numbered by increasing order of loss-sensitivity. Suppose that, at the NEP, user i refrains from using link l. It follows from Theorem 6 that so does any user j with a smaller loss-sensitivity, i.e., $C_i^l = 0$ implies $C_j^l = 0$ for $j < i$. Thus, at equilibrium, we can partition the set of links into a sequence of sets $\mathcal{L}_1, \mathcal{L}_2, \cdots, \mathcal{L}_N$, such that $\mathcal{L}_n \subseteq \mathcal{L}$ for $1 \leq n \leq N$ and $\mathcal{L}_n$ is the set of links that is used exclusively by users $n, n+1, \cdots, N$. We have that $\mathcal{L}_n \supset \mathcal{L}_{n+1}$; also, since each user, including the one with the smallest loss-sensitivity, should use some link, we have that $\mathcal{L}_1 \neq \emptyset$ (other sets may be empty).

B. General Topologies

In [9], we consider a network of general topology in which capacities are reserved on a multihop basis. We further extend the model and some of the results of Section V-A. In particular, we indicate that an NEP exists and that if users share the same source and destination nodes and employ the same cost functions, the NEP is unique and symmetrical.

VI. Experimental Results

In this section we present emulation results of the iterative, noncooperative schemes studied in Section IV, performed on a network of interconnected workstations.

The objective of our experimentation was twofold. First, through emulation, we could test the behavior of our schemes when employed in a practical setting, involving communication delays, protocol overhead, computational load, etc. Second, the experimentation allowed us to study additional properties, beyond what was available through formal analysis. In particular, we were interested in the dependence of the convergence rate on the form of the cost function, the scheme employed, the number of users, the level of link utilization and the arrival or departure of users to the system.

An issue of primary interest is to determine how these properties can be affected by the network manager. Therefore, in Section VI-B we introduce a parametrized cost function that allows architecting of the user's strategies. In Section VI-C we present the results of three experiments. The first explores the dependence of convergence rate on the parameters of the cost function. The second seeks to characterize the impact of the addition of new users to the system. Finally, the third experiment investigates the impact of the number of users on the convergence rate.

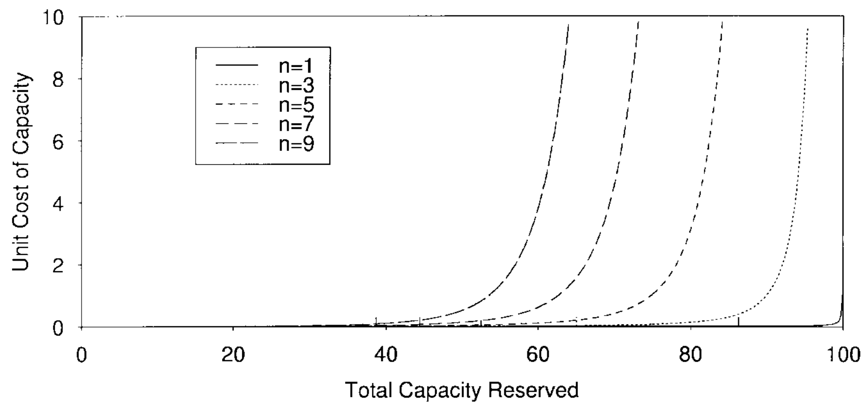

Fig. 1. Family of link capacity cost functions for different values of the n parameter.

A. Protocol Description

We assume that the capacity of each link is allocated by a link or switch controller, which is responsible for virtual circuit and virtual path setup and tear-down. Individual users communicate with the controller to inquire about the state of the link, but not among themselves.

In our implementation, controller and users were modeled as separate processes; communication with the controller was achieved using UDP (the IP user datagram protocol) sockets.

The asynchronous Gauss–Seidel iterative scheme requires a *locking* scheme to ensure that, at any time, at most one user can modify its reserved capacity. The iterative Jacobi scheme employs a "broadcasting" scheme to "synchronize" all capacity reservation requests. The controller broadcasts the current cost of capacity to all users in periodic time intervals.

In the following, we will be using the terms link capacity cost and link state (meaning total link capacity reserved) interchangeably, since one uniquely determines the other.

B. The Form of the Cost Function

The network manager can architect the strategies of the users by appropriately choosing the cost function $F_i(C_i, C)$. By changing a parameterized cost function, the manager pilots the users into achieving an NEP that has desirable characteristics such as the percentage of total capacity reserved and the rate of convergence.

In order to understand the capabilities of the manager, we have used in our emulation the G function defined in (4) and a modification of the F function of (3), as follows:

$$F_i(C_i, C) = C_i\left(\alpha_1 + \frac{\alpha_2}{\left(1-\frac{C}{B}\right)^n}\right) \tag{11}$$

where B is the total link capacity and α_1, α_2, n are positive real numbers. Note that $1/(1-\frac{C}{B})^n$ is an increasing function of C that approaches infinity as $C \rightarrow B$. The constants in (11) can be interpreted as follows: α_1 represents a fixed cost per unit of capacity, while α_2 represents a unit cost attributed to "congestion." Congestion is (loosely) defined as a situation in which the total reserved capacity C approaches the link capacity B. The parameter n determines how early congestion is detected; alternatively, it can express the degree of uncertainty in estimating the schedulable region of the link when shared by more than a single traffic class. We call n the *congestion avoidance parameter*. A sample of functions for the unit cost of link capacity is shown in Fig. 1 for different values of the parameter n. In this figure, $\alpha_1 = 0.01$, $\alpha_2 = 0.001$, and $B = 100$ Mb/s.

As shown in Fig. 1, the unit cost of capacity $F_i(C_i, C)/C_i$ exhibits two basic regions: one in which $F_i(C_i, C)/C_i \simeq (\alpha_1 + \alpha_2)$, i.e., the unit cost of capacity is almost constant, and one in which the unit cost increases rapidly with the total reserved capacity. We call the first the *low utilization region* and the second the *congestion region*. For small values of n, the congestion region starts at higher values of total reserved capacity, but its slope increases. Intuitively, when users operate in the first region, their strategies are decoupled and the cost of capacity is independent of other users' actions. When users operate in the congestion region, small changes in the capacity reserved by one user may affect dramatically the capacity cost of other users.

C. Results

As mentioned in the introduction to Section VI, we would like to explore how properties like the convergence rate can be influenced by the network manager. Using the cost function of (11), we conducted the following experiment in order to determine the dependence of the convergence rate on the parameter n. We set the total link capacity to $B = 100$ Mb/s, the capacity requirement of each call to 1 Mb/s and $\alpha_1 = 0.01$, $\alpha_2 = 0.001$. We then considered four users, with traffic loads $a_1 = 10$ Erl, $a_2 = 20$ Erl, $a_3 = 20$ Erl, and $a_4 = 30$ Erl, respectively. For various values of the parameter n, Fig. 2(a) and (b) shows the number of iterations for achieving convergence[3] for both the Gauss–Seidel and the Jacobi schemes, while Fig. 2(c) shows the value of total capacity reserved C^* at the NEP. As the plots indicate, small values of n lead to high capacity utilization at the expense of slow convergence rates. This is due to the fact that if the total demand is high, users will operate close to the congestion region and their strategies will be strongly coupled, thus leading to a large number of iterations. However, as the

[3]We assumed that the algorithm has converged when $|C_i - C^*|/C_i < 10^{-5}$, $\forall i$.

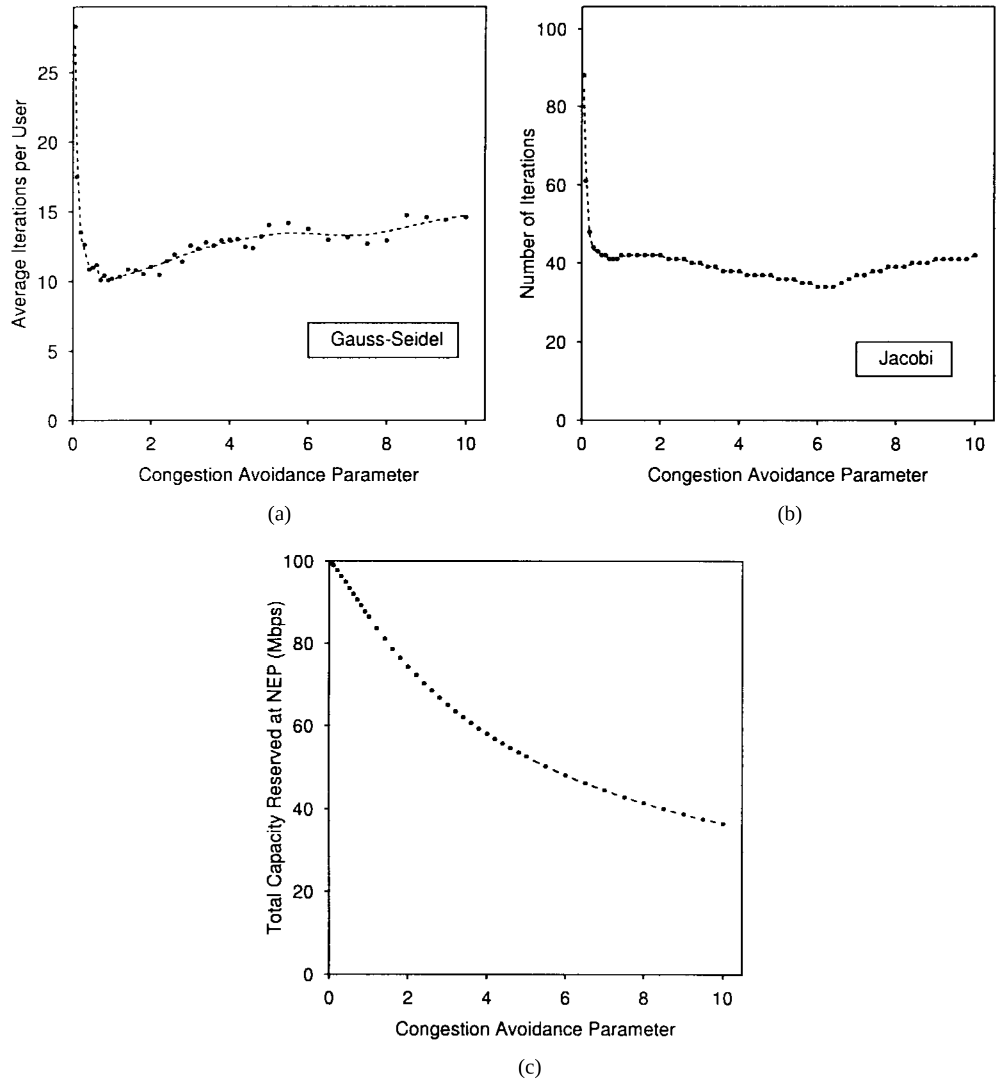

Fig. 2. Rate of convergence for (a) Gauss–Seidel, (b) Jacobi schemes, and (c) total capacity reserved at the NEP.

value of n increases beyond a certain value, the convergence speed starts to decrease again. This is due to the fact that as n increases, the total capacity at the NEP moves closer to the congestion region, leading, again, to a slow convergence rate. In Fig. 1 we depict this by showing the value of C^* that corresponds to each value of n by a small vertical line on the horizontal axis.

The above results point to a fundamental *tradeoff between capacity utilization and convergence rate*. Thus, the network manager is faced with the question of how to trade link capacity utilization with convergence speed. A balance between the two is achieved using an appropriate value of the congestion avoidance parameter n.

It is also observed that the Jacobi scheme exposes a significantly slower convergence rate than the Gauss–Seidel scheme. This is to be expected, since even the undamped Jacobi scheme is slower than the Gauss–Seidel for iterative solution of linear systems; in our case, the damping factor slows the Jacobi scheme further.

In our second experiment, we studied the dynamic behavior of the two iterative schemes in response to changes in user arrival statistics. We considered the setting of the previous experiment, initiated at its NEP, to which we added a new user with load $a_5 = 15$ Erl and zero initial capacity. Fig. 3 depicts the number of iterations that this system needed to reach the new NEP. As the figure shows, small values of n lead to a very slow convergence when the system operates under congestion. This means that, for small n, a user joining a highly utilized system experiences a long delay until it can reserve capacity equal to its NEP value.

Finally, we investigated the impact of the number of users on the convergence rate. In our last experiment, the congestion avoidance parameter was fixed to 1.0 and the aggregate call arrival rate of all users to 80.0 Erl. This rate was equally split among an increasing number of users. The results are shown in Fig. 4, where we observe that for both schemes the number of iterations increases as a constant total demand is split among an increasing number of (smaller) users. The Gauss–Seidel scheme appears to scale better with the number of users; this is due to the fact that the damping constant of the Jacobi scheme is equal to the number of users. It can also be observed that the rate of increase in the number of iterations per user in the

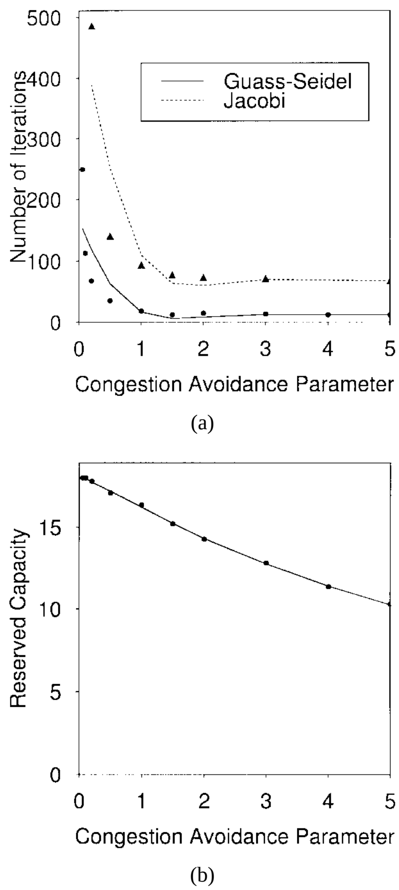

Fig. 3. (a) Rate of convergence and (b) C_5^* after fifth user added.

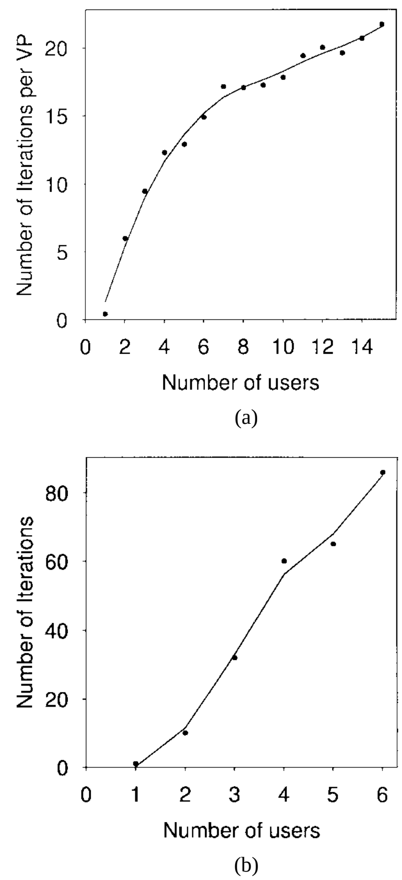

Fig. 4. Rate of convergence for (a) Gauss–Seidel and (b) Jacobi, as a function of the number of users.

Gauss–Seidel scheme decreases for larger number of users. This could be attributed to the fact that, as the number of users increases, the capacity at the NEP for each one decreases; therefore, the NEP gets progressively "closer" to the starting point, which is 0 for all users.

Our results highlight the differences between the two schemes. In particular, while the Gauss–Seidel scheme is shown to be significantly faster in convergence, its protocol structure is more complex since it involves a two-phase procedure for capacity reservation. It needs to implement locking, queueing of messages, and timeouts in order to avoid deadlocks, all of which are not needed for the Jacobi scheme. On the other hand, the speed of the damped Jacobi scheme decreases with the number of users, thus making it less appropriate for implementation in large-scale networks with a potentially large number of noncooperative users. In addition, in the Jacobi scheme, users need extra information to calculate their updates, namely the damping constant.

We plan to further explore these results, as well as to broaden their scope to more general network topologies.

VII. Conclusions

In this work we have investigated the virtual path capacity allocation problem in a noncooperative setting. In our formulation, users try to optimize individual cost functions that reflect the inherent tradeoff between user performance and the limited availability of resources. We have shown that the canonical problem has a unique NEP. We have also indicated that the unique NEP possesses desirable properties such as fairness. Furthermore, we investigated the dynamics of such a system and showed that convergence to the unique NEP is guaranteed under both the Gauss–Seidel and Jacobi iterative schemes. These convergence results are of particular importance from a practical standpoint, as they indicate how such a system behaves and stabilizes under dynamic conditions. We extended some of these results to more general settings. We indicated how to apply our insights to the case of multiple traffic classes when the capacity of the links is represented by schedulable regions [6]. A possible direction for further research is to consider competition among users *independently* on each dimension of the (multidimensional) schedulable region. We note that our contribution on convergence to the NEP can also be applied to the problem of noncooperative routing, thus enhancing the existence and uniqueness results presented in [16].

While the family of cost functions that we have examined in this paper is general enough to incorporate a considerable range of interesting functions, further investigation of other classes of functions and structures is of interest.

The formal model was emulated on an experimental setting. Apart from testing the behavior of our schemes in a practical scenario, these experiments indicate how the network manager can pilot the operating point and the dynamics of the system through a proper choice of the cost function parameters.

We believe that the properties demonstrated in this study, namely the existence, uniqueness, and convergence to the Nash equilibrium, support the possibility of effectively operating multiple users in virtual path-based networks.

APPENDIX

Proof of Lemma 3: Let user i be the one who updates its value of reserved capacity C_i at step $(n+1)$ and assume, w.l.o.g., that $C(n+1) \geq C^*$. This implies, by Lemma 2, that $C_i(n+1) \leq C_i^*$. (By symmetry, all arguments will apply for the reverse case as well.) Define

$$A(n) = \sum_{j=1}^{N} \left| C_j(n) - C_j^x \right|$$

and

$$B(n) = |C(n) - C^*|.$$

We distinguish among the following cases:

1) $C_i(n) > C_i^*$

Since $C_i(n+1) \leq C_i^*$, C_i "crosses" C_i^* at step $n+1$. Also, since $C_i(n+1) \leq C_i^* < C_i(n)$, we infer that $C(n) > C(n+1)$ $(\geq C^*)$. Hence

$$A(n+1) - A(n) = (C_i^* - C_i(n+1)) - (C_i(n) - C_i^*)$$

while

$$B(n+1) - B(n) = C_i(n+1) - C_i(n).$$

Thus,

$$\begin{aligned} &S(\mathbf{C}(n+1)) - S(\mathbf{C}(n)) \\ &\quad = (A(n+1) + B(n+1)) - (B(n) + A(n)) \\ &\quad = 2(C_i^* - C_i(n)) < 0 \end{aligned}$$

which implies that $S(\mathbf{C}(n+1)) < S(\mathbf{C}(n))$.

2) $C_i(n) \leq C_i^*$

Since $C_i(n+1) \leq C_i^*$, $C_i(n)$ and $C_i(n+1)$ are on the "same side" of C_i^*. We distinguish between the following two subcases:

a) $C_i(n+1) \leq C_i(n)$. Then, by Lemma 2, $C(n) \geq C(n+1)$ $(\geq C^*)$. Hence

$$\begin{aligned} &A(n+1) - A(n) \\ &\quad = |C_i(n+1) - C_i^*| - |C_i(n) - C_i^*| \\ &\quad = C_i(n) - C_i(n+1) \end{aligned}$$

and

$$\begin{aligned} &B(n+1) - B(n) \\ &\quad = |C(n+1) - C^*| - |C(n) - C^*| \\ &\quad = C(n+1) - C(n) = C_i(n+1) - C_i(n). \end{aligned}$$

Thus, $S(\mathbf{C}(n+1)) = S(\mathbf{C}(n))$.

b) $C_i(n+1) > C_i(n)$. Since $C_i(n+1) \leq C_i^*$, we have that $C_i(n) < C_i^*$. In this case, we cannot tell whether $C(n) \geq C^*$ or $C(n) < C^*$. Therefore, term $B(n)$ might either increase or decrease after step $(n+1)$. We can, however, give an upper bound on the increase of $B(n)$.

$$\begin{aligned} &A(n+1) - A(n) \\ &\quad = |C_i(n+1) - C_i^*| - |C_i(n) - C_i^*| \\ &\quad = C_i(n) - C_i(n+1) \end{aligned}$$

and

$$\begin{aligned} &B(n+1) - B(n) \\ &\quad = |C(n+1) - C^*| - |C(n) - C^*| \\ &\quad \leq C(n+1) - C(n) \\ &\quad = C_i(n+1) - C_i(n) \end{aligned}$$

which shows that $S(\mathbf{C}(n+1)) \leq S(\mathbf{C}(n))$. This concludes the proof. □

ACKNOWLEDGMENT

The authors would like to thank the reviewers and the Editor whose comments greatly improved the quality of this paper.

REFERENCES

[1] D. Bertsekas and R. Gallager, *Data Networks,* 2nd ed. Englewood Cliffs, NJ: Prentice-Hall, 1992.
[2] R. Cocchi, D. Estrin, S. Shenker, and L. Zhang, "Pricing in computer networks: Motivation, formulation, and example," *IEEE/ACM Trans. Networking*, vol. 1, pp. 614–627, Feb. 1993.
[3] A. A. Economides and J. A. Silvester, "Multi-objective routing in integrated services networks: A game theory approach," in *Proc. INFOCOM'91*, pp. 1220–1225, 1991.
[4] J. J. Garrahan, P. A. Russo, K. Kitami, and R. Kung, "Intelligent network overview," *IEEE Commun. Mag.*, vol. 31, no. 3, pp. 30–36, Mar. 1993.
[5] M.-T. T. Hsiao and A. A. Lazar, "Optimal decentralized flow control of Markovian queueing networks with multiple controllers," *Performance Evaluation*, vol. 13, no. 3, pp. 181–204, 1991.
[6] J. Hyman, A. A. Lazar, and G. Pacifici, "A methodology for virtual path and virtual network bandwidth assignment in broadband networks with QOS guarantees," in *Proc. IEEE Int. Conf. Communications,* New Orleans, LA, May 1994, pp. 165–169.
[7] Y. A. Korilis and A. A. Lazar, "On the existence of equilibria in noncooperative optimal flow control," *J. ACM,* vol. 42, no. 3, pp. 584–613, May 1995.
[8] Y. A. Korilis, A. A. Lazar, and A. Orda, "Architecting noncooperative networks," *IEEE J. Select. Areas Commun.*, vol. 13, no. 7, pp. 1241–1251, Sept. 1995.
[9] A. A. Lazar, A. Orda, and D. E. Pendarakis, "Virtual path bandwidth allocation in multiuser networks," Center Commun. Res., Columbia Univ., New York, NY, Tech. Rep. 425-95-31, rev., Sept. 1995. Available FTP: //ftp.ctr.columbia.edu/CTR-Research/comet/public/papers/95/LAZ95c.ps.gz
[10] D. G. Luenberger, *Linear and Nonlinear Programming,* 2nd ed. Reading, MA: Addison Wesley, 1984.
[11] R. Mazumdar, L. Mason, and C. Douligeris, "Fairness in network optimal flow control: Optimality of product forms," *IEEE Trans. Commun.*, vol. 39, pp. 775–781, May 1991.
[12] E. J. Messerli, "Proof of a convexity property of the Erlang B formula," *Bell Syst. Tech. J.*, vol. 51, no. 4, pp. 951–953, 1972.
[13] R. B. Myerson, *Game Theory: Analysis of Conflict.* Cambridge, MA: Harvard University Press, 1991.
[14] H. V. Norman, *LAN/WAN Optimization Techniques.* Norwood, MA: Artech House, 1992.
[15] S. Ohta and K.-I. Sato, "Dynamic bandwidth control of the virtual path in an asynchronous transfer node network," *IEEE Trans. Commun.*, vol. 40, July 1992.
[16] A. Orda, R. Rom, and N. Shimkin, "Competitive routing in multiuser communication networks," *IEEE/ACM Trans. Networking*, vol. 1, pp. 510–521, Oct. 1993.
[17] J. B. Rosen, "Existence and uniqueness of equilibrium points for concave n-person games," *Econometrica*, vol. 33, pp. 520–534, July 1965.

[18] H. Saito, K. Kawashima, and K.-I. Sato, "Traffic control technologies in ATM networks," *IEICE Trans.*, pt. E, vol. 74, Apr. 1991.

[19] K. Schulz, H. Kaufer, and B. Heilig, "The Deutsche bundespost telekom implements IN," *IEEE Commun. Mag.*, vol. 2, pp. 90–96, Feb. 1992.

[20] Z. Zhang and C. Douligeris, "Convergence of synchronous and asynchronous greedy algorithms in a multiclass telecommunications environment," *IEEE Trans. Commun.*, vol. 8, pp. 1277–1281, Aug. 1992.

Aurel A. Lazar (S'77–M'80–SM'90–F'93) is a professor of Electrical Engineering at Columbia University. 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 upon 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. In the early 1990's his research efforts shifted to the foundations of the control, management, and telemedia architecture of future multimedia networks. His involvement with gigabit networking research lead to the first fully operational service management system on ATM-based broadband networks. The system was implemented on top of AT&T's XUNET III gigabit platform spanning the continental United States. His management and control research pioneered the application of virtual reality to the management of ATM-based broadband networks. His current research in broadband networking with quality of service guarantees focuses on modeling of video streams and analyzing their multiplexing behavior, with emphasis on multiple time scales and subexponentiality. He is also leading investigations into multimedia 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 a broadband kernel (also known as **xbind**) that enables the rapid creation, deployment, and management of multimedia services. Resource allocation and networking games algorithms form the core of this architecture. He was instrumental in establishing the OPENSIG (http://www.ctr.columbia.edu/opensig) international working group with the goal of exploring network programmability and next-generation signalling technology.

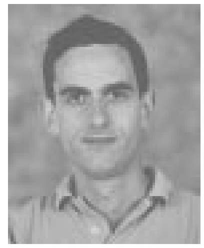

Ariel Orda (S'84–M'92–SM'97) received the B.Sc. (summa cum laude), M.Sc., and D.Sc. degrees in electrical engineering from the Technion—Israel Institute of Technology, Haifa, Israel, in 1983, 1985, and 1991, respectively.

Since 1994, he has been with the Department of Electrical Engineering, Technion, Haifa, Israel, as a Senior Lecturer. In 1993 and 1994, he was with the Center for Telecommunications Research, Columbia University, New York, NY, as a Visiting Scientist. In 1995 and 1996, he was with the IBM Thomas J. Watson Research Center, Hawthorne, NY, as an Academic Visitor. Since 1991, he has held consulting positions at Bell Laboratories, NJ, and in the Israeli Industry. His current research interests include the control and management of broadband networks, the application of game theory to computer networking, QoS routing, and online network algorithms.

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

Dimitrios E. Pendarakis (S'88–M'96) received the Diploma degree in electrical engineering from the National Technical University of Athens (NTUA), Athens, Greece, in 1990, and the M.S. and Ph.D. degrees in electrical engineering from Columbia University, New York, NY, in 1992 and 1996, respectively.

From 1991 to 1995 he was a Graduate Research Assistant with the NSF Center for Telecommunications Research, Columbia University, New York, NY, working on broadband network control algorithms. He spent the summer of 1992 at NEC C&C Research Laboratories, Princeton, NJ, investigating the modeling and analysis of video traffic, and the summers of 1993 and 1994 at AT&T Bell Laboratories, Murray Hill, NJ, where he developed tools for distributed monitoring of the XUNET ATM testbed. Since 1995 he has been a Research Staff Member with the IBM Thomas J. Watson Research Center, Yorktown Heights, NY. He holds an Adjunct Assistant Professor position at the Polytechnic University, NY. His current research interests are in control and management of multimedia networks, distributed algorithms, network security, and internet policy control.

Dr. Pendarakis served as the Technical Program Committee Chair for the OPENSIG Fall'97 Workshop.