Virtual Path Bandwidth Allocation in Multi-User Networks

Aurel A. Lazar*, Ariel Orda† and Dimitrios E. Pendarakis*

*Department of Electrical Engineering
and Center for Telecommunications Research
Columbia University
New York, NY 10027-6699, USA
E-mail: {aurel,dimitri}@ctr.columbia.edu

†Department of Electrical Engineering
Technion - Israel Institute of Technology
Haifa 32000, Israel
E-mail: ariel@obelix.technion.ac.il

Abstract

We consider a multi-user network that is shared by noncooperative users. Each user sets up virtual paths that optimize its own, selfish, performance measure. This measure accounts for both the guaranteed call level quality of service, as well as for the cost incurred for reserving the resource. The interaction between the user strategies is formalized as a game. We show that this 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.

1 Introduction

A trade-off encountered in communication networks is between bandwidth utilization and the complexity of call set-up. Consider for example, virtual circuit networks, in which resources are reserved upon call set up and released immediately upon call completion. This results in an optimal resource utilization but at the expense of a complex, per call, bandwidth management. This complexity might prove prohibitive in large networks, such as the emerging broadband ATM based networks which will support a variety of traffic classes and services. Indeed, recent studies of ATM testbeds have indicated that the performance of signalling and call processing entities can severely limit the ability of the network to accept calls. In such cases network users and operators will be willing to sacrifice resource utilization in order to reduce the set-up complexity. A common way to obtain that is offered by the virtual path approach [16]. 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 that the duration of a specific call.

An obvious implication of the virtual path scheme is that, typically, part of the reserved resources will be unused (hence the sacrifice in resource utilization). Yet another implication is that, at times, a user will experience lack of resources to accommodate all incoming calls through the virtual path, thus resulting in call blocking.

Therefore, the question arising in the virtual path approach is how the available capacity is to be distributed among users. 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 systemwide (single) administration is an impractical paradigm for the control of giant broadband networks. Indeed, the complexity of such systems calls for the distribution of the control tasks among various "local" controllers, each attempting 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]. Indeed, the application of game-theoretic tools to networking problems has been gaining increasing interest within the past few years [6, 4, 17, 14, 8, 3, 9]. A natural competition among users arises when they attempt to reserve network resources in order to establish their virtual paths.

Non cooperative games exhibit many times 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 multi-user networks, it calls for careful design and analysis. In the virtual path allocation problem, one possible scenario could be that each user attempts to reserve the maximum possible amount of resources for its own virtual path. Coupling this greedy behavior of users with the limited amount of resources 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 accounts for the limited availability of resources in the network.

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 [2]. With these results at hand, we proceed to discuss some implementation issues. In general, users may be sharing more that 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.

This paper is organized as follows. The model is formulated in Section 2. In Section 3 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 4, where we also consider some implementation details. Extensions to some more general configurations are presented in Section 5. Concluding remarks appear in Section 6. Due to space limitations, some of the proofs have been omitted from the conference proceedings. For a complete presentation of all the proofs, the reader may refer to [10].

2 Problem Formulation

We consider a set of N users, denoted by $\mathcal{I} = \{1, 2, \cdots, N\}$, that share a resource of total capacity of B units. We refer to this resource as to 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 on 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 5.

Each user reserves some of the resource capacity, in order to establish virtual paths 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 later case, the user faces the need to consume processing resources and waste time on call set-up. Thus, in either case, blocking at the virtual path level leads to performance degradation.

Users are noncooperative, meaning that each of them seeks to optimize its own performance. To that end, the user minimizes a cost function. This cost 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 for 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. 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 $\mathcal{C} = \{\mathbf{C} \mid \forall i \; 0 \leq C_i \leq B\}$. 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 i-th user, whereas the function G_i accounts for 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 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 (with respect to C_i).

G2 $G_i(C_i)$ is strictly decreasing (in C_i).

G3 $G_i(C_i)$ is convex (in C_i).

G4 $\lim_{C_i \to 0} G_i(C_i) = \infty$.

We shall denote by $F_i'(C_i, C) \triangleq \partial F_i(C_i, C)/\partial C_i$ 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] means that a user is not allowed to exhaust the resource. 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 dealt 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 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, $\hat{C}_i$, 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:

$$J_i(\mathbf{C}^*) = J_i(C_1^*, \ldots, C_{i-1}^*, C_i^*, C_{i+1}^*, \ldots, C_N^*) = \min_{\mathbf{C} \in \mathcal{C}} J_i(C_1^*, \ldots, C_{i-1}^*, C_i, C_{i+1}^*, \ldots, C_N^*) \qquad (2)$$

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]. In order to address this problem we introduce the following notation. 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 $\epsilon_i > 0$. The additional assumption is then defined as follows:

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

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 bandwidth available to user i is less than ϵ_i, user i will be charged *as if* the available bandwidth 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 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.

We note that the above formulation can accommodate 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 in computing the *schedulable region*[7]. 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 subdivided among the users according to the values of C_i, i.e., user i would get a portion of 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 types of calls.

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

Turning back to the general setting, we now illustrate it by the following example, motivated by [16]. Suppose that the function F_i is of the following form:

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

In other words, user i pays an amount of $1/(\max(B_i, \epsilon_i) - C_i)$ for each unit of reserved capacity. Turning now to the G_i function, we employ the well-known Erlang-B loss function [12]. To that end, we 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 the G_i function to be:

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

Note that the last assumption trivially holds for $\kappa_i \equiv 1$. 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 $G_i(C_i) < G_{i,min} < \infty$ (by assumption [E2]) and $F_i(C_i, C) \leq \infty$.

3 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, $J(C_i, C)$ 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 [14], Theorem 1 of [15] can be accommodated to guarantee that a Nash Equilibrium Point (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 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 established.

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

We proceed by proving an important property, namely, that at any NEP the total amount of reserved capacity C does not exceed the resource capacity B.

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

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 [11] then imply that a strategy vector $\mathbf{C} = (C_1, C_2, \ldots, C_N)$ is an NEP if and only if, for all users i, $1 \leq i \leq N$:

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

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

We proceed to prove 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. Nonetheless, for ease of exposition, we present the (much simpler) proof that applies to this case.

Theorem 1 *In a system of* N *noncooperative users sharing a resource of capacity* B*, as formalized in Section 2, the NEP is unique.*

Proof: Let $\mathbf{C} \in \mathcal{C}$ and $\hat{\mathbf{C}} \in \mathcal{C}$ be two NEPs. The optimality conditions (5) then imply:

$$\lambda_i = F_i'(C_i, C) = -G_i'(C_i)$$

and

$$\hat{\lambda}_i = F_i'(\hat{C}_i, \hat{C}) = -G_i'(\hat{C}_i)$$

Denote: $\lambda^+ = \{i \mid \hat{\lambda}_i > \lambda_i\}$ and $\lambda^- = \{i \mid \hat{\lambda}_i \leq \lambda_i\}$.

Suppose that $\hat{C} > C$. Due to the monotonicity of F_i' in each of its arguments, we have that, for $i \in \lambda^-$, $\hat{C}_i < C_i$. From the monotonicity of G_i' it follows that, for $i \in \lambda^-$, $G_i'(\hat{C}_i) < G_i'(C_i)$, which implies $\hat{\lambda}_i > \lambda_i$, contradicting $i \in \lambda^-$. Thus, $\lambda^- = \emptyset$, and $\hat{\lambda}_i > \lambda_i$ for all i. This means that $G_i'(\hat{C}_i) < G_i'(C_i)$ for all i, thus $\hat{C}_i < C_i$ for all i, hence $\hat{C} < C$, contradicting the assumption.

A symmetric argument contradicts the possibility $\hat{C} < C$, hence $\hat{C} = C$.

Consider a user i such that $\hat{\lambda}_i < \lambda_i$. Since $\hat{C} = C$, we have (from the monotonicity of F_i' and equation (5)) that $\hat{C}_i < C_i$. Thus, from the monotonicity of G_i', $G_i'(\hat{C}_i) < G_i'(C_i)$, hence (by (5)) $\hat{\lambda}_i > \lambda_i$, contradicting the assumption on i. A symmetric argument applies for the case $\hat{\lambda}_i > \lambda_i$, and therefore, for all i, $\hat{\lambda}_i = \lambda_i$. Since $\hat{C} = C$, (5) implies that $\hat{C}_i = C_i$ for all i, thus completing the proof. □

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 $i \in \mathcal{I}$, $F_i(C_i, C) = F(C_i, C)$. It should be observed that, even though users use the same cost functions, they still have different objectives since they use different arguments (i.e., C_i) in these functions. 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* $G_i'(\cdot) \leq G_j'(\cdot)$ *(correspondingly,* $G_i'(\cdot) < G_j'(\cdot)$*) for all (positive) values of* $\cdot$*.*

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

Definition 2 *An NEP* $C = (C_1, C_2, \cdots, C_L)$ *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* $C_i \geq C_j$ *(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: A particular case of Theorem 6. □

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 $i \in \mathcal{I}$.

4 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.

4.1 The Gauss-Seidel Scheme

In this subsection we consider a dynamic scheme in which users operate iteratively in an asynchronous fashion. More precisely, we assume that values of reserved capacities are recomputed in a sequence of steps. At each step n, $n = 1, 2, \ldots$, only one user updates the value of its reserved capacity based on information about *the total capacity reserved by all other users* up to the previous step, $n-1$. 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 of the same user is *finite*. Following common terminology, we will call this mode of operation the *asynchronous 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 asynchronous Gauss-Seidel scheme, the vector of capacities reserved by all users at step n, $\mathbf{C}(n) = (C_1(n), C_2(n), \ldots, C_N(n))$, converges to the unique NEP $(C_1^*, C_2^*, \ldots, C_N^*)$, as $n \to \infty$. If, in addition, we assume that the time interval between any two subsequent steps is *finite*, this implies that $\mathbf{C}$ 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:

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

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

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

$$C(n+1) \geq C^* \Rightarrow C_i(n+1) \leq C_i^*$$

(and, conversely)

$$C(n+1) \leq C^* \Rightarrow C_i(n+1) \geq C_i^*.$$

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

Corollary 1

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

We next prove that the sequence $S(n), n \in \mathbb{N}$ is non-increasing.

Lemma 3 *Under the asynchronous Gauss-Seidel scheme:*

$$S(n+1) \leq S(n)$$

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

We now show that $S(n)$ strictly decreases after a finite number of steps.

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

We are now ready to prove the convergence of the asynchronous Gauss-Seidel iteration to the unique NEP.

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

Proof: Observe that by Lemma 4, there is a subsequence of $S(n)$ that strictly decreases. Since $S(n) \geq 0$ for all n, we have that $lim_{n\to\infty} S(n) = 0$. This means that $lim_{n\to\infty} \mathbf{C}(n) = \mathbf{C}^*$. □

4.2 The Jacobi scheme

This subsection considers the convergence of a synchronous dynamic scheme. At each step n of the scheme, *all users* update simultaneously their value of reserved capacity, based on information on the capacity reserved by all users at the *previous*, $(n-1)$, step. We note that the assumption on simultaneous updating should not be interpreted as a strict timing synchronization requirement, but rather as the more relaxed assumption that, at each step, each user has available only information from the previous step. Using common terminology, we refer to this scheme as the *Jacobi scheme.*

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 iterative Jacobi scheme does converge to the unique NEP. This version is defined as follows: Denote by $\hat{C}_i(n+1)$ the capacity that user i would reserve at the nth step of the regular Jacobi scheme. Then, instead of 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 $\frac{1}{N}$ is a "damping" constant and N is the total number of users in the network.

As done with the Gauss-Seidel scheme, we define a metric $S(n)$, 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:

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

We note that $S(n)$ 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 their damped counterparts, calculated according to Equation (7). Also, recall that, at step n, *all* users update their capacity rather than just one.

We first prove the following:

Lemma 5 *If* $C(n) \geq C^*$ *and* $C_i(n) > C_i^*$, *then* $\hat{C}_i(n+1) < C_i(n)$.

Lemma 6 *If* $C(n) \geq C^*$ *and* $C_i(n) \leq C_i^*$, *then* $\hat{C}_i(n+1) \leq C_i^*$.

In the following, we assume, w.l.o.g., that $C(n) \geq C^*$, where n is the iteration that is considered. The complimentary case can be treated by symmetrical arguments. We note that, due to the above lemmata and the assumption $C(n) \geq C^*$, after the $(n+1)$'st iteration users can be classified into the following exhaustive sets, depending on the relation between $C_i(n)$, $\hat{C}_i(n+1)$ *and* C_i^*.

For $C_i(n) > C_i^*$, Lemma 5 yields the following two sets:

$$\mathcal{A}: \; C_i^* \leq \hat{C}_i(n+1) < C_i(n) \;,$$

$$\mathcal{B}: \; \hat{C}_i(n+1) \leq C_i^* < C_i(n) \;.$$

For $C_i(n) \leq C_i^*$, Lemma 6 yields the following sets:

$$\mathcal{C}: \hat{C}_i(n+1) \leq C_i(n) \leq C_i^* ,$$

$$\mathcal{D}: C_i(n) \leq \hat{C}_i(n+1) \leq C_i^* .$$

Observe that $\mathbf{C}(n) \neq \mathbf{C}^*$ and $C(n) \geq C^*$ implies $\mathcal{A} \cup \mathcal{B} \neq \emptyset$. Using the above sets, we will show that $S(n)$ decreases by considering the contribution of users in different sets separately, in a sense that will be made clear in the sequel. We begin with the following two lemmata.

Lemma 7 *For $C(n) \geq C^*$, if also $C(n+1) \geq C^*$ then $S(n+1) \leq S(n)$. Moreover, if $\mathcal{B} \neq \emptyset$, then $S(n+1) < S(n)$.*

Lemma 8 *Suppose that set $\mathcal{A}$ is empty, i.e., $i \in \mathcal{B} \cup \mathcal{C} \cup \mathcal{D}$ for all i. Then, for $C(n) \geq C^*$, we have that $C(n+1) \geq C^*$. Moreover, if $|\mathcal{B} \cup \mathcal{C}| < N$, then $C(n+1) > C^*$.*

We are now ready to show that $S(n)$ is descending:

Lemma 9 *If $\mathbf{C}(n) \neq \mathbf{C}^*$, then $S(n+1) < S(n)$.*

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

Proof: Follows from Lemma 9. □

4.3 Implementation Considerations

In this subsection we discuss implementation issues concerning the dynamic schemes presented previously. In practical scenaria, including recent ATM testbeds ([1]), link capacity is allocated by a link or switch controller, which is responsible for virtual circuit and virtual path set up and tear down. It is therefore reasonable to expect that for a specific link, all requests for obtaining or releasing capacity are to be handled by a common controller. We point out that this mode of operation differs from a truly centralized one in that all users operate *independently*, in a *non-cooperative* fashion; they only exchange information with the controller, regarding their capacity requests and its current availability (i.e., cost).

In this setting, the asynchronous iterative Gauss-Seidel scheme can be implemented by a "locking" mechanism in which, at any time, *at most one* user can modify its value of reserved capacity. A user wishing to update its value, requests the controller to provide it the current cost of a unit of capacity. If no other updating request is in progress, the controller responds and the updating can proceed; otherwise, the user will have to wait until the previous updating is completed. An obvious requirement in such a scenario is that users take finite time to calculate their optimal value of reserved capacity and complete an updating request. Numerous protocols implementing locking schemes have been considered in the literature. (See, for example [5]).

The modified Jacobi scheme can be implemented using a mechanism in which the controller "broadcasts" the current cost of capacity to all users and then waits for responses from users in the form of updating requests. The time interval between broadcasts could be constant or variable. In any case, responses to a broadcast that are received after the next broadcast are ignored. The convergence proof presented above for the modified Jacobi algorithm applies to this mechanism, since it can be verified that the proof still holds even when only a subset of the set of users update their values at a given step.

We observe that, in the way that both schemes have been defined, it is possible that, at some intermediate steps, the total reserved capacity exceeds the available amount B. While the schemes still guarantee convergence to the NEP (for which it has been shown that the total reserved capacity is less than B), in practice the controller should make sure that at all times the total reserved capacity is less than B. This could be done by separating the actual assignment of capacities from their computation. In other words, if, at step n, user i requests an amount that would make $C(n)$ be more than B, the controller allocates to that user a smaller amount, which is less than the current available amount in the resource (i.e., less than $B_i(n)$). It is easy to verify that such a user i must be such that $C_i(n-1) < \epsilon_i$ and also $C_i(n) < \epsilon_i$ (here, $C_i(n)$ denotes the amount *requested* by the user). In other words, such a user is currently requesting a small amount of capacity in any case. Also, at the NEP, all users are allocated the exact amounts that they request. While the controller may allocate amounts of capacity which are smaller than those requested by the users, the broadcasted values still correspond to the ones computed and requested.

5 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 take the form of a graph of general topology. Such extensions to our model are investigated in this section. We first concentrate on the simpler case of a parallel links topology. We then move to the case of general topologies, where resources (between source and destination) should be reserved on a multi-hop basis.

5.1 Parallel Links

We now 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 of which has a total capacity B^l. User i reserves an amount of capacity C_i^l on a link l, subject to the following constraints: $C_i^l \geq 0$ and $C_i^l \leq B^l$. 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. 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. 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 bandwidth 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:

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

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 Theorem 1 in [15] 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.

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.*

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 the case of a single link. All references to capacity values are to those at the NEP. In the following, we assume that, for each link l, all cost functions F_i^l are identical (for all users). We begin with the following auxiliary lemma, which demonstrates that, at the NEP, users reserve capacities at the various links in a "consistent" way.

Lemma 10 *Suppose that $C_i^{\hat{l}} > C_j^{\hat{l}}$ holds for some link $\hat{l}$ and users i and j. Then $C_i^l \geq C_j^l$ for all $l \in \mathcal{L}$; moreover, the last inequality is strict if $C_j^l > 0$.*

The fairness of the NEP is now established:

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.*

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 $\mathcal{L}_1, \mathcal{L}_2, ..., \mathcal{L}_N$, 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 $n, n+1, ..., N$. 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).

5.2 General Topologies

5.2.1 Model

We consider now a network $\mathcal{G}(\mathcal{V},\mathcal{L})$, where $\mathcal{V}$ is a finite set of nodes and $\mathcal{L} \subseteq \mathcal{V}\times\mathcal{V}$ is a set of directed links, such that each link $l \in \mathcal{L}$ has a bandwidth B^l. In other words, the resources form a general topology, and capacities should be reserved on a multi-hop basis. For simplicity of notation, we assume that at most one link exists between each pair of nodes (in each direction). This assumption involves no loss of generality, since any network can be reduced to this form by introducing fictitious nodes. For a link $l \in \mathcal{L}$, we denote by $S(l)$ the identity of the node at the starting point of l and $D(l)$ is the node at the ending point. Considering now a node $v \in \mathcal{V}$, we denote by $In(v)$ (correspondingly, $Out(v)$) the set of node v's in-going links (correspondingly, out-going links) i.e.: $In(v) = \{l|D(l) = v\}$ ($Out(v) = \{l|S(l) = v\}$).

As before, we are given a set $\mathcal{I} = \{1, 2, \ldots, N\}$ of users, which now share the bandwidth of the network $\mathcal{G}$. Users reserve capacity for calls to be set between a source node s and a destination node d. A user i reserves C_i^{uv} on a link (u, v), subject to the following:

- $C_i^{uv} \geq 0$,
- $C_i^{uv} \leq B^{uv}$,
- $\forall u \neq s, d$: $\sum_{v\in Out(u)} C_i^{vu} = \sum_{v\in In(u)} C_i^{uv}$, i.e., the reserved capacity outgoing from a tandem node must be equal to that ingoing to that node (note that an excess in any direction can be of no benefit).

As before, F_i^{uv} accounts for the cost incurred on i for reserving capacity on link (u, v); F_i^{uv} is a function of two arguments, namely the capacity reserved by i on link (u, v) (i.e., C_i^{uv}) and the total capacity reserved there by all users (C^{uv}). The assumptions made on these functions are identical to those made in the case of parallel links, except for the following. For parallel links, we assumed that, for each user i, there must be at least one link n for which $\epsilon_i^n > 0$; we then defined $\epsilon_i = \sum_l \epsilon_i^l > 0$. For general topologies, these assumption and definition are formulated differently. For each user i, we require that there is at least one path p_i between the source and the destination, such that, on all links $l \in p_i$, $\epsilon_i^l > 0$. Denote by ϵ_i the maximal "flow" that can be obtained in the network $\mathcal{G}$, when the values of ϵ_i^l are assumed to be the "capacities" of the links l. Since ϵ_i is no less than the minimal ϵ_i^l along any path from the source to the destination, the above assumption (on the existence of a path p_i) implies that $\epsilon_i > 0$ for all i. The reader can check that, with this formulation, all assumptions, and in particular [PE2], are now applicable to general topologies.

A user can assign an incoming call to any path (or combination of paths) between the source s and destination d, on which the amount of reserved and unused capacity can accommodate that call. As a consequence, the cost function G_i takes as its argument the total capacity C_i reserved by user i. Due to the nodal conservation of capacities, it is easy to verify that this total capacity equals to the sum of capacities reserved on links outgoing from the source node s (equivalently: the sum of capacities reserved on links ingoing into the destination node d). Thus, the function G_i (of a user i) takes as its argument the total capacity reserved for user i at the source node, i.e.: $G_i = G_i(C_i)$, where $C_i = \sum_v C_i^{sv}$. The assumptions made on G_i are the same as before.

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

$$J_i = \sum_{(u,v)} F_i^{uv}(C_i^{uv}, C^{uv}) + G_i(C_i)$$

The existence of the NEP follows from Theorem 1 in [15] 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 for Lemma 1. Nonetheless, in general the NEP may not be unique, unless more specific assumptions are made. In the following we discuss some cases for which uniqueness of the NEP can be shown.

5.2.2 Identical Users

Suppose that all users have the same loss sensitivity (following the terminology of the previous subsection) and they use the same F_i^l functions. We call such users *identical*. Note that, for identical users, the derivatives of the various cost functions are identical, i.e.: for all $i, j \in \mathcal{I}$ we have $F_i^{l'}(\cdot,\cdot) = F_j^{l'}(\cdot,\cdot)$ and $G_i'(\cdot) = G_j'(\cdot)$.

Lemma 11 *In a network with identical users, the capacity values at an NEP are such that*

$$C_i^l \equiv \frac{C^l}{N}$$

for all $i \in \mathcal{I}$ and $l \in \mathcal{L}$.

Proof: Similar to that of Lemma 4.1 in [14]. □

Theorem 7 *A network with identical users has a unique NEP.*

5.2.3 All-Positive Capacities

In this subsection we assume that the F cost functions are of the form

$$F_i^l(C_i^l, C^l) = C_i^l \cdot P^l(C^l)$$

where P^l can be thought to be a global price per unit of capacity in link l (homogeneous for all users), which depends on the state of the link (i.e., on its congestion). We call such link cost functions *homogeneous.*

Theorem 8 *Consider a system of N noncooperative users sharing a network of general topology, as formalized above, and assume that link cost functions are homogeneous. Let $\mathbf{C}$ and $\hat{\mathbf{C}}$ be two NEPs such that, in both, all users reserve capacities in the same set of links, i.e., there exists a set of links $\mathcal{L}_1 \subset \mathcal{L}$ for which $\{C_i^l > 0$ and $\hat{C}_i^l > 0,\ i \in \mathcal{I}\}$ for $l \in \mathcal{L}_1$, and $\{C_i^l = \hat{C}_i^l = 0,\ i \in \mathcal{I}\}$ for $l \notin \mathcal{L}_1$. Then $\mathbf{C} = \hat{\mathbf{C}}$.*

Proof: Similar to that of Theorem 3.3 in [14]. □

The above result implies that at most one NEP exists with the property that all users reserve strictly positive capacities on each link of the network. This may be the case in heavy loaded networks, where users typically exhaust all available resources.

6 Conclusions

In this work we have investigated the virtual path capacity allocation problem in a non-cooperative setting. In our formulation, users try to optimize individual cost functions which manifest the inherent trade-off between user performance and the limited availability of resources. We have shown that the basic problem has a unique Nash equilibrium point. We also have indicated that this 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. We note that, in general, proving convergence to a Nash equilibrium is very hard, if not impossible. These convergence results are of particular importance from a practical standpoint, since they indicate how such a system behaves and stabilizes.

We extended some of these results to apply to more general settings in which users share resources in more complex settings. As previously mentioned, our work forms part of a currently emerging effort to investigate non-cooperation in networks. We indicated how to apply our setting to the case of multiple traffic classes, in terms of *schedulable regions*[7]. 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 be applied also to the problem of noncooperative routing, thus enhancing the existence and uniqueness results present in [14].

We believe that the properties demonstrated in this study, namely the existence, uniqueness and convergence to the Nash equilibrium, give encouraging support to the possibility of effectively operating multiple users in giant, virtual-path based, networks under a non-cooperative paradigm.

References

[1] Nikos G. Aneroussis, Aurel A. Lazar, and Dimitrios E. Pendarakis. Taming XUNET III, 1994. Submitted for publication.

[2] Dimitri Bertsekas and Robert Gallager. *Data Networks*. Prentice-Hall, Englewood Cliffs, NJ, 2nd edition, 1992.

[3] R. Cocchi, D. Estrin, S. Shenker, and L. Zhang. Pricing in Computer Networks: Motivation, Formulation, and Example. *IEEE/ACM Transactions on Networking*, 1:614–627, 1993.

[4] A. A. Economides and J. A. Silvester. Multi-objective routing in integrated services networks: A game theory approach. In *Proceedings of INFOCOM'91*, pages 1220–1225, 1991.

[5] Ramez Elmasri and Shamkant B. Navathe. *Fundamentals of Database Systems*. The Benjamin/Cummings Publishing Company, Inc, 1989.

[6] Man-Tung T. Hsiao and Aurel A. Lazar. Optimal Decentralized Flow Control of Markovian Queueing Networks with Multiple Controllers. *Performance Evaluation*, 13(3):181–204, 1991.

[7] Jay Hyman, Aurel A. Lazar, and Giovanni Pacifici. A methodology for virtual path and virtual network bandwidth assignment in broadband networks with QOS guarantees. In *Proceedings of the IEEE International Conference on Communications*, pages 165–169, New Orleans, LA, May 1994.

[8] Yannis A. Korilis and Aurel A. Lazar. On the Existence of Equilibria in Noncooperative Optimal Flow Control. In *Proceedings of the ITC Sponsored Seminar*, pages 243–251, Bangalore, India, November 1993. Indian Institute of Science. Under revision for the *Journal of the ACM*.

[9] Yiannis A. Korilis, Aurel A. Lazar, and Ariel Orda. The designer's perspective to noncooperative networks. In *Proceedings of the IEEE INFOCOM '95*, Boston, MA, April 1995.

[10] Aurel A. Lazar, Ariel Orda, and Dimitrios E. Pendarakis. Virtual path bandwidth allocation in multi-user networks. CTR Technical Report 368-94-15, Center for Telecommunications Research, Columbia University, New York, July 1994. URL:"ftp://ftp.ctr.columbia.edu/CTR-Research/comet/public/papers/94/LAZ94d.ps.gz".

[11] David G. Luenberger. *Linear and Nonlinear Programming*. Addison Wesley, 2nd edition, 1984.

[12] E. J. Messerli. Proof of a convexity property of the Erlang B formula. *The Bell System technical journal*, 51(4):951–953, 1972.

[13] Roger B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, Massachusetts, 1991.

[14] Ariel Orda, Raphael Rom, and Nahum Shimkin. Competitive Routing in Multiuser Communication Networks. *IEEE/ACM Transactions on Networking*, 1(5):510–521, October 1993.

[15] J. B. Rosen. Existence and Uniqueness of Equilibrium Points for Concave n-Person Games. *Econometrica*, 33(3):520–534, July 1965.

[16] Hiroshi Saito, Konosuke Kawashima, and Kenichi Sato. Traffic control technologies in ATM networks. *IEICE Transactions*, E 74(4), April 1991.

[17] Zhensheng Zhang and Christos Douligeris. Convergence of Synchronous and Asynchronous Greedy Algorithms in a Multiclass Telecommunications Environment. *IEEE Transactions on Communications*, 40(8):1277–1281, August 1992.