Multi-user Rate-Based Flow Control

Eitan Altman, *Member, IEEE*, and **Tamer Başar**, *Fellow, IEEE*

Abstract— **Flow and congestion control allow the users of a telecommunication network to regulate the traffic they send into the network in accordance with the quality of service they require. Flow control may be performed by the network, as is the case in ATM networks (the Available Bit Rate transfer capacity), or by the users themselves, as is the case in the Internet (TCP/IP). We study in this paper both situations using optimal control and dynamic game techniques. The first situation leads to the formulation of a dynamic team problem, while the second one leads to a dynamic noncooperative game, for which we establish the existence and uniqueness of a linear Nash equilibrium, and obtain a characterization of the corresponding equilibrium policies along with the performance costs. We further show that when the users update their policies in a greedy manner, not knowing *a priori* the utilities of the other players, the sequence of policies thus generated converges to the Nash equilibrium. Finally, we study an extension of the model that accommodates multiple traffic types for each user, with the switching from one type of traffic to another being governed by a Markov jump process. Presentation of some numerical results complements this study.**

Keywords: Multi-user rate-based flow control; High speed networks; Linear-quadratic control; Linear-quadratic differential games; Nash equilibria.

I. Introduction

We consider M users that share a common bottleneck queue in a telecommunications network. The input flow of information from the users is controlled so as to achieve the best quality of service. As is the case in many proposed flow-control schemes [2], [16], we assume that there is some target value of the queue length which the users try to track; this value and the control policies are chosen on the one hand to avoid the buffer to be full (in order to minimize losses), and on the other, to avoid the queue to empty, which would lead to loss in the potential throughput. A second objective of each user is related to the input rate: we assume that some fraction of the available bandwidth is allocated to each user, and the user tries to minimize deviation from this allocated bandwidth.

Flow control is typically performed *dynamically*: some feedback information on the congestion or on round trip delay is used to update the input rate. For example, in the Tahoe version of TCP/IP congestion control [19], congestion is detected through losses or through time-out mechanisms. In the Vegas version of TCP/IP [11], the available bandwidth is also used as feedback information (it is obtained through estimation of round-trip delays). More detailed queue length information may also be available as feedback information. In the Available Bit Rate (ABR) transfer capacity of ATM, both queue length information as well as information on the rate may be conveyed from the switches to the sources through special information cells that are called RM (Resource Management) cells.

Flow control is often performed in a decentralized way in telecommunication networks: each user controls its own flow. This is the case in the Internet, see [19], or in some best-effort type traffic in ATM (the Unspecified Bit Rate transfer capacity, see [1]). This gives rise to a noncooperative dynamic game: each user has its own objectives, but the actions of the different users influence the quality of service of other users. Controllers that have been implemented in large scale, such as the TCP/IP, have heretofore typically been designed using heuristic techniques based on growing experience (and on simulation studies) and have not involved a game-theoretic analysis.

On the other hand, there has been some work on the use of noncooperative game theoretical techniques to design simple flow controls. The problem of choosing fixed (non-state dependent) rates has been investigated in [10], [12], [14], [15], [26], [30] and some of the references therein; in this context the existence and uniqueness of Nash equilibrium have been established and convergence of synchronous and asynchronous implementations of the flow control have been obtained. For state-dependent flow control, on the other hand, only structural results on the Nash equilibria have been established [17], [22]. Another type of a non-

Paper approved by A. Pattavina, the Editor for Switching Architecture Performance of the IEEE Communications Society. Manuscript received November 6, 1996; revised October 8, 1997. A shorter version was presented at the *36th IEEE Conference on Decision and Control*, San Diego, California, Dec. 1997.

Eitan Altman is with INRIA B.P. 93, 06902 Sophia Antipolis Cedex, France, altman@sophia.inria.fr. His work was supported by the CNET through research grant CTI 97 1B 206.

Tamer Başar is with the Coordinated Science Laboratory, University of Illinois, 1308 West Main Street, Urbana, IL 61801/USA. Tel: (217) 3333607; Fax: (217) 2441653. His research was supported in part by NSF Grants ECS 92-20632, ECS 93-12807, and in part by the AFOSR MURI Grant AF-DC-5-36128.

cooperative flow control occurs when a central controller seeks a globally optimal performance measure which is a function of utilities declared by the users. The true utilities of the users are, however, private information, and the declared utilities may not coincide with the true ones. In that case, the game aspect of the problem arises in the revealing policy of the utilities. The design of a (non-state dependent) flow control that induces a truthful revealing policy for all users has been obtained in [13] and some of the references therein. For other references on networking games, see [23], [24], [25].

In this paper we study *state-dependent* decentralized control and game problems. For both the noncooperative problem (that occurs when the controls are at end-points) as well as the team problem (occurring when the controllers are at the switches), we obtain explicitly the unique optimal (or equilibrium) controllers and the associated values. We then study asynchronous and synchronous update algorithms that the users might implement in order to compute their policies *on line*, since in practice, a user may not have access to full information on the utilities of other users, and thus, may not be able to construct its own Nash equilibrium solution *off line*. Instead, it is natural to assume that such a user would follow a "greedy approach" of optimizing from time to time its response against the current policies of other users. We present four such algorithms, establish convergence properties for two of them, and present numerical results illustrating convergence as well as its rate for three of the algorithms.

The balance of the paper has been organized as follows. In the next section we introduce and motivate the basic mathematical model, and obtain characterizations of equilibrium policies under both team and noncooperative game scenarios. We then introduce, in Section III, four algorithms for on line computation of Nash equilibria, and present some convergence analysis. Section IV deals with an extension of the basic model to the case where users have different traffic types, with switching from one type to another governed by a Markov jump process. For this case, we obtain the linear Nash equilibrium. Section V is devoted to presentation of some numerical results, and Section VI to a discussion on the asymptotic behavior of various solutions as the number of users grow, and comparisons between the Nash and the team solutions. The paper ends with the concluding remarks of Section VII.

II. The model and basic existence and uniqueness results

The basic model we work with in this paper is based on three simplifying (but realistic) assumptions:

1. *Fluid approximation.* We replace a discrete number of packets by a continuous fluid. The fluid approximation is justified by the fact that in today's technology, buffering capabilities are very large (several thousands) in terms of the number of packets that they can store, so that the error of replacing an integer number of packets by a real number is small relative to the size of buffers. This type of approximation is common both in the design of controllers in high-speed telecommunication networks (see e.g. [8]) as well as in performance evaluation of existing controllers [27], [29].

2. *Linearized dynamics.* It is assumed that the network has linearized dynamics for the control of queue length; see (1) below. This means in particular that we neglect losses when the buffer is full, and we neglect the boundary effect of an empty queue. To motivate this linearization, we use the fact that the controllers that we derive operate in a region close to $\rho = 1$ (full throughput utilization), so that the queue will almost never empty. Operating at $\rho = 1$ is common in the control of ABR switches, see e.g., [20]. The full utilization ($\rho = 1$) is possible by *regulating* the (controlled) input rate and adapting it to the available capacity. As discussed below, we set some desirable threshold on the queue length which we attempt to track, precisely so as to avoid large queues (which might result in losses) or empty queues (which might result in loss of potential throughput). When a control mechanism has a full utilization, then the nonlinearity around zero disappears. For similar models with a single controller, simulations have confirmed [3], [2] that *controlled* linearized models lead to trajectories that are very close to the original one. The fact that the other boundary is ignored is motivated by similar arguments, since our optimal control will be shown to be symmetric with respect to positive or negative deviations around the target queue value.

3. *Bottleneck assumption.* We assume that all performance measures (such as throughput, delays, loss probabilities, etc.) are determined essentially by a bottleneck node. This assumption admits theoretical as well as experimental justifications; see [7].

We now introduce the model we adopt in this paper. Let $q(t)$ denote the queue length at a bottleneck link, and $s(t)$ denote the total effective service rate available at that link. Assume that each user is assigned a fix proportion of the available bandwidth; more specifically, the traffic of source m at that link at time t has an available bandwidth of $a_m s(t)$. We assume throughout that $\sum_{m=1}^{M} a_m = 1$. We let $s(t)$ be arbitrary, but assume that the controllers have perfect measurements of it. Let $r_m(t)$ denote the (controlled) rate of source m at time t, $m \in \mathcal{M} := \{1, \ldots, M\}$, and $u_m(t) := r_m(t) - a_m s(t)$ be its shifted version. Consider the following dynamics for the queue length:

$$\frac{dq}{dt} = \sum_{m=1}^{M} (r_m - a_m s) = \sum_{m=1}^{M} u_m, \tag{1}$$

which is *idealized* because the end-point effects have been ignored. The objectives of the flow controllers are *(i)* to ensure that the bottleneck queue size stays around some desired level $\overline{Q}$, and *(ii)* to achieve good tracking between input and output rates. In particular, the choice of $\overline{Q}$ and the variability around it have a direct impact on loss probabilities and throughput. We therefore define a shifted version of q: $x(t) := q(t) - \overline{Q}$, in view of which (1) now becomes

$$\frac{dx}{dt} = \sum_{m=1}^{M} u_m. \tag{2}$$

An appropriate local cost function that is compatible with the objectives stated above would be the one that penalizes variations in $x(t)$ and $u_m(t)$ around *zero* — a candidate for which is the weighted quadratic cost function. Associated with user m is a positive constant c_m appearing in its immediate cost, as described below.

We shall first consider two noncooperative scenarios, formulated as noncooperative differential games, in which each user minimizes its own individual cost function. Then we shall study two cooperative scenarios, formulated as team control problems, in which all users have the same (common) objective.

We fix an initial state $x(0)$, and assume that the actions of controller m ($m \in \mathcal{M}$) are determined by a control policy $\mu_m \in \mathcal{U}_m$, where $u_m(t) = \mu_m(t, x_{[0,t]})$, $t \in [0, \infty)$. Here, μ_m is taken to be piecewise continuous in its first argument, and piecewise Lipschitz continuous in its second argument. We denote the class of all such policies for user m by $\mathcal{U}_m$. It will soon turn out that it will be sufficient to restrict attention to a subclass of $\mathcal{U}_m$, comprising policies that are linear in the current value of x; by an abuse of notation, we will denote such policies also by μ_m, i.e., $u_m(t) = \mu_m(x(t))$.

The two noncooperative cases are defined as follows, where $c_m > 0, m \in \mathcal{M}$:

N1: the individual cost to be minimized by controller m ($m \in \mathcal{M}$) is

$$J_m^{N1}(u) = \int_0^\infty \left(|x(t)|^2 + \frac{1}{c_m} |u_m(t)|^2 \right) dt. \tag{3}$$

N2: the individual cost to be minimized by controller m ($m \in \mathcal{M}$) is

$$J_m^{N2}(u) = \int_0^\infty \left(\frac{1}{M} |x(t)|^2 + \frac{1}{c_m} |u_m(t)|^2 \right) dt. \tag{4}$$

Note that in case N2 the "effort" for keeping the deviations of the queue length from the desired value is split equally between the users.

The solution concept we adopt for N1 and N2 is the Nash equilibrium: We seek a multi-policy $\mu^* := (\mu_1^*, ..., \mu_M^*)$ that no user has an incentive to deviate from, i.e.,

$$J_m^{N1}(\mu^*) = \inf_{\mu_m \in \mathcal{U}_m} J_m^{N1}([\mu_m | \mu_{-m}^*]) \tag{5}$$

where $[\mu_m | \mu_{-m}^*]$ is the policy obtained when for each $j \neq m$, player j uses policy μ_j^*, and player m uses μ_m. The definition of a Nash equilibrium for N2 is similar.

We next introduce two related team problems:

T1: the common (team) cost to be minimized is

$$J^{T1}(u) = \int_0^\infty \left(M |x(t)|^2 + \sum_{m=1}^{M} \frac{1}{c_m} |u_m(t)|^2 \right) dt. \tag{6}$$

T2: the common (team) cost to be minimized is

$$J^{T2}(u) := \int_0^\infty \left(|x(t)|^2 + \sum_{m=1}^{M} \frac{1}{c_m} |u_m(t)|^2 \right) dt. \tag{7}$$

We first have the following result, for N1 and N2:

Theorem 1: For the noncooperative differential game Ni $(i = 1, 2)$, there exists a state-dependent Nash equilibrium policy given by

$$\mu_{Ni,m}^*(x) = -\beta_m^{Ni} x, \; m \in \mathcal{M}, \text{ where} \tag{8}$$

$$\beta_m^{N1} = \overline{\beta}^{(N1)} - \sqrt{\left[\overline{\beta}^{(N1)}\right]^2 - c_m} \text{ and}$$
$$\beta_m^{N2} = \overline{\beta}^{(N2)} - \sqrt{\left[\overline{\beta}^{(N2)}\right]^2 - \frac{c_m}{M}}, \tag{9}$$

for cases N1 and N2, respectively, where

$$\overline{\beta}^{(Ni)} := \sum_{m=1}^{M} \beta_m^{Ni}, \;\; i = 1, 2,$$

are the unique solutions of

$$\overline{\beta}^{(N1)} = \frac{\sum_{m=1}^{M} \sqrt{\left[\overline{\beta}^{(N1)}\right]^2 - c_m}}{M-1} \text{ and}$$

$$\overline{\beta}^{(N2)} = \frac{\sum_{m=1}^{M} \sqrt{\left[\overline{\beta}^{(N2)}\right]^2 - \frac{c_m}{M}}}{M-1} = \frac{\overline{\beta}^{(N1)}}{\sqrt{M}}. \qquad (10)$$

Moreover,

$$\beta_m^{N1} = \beta_m^{N2}\sqrt{M}. \qquad (11)$$

The costs accruing to user m under the two Nash equilibria above, when the initial state is x, are given by

$$J_m^{N1}(\mu_{N1}^*) = \frac{\beta_m^{N1}}{c_m}x^2$$

and

$$J_m^{N2}(\mu_{N2}^*) = \frac{\beta_m^{N2}}{c_m}x^2 = \frac{1}{\sqrt{M}}J_m^{N1}(\mu_{N1}^*).$$

For each case, (8) is the unique equilibrium solution among linear memoryless stationary policies and is time-consistent. In particular,
(i) for the symmetric case $c_m = c_j =: c$ for all $m, j \in \mathcal{M}$, we have $\forall m \in \mathcal{M}$,

$$\beta_m^{N1} = \sqrt{\frac{c}{2M-1}}, \beta_m^{N2} = \sqrt{\frac{c}{M(2M-1)}}, \qquad (12)$$

(ii) in the case of $M = 2$, with general c_m's, we have for $m = 1, 2,\ j \neq m,\ j = 1, 2,$

$$\beta_m^{N1} = \left[-\frac{2c_j - c_m}{3} + 2\frac{\sqrt{c_1^2 - c_1c_2 + c_2^2}}{3}\right]^{1/2}, \beta_m^{N2} = \frac{\beta_m^{N1}}{\sqrt{2}},$$

and if moreover, $c_1 = c_2 = c$, then $\beta_m^{N1} = \sqrt{c/3}$, $\beta_m^{N2} = \sqrt{c/6}$.

Proof: We prove the result only for $N1$; its counterpart $N2$ can be proven similarly. Also for ease of notation, we drop all superscripts pertaining to the case considered.

We first choose a candidate solution of the form (8) for each player, and consider the optimal response of an arbitrary player m to the fixed policy $\mu_j,\ j \neq m$ of the other players. Let $\beta_{-m} := \sum_{j\neq m}\beta_j$. Player m is then faced with a linear-quadratic optimal control problem with the dynamics $dx/dt = u_m - \beta_{-m}x$, and cost $J_m^{N1}(u)$ that is strictly convex in u_m. By a standard result in optimal control [5], there exists a unique optimal response for player m, of the form $u_m = -c_mP_mx$, where P_m is the unique positive solution of the quadratic equation

$$-2\beta_{-m}P_m - P_m^2c_m + 1 = 0. \qquad (13)$$

The optimal cost to player m is then $J_m^{N1} = P_mx^2$. Introducing $\beta'_m := c_mP_m$, we obtain from (13)

$$\beta'_m = f_m(\beta_{-m}) := -\beta_{-m} + \sqrt{\beta_{-m}^2 + c_m}. \qquad (14)$$

A necessary and sufficient condition for μ to be in equilibrium is then that $\beta'_m = \beta_m$, or equivalently, that $\overline{\beta}^2 = \beta_{-m}^2 + c_m$. This yields the expressions (9)-(10) ((10) is obtained by summing (9) over $m \in \mathcal{M}$). The fact that (10) admits a unique solution follows from the fact that the left-hand side minus the right-hand side of (10) is strictly decreasing in $\overline{\beta}$ over the interval $[\max_m \sqrt{c_m}, \infty)$, it is positive at $\overline{\beta} = \max_m \sqrt{c_m}$ and it tends to $-\infty$ as $\beta \to \infty$. (i) and (ii) are obtained by solving for $\overline{\beta}$ from (10) for these various cases. ◇

Remark 1: The function f_m defined in (14) has the interpretation of the optimal response of user m to the rates of all other users. As seen here, the actions of all other users determine this response only through their sum. This feature appears quite frequently in the literature on networking games; see [10], [14], [15], [17], [22], [28]. In practice, the users need not exchange information on their β_i's; a possible way for a user to estimate β_{-m} is to infer it from the queue length information. Indeed, we have by substitution into (1), $dq/dt = -\overline{\beta}(q - \overline{Q})$. A possible way to estimate the queue length is by estimating the round trip delay, as was done in [11]. ◇

Corollary 1: Let

$$\max_m c_m << \overline{c} := \sum_{m=1}^{M} c_m. \qquad (15)$$

Then, $\beta_m^{Ni}, i = 1, 2$, in (9) can be approximated by

$$\beta_m^{N1} \sim c_m/\sqrt{2\overline{c}},\ \beta_m^{N2} \sim c_m/\sqrt{2M\overline{c}},$$

and the Nash equilibrium costs per user m are approximated by

$$J_m^{N1}(\mu_{N1}^*) \sim x^2/\sqrt{2\overline{c}},\ J_m^{N2}(\mu_{N2}^*) \sim x^2/\sqrt{2M\overline{c}},$$

which are independent of m.

Proof: We again consider only N1; the result for N2 then follows from (11). Furthermore, as before we drop the superscripts identifying the two cases. Assume that (15) holds. Then, (9) can be written as:

$$\beta_m = \overline{\beta}\left[1 - \sqrt{1 - \frac{c_m}{\overline{\beta}^2}}\right] = \frac{c_m}{2\overline{\beta}}\left[1 + O\left(\frac{c_m}{\overline{\beta}^2}\right)\right], \quad (16)$$

where $O(\cdot)$ is a function that satisfies $\lim_{x\to 0} O(x) = 0$. The right hand side of the expression for $\overline{\beta}^{N1}$ in (10) can be written as:

$$\begin{aligned} &\frac{\overline{\beta}}{M-1}\sum_{m=1}^{M}\sqrt{1 - \frac{c_m}{\overline{\beta}^2}} \\ &= \frac{\overline{\beta}}{M-1}\sum_{m=1}^{M}\left[1 - \frac{c_m}{2\overline{\beta}^2}\left(1 + O\left(\frac{c_m}{\overline{\beta}^2}\right)\right)\right] \\ &= \frac{\overline{\beta}}{M-1}\left[M - \frac{\overline{c}}{2\overline{\beta}^2}\left(1 + O\left(\frac{\max_m c_m}{\overline{\beta}^2}\right)\right)\right]. \end{aligned}$$

Substituting this in (10) yields

$$\overline{\beta} = \sqrt{\overline{c}/2}\left[1 + O\left(\max_m c_m / 2\overline{\beta}^2\right)\right]. \quad (17)$$

A solution of (17) (and thus of (10)) is

$$\overline{\beta} = \sqrt{\overline{c}/2}\left[1 + O\left(\max_m c_m / \overline{c}\right)\right]. \quad (18)$$

Since, by Theorem 1, the solution of (10) is unique, this is indeed the required expression for $\overline{\beta}$. Substituting this into (16) yields the approximation for β_m. $\diamond$

Theorem 2: Consider the problems T1 and T2. For each case, there exists a unique state-dependent optimal policy which is stationary and is given by

$$\mu^*_{Ti,m}(x) = -\beta_m^{Ti} x, \text{ where} \quad (19)$$

$$\beta_m^{T1} = c_m\sqrt{M/\overline{c}},\ \beta_m^{T2} = c_m/\sqrt{\overline{c}},\ m \in \mathcal{M}; \quad \overline{c} := \sum_{m=1}^{M} c_m. \quad (20)$$

The optimum team cost values for an initial state x are given by

$$J^{T1} := \inf_\mu J^{T1}(\mu) = (\sqrt{M/\overline{c}})x^2,\ J^{T2} := \inf_\mu J^{T2}(\mu) = x^2/\sqrt{\overline{c}}. \quad (21)$$

Proof: Consider case T1. What we have here is a standard linear quadratic regulator problem with cost

$$\int_0^\infty \left(M|x(t)|^2 + |u|^2_{R^{-1}}\right) dt, R = diag(c_1, ..., c_M),$$

and state dynamics

$$\dot{x} = b^T u, \quad u := (u_1, ..., u_N)^T, \quad b^T := (1, ..., 1).$$

Since controllability and observability conditions are satisfied, this problem admits a unique solution: $u = -RbPx \equiv -(c_1, c_2, ..., c_M)^T Px$, where P is the unique positive solution of the algebraic Riccati equation (which is scalar in this case): $-P^2 b^T R b + M = 0 \Rightarrow P = \sqrt{M/\overline{c}}$. This immediately leads to the expression for β_m^{T1} in (20). The corresponding optimal team cost is $J^{T1} = Px^2 \equiv (\sqrt{M}/\overline{c})x^2$ which verifies the first expression of (21). Counterparts for T2 follow by appropriate rescaling of c_m's (by replacing c_m by c_m/M, and by dividing the cost by M). $\diamond$

III. Greedy decentralized algorithms

Although the Nash equilibrium is a natural solution concept for the noncooperative decentralized control problem formulated here, its computation might yet require some coordination (and thus centralization and cooperation) between the users since it involves the individual utilities of all players, as captured by the constants c_m, $m \in \mathcal{M}$. In practice, however, these utilities are typically private information, and communicating these might result in unacceptable additional complexity. It is thus natural to investigate whether simple *greedy* "best response" algorithms could be used for updating the users' control policies in a decentralized way, thus avoiding the need for communication, coordination, and computation of the Nash equilibrium. We show in this section that this is indeed the case, and moreover, sequences generated by two such algorithms converge to the Nash equilibrium.

A general greedy "best response" algorithm is defined by the following four conditions [6]:

(i) Each user updates from time to time its policy by computing the best response against the most recently announced policies of the other users.

(ii) The time between updates is sufficiently large, so that the control problem faced by a user when it updates its policy is well approximated by the original infinite horizon problem.

(iii) The order of updates could be arbitrary, but each user performs updates infinitely often.

(iv) When the nth update occurs, a subset $K_n \subset \{1, \dots, M\}$ of users simultaneously update their policies.

Hence, assuming that an initial set of policies is linear, a general greedy best response algorithm will have the form

$$\beta_m^{(n)} = \begin{cases} f_m(\beta_{-m}^{(n-1)}) & \text{if } m \in K_n \\ \beta_m^{(n-1)} & \text{otherwise }, \end{cases} \tag{22}$$

where f_m is as defined in (14). What distinguishes one particular algorithm from another one is the choice of the set K_n. We identify here four such choices:

1. Parallel Update Algorithm (PUA): $K_n = \{1, ..., M\} =: \mathcal{M}$ for all n.

2. Round Robin Algorithm (RRA): K_n is a singleton for all n and equals
$(n+k)\mathrm{mod}M + 1$, where k is an arbitrary positive integer.

3. Random Polling Algorithm (RPA): $K_n = \{\xi_n\}$, where ξ_n, $n \geq 1$, is an independent discrete random process, where for each fixed n, the random variable ξ_n takes values in $\mathcal{M}$, with every possible outcome $(1, ..., M)$ having positive probability, uniformly bounded away from zero. Hence, here only one user acts at every update time, and this user is picked according to the outcome of a chance mechanism.

4. Stochastic Asynchronous Algorithm (SAA): $K_n = \Theta_n$, where Θ_n, $n \geq 1$, is an independent set-valued random process, where for each fixed n, the random quantity Θ_n takes values in the set of all subsets of $\mathcal{M}$, with the property that every player is in an outcome with positive probability, uniformly bounded away from zero. Hence, here more than one user is allowed to act at every update time, but the users (as well as their number) are picked randomly.

We present a numerical analysis of the first three of these algorithms in Section V. Below we study the convergence of PUA and RRA.

Theorem 3: Consider PUA.

(i.a) Let $\beta_k^{(1)} = 0$ for all k. Then $\beta_k^{(2n)}$ monotonically decrease in n and $\beta_k^{(2n+1)}$ monotonically increase in n, for every player k, and thus, the following limits exist:

$$\hat{\beta}_k := \lim_{n\to\infty} \beta_k^{(2n)}, \quad \tilde{\beta}_k := \lim_{n\to\infty} \beta_k^{(2n+1)}.$$

(i.b) Assume that $\hat{\beta}_k = \tilde{\beta}_k$ (defined as above, with $\beta_k^{(1)} = 0$ for all k). Consider now a different initial condition satisfying either $\beta_k^{(1)} \leq \beta_k$ for all k, where β_k is as given in (9), or $\beta_k^{(1)} \geq \beta_k$ for all k. Then for all k, $\lim_{n\to\infty} \beta_k^{(n)} = \beta_k$.

(ii) Global convergence: If

– (ii.a) $M = 2$, and either $\beta_k^{(1)} \leq \beta_k$ for all k, or $\beta_k^{(1)} \geq \beta_k$ for all k;

or if

– (ii.b) $\beta_k^{(1)}$ and $c := c_k$ are the same for all k,

then the sequence $\{\beta^{(n)}\}_{n\geq 1}$ converges to the unique equilibrium β^*.

(iii) Local convergence: For arbitrary c_k, there exists some open neighborhood V of β^* where β^* is the unique solution of (9), such that if $\beta_k^{(1)} \in V$ then $\beta^{(n)}$ converges to the unique equilibrium β^*.

Proof: (i) We note that for all m, f_m is positive, and is strictly decreasing in its argument:

$$\frac{df_m(\beta_{-m})}{d\beta_{-m}} < 0. \tag{23}$$

Consider first the case when $\beta_k^{(1)} = 0$ for all k. Since f_k is positive, $\beta_k^{(3)} > \beta_k^{(1)} = 0$ for all k so that $\beta_{-k}^{(3)} > \beta_{-k}^{(1)} = 0$ for all k. (23) then implies that $\beta_k^{(3)} < \beta_k^{(2)}$ and $\beta_k^{(4)} < \beta_k^{(2)}$. By an inductive argument it then follows that $\beta_k^{(2n)}$ is a decreasing sequence and $\beta_k^{(2n+1)}$ is an increasing sequence in n, for all k. This leads to the result (i) for the initial condition $\beta_k^{(1)} = 0$.

Denote by $\beta_k^{(n)}(0)$ the above sequence when obtained with the initial condition $\beta_k^{(1)} = 0$. Consider now an arbitrary initial condition $\beta_k^{(1)}$, satisfying $0 < \beta_k^{(1)} \leq \beta_k$ for all k. This condition on $\beta_k^{(1)}$ implies that $\beta_{-k}^{(1)} \leq \beta_{-k}$, and hence by (23), $\beta_k^{(2)} \geq \beta_k$. Proceeding by induction, we obtain for all integers n and for all k,

$$\beta_k^{(2n)} \geq \beta_k, \quad \beta_k^{(2n+1)} \leq \beta_k. \tag{24}$$

On the other hand, since $\beta_k^{(1)} > 0 = \beta_k^{(1)}(0)$, we have by (23): $\beta_k^{(2)} < \beta_k^{(2)}(0)$, and thus $\beta_k^{(3)} > \beta_k^{(3)}(0)$. Proceeding by induction, we get for all integers n and for all k :

$$\beta_k^{(2n)} < \beta_k^{(2n)}(0), \quad \beta_k^{(2n+1)} > \beta_k^{(2n+1)}(0). \tag{25}$$

Combining (24) with (25) establishes (i.b).

(ii.a) Let $M = 2$, and assume first that $\beta_k^{(1)} = 0$, $k = 1, 2$. Then it follows that both $(\hat{\beta}_1, \tilde{\beta}_2)$ as well as $(\tilde{\beta}_1, \hat{\beta}_2)$ (defined in the first part of the Theorem) are in Nash equilibrium (since for any integer n, $\beta_1^{(n+1)}$ is the optimal response against $\beta_2^{(n)}$ and $\beta_2^{(n+1)}$ is the optimal response against $\beta_1^{(n)}$). Since the Nash equilibrium is unique, we have $\tilde{\beta} = \hat{\beta}_2$. The proof of (ii.a) (for nonzero initial conditions as well) now follows from part (i.b) of the Theorem.

(ii.b) Assume first that $\beta_k^{(1)} = 0$ for all k. Due to the symmetry, $\overline{\beta}^{(n)} := \sum_k \beta_k^{(n)}$ is given by

$$\overline{\beta}^{(n+1)} = -(M-1)\overline{\beta}^{(n)} + \sqrt{(M-1)^2(\overline{\beta}^{(n)})^2 + M^2 c}.$$

Hence, both $\overline{\hat{\beta}} := \sum_m \hat{\beta}_m$ and $\overline{\tilde{\beta}} := \sum_m \tilde{\beta}_m$ satisfy

$$\overline{\hat{\beta}} = -(M-1)\overline{\hat{\beta}} + \sqrt{(M-1)^2(\overline{\hat{\beta}})^2 + M^2 c},$$

$$\overline{\tilde{\beta}} = -(M-1)\overline{\tilde{\beta}} + \sqrt{(M-1)^2(\overline{\tilde{\beta}})^2 + M^2 c}.$$

Since these two equations are identical, this leads to the conclusion that $\overline{\hat{\beta}} = \overline{\tilde{\beta}} =: \overline{\beta}$, where $\overline{\beta}$ satisfies $\overline{\beta} = -(M-1)\overline{\beta} + \sqrt{(M-1)^2\overline{\beta}^2 + M^2 c}$, whose only positive solution is $\overline{\beta} = \left\{M^2 c/(2M-1)\right\}^{1/2}$, which corresponds to the Nash equilibrium solution β^{M1} in (12). Hence $\hat{\beta}_k = \tilde{\beta}_k$ for all k. (ii.b) is now established by applying the result of part (i.b).

(iii) Let $\Delta\beta_k^{(n)} := \beta_k^{(n)} - \beta_k^*$ for all integers n and for all players k. The proof is established by showing that there exists some open neighborhood V of β such that $f = (f_1, ..., f_M)$ is a contraction on V, where f is defined in (14) (and is used in (22)). In other words, we have to show that there exists some matrix B whose eigenvalues are in the interior of the unit disk, such that

$$\Delta\beta^{(n+1)} = B\Delta\beta^{(n)} + o(\Delta\beta^{(n)})\underline{1}_M \qquad (26)$$

where $\underline{1}_M$ is the M-dimensional vector $(1, 1, ..., 1)^T$, $o(\cdot)$ is some function satisfying
$\lim_{x\to 0} o(x)/|x| = 0$, and $\Delta\beta^{(\cdot)}$ is an M-dimensional vector whose kth component is $\Delta\beta_k^{(\cdot)}$. Using (14) and (22), B is obtained as follows. The mk-th entry of B is given by

$$B_{mk} = \left.\frac{\partial f_m(\beta'_{-m})}{\partial \beta'_k}\right|_{\beta'=\beta^*} = b_m[1 - \delta_{km}], \qquad (27)$$

where

$$b_m := -1 + \frac{\beta^*_{-m}}{\sqrt{\beta^{*2}_{-m} + c_m}} = -1 + \frac{\beta^*_{-m}}{\overline{\beta}^*} \equiv -\frac{\beta^*_m}{\overline{\beta}^*} \qquad (28)$$

and δ_{km} is the Dirac delta function. A sufficient condition for all eigenvalues of B to be in the interior of the unit disk is that $\sum_{k\neq m} |b_k| < 1$, which follows directly from Gersgorin's Theorem (see [18] p. 344). This condition is indeed satisfied, as $\sum_{k\neq m} |b_k| = \beta^*_{-m}/\overline{\beta}^*$. ⋄

Remark 2: (Rate of convergence)
Numerical experimentation has shown that PUA has a slow rate of convergence to the Nash equilibrium. To illustrate this (analytically), consider the symmetric case, where $c_m = 1$ for all players, and the initial $\beta_m^{(1)}$ are the same. Combining (26), (27) and (28) yields

$$\Delta\beta_k^{(n+1)} = -\frac{M-1}{M}\Delta\beta_k^{(n)} + o(\Delta\beta_k^{(n)}).$$

Thus the difference between $\beta_k^{(n)}$ and its limit β_k^* decreases in a neighborhood of the equilibrium approximately by a multiplicative factor of $(M-1)/M$, and changes sign at each iteration. ⋄

Theorem 4: Consider the RRA. Then β^* given in Theorem 1 is locally stable, that is there exists some open neighborhood of β^* such that for any initial multi-policy in this neighborhood, RRA converges to β^*.

Proof: Define $\Delta\beta_m = \beta_m - \beta_m^*$. In order to establish the convergence of RRA, it suffices to show (as in the proof of part (iii) of Theorem 3) that the *linearized* version of the updating rule (14) in a small open neighborhood of β^* is asymptotically stable. This update (14) can be represented as

$$\Delta\beta'_m = f_m(\beta_{-m}) - \beta^*_m = b_m\Delta\beta_{-m} + o(\Delta\beta_{-m}) \qquad (29)$$

where $b_m := -\beta^*_m/\beta^*$, for both cases N1 and N2.

Let $\beta_m^{(n)}$ denote the value of β_m at the nth time that it is updated. Assume, without any loss of generality, that the first to iterate at time $n = 1$ is user 1. Then the linearized dynamics for the iteration are given by:

$$\begin{aligned}
\Delta\beta_1^{(n+1)} &= b_1 \sum_{i=2}^{M} \Delta\beta_i^{(n)} \\
b_m^{(n+1)} &= b_m \sum_{i=1}^{m-1} \Delta\beta_i^{(n+1)} + b_m \sum_{i=m+1}^{M} \Delta\beta_i^{(n)},
\end{aligned}$$

$m = 2, ..., M$. We can write it in matrix form as

$$\Delta\beta^{(n+1)} = A\Delta\beta^{(n)} + C\Delta\beta^{(n+1)},$$

where

$$A = \begin{pmatrix} 0 & b_1 & b_1 & .... & b_1 \\ 0 & 0 & b_2 & .... & b_2 \\ . & . & . & .... & . \\ 0 & 0 & 0 & .... & b_{M-1} \\ 0 & 0 & 0 & .... & 0 \end{pmatrix}$$

$$C = \begin{pmatrix} 0 & 0 & 0 & .... & 0 & 0 \\ b_2 & 0 & 0 & .... & 0 & 0 \\ b_3 & b_3 & 0 & .... & 0 & 0 \\ . & . & . & .... & . & . \\ b_{M-1} & b_{M-1} & b_{M-1} & .... & 0 & 0 \\ b_M & b_M & b_M & .... & b_M & 0 \end{pmatrix}$$

The matrix $(I - C)$ is invertible, so that $\Delta\beta^{(n+1)} = (I - C)^{-1} A \Delta\beta^{(n)}$. In order to establish the asymptotic stability of the linearized system it suffices to show that the eigenvalues of $(I - C)^{-1} A$ are in the interior of the unit disk, or equivalently that those of $D := A(I - C)^{-1}$ are in the interior of the unit disk. We have

$$(I - C)^{-1} =$$

$$\begin{pmatrix} 1 & 0 & 0 & .... & 0 & 0 \\ b_2 & 1 & 0 & .... & 0 & 0 \\ b_3(1+b_2) & b_3 & 1 & .... & 0 & 0 \\ b_4(1+b_2)+b_4b_3(1+b_2) & b_4(1+b_3) & b_4 & .... & 0 & 0 \\ . & . & . & .... & . & . \\ b_M(b_{M-1}+1)...(1+b_2) & . & . & .... & . & 1 \end{pmatrix}$$

and the entries of D are:

$$\begin{aligned} d_{jm} &= b_j[1 + b_{m+1} + b_{m+2}(1 + b_{m+1}) + ... \\ &\quad + b_M(1 + b_{M-1})...(1 + b_{m+1})], j < m \\ d_{mm} &= b_m[b_{m+1} + b_{m+2}(1 + b_{m+1}) + ... \\ &\quad + b_M(1 + b_{M-1})...(1 + b_{m+1})], \\ d_{jm} &= b_j[b_{j+1}(1 + b_j)...(1 + b_{m+1}) + ... \\ &\quad + b_M(1 + b_{M-1})...(1 + b_{m+1})], j > m \end{aligned}$$

(and $d_{M,m} = 0$). Some extensive but rather straightforward manipulations lead to the result that for all $m = 1, ..., M$,

$$\sum_{i=1}^{M} |d_{im}| < -(b_M + b_{M-1} + ... + b_{m+1})(1 + b_{m+1}) < 1$$

This implies that the eigenvalues of D are in the interior of the unit disk, and hence β^* is a locally stable equilibrium for RRA. ◇

IV. Multi-type traffic

We consider in this section the following extension of the model described in Section II. As before, each user has to control its own transmission rate; however, the type of traffic to be transmitted is not fixed anymore. Different traffic classes may correspond to different applications that may have different performance requirements. The state of the system, which determines the dynamics and upon which users base their control decisions, will now include not only the queue length, but also the traffic types. Each user transmits just one traffic type at a time.

To make this precise, we now introduce some notation and make some simplifying assumptions. Source m $(m \in \mathcal{M})$ may have several types (say s_m) of possible traffic, with different kinds of requirements on the performance measures. Associated with type i_m traffic of user m, we have i-dependent positive constants, $c_m(i)$, appearing in the immediate cost (instead of the constants c_m we had before). Typically, traffic requiring higher quality of service (QoS) might have a larger $c_m(i)$, which reflects the fact that it might require lower loss probabilities and higher throughput. It could be receiving a higher priority from the network in the sense that larger variations in $u_m(i)$ will be tolerated so as to achieve the required QoS.

The occurrence of these different types of traffic is governed by a *continuous-time Markov jump process* taking values in a finite state space $\mathcal{S}$; an element θ in $\mathcal{S}$ describes a possible traffic constellation of the s different sources. Transition in this underlying Markov chain may reflect the beginning and the end of sessions of different types. The rate matrix (of transitions within $\mathcal{S}$) is $\Lambda = \{\lambda_{ij}\}$, $i, j \in \mathcal{S}$, where the λ_{ij}'s are real numbers such that for any $i \neq j$, $\lambda_{ij} \geq 0$, and for all $i \in \mathcal{S}$, $\lambda_{ii} = -\sum_{j \neq i} \lambda_{ij}$.

Fix some initial state i_0 of the Markov chain $\mathcal{S}$. Consider the class of policies $\mu_m \in \mathcal{U}_m$ for controller m, whose elements are of the form $u_m(t) = \mu_m(t, x_{[0,t]}; \theta_{[0,t]})$, $t \in [0, \infty)$. Here, μ_m is taken to be piecewise continuous in its first argument, and piecewise Lipschitz continuous in its second argument. Then, we have the following counterpart of Theorem 1, with the cost per user being of type N1. The counterpart of this result for N2 follows by an appropriate rescaling of the c_m's.

Theorem 5: Consider the noncooperative framework, where user m minimizes the cost

$$J_m(\mu) := E^{\mu} \left[\int_0^{\infty} \left(|x(t)|^2 + \frac{1}{c_m(\theta(t))} |u_m(t)|^2 \right) dt \right| \\ x(0) = x_0, \theta(0) = i_0 \Big],$$

where E^{μ} denotes the expectation operator under multi-policy μ. Then, we have:

(i) There exists a Nash equilibrium in state-dependent policies, given by

$$\mu_m(x, i) = -\beta_m(i) x, \; m \in \mathcal{M}, \tag{30}$$

where

$$\beta_m(i) = c_m(i) P_m(i), \quad \beta_{-m} := \sum_{j \neq m} \beta_j, \tag{31}$$

and $\{P_m(i),\ i \in \mathcal{S},\ m \in \mathcal{M}\}$ is a positive solution of the coupled set of equations

$$-2\beta_{-m}(i)P_m(i) - P_m(i)^2 c_m(i) + 1 + \textstyle\sum_{j\in\mathcal{S}} \lambda_{ij} P_m(j) = 0,$$
$$i \in \mathcal{S}, m \in \mathcal{M}. \tag{32}$$

(ii) $\beta_m(i)$ satisfies the bound:

$$0 \le \beta_m(i) \le \frac{c_m(i)}{\min_j \sqrt{c_m(j)}}, \quad i \in \mathcal{S},\ m \in \mathcal{M}. \tag{33}$$

(iii) The cost for player m, corresponding to the Nash equilibrium above, for an initial state x, is given by $J_m(x,i) = [\beta_m(i)/c_m(i)]x^2$.

Proof: We first show that a Nash equilibrium exists among policies of the form (30). To see this, first note that since J_m is positive-quadratic in x and u_m, and x is linear in u_m for each fixed μ_k, $k \in \mathcal{M}$, $k \neq m$, the costs $J_m^{N1}(x,i;\mu)$ are strictly convex in μ_m for each user. This implies, in particular, that the cost of user m is strictly convex in β_m. We shall show, next, that β can be restricted without any loss of generality to a compact set.

Assume now that all players other than player m use policies of the form (30), where $\beta_k(i)$'s are positive. Then player m is faced with a linear quadratic control problem with jump parameters, and has an optimal response given by (30)-(32) (see [21]). We show that this response satisfies the bound (33). First note that under (31), (32) can be rewritten as

$$-2\beta_{-m}(i)\beta_m(i) - \beta_m(i)^2 + c_m(i) + \sum_{j\in\mathcal{S}} \lambda_{ij} \frac{\beta_m(j)}{c_m(j)} c_m(i) = 0,$$

which yields

$$\beta_m(i)[2\overline{\beta}(i) - \beta_m(i)] = c_m(i) + c_m(i) \sum_{j\in\mathcal{S}} \lambda_{ij} \frac{\beta_m(j)}{c_m(j)}. \tag{34}$$

Let i^* be the state for which $\beta_m(i)/c_m(i)$ is maximized. We have $\sum_{j\in\mathcal{S}} \lambda_{i^*j}\beta_m(j)[c_m(j)]^{-1} \le 0$, so that (34) yields $[\beta_m(i^*)/c_m(i^*)][2\overline{\beta}(i^*) - \beta_m(i^*)] \le 1$. This implies (since $\overline{\beta}(i^*) \ge \beta_m(i^*)$) that $\beta_m(i^*) \le \sqrt{c_m(i^*)}$. From the definition of i^* we have, for each $i \in \mathcal{S}$,

$$\frac{\beta_m(i)}{c_m(i)} \le \frac{\beta_m(i^*)}{c_m(i^*)} \le \frac{1}{\sqrt{c_m(i^*)}}$$

from which we obtain $0 \le \beta_m(i) \le c_m(i)/\sqrt{c_m(i^*)}$. This implies that the optimal response of player m also satisfies (33).

Consider now a constrained version of the game in which player k is restricted to use $\beta_k(i) \in [0, B_k(i)]$, where $B_k(i)$ is an arbitrary constant larger than the bound $c_k(i)/\min_j \sqrt{c_k(j)}$. Since the cost of a typical player k in this constrained game is strictly convex in the policy of that player, and since the policy space is now compact and convex, there exists a Nash equilibrium β^* by a standard existence result for static games (see [6], Theorem 4.3, p.179). This turns out to be, however, an equilibrium in the unconstrained game as well, because as shown above, the unique response of any player to any set of policies of the other players picked as in (30) is also of the form (30), satisfying the bound (33). Hence, this establishes the existence of a Nash equilibrium of the form (30) for the original game, which satisfies (31) and (32). $\diamond$

V. Numerical results

We present in this section results of a numerical study on the behavior of the first three greedy decentralized algorithms presented in Section III. As pointed out in Remark 2, the rate of convergence of PUA becomes slower as the number of users become large. We have therefore tested the convergence of the algorithms for two cases: M moderate ($M = 4$) and M relatively large ($M = 10$). We have focused on the symmetric case $c_m = 1$, for all m, and started the iterations with zero initial conditions. All three algorithms converged.

Figures 1 and 2 depict a comparison between the convergences of PUA and RRA, for $M = 10$ and $M = 4$, respectively. In both figures we have focused on user number 1, and have computed its β after each cycle, i.e., each time all users have updated their policies. The PUA is seen to converge quite slowly: it takes around 40 cycles for convergence, and it takes longer as the number of users grow. The RRA converges almost instantaneously to the Nash equilibrium.

Figures 3 and 4 pertain to RPA for 4 and 10 users, respectively. The user that updates at a given iteration is chosen with equal probabilities, and the choices are independent. Thus, the time between updates of user m are geometrically distributed with parameter $1/M$. Figure 3 depicts the behavior of each of the users, whereas Figure 4 depicts the behavior of only players 1,3,5,7,9. The basic unit of the x axis is one iteration. We can see that there are iterations where no update occurs: this happens when the random mechanism picks the same user at two consecutive iterations. The rate of convergence is seen to be faster than PUA and slower than RRA.

VI. Asymptotic behavior and comparisons

We present, in this section, some observations on the asymptotic behaviors of the various solutions obtained, as the number of users grow, as well as the relative rates of convergence of the proposed algorithms — all for the single-type traffic.

Asymptotic behavior for large M
In order to come up with meaningful comparisons between systems with different numbers of users, we have to make some unifying assumptions on the cost functions. Natural assumptions in this context are:

(A1) $\overline{c} := \sum_{m=1}^{M} c_m$ tends to infinity as $M \to \infty$,
(A2) Assumption (A1) holds and the rate at which $\overline{c} \to \infty$ is at least linear in M (which would be the case if c_m's are uniformly bounded away from zero).

By Theorem 2, (A1) implies that the equilibrium value J^{T2} goes to 0 as M grows to infinity. If (A2) holds then J^{T1} remains bounded as M grows to infinity. The fact that for large M the value is very low in T2, and possibly in T1, might mean that the resources are under used in some sense: we have very small deviations from the target queue length which is obtained by minor effort on the part of the controllers (both the term that corresponds to x^2 and the term corresponding to u_m^2 are small when the total cost is low). This behavior is due to the fact that the dynamics of the queue length depend only on the sum of the u_m's; however, the total cost takes into account the sum of the *squares* of the u_m's. For the same (fixed) value of the sum of u_m's, the sum of the squares of the u_m's decreases as the number of users grow.

A problem that may arise due to this situation is that the users might choose non-optimal policies, which might cause an overall inefficient use of the network (for example, large variations of the sum of u_m's) and still have a correspondingly low cost. A way to circumvent this situation is to choose a network pricing policy that does not satisfy (A1) and hence (A2). In other words, the individual costs $1/c_m$ may be chosen by the network according to the expected number of users; if the network is designed for a large number of users then the c_m should be smaller than those used in pricing in a network with a smaller expected number of users.

The situation is different, however, in the noncooperative case: the total cost (summing over all users) need not go to zero, as can be seen from Corollary 1. Assume that $\overline{c}$ grows linearly with M, i.e., $\overline{c} \sim Mc$. In the case N1, the sum of the values of the individual users goes to infinity as M goes to infinity, whereas in case N2 it converges to a constant: $x^2/\sqrt{2c}$.

Yet, even in the noncooperative case, we see that the cost *per user* tends to 0 as M tends to infinity, if $\overline{c} \sim Mc$. This again may result in the problems discussed above for the team case, but a way to circumvent this situation is again to choose a network pricing policy for which $\overline{c}$ does not grow linearly in M. For example, if we choose $\overline{c}$ to be constant in M then the value per user in case N1 will tend to a constant as M goes to infinity. For the case N2, $\overline{c}$ has to be chosen to be decreasing (like $1/M$) in order to achieve this same behavior of the value.

Comparing values and gains in T1,T2, N1,N2

We first note that the optimum (equilibrium) costs in cases T1 and N1 are, as can be expected, larger than those in the corresponding cases T2 and N2, with the multiplicative factor being $\sqrt{M}$ in each case. If we want to make a valid comparison between the team and the Nash cost value, on the other hand, say between T1 and N1, we have to sum J_m^{N1} over all M users. This is because, for a fixed policy μ for all users, $J_m^{T1}(\mu) = \sum_m J_m^{N1}(\mu)$. A comparison of Corollary 1 and Theorem 2 reveals that when $\overline{c} >> \max_m c_m$,

$$\sum_m J_m^{N1} \sim (M/\sqrt{2\overline{c}})x^2 = (\sqrt{M/2})J^{T1}.$$

As can be expected, the value is indeed higher under the noncooperative mode of play.

For the symmetric case, where $c_m = c$ for all m, this conclusion is also confirmed:

$$\sum_m J_m^{N1} = M\sqrt{\frac{1}{(2M-1)c}}x^2 = \left(\sqrt{\frac{M^2}{2M-1}}\right) J^{T1}.$$

A comparison of the control gains (β) in the symmetric team and Nash equilibrium solutions also reveal a consistent pattern, but this time in the opposite direction. Team control gains are now higher in both cases, by the same factor: $\sqrt{2M-1}$.

VII. Concluding Remarks

We have presented a fairly simple model of a communication network where users control their individual rates of transmission. This has led to very simple and appealing transmission policies under both team and noncooperative dynamic game scenarios, and has allowed us to prove powerful existence, uniqueness and convergence results in the latter case, which are reported for the first time in the literature. The simplicity of the model adopted has, inescapably, also led to some limitations in the direct application of these results to real communication networks. These limitations are:

(i) We have not taken into account delays, and the presence of noise in the information available to the users; taking these into account would have required incorporation into the model some stochasticity assumptions on the fluctuations of the available bandwidth (or consideration of robust controller design). First steps in this directions have been taken in [3] for the case of a single controller, and in [4] for a team situation (where only a suboptimal control policy has been derived). Another feature that was ignored here is the fact that information might not be available continuously, but only at some sampling instants (such as the arrival of RM cells to a source), unless the control is performed *by the bottleneck switch.*

(ii) We have not taken into account here constraints such as PCR (peak cell rate) and MCR (minimum cell rate). The latter is a feature that can be offered in the ABR service [1], although it does not exist on the internet. When considering saturated sources, one can ignore the MCR and apply the analysis of this paper only to the part that exceeds it.

(iii) As explained at the beginning of Section II, we have ignored the nonlinearities in the dynamics due to the boundary effects (empty or full buffer). We have further used a fluid model approximation, and have restricted the analysis to a simple bottleneck node.

Incorporating all these features into the model would lead to a much more complicated model and naturally to more complex control policies. Nevertheless, this is a direction worth pursuing, under the guidance of the results presented in this paper.

References

[1] The ATM Forum Technical Committee, *Traffic Management Specification*, Vers. 4.0, 95-0013R8, Apr. 1996.

[2] E. Altman, F. Baccelli and J. C. Bolot, "Discrete-time analysis of adaptive rate control mechanisms", *High Speed Networks and their performance*, H.G. Perros and Y. Viniotis Eds., North Holland, 121-140, 1994.

[3] E. Altman and T. Başar , "Optimal rate control for high speed telecommunication networks, *Proc. 34th IEEE Conference on Decision and Control*, New Orleans, Louisiana, Dec. 1995.

[4] E. Altman, T. Başar and R. Srikant, "Multi-user rate-based flow control with action delays: A team-theoretic approach", *Proc. 36th IEEE Conference on Decision and Control*, San Diego, California, Dec. 1997.

[5] B. D. O. Anderson and J. B. Moore, *Optimal Control: Linear Quadratic Methods*, Prentice-Hall, NJ, 1990.

[6] T. Başar and G.J. Olsder, *Dynamic Noncooperative Game Theory*, Academic Press, London, (2nd ed.) 1995.

[7] J-C. Bolot, "End-to-end packet delay and loss behavior in the Internet", *Proc. ACM Sigcomm '93, pp. 289-298*, San Francisco, CA, Sept. 1993.

[8] J-C. Bolot, A. U. Shankar, "Analysis of a fluid approximation to flow control dynamics", *Proc. IEEE Infocom*'92, pp. 2398-2407, Florence, Italy, May 1992.

[9] F. Bonomi and K. W. Fendick, "The rate-base flow control framework for the Available Bit Rate ATM service", *IEEE Network*, 25-39, March/Apr. 1995.

[10] A. D. Bovopoulos and A. A. Lazar, "Asynchronous iterative algorithms for optimal load balancing", *Proc. 22nd Annual Conference on Information Sciences and Systems*, Princeton Univ., March 1988.

[11] L.S Brakmo, L.L. Peterson, "TCP Vegas : End to end congestion avoidance on a global Internet," *IEEE Journal on selected Areas in communications,* vol. 13, 1465-1480, 1995.

[12] D. Cansever, "Decentralized algorithms for flow control in networks", *Proc. 25th IEEE Conference on Decision and Control*, Athens, Greece, Dec. 1986.

[13] B. Chakravorti, "Optimal flow control of an M/M/1 queue with a balanced budget", *IEEE Trans. Automatic Control*, **39**, 1918-1921, 1994.

[14] C. Douligeris and R. Mazumdar, "User optimal flow control in an integrated environment", *Proc. of the Indo-US Workshop on Signals and Systems*, Bangalore, India, Jan. 1988.

[15] C. Douligeris and R. Mazumdar, "A game theoretic approach to flow control in an integrated environment", *Journal of the Franklin Institute*, **329**, No. 3, March, pp. 383-4002, 1992.

[16] K. Fendick, M. A. Rodrigues and A. Weiss, "Analysis of a rate-based control strategy with delayed feedback," *ACM SIGCOMM '92*, Baltimore, MD, August 1992, 136-148.

[17] M. T. Hsiao and A. A. Lazar, "Optimal decentralized flow control of Markovian queueing networks with multiple controllers", *Perf. Eval.*, **13**, 181-204, 1991.

[18] R. Horn and C. R. Johnson, *Matrix Analysis*, Cambridge Univ. Press, 1985.

[19] V. Jacobson, "Congestion avoidance and control", *ACM SIGCOMM*, Stanford, CA, 314-329, Aug. 1988.

[20] R. Jain, "ERICA+", contribution ATM Forum/95-1346.

[21] Y. Ji and H.J. Chizeck, "Controllability, stabilizability, and continuous-time Markovian jump linear quadratic control," *IEEE Trans. on Automatic Control*, vol. AC-35,

[22] Y. A. Korilis and A. Lazar, "On the existence of equilibria in noncooperative optimal flow control", *J. Assoc. Comput. Mach.* **42**, pp. 584-613, May 1995.

[23] Y. A. Korilis, A. A. Lazar and A. Orda, "Architecting noncooperative networks", *IEEE J. on Selected Areas in Communications*, **13**, No. 7, pp. 1241-1251. 1995.

[24] Y. A. Korilis, A. A. Lazar, and A. Orda, "Achieving network optima using Stackelberg routing games", *ACM/IEEE Trans. on Networking*, **5**, No. 1, pp. 161-173, 1997.

[25] A. A. Lazar, A. Orda, and D. E. Pendarakis, "Virtual path bandwidth allocation in multi-user networks", *IEEE Transactions on Networking*, Vol. 5, December 1997.

[26] R. Mazumdar, L. Mason and C. Douligeris, "Fairness in network

optimal flow control: optimality of product forms", *IEEE Trans. on Communications*, **39**, 775-782, 1991.

[27] H. Ohsaki, M. Murata, H. Suzuki, C. Ikeda and H. Miyahara, "Rate-based control for ATM Networks," *Computer Communication Review, ACM-Sigcomm,* special issue in ATM, ed. R. Jain and K.Y. Siu, Vol 25, No. 2, pp. 60-71, 1995.

[28] A. Orda, R. Rom and N. Shimkin, "Competitive routing in multiuser communication networks", *IEEE/ACM Transactions on Networking*, **1**, pp. 614-627, 1993.

[29] M. Ritter, "Network buffer requirements of the rate-based control mechanism for ABR services," *IEEE INFOCOM'96,* San Francisco, USA, March 1996.

[30] Z. Zhang and C. Douligeris, "Convergence of synchronous and asynchronous algorithms in multi-class networks", *Proceedings of the IEEE INFOCOM '91*, Bal Harbor, FL, 1991.

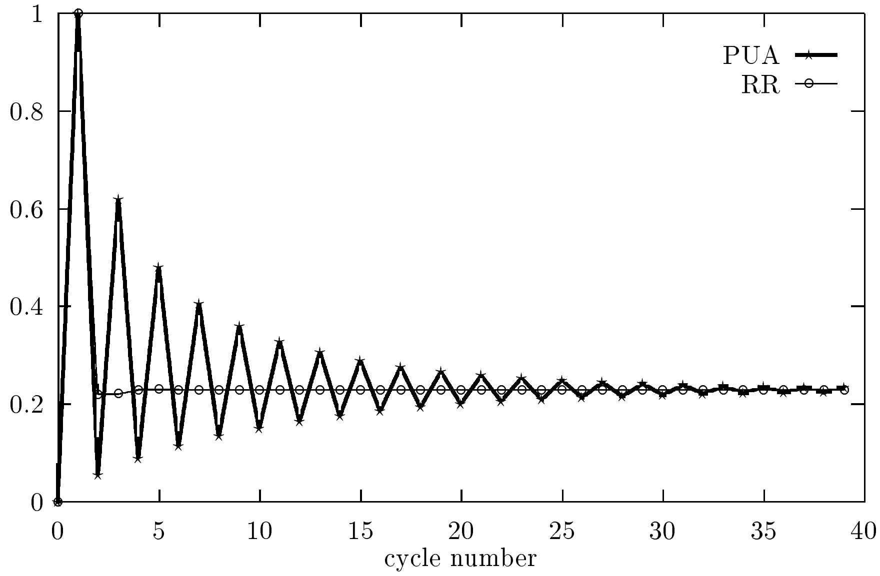

Fig. 1. PUA versus RRA for $M = 10$

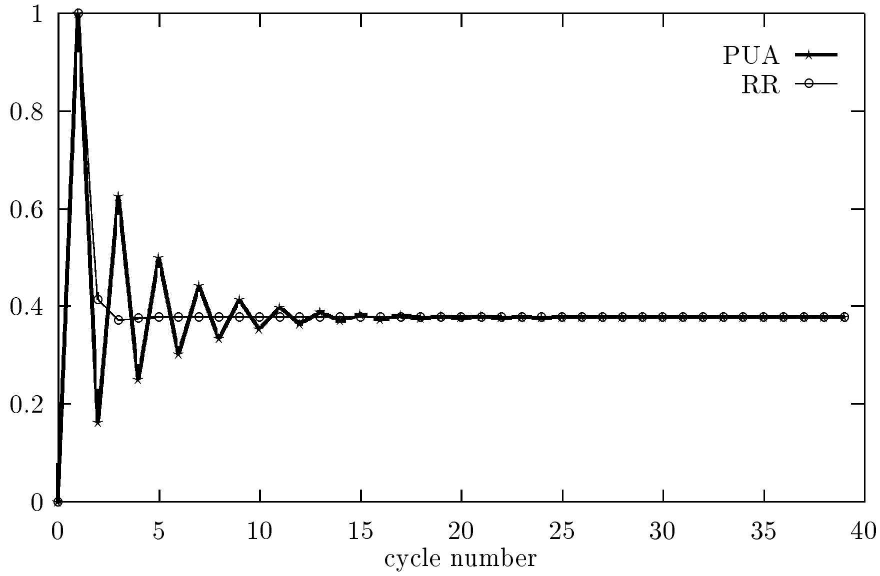

Fig. 2. PUA versus RRA for $M = 4$

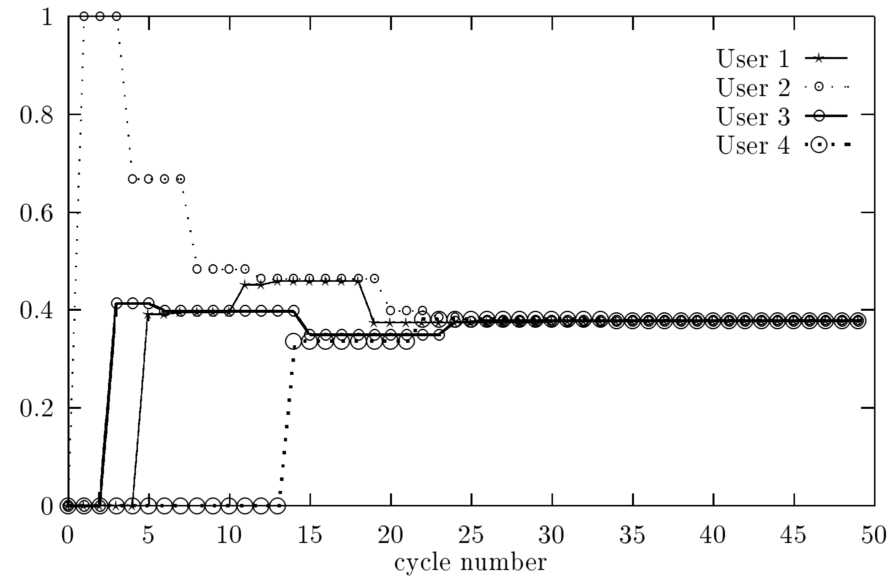

Fig. 3. RPA for $M = 4$

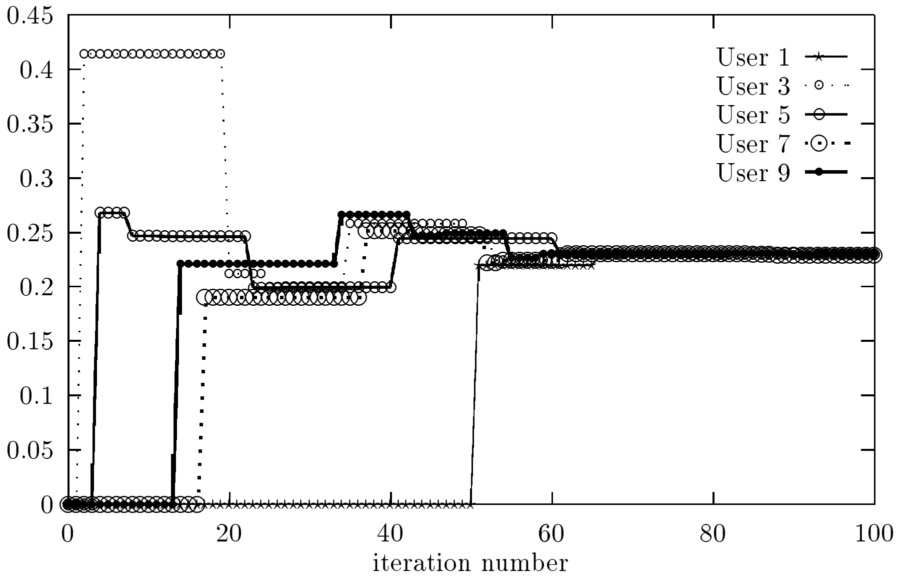

Fig. 4. RPA for $M = 10$