Robust Rate Control for Integrated Services Packet Networks

Franco Blanchini, *Member, IEEE*, Renato Lo Cigno, *Associate Member, IEEE*, and Roberto Tempo, *Fellow, IEEE*

***Abstract*—Research on congestion-control algorithms has traditionally focused more on performance than on robustness of the closed-loop system to changes in network conditions. As the performance of the control loop is strictly connected with the quality of service, these systems are natural candidates to be approached by the optimal control theory. Unfortunately, this approach may fail in the presence of transmission delay variations, which are unavoidable in telecommunication systems.**

In this paper, we first show the fragility of optimal controllers and demonstrate their instability when the control delay is not known exactly. Then we propose a robust control algorithm based on a classical proportional integral derivative scheme which does not suffer from this fragility phenomenon. Its stability versus the control delay variations, as well as versus sources that transmit less than their computed share, is studied with Nyquist analysis. The control algorithm is implemented within a simulator in the framework of the asynchronous transfer mode (ATM) ABR transfer capability. The final part of the paper shows some selected results assessing the performance of the control algorithm in a realistic network environment. ABR was chosen as an example, but the control studied here can be applied in any data network to obtain a robust and reliable congestion-control scheme.

***Index Terms*—Integrated networks, PID compensators, robust congestion control, uncertain transmission delay.**

I. Introduction

INTEGRATED services packet networks (ISPNs) are a broad category of telecommunication networks that support several different services on a common packet-switched network layer. The definition of ISPN is so broad that most of the future telecommunication networks fall within its definition. Most naturally, asynchronous transfer mode (ATM) networks [1] are ISPNs, but the IntServ [2] and DiffServ [3] architectures, which are defining the infrastructure of the future Internet, are also based on an ISPN approach.

Many papers have appeared in the literature dealing with the congestion-control problem in ISPNs. Few of these are based on a high-level model [4]–[6], while the vast majority concentrates on a specific network architecture, such as TCP/IP [7], [8] or ATM [9]–[15]. The attention on communication networks control is far from dormant, as proven also by [16], [17]. The papers in these special issues cover a wide range of problems of current interest in communication networks, with the common thread being a focus on control and optimization issues.

Congestion control in telecommunication networks struggles with two major problems that are not completely solved. The first one is the unknown time-varying delay between the control point and the traffic sources. The second is the possibility that the traffic sources do not follow the feedback signal. It may always happen that some sources are silent because they have nothing to transmit. A global optimization approach of the dynamic system is, to our best knowledge, not feasible. Recently, a distributed iterative algorithm for the global optimization of the network working point has been proposed [8]; however, the algorithm operates in steady state, while here we are interested in the transient behavior of the network.

In this paper, we consider a classical approach to design a proportional integral derivative (PID) control. The class of systems we consider is characterized by dominant delays. It is known [18] that for this class of systems, more sophisticated control laws than PID, such as optimal controllers, may provide better performance as long as the delay is known. Unfortunately, when the delay is unknown, as in congestion control, these controllers suffer from fragility; namely, they may destabilize the system for any deviation of the delay from the nominal value. Delay measurements do not help in this context, since delays may change and contain randomness that prevent exact measures. Conversely, PID compensators do not suffer from this kind of fragility phenomena.

The main features of this paper are summarized as follows.

- We deal with delays which can be unknown and time-varying.
- We show that the optimal control does not adequately solve the problem because the resulting controllers are fragile. This can be seen even in the simplest case of a single link, where *any* variation of the time delay from the nominal one renders the closed-loop system unstable.
- We propose a classical control design approach based on a PID controller for which well-established design tools are available. For the single-link case, we provide some *analytic* conditions which assure that the controller stabilizes the system for *all* delay values below an assigned upper limit.
- We finally consider the multisource case and we prove the following *equivalence property*: the multisource case reduces, in our framework, to the single-source case in which the maximum admissible delay is equal to that of the source at the maximum distance.

Manuscript received June 6, 2000; revised April 7, 2001 and February 6, 2002; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor S. Low. This work was supported by the Italian National Research Council (CNR) and the Italian Ministry of University under the PLANET-IP Project. A preliminary version of this paper was presented at the 1998 American Control Conference, Philadelphia, PA.

F. Blanchini is with the Dipartimento di Matematica e Informatica, Università degli Studi di Udine, 33100 Udine, Italy (e-mail: blanchini@uniud.it).

R. Lo Cigno is with the Dipartimento di Elettronica, Politecnico di Torino, 10129 Torino, Italy (e-mail: locigno@polito.it).

R. Tempo is with IRITI-CNR, Politecnico di Torino, 10129 Torino, Italy (e-mail: tempo@polito.it).

Digital Object Identifier 10.1109/TNET.2002.803896.

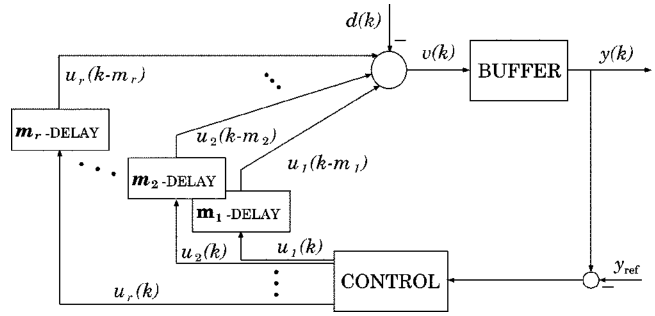

Fig. 1. Dynamic model for the congestion control of ISPNs.

A. Problem Modeling and Formulation

We restrict our analysis to the case of unicast connections, nonetheless, the considered system is extremely complex, being a multisource system with fast dynamics. We notice that any connection has several controllers, one at each node; however, at any given time only one controller, the one with the least resources, will drive the source. The bottleneck can jump from one node to another depending on the traffic conditions.

The analysis is based on a discrete time-modeling approach, which is deemed more suitable for communication networks and leads to straightforward implementations.

The dynamic system to be controlled is described by

$$y(k+1) = y(k) - d(k) + \sum_{i=1}^{r} u_i(k - m_i), \qquad (1)$$

where k is the discrete time, $y(\cdot)$ is the buffer level, $d(\cdot)$ is a compound disturb term that takes into account the background traffic as well as the link capacity, r is the number of controlled sources or connections, $u_i(\cdot)$ is the control signal for the ith source, and m_i is the (unknown) control delay for the ith source. Fig. 1 shows the dynamic system represented by (1), with a generic controller that generates the control signals. This is the same dynamic model that is often discussed throughout the literature on ISPN congestion control [4], [9], [14], [15]. The addressed problem is the definition of an appropriate form for the controller.

II. Instability of Optimal Control

We begin our analysis by considering the case of a single source. With a single source, (1) reduces to

$$y(k+1) = y(k) + u(k-m) - d(k).$$

The goal of the control is to keep the buffer level y as close as possible to the nominal value $y_{\rm ref}$ despite the action of the disturbance d. There is no restriction in setting $y_{\rm ref} = 0$. This is a disturbance rejection problem, for which two approaches appear meaningful in our case.

- The $\mathcal{H}_2$ approach or *stochastic approach.* The disturbance is assumed to be a Gaussian noise. A stabilizing compensator is designed to minimize the output power or the root mean square value [20], [22]. This is known to be equivalent to the minimization of the $\mathcal{H}_2$ norm of the closed-loop transfer function or the l_2 norm of the closed-loop impulse response.
- The l_1 approach or *worst case approach.* The disturbance $d(k)$ is assumed unknown but bounded $|d(k)| \leq 1$. The stabilizing compensator is designed to minimize the worst case output amplitude $\sup_{k\geq 0} |y(k)|$. In the linear compensator case, this is known to be equivalent to the minimization the l_1 norm of the closed-loop transfer function [20].

Although the above mentioned criteria appear the most meaningful ones for our problem, there are several other methods available, such as $\mathcal{H}_\infty$ (see [20]–[22] for details).

Proposition 1: Both the $\mathcal{H}_2$ and l_1 controllers for the considered system coincide and are equal to

$$K_{\rm opt}(z) = -\frac{z^m}{1 + z + \cdots + z^m}.$$

The closed-loop system with this compensator is

$$\Phi_{\rm opt}(z) = -\frac{1 + z + \cdots + z^m}{z^{m+1}}.$$

Proof: Note that the considered closed-loop system is FIR and that $||\Phi_{\rm opt}||_{\mathcal{H}_2} = \sqrt{m+1}$ and $||\Phi_{\rm opt}||_{l_1} = m+1$; the latter condition implies that $\sup_{h\geq 0} |y(h)| \leq m+1$ for $|d(k)| \leq 1$. Thus, we need to show that there is no compensator which improves this performance.

The $\mathcal{H}_2$ Case: Consider the impulse $d = \{1, 0, 0, \ldots\}$. The corresponding output for $k = 1, 2, \ldots, m$ is $y(k) = -1$, because there is no effect of the control which is delayed of m. Now $y(m+1) = -1 + u(0)$. Since we assume zero initial conditions and the controller ignores $d(k)$ at time k, then we should have $u(0) = 0$ and, consequently, $y(m+1) = 1$. This implies that the $\mathcal{H}_2$ norm of the output impulse response is $\sqrt{\sum_{k=0}^{\infty} y^2(k)} \geq \sqrt{m+1}$.

The l_1 Case: Consider any bounded input $|d(k)| \leq 1$, then the corresponding output at $k = m+1$ is $y(m+1) = -d(0) - d(1) - \cdots - d(m) + u(0)$. As in the $\mathcal{H}_2$ case, since the controller ignores $d(k)$ at time k, it follows that $u(0) = 0$. Therefore, for $d(k) \equiv 1$, we obtain $|y(m+1)| \geq m+1$. □

Thus, the above proposition might lead us to conclude that the optimal disturbance rejection problem for the considered plant can be easily solved, at least under the considered criteria. Unfortunately, there is a serious drawback with the solutions above. If the delay time m is not known exactly, but it can be determined only with a certain approximation, then the compensator is implemented using the estimated value $\hat{m}$ of the delay. Consequently, the resulting compensator is

$$K(z) = -\frac{z^{\hat{m}}}{1 + z + \cdots + z^{\hat{m}}}.$$

Straightforward computations show that if $\hat{m}$ is not the true value of the delay with this new compensator, we get one of the following closed-loop polynomials:

$$\begin{aligned} p_{cl}(z) &= z^{\hat{m}+1} + z^{\hat{m}-m} - 1, \qquad \text{if } \hat{m} > m \\ p_{cl}(z) &= z^{m+1} - z^{m-\hat{m}} + 1, \qquad \text{if } \hat{m} < m. \end{aligned} \qquad (2)$$

Proposition 2: The polynomials in (2) are *unstable.*

Proof: The absolute value of the known term of both polynomials is 1. Thus, denoting by λ_i the roots of the polynomials, we have for all $m \neq \hat{m}$, $\prod |\lambda_i| = 1$. Then $\max_i |\lambda_i| \geq 1$. □

From a practical point of view, this fact clearly has a negative consequence on the "expected optimal" controller $K_{\rm opt}$. This fragility phenomenon of the optimal control is not new and it has already been pointed out in the literature [23], [24], although, to our best knowledge, it has never been considered in telecommunication networks congestion control. The instability of optimal controllers based on other metrics than $\mathcal{H}_2$ and l_1 can be easily proven.

A recent paper [25] shows how to stabilize an H_∞ optimal rate controller for high-speed networks given an upper bound to the control delay is known; however, the same paper shows that the controller performance gets worse as the real delay becomes smaller, which is a property that may have undesirable consequences.

III. Classical Approach Based on PID Compensators

Motivated by the difficulties discussed in the previous section, we propose a different synthesis approach that guarantees robustness against time-delay variations. Henceforth, we assume that the delay m is unknown but bounded by a known quantity

$$0 \le m \le M.$$

We consider PID compensators. The PID controller is a standard compensator which is widely used in many applications. Its basic properties can be found in textbooks (see, e.g., [27]). A PID compensator consists of the three terms proportional, integral, and derivative, and has the form

$$\begin{aligned}\xi(k) &= \xi(k-1) + y(k)\\ u(k) &= C_P y(k) + C_D[y(k) - y(k-1)] + C_I \xi(k)\end{aligned} \tag{3}$$

where C_P, C_D, and C_I are positive design parameters. The corresponding z transform is

$$u(z) = \left(C_P + C_D\,\frac{z-1}{z} + C_I\,\frac{z}{z-1}\right) y(z).$$

The proposed synthesis method is based on the Nyquist diagram of the open-loop transfer function

$$\begin{aligned}F(z) &= \frac{1}{(z-1)z^m}\left[C_P + C_D\,\frac{z-1}{z} + C_I\,\frac{z}{z-1}\right]\\ &= \frac{1}{z^{m+1}}\left[C_T + \frac{(C_P+2C_I)}{z-1} + \frac{C_I}{(z-1)^2}\right]\end{aligned} \tag{4}$$

where $C_T = (C_P + C_D + C_I)$. The unit disk is the stability domain. Setting $z = e^{j\theta}$, $0 \le \theta \le \pi$ and noting that

$$\frac{1}{e^{j\theta}-1} = -\frac{1}{2}\left[1 + j\cot\left(\frac{\theta}{2}\right)\right]$$

simple computations yield

$$F(e^{j\theta}) = \underbrace{[(p - rx^2) - jqx]}_{G(x)}\, e^{-j\theta(m+1)} \tag{5}$$

where

$$\begin{aligned}x(\theta) &\doteq \cot\left(\frac{\theta}{2}\right)\\ p &\doteq C_P/2 + C_D + C_I/4\\ q &\doteq C_P/2 + C_I/2\\ r &\doteq C_I/4.\end{aligned}$$

Note that as θ ranges over $(0,\,\pi]$, $x(\theta)$ ranges over $[0,\,\infty)$. Thus, we can easily see that the Nyquist plot of the function $G(x)$ as defined in (5) is a parabola, as represented in Fig. 2.

Denote by $\overline{x} = \cot\left(\overline{\theta}/2\right)$ the (unique) value of $x \ge 0$ such that $|G(x)| = 1$ ($G(\overline{x})$ is the point A in Fig. 2). Since the points of the Nyquist diagram of $F(e^{j\theta})$ are the points of the Nyquist diagram of $G(x)$ rotated by $(m+1)\theta$ in the clockwise direction, the classical Nyquist stability condition requires that $(m+1)\overline{\theta} \le \phi$, where $\phi = \arccos(-\mathrm{Re}[G(\overline{x})])$ is the (positive) angle represented in Fig. 2. The robust stability condition is then formalized in the following theorem.

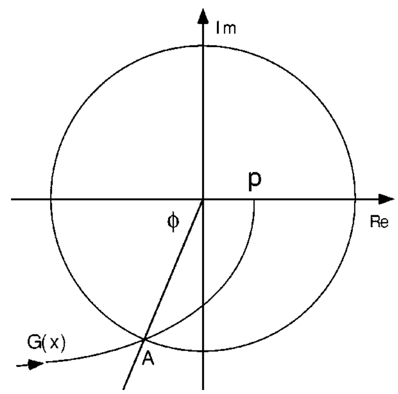

Fig. 2. Nyquist plot of $G(x)$.

Theorem 1: The closed-loop system is stable if and only if $|p| < 1$ and the delay is bounded by

$$0 \le m < \frac{\arccos(r\overline{x}^2 - p)}{2\arctan(1/\overline{x})} - 1 \doteq M.$$

Proof: Sufficiency. Assume $|p| < 1$. For all $R \ge 1$ there exists a unique solution in $0 < \theta \le \pi$ to the equation $|F(j\theta)| = R^2$. Indeed

$$|F(j\theta)|^2 = |G(j\theta)|^2 = [(p - rx^2)^2 + (qx)^2] = R^2$$

is equivalent to

$$r^2x^4 + (q^2 - 2pr)x^2 + (p^2 - m^2) = 0.$$

Thus, if $R > 1$, the above biquadratic equation admits a single positive root $\overline{x}$. This means that, in particular, for $0 < \theta \le \pi$ the finite part of the Nyquist plot from $-\infty^2 - j\infty$ reaches the unit disk and never leaves it. The unique intersection with the unit circle is $F(j\overline{\theta})$. Under the above conditions, this intersection is such that its complex phase $\angle F(j\overline{\theta}) > -\pi$. $z = 1$ is an open-loop pole on the unit circle, and we manage this case with the boundary deviation technique by adding an infinitesimal half-circle

$$\mathcal{C}_\epsilon = \{z = \epsilon e^{j\alpha},\ -\pi/2 \le \alpha \le \pi/2\}$$

to the unit disk. This choice is equivalent to treat $z = 1$ as a stable open-loop pole. When α passes from $-\pi/2$ to $\pi/2$, the Nyquist diagram describes a circle at infinity in the clockwise direction. Then the resulting complete Nyquist plot (i.e., including $z = \infty$) does not encircle the point $1 - j0$. According to the Nyquist stability criterion, we have closed-loop stability.

Necessity can be proved by noticing that if either $p \ge 1$ or $M \le m$ the origin is touched or encircled by the Nyquist plot. □

By elementary geometric considerations, $\overline{x}^2$ can be computed as the positive roots of

$$(p - rx^2)^2 + q^2x^2 = 1.$$

Note that for $x = 0$, $G(0) = p$ (see Fig. 2). Since $|p| \le 1$ is a necessary condition, the equation above is solvable. We studied the plot of the maximum delay M for varying values of the control parameters C_P, C_D, and C_I. $M(C_P,\,C_D,\,C_I)$ is indeed a three-dimensional (3-D) surface not shown here for its difficult

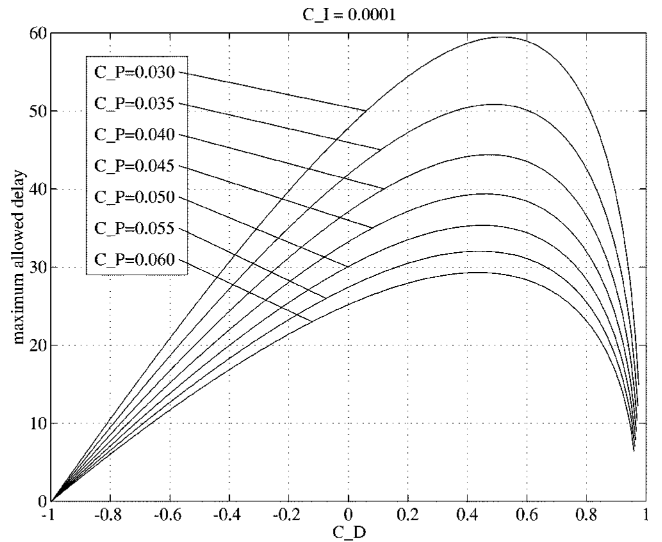

Fig. 3. Value of the maximum allowed delay M versus C_D for $C_I = 0.0001$ and different values of C_P.

graphical representation. This surface has some peculiar characteristics: 1) M monotonically increases while C_P decreases; 2) Given C_P and C_I, M has a single maximum in $\overline{C}_D$, $\overline{C}_D$ increases as the ratio C_P/C_I decreases, eventually becoming coincident with $C_D^{\max} = 1 - C_P/2 - C_I/4$. Fig. 3 shows several cross sections of the surface $M(C_P, C_D, C_I)$ for a fixed value of C_I and different values of C_P.

IV. MULTISOURCE CASE

Let us consider now the general model of several sources sharing the same buffer as described by (1). We recall that $u_i(k)$ is the speed rate assigned at the time k to the ith source and m_i is the unknown control delay for connection i. We assume that all the delays are bounded by the same maximum delay $0 \le m_i \le M$. We work under the assumption that $u_i(k)$ is

$$u_i(k) = \rho_i u(k)$$

where $\rho_i > 0$ are assigned weights such that

$$\sum \rho_i \le 1$$

and $u(k)$ is a transmission rate required by the node congestion controller. In this way, we define a partition of the signal $u(k)$ among the sources with assigned weights. Each weight may be selected, for instance, on a priority criterion. We assume that this priority criterion is *unknown* to the controller being the choice possibly made at a different level. The inequality takes into account the case in which some sources are idle or underutilize resources.

By performing the same computations that lead to (4) in Section III, the open-loop transfer function that we need to consider for closed-loop stability of the multisource case is

$$F(z) = \frac{1}{z-1} \sum_{i=1}^{r} \rho_i \frac{1}{z^{m_i}} \left[C_P + C_D \frac{z-1}{z} + C_I \frac{z}{z-1} \right]$$

where r is the number of sources. For $z = e^{j\theta}$, we write

$$F(j\theta, \rho_i, m_i) = \sum_{i=1}^{r} \rho_i e^{-j\theta(m_i+1)} G(j\theta) \doteq z(\theta, \rho_i, m_i) G(j\theta)$$

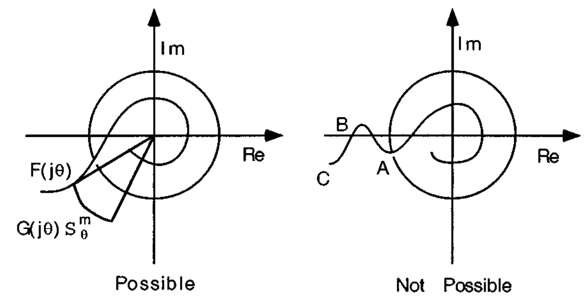

Fig. 4. Nyquist plots showing the two cases considered in Proposition 3.

where $G(j\theta)$ has the same expression as above. Note that for $m_1 = M$ and $\rho_1 = 1$ and $\rho_i = 0$ $i > 1$, we immediately obtain the single-source case. The compensator chosen in terms of C_P, C_I, C_D guarantees stability for all $\rho_i \ge 0$ with $\sum \rho_i \le 1$ and all $0 \le m_i \le M$.

The next result is what we call the property of the *equivalent single source*, which extends in a significant way the results of Section III.

Theorem 2: Given (C_P, C_I, C_D), the PID compensator stabilizes the system for all $0 \le m_i \le M$ and $\sum_{i=1}^{r} \rho_i \le 1$, $\rho_i \ge 0$, if and only if it stabilizes the single-source system with delay M.

Proof of Necessity: The single-input case with delay M is equivalent to the possible choice $m_1 = M$, $\rho_1 = 1$ and $\rho_i = 0$ for $i \ge 2$ (single active source). □

To prove sufficiency, we need first to examine two propositions. The first proposition states that the portion of the Nyquist plot of $F(j\theta, \rho_i, m_i)$ which is outside the unit circle lies in the union of the third and fourth sectors (see Fig. 4).

Proposition 3: Assume that the single-source closed-loop system is stable. Then, denoting by $\overline{\theta}$ the unique value of θ such that $|F(j\overline{\theta})| = 1$, we have $\mathrm{Im}[F(j\theta)] < 0$ for $0 < \theta \le \overline{\theta}$.

Proof: See the Appendix.

Proposition 4 ($-1 + j0$ Exclusion Principle): For a fixed frequency θ define the value set

$$V(\theta) = \left\{ z = F(j\theta, \rho_i, m_i), \sum_{i=1}^{r} \rho_i \le 1, 0 \le m_i \le M \right\}.$$

The system is stable for all ρ_i and m_i if and only if the following two conditions hold:

- There is a choice of the ρ_i and m_i such that the system is stable.
- For $0 \le \theta \le \pi$ the value set does not include the point $-1 + j0$.

Proof: See, e.g., [26].

As already mentioned, the case $\rho_1 = 1$, $\rho_i = 0$, $i \ge 2$, and $m_1 = M$, corresponds to the single-source case. Therefore, the first condition is immediately satisfied. Next, we show that the second condition is also satisfied. Indeed, the basic idea of the proof is to observe that the value set is included in the circular sector denoted by $G(j\theta)\mathcal{S}_\theta^M$ in Fig. 4 having $F(j\theta)$ as the left corner. Then we consider the fact that the "possible" Nyquist plot is as in Fig. 4, thus, such sector cannot include the point $-1 + j0$. Conversely, note that in the "not possible" case of Fig. 4, the same sector could include -1. Nevertheless, the compensator could stabilize $F(z)$.

Proof of Sufficiency of Theorem 2: First, note that with the assigned constraints and for given θ, the following inclusion holds:

$$z(\theta, \rho_i, m_i) \doteq \sum_{i=1}^{r} \rho_i e^{-j\theta(m_i+1)} \in \mathcal{S}_\theta^M \tag{6}$$

where, denoting by $\angle z$ the complex phase of z, the set $\mathcal{S}_\theta^M$ is the circular sector of radius one included between the angles $-(M+1)\theta$ and $-\theta$:

$$\mathcal{S}_\theta^M = \{z \in \mathcal{C}\colon |z| \le 1, \text{ and } -(M+1)\theta \le \angle z \le -\theta\}.$$

This fact is immediate since $\mathcal{S}_\theta^M$ is a convex set and

$$z(\theta, \rho_i, m_i) = \sum_{i=1}^{r} \rho_i e^{-j\theta(m_i+1)} + \left(1 - \sum_{i=1}^{r} \rho_i\right) 0$$

namely, $z(\theta, \rho_i, m_i)$ is a convex combination of the origin and the unit vectors $e^{-j\theta(m_i+1)} \in \mathcal{S}_\theta^M$. Note also that each point of the boundary of $\mathcal{S}_\theta^M$ is reached by an appropriate choice of ρ_i and m_i. Then, to prove sufficiency of Theorem 2, we simply need to apply the inclusion (6) and the claim of Proposition 3 and to invoke the $-1+j0$ exclusion principle of Proposition 4. For $0 < \theta \le \pi$, the following inclusion holds for all m_i and ρ_i:

$$F(j\theta, \rho_i, m_1) = G(j\theta) z(\theta, \rho_i, m_i) \in G(j\theta)\mathcal{S}_\theta^M.$$

Thus, the value set $V(\theta)$ is included in $G(j\theta)\mathcal{S}_\theta^M$, for all θ. Now, we have that the circular sector $G(j\theta)\mathcal{S}_\theta^M$ is bounded by the two vectors $F(j\theta) = G(j\theta)e^{-j(m+1)\theta}$ and $G(j\theta)e^{-j\theta}$, of magnitude $|F(j\theta)| = |G(j\theta)|$. Since $F(j\theta) = G(j\theta)e^{-(m+1)\theta}$, the vertex of the smallest phase of $G(j\theta)\mathcal{S}_\theta^M$ is $F(j\theta)$. In fact, the sector $G(j\theta)\mathcal{S}_\theta^M$ (see Fig. 4) has the following expression:

$$G(j\theta)\mathcal{S}_\theta^M = \{z\colon \angle F(j\theta) \le \angle z \le \angle G(j\theta) - \theta, \, |z| \le |G(j\theta)|\}.$$

For $\overline{\theta} < \theta \le \pi$, $|G(j\theta)| = |F(j\theta)| < 1$, the sector $G(j\theta)\mathcal{S}_\theta^M$ is in the interior of unit disk. Thus, the value set $V(\theta) \subset G(j\theta)\mathcal{S}_\theta^M$ is also in the interior of the unit disk and, therefore, it does not include -1.

For $0 < \theta \le \overline{\theta}$, we have by Proposition 3 that $\mathrm{Im}[F(j\theta)] < 0$; thus, $-\pi < \angle F(j\theta) < 0$. Furthermore, $\angle G(j\theta) - \theta < \angle G(j\theta) < 0$. Then the sector $G(j\theta)\mathcal{S}_\theta^M$ is strictly included in the open lower half plane (i.e., $\{z\colon \mathrm{Im}[z] < 0\}$), because there lie its extreme vectors and, therefore, it does not include -1. For $\theta = 0$, any transfer function has a singularity, thus, the value set is at infinity. Since the plant $F(j\theta)$ is stable, by applying the $-1+j0$ exclusion principle, we have stability. □

V. Design, Implementation, and Results

As discussed in Section III, the PID compensator consists of three terms. The term C_P produces a feedback signal proportional to the relative buffer level y. The derivative term $C_D((z-1)/z)$ produces a signal proportional to the variation of y, and it gives benefit to the closed-loop system stability. The integral action, whose effects are typically negative for stability, has the following role: under constant background traffic conditions, the steady-state buffer is forced to be at its prescribed level. In our context, the integral action has a *fundamental property* concerning the presence of idle or nonpersistent sources. Suppose that a node controls r sources with equally partitioned transmission rate, namely, $\rho_i = 1/r$ for all i, and consider a fixed disturbance $d_{ss} > 0$. Under steady-state conditions, each source transmits at u_{ss}/r. The equilibrium condition is given by

$$\overline{y} = \overline{y} - d_{ss} + u_{ss}$$

where $\overline{y}$ is the constant buffer level, which implies $u_{ss} = d_{ss}$. Assume now that r_{id} sources become idle. Then, this equilibrium does not hold anymore. The new equilibrium condition is, therefore, characterized by

$$\overline{y} = \overline{y} - d_{ss} + (r - r_{id})\frac{u_{ss}^*}{r}$$

then $u_{ss}^* = d_{ss}r/(r - r_{id})$. Thus, each active source transmits at $u_{ss}^*/r = u_{ss}/(r - r_{id}) = d_{ss}/(r - r_{id}) > d_{ss}/r$. This means that the transmission rate devoted to the idle sources is automatically directed to the active ones. Using a proportional-derivative control, the important property of distributing unused resources may fail. Indeed, under steady-state conditions, the derivative action is no longer active and, therefore, $u_{ss}^* = -C_P\overline{y}$. Consequently, the equilibrium implies

$$-d_{ss} - \frac{r - r_{id}}{r} C_P \overline{y} = 0.$$

Since $d_{ss} > 0$, then $\overline{y} < 0$. However, in practice, the buffer is lower saturated as $y \ge -y^-$ for some $y^- > 0$. Then, we must have

$$d_{ss} \le \frac{r - r_{id}}{r} C_P y^-.$$

If such condition is not satisfied, then the buffer saturates to the level $\overline{y} = -y^-$. In this case, the transmission rate of each active source remains equal to $C_P y^-/r < d_{ss}/(r - r_{id})$, which is smaller than that achieved at a different level $\overline{y} > -y^-$. The integral action avoids this situation. As long as $\overline{y} = -y^-$, we have that $u(k)$ grows at the rate $C_I y^-$. The effect is that y is removed from its lower bound and is eventually driven to the desired steady-state condition $y = 0$, ensuring full link utilization and max–min fairness.

The presence of the integrator requires special attention concerning its implementation. The integral component of the control must be saturated. Define the saturation function as

$$\mathrm{sat}_{\xi^+,\xi^-}[\xi] = \begin{cases} \xi^+, & \text{if } \xi^+ < \xi \\ \xi, & \text{if } \xi^- \le \xi \le \xi^+ \\ \xi^-, & \text{if } \xi < \xi^-. \end{cases}$$

The first equation in (3) should be rewritten as

$$\xi(k) = \mathrm{sat}_{\xi^+,\xi^-}[\xi(k-1) + y(k)]$$

for appropriate ξ^+ and ξ^-. This saturation process avoids excessive and unnecessary integral actions. Furthermore, it avoids divergence of the integral action in the case in which a node buffer remains empty $y = -y^-$ due to the absence of traffic in the node.

As far as the problem of designing the parameters C_P, C_D, and C_I is concerned, we stress that there are many tuning methods available (see, e.g., [27]). In our case, we have to impose an additional constraint, namely, robustness against all possible delay variations. According to the results of Section IV, we can synthesize the parameters by taking into account a single-source model whose delay is larger than the largest delay

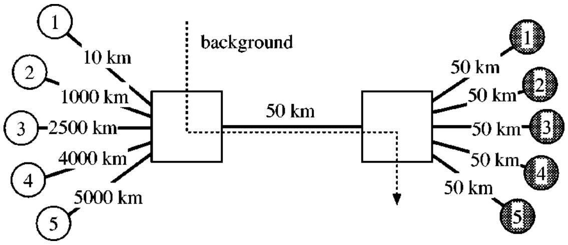

Fig. 5. Experimental setup.

allowed in the network. To assure robust stability, C_P, C_D, and C_I parameters must satisfy the conditions given in Theorem 1.

A. Simulation Results

The analysis and design carried out are very general and can be applied to any communication network that provides a general framework for rate control. We use a standard ABR implementation, to show the performance of the closed-loop control algorithm proposed in this paper. We added our control algorithm on top of the Cell Level ATM Services Simulator (CLASS), which is a detailed tool for ATM network analysis and design [31].

The description of ABR is beyond the scope of this paper. The interested reader can find all the required information either in standard documents [28] or in textbooks such as [1]. In order to provide some comparisons, we report results obtained with the ERICA+ control algorithm [11], [29]. This algorithm is reported as a sample algorithm in ATM Forum documents [28] and has received considerable attention in recent years.

First, we present results for a simple bottleneck topology, depicted in Fig. 5. All link speeds are 150 Mb/s. Connections are unidirectional. The sampling interval is set to 500 slots $T_s \simeq 1.415$ ms. The choice of T_s is not critical, since anything larger than the slot size and smaller than the maximum round-trip time (RTT) is suitable and does not influence the performance of the system. We consider wide-area networks spanning a continent, hence, limiting the connection length to roughly 5000 km. These parameters yield $M = 34$; we set $M = 40$ to stay on the safe side and account for possible additional delays. We set the buffer reference to $y_{\rm ref} = 0.53$ MB, searching for a compromise between the delay and the link utilization.

Fig. 6 shows the basic convergence behavior of the PID controller when five connections compete for the resources, and two connections have a small peak cell rate (PCR). The connection lengths are distributed as described in Fig. 5. The PCR-limited connections are the 1000-km connection that has PCR $=$ 15 Mb/s, and the 4000-km connection that has PCR $= 20$ Mb/s. The max–min fairness criterion states that the other three connection should receive 38.3 Mb/s each at steady state. Similar results without PCR-limited connections show analogous convergence properties. Fig. 6 confirms that the proposed control algorithm is capable of coping with multiple connections, each one with a different propagation delay, and some of them possibly underutilizing resources.

The real goal of congestion control is the adaptation of source rates to network changes, i.e., its behavior in presence of background traffic or disturbs. Fig. 7 reports the same performance

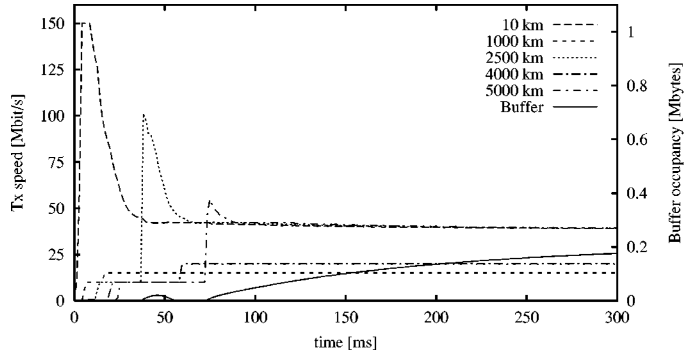

Fig. 6. Transmission rate and bottleneck buffer size. PID controller with five connections at different distances; two are PCR limited.

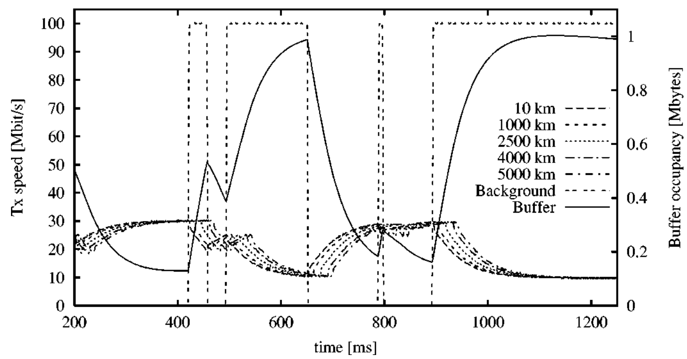

Fig. 7. Transmission rate and bottleneck buffer size. PID controller with five connections at different distances and ON–OFF background traffic.

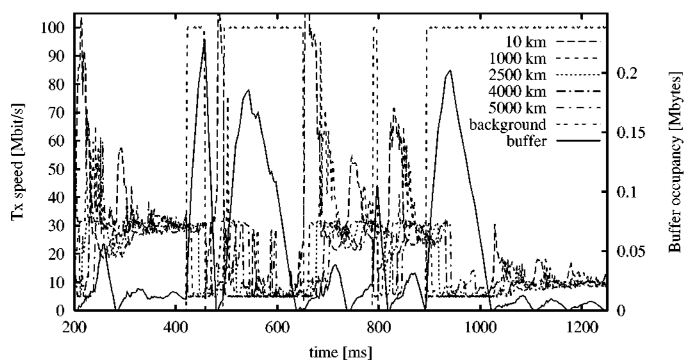

Fig. 8. Transmission rate and bottleneck buffer size. ERICA+ controller with five connections at different distances and ON–OFF background traffic.

figures as the previous charts, when ON–OFF background traffic modifies the available link capacity for the controlled connections. ON and OFF periods are exponentially distributed with average 100 ms; the ON transmission rate is 100 Mb/s. The figure reports 1 s of network operation after the initial transient is exhausted. The controller behavior is smooth and satisfactory, with connections receiving the same amount of resources, apart from small differences during transients due to the different time constant of connections with different length. Notice that these transient unfairnesses compensate one another since the shortest connections always react faster, both in increasing and decreasing the transmission rate.

In Fig. 8 we report, for comparison purposes, the results we obtained with the ERICA+ algorithm. The reaction of this algorithm to traffic changes is much faster than the one we propose;

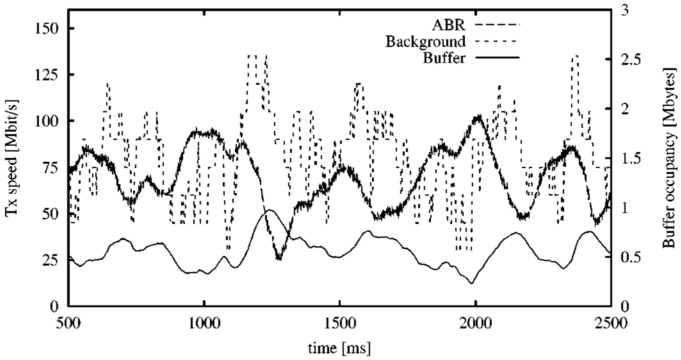

Fig. 9. Total transmission rate and bottleneck buffer size. PID controller with 100 connections at different distances and ten ON–OFF background superimposed sources.

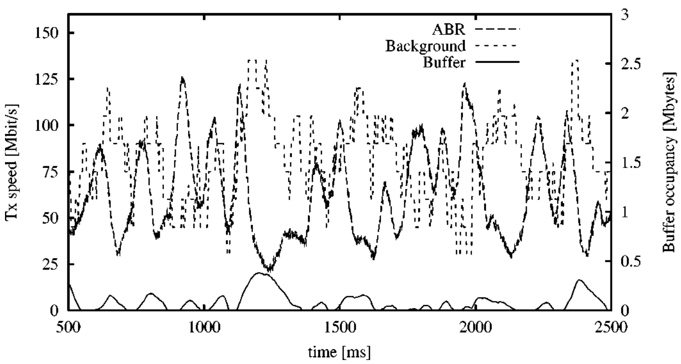

Fig. 10. Total transmission rate and bottleneck buffer size. ERICA+ controller with 100 connections at different distances and ten ON–OFF background superimposed sources.

however, this results in somewhat chaotic behavior of the connections. One advantage is the smaller buffer occupancy, though during transients the buffer can be empty, hinting to possible underutilization of resources. Moreover, every time the background goes from ON to OFF, the shortest connection transmission rate has spikes exceeding 100 Mb/s, which is not desirable.

Let us now increase the number of ABR connections to 100; their lengths are evenly distributed between 50 and 5000 km. The background traffic is obtained as a superposition of ten ON–OFF sources with the same ON and OFF periods distribution as before, but with the ON transmission rate equal to 15 Mb/s. The total background transmission speed ranges from 0 to 150 Mb/s, with an average of 75 Mb/s. Fig. 9 reports the results for the proposed algorithm, while Fig. 10 reports the results obtained with ERICA+ for comparison. The plots report the queue level and the total ABR and background measured at the bottleneck.

In order to assess the stability and performance of the proposed algorithm in a more realistic scenario, we have implemented the controller in all the interfaces of a four-node parking-lot topology. The experimental setup is shown in Fig. 11. As for results shown before, all link capacities are 150 Mb/s, while the length of links adducing from nodes to sources and receivers are only a few kilometers long. Two sources are attached to nodes A, B, and C, while all receivers are connected to node D, as is usual in parking-lot topologies. An ON–OFF background is active between nodes B and C.

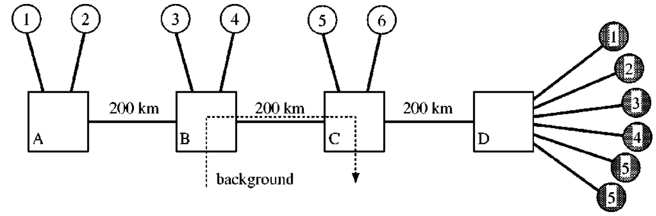

Fig. 11. Parking-lot topology.

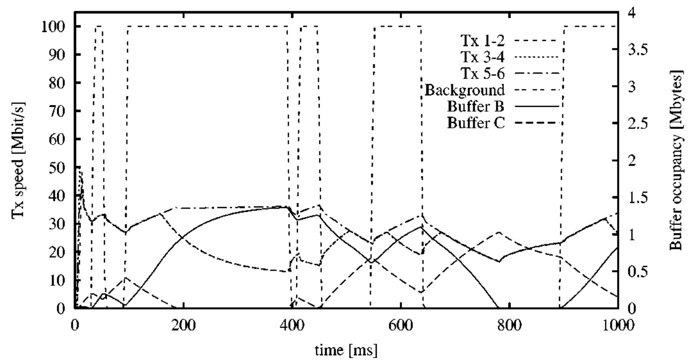

Fig. 12. Parking-lot topology. Transmission rate of the six connections and buffer level at node B and C with ON–OFF background between nodes B and C.

Fig. 12 presents the results. When the background is active, connections 1–4 are throttled in node B, hence, connections from node C can exploit the additional resources left free. The figure reports the transmission rates of the six connections (only one curve is shown for each pair of connections sharing the same input node, since the performance is identical) and the buffer level at nodes B and C (the buffer in node A is always empty). As foreseen by the theory, when the background is OFF, only the controller in node C drives the performance of connections. When the background is active, instead, the controller in B drives the connections upstream (1–4) and the controller in C allows connections 5 and 6 to receive more resources approaching the max–min criterion as time elapses. Buffers in nodes B and C fill up and empty depending on where the bottleneck is located.

VI. Conclusion

This paper presented a control-theoretical approach to network congestion control that privileges stable and robust network operation. In view of the fragility of optimal control when the control delay is unknown or time-varying, we have proposed a control scheme based on the classical PID controller which is robust versus control-delay variations. Such a scheme, thought for the single-link case, applies to the multisource when a properly defined equivalent single source is defined.

The proposed control was implemented on top of the ATM ABR transfer capability, in order to verify its performance in a realistic scenario with the approximation introduced by implementations and protocol constraints. Results are shown and compared with the commonly used algorithm ERICA+, showing that a robust control based on sound theory can achieve performance that, in some respects, is superior to that of performance-oriented heuristic algorithms. Our final conjecture is that heuristic and theoretical approaches can be combined in order to obtain high-performance yet robust controllers, especially in view of developing new congestion-control schemes for the Internet.

APPENDIX

In this proof, we often use the change of variable

$$\begin{cases} x(\theta) = \cot(\theta/2), & 0 < \theta < \pi \\ \theta(x) = 2\,\mathrm{acot}(x), & 0 < x < \infty. \end{cases}$$

Since there exists a one-to-one correspondence between x and θ in the sequel we will exchange x and θ according to the above relation. For $\theta \to 0$, $F(j\theta) \to -\infty^2 - j\infty$, then there exists θ^* such that for $0 < \theta \leq \theta^*$, $F(j\theta)$ is in the third open sector (i.e., $\mathrm{Re}[F(j\theta)] < 0$ and $\mathrm{Im}[F(j\theta)] < 0$). $F(j\theta^*)$ is the point C in Fig. 4. Note that $\angle F(j\theta^*) > -\pi$.

Now assume that the single-source loop with delay m is stable. This means, according to Theorem 1, that the Nyquist plot of $F(j\theta)$ does not encircle the point $-1 + j0$ and that the unique intersection of the Nyquist plot with the unit circle, say, $F(j\overline{\theta})$, must be on the lower half plane ($\mathrm{Im}[F(j\overline{\theta})] < 0$). The point $F(j\overline{\theta})$ is denoted by A in Fig. 4.

To prove the Proposition, we need to show that $\forall\, \theta^* \leq \theta \leq \overline{\theta} \mathrm{Im}[F(j\theta)] < 0$. By contradiction, assume that this is not the case. Denote by $\hat{\theta}$ the real-axis crossing point, that is, the smallest value of θ such that $\mathrm{Im}[F(j\theta)] = 0$. Since $|F(j\hat{\theta})| > 1$, there are two cases: $\mathrm{Re}[F(j\hat{\theta})] > 1$ or $\mathrm{Re}[F(j\hat{\theta})] < 1$. The former is not possible (this roughly means that the Nyquist plot cannot reach the upper plane outside the unit disk, by crossing on the right). Indeed, in this case we would have $\angle F(j\hat{\theta}) = 0$, while $F(j\theta)$ has negative phase. Thus, consider the second case, that is, $\mathrm{Im}[F(j\hat{\theta})] = 0$ with $\mathrm{Re}[F(j\hat{\theta})] < 1$. ($F(j\hat{\theta}) = B$ in Fig. 4.) Then $\angle F(j\hat{\theta}) = -\pi$. Summarizing, we have the following conditions for $0 < \theta^* < \hat{\theta} < \overline{\theta} < \pi$:

$$\begin{cases} \angle F(j0^+) = -\pi \\ \angle F(j\theta^*) > -\pi \\ \angle F(j\hat{\theta}) = -\pi \\ \angle F(j\hat{\theta}) > -\pi \\ \angle F(j\pi) = (m+1)\pi. \end{cases}$$

Since the phase function is continuous on $(0, \pi]$ the above conditions imply the existence of *at least* three local extrema (precisely, at least two maxima and one minimum). Constructing the phase function, we see a contradiction. From

$$F(j\theta) = [(p - rx^2) - j(qx)]e^{-(m+1)\theta}$$

we have

$$\phi(x) = \angle F(j\theta(x)) = \arctan\left(\frac{p-rx^2}{qx}\right) - \frac{\pi}{2} - 2(m+1)\,\mathrm{acot}(x)$$

which is continuously differentiable on $0 < x < \infty$ and its derivative is given by

$$\phi'(x) = \frac{-(1+x^2)q(q+rx^2) + 2(m+1)[(qx)^2 + (p-rx^2)^2]}{[(qx)^2 + (p-rx^2)^2][1+x^2]}.$$

The numerator of ϕ' is a biquadratic function of x. This means that $\phi'(x) = 0$ has at most two solutions for $0 < x < \infty$. This is in contradiction with the fact that $\phi(x)$ has at least three local extrema. □

ACKNOWLEDGMENT

The authors would like to thank Prof. M. Ajmone Marsan and Prof. M. Sznaier for helpful discussions and Dr. M. Munafò for the development, maintenance, and useful usage hints of CLASS.

REFERENCES

[1] M. De Prycker, *Asynchronous Transfer Mode: Solution for Broadband ISDN*. Englewood Cliffs, NJ: Prentice-Hall, 1995.

[2] R. Braden, D. Clark, and S. Shenker, "Integrated services in the Internet architecture: An overview," Internet Engineering Task Force (IETF), RFC 1633, June 1994.

[3] S. Blake *et al.*, "An architecture for differentiated services," Internet Engineering Task Force (IETF), RFC 2475, Dec. 1998.

[4] L. Benmohamed and S. M. Meerkov, "Feedback control of congestion in packet switching networks: The case of a single congested node," *IEEE/ACM Trans. Networking*, vol. 1, pp. 693–708, Dec. 1993.

[5] E. Altman, T. Basar, and R. Srikant, "Control methods for communication networks," in *Proc. 36th Conf. Decision and Control*, San Diego, CA, 1997, pp. TA 31774–1809, TM3 2368–2404, TP3 2903–2945.

[6] S. H. Low and D. E. Lapsey, "Optimization flow control—I: Basic algorithms and convergence," *IEEE/ACM Trans. Networking*, vol. 7, pp. 861–874, Dec. 1999.

[7] S. Floyd and K. Fall, "Promoting the use of end-to-end congestion control in the Internet," *IEEE/ACM Trans. Networking*, vol. 7, pp. 458–472, Aug. 1999.

[8] S. J. Golestani and S. Bhattacharyya, "A class of end-to-end congestion-control algorithms for the Internet," in *Proc. Int. Conf. Network Protocols (ICNP)*, Austin, TX, Oct. 1998, pp. 137–150.

[9] A. Kolarov and G. Ramamurthy, "A control theoretic Approach to the design of explicit rate controller for ABR service," *IEEE/ACM Trans. Networking*, vol. 7, pp. 781–753, Oct. 1999.

[10] N. Ghani and J. W. Mark, "Enhanced distributed explicit rate allocation for ABR services in ATM networks," *IEEE/ACM Trans. Networking*, vol. 7, pp. 710–723, Oct. 1999.

[11] S. Kalyanaraman, R. Jain, R. Goyal, S. Fahmy, and B. Vandalore, "The ERICA switch algorithm for ABR traffic management in ATM networks," *IEEE/ACM Trans. Networking*, vol. 8, pp. 87–98, Feb. 2000.

[12] T. V. Lakshman, P. P. Mishra, and K. K. Ramakrishan, "Transporting compressed video over ATM networks with explicit-rate feedback control," *IEEE/ACM Trans. Networking*, vol. 8, pp. 71–86, Feb. 2000.

[13] R. Muthukrishnan, S. Dasgupta, A. Varma, L. Kalampoukas, and K. K. Ramakrishnan, "Design, implementation and evaluation of an explicit rate allocation algorithm in an ATM switch," in *Proc. IEEE INFOCOM*, Tel Aviv, Israel, Apr. 2000, pp. 1313–1322.

[14] Y. Zhao, S. Q. Li, and S. Sigarto, "A linear dynamic model for design of stable explicit-rate ABR control schemes," in *Proc. IEEE INFOCOM*, Kobe, Japan, Apr. 1997, pp. 283–292.

[15] S. Mascolo, D. Cavendish, and M. Gerla, "ATM rate-based congestion control using a Smith predictor: An EPRCA implementation," in *Proc. IEEE INFOCOM'96*, San Francisco, CA, 1996, pp. 569–576.

[16] "Special issue on control methods for communication networks," *Automat.*, vol. 35, no. 12, Dec. 1999.

[17] "Special selection on networks and control," *IEEE Control Syst. Mag.*, vol. 21, no. 1, Feb. 2001.

[18] K. J. Aström and T. Hägglund, *PID Controllers: Theory, Design and Tuning*. Research Triangle Park, NC: Instrument Soc. Amer., 1995.

[19] F. Blanchini, R. Lo Cigno, and R. Tempo, "A robust feedback strategy for rate control in ATM networks," in *Proc. 1998 Amer. Control Conf.*, Philadelphia, PA, June 1998, pp. 2852–2857.

[20] M. A. Dahleh and I. J. Diaz Bobillo, *Control of Uncertain Systems: A Linear Programming Approach*. Englewood Cliffs, NJ: Prentice-Hall, 1995.

[21] S. P. Boyd and C. H. Barrat, *Linear Controller Design, Limits of Performance*. Englewood Cliffs, NJ: Prentice-Hall, 1991.

[22] K. Zhou, J. C. Doyle, and K. Glover, *Robust and Optimal Control*. Englewood Cliffs, NJ: Prentice-Hall, 1996.

[23] L. H. Keel and S. P. Bhattacharyya, "Robust, fragile, or optimal?," *IEEE Trans. Automat. Contr.*, vol. 42, pp. 1098–1105, Aug. 1997.

[24] J. C. Doyle, "Guaranteed margins for LQG regulators," *IEEE Trans. Automat. Contr.*, vol. AC-23, pp. 756–757, 1978.

[25] H. Özbay, S. Kalyanaraman, and A. Îftar, "On rate-based congestion control in high-speed networks: Design of an H^∞ based flow controller for single bottleneck," in *Proc. American Control Conf.*, Philadelphia, PA, June 1998, pp. 2381–2385.
[26] B. R. Barmish, *New Tools for Robustness of Linear Systems*. New York: MacMillan, 1994.
[27] G. F. Franklin, J. D. Powell, and M. L. Workman, *Digital Control Systems*. Reading, MA: Addison-Wesley, 1990.
[28] "ATM forum traffic management specification, Ver. 4.0," ATM Forum, TM-0056.000, Apr. 1996.
[29] R. Jain, S. Kalyanaraman, S. Fahmy, R. Goyal, and R. Viswanathan, "ERICA switch algorithm: A complete description," ATM Forum, TM 96-1172, Aug. 1996.
[30] M. Gerla, W. Weng, and R. Lo Cigno, "Bandwidth feedback control of TCP and real time sources in the Internet," in *Proc. IEEE Globecom, Session HS-7, High Speed Networks Symp.*, San Francisco, CA, Nov./Dec. 2000.
[31] M. Ajmone Marsan, A. Bianco, T. V. Do, L. Jereb, R. Lo Cigno, and M. Munafò, "ATM simulation with CLASS," in *Performance Modeling Tools, Performance Evaluation*. New York: Elsevier Science, 1995, vol. 24, pp. 137–159.

Franco Blanchini (M'92) was born on December 29, 1959, in Legnano, Italy. He received the Laurea degree in electrical engineering from the University of Trieste, Trieste, Italy, in 1984.

In 1985, he was a Lecturer in numerical analysis with the Faculty of Science, University of Udine, Udine, Italy, and was a Research Associate in system theory from 1986 to 1991. In 1992, he became an Associate Professor of automatic control with the Faculty of Engineering, University of Udine. He has been a Full Professor of automatic control since November 2000. He is also affiliated with the Department of Mathematics and Computer Science, University of Udine, and is the Director of the Laboratory of System Dynamics of that Department. He is currently an Associate Editor of *Automatica*. His previous research activity was in numerical methods for analysis and synthesis of linear systems and in the theory of generalized linear systems. Currently, his main research activity is in the field of robust control, especially Lyapunov methods and $\mathcal{L}^1$ control theory. His research interests also include the control of constrained systems, mechanical systems, and the control of distribution networks.

Dr. Blanchini was a Member of the Program Committee of the IEEE Conference on Decision and Control in 1997, 1999, and 2001.

Renato Lo Cigno (A'94) was born in Ivrea, Italy, in 1963. He received the Dr. Ing. degree in electronic engineering from the Politecnico di Torino, Torino, Italy, in 1988.

He is currently an Assistant Professor with the Department of Electronics, Politecnico di Torino. Since 1988, he has been with the Telecommunication Research Group of the Department of Electronics, first as a Research Engineer and then as an Assistant Professor. From June 1998 to February 1999, he was with the Department of Computer Science, University of California at Los Angeles, as a Visiting Scholar. He is the coauthor of about 80 journal and conference papers in the area of communication networks and systems. His current research interests are in performance evaluation of wired and wireless networks, modeling and simulation techniques, and flow and congestion control.

Dr. Lo Cigno has been a Member of the Program Committees of the IEEE Globecom and the IEEE International Conference on Network Protocol.

Roberto Tempo (M'90–SM'98–F'00) was born in Cuorgné, Italy, in 1956. In 1980, he graduated in electrical engineering from the Politecnico di Torino, Torino, Italy.

From 1981 to 1983, he was with the Department of Automation and Information, Politecnico di Torino. In 1984, he joined the National Research Council of Italy (CNR) at the research institute IRITI, Torino, where he has been a Director of Research of Systems and Computer Engineering since 1991 and an elected Member of the Scientific Council. He has held visiting and research positions at the University of Illinois at Urbana-Champaign, the German Aerospace Research Organization, Oberpfaffenhofen, Germany, and Columbia University, New York. He has been an Associate Editor of *Systems and Control Letters* and is currently an Editor of *Automatica*. His research activities are mainly focused on robustness analysis and control of uncertain systems and identification of complex systems subject to bounded errors.

Dr. Tempo received the Outstanding Paper Prize Award in 1993 from the International Federation of Automatic Control (IFAC) for a paper published in the IFAC journal *Automatica*. He has been an Associate Editor of the IEEE TRANSACTIONS ON AUTOMATIC CONTROL. He is Vice-President for Conference Activities of the IEEE Control Systems Society. He is also a member of the European Union Control Association (EUCA) Council.