

Explicit Rate Flow Control for ABR Services in ATM Networks

Ching-Fong Su, *Member, IEEE*, Gustavo de Veciana, *Member, IEEE*, and Jean Walrand, *Fellow, IEEE*

Abstract—We propose a novel explicit rate-flow-control algorithm intended for available-bit-rate (ABR) service on an ATM network subject to loss and fairness constraints. The goal is to guarantee low cell loss in order to avoid throughput collapse due to retransmission by higher level protocols. The mechanism draws on measuring the current queue length and bandwidth availability, as well as tracking the current number of active sessions contending for capacity, to adjust an explicit bound on the source transmission rates. We identify the factors that affect queue overflows and propose simple design rules aimed at achieving transmission with controlled loss in a dynamic environment. We also discuss how conservative design rules might be relaxed by accounting for statistical multiplexing in bandwidth sharing among bursty ABR sources and variable-bit-rate (VBR) sources.

Index Terms—ABR service, ATM networks, delay differential equations, explicit rate flow control.

I. INTRODUCTION

ASYNCHRONOUS transfer mode (ATM) networks are geared at supporting and integrating a variety of communication services. These services may be divided into those based on *reservation*, e.g., constant- and variable-bit-rate (CBR, VBR) services, and *best effort* services, such as unspecified-bit-rate (UBR) and to some extent available-bit-rate (ABR) services. Among the latter, ABR service promises to play an important role in supporting high-bandwidth data as well as Internet traffic, such as TCP traffic. The rationale for including ABR is to provide an economical and flexible way to carry data traffic, as might be needed to simplify adoption of ATM to support delay adaptive real-time applications [6]. From the service provider's point of view, ABR traffic can be used to enhance utilization by directing sources to make the most of the network's available capacity subject to minimum cell rate and cell loss guarantees.

It has been shown in [22], [23] that TCP performs poorly over ATM networks when congestion leads to lost ATM cells which in turn result in multiple corrupted IP packets. Significant performance degradation results from corrupted IP packets since

they waste network bandwidth because they are useless upon arriving at their destination hosts, and trigger retransmission by sources making the effective throughput even lower. Several algorithms have been proposed for dealing with this throughput collapse problem, including packet discard strategies and tuning of the TCP flow control mechanism [23]. However, instead of tuning the concurrent feedback loops of TCP and ABR, we believe that it is important to ensure low cell loss inside the network, so as to avoid the throughput collapse altogether. One of the aims of this work is to design and analyze a rate adaptation mechanism for ABR service which permits controlling losses.

ABR service is likely to be based on rate-based feedback flow control,¹ i.e., adjusting the transmission rates of sources based on the current network state. There are two types of rate control mechanisms. The network can either determine and enforce a bound on the transmission rate for each ABR session or rely on exchanging minimal (binary) congestion indications to incrementally adjust source transmission rates; for a survey see [6], [16], [21] and for a representative analysis see [7]. These two mechanisms are not incompatible, and in fact future networks might use a natural combination of binary feedback adjustments with explicit rate bounds to adjust source transmission rates, e.g., the proportional rate control algorithm (PRCA) and enhanced PRCA (EPRCA) discussed in [6]. However, slow adaptation to the network's state and oscillating queue (traffic flow) are two problems with the algorithms using single bit feedback, see for example [7]. By contrast, using explicit rate feedback allows switches to specify a desirable traffic rate, so sources can rapidly adapt their traffic transmissions. In this paper we consider ABR flow control based on explicit rate feedback.

Several algorithms have been suggested for computing explicit rates. In general, these algorithms are of two types: the computation is based on the queue length and/or the arrival rate to the queue. Algorithms based on monitoring the arrival rate to the system may, for example, compute the explicit rate by dividing the available capacity among ongoing sessions in a "fair" manner. In this case, typically switches maintain rate/state information for each session, see for example [2], [9], [19]. The computational complexity incurred by such per-source accounting is however an issue in implementation. Rather than computing the fair rate allocation to sources by division, other algorithms attempt to *estimate* it, see for example [10], [17], [24]. In most cases these algorithms are not shown to be asymptotically stable in steady-state, nor are the the resources required (bandwidth and buffer) to avoid losses during transient overloads quanti-

Manuscript received June 5, 1997; revised May 7, 1999 and January 20, 2000; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor J. N. Daigle. This work was supported by a National Science Foundation RIA Grant NCR-9409722, Career Grant NCR-9624230, and Southwestern Bell Corporation/TRC.

C.-F. Su is with Fujitsu Laboratories of America, Inc., Sunnyvale, CA 94086 USA (e-mail: csu@fla.fujitsu.com).

G. de Veciana is with the Electrical and Computer Engineering Department, University of Texas, Austin, TX 78712 USA (e-mail: gustavo@ece.utexas.edu).

J. Walrand is with the Department of Electrical Engineering and Computer Science, University of California, Berkeley, CA 94720 (e-mail: wlr@eecs.berkeley.edu).

Publisher Item Identifier S 1063-6692(00)04999-2.

¹Rate-based feedback flow control was chosen by ATM Forum in 1995.

fied. Charny *et al.* [9] proposed an algorithm using arrival rate information and proved its stability. However in their scheme, stability is achieved because sources are required to delay rate increases by approximately twice the maximum round-trip time. Unfortunately, this scheme may cause significant transient link underutilization.

Alternatively, one can use queue length information at bottleneck links to compute the explicit rates, see for example [3], [5], [11], [12], [20]. In these algorithms, feedback control can sometimes be modeled by differential equations whose asymptotic stability can be investigated. However, in order to obtain a simple feedback control model, it is usually assumed that traffic arrival rates are equal to possibly delayed explicit rate advertised by the network. These works do not capture bursty source traffic or constrained source dynamics.

This paper extends an approach first proposed in [8] which was inspired from [14]. The proposed explicit rate computation is based on the available capacity, the queue's state, and an estimate of the number of active ongoing sessions. Unlike the aforementioned algorithms, our algorithm and model capture constraints on source rate increases and available capacity drops.

Feedback control in the context of wide and even local area networks is plagued by the potentially large source transmission rates relative to the propagation (as well as processing and queuing) delays in the system, making the responsiveness of such mechanisms sluggish. In general, when ABR sessions have relatively long bursts of traffic to send, one might hope to have enough time to properly adapt their transmission rates—such sessions are said to be *greedy*. By contrast, traffic with small burst sizes relative to the control time scales, e.g., some WWW connections, are said to be *bursty* sessions. Feedback control would typically be ineffective for individual bursty connections though it might still work reasonably well on an aggregated basis. In practice one might expect a mix of traffic with various burst scale properties to use ABR service, and it is of interest to understand the impact that both the unpredictable nature of source transmissions and spare capacity in the network will have on flow control. Another novelty of this work is a worst-case analysis of resource requirements to achieve loss-free transmission when sources are bursty.

The paper is organized as follows. In Section II we consider a simple model for a flow control mechanism which accounts for such fluctuations. In Section III we analyze the asymptotic stability of the proposed model and show that the lossless guarantee can be met by reserving a minimum capacity and buffer in the bottleneck link. In principle one might argue that by statistically multiplexing a large number of *bursty* ABR sessions on a given link, one can achieve relatively high utilization. In Section IV we articulate this point of view and suggest how one might hope to optimistically use this to deal with bursty ABR sessions. We also briefly discuss a possible implementation of the proposed mechanism and its relation to standardized mechanism and system parameters. The balance of the paper includes preliminary simulations of the proposed algorithm in Section V, and concluding remarks in Section VI.

II. EXPLICIT RATE FLOW CONTROL—A FLUID MODEL

In this section we describe and motivate an explicit rate control mechanism based on a simplified “fluid model.” That is, we assume that the instantaneous transmission rates of sources and links are well defined. In practice such rates would correspond to windowed estimates of the cell rates in the system. The key elements and assumptions underlying the model are subdivided into the network behavior, the source behavior and the rate control mechanism. A discussion of how our idealized model's parameters translate to practice as well as related ATM Forum standardized system parameters is included in Section IV. Note, however, that for ease of discussion we will assume the minimum cell rates (MCR) of ABR sessions are set to zero. If the MCR's are positive, we can reserve the MCR, m_j for each session and make the worst-case assumption that a session will attempt to send at a rate no smaller than its reserved MCR. In other words, the available bandwidth, the explicit rate, and the transmission rate in the following discussion are associated with the transmissions in excess of the reserved minimum cell rates.

Network Model: Three key features characterize our network model. First, rather than attempting to model the overall network dynamics, we consider a single “bottleneck” buffered link shared by at most n_{\max} concurrent ABR sessions. This simplifying assumption, which is often used in the literature, e.g., [3], [7], renders the problem tractable and provides quite a bit of insight into the network dynamics. A notable exception for a work considering a full network level analysis is [4]. In Section III-C, we consider the impact of multiple bottleneck links on the steady-state behavior of the network.

Second, we shall assume that the worst-case delay, including both propagation and queuing time, from the j th source to the bottleneck is τ_j^f and then back is τ_j^b for a total round-trip delay of τ_j . Let $\tau = \max_{j=1}^{n_{\max}} \tau_j$ and $\Delta\tau = \max_{j=1}^{n_{\max}} [\tau - \tau_j]$ denote the worst-case round-trip delay and worst-case delay discrepancy respectively. We follow [3], [7] in assuming worst-case delays per source, however note that in practice, such feedback delays are likely to fluctuate, thus in Section III-B we shall re-examine the robustness of this modeling choice.

The bottleneck model is shown in Fig. 1, where $c(t)$ denotes the instantaneous capacity available at a bottleneck link. Our third and final assumption is that the rate at which the available capacity can decrease is lower-bounded, i.e., $dc(t)/dt \geq -\rho$. Changes in $c(t)$ are primarily due to fluctuations in the aggregate bandwidth requirements of current reserved services, e.g., VBR/CBR sessions sharing the link, as well as changes in the *number* of such sessions. The impact of statistical multiplexing among VBR sessions and the changes of the number of ongoing connections are further discussed in Section IV.

Source Model: Let $r_j(t)$ denote the instantaneous transmission rate for the j th session at time t . Throughout its lifetime a source's transmission rate can never exceed the most current explicit rate indication $e(t)$ received from the network, i.e., $r_j(t) \leq e(t - \tau_j^b)$. We introduce a threshold r^* to discriminate among sources with different “activity” levels as follows:

- If $e(t - \tau_j^b) \geq r^*$, we say a session is “on” if its current transmission rate exceeds r^* , i.e., $r_j(t) > r^*$, otherwise




Fig. 1. Network bottleneck model.

the session is said to be “off.” Sessions which are “on” are contending for available capacity.

- If $e(t - \tau_j^b) < r^*$, then the available capacity of the link is low, that is, the link appears to be congested and all sessions are considered to be “on.”

We shall also assume that once a session’s transmission rate exceeds r^* , at most a linear rate of increase, g , can be supported. Fig. 2 shows the characteristics of the source transmission rate as discussed above.

Our source model captures a possibly desirable mechanism wherein idle sources can jump-start their transmissions up to an initial cell rate r^* and may thereafter ramp up linearly. Note that sources desiring to transmit at a rate below the threshold r^* may do so freely, which should expedite short bursty transmissions. By contrast a persistent session wishing to transmit at a high rate may certainly do so but must give the network time to detect that it is becoming a major contender for capacity in the network, hence the ramp-up above r^* is constrained.

Flow Control Mechanism: The dynamics of the bottleneck queue are given by

$$\dot{q}(t) = \begin{cases} \sum_{j=1}^{n_{\max}} r_j(t - \tau_j^f) - c(t), & q(t) > 0 \\ \left[\sum_{j=1}^{n_{\max}} r_j(t - \tau_j^f) - c(t) \right]^+, & q(t) = 0 \end{cases} \quad (1)$$

with $[x]^+ = \max[x, 0]$ and where the transmission rate of each session is bounded by the latest explicit rate received, i.e., $r_j(t) \leq e(t - \tau_j^b)$.

The explicit rate $e(t)$ is computed based on the current network state which includes the queue length, $q(t)$, the current available capacity, $c(t)$, and *delayed information* about the current number of sources that are “on.” The number of “on” sources is given by

$$\hat{n}(t) = \sum_{j=1}^{n_{\max}} 1\{r_j(t - \tau_j^f) \geq r^*\} + 1\{e(t - \tau_j) < r^*\} \quad (2)$$

where $1\{\cdot\}$ is the indicator function. The first term corresponds to the sources with transmission rates exceeding r^* , while the second term corresponds to the case where the bottleneck link appeared to be congested, indicated by $e(t - \tau_j) < r^*$. Recall that in this second case a session is assumed to be “on” regardless of its rate.




Fig. 2. Source model characteristics.

The explicit rate is computed so that the network approximately tracks a delay-free reference model for the queue dynamics given by

$$\dot{q}(t) = f(q(t)) \quad \text{e.g., } f(q(t)) = -k(q(t) - q^*),$$

where the drift $f(q(t))$ is selected to drive the reference queue toward a target level q^* . For the remainder of this paper, unless explicitly mentioned, we assume the drift function is linear, i.e., $f(q(t)) = -k(q(t) - q^*)$, where $k > 0$. The bottleneck link computes $e(t)$ so as to approximate this drift assuming the sources that were “on” [with at least one being on, i.e., $\hat{n}(t) \vee 1$] will transmit at this new rate, i.e.,

$$f(q(t)) = e(t)(\hat{n}(t) \vee 1) - c(t) \Rightarrow e(t) = \frac{f(q(t)) + c(t)}{\hat{n}(t) \vee 1} \quad (3)$$

Thus the explicit rate is based on the available capacity as well as the queue’s state, which as suggested in [1] are necessary to ensure stability. Note that a single $e(t)$ is computed for all ABR sessions carried by the bottleneck link, which significantly reduces the implementation complexity of this algorithm.

To summarize, we highlight several novel characteristics of both the model and flow control mechanism we have proposed. First, we explicitly model the drops of available capacity for ABR service, through a deterministic bound $-\rho$ on the rate of decrease. Next we introduce an on/off threshold rate r^* under which sources can burst, and above which they can ramp up their transmissions at a rate not exceeding g . Finally we propose an explicit rate control mechanism which, based on estimating the number of active sources, attempts to bring the queue to a target level q^* .

III. ANALYSIS OF PROPOSED MODEL

We shall investigate both the worst-case transient behavior as well as asymptotic stability of the proposed rate control mechanism. On one hand, we show that by exerting call admission,

i.e., ensuring no more than n_{\max} sources are present, and by appropriately provisioning the system, we can guarantee no buffer overflows, even with bursty source transmissions. On the other hand, when sources are greedy, we show conditions for asymptotic stability of the system, and characterize the steady-state regime.

A. Transient Behavior: Guaranteeing No Loss and Positive $e(t)$

Below we will show that by provisioning the potential bottleneck link with a buffer of size b_{\min} one can ensure no loss. In addition, to guarantee that the proposed rate control mechanism is well defined, i.e., $e(t) \geq 0$, a guaranteed minimal service rate c_{\min} needs to be reserved. These two quantities are given in terms of the system parameters as follows:

$$\boxed{b_{\min} = q^* + \frac{w}{k} + (kq^* + w)\tau}$$

$$\boxed{c_{\min} = k(b_{\min} - q^*)} \quad (4)$$

where $w = \tau\rho + n_{\max}[r^* + g\tau]$.

To obtain these bounds we investigate the transient behavior, or overshoots, of the bottleneck queue dynamics. The following lemma, proved in Appendix I, gives a bound on the aggregate arrival rate to that queue.

Lemma 1: The aggregate arrival rate to the queue at time t is bounded by

$$\sum_{j=1}^{n_{\max}} r_j(t - \tau_j^f) \leq f(q(t - \tau)) + c(t - \tau) + n_{\max}[r^* + g\tau]. \quad (5)$$

Intuitively the term $n_{\max}[r^* + g\tau]$ accounts for the transient bursts caused by sessions turning “on” between $t - \tau$ and t , whereas the remaining terms account for the mismatch due to feedback delays in the computed and desired explicit rates for the sources. Note that the feedback mechanism guarantees that the arrival rates of bursting sessions will be regulated no later than τ seconds after the bottleneck link detects them.

Using (5) and (1), as well as the variability constraint on the available capacity, we can determine the following differential inequality:

$$\begin{aligned} \dot{q}(t) &\leq f(q(t - \tau)) + c(t - \tau) - c(t) + n_{\max}[r^* + g\tau] \\ &\leq f(q(t - \tau)) + \tau\rho + n_{\max}[r^* + g\tau] \\ &\leq f(q(t - \tau)) + w \end{aligned} \quad (6)$$

where $w = \tau\rho + n_{\max}[r^* + g\tau]$. Note that ρ is the variation of available capacity, or equivalently the burstiness of VBR connections sharing the same link. In the worst case, $\tau\rho$ corresponds to the maximum variation in the arrival rate of VBR traffic over τ seconds. Thus the worst-case queue growth is driven by three factors: 1) the drift computed from delayed queue information; 2) the *unexpected* burstiness of VBR connections sharing the link; and 3) the *unexpected* bursting of ABR connections.

The bound in (6) is conservative as it is based on the assumption that capacity is dropping by $\rho\tau$ and n_{\max} sessions are

turning “on” at the same time. Nevertheless, it can be used to derive the following upper bound on the queue length. (See Appendix I.)

Lemma 2: Assuming the queue is empty at the beginning, its queue length is thereafter upper-bounded by $q_{\max} = q^* + (w/k) + (kq^* + w)\tau$.

Thus in order to guarantee that no loss occurs, we need only reserve a buffer of size exceeding q_{\max} as given in (4). As mentioned earlier, to guarantee a nonnegative $e(t)$, we should reserve a minimum capacity c_{\min} . In particular, to ensure $e(t) \geq 0$, we require that $f(q(t)) + c(t) \geq 0$. Since $f(q(t))$ is minimized when the queue is at a maximum, a capacity $c_{\min} = -f(q_{\max}) = k(b_{\min} - q^*)$ is sufficient to ensure that $e(t)$ is nonnegative. This translates to the second requirement shown in (4). Note that by modifying the drift function $f(\cdot)$, we can change the minimum buffer/capacity requirements. This is further discussed in Section IV-C.

B. Asymptotic Stability

Next we consider the asymptotic stability of the proposed algorithm in a *fixed* environment. That is, the number of sessions n and bottleneck capacity $c(t) = c$ are fixed and the sessions are all greedy. Note that greedy sessions track the latest explicit rate indications and thus in the following analysis r^* plays no role in distinguishing source activity—sources are always “on.” We also assume that the bottleneck buffer is large enough so that there is no loss and the bottleneck capacity c is large enough so that the explicit rate $e(t)$ is well-defined—see the previous subsection for discussion. Our goal is to analyze the impact of the round-trip delay, drift function $f(\cdot)$, and the ramp-up constraint on the stability and stationary regime of the system.

For ease of discussion, we shall first relax the ramp-up constraint on the sources and discuss the stability of the system. If we take the derivative of “ $e(t) = (f(q(t)) + c)/n$ ” with respect to t , it follows that $\dot{e}(t) = -k\dot{q}(t)/n$. Now substituting into (1), we find that $e(t)$ is governed by a delay-differential equation given by

$$\dot{e}(t) = -\frac{k}{n} \left[\sum_{i=1}^n e(t - \tau_i) - c \right]. \quad (7)$$

An equivalent system for (7) is shown in Fig. 3, where $G(s) = (1/s)D(s)$, and $D(s) = \sum_{i=1}^n e^{-s\tau_i}$ models the feedback delays.

The model shown in Fig. 3(b) is a linear feedback control system and its stability can be investigated based on the frequency response of $G(s)$ using the Nyquist criterion [25]. Alternatively, one can take a transfer function approach and consider the location of the transfer function’s poles, i.e., the roots of $1 + (k/n)G(s)$, see for example [3], [11], [12]. The following lemma, proved in Appendix I, specifies a condition for asymptotic stability of this system. We note that this result is identical to that in [11].

Lemma 3: For a bottleneck system (see Fig. 3) with n greedy sources and fixed capacity c sufficiently large to ensure $e(t) \geq 0$ but with no source ramp-up constraint, the condition $(k/n) \sum_{i=1}^n \tau_i < 1$ is sufficient to guarantee asymptotic stability.

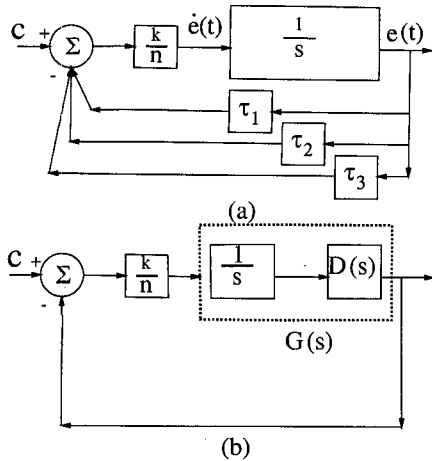


Fig. 3. Control system model and its equivalent.




Fig. 4. Nonlinear model with constraint.

Next we discuss the stability of the proposed mechanism including the source ramp-up constraint. Consider the nonlinear system shown in Fig. 4. It is similar to the one in Fig. 3 except for the nonlinear block ensuring that $\dot{r}(t) = \min[\dot{e}(t), g]$, i.e., the ramp-up of sources is constrained.

Based on a generalized Nyquist criterion—Circle criterion [25, p. 344], we prove the following stability result in Appendix I.

Lemma 4: For a bottleneck system (see Fig. 4) with n greedy sources and fixed capacity c sufficiently large to ensure $e(t) \geq 0$, the condition $(k/n) \sum_{i=1}^n \tau_i < 1$ is sufficient to guarantee asymptotic stability.

Thus our sufficient condition for asymptotic stability is not altered by the ramp-up constraint. Note, however, that the ramp-up constraint g will affect the minimum buffer requirement in Section III-A. Hence we can change g to meet different buffer requirements without affecting the stability. Also note that this condition only depends on the *average* round-trip delay of sessions sharing the bottleneck link. Thus, stability can be maintained, even if sessions have different round-trip delays as long as the *average* does not exceed $1/k$. However, unlike our model, in practice one would expect a session's individual round-trip delay to vary over time. Fortunately our results suggest that stability depends on the overall average round-trip delays seen by sessions sharing a link. This lends credence to the robustness of this result if delays were to vary. In addition, one might argue, that in the limit as the number of sessions on network links grows, a central limit type of effect would ensure that the constraint is met, if it is met for the average round-trip delays.

C. Stationary Regime, Fairness, and Constrained Flows

When the system is asymptotically stable, it will eventually reach a stationary regime. It is thus of interest to establish how bandwidth is allocated among ongoing greedy sessions when the system reaches this steady-state. In steady-state the derivatives of system variables go to zero, i.e.,

$$\dot{e}(\infty) = \lim_{t \rightarrow \infty} \dot{e}(t) = 0 \quad \text{and} \quad \dot{q}(\infty) = \lim_{t \rightarrow \infty} \dot{q}(t) = 0.$$

If $\dot{e}(\infty) = 0$, then from (7) it follows that $ne(\infty) = c$. Hence $f(q(\infty)) = 0$ and $q(\infty) = q^*$. As a result, the system converges to $e(\infty) = c/n$ and $q(\infty) = q^*$. Thus the available capacity at the bottleneck link is partitioned fairly among sources which are “on” and greedy.

The bottleneck system model we have considered heretofore represents a limited view of the interactions among network links. A crude model for the impact that other network links have on the system is to assume that the allocated explicit rates at other links are fixed and further constrain the flows. Therefore, such constrained sessions are not able to send traffic at the allocated explicit rate, which will cause the aggregate arrival rate to be smaller than expected, and the queue to be below target at the bottleneck link being considered. Now, the drift function $f(\cdot)$ in (3) aims to bring the queue length to the target level q^* , so it will compensate for low queue levels by further increasing $e(t)$. The increased $e(t)$ allows other sessions which are unconstrained to send traffic at even higher rates, so the available capacity will not be wasted. In other words, the unused bandwidth of constrained sessions is re-allocated among greedy sessions. Suppose m sessions are constrained by peak rates $p_i > r^*$ elsewhere in the network. In steady-state the following equation would hold:

$$\frac{c + f(q(\infty))}{n} * (n - m) = c - \sum_{i=1}^m p_i.$$

It follows that $f(q(\infty)) = (m/(n - m))c - (n/(n - m)) \sum_{i=1}^m p_i$ and $q(\infty) = [0, q^* - (1/(k(n - m)))(mc - n \sum_{i=1}^m p_i)]^+$, so the steady-state queue length lies between q^* and 0 as a result of constrained traffic. Thus in order to fully re-allocate unused bandwidth, we need q^* to be greater than $(1/(k(n - m)))(mc - n \sum_{i=1}^m p_i)$. This heuristic argument reveals an important property. Specifically it shows that the queue target level q^* is a critical parameter in determining the ability of the control mechanism to support reallocation of available bandwidth among constrained session rates.

IV. IMPLEMENTATION AND DESIGN ISSUES

In this section we first briefly discuss how the somewhat stringent no-loss scenario considered above might be relaxed to ensure infrequent loss while at the same time allowing more efficient use of network resources. Next we discuss a possible implementation of the proposed control mechanism, and its relation to standardized mechanisms and system parameters.

A. ABR Call Admission and Statistical Multiplexing

Feedback control is known to be ineffective in controlling sessions which send small bursts whose durations are shorter than the control time scale. In turn, resources reserved to account for

worst-case bursting of sources, are likely to have a low utilization. Based on (4), when at most n_{\max} sessions are ongoing, we can find the buffer size ensuring no losses will occur. Alternately, given a fixed network buffer size, (4) can also be used to determine an admissible number n_{\max} of ABR sessions such that no losses occur. However the analysis in Section III-A is based on the pessimistic assumption that n_{\max} ABR sessions are perpetually and concurrently “bursting.” Thus it is desirable to consider the benefit of “statistically multiplexing” independent ABR sessions. In particular by relaxing the loss requirement, we can increase the admissible number of sessions from n_{\max} to n_{\max}^b .

An approximate way of doing this would be to control the probability that n_{\max} sessions out of larger pool of n_{\max}^b burst concurrently within τ seconds. As a simple example, suppose we can characterize bursty ABR sessions as alternating between “off” and “on” modes throughout their lifetime and the distribution of their off times are independent and identically modeled by an exponential distribution with a parameter λ . Let I_j be defined as follows:

$$I_j = \begin{cases} 1, & \text{if stream } j \text{ jumps from “off” to “on” within } \tau \\ 0, & \text{otherwise.} \end{cases}$$

Therefore, a *conservative* admissible number of ABR sessions subject to controlled loss would be

$$n_{\max}^b = \max \left\{ n \mid \mathbb{P} \left(\sum_{j=1}^n I_j > n_{\max} \right) < \delta \right\}$$

where δ is a design parameter which roughly characterizes the desired quality of service at the buffer of the bottleneck link. From the exponential distribution, it follows that $\mathbb{P}(I_j = 0) = e^{-\lambda\tau}$, and n_{\max}^b can be determined by a binomial distribution. The additional number of sources $n_{\max}^b - n_{\max}$ that can be admitted on the link is due to statistical multiplexing of ABR bursts. The effectiveness of such multiplexing depends on the average interval between “bursts,” $1/\lambda$, and the round-trip delay τ . Some examples are considered in Section V.

B. Protocol Implementation and Complexity

Feedback rate control depends on regular information exchange between the network and sources. This information is carried in special resource management (RM) cells which are periodically (every Nrm data cells) generated by the source and sent along the session’s route to the destination where they are looped back to the source. The RM cells carry various types of information; of particular interest herein, will be the current cell rate (CCR), the minimum cell rate (MCR), the explicit rate (ER), and a congestion indication (CI) bit, see for example [13].

We envisage a setting where explicit rates are computed at some or all of the switches a session traverses, and the minimum of the computed ER’s is stamped on the returning RM cells. Each source would receive a returned RM cell with an explicit rate, which it would for example add to its reserved MCR to determine its allowed cell rate. This algorithm has the advantage that the same $e(t)$ is computed for all sources sharing a

given link, which significantly reduces the complexity of computing explicit rates and stamping RM cells. *Pseudocode* of the proposed algorithm can be found in Appendix II.

Estimation of $\hat{n}(t)$: The proposed algorithm uses the available capacity $c(t)$, the queue length $q(t)$, and the number of “on” ABR sessions $\hat{n}(t)$ to compute $e(t)$. Notice that, to estimate $\hat{n}(t)$, an assessment of session activity could be done at the source/policer end and the result encoded in the forward RM cells. The switch does not need to monitor the rate/state information for each session, which otherwise could be prohibitively expensive for switches carrying a large number of ABR sessions. In the following, we propose an algorithm for estimating $\hat{n}(t)$ without doing per-source accounting.

Suppose an “on” source’s CCR is fixed, it will send RM cells every Nrm/CCR seconds. Within l seconds, $l * CCR/Nrm$ RM cells are expected to arrive at the bottleneck link. Thus the switch can monitor the RM cell arrivals in a synchronous fashion over fixed-length intervals of l seconds. Assume CCR_i and the status $s_i \in \{0, 1\}$, i.e., “off” or “on,” are carried by RM cell i which arrives at the bottleneck link at time t_i . For the j th interval of measurement, the number of “on” sources can be approximated by

$$\delta n_j = \sum_{i \in L_j} \frac{Nrm}{l * CCR_i} s_i,$$

where

$$L_j = \{i \mid jl < t_i \leq (j+1)l\}. \quad \hat{n}((j+1)l) = \hat{n}(jl) * \alpha + \delta n_j * (1 - \alpha) \quad (8)$$

where α is an averaging factor.²

As a result, $\hat{n}(t)$ is a piece-wise step function, which is continuous from the right-hand side at the points $t = j * l$, $j \in \mathcal{N}$. Fig. 5 shows an example of estimated $\hat{n}(t)$ and real $n(t)$ when a link carries 100 bursty ABR sessions. Clearly $\hat{n}(t)$ can track the number of “on” sessions quite nicely. In general the measurement interval l should be short so as to track $n(t)$ effectively. However l should be long enough such that the bottleneck link can collect at least one RM cell from each “on” session during each measurement period.

C. Design Parameters

From the analysis in Section III-A, the buffer requirement is a function of several system parameters. In the following we discuss the design tradeoffs in selecting these parameters. Those parameters related to source behavior in this paper can be translated to the ABR parameters [13], as shown in Table I.

Drift Function $f(\cdot)$: In Section III-A we considered the linear drift function $f(q(t)) = -k(q(t) - q^*)$ and used the fact that $f(q(t)) \leq kq^*$ to derive an upper-bound on the queue length. In practice one may want to saturate the maximum value of the drift $f(\cdot)$ in order to control b_{\min} to a desirable value, see for example Fig. 6. Moreover, if the minimum value

²Currently, the on/off status bit is not in the standard definition of RM cell fields. As suggested in [26], one can use an alternative estimator $\delta n_j = \sum_{i \in L_j} (Nrm/l * CCR_i)$ without relying on s_i . The key insight is that the “off” sources whose real cell rates are smaller than its allowed cell rate (ACR) (or the CCR in the RM cells) have insignificant contribution to the summation.




Fig. 5. Estimation of $\hat{n}(t)$.

TABLE I
PARAMETERS RELATED TO ABR STANDARD

parameters in Section III-A	ABR standard
$e(t)$	ER
$r(t)$	ACR (CCR)
g	$\frac{PCR}{Nrm \cdot RIF}$
r^*	ICR




Fig. 6. Saturated function $f(\cdot)$.

of $f(\cdot)$ is clamped by f_{min} and $-k(b_{min} - q^*) < f_{min} < 0$, the minimum capacity requirement c_{min} can be reduced.

Ramp-Up Constraint g : In practice source rate adjustments might occur according to the positive feedback CI mechanism using additive rate increase and proportional rate decrease factors (RIF, RDF) discussed in [6], [18]. Note that using this type of proportional rate control mechanism, one can bound a source's growth g by $PCR/Nrm \times RIF$. Thus these parameters can be used to optimize the operation and determine the growth rate that should be assumed in dimensioning of the network resources. Linear ramp-up also facilitates the integration of CI-based and ER-based mechanisms in a heterogeneous environment.

Rate Threshold r^* : This parameter allows the sessions some degree of freedom in sending small bursts. It also captures the impact of such bursts on the network operation. The threshold r^* prevents the switch from incorrectly believing a session sending small bursts is becoming a major contender for spare capacity, thus reducing the overhead of doing unnecessary adjustments on the explicit rate computation. Moreover, r^* can be interpreted as the initial cell rate (ICR) in the ABR framework [6], which limits the initial transmission rate after an idle period.

Capacity Variation ρ : Since ABR sessions might exploit unused bandwidth allocated to VBR sessions sharing the same link, the variation of the aggregate VBR bandwidth requirement will affect the available capacity of ABR sessions. Fig. 7 shows what might be the aggregate bandwidth increment for a fixed

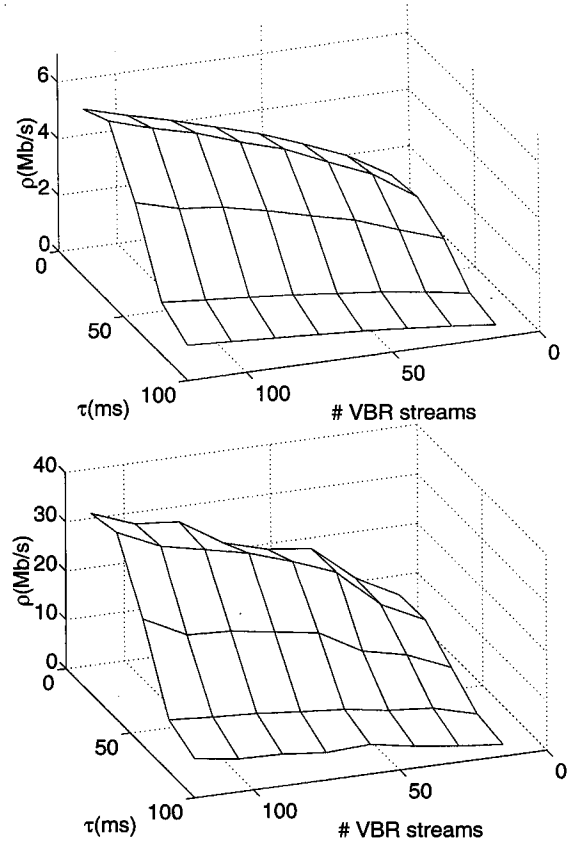


Fig. 7. Average and maximum bandwidth increments for aggregated VBR sources.

number of heterogeneous VBR (MPEG 1 video) sessions over various time scales τ . In particular we approximate ρ by considering consecutive time intervals of size τ and compute the difference in the rate of arrivals over these intervals versus the associated times scale τ .

Measurements of the worst-case and average ρ show that the variability is on the order of 10's of Mb/s/s and highly dependent on the time scale of interest. Note that there is a large discrepancy between the average and worst-case variability, which makes buffer dimensioning difficult. In practice, one might consider the distribution of this quantity, represented by a random variable R , and let ρ satisfy $P(R > \rho) \leq 10^{-6}$, so that the probability of failure for the control is small, and the quality of service is roughly maintained.

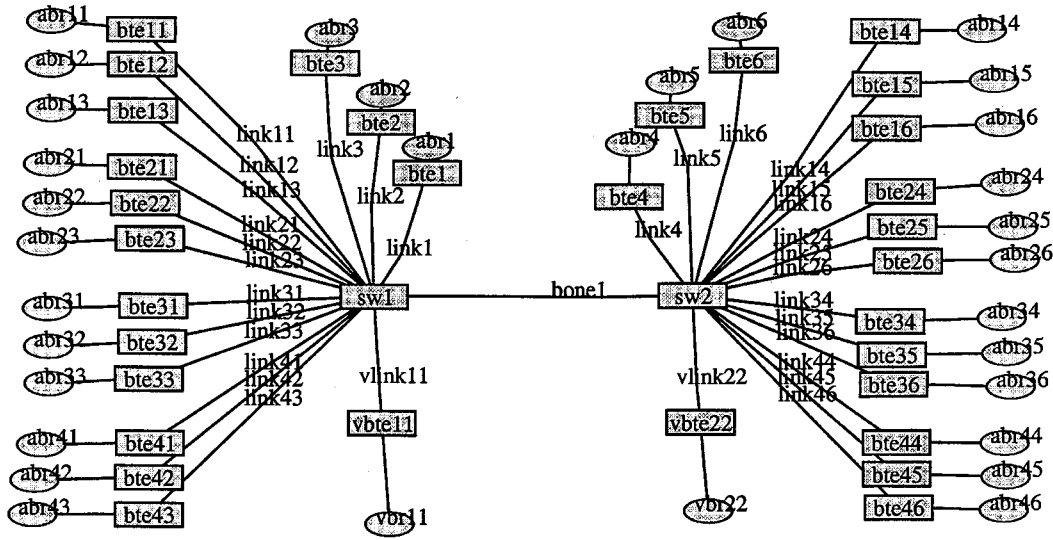


Fig. 8. Bottleneck link shared by ABR and VBR connections.

In addition to fluctuations in available capacity due to statistical multiplexing of VBR flows, a further contributor to changes in available capacity would be the admission/departure of CBR/VBR calls into the system. Note that switch hardware and traffic demands would limit the rate at which new calls are admitted into the system. Overall we believe it is not unreasonable to assume that once the operation regime and traffic on a link is known, the variability ρ can be assessed by combining empirical evaluation of VBR traffic fluctuations and admission control on the connection process.

Queue Threshold q^* : The target queue level q^* will determine the overall utilization of the system. Intuitively the larger q^* , the greater the ability of the system to buffer ABR traffic, and thus to exploit available capacity if it suddenly becomes available. However, a larger q^* means a larger queuing delay in steady-state, so a tradeoff between utilization and delay needs to be made in selecting q^* . In addition, if some ABR sessions are constrained and can not fully utilize the allocated rate, we showed in Section III-C that a large enough q^* is necessary to allow reallocating the unused capacity of constrained connections. In essence, q^* determines the “dynamic range” for the explicit rate that the link can support when sessions are constrained elsewhere.

V. SIMULATION AND PERFORMANCE EVALUATION

In this section we present some simulation results to verify the analysis in previous sections. Our network configuration, shown in Fig. 8, contains 15 ABR connections and aggregated VBR connections sharing a bottleneck link. We are interested in the interaction between ABR feedback control and the rate variation of VBR connections, as well as their impact on the bottleneck link’s queue length.

Stability: In Section III-B we proved the bottleneck queue length will converge to the target level q^* in a fixed environment if the drift function gain k is chosen such that $(k/n) \sum_{i=1}^n \tau_i < 1$. We first consider the case where ABR connections are greedy and the VBR connections are off, hence the available capacity and the number of “on” sources




Fig. 9. Queue dynamics for different drift function gains k .

are both fixed. The queue threshold q^* is 200 and the largest round-trip delay τ for the ABR connections is set to be 20 ms, which means that the largest k guaranteed to ensure stability is 0.0207 Mb/(Cells * s). The queue dynamics for $k = 0.02$ and $k = 0.06$ are shown in Fig. 9, illustrating that $k = 0.06$ may result in instability.

Queue Response to Varying Available Capacity: Next we study the bottleneck queue response when the available capacity changes. We feed greedy ABR connections and an on-off VBR connection into the bottleneck link. The arrival rate of the VBR connection, ACR of an ABR connection, and the queue length are shown in Fig. 10. The figure shows a jump in the queue length when the VBR flow starts bursting. However, the queue length goes back to the target level and available capacity is reallocated after the bottleneck link responds to the change in the available capacity. Since the queue length is controlled around the target level, the available capacity is fully utilized.

Statistical Multiplexing: In the lossless case the queue overshooting contributed from bursting ABR sources is determined by the total number of ABR sessions, n_{max} . In Section IV-A this contribution was reconsidered because the effective number of bursting ABR sessions within a round-trip delay time is smaller

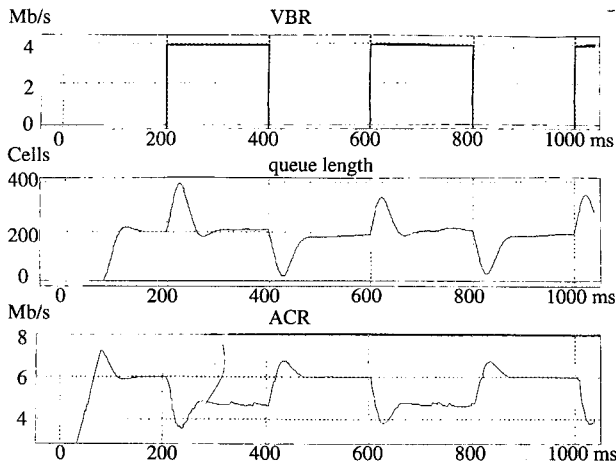


Fig. 10. Queue length and ACR in a changing environment.

TABLE II
COMPARISON OF ADMISSIBLE NUMBERS

"QoS" (δ)	0	1e-7	1.0e-6	1.0e-5	1.0e-4	1.0e-3
n_{max}	30	30	30	30	30	30
n_{max}^b	30	63	68	75	84	94

than the total number of ABR connections due to statistical multiplexing. Hence, in a controlled loss scenario, the buffer requirement can be reduced, or alternatively the admissible number can be increased. We assume the average idle time of ABR connections is 100 ms and use the result in Section IV-A to compute the n_{max}^b for different "quality-of-service" (QoS) when $n_{max} = 30$. The results, shown in Table II, indicate that the admissible number of ABR sessions is increased significantly due to statistical multiplexing of their bursts.

Utilization Improvement: An important advantage of introducing ABR service is to improve the network utilization. Because of the stringent quality of service requirement and bursty behavior of VBR traffic, the network utilization is usually low if only VBR connections are carried. One can let ABR connections use the momentarily unused bandwidth inside the network, so as to improve utilization. In our final simulation we let the bottleneck link carry 45 VBR connections from the video traces for obtaining Fig. 7. In addition, we introduce three greedy ABR connections to exploit the unused bandwidth. We found the utilization of the bottleneck link increased from 70% to 95%. The plots of aggregate VBR arrival rate, ACR of an ABR connection, and the queue length are shown in Fig. 11. It shows that the ACR's of ABR connections are varying according to the changes in the VBR arrival rates.

VI. CONCLUDING REMARKS

In order to avoid throughput collapse, ABR service needs to be implemented so as to provide some control on cell loss. To achieve this, flow control mechanisms need to be designed, so that by making appropriate resource reservations and performing call admission, the network can ensure that losses are low. We have proposed a *simple* algorithm to compute an explicit rate bound on source traffic. By accounting for the rate




Fig. 11. 45 VBR and 3 ABR connections.

variability of the interfering (VBR) traffic and the source update behavior, we analyzed the queue dynamics at the bottleneck link and derived the minimum buffer and capacity requirements for guaranteeing lossless service to ABR connections.

In general, resources need to be reserved to absorb such traffic variability and to control loss, but doing so would typically reduce link utilization. By accounting for statistical multiplexing of source bursts, one can reduce the required reservations to achieve the desired quality of service, or alternatively one can allow for a larger number of concurrent ABR connections for a given reservation. Finally, we have discussed design parameters for our proposed algorithm in the context of standard ABR rate control mechanisms. We identified the factors that would affect queue overflows, such as the source behavior, variability of available capacity, and of course round-trip delays. We believe that our analysis provides some novel insights into the capacity/buffer requirements for ABR rate control mechanisms.

APPENDIX I PROOF OF LEMMAS

A. Proof of Lemma 1

Consider the aggregate traffic rate reaching the bottleneck queue at time t . We can subdivide the contributing sources into the $\hat{n}(t - \tau)$ that were thought to be "on" at time $t - \tau$, and those that were thought to be "off" as shown in the equation at the bottom of the next page, where the arguments of indicator functions are based on the two cases captured in (2).

Using the explicit rate constraint and the bound on the ramp-up of sources' transmission rates, we can establish that

$$\begin{aligned} r_j(t - \tau_j^f) &\leq r_j(t - \tau + \tau_j^b) + g[(t - \tau_j^f) - (t - \tau + \tau_j^b)] \\ &\leq e(t - \tau) + g\Delta\tau. \end{aligned}$$

Now distinguishing between the sources which were "on" and accounting for the worst-case linear growth of "off" sessions

from r^* , we get the following bound on the aggregate rate into the queue at time t :

$$\begin{aligned}
 \sum_{j=1}^{n_{\max}} r_j(t - \tau_j^f) &\leq \hat{n}(t - \tau)[e(t - \tau) + g\Delta\tau] \\
 &\quad + [n_{\max} - \hat{n}(t - \tau)][r^* + g\tau] \\
 &\leq f(q(t - \tau)) + c(t - \tau) + \hat{n}(t - \tau)g\Delta\tau \\
 &\quad + [n_{\max} - \hat{n}(t - \tau)][r^* + g\tau] \\
 &\leq f(q(t - \tau)) + c(t - \tau) \\
 &\quad + n_{\max}[r^* + g\tau]. \tag{9}
 \end{aligned}$$

B. Proof of Lemma 2

Note that the upper bound (6) on $\dot{q}(t)$ depends on the queue length at $t - \tau$. If $q(t - \tau) \geq q^* + (w/k)$, it follows that $f(q(t - \tau)) \leq -w$ and $\dot{q}(t) \leq 0$. In other words, the queue length has to stop increasing τ seconds after it exceeds $q^* + (w/k)$. Consider arbitrary time intervals during which the queue length exceeds $q^* + (w/k)$; we call such periods ‘‘overshooting cycles.’’ We first show an upper bound on the queue length over such overshooting cycles.

Without loss of generality, let $t = 0$ be the beginning of an overshooting cycle and $q(0) = q^* + (w/k)$. Since the queue length exceeds $q^* + (w/k)$, we know $\dot{q}(t) \leq 0$, for $t > \tau$ on the overshooting cycle. To compute the maximum queue length q_{\max} over an overshooting cycle, it suffices to consider the worst-case queue growth, see Fig. 12, on the interval $[0, \tau]$:

$$\begin{aligned}
 q(t) &\leq q^* + \frac{w}{k} + \int_0^t f(q(t - \tau)) + w \, dt \\
 &\leq q^* + \frac{w}{k} + \int_0^t (kq^* + w) \, dt \\
 &= q^* + \frac{w}{k} + (kq^* + w)\tau = q_{\max}. \tag{10}
 \end{aligned}$$

We have shown that q_{\max} is an upper bound on the queue length over an overshooting cycle. For intervals other than overshooting cycles, the queue length does not exceed $q^* + (w/k)$, thus q_{\max} is an upper bound on the queue length. ■

C. Proof of Lemma 3

We can show the stability of this system based on the Nyquist criterion [25]. In particular, we consider how the frequency re-




Fig. 12. Upper bound on the queue length.

sponse of $G(s)$ encircles the Nyquist point $z = -(n/k)$. The Nyquist plot of $G(s)$ is determined by

$$G(j\omega) = \sum_{i=1}^n \frac{e^{-j\omega\tau_i}}{j\omega} = \sum_{i=1}^n \left[-\frac{\sin(\tau_i\omega)}{\omega} - j \frac{\cos(\tau_i\omega)}{\omega} \right].$$

Since $|\sin(\omega\tau_i)/\omega| \leq \tau_i$, it follows that $\text{Re}(G(j\omega)) \geq -\sum_{i=1}^n \tau_i$. Thus the Nyquist plot of $G(j\omega)$ always resides on the right-hand side of the vertical line $z = -\sum_{i=1}^n \tau_i$ on the complex plane. If we choose the Nyquist point to be $z = -(n/k) < -\sum_{i=1}^n \tau_i$, then $G(j\omega)$ will stay away from circling the Nyquist point at all. As a result, $(k/n) \sum_{i=1}^n \tau_i < 1$ is a sufficient condition to ensure the asymptotic stability of the system. ■

D. Proof of Lemma 4


In Fig. 4 we have a controller $k(x)$ with input $x = c - \sum_{i=1}^n r(t - \tau_i)$ in the system. Since the ramp-up of sources is constrained, i.e., $\dot{r}(t) = \min[\dot{c}(t), g]$, the controller $k(x)$ is nonlinear and $k(x) = \min[(k/n)x, g]$, see Fig. 13. A generalized Nyquist criterion—Circle criterion [25, p. 344] is useful in determining the stability of a nonlinear system. For reference, an abridged and rephrased version of the theorem is given below.

Theorem 1 (Circle Criterion): Consider a feedback control system consisting of a nonlinear controller (memoryless gain function) $k(x)$ and a LTI system $G(s)$, e.g., the system in Fig. 4. The system is asymptotically stable if

1. $k(x)$ lies in the sector $[a, b]$, i.e., $a \leq (k(x)/x) \leq b$, $\forall x \neq 0$, where $0 < a < b$,
2. $G(s) = G_a(s) + G_r(s)$, where $G_r(s)$ is strictly proper and $G_a(s)$ is the Laplace transform of a function in the space $L_1[0, \infty)$ augmented by delayed impulses,

and the Nyquist theorem is satisfied for the $G(j\omega)$ locus with respect to the circle with diameter on the negative real axis of the complex plane from $-(1/a)$ to $-(1/b)$, i.e., the locus stays away from the circle and encircles it an appropriate number of times [25] according to the Nyquist theorem.

$$\begin{aligned}
 \sum_{j=1}^{n_{\max}} r_j(t - \tau_j^f) &= \underbrace{\sum_{j=1}^{n_{\max}} 1\{r_j(t - \tau - \tau_j^f) \geq r^* \text{ OR } e(t - \tau - \tau_j) < r^*\} r_j(t - \tau_j^f)}_{\text{‘‘on’’ sources}} \\
 &\quad + \underbrace{\sum_{j=1}^{n_{\max}} 1\{r_j(t - \tau - \tau_j^f) < r^* \text{ AND } e(t - \tau - \tau_j) \geq r^*\} r_j(t - \tau_j^f)}_{\text{‘‘off’’ sources}}
 \end{aligned}$$

Fig. 13. $k(x)$ lies in a sector.

We first show that $k(x)$ lies in a sector. Since $r(t)$ is nonnegative, hence $x = c - \sum_{i=1}^n r(t - \tau_i) \leq c$. In addition, we assume that $k(c) = \min[(k/n)c, g] = g < (k/n)c$, otherwise $k(x)$ is a linear function and the stability of such system can be considered based on the Nyquist criterion, see Lemma Section III. We consider the possible values of $k(x)/x$ as follows.

Case 1: $0 < x \leq c$

$$\frac{k(x)}{x} = \begin{cases} \frac{k}{n}, & \text{if } 0 < x < g \frac{n}{k}, \\ \frac{g}{x} \geq \frac{g}{c}, & \text{if } g \frac{n}{k} \leq x \leq c. \end{cases}$$

Case 2: $x < 0$

$$\frac{k(x)}{x} = \frac{\min\left[\frac{k}{n}x, g\right]}{x} = \frac{k}{n}.$$

Note that $(g/c) < (k/n)$ because $k(c) = g < (kc/n)$. Thus we can conclude that $(g/c) \leq (k(x)/x) \leq (k/n)$.

Next we verify that whether $G(s)$ can be decomposed in the form of $G_a(s) + G_r(s)$. Let us choose a single term $(1/s)e^{-\tau_1 s}$ in $G(s)$ as an example. We can decompose $(1/s)e^{-\tau_1 s}$ in the following way:

$$\frac{e^{-\tau_1 s}}{s} = \frac{e^{-\tau_1 s} - 1}{s} + \frac{1}{s}.$$

The second term is strictly proper, which satisfies the condition of Theorem 1. The inverse Laplace transform of the first term is $u[t - \tau_1] - u[t]$, where $u[t]$ is a unit step function, thus it is clearly in $L_1[0, \infty)$.


Finally, let us consider the Nyquist plot of $G(j\omega)$. As shown in the proof of Lemma 3, the real part of $G(j\omega)$ always resides on the right hand side of the vertical line $z = -\sum_{i=1}^n \tau_i$ on the complex plane. If $(k/n) \sum_{i=1}^n \tau_i < 1$, the Nyquist plot of $G(j\omega)$ will stay away from the circle determined by $-(n/k)$ and $-(c/g)$, see Fig. 14. Therefore, $(k/n) \sum_{i=1}^n \tau_i < 1$ is a sufficient condition to guarantee the stability of the system in Fig. 4. ■

APPENDIX II

PSEUDOCODE FOR SOURCES AND BOTTLENECK LINKS

A. ABR Source

```
/* ABR SOURCE */
/* initialization */
```

Fig. 14. Nyquist plot of $G(j\omega)$ and the circle.

```
count= 0;
acr = MCR;
/* upon receiving a returned RM */
er =RM_cell(ER);
acr = acr + g*(now() - last_RM);
acr= min(acr, er + MCR, PCR);
/* send cells */
if(now() >= scheduled_cell_time)
{
  if (count== 0)
  {
    send_RM_cell();
    last_RM = now();
    count++;
  }
  else
  {
    send_data_cell();
    count = (count +1) mod Nrm;
  }
  scheduled_cell_time = now() +1/acr;
}
```

B. Bottleneck Links

```
/* LINKS */
/* upon receiving a forward RM cell */
delta_n = delta_n + Nrm/(1*CCR);
drift = -k*(q_length - q_target);
er = (link_rate + drift - measured_VBR_rate)
/(min(1, n_on));
/* upon receiving a backward RM cell */
RM_cell(ER) = min(er, RM_cell(ER));
/* at the end of each interval l */
n_on = n_on * alpha + delta_n(1-alpha);
delta_n=0;
```

ACKNOWLEDGMENT

The authors used the NIST ATM network simulator for the simulation in this paper. The pseudocodes of switches and ABR sources were implemented based on [15]. The authors would like to thank the anonymous reviewers and editor J. N. Daigle for their constructive comments.


REFERENCES

- [1] E. Altman, F. Baccelli, and J. C. Bolot, "Discrete-time analysis of adaptive rate control mechanisms," in *Proc. 5th Int. Conf. Data and Communications*, 1993, pp. 121–140.
- [2] A. Arulambalam, X. Chen, and N. Ansari, "An intelligent explicit rate control algorithm for ABR service in ATM networks," in *Proc. IEEE ICC*, Montreal, Canada, 1997.
- [3] L. Benmohamed and S. M. Meerkov, "Feedback control of congestion in packet switching network: The case of a single congested node," *IEEE/ACM Trans. Networking*, vol. 1, pp. 694–708, Dec. 1993.
- [4] —, "Feedback control of congestion in packet switching networks: The case of multiple congested nodes," in *Proc. American Control Conf.*, 1994, pp. 1104–1108.
- [5] L. Benmohamed and Y. T. Wang, "A control-theoretic ABR explicit rate algorithm for ATM switches with Per-VC queueing," in *Proc. IEEE Infocom*, San Francisco, CA, 1998.
- [6] F. Bonomi and K. W. Fendick, "The rate-based flow control framework for the available bit rate ATM service," *IEEE Network Mag.*, vol. 9, no. 2, pp. 25–39, 1995.
- [7] F. Bonomi, D. Mitra, and J. B. Serry, "Adaptive algorithms for feedback-based flow control in high-speed wide-area ATM networks," *IEEE J. Select. Areas Commun.*, vol. 13, no. 7, pp. 1267–1283, 1995.
- [8] J.-Y. Le Boudec, G. de Veciana, and J. Walrand, "QoS in ATM: Theory and practice," in *Proc. 35th IEEE CDC*, 1996, pp. 773–778.
- [9] A. Charny, K. K. Ramakrishnan, and A. Lauck, "Time scale analysis and scalability issue for explicit rate allocation in ATM networks," *IEEE/ACM Trans. Networking*, vol. 4, pp. 569–581, Aug. 1996.
- [10] F. Chiussi, Y. Xia, and V. Kumar, "Virtual queueing techniques for ABR service: Improving ABR/VBR interaction," in *Proc. IEEE Infocom*, Kobe, Japan, 1997.
- [11] S. Chong, R. Nagarajan, and Y. T. Wang, "First-order rate-based flow control with dynamic queue threshold for high-speed wide-area ATM networks," in *Proc. SPIE Conf. Performance and Control of Network Systems*, vol. 3231, 1997, pp. 259–270.
- [12] A. I. Elwalid, "Analysis of adaptive rate-based congestion control for high-speed wide-area networks," in *Proc. IEEE ICC'95*, Seattle, WA.
- [13] *ATM Traffic Management Specification Version 4.0*, The ATM Forum, 1996.
- [14] C. Fulton and S.-Q. Li, "UT: ABR feedback control with tracking," in *Proc. IEEE Infocom'97*, Kobe, Japan.
- [15] N. Golmie, A. Koenig, and D. Su, "The NIST ATM Network Simulator Operation and Programming," NIST, Gaithersburg, MD, 1995.
- [16] R. Jain, "Congestion control and traffic management in ATM networks: Recent advances and a survey," *Comput. Networks and ISDN Syst.*, pp. 1723–1738, Oct. 1996.
- [17] R. Jain, S. Fahmy, S. Kalyanaraman, and R. Goyal, "The ERICA switch algorithm for ABR traffic management in ATM networks," to be published.
- [18] R. Jain, S. Kalyanaraman, S. Fahmy, R. Goyal, and S. C. Kim, "Source behavior for ATM ABR traffic management: An explanation," *IEEE Commun. Mag.*, vol. 34, pp. 50–57, 1996.
- [19] L. Kalampoukas and A. Varma, "Dynamics of an explicit rate allocation algorithm for available bit rate service in ATM networks," UCSS-CRL-95-54, 1995.
- [20] Y. Moret and S. Fdida, "ERAQLES: An efficient explicit rate algorithm for ABR," in *Proc. IEEE Globecom*, Phoenix, AZ, 1997.
- [21] H. Ohsaki, M. Murata, H. Suzuki, C. Ikeda, and H. Miyahara, "Rate-based congestion control for ATM networks," in *ACM SIGCOMM*, 1995, pp. 60–72.
- [22] A. Romanrow, "TCP over ATM: Some performance results," ATM Forum/93-784, 1993.
- [23] A. Romanrow and S. Floyd, "The dynamics of TCP traffic over ATM networks," *IEEE J. Select. Areas Commun.*, vol. 13, May 1995.
- [24] K. Y. Siu and H. Y. Tzeng, "Intelligent congestion control for ABR services in ATM networks," *Comput. Commun. Rev.*, vol. 24, no. 5, pp. 81–106, 1994.
- [25] M. Vidyasagar, *Nonlinear Systems Analysis*. Englewood Cliffs, NJ: Prentice-Hall, 1993.
- [26] M. Wong and F. Bonomi, "Z: A novel explicit rate congestion control algorithm," to be published.



Ching-Fong Su (S'95–M'98) received the B.S. degree in electrical engineering from National Tsing-Hua University, Hsinchu, Taiwan, R.O.C., in 1991, and the M.S. and Ph.D. degrees in electrical and computer engineering from University of Texas at Austin in 1995 and 1998 respectively.

Since 1998, he has been a Member of Research Staff with Fujitsu Labs of America Inc., Sunnyvale, CA. His current research interests include routing in high-speed networks, multimedia applications, and performance evaluation of communication systems.



Gustavo de Veciana (S'88–M'94) received the B.S., M.S., and Ph.D. degrees in electrical engineering from the University of California at Berkeley in 1987, 1990, and 1993, respectively.

In 1993, he joined the Department of Electrical and Computer Engineering, University of Texas at Austin, where he is currently an Associate Professor. His research focuses on issues in the design and control of telecommunication networks.

Dr. de Veciana is an editor for the *IEEE/ACM TRANSACTIONS ON NETWORKING*. He is the recipient of a General Motors Foundation Centennial Fellowship in Electrical Engineering and a 1996 National Science Foundation CAREER Award.

Jean Walrand (S'71–M'74–SM'90–F'93) received the Ph.D. degree in electrical engineering from the University of California at Berkeley in 1979.

He is currently a Professor in the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley. His research interests include communication networks and applied probability. He is the author of three books and many technical papers.

Dr. Walrand is a recipient of the Lanchester Prize.