

Efficient Network QoS Provisioning Based on per Node Traffic Shaping

Leonidas Georgiadis, *Senior Member, IEEE*, Roch Guérin, *Senior Member, IEEE*, Vinod Peris, *Student Member, IEEE*, and Kumar N. Sivarajan, *Member, IEEE*

Abstract—This paper addresses the problem of providing per-connection end-to-end delay guarantees in a high-speed network. We assume that the network is connection oriented and enforces some admission control which ensures that the source traffic conforms to specified traffic characteristics. We concentrate on the class of rate-controlled service (RCS) disciplines, in which traffic from each connection is reshaped at every hop, and develop end-to-end delay bounds for the general case where different shapers are used at each hop. In addition, we establish that these bounds can also be achieved when the shapers at each hop have the same “minimal” envelope.

The main disadvantage of this class of service discipline is that the end-to-end delay guarantees are obtained as the sum of the worst-case delays at each node, but we show that this problem can be alleviated through “proper” reshaping of the traffic. We illustrate the impact of this reshaping by demonstrating its use in designing RCS disciplines that outperform service disciplines that are based on generalized processor sharing (GPS). Furthermore, we show that we can restrict the space of “good” shapers to a family which is characterized by only one parameter. We also describe extensions to the service discipline that make it work conserving and as a result reduce the average end-to-end delays.

Index Terms—QoS provisioning, real-time traffic, traffic shaping, ATM, scheduling, end-to-end delay guarantees.

I. INTRODUCTION

IN this paper, we consider the problem of providing per connection end-to-end delay (and throughput) guarantees in high speed networks. Various scheduling policies have been suggested in the literature for this purpose. Among them, policies based on fair queueing, alternatively known as generalized processor sharing (GPS) [7], [11]–[13], have attracted special attention since they guarantee throughput to individual connections and provide smaller end-to-end delay bounds than other policies for connections that cross several nodes. A key factor in obtaining these smaller delay bounds is the ability to take into account (delay) dependencies in the

successive nodes that a connection has to cross, which is in general very difficult to do with other policies.

One notable attempt at addressing this general problem is that of [6], which introduced the concept of service burstiness, and used it to provide a framework to characterize service disciplines and evaluate their end-to-end delay performance. However, the generality of the framework in [6] did not result in as tight end-to-end delay bounds as those obtained by focusing on a specific policy. For example, the bounds available based on the techniques of [6] are no better than the looser bounds found in [12].

In this paper we concentrate on rate-controlled service (RCS) disciplines, which have also been proposed in the literature [18] to provide performance guarantees to individual connections. In this class of service disciplines, the traffic of each connection is reshaped at every node to ensure that the traffic offered to the scheduler arbitrating local packet transmissions conforms to specific characteristics. In particular, it is typically used to enforce, at a node inside the network, the same traffic parameter control as the one performed at the network access point, which is based on the parameters negotiated during connection establishment. Reshaping makes the traffic at each node more predictable and, therefore, simplifies the task of guaranteeing performance to individual connections; when used with a particular scheduling policy, it allows the specification of worst case delay bounds at each node [18]. End-to-end delay bounds can then be computed as the sum of the worst case delay bounds at each node along the path.

The main advantages of an RCS discipline, especially when compared to GPS, are flexibility, lower buffer requirements at intermediate nodes, and typically simpler implementation [17]. In addition, in the single node case the RCS discipline that uses the non-preemptive earliest deadline first (NPEDF) scheduling policy, is known to be optimal (the optimality criterion is defined in Section IV) [8]. However, for the more interesting case of general networks with many nodes, optimality does not hold. Furthermore, in Section IV-A we show with simple examples that when a connection has to cross many nodes, GPS outperforms the “naive” rate-controlled NPEDF discipline. As a result, it has been argued that despite its potentially greater complexity, a GPS-based service discipline should be the solution of choice to provide performance guarantees to individual connections (see, for example, [3]).

A key result of this paper is to establish that RCS disciplines can be designed so as to outperform GPS-based ones, even in

Manuscript received December 22, 1995; revised April 26, 1996; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor B. Mukherjee. An earlier version of this paper was presented at the IEEE INFOCOM'96 Conference, San Francisco, CA, March 26–28, 1996. The paper was selected by the Conference as one of its top papers and referred to the TRANSACTIONS for possible publication after the TRANSACTIONS' own independent review.

L. Georgiadis is with the School of Engineering, ECE Department, Aristotle University, Thessaloniki 54006, Greece (email: leonid@eng.auth.gr).

R. Guérin and V. Peris are with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 USA (email: guerin@watson.ibm.com; vperis@watson.ibm.com).

K. N. Sivarajan is with the ECE Department, Indian Institute of Science, Bangalore 560012, India (email: kumar@ece.iisc.ernet.in).

Publisher Item Identifier S 1063-6692(96)06091-8.

a network environment. This is achieved by proper selection of the traffic reshaping performed at each node. Specifically, any end-to-end delay bounds that can be guaranteed by the GPS discipline can also be achieved by an RCS discipline by using a simple algorithm to determine how to reshape the traffic and then specify worst-case delay bounds at each node. The sum of the worst-case delay bounds of this RCS discipline is then no larger than the delay guarantees provided by the GPS discipline. We also show that RCS disciplines have the additional flexibility of providing end-to-end delay bounds that cannot be guaranteed by the GPS discipline. Furthermore, because of traffic reshaping, the network buffer requirements of RCS disciplines are in general significantly smaller than those of the GPS discipline (see [6] for related discussions). Based on these advantages and their potential implementation simplicity [17], we believe that RCS disciplines are very effective candidates for providing end-to-end performance guarantees to individual connections in integrated services networks.

The paper is structured as follows. In Section II, we introduce our traffic model, and in particular our assumptions concerning properties of the envelope of the input traffic, and the general structure of our shapers. Section III is dedicated to the description of RCS disciplines and to the derivation of several results concerning the delay guarantees they can provide given the traffic and shaper models of Section II. Section IV is devoted to a comparison with the GPS service discipline. Section IV-A considers first the simpler version of GPS, i.e., rate proportional processor sharing (RPPS), as it is of greater practical significance. Section IV-B considers the more complicated case of general GPS for which similar results are established. Various properties of traffic shapers are investigated in Section V and used to establish that the reshaping needed for RCS disciplines to perform well can be achieved using “simple” shapers. Finally, the important extension demonstrating that the results of the paper hold when reshaping is performed only in case of congestion is the topic of Section VI. A brief conclusion summarizes the main findings of the paper. The Appendixes contain proofs of the lemmas, as well as an extension to the more general case of subadditive traffic envelopes.

II. SYSTEM MODEL AND DEFINITIONS

We consider a network comprised of store-and-forward packet switches, in which a packet scheduler is available at each output link. Traffic from a particular connection entering the switch passes through a packetizer and a traffic shaper before being delivered to the scheduler, as indicated in Fig. 1. The traffic shaper regulates traffic, so that the output of the shaper satisfies certain prespecified traffic characteristics.

In this paper, we use a deterministic approach to specify the traffic characteristics of a connection. Modeling traffic as a fluid, $U[t, t + \tau]$ is used to denote the amount of traffic arriving at the network ingress in the interval $[t, t + \tau]$. However, a network element typically operates on packets and so there is a packetizer (see Fig. 1) that reassembles the packets. These packets are then regulated by the traffic shaper before reaching

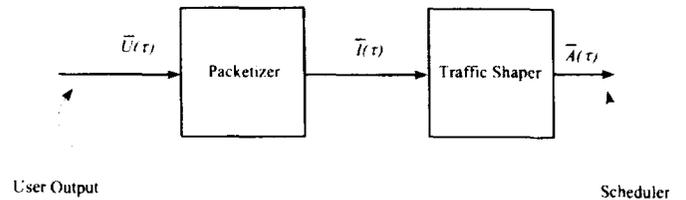


Fig. 1. Connection traffic flow.

the link scheduler which arbitrates the transmission of packets on the link.

We assume that $U(\tau) := U[0, \tau]$ is right-continuous and that there is a nonnegative function $\bar{U}(\tau)$ called *envelope* of $U[0, t]$, such that

$$U[t, t + \tau] \leq \bar{U}(\tau), \quad t \geq 0, \tau \geq 0.$$

The envelope function is not unique; without loss of generality (see [2]) we can assume that $\bar{U}(\tau)$ is right-continuous, nondecreasing, and subadditive.

The packetizer spits out packets of maximum length L , which are instantaneously delivered to the shaper when the last bit of the packet is received. We denote the traffic at the output of the packetizer in the interval $[t, t + \tau]$ as $I[t, t + \tau]$. It is easy to see that, for any nonnegative t and τ

$$I[t, t + \tau] \leq U[t, t + \tau] + L \leq \bar{U}(\tau) + L =: \bar{I}(\tau). \quad (1)$$

Thus $\bar{I}(\tau)$ is an envelope of the traffic that is input to the traffic shaper. The traffic shaper reshapes the incoming traffic by delaying the packets according to the rules described next, and then delivers them to the scheduler. The traffic shaper is characterized by a traffic envelope, $\bar{A}(\tau)$, which provides an upper bound on the amount of traffic that is output by the shaper in any interval of length τ . If $A[t, t + \tau]$ denotes the traffic that is output from the shaper in the interval $[t, t + \tau]$, then the shaper ensures that $A[t, t + \tau] \leq \bar{A}(\tau)$.

More precisely, the traffic shaper outputs packets in order with each packet being released at the earliest time, t , such that

$$A[t - \tau, t] \leq \bar{A}(\tau), \quad 0 \leq \tau \leq t. \quad (2)$$

The traffic shapers that we use in this paper can be constructed from the simple (σ, ρ) traffic regulators of [4] which are described in terms of the backlog in a hypothetical queue that is served at the rate ρ . Assume that traffic $I[0, t]$ is fed to a queue that is served at the rate of ρ . Then, the backlog in this queue at time t , denoted by $W_\rho(I)(t)$, is given by [4]

$$W_\rho(I)(t) := \max_{0 \leq s \leq t} \{I[s, t] - \rho(t - s)\}. \quad (3)$$

Note that the backlog $W_\rho(I)(t)$ differs from that defined in [4] in that it includes the packet that may have arrived at time t . The operation of the (σ, ρ) traffic shaper, can be described in terms of the backlog: The i th packet arriving at time s_i is released at the earliest time $f_i \geq s_i$ such that the shaper output traffic, $A[0, f_i]$, satisfies the condition

$$W_\rho(A)(f_i) \leq \sigma = L + \delta, \quad \delta \geq 0.$$

Note that the condition $\delta \geq 0$ is necessary in order to allow packets of size L to pass through the shaper. This shaper

corresponds to the operation of a leaky bucket in a store-and-forward network [1], which differs from the (σ, ρ) -regulator defined in [4] in two minor respects:

- 1) packets are entering and exiting the shaper instantaneously and not at a constant rate C ;
- 2) the length of the packet that exits the traffic shaper at time f_i is taken into account in the calculation of $W_{\rho}(A)(f_i)$.

However, with $d_i := f_i - s_i$ denoting the delay that the i th packet experiences in a shaper, the analysis in [4] can be repeated with minor modifications to show that

$$d_i = \frac{1}{\rho} (W_{\rho}(I)(s_i) - \sigma)^+ \text{ and}$$

$$A[t, t + \tau] \leq \sigma + \rho\tau$$

where $x^+ \equiv \max(0, x)$. The (σ, ρ) shaper has also been described in the literature in terms of a token bucket (or leaky bucket), with ρ being the rate of token accumulation and σ being the bucket depth.

A. Delay in the Traffic Shaper

In general, we will be using shapers whose output is a concave, increasing (i.e., $f(t_1) < f(t_2)$ whenever $t_1 < t_2$), piecewise linear function with finite number (say K) of slopes. We are interested in these types of shapers because they are a generalization of the shapers adopted by the the Internet [16] and asynchronous transfer mode (ATM) standards [1]. These shapers can be realized by passing the traffic through a series of K (σ_m, ρ_m) -shapers, $m = 1, 2, \dots, k$ [5]. Let \mathcal{A} denote the aforementioned series and for the input traffic model described earlier, the delay of packet i through \mathcal{A} is [5, Theorem 5.1]

$$d_i = \max_{m=1,2,\dots,K} \left\{ \frac{1}{\rho_m} (W_{\rho_m}(I)(s_i) - \sigma_m)^+ \right\}. \quad (4)$$

It can be readily verified that \mathcal{A} is indeed a shaper with an envelope [5], [14]

$$\bar{A}(\tau) := \min_{m=1,2,\dots,K} \{ \sigma_m + \rho_m \tau \}. \quad (5)$$

We can develop an upper bound on the delay encountered in the shaper \mathcal{A} by traffic that has an envelope $\bar{I}(\tau)$. Taking into account (3) we have that

$$d_i \leq \max_{m=1,2,\dots,K} \left\{ \frac{1}{\rho_m} \left(\max_{0 \leq s \leq s_i} \{ \bar{I}(s_i - s) - \rho_m(s_i - s) - \sigma_m \} \right)^+ \right\}$$

$$\leq \max_{m=1,2,\dots,K} \left\{ \left(\max_{\tau \geq 0} \left\{ \frac{\bar{I}(\tau) - \sigma_m - \rho_m \tau}{\rho_m} \right\} \right)^+ \right\} \quad (6)$$

$$= \max_{\tau \geq 0} \left\{ \left(\max_{m=1,2,\dots,K} \left\{ \frac{\bar{I}(\tau) - \sigma_m}{\rho_m} \right\} - \tau \right)^+ \right\}. \quad (7)$$

Equation (7) provides an upper bound on the delay encountered by traffic with envelope $\bar{I}(\tau)$ through shaper \mathcal{A} that has

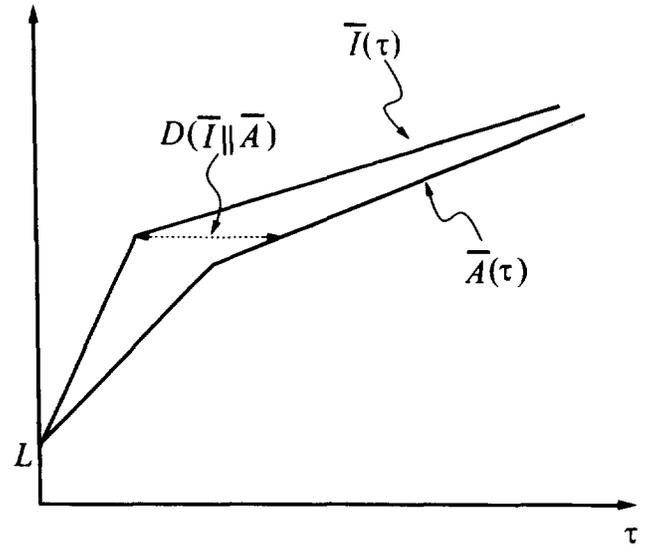


Fig. 2. Graphical illustration of $D(\bar{I}||\bar{A})$.

envelope $\bar{A}(\tau)$ and we denote this by

$$D(\bar{I}||\bar{A}) := \max_{\tau \geq 0} \left\{ \left(\max_{m=1,2,\dots,K} \left\{ \frac{\bar{I}(\tau) - \sigma_m}{\rho_m} \right\} - \tau \right)^+ \right\}. \quad (8)$$

We can write $D(\bar{I}||\bar{A})$ in another form that will be useful in the sequel. The range of $\bar{A}(\tau)$ is $[\min_m \sigma_m, \infty)$ and the inverse of $\bar{A}(\tau)$ is given by

$$\bar{A}^{(-1)}(y) = \max_{m=1,\dots,K} \left\{ \frac{y - \sigma_m}{\rho_m} \right\}, \quad \min_m \sigma_m \leq y < \infty. \quad (9)$$

Extending the definition of $\bar{A}^{(-1)}(y)$ by setting $\bar{A}^{(-1)}(y) = 0$ whenever $0 \leq y < \min_m \sigma_m$, it can be seen from (8) and (9) that

$$D(\bar{I}||\bar{A}) = \max_{\tau \geq 0} \{ (\bar{A}^{(-1)}(\bar{I}(\tau)) - \tau)^+ \}. \quad (10)$$

Graphically, (10) represents the maximum horizontal distance between the envelopes $\bar{I}(\tau)$ and $\bar{A}(\tau)$ as illustrated in Fig. 2.

A word about the notation used in this paper is probably overdue. We use a bar on top to denote a traffic envelope function, with the exception of \bar{D}_n which denotes the end-to-end delay guarantee that can be provided to connection n . Traffic shapers are denoted by the calligraphic alphabet and their envelopes are denoted by the corresponding regular alphabet with a bar on top. When the traffic entering shaper \mathcal{A}_2 with envelope $\bar{A}_2(\tau)$ is the output of a shaper \mathcal{A}_1 with envelope $\bar{A}_1(\tau)$, we denote the bound on the delay in shaper \mathcal{A}_2 , with a slight abuse of notation, by any of the following:

$$D(\mathcal{A}_1||\mathcal{A}_2) \equiv D(\bar{A}_1||\mathcal{A}_2) \equiv D(\bar{A}_1||\bar{A}_2).$$

Notice that if the input traffic envelope, $\bar{I}(\tau) \leq \bar{A}(\tau)$, $\tau \geq 0$, then from (6) (or alternatively Fig. 2) we have that $D(\bar{I}||\bar{A}) = 0$ which implies that no packet is delayed in the shaper \mathcal{A} . In particular, $D(\mathcal{A}||\mathcal{A}) = 0$.

Consider next two shapers $\mathcal{A}_1, \mathcal{A}_2$ in series. Equations (4) and (5) imply that this arrangement is equivalent to a traffic

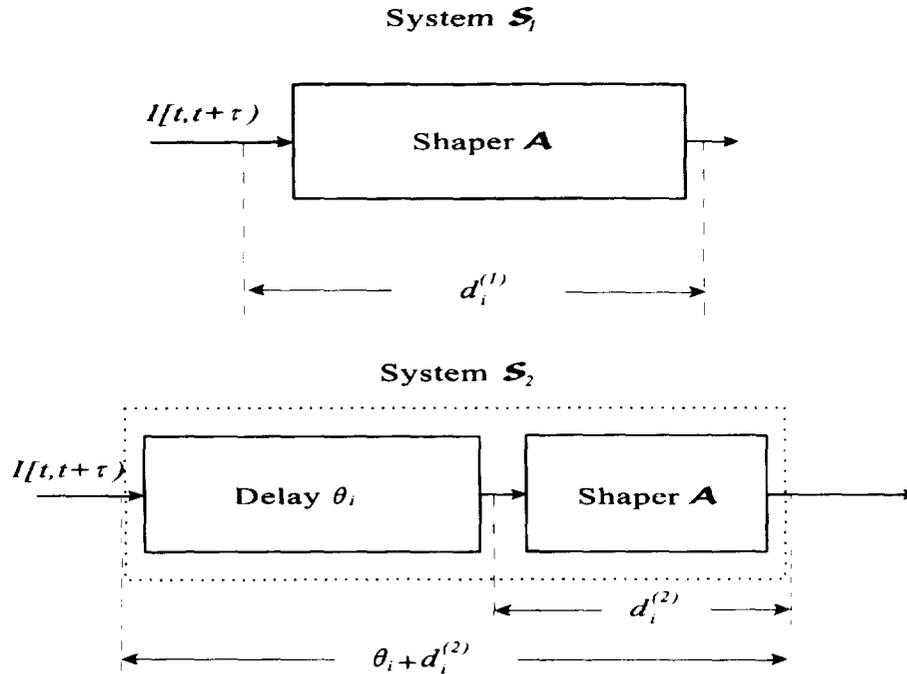


Fig. 3. The systems \mathcal{S}_1 and \mathcal{S}_2 .

shaper \mathcal{A}_3 with envelope

$$\bar{A}_3(\tau) = \min\{\bar{A}_1(\tau), \bar{A}_2(\tau)\}. \quad (11)$$

Equivalence here means that for any input traffic pattern, the delay of every packet from the time it enters \mathcal{A}_1 to the time it exits \mathcal{A}_2 is identical to the delay of the packet in \mathcal{A}_3 .

Next, we state a useful lemma, that relates the packet delays in the two systems \mathcal{S}_1 and \mathcal{S}_2 of Fig. 3. System \mathcal{S}_1 consists of a traffic shaper, \mathcal{A} . System \mathcal{S}_2 consists of a “delay” subsystem and an identical shaper \mathcal{A} connected in series. The delay subsystem delays the i th arriving packet by an arbitrary amount, $\theta_i \geq 0$, before delivering it to \mathcal{A} .

Lemma 1: Assume that packets arrive to systems \mathcal{S}_1 and \mathcal{S}_2 according to the same arrival process $I[0, t]$. If $d_i^{(1)}$ and $d_i^{(2)}$ denote the delay experienced by packet i in the traffic shaper in systems \mathcal{S}_1 and \mathcal{S}_2 , respectively, then

$$d_i^{(1)} \leq d_i^{(2)} + \theta_i, \quad i = 1, 2, \dots$$

that is, the delays of all packets in system \mathcal{S}_1 are smaller than their corresponding delays in system \mathcal{S}_2 .

Lemma 1 identifies the monotonicity property of the shaper with respect to the arrival process and the proof can be found in Appendix A. This is an important property of the traffic shapers considered in this paper and is key to establishing the general end-to-end delay bounds for RCS disciplines.

III. RATE-CONTROLLED SERVICE DISCIPLINES

We are interested in a generalized form of the class of RCS disciplines introduced in [18]. In that work, it is assumed that connections whose traffic satisfies certain burstiness constraints enter the network at various nodes. A node can have several output links, each of which contains a scheduler that decides the order in which packets are transmitted. At each

node along the path of a connection, traffic is reshaped to conform to its original envelope before it enters the appropriate scheduler. Based on the traffic envelope of the connection, upper bounds on the scheduling delays at each node can be guaranteed. It is also shown in [18] that for the traffic shapers considered there, reshaping the traffic to its original envelope does not introduce extra delays. Therefore, an upper bound on the end-to-end packet delay is obtained as the sum of the scheduling and propagation delays.

In this paper, we study the following general class of service disciplines. The traffic of connection n entering the network has an envelope function $\bar{U}_n(\tau)$. At node m , the traffic of connection n is shaped by a traffic shaper \mathcal{A}_n^m . Traffic shapers \mathcal{A}_n^m are of the general type considered in Section II, and different traffic shapers can be used for the same connection at different nodes. The connection traffic exiting \mathcal{A}_n^m enters a scheduler at the appropriate output link at node m , and is scheduled for transmission to the next node or to its destination. We develop end-to-end delay bounds based on the scheduling policies at each node as well as the form of the traffic shapers \mathcal{A}_n^m . These bounds are then used to provide delay guarantees to each connection. In the rest of this paper, we use the term *service discipline* to denote the operation of the system consisting of the traffic shaper as well as the scheduler. We are interested in designing service disciplines of the type described above, so that end-to-end delay guarantees can be provided as efficiently as possible.

We assume that the nodes are output queuing switches, and without loss of generality, that there is no delay inside the switch. In other words, the only delay that a packet incurs at a switch is due to queuing at the output link. Let $C^{m,l}$ be the set of connections passing through output link l of node m . Given \mathcal{A}_n^m , $n \in C^{m,l}$, and the scheduling policy employed at link l of node m , we assume that a bound on the scheduling delay,

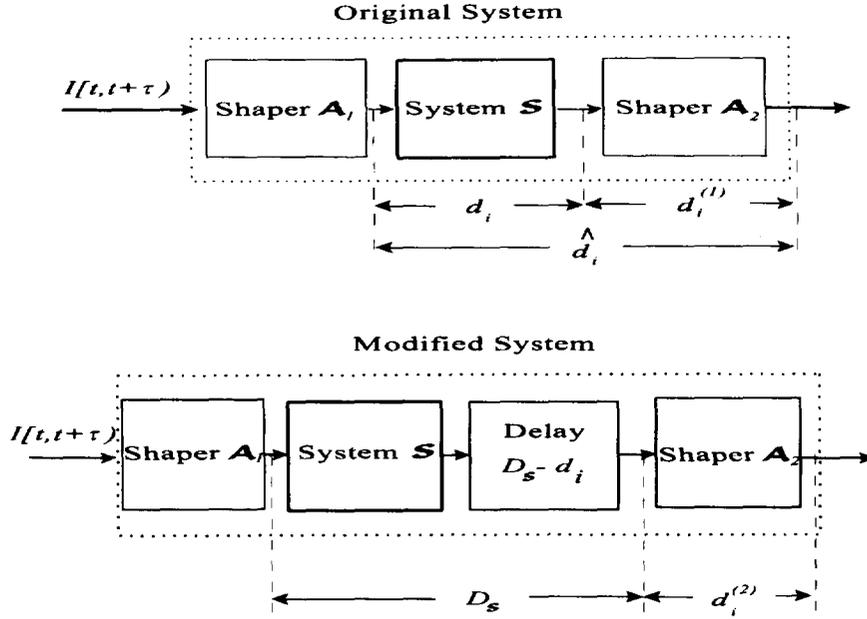


Fig. 4. Original and modified system.

D_n^m , is known for each connection $n \in C^{m,l}$. The scheduling delay includes both the queueing delay and the transmission time of a packet. For example, bounds of this form can be developed for the general traffic shapers of this paper, when the earliest deadline first (EDF) scheduling policy is employed, by a straightforward extension of the method in [19], [8] (see also Theorem 1 in Section IV in this paper). We also assume that an upper bound on the propagation delay for link l is T^l . Knowledge of the bounds D_n^m and T^l alone are not enough to provide bounds on the end-to-end packet delays. We still have to account for any additional delays incurred in the traffic shapers and this is done based on the following proposition.

Proposition 1: Assume that the output of traffic shaper \mathcal{A}_1 enters a system S where it is known that the delay experienced by these packets is bounded above by D_S . The output of system S enters shaper \mathcal{A}_2 . The total delay, \hat{d}_i , that packet i experiences from the time it exits \mathcal{A}_1 to the time it exits \mathcal{A}_2 is bounded above by

$$\hat{d}_i \leq D_S + D(\mathcal{A}_1 \parallel \mathcal{A}_2).$$

Proof: Let d_i be the delay of packet i in system S , and let $d_i^{(1)}$ be its delay in \mathcal{A}_2 . By definition, $\hat{d}_i = d_i + d_i^{(1)}$. Consider next a modified system where a delay system that delays the i th packet by $\theta_i = D_S - d_i$, is inserted between S and \mathcal{A}_2 (see Fig. 4). Now let $d_i^{(2)}$ denote the delay of packet i in \mathcal{A}_2 under this new arrangement. Note that by the definition of D_S , $\theta_i \geq 0$. Applying Lemma 1 we conclude that

$$d_i^{(1)} \leq D_S - d_i + d_i^{(2)}$$

and therefore

$$\hat{d}_i \leq D_S + d_i^{(2)}.$$

Observe now that since the delay of every packet between its entrance time to S and its exit from the delay system is $d_i + \theta_i = D_S$, the traffic entering shaper \mathcal{A}_2 from the delay

system, is a time-shifted version of the traffic exiting \mathcal{A}_1 , and therefore it has envelope $\bar{\mathcal{A}}_1(\tau)$. Hence, $d_i^{(2)} \leq D(\bar{\mathcal{A}}_1 \parallel \mathcal{A}_2)$. \square

From the proof we see that any shaper that has the property described in Lemma 1 satisfies Proposition 1 as well. In particular, the shaper of [18] can easily be seen to satisfy Lemma 1.

Let us now focus our attention on a single connection n that passes through M network nodes, numbered from one to M , with $M+1$ denoting the destination and let l_m denote the link that it traverses in going from node m to node $m+1$. We then apply Proposition 1 with the system S consisting of both the scheduler at node m and the link $l_m = (m, m+1)$, and replacing the shapers \mathcal{A}_1 with \mathcal{A}_n^m , and \mathcal{A}_2 with \mathcal{A}_n^{m+1} . We conclude that the delay that a packet from connection n experiences between the time it exits shaper \mathcal{A}_n^m and the time it exits \mathcal{A}_n^{m+1} is bounded above by

$$D_n^{(m,m+1)} = D_n^m + T^{l_m} + D(\mathcal{A}_n^m \parallel \mathcal{A}_n^{m+1}). \quad (12)$$

Taking (12) into account, we then have the following guaranteed upper bound on the end-to-end delay:

$$\begin{aligned} \bar{D}_n &= D(\bar{I}_n \parallel \mathcal{A}_n^1) + \sum_{m=1}^{M-1} D_n^{(m,m+1)} + D_n^M + T^{l_M} \\ &= D(\bar{I}_n \parallel \mathcal{A}_n^1) + \sum_{m=1}^{M-1} D(\mathcal{A}_n^m \parallel \mathcal{A}_n^{m+1}) \\ &\quad + \sum_{m=1}^M D_n^m + \sum_{m=1}^M T^{l_m}. \end{aligned} \quad (13)$$

While the above framework does provide the flexibility of specifying different shaper envelopes at different nodes along the path of the connection, it is not clear that there is any benefit in doing so. In the next proposition, we show that for the same connection, it does not pay to have different shapers at different nodes. First we introduce some notation to denote

a partial ordering among shapers. We write $\mathcal{A}_1 \preceq \mathcal{A}_2$ (or $\mathcal{A}_2 \succeq \mathcal{A}_1$) whenever $\bar{A}_1(\tau) \leq \bar{A}_2(\tau)$, $\tau \geq 0$. We then make the following simple but important observation.

Lemma 2: If at node m we replace \mathcal{A}_n^m with a shaper \mathcal{B}_n^m such that $\mathcal{B}_n^m \preceq \mathcal{A}_n^m$, then the scheduler delay guarantee of D_k^m can still be provided for any connection k (including connection n).

Proof: Observe that since $\bar{B}_n^m(\tau) \leq \bar{A}_n^m(\tau)$, $\tau \geq 0$, $\bar{A}_n^m(\tau)$ is also an envelope for the traffic exiting \mathcal{B}_n^m . By definition, D_k^m remains an upper bound on the delay of any connection k traffic as long as connection n still has envelope $\bar{A}_n^m(\tau)$. \square

Another simple observation can be made with regard to a series connection of traffic shapers. Let $\mathcal{A}_1 \wedge \mathcal{A}_2$ denote the arrangement of \mathcal{A}_1 followed by \mathcal{A}_2 in series. Since the output of shaper \mathcal{A}_1 has envelope $\bar{A}_1(\tau)$, it follows that:

$$D(\mathcal{A} \parallel \mathcal{A}_1 \wedge \mathcal{A}_2) \leq D(\mathcal{A} \parallel \mathcal{A}_1) + D(\mathcal{A}_1 \parallel \mathcal{A}_2). \quad (14)$$

Also, observe that by (11)

$$\mathcal{A}_i \succeq \mathcal{A}_1 \wedge \mathcal{A}_2, \quad i = 1, 2. \quad (15)$$

Proposition 2: Consider connection n that traverses nodes $1, 2, \dots, M$. Given any RCS discipline π that uses shapers \mathcal{A}_n^m , and guarantees scheduler delay D_n^m , at node m , $1 \leq m \leq M$, the RCS discipline π' that uses the same scheduling policy at all nodes as π , but with the shapers \mathcal{A}_n^m replaced by the shapers

$$\mathcal{B}_n^m = \mathcal{B} = \bigwedge_{m=1}^M \mathcal{A}_n^m$$

can provide the same end-to-end delay guarantees as π , to all connections.

Proof: By (15), we have that $\mathcal{A}_n^m \succeq \mathcal{B}_n^m$ and, therefore, by Lemma 2, π' can guarantee the same scheduling delays to all connections. Since for any connection $k \neq n$, the shapers remain the same, it follows that for these connections policy π' guarantees the same end-to-end delays as π . Consider next connection n , whose envelope at the first shaper is denoted by $\bar{I}_n(\tau)$. Let the end-to-end delay guarantee provided by π and π' to connection n be denoted by \bar{D}_n^π and $\bar{D}_n^{\pi'}$, respectively. Taking into account the fact that $D(\mathcal{B}_n^m \parallel \mathcal{B}_n^{m+1}) = D(\mathcal{B} \parallel \mathcal{B}) = 0$, we conclude from (13) that

$$\bar{D}_n^{\pi'} = D(\bar{I}_n \parallel \mathcal{B}) + \sum_{m=1}^M D_n^m + \sum_{m=1}^M T_n^{(m, m+1)}.$$

Finally, observe that by (14)

$$D(\bar{I}_n, \mathcal{B}) \leq D(\bar{I}_n \parallel \mathcal{A}_n^1) + \sum_{m=1}^{M-1} D(\mathcal{A}_n^m \parallel \mathcal{A}_n^{m+1}).$$

Using (13) again we conclude that $\bar{D}_n^{\pi'} \leq \bar{D}_n^\pi$. \square

Note: According to Proposition 2 we can restrict our attention to disciplines that use identical shapers at all nodes. In the rest of this paper, we consider RCS disciplines that for any given connection use identical shapers at each node, i.e., $\mathcal{A}_n^m = \mathcal{A}_n$. Then, the end-to-end delay guarantee for connection n becomes

$$\bar{D}_n = D(\bar{I}_n \parallel \mathcal{A}_n) + \sum_{m=1}^M D_n^m + \sum_{m=1}^M T_n^m. \quad (16)$$

A word of caution is warranted, as one should not conclude from (16) that the end-to-end delay guarantees are minimized by choosing $\bar{I}_n(\tau)$ as the envelope for all the traffic shapers. While this choice will result in $D(\bar{I}_n \parallel \mathcal{A}_n) = 0$, the scheduler delay bounds, D_n^m , may increase because they depend on the choice of the traffic shapers \mathcal{A}_n^m . In fact, as we will see in the next section, choosing the shaper envelopes to be identical to the input traffic envelope may be quite inappropriate.

As in the policies proposed in [18], the delay bounds in (13) are basically a sum of the worst case delays at *each* node along the path of a connection. However, an individual packet may not encounter the worst case delay at each node. Therefore, one may suspect that these bounds are overly pessimistic and lead to inefficient resource allocations when compared to bounds for other disciplines that take into account delay dependencies between nodes along the path. As mentioned earlier, the impact of delay dependencies is in general difficult to evaluate but can be accounted for in some instances. In particular, these delay dependencies can be accounted for in the case of GPS disciplines [7], [11], [12], which is one of the reasons why tight end-to-end bounds can be obtained. This argument about the inefficiency of worst-case delay assignment relative to GPS was also mentioned in [18].

In the next section, we address this issue by demonstrating that with a suitable choice of shaper envelopes the RCS discipline can provide the same end-to-end delay guarantees that the best delay bounds for GPS can provide. More specifically, we show that for a given set of connections, and their associated paths, the RCS discipline can provide the same end-to-end delay bounds as the GPS discipline. In addition, we show that the RCS discipline can accept a set of connections with associated delay requirements, that cannot be accepted by GPS. This demonstrates the advantage of RCS over GPS in providing efficient end-to-end delay guarantees.

IV. COMPARISON WITH GPS

In this section, we compare the performance of the GPS service discipline with the performance of the RCS disciplines introduced in the previous section. In order to compare two service disciplines, we need to define the performance measure which is of interest to us. The ability of a discipline to provide efficient end-to-end delay guarantees to a given set of connections, is best quantified by the notion of *schedulable region*. Assume that we have N_T connections in a communication network, with the same scheduling discipline, π , operating at all the links in the network. The input traffic of connection n has envelope function $\bar{I}_n(\tau)$, and traverses path P_n of the network, $1 \leq n \leq N_T$. Under these assumptions, we require

that the packets of connection n have an upper bound on their end-to-end delay (delay guarantee), \bar{D}_n , $1 \leq n \leq N_T$. The vector $\bar{\mathbf{D}} = (\bar{D}_1, \dots, \bar{D}_{N_T})$ is *schedulable* under discipline π if the delay bound \bar{D}_n can be guaranteed under π for all packets of connection n , $1 \leq n \leq N_T$. The *schedulable region* of discipline π is the set of all vectors $\bar{\mathbf{D}}$ that are schedulable under π . Note that the schedulable region of a service discipline depends on the envelope functions $\bar{I}_n(\tau)$ and the paths P_n , $n = 1, 2, \dots, N_T$. We say that service discipline π_1 is *at least as good as* the discipline π_2 , if the schedulable region of π_1 is a superset of π_2 , for any given set of connections and paths. If, in addition, there is a set of connections, paths and associated delay bounds that can be guaranteed by π_1 , but not by π_2 , we say that π_1 is *better* than π_2 .

Note that the schedulable region is defined in terms of delay bounds that can be guaranteed *a priori*. These bounds are an integral part of the service discipline and may in fact be significantly worse than the delays actually experienced by packets. From the point of view of admission control, it is irrelevant if in the actual operation of a policy smaller delays are observed, since what is required at the time of connection establishment, is to know whether the delay bounds can be *guaranteed* or not.

Before we proceed with the comparison of RCS and GPS disciplines, we need to recall some preliminary results regarding the NPEDF scheduling policy. This policy has the largest schedulable region among the class of nonpreemptive policies in the single-node case [8] and is therefore a natural choice when considering RCS disciplines. The schedulable region is defined here with respect to scheduler delays only. The schedulable region for N connections that are entering the scheduler through traffic shapers with envelopes $\bar{A}_n(\tau) = L + \delta_n + \rho_n\tau$, $1 \leq n \leq N$, and contending for an output link of speed r , is given by Theorem 4 in [8], which we repeat here for convenience, slightly rephrased to conform to our definitions and notation.

Theorem 1: The NPEDF policy is optimal among the class of nonpreemptive scheduling policies when the connection n traffic entering the scheduler has envelope $\bar{A}_n(\tau) = L + \delta_n + \rho_n\tau$, $1 \leq n \leq N$. Under the stability condition $\sum_{n=1}^N \rho_n \leq r$, the schedulable region of NPEDF consists of the set of vectors (D_1, \dots, D_N) that satisfy the constraints

$$\begin{aligned} & \min\{k+1, N\}L + \sum_{n=1}^k \delta_{i_n} \\ & \leq D_{i_k} \left(r - \sum_{n=1}^{k-1} \rho_{i_n} \right) + \sum_{n=1}^{k-1} \rho_{i_n} D_{i_n}, \quad 1 \leq k \leq N \end{aligned} \quad (17)$$

whenever $D_{i_1} \leq \dots \leq D_{i_N}$.

We note that while the optimality, i.e., largest schedulable region, of NPEDF was established in [8] for envelopes of the form $\bar{I}_n(\tau) = L + \delta_n + \rho_n\tau$, it is straightforward to see that all the arguments used in [8] to derive Theorem 1, go through by simply replacing $L + \delta_n + \rho_n\tau$ with a general envelope $\bar{A}_n(\tau)$ of the type considered here. For these general envelopes, the

appropriate analogue of Theorem 1 can be easily derived by simply rephrasing Lemmas 1 and 2 in [8].

A. Achieving RPPS Delay Guarantees

In this and the next section, we assume for comparison purposes that the traffic of connection n , entering the first node packetizer has envelope $\bar{U}_n(\tau) = \delta_n + \rho_n\tau$. Therefore, the envelope of the traffic that enters the first traffic shaper is $\bar{I}_n(\tau) = L + \delta_n + \rho_n\tau$. We also assume that connection n traverses nodes $1, 2, \dots, M$ and that all the propagation delays are zero. For definitions and notations relating to GPS the reader is referred to [11] and [12]. Recall from Section III, that $C^{m,l}$ is the set of connections that pass through the output link l of node m . Denoting the speed of this link as $r^{m,l}$, we will assume throughout the rest of this section the stability condition

$$\sum_{n \in C^{m,l}} \rho_n \leq r^{m,l}.$$

The GPS policy operates by allocating weight ϕ_n^m for connection n whose traffic crosses node m . These weights are used to determine the rate at which traffic from connection n is served when a set $B^{m,l}$ of connections is backlogged at the output link l of node m through which connection n passes. Specifically, the service rate of connection n is given by

$$g_n^m = \frac{\phi_n^m}{\sum_{k \in B^{m,l}} \phi_k^m} r^m$$

where for simplicity in notation, we denote $r^{m,l}$ as r^m and $B^{m,l}$ as B^m when there is no possibility of confusion. PGPS is a nonpreemptive policy that tracks GPS. In general the procedure developed in [11] to obtain delay bounds given the weights, ϕ_n^m , is complicated and imposes certain restrictions on the ϕ_n^m . Moreover, the practically more important inverse procedure of specifying appropriate weights, that satisfy predetermined delay bounds, is even more cumbersome. However, a simple bound can be obtained in the special case of nonpreemptive RPPS, where $\phi_n^m = \rho_n$ at all nodes through which the connection passes. Specifically, the end-to-end delay bound, \bar{D}_n^* , obtained under nonpreemptive RPPS is given by [9], [12]

$$\bar{D}_n^* = \frac{\delta_n + ML}{\rho_n} + \sum_{m=1}^M \frac{L}{r^m}. \quad (18)$$

From (18), we can already see the weakness of the RCS disciplines relative to RPPS, if the traffic shapers for connection n at every node have envelopes identical to the input envelope $\bar{I}_n(\tau)$. In this case $D(\bar{I}_n || \mathcal{A}_n) = 0$. Since propagation delays are assumed to be zero, from (16), we obtain

$$\bar{D}_n = \sum_{m=1}^M D_n^m.$$

At node m , even if the entire link bandwidth of r^m is somehow dedicated to connection n , the scheduler delay bound, D_n^m , can at best be, $(\delta_n + L)/r^m$. Therefore, the end-to-end delay

bound guaranteed by the RCS discipline satisfies the following inequality:

$$\bar{D}_n \geq \sum_{m=1}^M \frac{\delta_n}{r^m} + \sum_{m=1}^M \frac{L}{r^m}.$$

Since δ_n can be much larger than L , the bounds provided by the RCS discipline under the scenario considered here can be much worse than those obtained under RPPS. For example, assume that all the link speeds are the same, i.e., $r^m = r$, $1 \leq m \leq M$. If $\delta_n = 50L$ and $\rho_n = 0.8r$, we have

$$\frac{\bar{D}_n}{\bar{D}_n^*} \geq \frac{40.8 \times M}{50 + 1.8 \times M}.$$

Therefore, when $M = 2$ we already have $\bar{D}_n/\bar{D}_n^* \geq 1.52$, and for large M , $\bar{D}_n/\bar{D}_n^* \geq 22.67$. As was mentioned in Section III, this discrepancy is due to the fact that the bounds for RPPS take into account delay dependencies at the various nodes, while the bounds for the RCS disciplines are based on independently summing the worst case bounds at each node.

The previous example notwithstanding, we show next that we can design RCS disciplines that provide the same delay guarantees as RPPS by employing traffic shapers with envelopes that are, in general, different from that of the input traffic.

We design the RCS discipline π as follows. For each link we use the NPEDF scheduling policy. We choose the same traffic shaper \mathcal{A}_n for connection n at each node along its path, with its envelope being

$$\bar{A}_n(\tau) = L + \rho_n \tau.$$

Assume that connection n is routed through output link l at node m and let r^m denote the speed of this link. For connection n , we specify the delay bounds for the NPEDF scheduling policy, at node m as

$$D_n^m = L/\rho_n + L/r^m. \quad (19)$$

Let us first show that these bounds can be guaranteed by the NPEDF policy at every node. Consider output link l at node m . Denote by N the total number of connections multiplexed on this link, and index the connections by i_1, i_2, \dots, i_N such that $D_{i_1}^m \leq D_{i_2}^m \leq \dots \leq D_{i_N}^m$. Based on Theorem 1, it is clear that we only need to verify (17).

Using (19) we have

$$\begin{aligned} & D_{i_k}^m \left(r^m - \sum_{n=1}^{k-1} \rho_{i_n} \right) + \sum_{n=1}^{k-1} \rho_{i_n} D_{i_n}^m \\ &= L \frac{r^m - \sum_{n=1}^{k-1} \rho_{i_n}}{\rho_{i_k}} + L \frac{r^m - \sum_{n=1}^{k-1} \rho_{i_n}}{r^m} \\ & \quad + (k-1)L + L \frac{\sum_{n=1}^{k-1} \rho_{i_n}}{r^m} \\ &= L \frac{r^m - \sum_{n=1}^{k-1} \rho_{i_n}}{\rho_{i_k}} + kL \\ & \geq (k+1)L \end{aligned}$$

where the last inequality follows from the stability condition, $\sum_{n=1}^N \rho_{i_n} \leq r^m$. Since by design, the traffic shapers have $\delta_n^m = 0$, (17) is verified.

We now proceed to derive the end-to-end delay bounds for the connections. Recall that we have assumed zero propagation delays, so from (16) we obtain

$$\bar{D}_n = D(\bar{I}_n \| \mathcal{A}_n) + \sum_{m=1}^M D_n^m.$$

For the delay $D(\bar{I}_n \| \mathcal{A}_n)$, using (6), we have

$$D(\bar{I}_n \| \mathcal{A}_n) = \max_{\tau \geq 0} \left\{ \frac{\bar{I}_n(\tau) - L - \rho_n \tau}{\rho_n} \right\} = \frac{\delta_n}{\rho_n}.$$

Therefore, taking into account (19) we obtain

$$\begin{aligned} \bar{D}_n &= \frac{\delta_n}{\rho_n} + \sum_{m=1}^M \frac{L}{\rho_n} + \sum_{m=1}^M \frac{L}{r^m} \\ &= \frac{\delta_n + ML}{\rho_n} + \sum_{m=1}^M \frac{L}{r^m}. \end{aligned} \quad (20)$$

Since (18) is identical to (20), we see that the proposed RCS discipline π can guarantee the same end-to-end delays as RPPS.

From the above argument we see that if the delay bounds in (18) are required by the connections in the network, then the RCS discipline π , proposed above can be used. It provides the flexibility of easily specifying other delay bounds, whereas the bounds in RPPS are tied to the rate ρ_n of a connection. In addition, since reshaping is performed at each node, buffer requirements will typically be lower than those of RPPS, and its implementation may also be simpler.

If the end-to-end delay requirements of connection n are smaller than (18), a slightly more general version of RPPS can be used. Rather than providing a rate of ρ_n to connection n , better delay performance can be obtained by giving it a rate of $g_n \geq \rho_n$, at each node. The end-to-end delay bound is then given by

$$\bar{D}_n^* = \frac{\delta_n + ML}{g_n} + \sum_{m=1}^M \frac{L}{r^m}. \quad (21)$$

The previous analysis still applies with very little modification and can be used to specify an RCS discipline that guarantees the bounds in (21). In this case, all traffic shapers have envelopes $\bar{A}_n^m(\tau) = L + g_n \tau$ and the delay guarantees at the scheduler of node m are

$$D_n^m = L/g_n + L/r^m.$$

The intuition behind choosing traffic shapers of this kind is as follows. If the RPPS discipline guarantees a clearing rate of g_n to connection n , then somewhere along the path, say at node m , the connection n may only receive a service rate of g_n . This congested link behaves like a traffic shaper that has an envelope of $\bar{A}_n^m(\tau) = L + g_n \tau$. Based on Proposition 2, we know that for an RCS discipline it is beneficial to choose the “smallest” shaper at all the nodes, so that they can all take advantage of the smaller traffic envelope. Since in RPPS the smallest rate that a connection can be given at any node is g_n , a natural choice for the shaper envelopes of the RCS discipline is then $L + g_n \tau$.

In addition to being able to provide the same bounds as RPPS, the RCS discipline also has the advantage of allowing additional connections to be accepted, albeit with looser delay requirements. Specifically, observe that the schedulability check for RPPS is now $\sum_{l \in C_n^m} g_l \leq r^m$, $m = 1, \dots, M$, where C_n^m denotes the set of connections that are multiplexed on the same link as connection n at node m . This implies that some amount of bandwidth *viz.* $r^m - \sum_{l \in C_n^m} \rho_l$, cannot be utilized by RPPS. This bandwidth can be used by an RCS discipline to accept additional connections that require relatively larger end-to-end delay guarantees. At the end of this section we provide a specific example of this benefit of RCS disciplines over the more general GPS disciplines.

B. Achieving GPS Delay Guarantees

In [12, Sec. VIII], tight bounds on per connection packet delays are developed for GPS under a fairly general assignment of weights, ϕ_i^m , called consistent relative session treatment (CRST). These bounds are achieved in certain node configurations, and even in the special case of RPPS they can be much tighter than those provided by (18). However, the calculation of the bounds is much more cumbersome as they take into account the effect of all the other connections along a connection's path. We will show that even with these tight bounds, an RCS discipline can be designed that guarantees the same delay bounds.

To simplify the discussion and to avoid obscuring the main idea of the argument, we assume a continuous flow model, *i.e.*, packetization is not taken into account. Therefore, we consider the GPS policy (instead of PGPS) and assume that the RCS discipline uses the EDF scheduling policy (instead of NPEDF). As far as the design of traffic shapers is concerned, this assumption basically amounts to setting $L = 0$.

Before proceeding with the design of the RCS discipline, we need some preliminary results. Consider a single link on which N connections are multiplexed, and assume that all of them are "greedy," *i.e.*, the amount of connection v traffic, $1 \leq v \leq N$, arriving in the interval $[0, t]$ is $\tilde{\delta}_v + \rho_v t$. Then, the service function (see [11] for a precise definition) of connection n , $S_n(t)$, is the amount of connection n traffic that is served in the interval $[0, t]$. In [11, p. 355] $S_n(t)$ is used to derive delay bounds for connection n traffic whose envelope is $\tilde{\delta}_n + \rho_n \tau$. The next lemma improves these delay bounds for connection n , when it has a smaller envelope $\bar{I}_n(\tau)$ such that $\bar{I}_n(\tau) \leq \tilde{\delta}_n + \rho_n \tau$, $\tau \geq 0$.

Lemma 3: Assume that the connection n traffic satisfies $I_n[t, t+\tau] \leq \bar{I}_n(\tau) \leq \tilde{\delta}_n + \rho_n \tau$, $t, \tau \geq 0$, for every connection n that is multiplexed on a given link. If the system starts empty, then an upper bound on connection n delay under GPS is

$$D_n^* = \max_{\tau \geq 0} \left\{ \min_{t \geq \tau} \{t : S_n(t) \geq \bar{I}_n(\tau)\} - \tau \right\}.$$

The proof can be found in Appendix A. For our purposes, the case where $\bar{I}_n(\tau) = \min\{c_n \tau, \tilde{\delta}_n + \rho_n \tau\}$, $c_n \geq \rho_n$, $\tilde{\delta}_n \leq \bar{\delta}_n$ will be of interest. For convenience, we summarize in the next corollary two specific cases of Lemma 3 that will be useful in the rest of this section. As stated in [11], $S_n(t)$ is

a piecewise linear function, convex in the range $[0, t^B]$, where t^B is the end of the first busy period of connection n , when all the N connections are greedy. In this range, $S_n(t)$ is characterized by the pairs $(s_k, b_k)_{k=1}^{k_n}$ where s_k is the slope of the k th segment, b_k its duration, and k_n is the number of line segments in $S_n(t)$. Because of the convexity of $S_n(t)$ we have

$$s_1 \leq s_2 \leq \dots \leq s_{k_n}.$$

Corollary 1: Assume that the conditions of Lemma 3 hold, so that $\bar{I}_l(\tau) \leq \tilde{\delta}_l + \rho_l \tau$, $\tau \geq 0$, $1 \leq l \leq N$, and furthermore let $\bar{I}_n(\tau) = \min\{c_n \tau, \tilde{\delta}_n + \rho_n \tau\} \leq \bar{\delta}_n + \rho_n \tau$, $c_n \geq \rho_n$, $\tau \geq 0$.

- 1) If $s_1 \geq c_n$, then $D_n^* = 0$.
- 2) If $s_k < c_n$, $k = 1, \dots, j-1$, $s_j \geq c_n$, and

$$S_n \left(\sum_{k=1}^{j-1} b_k \right) \leq e_n := (c_n \tilde{\delta}_n) / (c_n - \rho_n)$$

then $D_n^* = \sum_{k=1}^{j-1} b_k - \eta$, where $\eta = S_n(\sum_{k=1}^{j-1} b_k) / c_n$.

The first part of the corollary follows by observing that $s_1 \geq c_n$ implies that $\bar{I}_n(\tau) \leq S_n(\tau)$ and therefore

$$\min_{t \geq \tau} \{t : S_n(0, t) \geq \bar{I}_n(\tau)\} = \tau.$$

A geometric interpretation of the second part is given in Fig. 5.

The development of GPS bounds for connection n is based on the universal service curve (USC) for that connection [12, Sec. VIII]. Just as $S_n(t)$ characterizes the service that connection n receives at a single node, the USC of a connection characterizes the end to end service that it receives. We summarize here the method by which the USC is obtained when all the nodes use a GPS discipline [12].

- 1) Under a CRST weight assignment, an algorithm is developed by which an envelope function, $\tilde{\delta}_n^m + \rho_n \tau$, is guaranteed for every connection n traffic entering node m [12, p. 142]. For our purposes, it is important to note that

$$\tilde{\delta}_n^1 = \tilde{\delta}_n, \quad \tilde{\delta}_n^m \geq \tilde{\delta}_n, \quad 2 \leq m \leq M. \quad (22)$$

- 2) Given envelope functions of the form $\tilde{\delta}_v^m + \rho_v \tau$, for any connection v that is multiplexed with connection n at node m (v and n are on the same output link at node m), the service function for connection n , $S_n^m(\tau)$, is calculated. Let (s_k^m, b_k^m) , $k = 1, \dots, k_n^m$ be the set of slopes that characterize $S_n^m(\tau)$.
- 3) The USC, $\hat{S}_n(\tau)$, for connection n is given by the formula

$$\hat{S}_n(\tau) = \min \{G_n^M(\tau), \bar{I}_n(\tau)\}$$

where $G_n^M(\tau)$ is defined as ∞ for $\tau > \sum_{m=1}^M \sum_{k=1}^{k_n^m} b_k^m$, and for $\tau \leq \sum_{m=1}^M \sum_{k=1}^{k_n^m} b_k^m$ it is composed of the segments (s_k^m, b_k^m) , $m = 1, \dots, M$, $k = 1, \dots, k_n^m$ of $S_n^m(\tau)$, arranged in a nondecreasing order of slopes [12, p. 144]. We denote by (\hat{s}_k, \hat{b}_k) , $k = 1, \dots, \sum_{m=1}^M k_n^m$ this nondecreasing order.

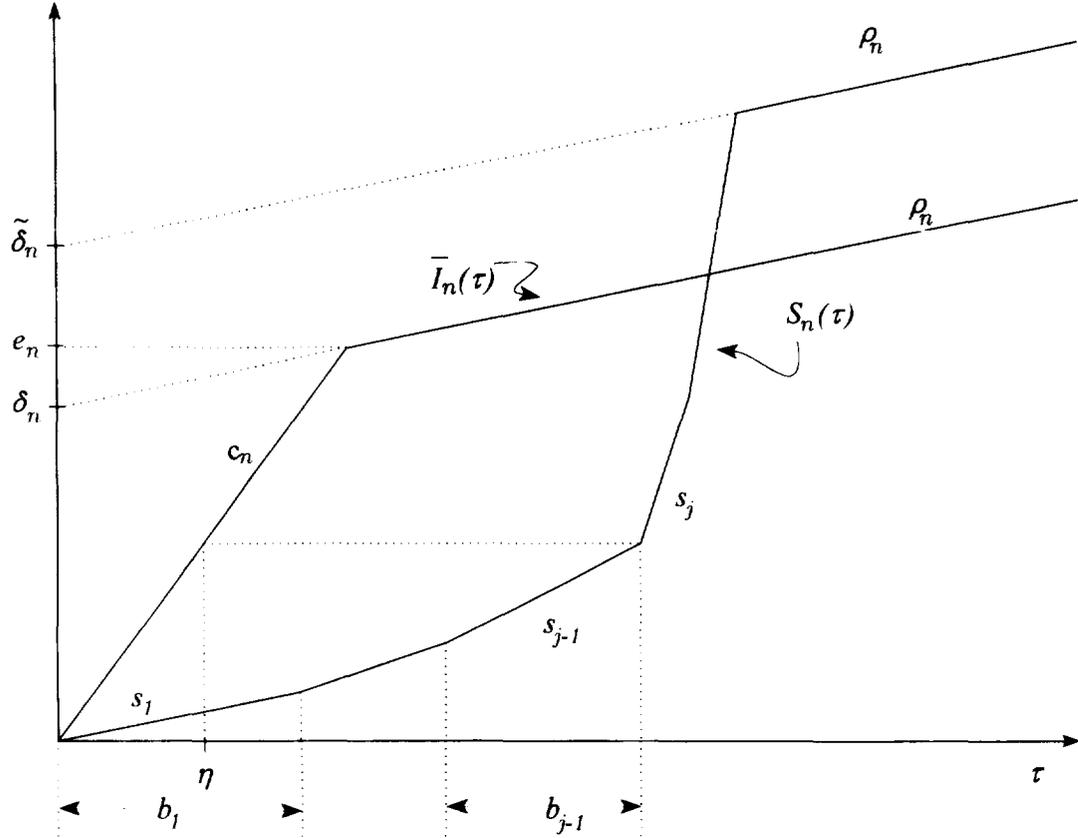


Fig. 5. Delays under GPS.

Let k_q be such that $\hat{s}_{k_q} \geq \rho_n$, and

$$\hat{s}_k < \rho_n, \quad k = 1, \dots, k_q - 1. \quad (23)$$

We are now ready to design an RCS discipline that is at least as good as GPS. Consider first the design of traffic shapers. Recall from the beginning of Section IV-A, that for the purpose of comparison with GPS we assume that the envelope of connection n traffic entering the first traffic shaper is of the form $\bar{I}_n(\tau) = \delta_n + \rho_n \tau$ ($L = 0$). For connection n , at each node m on the path, we choose traffic shapers that have the same envelope, i.e., $\bar{A}_n^m(\tau) = \min\{c_n \tau, \delta_n + \rho_n \tau\}$, $m = 1, 2, \dots, M$. To specify how the parameter c_n is picked, we need to distinguish between two classes of connections.

1) *Class a*): Connection n belongs to this class when

$$\hat{S}_n \left(\sum_{k=1}^{k_q-1} \hat{b}_k \right) < \delta_n \quad (24)$$

where the USC, \hat{S}_n is defined as above. In this case, the delay bound for connection n traffic under GPS is given by the solution of the equation [13, p. 136] [see Fig. 6(a)],

$$\bar{D}_n^* : \hat{S}(\bar{D}_n^*) = \delta_n.$$

Let $k^* \geq k_q$, be the index of the slope of the USC at time \bar{D}_n^* . If at time \bar{D}_n^* there is a change in slope, then define k^* as the index of the smaller of the two slopes (in fact either slope would work). We set $c_n = \hat{s}_{k^*}$.

2) *Class b*): Connection n belongs to this class when

$$\hat{S}_n \left(\sum_{k=1}^{k_q-1} \hat{b}_k \right) \geq \delta_n.$$

In this case, the delay bound for connection n traffic under GPS is [13, p. 136] [see Fig. 7(a)]

$$\bar{D}_n^* = \sum_{k=1}^{k_q-1} \hat{b}_k - \frac{\hat{S}_n \left(\sum_{k=1}^{k_q-1} \hat{b}_k \right) - \delta_n}{\rho_n}.$$

We then set $c_n = \rho_n$.

For connection n , we assign the scheduler delay at node m , D_n^m , to be equal to the maximum delay that would be experienced by the connection under the GPS scheduling policy at that node, when the conditions of Corollary 1 are satisfied. This amounts to the following assignment.

- If at node m , $s_1^m \geq c_n$, then set $D_n^m = 0$.
- If at node m , $s_k^m < c_n$, $k = 1, \dots, j_m - 1$, $s_{j_m}^m \geq c_n$, then assign

$$D_n^m = \sum_{k=1}^{j_m-1} b_k^m - \eta, \quad \text{where } \eta = S_n^m \left(\sum_{k=1}^{j_m-1} b_k^m \right) / c_n.$$

We first establish that the specified delays can be guaranteed by the EDF policy at each node. Instead of using the extension of Theorem 1 to general shaper envelopes, it will be simpler to argue indirectly as follows: We will show that the specified delays are guaranteed when the RCS discipline uses GPS as

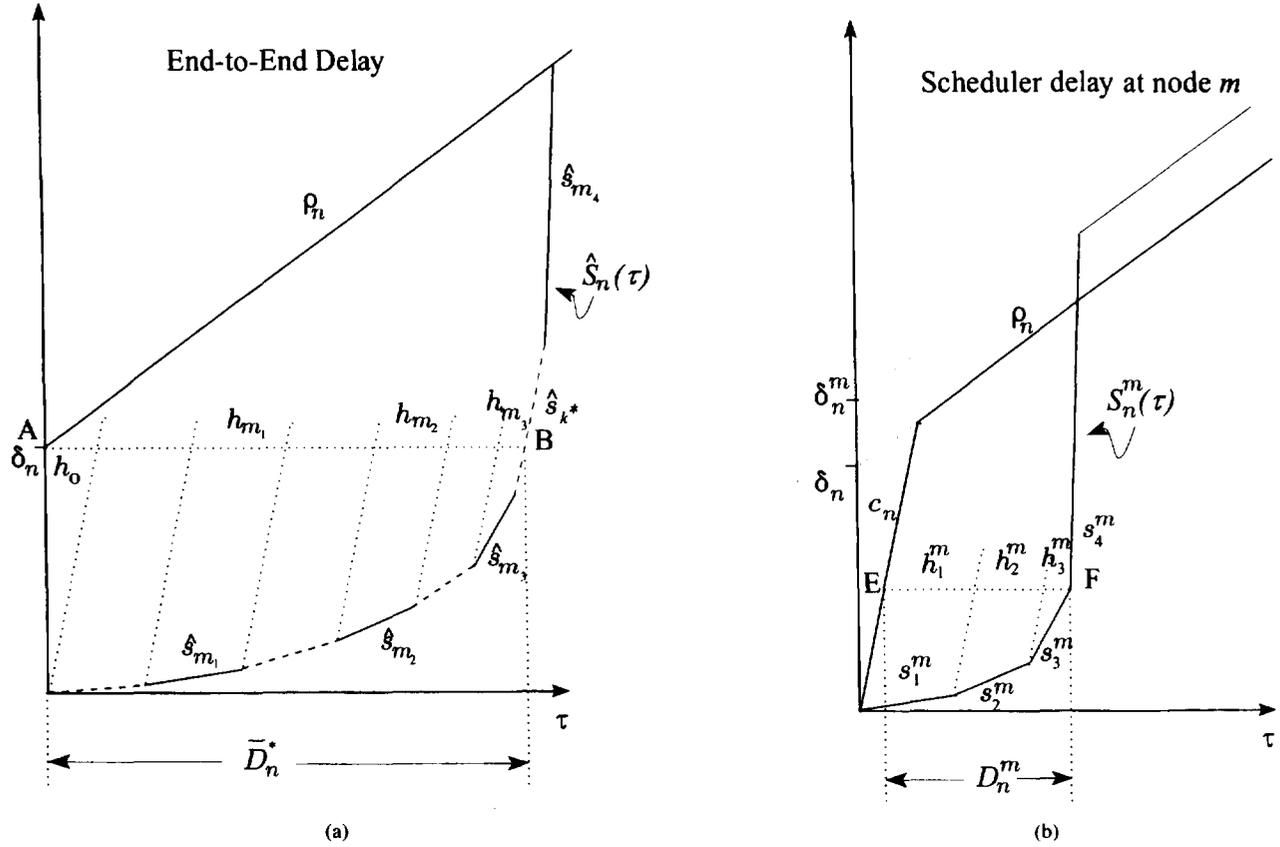


Fig. 6. Delay decomposition of a class a) connection.

the scheduling policy at each node. Since EDF is better than GPS in the single node case, it will follow that the same delay guarantees can, at a minimum, be provided when the EDF scheduling policy is employed.

Observe that according to (22), we have that $\bar{A}_v(\tau) \leq \delta_v^m + \rho_v \tau$ for any connection v that is multiplexed with connection n on the same output link of node m . It is also true that $c_n \geq \rho_n$. This follows by definition for a connection in class b). For a connection in class a), observe that because of (23) and the fact that \hat{s}_k , $k = 1, 2, \dots$, is nondecreasing we have $c_n = \hat{s}_k \geq \hat{s}_{k_q} \geq \rho_n$. Applying Corollary 1 (where we replace $\tilde{\delta}_n \leftarrow \delta_n^m$), we conclude that the delay bound $D_n^m = 0$ can be guaranteed under the GPS policy for any node m for which $s_1^m \geq c_n$. For a node m , where $s_k^m < c_n$, $k = 1, \dots, j_m - 1$, $s_{j_m}^m \geq c_n$, we apply part 2) of Corollary 1 and, therefore, we first need to show that

$$S_n^m \left(\sum_{k=1}^{j_m-1} b_k^m \right) \leq \frac{c_n \delta_n}{c_n - \rho_n}.$$

This is trivially true for a connection in class b) since $c_n \delta_n / (c_n - \rho_n) = \infty$. If connection n belongs to class a), observe that from the definition of \hat{s}_k , j_m and $\hat{S}_n(\tau)$, we have (see Fig. 6)

$$\begin{aligned} S_n^m \left(\sum_{k=1}^{j_m-1} b_k^m \right) &\leq \hat{S}_n \left(\sum_{k=1}^{k^*-1} \hat{b}_k \right) \\ &\leq \delta_n \\ &\leq \frac{\hat{s}_{k^*} \delta_n}{\hat{s}_{k^*} - \rho_n}. \end{aligned}$$

Thus, we have established that in both cases a) and b), the specified delay bound can be guaranteed at node m .

Next, we need to establish that the end-to-end delay guarantee of the RCS discipline as given by (16), does not exceed \bar{D}_n^* . Recall that the input traffic envelope for connection n , $\bar{I}_n(\tau) = \delta_n + \rho_n \tau$, and so from (8), the delay in the first traffic shaper is

$$D(\bar{I}_n \| A_n) = \frac{\delta_n}{c_n}.$$

Therefore, it suffices to show that

$$\frac{\delta_n}{c_n} + \sum_{m=1}^M D_n^m = \bar{D}_n^*. \quad (25)$$

Let M_0 be the set of nodes for which $D_n^m > 0$. Obviously then, $\sum_{m=1}^M D_n^m = \sum_{m \in M_0} D_n^m$. Assume first that connection n belongs to class a). Observe that the set of slopes \hat{s}_k , $k = 1, \dots, k^* - 1$, can be partitioned into subsets F_m , $m \in M_0$, where

$$F_m = \{\hat{s}_l : \hat{s}_l = s_k^m, \text{ for some } k = 1, \dots, j_m - 1\}.$$

We denote by m_k the index l for which $\hat{s}_l = s_k^m$, i.e., $\hat{s}_{m_k} = s_k^m$. For the rest of the discussion, it is best to use geometric arguments. Referring to Fig. 6(a), draw lines with slope \hat{s}_k from all the points in $\hat{S}_n(\tau)$ where the slope changes and remains less than \hat{s}_{k^*} . These lines intersect segment AB (which corresponds to the delay \bar{D}_n^*) and divide it into segments of length h_k , $0 \leq k \leq k^* - 1$, where segment h_k corresponds to slope \hat{s}_k , $1 \leq k \leq k^* - 1$. Denote by h_{m_k}

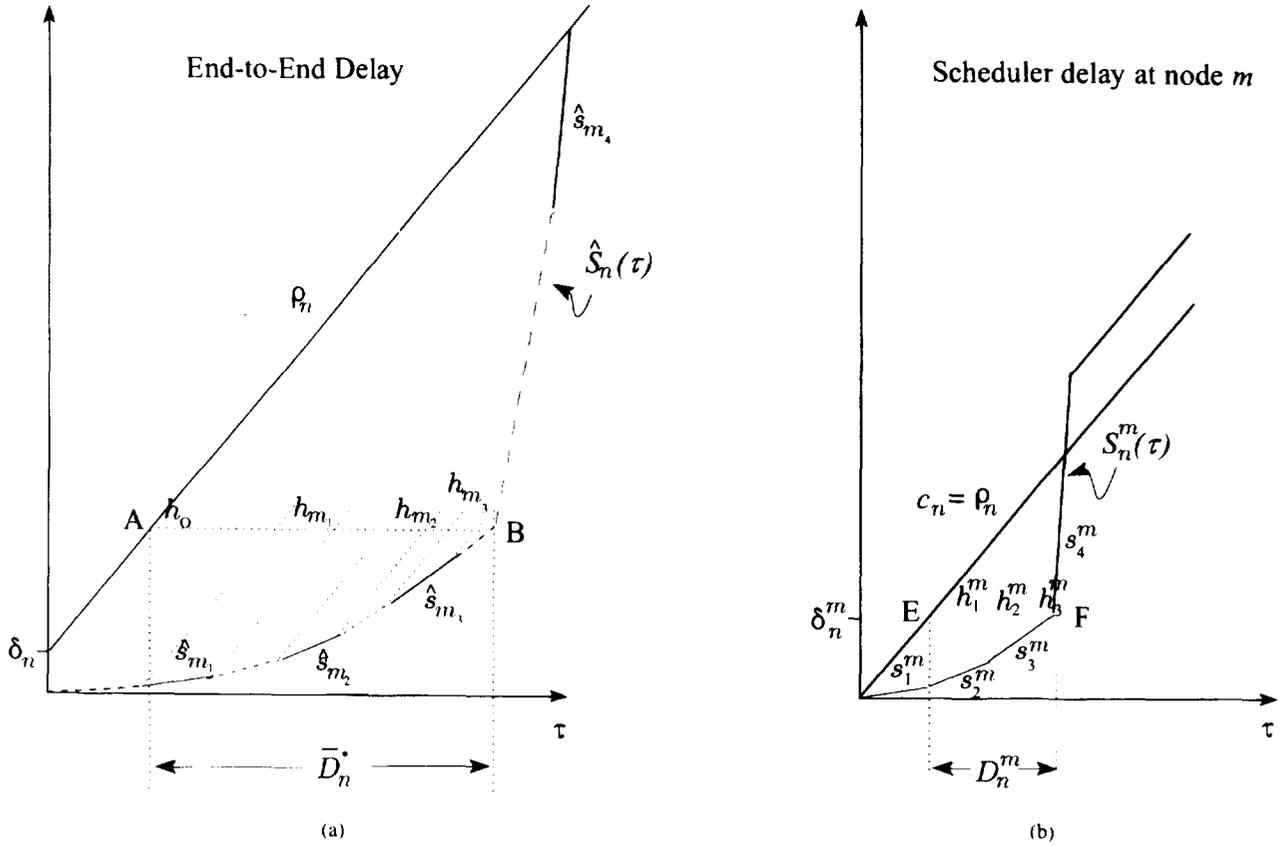


Fig. 7. Delay decomposition of a class b) connection.

the segment that corresponds to \hat{s}_{m_k} . Since by construction $h_0 = \delta_n / c_n$, we have

$$\bar{D}_n^* = \frac{\delta_n}{c_n} + \sum_{m \in M_0} \sum_{k=1}^{j_m-1} h_{m_k}. \quad (26)$$

Similarly, in Fig. 6(b), draw lines with slope \hat{s}_k from all the points in $S_n^m(\tau)$ where the slope changes and remains less than \hat{s}_k . These lines intersect segment EF (which corresponds to the delay D_n^m) and divide it into segments h_k^m , $1 \leq k \leq j_m - 1$ (in the figure we have $j_m - 1 = 3$). We can then write

$$D_n^m = \sum_{k=1}^{j_m-1} h_k^m. \quad (27)$$

Using the facts that $\hat{s}_{m_k} = s_k^m$ and that $\hat{b}_{m_k} = b_k^m$, it can be easily seen that $h_{m_k} = h_k^m$. Taking into account (26) and (27), we conclude the correctness of (25).

Similar arguments can be made for a connection that belongs to class b). The main difference is that we now draw lines with slope ρ_n . Fig. 7 illustrates the construction in this case.

Note 1: The above derivations established that an RCS discipline that provides the same delay bounds as GPS can be constructed, but the arguments used were more involved than for the simpler case of RPPS. As a result, it is much harder to gain some intuition into why and how this is achieved. A possible (and partial) explanation is that the reshaping peak rate, c_n , for connection n , should be set to the service rate in the USC of the GPS policy, that corresponds to the maximum

delay value. Using a larger value will not help since service, and hence reshaping, at that rate will be encountered. Using a smaller value will result in higher delays.

Note 2: In the course of the previous argument, we showed that the delay guarantees provided by a pure GPS policy can also be achieved by an RCS discipline working with worst case delays at each node, where the scheduling policy at each node is GPS. If we replace GPS with the (simpler) EDF scheduling policy at each node, we are not only assured that we can still guarantee the GPS end-to-end delay bounds, but we also create a service discipline that is better than GPS. This is due to the fact that in the single node case, EDF is better than GPS [8]. That is, there are delay vectors that can be guaranteed by EDF but cannot be guaranteed by GPS no matter what weights are chosen. For example, consider a link of capacity r , where two connections are multiplexed and $\bar{I}_n(\tau) = \delta_n + \rho_n \tau$, $n = 1, 2$, with $\rho_1 + \rho_2 \leq r$. Using Theorem 1 with $L = 0$, we can see that the delays that can be guaranteed by the EDF policy are

$$D_1 = \frac{\delta_1}{r}, \quad D_2 = \frac{\delta_2}{r - \rho_1} + \frac{\delta_1}{r}.$$

For GPS on the other hand, it can be seen from the construction in [11, Sec. VI-C], that in order to guarantee $D_1^* = \delta_1 / r$ we need to specify $\phi_2 = 0$ and then the minimum guaranteed delay for connection 2 is

$$D_2^* = \frac{\delta_2}{r - \rho_1} + \frac{\delta_1}{r - \rho_1}. \quad (28)$$

The difference between the GPS and EDF delay guarantees for connection 2 is

$$D_2^* - D_2 = \frac{\rho_1 \delta_1}{r(r - \rho_1)}$$

which can be quite large. Similar examples can be given for the packetized model when comparing PGPS to NPEDF. The better bounds of EDF in this simple example, are essentially a reflection of the fact that, in the single node case, EDF is the optimal policy. This is in part due to EDF's ability to, unlike GPS (or its variants), decouple delay and rate guarantees. In the above example, this difference is expressed in the $\delta_1/(r - \rho_1)$ term of (28). This term reflects the behavior of GPS, which serves all new packets of connection 1 at rate r , irrespective of the fact that they may have just arrived and, therefore, are in no danger of being excessively delayed. In contrast, the EDF policy exploits this knowledge to improve the delay guarantee it gives to connection 2. In the multiple node case, the benefit of decoupling delay and rate guarantees is still obtained, while the problem of summing up worst case node delays has been alleviated by suitably reshaping the connection traffic.

To summarize, in this section, we have shown how "proper" selection of the traffic shapers allows us to construct an RCS discipline that outperforms GPS. In the next section, we provide results that can be used to narrow the search for "good" shapers for RCS disciplines.

V. TRAFFIC SHAPER PROPERTIES

In this section, we discuss some interesting properties of traffic shapers in the context of RCS disciplines. First, we consider the problem of constructing the "smallest" shaper that causes a specified maximum delay on the input traffic $\bar{I}_n(\tau)$. Specifically, given $d \geq 0$, we want to construct a shaper $\mathcal{A}_n(d)$ such that $D(\bar{I}_n \| \mathcal{A}_n(d)) \leq d$ with the additional requirement that $\mathcal{A}_n(d) \preceq \mathcal{A}$, for any shaper \mathcal{A} that satisfies $D(\bar{I}_n \| \mathcal{A}) \leq d$. Recall that $\bar{U}_n(\tau)$ denotes the input traffic envelope of connection n before the first packetizer in the network (see Fig. 1). We further assume that the input traffic envelope, $\bar{U}_n(\tau)$, is an increasing, concave, piecewise linear function with a finite number of slopes. In Appendix B, we show that these assumptions on the input traffic envelope do not entail any essential loss of generality. We can write $\bar{U}_n(\tau)$ in the form (see Fig. 8)

$$\bar{U}_n(\tau) = \min_{k=1, \dots, K} \{\delta_{n,k} + \rho_{n,k} \tau\} \quad (29)$$

where $\rho_{n,k} > \rho_{n,k+1}$, $\delta_{n,k} < \delta_{n,k+1}$, and when $K \geq 3$

$$\frac{\delta_{n,k} - \delta_{n,k-1}}{\rho_{n,k-1} - \rho_{n,k}} < \frac{\delta_{n,k+1} - \delta_{n,k}}{\rho_{n,k} - \rho_{n,k+1}}, \quad k = 2, \dots, K-1.$$

Let $\tau_{n,1} = 0$ and $\tau_{n,k} = (\delta_{n,k} - \delta_{n,k-1})/(\rho_{n,k-1} - \rho_{n,k})$, $2 \leq k \leq K$. At the point $P_k = (\tau_{n,k}, \delta_{n,k} + \rho_{n,k} \tau_{n,k})$, the slope of the envelope, $\bar{U}_n(\tau)$, changes from $\rho_{n,k-1}$ to $\rho_{n,k}$. According to (1), the envelope of the traffic entering the first shaper is

$$\bar{I}_n(\tau) = L + \min_{k=1, \dots, K} \{\delta_{n,k} + \rho_{n,k} \tau\}.$$

Now, let $\bar{A}(\tau) = L + \min_{j=1, \dots, J} \{\delta_j + \rho_j \tau\}$ be the envelope of \mathcal{A} . According to (8), $D(\bar{I}_n \| \mathcal{A}) = \infty$ when $\min_{j=1, \dots, J} \{\rho_j\} < \rho_{n,K}$, while $D(\bar{I}_n \| \mathcal{A}) \leq \delta_{n,K}/\rho_{n,K}$ whenever $\min_{j=1, \dots, J} \{\rho_j\} \geq \rho_{n,K}$. Therefore, it is sufficient to restrict our attention to the range $0 \leq d \leq \delta_{n,K}/\rho_{n,K}$. For the next proposition, it will be helpful to refer to Fig. 8.

Proposition 3: Let $0 \leq d \leq \delta_{n,K}/\rho_{n,K}$ and let k^* be given by

$$k^* = \min_{k=1, \dots, K} \{k : \bar{U}_n(\tau_{n,k}) - \rho_{n,k}(\tau_{n,k} + d) \geq 0\}.$$

Then, the envelope of the smallest shaper $\mathcal{A}_n(d)$, such that $D(\bar{I}_n \| \mathcal{A}_n(d)) \leq d$ is

$$\bar{A}_n(d)(\tau) = L + \bar{a}_n(d)(\tau)$$

where denoting $c_n^*(d) := \bar{U}_n(\tau_{n,k^*})/(\tau_{n,k^*} + d)$, we set

$$\bar{a}_n(d)(\tau) = \begin{cases} c_n^*(d)\tau, & \text{if } 0 \leq \tau < \tau_{n,k^*} + d \\ \bar{U}_n(\tau - d), & \text{if } \tau \geq \tau_{n,k^*} + d, \end{cases}$$

Proof: From Fig. 8, it can be seen that k^* is the smallest index k such that the line with slope $\rho_{n,k}$ passing through the point $Q_k = (\tau_{n,k} + d, \bar{U}_n(\tau_{n,k}))$ has a nonnegative y -intercept. Observe that the index k^* always exists since

$$\begin{aligned} \bar{U}_n(\tau_{n,K}) - \rho_{n,K}(\tau_{n,K} + d) &= \delta_{n,K} + \rho_{n,K} \tau_{n,K} - \rho_{n,K}(\tau_{n,K} + d) \\ &= \delta_{n,K} - \rho_{n,K} d \geq 0. \end{aligned}$$

Next, we show that $\bar{A}_n(d)(\tau)$ corresponds to a shaper envelope function. For this, it is sufficient to show that $\bar{A}_n(d)(\tau)$ is a concave function, which will follow by construction, if we show that $c_n^*(d) \geq \rho_{n,k^*}$; but this is a consequence of the definition of k^* . To show that $D(\bar{I}_n \| \mathcal{A}_n(d)) \leq d$, recall that according to (10) we can write

$$D(\bar{I}_n, \mathcal{A}_n(d)) = \max_{\tau \geq 0} \{(\bar{A}_n^{(-1)}(d)(\bar{I}_n(\tau)) - \tau)^+\}$$

and that by construction, $\bar{A}_n^{(-1)}(d)(\bar{I}_n(\tau)) - \tau \leq d$, for all $\tau \geq 0$. Finally, we need to show that $\bar{A}_n(d) \preceq \mathcal{A}$, for any other shaper \mathcal{A} that satisfies $D(\bar{I}_n \| \mathcal{A}) \leq d$. To see this, observe that if $\bar{A}_n(d)(\tau) > \bar{A}(\tau)$ for some $\tau \geq \tau_{n,k^*} + d$, then $\bar{A}^{(-1)}(\bar{A}_n(d)(\tau)) > \tau$. Also, by construction, $\bar{A}_n(d)(\tau) = \bar{I}_n(\tau - d)$, $\tau \geq \tau_{n,k^*} + d$. From (10), for all $\tau \geq \tau_{n,k^*} + d$ we have

$$\begin{aligned} D(\bar{I}_n, \mathcal{A}) &\geq \bar{A}^{(-1)}(\bar{I}_n(\tau - d)) - (\tau - d) \\ &> \tau - (\tau - d) = d \end{aligned}$$

a contradiction. Therefore, $\bar{A}_n(d)(\tau) \leq \bar{A}(\tau)$ for all $\tau \geq \tau_{n,k^*} + d$. Using the inequalities $\bar{A}_n(d)(0) = L \leq \bar{A}(0)$, $\bar{A}_n(d)(\tau_{n,k^*} + d) \leq \bar{A}(\tau_{n,k^*} + d)$ and the concavity of $\bar{A}(\tau)$, we conclude that we also have $\bar{A}_n(d)(\tau) \leq \bar{A}(\tau)$ for $0 \leq \tau \leq \tau_{n,k^*} + d$. \square

Using Lemma 2 and Proposition 3, we conclude that given the input envelope function $\bar{I}_n(\tau)$, and a maximum shaper delay d , it is sufficient to restrict our attention to RCS disciplines that use shapers with envelopes of the form $\bar{A}_n(d)(\tau)$. We state this as a Corollary.

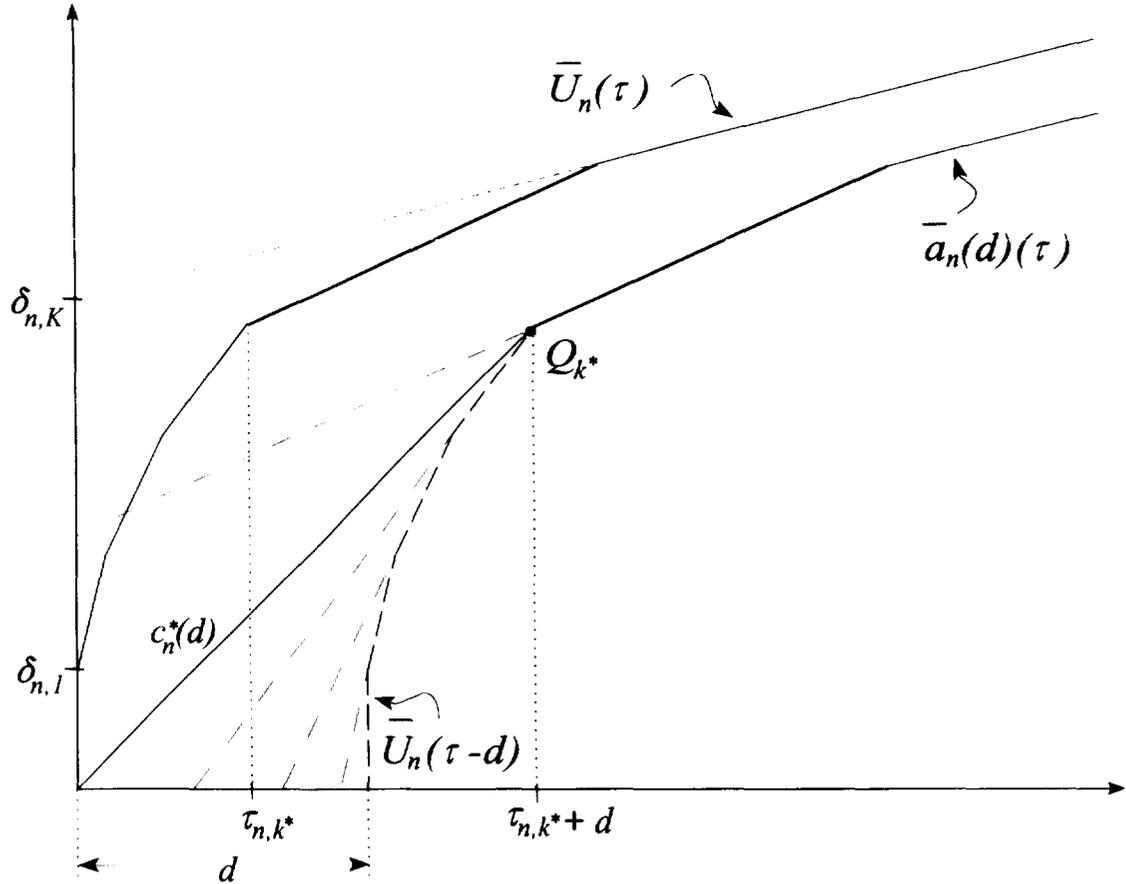


Fig. 8. Construction of smallest envelope function.

Corollary 2: Given an RCS discipline that for connection n uses shaper \mathcal{A} at all nodes, the RCS discipline that uses the shaper with envelope $\bar{A}_n(d)(\tau)$, where $d = D(\bar{I}_n \parallel \mathcal{A})$, can guarantee the same end-to-end delays to all connections.

From the above discussion we see that given $\bar{I}_n(\tau)$, the search for the appropriate shaper envelope, is reduced to the one-parameter family $\bar{A}_n(d)(\tau)$. We can further constrain the range of the parameter d by taking into account the link speeds along the path of connection n . In the next proposition, we show that it is sufficient to restrict attention to envelopes $\bar{A}_n(d)(\tau)$ whose maximum slope $c_n^*(d)$ (peak rate) is not larger than the minimum link speed along the path of the connection.

Proposition 4: Consider connection n with input traffic envelope $\bar{U}_n(\tau)$ that traverses nodes $1, \dots, M$ with corresponding output link speeds r^m . Then, given an RCS discipline that uses shaper envelope $\bar{A}_n(d)(\tau)$, there is an RCS discipline using shaper envelope $\bar{A}_n(d')(\tau)$, $d' \geq d$ which guarantees the same end-to-end delays to all connections and whose peak rate $c_n^*(d')$ is such that

$$\rho_{n,K} \leq c_n^*(d') \leq \min\{r^{\min}, c_n^*(d)\} \leq \min\{r^{\min}, c_n\}$$

where $r^{\min} = \min_{m=1, \dots, M}\{r^m\}$ and c_n is the peak rate of $\bar{U}_n(\tau)$, i.e., $c_n = \rho_{n,1}$ if $\delta_{n,1} = 0$ and $c_n = \infty$ otherwise.

Proof: Observe first that by the design of $\bar{A}_n(d)(\tau)$, for all d , $0 \leq d \leq \delta_{n,K}/\rho_{n,K}$, we have $\rho_{n,K} \leq c_n^*(d)$ and

$$c_n^*(d) \leq c_n.$$

Denote by $U_n^{m+1}[t, t + \tau]$ the connection n traffic entering node $m + 1$ in the interval $[t, t + \tau]$. Then, since the output link of node m has speed r^m

$$U_n^{m+1}[t, t + \tau] \leq r^m \tau.$$

Therefore, for the traffic exiting the packetizer at node $m + 1$, we have [see (1)]

$$I_n^{m+1}[t, t + \tau] \leq L + r^m \tau =: \bar{B}^m(\tau).$$

Therefore, we can replace $\mathcal{A}_n(d)$ with $B^m \wedge \mathcal{A}_n(d)$, without altering the shaper delay or the scheduler delay at node $m + 1$, $m = 1, \dots, M - 1$. Also, by introducing a shaper with envelope $B^M(\tau)$ at the exit point of connection n , we do not affect the end-to-end delay guarantees. Using Proposition 2, we conclude that the delay guarantees are not affected if we employ the RCS discipline that uses the shaper

$$B_n = \bigwedge_{m=1}^M B^m \wedge \mathcal{A}_n(d)$$

in place of $\mathcal{A}_n(d)$ at each node along the path of connection n . But then, for the peak rate c'_n of the envelope of shaper B_n , we have

$$c'_n \leq \min\{r^{\min}, c_n^*(d)\} \leq \min\{r^{\min}, c_n\}.$$

Let $d' = D(\bar{I}_n \parallel B_n)$. Using Corollary 2, we can replace B_n with shaper $\mathcal{A}_n(d')$ without altering the delay guarantees for

any connection. Since by design $\mathcal{A}_n(d') \preceq \mathcal{B}_n$, we must have $c_n^*(d') \leq c'_n$ and the proposition follows. \square

In the important special case of shapers used in the ATM standards [1] and those proposed for the Internet as well [16], we have

$$\bar{U}_n(\tau) = \min\{c_n\tau, \delta_n + \rho_n\tau\}, \quad c_n > \rho_n.$$

In this case, $\tau_{n,2} = \delta_n/(c_n - \rho_n)$, $k^* = 2$

$$\bar{a}_n(d)(\tau) = \begin{cases} \tau \frac{c_n \delta_n}{\delta_n + d(c_n - \rho_n)}, & \text{if } 0 \leq \tau < \tau_{n,2} + d \\ \delta_n + \rho_n(\tau - d), & \text{if } \tau \geq \tau_{n,2} + d. \end{cases}$$

and the range of d is determined by the inequalities

$$\rho_n \leq \frac{c_n \delta_n}{\delta_n + d(c_n - \rho_n)} \leq \min\{r^{\min}, c_n\}.$$

Therefore, to specify an RCS discipline, one has to determine the single parameter d as well as the scheduler delays, D_n^m for each node m along the path of connection n . The determination of these parameters is an interesting design problem, which is the subject of ongoing research and is not addressed in this paper. However, note that a simple and reasonable choice is to use the parameters induced by the GPS (or RPPS) policy. Such a choice gives reasonably good performance, i.e., as good as GPS, and still allows an RCS discipline to accept some additional calls.

The use of traffic shapers at each hop can introduce extra delays for the traffic of connection n , even if there is no congestion in the network. While this leads to a reduction of jitter and buffer requirements at each node in the network, there may be instances where the resulting increase in the average delay is undesirable. In the next section we describe some simple modifications to the RCS discipline that make it work conserving, without compromising the end-to-end delay guarantees that can be provided.

VI. WORK CONSERVING SYSTEM

In the RCS disciplines considered in this paper and in [18], the link schedulers by themselves are assumed to be work-conserving. However, if we consider the traffic shaper and the scheduler as a single system, it is evident that this system is not work-conserving since there may be instants in time when the link is idle even though there are packets waiting in the system (actually the shaper). In what follows, we outline a modification to the system that will make it work conserving, and yet provide the same delay guarantees to all connections. As a result, the output link will no longer be idle when there are packets in the system, thus improving the average delay seen by the packets. A similar approach and motivation can be found in [10] for a system where reshaping is performed based on timestamps carried in each packet, and in [6] for the AIRPORT policy proposed in that paper.

To clarify the exposition, we use the model of [10] to represent both the shaper and the scheduler at an output link of node m . Instead of holding up the packets in the shaper, we maintain two queues in the system: Q_e^m is a queue of packets that are eligible for scheduling, i.e., have been reshaped, and Q_γ^m is a queue of not yet eligible packets. Eligibility is determined by the shaper which stamps an eligibility time on the packets and enqueues them in Q_γ^m . The *eligibility time*

is the earliest time the packet could have left the shaper, for the output of the shaper to be in conformance with the pre-specified traffic envelope. The delay of the packet in the scheduler is calculated with respect to its eligibility time. When a packet in Q_γ^m becomes eligible for scheduling, viz. its eligibility time equals the current time, it is promoted to Q_e^m . The scheduler in the nonwork conserving discipline, π_{NW} , only selects packets in Q_e^m for transmission on the output link. Once a packet has completed its transmission, it is removed from Q_e^m and the scheduler repeats the above process.

Packets from each of the connections at node m enter Q_e^m in conformance with their respective traffic envelopes. The call admission criteria ensure that packets in Q_e^m can be scheduled without violating their deadlines. Note that π_{NW} is nonwork conserving since there can be packets queued in Q_γ^m , that are not considered for transmission by the scheduler, even though the link may be idle.

We now develop a work-conserving discipline π_W , by modifying the scheduler in π_{NW} as follows: There is a nonpreemptive priority mechanism that arbitrates the service of queues Q_e^m and Q_γ^m , with Q_e^m being the higher priority. Ineligible packets from Q_γ^m are selected for transmission (in any order) only when Q_e^m is empty. When there are packets in Q_e^m , the same scheduling discipline as that used in π_{NW} is used to select packets for transmission.

We can specify the operation of the scheduler in π_W during periods when Q_e^m is nonempty in a more precise manner. If Q_e^m is nonempty at time t , define a Q_e^m -busy period at t to be the largest closed interval containing t , during which Q_e^m is nonempty. Let t_0 denote the start of one such busy period. Note that at time t_0 it is possible that an ineligible packet is being served, in which case let t_s denote the time that the ineligible packet begins transmission; otherwise define $t_s := t_0$. Let q be the packet that begins transmission at time t_s . Consider the sequence of packet arrivals consisting of packet q , whose arrival time is set to t_s , along with the other packets that arrive to Q_e^m during the corresponding Q_e^m -busy period. Assume that this sequence is fed to the scheduler in π_{NW} with packet q being the first packet to ever arrive at that scheduler. The scheduler in π_W then schedules this sequence of packets in the same manner as the scheduler in π_{NW} . Note that if the scheduler in π_{NW} is NPEDF (FCFS, PGPS, fixed-priority scheduler, etc.), then during a Q_e^m -busy period the corresponding scheduler in π_W is again NPEDF (FCFS, PGPS, fixed-priority scheduler, etc.). The next proposition shows that the end-to-end delay guarantees are not affected when the service discipline at each node is modified to be work-conserving, as defined above.

Proposition 5: Let the nodes traversed by connection n be numbered $1, \dots, M$. The above modification to the service discipline does not increase the guaranteed upper bound on the end-to-end delays. If \bar{D}_n is the end-to-end delay guarantee for connection n , we still have

$$\begin{aligned} \bar{D}_n \leq & D(\bar{I}_n \| \mathcal{A}_n^1) + \sum_{m=1}^{M-1} D(\mathcal{A}_n^m \| \mathcal{A}_n^{m+1}) \\ & + \sum_{m=1}^M D_n^m + \sum_{m=1}^M T_n^m. \end{aligned}$$

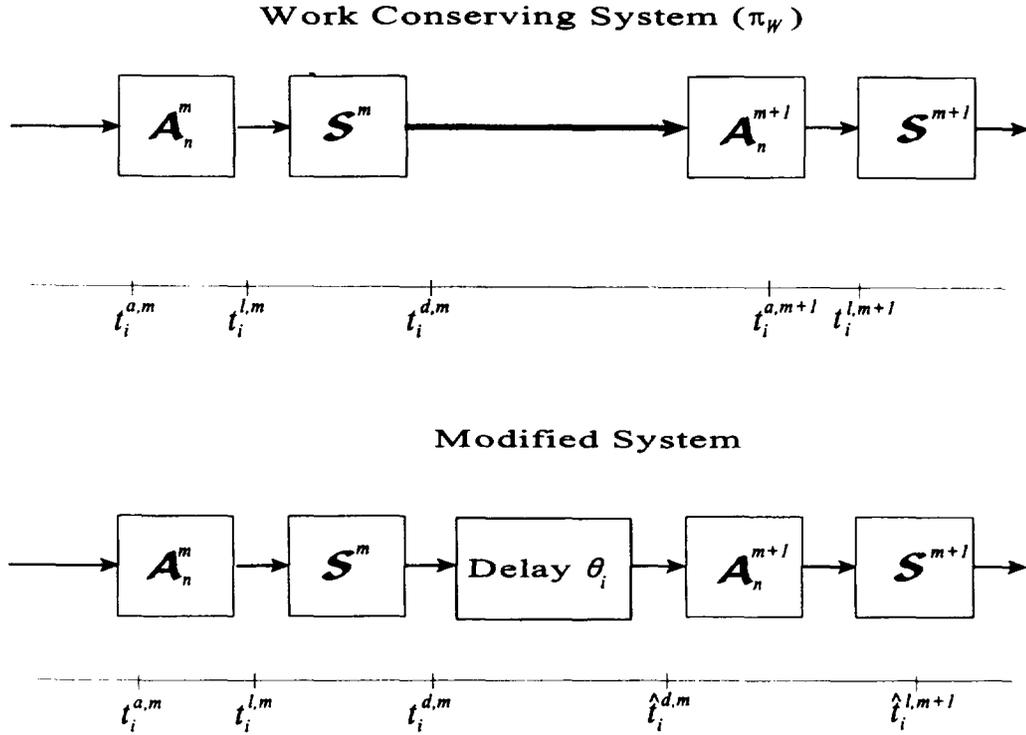


Fig. 9. Original (work conserving) system and the modified system.

Proof: Assume for clarity in the exposition that the propagation delays, $T_n^{l,m}$, $1 \leq m \leq M$, are zero. We first establish that with the above modification to the service discipline, the scheduler delays at node m , $1 \leq m \leq M$, are still bounded above by D_n^m .

Lemma 4: Under discipline π_W , packets of any connection n , are not delayed by more than D_n^m at the scheduler in node m , $1 \leq m \leq M$.

The proof of the above lemma can be found in Appendix A. We denote by $t_i^{l,m}$, the timestamp with which the i th packet is enqueued in Q_n^m ; $t_i^{l,m}$ is the time that the i th packet would leave shaper A_n^m in conformance with the traffic envelope $\bar{A}_n^m(\tau)$. The time at which the packet leaves Q_n^m (to be transmitted on the link or promoted to Q_n^m) is denoted by $t_i^{a,m}$. If the link is idle, the packet may be transmitted before it becomes eligible, i.e., $t_i^{a,m} \leq t_i^{l,m}$. The departure time of the i th packet from the scheduler is denoted as $t_i^{d,m}$. Similarly, let $\tau_{i,a}$ be the arrival time of the i th packet of connection n to the first traffic shaper, and let $\tau_{i,d}$ be the time it arrives at its destination. Since $t_n^{a,m} \leq t_n^{l,m}$, $1 \leq m \leq M$, we can write

$$\begin{aligned} \tau_{i,d} - \tau_{i,a} &\leq t_n^{l,M} - \tau_{i,a} + \tau_{i,d} - t_n^{a,M} \\ &= \sum_{m=1}^{M-1} (t_n^{l,m+1} - t_n^{l,m}) + t_n^{l,1} - \tau_{i,a} + \tau_{i,d} - t_n^{a,M}. \end{aligned}$$

Since by definition, $t_n^{l,1} - \tau_{i,a} \leq D(\bar{I}_n \| A_n^1)$, and according to Lemma 4, $\tau_{i,d} - t_n^{a,M} \leq D_n^M$, it suffices to show that

$$t_i^{l,m+1} - t_i^{l,m} \leq D_n^m + D(A_n^m \| A_n^{m+1}), \quad 1 \leq m \leq M-1. \quad (30)$$

Let S^m be the system consisting of the scheduler at node m . Consider the modified system which is same as the work

conserving system operating under π_W except for a delay system inserted between S^m and shaper A_n^{m+1} as shown in Fig. 9. The delay system delays packet i by $\theta_i = D_n^m + t_i^{l,m} - t_i^{d,m}$; therefore, packet i departs the delay system at time $\hat{t}_i^{d,m} = D_n^m + t_i^{l,m}$. First we verify that $\theta_i \geq 0$.

$$\begin{aligned} \theta_i &= D_n^m + t_i^{l,m} - t_i^{d,m} \\ &\geq D_n^m + t_i^{a,m} - t_i^{d,m} \end{aligned} \quad (31)$$

$$\geq 0. \quad (32)$$

Inequality (31) follows because packets never depart the shaper later than they are supposed to, i.e., $t_i^{l,m} \geq t_i^{a,m}$, and (32) follows from Lemma 4. Let $\hat{t}_i^{l,m+1}$ be the timestamp with which packet i is enqueued in Q_n^{m+1} in the modified system. From Lemma 1, we conclude that

$$t_i^{l,m+1} - t_i^{d,m} \leq \hat{t}_i^{l,m+1} - t_i^{d,m}. \quad (33)$$

Adding $t_i^{d,m} - t_i^{l,m}$ to both sides of (33) we have

$$\begin{aligned} t_i^{l,m+1} - t_i^{l,m} &\leq \hat{t}_i^{l,m+1} - t_i^{l,m} \\ &= \hat{t}_i^{l,m+1} - \hat{t}_i^{d,m} + \hat{t}_i^{d,m} - t_i^{l,m}. \end{aligned}$$

Since for all i , $\hat{t}_i^{d,m} = t_i^{l,m} + D_n^m$, the traffic exiting the delay system has envelope $\bar{A}_n^m(\tau)$. Therefore

$$\hat{t}_i^{l,m+1} - \hat{t}_i^{d,m} \leq D(A_n^m \| A_n^{m+1}).$$

It follows that

$$t_i^{l,m+1} - t_i^{l,m} \leq D(A_n^m \| A_n^{m+1}) + D_n^m$$

as desired. When $T_n^{l,m} > 0$, the same reasoning applies, provided that system S^m consists of the scheduler at node

m , and the link l_m , i.e., the bound on the delay at S^m is now $D_n^m + T_n^{l_m}$. \square

We have described a general procedure for making the RCS disciplines work conserving that applies regardless of the scheduling discipline that is used—of course, the scheduling discipline, by itself needs to be work conserving. In keeping with this generality, we have not specified the exact order in which packets from the ineligible queue can be serviced. Based on the scheduling discipline that is used, as well as the behavior that is desired for the flows in general, some specific order may be appropriate. This is another topic and is the subject of ongoing and future work. Note that while the modification introduced in this section to make the RCS disciplines work-conserving does not affect delay bounds, it will affect both jitter and buffer requirements and this may be undesirable in some cases. For example, jitter sensitive applications may be better served using the nonwork-conserving version.

VII. CONCLUSION

In this paper, we have established that RCS disciplines offer a powerful solution to provide end-to-end delay and throughput guarantees in high speed networks. We showed that the main disadvantage of these service disciplines, namely that of summing worst case delays at each node to determine end-to-end delay bounds, can be overcome through “proper” reshaping of the source traffic. In particular, we have shown that controlling the peak rate of a connection as a function of its delay requirements is critical to efficient network quality-of-service (QoS) provisioning. How to perform this reshaping was also investigated in the paper, and illustrated by designing RCS disciplines that outperform GPS. This is significant since the bounds available for these policies take dependencies between nodes into account.

In addition to their efficiency, RCS disciplines are also relatively simple to implement, and offer the flexibility to accommodate a wide range of implementation constraints. For example, it is possible to use different schedulers and shapers at different nodes depending on the capabilities available locally. Furthermore, because we also showed that delay guarantees are not affected when operating in a work conserving manner, i.e., reshaping traffic only in case of congestion, RCS disciplines also enable us to offer low average delays when the network is not congested. Finally, note that the greater flexibility of RCS disciplines introduces new and interesting problems, e.g., how to best split a given end-to-end delay budget into local delay bounds, and addressing them is the topic of ongoing work.

APPENDIX A LEMMA PROOFS

Proof of Lemma 1: Let $\bar{A}(\tau)$ denote the envelope of the shaper with $\bar{A}(0) \geq L$. Also, let $s_i^{(1)}, s_i^{(2)}$ denote the arrival times, and $f_i^{(1)}, f_i^{(2)}$ denote the departure times of the i th packet at the shapers of system S_1 and S_2 , respectively (see Fig. 3). By definition, $s_i^{(2)} = s_i^{(1)} + \theta_i$, with $\theta_i \geq 0$, and

therefore it suffices to show that

$$f_i^{(1)} \leq f_i^{(2)}, \quad i = 1, 2, \dots$$

Since $\bar{A}(0) \geq L$, the first packet leaves the shaper instantaneously in both the systems, i.e., $f_1^{(1)} = s_1^{(1)}$ and $f_1^{(2)} = s_1^{(2)}$. Therefore, we have $f_1^{(1)} \leq f_1^{(2)}$.

Let l_i denote the length of the i th packet, and assume that

$$f_i^{(1)} \leq f_i^{(2)}, \quad i = 1, 2, \dots, m. \quad (34)$$

From (2) we can compute the departure time for the $(m+1)$ th packet in system S_1 as

$$\begin{aligned} f_{m+1}^{(1)} &= \min_{t \geq s_{m+1}^{(1)}} \left\{ t : \bar{A}(t - f_i^{(1)}) \geq \sum_{k=i}^{m+1} l_k, \quad i = 1, \dots, m \right\} \\ &\leq \min_{t \geq s_{m+1}^{(2)}} \left\{ t : \bar{A}(t - f_i^{(1)}) \geq \sum_{k=i}^{m+1} l_k, \quad i = 1, \dots, m \right\} \\ &\leq \min_{t \geq s_{m+1}^{(2)}} \left\{ t : \bar{A}(t - f_i^{(2)}) \geq \sum_{k=i}^{m+1} l_k, \quad i = 1, \dots, m \right\} \\ &= f_{m+1}^{(2)} \end{aligned} \quad (35)$$

where (35) is because $s_{m+1}^{(2)} \geq s_{m+1}^{(1)}$, and (36) follows from the nondecreasing nature of the shaper envelope, $\bar{A}(\tau)$, and the induction hypothesis (34). \square

Proof of Lemma 3: Let $S_n[t_1, t_2]$ denote the amount of connection n traffic served under GPS in the interval $[t_1, t_2]$. Assuming that the system starts empty, the backlog of connection n traffic at time t is defined as the difference $I_n[0, t] - S_n[0, t]$. Now, the busy period containing t is defined as the largest closed interval containing t , during which the backlog of connection n is positive. Note that only traffic arriving during a busy period can have positive delays and, therefore, we only need to consider the delays of such traffic (since by definition $D_n^* \geq 0$).

Assume that connection n traffic arrives at time t_b which is within a busy period of connection n that starts at time t_0 . Then, since packets from the same connection are served in order, the delay of the connection n traffic arriving at time t_b is given by

$$\begin{aligned} d(t_b) &= \min_{t' \geq t_b} \{ t' : S_n[0, t'] \geq I_n[0, t_b] \} - t_b \\ &= \min_{t' \geq t_b} \{ t' : S_n[0, t_0] + S_n[t_0, t'] \geq \\ &\quad I_n[0, t_0] + I_n[t_0, t_b] \} - t_b \\ &= \min_{t' \geq t_b} \{ t' : S_n[t_0, t'] \geq I_n[t_0, t_b] \} - t_0 - (t_b - t_0). \end{aligned}$$

The last equality follows from the fact that since t_0 is the start of a connection n busy period, $S_n[0, t_0] = I_n[0, t_0]$. Setting $t = t' - t_0$ and $\tau = t_b - t_0$, and observing that

$$\begin{aligned} &\min_{t' \geq t_b} \{ t' : S_n[t_0, t'] \geq I_n[t_0, t_b] \} \\ &= \min_{t \geq \tau} \{ t : S_n[t_0, t_0 + t] \geq I_n[t_0, t_b] \} + t_0 \end{aligned}$$

we get

$$d(t_b) = \min_{t \geq \tau} \{t : S_n[t_0, t + t_0] \geq I_n[t_0, t_b]\} - \tau.$$

Now, since the connection n traffic satisfies $I[t, t + \tau] \leq \delta_n + \rho_n \tau$, $t, \tau \geq 0$, from Lemma 10 in [11], we conclude that

$$S_n[t_0, t + t_0] \geq S_n(t).$$

Since by definition we also have

$$I_n[t_0, t_b] \leq \bar{I}_n(t_b - t_0)$$

it follows that

$$\{t : S_n[t_0, t + t_0] \geq I_n[t_0, t_b]\} \supseteq \{t : S_n(t) \geq \bar{I}_n(t_b - t_0)\}$$

and therefore

$$\begin{aligned} \min_{t \geq \tau} \{t : S_n[t_0, t + t_0] \geq I_n[t_0, t_b]\} \\ \leq \min_{t \geq \tau} \{t : S_n(t) \geq \bar{I}_n(t_b - t_0)\}. \end{aligned}$$

Recalling that $\tau = t_b - t_0$, we finally get

$$d(t_b) \leq \min_{t \geq \tau} \{t : S_n(t) \geq \bar{I}_n(\tau)\} - \tau. \quad \square$$

Proof of Lemma 4: We concentrate on the system operating according to π_W as defined in Section VI and repeat some notation for the sake of clarity. We denote by $t_i^{l,m}$, the timestamp with which the i th packet is enqueued in Q_γ^m ; recall that $t_i^{l,m}$ is the time that the i th packet would leave shaper A_n^m in conformance with the traffic envelope $\bar{A}_n^m(\tau)$. The time at which the packet leaves Q_γ^m (to be transmitted on the link or promoted to Q_e^m) is denoted as $t_i^{a,m}$ and we say that the packet *arrives at the scheduler* at time $t_i^{a,m}$. If the link is idle, the packet may be transmitted before it becomes eligible, i.e., $t_i^{a,m} \leq t_i^{l,m}$. The departure time of the i th packet from the scheduler is denoted as $t_i^{d,m}$. We need to show that for any packet i

$$t_i^{d,m} - t_i^{a,m} \leq D_n^m.$$

D_n^m has to be larger than the time it takes to transmit a complete packet, and so packets that are scheduled before they become eligible can never miss their deadline. All the eligible packets are scheduled in Q_e^m -busy periods, and so it suffices to show that packets that enter Q_e^m are never delayed by more than D_n^m .

Let $[t_0, t_f]$ be a Q_e^m -busy period. At time t_0 a packet either starts transmission or is in the process of being transmitted and recall that $t_s \leq t_0$ is the time that this packet starts transmission. If $t_s = t_0$ (in this case a packet from Q_e^m starts transmission at t_0), then the traffic of all connections arriving to the scheduler in $[t_s, t_f]$ is in conformance with the respective traffic shapers. In addition, in the interval $[t_s, t_f]$ the operation of the scheduler in π_W is identical to the one in π_{NW} if t_s were the start of the first busy period. Thus, for $t_s = t_0$ the result is true by the definition of D_n^m .

Now assume that $t_s < t_0$, i.e., an ineligible packet from connection j starts transmission at t_s . Observe that in a busy period $[t_0, t_f]$, the scheduler in π_W only schedules packets

that are in Q_e^m , except for the packet that is being transmitted at the start of the busy period. We will show next that the packets of all connections that have been transmitted in the interval $[t_s, t_f]$ have arrived to the scheduler in conformance with their respective traffic envelopes. The truth of the lemma will then follow as before.

Recall that $A_n^m[t_1, t_2]$ denotes the traffic from connection n that is promoted to Q_e^m in the interval $[t_1, t_2]$. Let $\hat{A}_n^m[t_1, t_2]$ denote the connection n traffic that arrives at the scheduler in the interval $[t_1, t_2]$. We need to show that

$$\hat{A}_n^m[t_1, t_2] \leq \bar{A}_n^m(t_2 - t_1), \quad t_s \leq t_1 \leq t_2 \leq t_f.$$

Since we are only concerned with node m here, we drop the superscript m for the sake of clarity. By the definition of t_0 , we have that for any connection n

$$A_n[t_1, t_2] \leq \bar{A}_n(t_2 - t_1), \quad t_0 \leq t_1 \leq t_2 \leq t_f.$$

For a connection $n \neq j$, we have in addition, $A_n[t_s, t_0] = 0$, and therefore

$$A_n[t_1, t_2] \leq \bar{A}_n(t_2 - t_1), \quad t_s \leq t_1 \leq t_2 \leq t_f, \quad n \neq j. \quad (37)$$

Note also that by definition $\hat{A}_n[t_1, t_2] = A_n[t_1, t_2]$, $t_s \leq t_1 \leq t_2 \leq t_f$, holds for $n \neq j$ since connection j is the only connection which transmits an ineligible packet in $[t_s, t_f]$.

Consider next connection j , and let p_j be the packet that starts transmission at t_s . Let τ_e denote the eligibility time of packet p_j . If $\tau_e \geq t_f$, then clearly $\hat{A}_j[t_1, t_2] \leq L \leq \bar{A}_j(0)$, $t_s \leq t_1 \leq t_2 \leq t_f$, since no more packets from connection j will be transmitted in $[t_s, t_f]$. Now suppose $t_s \leq \tau_e < t_f$. Then, all other packets of connection j will arrive after τ_e . For the case when $t_s \leq t_1 \leq \tau_e$ and $\tau_e \leq t_2 \leq t_f$, we have

$$\hat{A}_j[t_1, t_2] \leq \bar{A}_j(t_2 - \tau_e) \leq \bar{A}_j(t_2 - t_1).$$

The other cases can be similarly checked. \square

APPENDIX B

SUBADDITIVE TRAFFIC ENVELOPES

Proof of Lemma 1: In Proposition 3, we presented a method to obtain the minimal shaper envelope, for a given input traffic envelope and a shaper delay bound. There, we made the assumption that the input traffic envelope $\bar{I}_n(\tau)$ is a concave, increasing, piecewise linear function with finite number of slopes. Using such functions, we can approximate arbitrarily closely any concave increasing envelope. This means that by using the construction in Proposition 2 we can construct shapers that, for given concave increasing input envelope and shaper delay, are arbitrarily close to optimal. However, a general input envelope satisfies a weaker property than concavity, namely subadditivity [2]. In this appendix we outline how the method of Proposition 2 can be applied to subadditive input envelopes as well.

Consider a nondecreasing, piecewise linear function, $\bar{I}_n(\tau)$, with finite number of slopes (say K) and such that $\lim_{\tau \rightarrow \infty} \bar{I}_n(\tau)/\tau > 0$. Such functions can approximate arbitrarily closely any nondecreasing function in an appropriate

sense (using the Skorohod metric [15, Chapter VI]). Let $\hat{I}_n(\tau)$ be the minimal concave function such that $\bar{I}_n(\tau) \leq \hat{I}_n(\tau)$, $\tau \geq 0$. It can be seen that $\hat{I}_n(\tau)$ is increasing, piecewise linear with finite number of slopes and, therefore, can be written in the form of (29)

$$\hat{I}_n(\tau) = \min_{k=1, \dots, K} \{\hat{\delta}_{n,k} + \hat{\rho}_{n,k}\tau\}, \quad \tau \geq 0.$$

Given d such that $0 \leq d \leq \hat{\delta}_{n,K}/\hat{\rho}_{n,K}$, we construct the minimal envelope $\hat{A}_n(d)(\tau)$ corresponding to $\hat{I}_n(\tau)$ using Proposition 3. We claim that $\hat{A}_n(d)(\tau)$ is also the envelope of the minimal shaper $\mathcal{A}_n(d)$ that provides delay bound, d , to input traffic with envelope $\bar{I}_n(\tau)$.

To see this, consider another shaper \mathcal{A} such that $D(\bar{I}_n|\mathcal{A}) \leq d$, and assume that there exists some $\tau' \geq 0$, such that

$$\bar{A}(\tau') < \hat{A}_n(d)(\tau'). \quad (38)$$

Interpreting the inequality $D(\bar{I}_n|\mathcal{A}) \leq d$ using (10) (alternatively, see Fig. 2), we know that

$$\bar{I}_n(\tau) \leq \bar{A}(\tau + d), \quad \tau \geq 0.$$

By assumption, $\bar{A}(\tau)$ is a concave function; however, $\hat{I}_n(\tau)$ is the minimal concave function such that $\bar{I}_n(\tau) \leq \hat{I}_n(\tau)$. Therefore

$$\hat{I}_n(\tau) \leq \bar{A}(\tau + d), \quad \tau \geq 0.$$

Using (10) again, we conclude that $D(\hat{I}_n|\mathcal{A}) \leq d$. Now, using (8) we have

$$D(\hat{I}_n|\mathcal{A}_n(d) \wedge \mathcal{A}) = \max\{D(\hat{I}_n|\mathcal{A}_n(d)), D(\hat{I}_n|\mathcal{A})\} \leq d.$$

But the shaper $\mathcal{A}_n(d) \wedge \mathcal{A}$ has envelope

$$\bar{B}(\tau) = \min\{\hat{A}_n(d)(\tau), \bar{A}(\tau)\}, \quad \tau \geq 0. \quad (39)$$

From (39) and (38) we know that envelope $\bar{B}(\tau)$ is strictly smaller than $\hat{A}_n(d)(\tau)$, which contradicts the optimality of $\mathcal{A}_n(d)$.

ACKNOWLEDGMENT

The authors wish to thank R. Rajan for suggesting the proof to Lemma 1 as well as numerous insightful comments on early versions of this paper. They would also like to thank A. Parekh for many fruitful discussions regarding the comparative performance of GPS and EDF and S. Varma for prompting them to look into the issue of the efficiency of RCS disciplines as well as for many comments which helped improve the paper. They would also like to thank A. Makowski for many helpful discussions and comments on an earlier draft.

REFERENCES

- [1] *ATM UNI Specification, Version 3.1*, ATM Forum, Sept. 1994.
- [2] C.-S. Chang, "Stability, queue length and delay of deterministic and stochastic queueing networks," *IEEE Trans. Automat. Contr.*, vol. 39, no. 5, pp. 913-931, May 1994.
- [3] D. Clark, S. Shenker, and L. Zhang, "Supporting real-time applications in an integrated packet network: Architecture and mechanisms," in *Proc. ACM SIGCOMM'92*, Baltimore, MD, Aug. 1992.

- [4] R. L. Cruz, "A calculus for network delay. Part I: Network elements in isolation," *IEEE Trans. Inform. Theory*, vol. 37, pp. 114-131, Jan. 1991.
- [5] R. L. Cruz, "A calculus for network delay. Part II: Network analysis," *IEEE Trans. Inform. Theory*, vol. 37, pp. 132-141, Jan. 1991.
- [6] R. L. Cruz, "Service burstiness and dynamic burstiness measures: A framework," *J. High Speed Networks*, vol. 1, no. 2, pp. 105-127, 1992.
- [7] A. Demers, S. Keshav, and S. Shenker, "Analysis and simulation of a fair queueing algorithm," *J. Internetworking: Research Experience*, vol. 1, pp. 3-26, Jan. 1990.
- [8] L. Georgiadis, R. Guérin, and A. K. Parekh, "Optimal multiplexing on a single link: Delay and buffer requirements," IBM, T. J. Watson Research Center, Res. Rep. RC 19711 (97393), Aug. 1994. Short version appeared in *Proc. INFOCOM'94*.
- [9] P. Goyal and H. Vin, "Generalized guaranteed rate scheduling algorithms: A framework," Department of Computer Sciences, University of Texas at Austin, Tech. Rep. TR-95-30, 1995.
- [10] D. D. Kandlur, K. G. Shin, and D. Ferrari, "Real-time communication in multi-hop networks," in *Proc. INFOCOM'91*, 1991, pp. 300-307.
- [11] A. K. Parekh and R. G. Gallager, "A generalized processor sharing approach to flow control in integrated services networks: The single-node case," *IEEE/ACM Trans. Networking*, vol. 1, no. 3, pp. 344-357, June 1993.
- [12] ———, "A generalized processor sharing approach to flow control in integrated services networks: The multiple node case," *IEEE/ACM Trans. Networking*, vol. 2, no. 2, pp. 137-150, Apr. 1994.
- [13] A. K. J. Parekh, "A generalized processor sharing approach to flow control in integrated services networks," Ph.D. dissertation, Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139, Feb. 1992.
- [14] V. Peris, "Architecture for guaranteed delay service in high speed networks," Ph.D. dissertation, Department of Electrical Engineering, University of Maryland, College Park, in preparation.
- [15] D. Pollard, *Convergence of Stochastic Processes*. New York: Springer-Verlag, 1984.
- [16] S. Shenker, C. Partridge, and R. Guérin, "Specification of guaranteed quality of service," Internet Draft. draft-ietf-intserv-guaranteed-svc-05.txt, Integrated Services WG, IETF, June 1996.
- [17] H. Zhang, "Service disciplines for guaranteed performance service in packet-switching networks," *Proc. IEEE*, vol. 83, no. 10, pp. 1374-1396, Oct. 1995.
- [18] H. Zhang and D. Ferrari, "Rate-controlled service disciplines," *J. High Speed Networks*, vol. 3, no. 4, pp. 389-412, 1994.
- [19] Q. Zheng and K. Shin, "On the ability of establishing real-time channels in point-to-point packet-switching networks," *IEEE Trans. Commun.*, vol. 42, no. 3, pp. 1096-1105, Mar. 1994.



Leonidas Georgiadis (M'86-SM'95) received the Diploma degree in electrical engineering from Aristotle University, Thessaloniki, Greece, in 1979, and the M.S. and Ph.D. degrees from the University of Connecticut, Storrs, both in electrical engineering, in 1981 and 1986, respectively.

From 1981 to 1983, he was with the Greek army. From 1986 to 1987, he was Research Assistant Professor at the University of Virginia, Charlottesville. In 1987 he joined the IBM T. J. Watson Research Center, Yorktown Heights, NY, as a Research Staff Member. Since October 1995, he has been with the Telecommunications Department of Aristotle University, Thessaloniki, Greece. His interests are in the area of high speed networks, congestion control, mobile communications, scheduling, modeling, and performance analysis.

Dr. Georgiadis received the IBM Outstanding Innovation Award for his work on goal-oriented workload management for multiclass systems, in 1992.



Roch Guérin (S'84-M'86-SM'91) received the Diplôme d'Ingénieur from the École Nationale Supérieure des Télécommunications, Paris, France, in 1983, and the M.S. and Ph.D. degrees from the California Institute of Technology, Pasadena, both in electrical engineering, in 1984 and 1986, respectively.

Since 1986, he has been with IBM at the Thomas J. Watson Research Center, Yorktown Heights, NY, where he manages the Broadband Networking Department in the Advanced Networking Labora-

tory. His current research interests are in the areas of performance analysis and traffic modeling, scheduling policies, and in general quality-of-service issues in high-speed networks. He is also interested in issues related to inter-operability and scalability of internets, in particular in the context of end-to-end guarantees and QoS sensitive routing algorithms. He is an editor for the *IEEE/ACM TRANSACTIONS ON NETWORKING* and was an editor for the *IEEE TRANSACTIONS ON COMMUNICATIONS* and the *IEEE COMMUNICATIONS MAGAZINE*.

Dr. Guérin received an IBM Outstanding Innovation Award for his work on traffic management in the broadband services network architecture in 1994. He is a member of Sigma Xi.



Vinod Peris (S'93) received the B.Tech. degree in electrical engineering from the Indian Institute of Technology, Kanpur, India, in 1989, and the M.S. degree in electrical engineering from the University of Maryland, College Park, in 1992, where he is currently a Ph.D. candidate.

Since January 1995, he has been with IBM at the T. J. Watson Research Center, Yorktown Heights, NY, in the Broadband Networking Department. His research interests are in the area of performance evaluation, scheduling policies, and quality-

of-service issues in high-speed networks.

Mr. Peris co-authored the paper that received the Best Paper Award at ACM SIGMETRICS'94.



Kumar N. Sivarajan (S'88-M'91) received the B. Tech. degree in electrical engineering (electronics) from the Indian Institute of Technology, Madras, India, in 1987, and the M.S. and Ph.D. degrees in electrical engineering from the California Institute of Technology, Pasadena, in 1988 and 1990, respectively.

From July 1990 to October 1994, he was with the IBM T. J. Watson Research Center, Yorktown Heights, NY. Since October 1994, he has been with the Indian Institute of Science, Bangalore. His

research interests are in the areas of optical, cellular, and ATM networks where he has worked on network architecture, algorithms and performance analysis.

Dr. Sivarajan was the recipient of the IEEE Charles LeGeyt Fortescue Fellowship for the academic year 1987-88, and co-recipient of the IEEE Communications Society William R. Bennet 1995 Prize Paper Award.