

Guaranteeing Statistical QoS to Regulated Traffic: The Multiple Node Case ¹

Martin Reisslein
GMD FOKUS
Kaiserin–Augusta–Allee 31
10589 Berlin, Germany
reisslein@fokus.gmd.de

Keith W. Ross
Institute Eurecom
2229 Route des Cretes
06904 Sophia–Antipolis, France
ross@eurecom.fr

Srini Rajagopal
Manugistics
Philadelphia
srajagop@manu.com

Abstract

We design and evaluate traffic management schemes for multimedia traffic in multihop packet-switched networks. Our traffic management schemes give multimedia connections statistical Quality of Service (QoS) guarantees while simultaneously permitting the network to have a large connection-carrying capacity. We consider a QoS requirement that is particularly appropriate for multimedia connections which typically can tolerate some loss but have rigid delay constraints. Roughly speaking, we require that the fraction of traffic that exceeds a specific end-to-end delay limit be less than a prescribed bound. To account for a diversity of traffic types, our design allows for connection specific end-to-end delay limits and loss bounds. In order for the network to be able to provide these statistical QoS guarantees, we assume that each connections' traffic is regulated. In order to support many connections, the links statistically multiplex the connections' traffic. We propose a pragmatic scheme for multimedia traffic which consists of (i) leaky buckets for traffic regulation, (ii) smoothers at the network edges, and (iii) bufferless statistical multiplexing within the network. We also develop a call admission procedure for our traffic management scheme. Finally, we compare the connection-carrying capacity of our scheme with that of guaranteed service schemes (i.e., no loss) using GPS and RCS. Our numerical work indicates that our scheme can support significantly more connections.

1 Introduction

The development of traffic management schemes for multimedia traffic in packet-switched networks has received a great deal of attention over the last decade. The challenge lies in designing traffic management schemes that give multimedia connections stringent Quality of Service (QoS) guarantees while efficiently utilizing the networking resources. Much of the research effort has focused on deterministic service disci-

plines. These deterministic service disciplines assume that the traffic sent into the network by each connection is regulated, for instance, by leaky buckets. The deterministic service disciplines guarantee lossless service with specific, deterministic end-to-end delay bounds. Generalized Processor Sharing (GPS) [9, 10] and the Rate Controlled Service (RCS) discipline [17, 16, 6] are the most prominent deterministic service disciplines. Both have been studied very carefully in the context of general multihop networks. The drawback of these deterministic service disciplines is that they typically imply a small connection-carrying capacity for networks with bursty multimedia traffic.

We propose a traffic management scheme that offers Quality of Service (QoS) guarantees that are particularly appropriate for multimedia traffic *and* allows for a large connection-carrying capacity of the network. The QoS guarantees provided by our scheme can be roughly stated as follows: the fraction of traffic that exceeds a specific end-to-end delay limit is below a prescribed bound. We allow each connection to have its own limit on the end-to-end delay and its own bound on the fraction of traffic that exceeds this delay limit. These *statistical* QoS guarantees are particularly appropriate for multimedia traffic, whereby timestamping and a playout buffer can ensure the continuous playout of video or audio without jitter. Our traffic management scheme works roughly as follows: we (i) smooth each connection's traffic at the connection's input to the network as much as allowed by the connection's delay constraint; (ii) employ bufferless statistical multiplexing within the network; and (iii) base admission control on the worst-case assumption that sources are adversarial to the extent permitted by the connection's regulator, while concurrently assuming the connections generate traffic independently.

The statistical multiplexing within the network increases the connection carrying capacity of the network significantly at the expense of miniscule losses in the network. We provide numerical examples that demonstrate that by allowing for very small losses of the order of 10^{-7} (which can be effectively hidden by error

¹Supported partially by NSF grant NCR96-12781.

concealment techniques [8]) our scheme can typically support two to three times the number of connections that deterministic service disciplines can support.

In [12] we have laid the groundwork for our traffic management scheme which relies on smoothing and bufferless multiplexing. The work in [12] has focused on a single node. In this paper we extend the traffic management scheme of [12] to general multihop networks.

This paper is organized as follows. In Section 2 we formally define the leaky-bucket regulators and the statistical QoS requirement. We also discuss the smoothers at the network ingresses. In Section 3 we describe our network model and analyze the multihop network. In Section 4 we compare the performance of our smoothing/bufferless multiplexing scheme with that of deterministic service disciplines.

2 Regulated Traffic and the Statistical QoS Requirement

In this paper we study networks consisting of interconnected bufferless nodes. We assume a virtual circuit, connection-oriented network and view traffic as fluid, that is, packets are infinitesimal. The fluid model, which closely approximates a packetized model with small packets, permits us to focus on the central issues and significantly simplifies notation.

Each connection j entering the network has an associated regulator function, denoted by $\mathcal{E}_j(t)$, $t \geq 0$. The regulator function constrains the amount of traffic that connection j can send into the network over all time intervals. Specifically, if $A_j(t)$ is the amount of traffic that connection j sends into the network over the interval $[0, t]$, then $A_j(t)$ is required to satisfy

$$A_j(t + \tau) - A_j(\tau) \leq \mathcal{E}_j(t) \quad \forall \tau \geq 0, \quad t \geq 0.$$

In this paper we focus on a regulator which consists of a single leaky bucket; for the single leaky bucket, the regulator function takes the following form:

$$\mathcal{E}_j(t) = \sigma_j + \rho_j t.$$

Note that the single leaky bucket regulator constrains the long-run average rate of connection j to be no greater than ρ_j . Our traffic management scheme can be extended in a straightforward manner to more complex regulators consisting, for instance, of a cascade of leaky buckets. However, GPS which we shall use as a benchmark to evaluate our scheme, has been analyzed extensively in [9, 10] for single leaky bucket regulators. We will make use of some of those analytical results in our performance evaluation and focus therefore on single leaky bucket regulators throughout this paper.

Each connection also has a QoS requirement. We consider a QoS requirement that is particularly appropriate for multimedia traffic that has stringent end-to-end delay requirements but can tolerate some loss. Specifically, each connection has a connection-specific delay limit and a connection-specific loss bound. Let d_j and ϵ_j denote the delay limit and loss bound for connection j . Any traffic that overflows at one of the bufferless links in the network is considered to have infinite delay, and therefore violates the delay limit. The QoS requirement is as follows: the long-run fraction of connection- j traffic that is delayed by more than d_j seconds must be less than ϵ_j .

This QoS requirement can assure continuous, uninterrupted playback for a multimedia connection as follows. Each bit (or packet for packetized traffic) is time-stamped at the source. If a bit from connection j is time-stamped with value x , the bit (if not lost in the network) arrives at the receiver no later than $x + d_j$. The receiver delays playout of the bit until time $x + d_j$. Thus, by including a buffer at each receiver, the receiver can playback a multimedia stream without jitter with a fixed delay of d_j and with a bit loss probability of at most ϵ_j .

The first aspect of our strategy is to pass each connection's traffic through a buffered smoother at the connection's input to the network. We design the smoother for connection j so that the connection- j traffic is *never* delayed by more than d_j in the smoother. After having smoothed a connection's traffic, we pass the smoothed traffic to the network, and the traffic follows its route through the network. At each link along its route, the connection's traffic is statistically multiplexed with traffic from other connections. The second aspect of our strategy is to remove all of the buffers inside the network; that is, we use bufferless statistical multiplexing rather than buffered multiplexing before each link in the network. In our fluid model, a connection's traffic that arrives to a bufferless link either flows through the link without any delay or overflows at the link, and therefore has infinite delay. The QoS requirement of a connection j is met if the fraction of connection- j traffic that overflows any of the links along the route of connection j is less than ϵ_j . Also, note that provided the loss at each link is small, we can reasonably approximate a connection's traffic at the output of the multiplexer as being identical to its traffic at the input of the multiplexer. In other words, a connection that satisfies a certain regulator constraint at the input of a node satisfies the same regulator constraint at the output of the node.

For the smoother at the input of connection j to the network we use a buffer which serves traffic at rate c_j^* . When the smoother buffer is nonempty, traffic is drained from the buffer at rate c_j^* . When the smoother




Figure 1: Multihop network with N nodes.

buffer is empty and connection- j 's traffic is arriving at a rate less than c_j^* , traffic leaves the buffer exactly at the rate at which it enters the buffer.

Using the theory developed in [2], it can be shown that the maximum delay in the smoother is

$$\max_{t \geq 0} \left\{ \frac{\mathcal{E}_j(t)}{c_j^*} - t \right\}.$$

We set the smoother rate to

$$c_j^* = \min \left\{ c_j \geq 0 : \max_{t \geq 0} \left\{ \frac{\mathcal{E}_j(t)}{c_j} - t \right\} \leq d_j \right\}, \quad (1)$$

where d_j is the delay requirement for connection j . For $\mathcal{E}_j(t) = \sigma_j + \rho_j t$ and finite d_j we obtain

$$c_j^* = \max\left(\frac{\sigma_j}{d_j}, \rho_j\right).$$

Since the bufferless nodes inside the network introduce no additional delay, traffic from connection j that flows through the network without loss has an end-to-end delay of no more than d_j .

3 Guaranteeing Statistical QoS in a Multihop Network

3.1 Network Model

We consider a multihop network with N nodes; see Figure 1. We assume that node 1 is a bufferless multiplexer (see Figure 2) and nodes 2 through N are composed of a demultiplexer and a bufferless multiplexer (see Figure 3). We assume that the demultiplexer works instantaneously and the links connecting demultiplexers and multiplexers have infinite bandwidth. Let $I(n)$, $n = 1, \dots, N$, denote the set of connections multiplexed onto the output link of node n . Let $|I(n)|$ denote the cardinality of the set $I(n)$, that is, the number of connections multiplexed onto the output link of node n . Let C_n denote the capacity of the output link of node n . Often in this paper we will focus on connection 0, which we assume traverses nodes 1 through N . The path of connection 0 is illustrated in Figures 1, 2 and 3. We shall suppose for the study in this paper that the traffic streams multiplexed onto any of the output links are independent. In order to ensure the independence of the traffic streams we assume that




Figure 2: Node 1 is a bufferless multiplexer. The independent smoothed streams in $I(1)$ are multiplexed onto the output link of capacity C_1 .




Figure 3: Nodes 2 through N consist of a demultiplexer and a bufferless multiplexer. A new set of independent streams is multiplexed with connection 0 at each node.

$|I(1)|$ independent traffic streams are passed through their respective smoothers and are then multiplexed onto link C_1 (see Figure 2). Suppose that stream 0 is one of those streams, that is, $0 \in I(1)$. As illustrated in Figure 3, we assume that connection 0 is the only connection in $I(n-1)$ that is multiplexed onto link C_n . We assume that connection 0 is joined by $|I(n)| - 1$ fresh smoothed streams. We furthermore assume that the new streams are mutually independent, independent of stream 0 and independent of the streams that were multiplexed with connection 0 before. We summarize our assumptions about the traffic streams in the following two statements: (i) all streams in $\cup_{1 \leq n \leq N} I(n)$ are mutually independent; and (ii) $I(m) \cap I(n) = \{0\} \quad \forall m \neq n, 1 \leq m, n \leq N$.

3.2 Single Node Analysis

We now extend the single node analysis of [12] to multihop networks. For this purpose we initially focus on a particular node n , $1 \leq n \leq N$. As shown in Fig-

ure 3, $|I(n)|$ smoothed streams are multiplexed onto the output link of capacity C_n . Each of the connections j , $j \in I(n)$, has a regulator function $\mathcal{E}_j(t)$ and QoS parameters d_j and ϵ_j . Now regard the arrival process of stream j to its smoother as a stochastic process. Let $(A_j(t), t \geq 0)$ denote the arrival process of the unsmoothed stream j , and let $(A_j(t, \omega), t \geq 0)$ denote a realization of the stochastic process. Also let $\mathbf{A}_n(t) = (A_j(t), j \in I(n))$, and let $(\mathbf{A}_n(t), t \geq 0)$ be the associated vector stochastic arrival process. We say that the vector arrival process $(\mathbf{A}_n(t), t \geq 0)$ is *feasible* if (i) the component arrival processes $(A_j(t), t \geq 0)$, $j \in I(n)$, are independent, and (ii) for each $j \in I(n)$, each realization $(A_j(t, \omega), t \geq 0)$ satisfies the regulator constraint

$$A_j(t + \tau, \omega) - A_j(\tau, \omega) \leq \mathcal{E}_j(t) \quad \forall \tau \geq 0, \quad t \geq 0.$$

Denote \mathcal{A}_n for the set of all feasible vector arrival processes $(\mathbf{A}_n(t), t \geq 0)$. For a fixed feasible vector arrival process $(\mathbf{A}_n(t), t \geq 0)$, let $U_j(t)$ be the rate at which traffic from connection j leaves the associated smoother at time t , and let U_j be the corresponding steady-state random variable. Note that the streams U_j , $j \in I(n)$, may have traversed a number of bufferless nodes before reaching node n . The bufferless nodes do not delay or alter the traffic streams (except for miniscule losses due to link overflow which are negligible in typical networking scenarios). Consider multiplexing the streams U_j , $j \in I(n)$, onto the bufferless link of capacity C_n . The long-run average fraction of traffic lost by connection j is

$$P_{\text{loss}}^{\text{info},n}(j) = \frac{E \left[(\sum_{i \in I(n)} U_i - C_n)^+ \frac{U_j}{\sum_{i \in I(n)} U_i} \right]}{E[U_j]}.$$

The definition of $P_{\text{loss}}^{\text{info},n}(j)$ relies on the natural assumption that traffic loss at multiplexer n is split between the sources in a manner proportional to the rate at which the sources send traffic into the multiplexer. Noting that the term in the expectation of the numerator is non-zero only when $\sum_{i \in I(n)} U_i > C_n$, we obtain the following mathematically more tractable bound on $P_{\text{loss}}^{\text{info},n}(j)$:

$$P_{\text{loss}}^{\text{info},n}(j) \leq \frac{E \left[(\sum_{i \in I(n)} U_i - C_n)^+ U_j \right]}{C_n \cdot E[U_j]} := P_{\text{loss}}^n(j).$$

The bound $P_{\text{loss}}^n(j)$ preserves the essential characteristics of the performance measure $P_{\text{loss}}^{\text{info},n}(j)$ and is very tight. In the rare event when the aggregate demand for bandwidth $\sum_{i \in I(n)} U_i$ exceeds the link capacity C_n , $\sum_{i \in I(n)} U_i$ is typically very close to C_n . We provide numerical results that demonstrate the tightness of the bound in [12]. Henceforth, we focus on the bound $P_{\text{loss}}^n(j)$, and we refer to $P_{\text{loss}}^n(j)$ as the *loss probability for connection j at node n* .

By taking the supremum over all the feasible vector stochastic processes, we obtain the following worst-case loss probability for connection j at node n :

$$\phi_j^{*n} = \sup_{\mathcal{A}_n} \frac{E \left[(\sum_{i \in I(n)} U_i - C_n)^+ U_j \right]}{C_n \cdot E[U_j]}. \quad (2)$$

The loss probability of connection j at node n is guaranteed to be bounded by ϕ_j^{*n} for all feasible vector arrival processes in \mathcal{A}_n , that is, for all independent arrival processes whose sample paths satisfy the regulator constraints.

It has been shown in [12] that the supremum in (2) is attained by the random variables U_j^* with distribution

$$U_j^* = \begin{cases} c_j^* & \text{with probability } \frac{\rho_j}{c_j^*} \\ 0 & \text{with probability } 1 - \frac{\rho_j}{c_j^*}. \end{cases}$$

The proof in [12] depends critically on the independence of the random variables U_j . For this reason we restrict our attention in this paper to the networking scenario depicted in Figure 3. The routing strategy depicted in Figure 3 ensures that all multiplexed streams are independent. In [12] we also construct feasible arrival processes $A_j(t)$ that produce the Bernoulli random variables U_j^* at the smoother outputs.

It has furthermore been shown in [12] that single buffers serving traffic at rates c_j^* given by (1) are the optimal smoothers. (This proof also exploits the independence of the random variables U_j .) The single buffer smoothers minimize the worst-case loss probabilities ϕ_j^{*n} for all connections $j \in I(n)$ passing through node n . More complex smoothers consisting of a cascade of leaky buckets do not reduce the ϕ_j^{*n} 's and can therefore not increase the connection carrying capacity of node n . It has also been shown in [12] that in order to maximize the admission region of node n , the traffic streams should be smoothed as much as possible at the network ingresses. In other words, the smoother rates should be chosen as small as the delay limits permit.

Exploiting the fact that the U_j^* 's are Bernoulli random variables, we can simplify the expression for ϕ_j^{*n} :

$$\phi_j^{*n} = \frac{E \left[(\sum_{i \in I(n) - \{j\}} U_i^* + c_j^* - C_n)^+ \right]}{C_n}. \quad (3)$$

These bounds can be computed by convolving the distributions of the independent random variables. An approximate convolution algorithm is described in [7]. However, convolution often leads to numerical problems. We therefore apply the Large Deviation (LD) approximation, which is known to be accurate and also computationally very efficient [13, 4, 5, 11], to the expectation in the numerator. Towards this end, let

$\mu_{U_i^*}(s)$ denote the logarithm of the moment generating of U_i^* :

$$\mu_{U_i^*}(s) := \ln E[e^{sU_i^*}].$$

We define

$$U^* := \sum_{i \in I(n) - \{j\}} U_i^*.$$

Note that

$$\mu_{U^*}(s) = \sum_{i \in I(n) - \{j\}} \mu_{U_i^*}(s)$$

by the independence of the U_i^* 's. The large deviation (LD) approximation gives the following approximation for ϕ_j^{*n} [13]

$$\frac{1}{C_n s^{*2} \sqrt{2\pi \mu_{U^*}''(s^*)}} e^{-s^*(C_n - c_j^*) + \mu_{U^*}(s^*)},$$

where s^* is the unique solution to

$$\mu_{U^*}'(s^*) = C_n - c_j^*.$$

This concludes our analysis of a particular node n . In summary, (3) is a simple expression for the worst-case loss probability of connection j at node n ; this expression involves the independent Bernoulli random variables U_j^* , $j \in I(n)$, whose distributions we know explicitly. The LD approximation for (3) is highly accurate and is easily calculated.

3.3 Multihop Analysis

We now turn our attention to the entire multihop network depicted in Figures 1, 2 and 3. Without loss of generality we focus on connection 0 traversing nodes 1 through N . At the output of any of the nodes connection 0 has a peak rate no larger than c_0^* and an average rate no larger than ρ_0 . We can therefore use (3) to calculate the worst-case loss probability ϕ_0^{*n} at any of the bufferless multiplexers n , $n = 1, \dots, N$. The end-to-end loss probability of connection 0 is bounded by the sum of the worst-case loss probabilities of the individual hops along connection 0's path, that is, the loss in the network is bounded by $\sum_{n=1}^N \phi_0^{*n}$.

We note here that the single buffer serving traffic at rate c_j^* which was shown to minimize ϕ_j^{*n} at a single node n also minimizes the sum of the ϕ_j^{*n} . To see this, recall that the design of the smoother for connection j depends only on the connection parameters (the regulator function $\mathcal{E}_j(t)$ and the delay limit d_j). Therefore, the same smoother minimizes the ϕ_j^{*n} at every node n along connection j 's path. As a consequence the single buffer smoother with rate c_0^* minimizes $\sum_{n=1}^N \phi_0^{*n}$, the bound on the overall fraction of overflowing connection-0 traffic in the network.

The end-to-end QoS requirement of connection 0 is met if

$$\sum_{n=1}^N \phi_0^{*n} \leq \epsilon_0. \quad (4)$$

For admission control, we must ensure that (4) holds for all connections. Specifically, we must partition — either statically or dynamically — the loss constraint ϵ_j among the nodes traversed by each of the connections. This problem is of independent interest and is discussed in Sections 5.10 and 5.11 of [14].

In summary, our scheme consists of input smoothers at the network ingress and bufferless statistical multiplexing within the network. Increasing the number of nodes a connection traverses increases the loss probability but not the delay. Roughly speaking, the network loss probability for a connection is approximately the loss probability of a typical node multiplied by the number of nodes through which a connection passes. Because the loss probability of a node is dimensioned to be on the order of 10^{-6} or less, the increased loss is only of minor importance.

We note at this juncture that $\sum_{n=1}^N \phi_0^{*n}$ also provides a bound on the probability that a bit of connection 0 experiences an end-to-end delay of more than d_0 in the network. More formally, with D_0 denoting the end-to-end delay incurred by a bit of connection 0 in the network, we have

$$P(D_0 > d_0) \leq \sum_{n=1}^N \phi_0^{*n}. \quad (5)$$

Recall from Section 2 that by design a bit of connection 0 is delayed by at most d_0 in the smoother. Bits that do not overflow at any of the bufferless links in the network incur no additional delay while bits that do overflow are considered to have infinite delay. The bound (5) follows by noting that $\sum_{n=1}^N \phi_0^{*n}$ is a bound on the fraction of bits that do overflow. We emphasize that the bound on the probability that a bit violates a given delay limit is minimized by smoothing as much as the delay limit permits at the network ingress. We compare the performance of our smoothing/bufferless multiplexing scheme with that of deterministic traffic management schemes in the next section. These deterministic schemes are lossless and guarantee that a specific delay limit d_j is never violated, that is, they guarantee that $D_j \leq d_j$ with probability one.

In order to facilitate the comparison of the performance with the deterministic benchmarks we make the following simplifying assumptions about the traffic streams and the network. We assume that all streams in the network are homogeneous, that is, all streams have the same leaky bucket parameters and QoS requirement.

We set $\sigma_j = \sigma$, $\rho_j = \rho$, $d_j = d$ and $\epsilon_j = \epsilon$ for all streams j in the network. This implies that all connections have the same smoother rates, that is, $c_j^* = c^*$ for all streams j . Also, all of the Bernoulli random variables U_j^* are now identically distributed (but still independent). When comparing the performance we again focus on connection 0 traversing nodes 1 through N . We assume that each of the nodes n , $n = 1, \dots, N$ serves J streams, that is, $|I(n)| = J \forall n = 1, \dots, N$. We also assume that all output links in the network have the same capacity C . With these simplifying assumptions the worst-case loss probability of connection 0 at a node is

$$\phi_0^{*n} = \frac{E \left[\left(\sum_{i=1}^{J-1} U_i^* + c^* - C \right)^+ \right]}{C} := \phi^*. \quad (6)$$

The end-to-end loss probability of connection 0 is given by $N\phi^*$. Now assume that connection 0 is new and requests a connection traversing nodes 1 through N . The QoS requirement of the new connection 0 is satisfied if $\phi^* \leq \epsilon/N$. Suppose that all other streams that traverse one of the nodes n , $n = 1, \dots, N$, have allocated a loss constraint larger than ϵ/N to that node n . With this assumption the QoS requirements of all other streams will continue to hold if $\phi^* \leq \epsilon/N$. Hence connection 0 can be admitted if $\phi^* \leq \epsilon/N$.

We use the maximum number of connections each of the multiplexers 1 through N can carry without violating any QoS commitment as a measure of the performance of our scheme. Let J^* denote this maximum number of connections. We clearly have:

$$J^* = \max_{J \in \mathbb{N}} \{ J : \phi^* \leq \frac{\epsilon}{N} \},$$

where \mathbb{N} denotes the set of natural numbers. Note that in the described networking scenario each of the multiplexers 1 through N is serving connection 0 and $J^* - 1$ fresh connections.

4 Comparison with Deterministic Service Disciplines

In this section we compare the performance of our smoothing/bufferless multiplexing scheme with that of deterministic service disciplines. These deterministic service disciplines provide lossless service and guarantee a deterministic end-to-end delay bound. Of the deterministic service disciplines discussed in the literature, the Generalized Processor Sharing (GPS) [9, 10] and Rate-Controlled Service (RCS) [6] disciplines guarantee the smallest delay bounds. GPS considers the route of a connection as a whole and is thus able to guarantee tighter bounds than are achievable by adding worst-case delays at each hop [2, 3]. RCS, which is at the

heart of the Guaranteed Service framework of the Internet [15], relies on traffic shaping at every hop and can guarantee the same delay bounds as GPS. In fact it is shown in [6] that RCS has the potential of providing tighter delay bounds than GPS. However, the problem of how to choose the parameters of the RCS discipline in order to achieve these tighter delay bounds is not addressed. Instead, the authors suggest to use the parameters induced by the GPS discipline. This ensures that RCS can accept as many connections as GPS (and some more in a heterogeneous network). With the networking scenario that we have chosen for the performance comparison — homogeneous connections, homogeneous nodes, fluid model — GPS and RCS have exactly the same performance. We shall therefore compare our scheme's performance with that of GPS. For this purpose we modify the network depicted in Figures 1, 2 and 3. We remove the buffered smoothers at the network ingress and replace the bufferless multiplexers with buffered GPS servers.

4.1 Review of GPS

First, we briefly review GPS [9, 10] and adapt the notation of [9, 10] to our network model. The GPS server n serving the streams in $I(n)$ is characterized by positive real numbers w_j^n , $j \in I(n)$. These numbers govern the allocation of service to each of the streams. Let $S_j^n(\tau, t)$ denote the amount of stream j traffic served by server n during an interval $[\tau, t]$. The GPS policy guarantees that for any connection $j \in I(n)$ that is continuously backlogged in the interval $[\tau, t]$, that is, has a positive amount of traffic in server n 's buffer throughout the interval $[\tau, t]$,

$$\frac{S_j^n(\tau, t)}{S_i^n(\tau, t)} \geq \frac{w_j^n}{w_i^n}, \quad i \in I(n).$$

A connection j that is backlogged is thus guaranteed a minimum service rate called *connection j backlog clearing rate* of

$$g_j^n = \frac{w_j^n}{\sum_{i \in I(n)} w_i^n} C_n$$

by server n . The minimum connection-0 backlog clearing rate along its route traversing nodes 1 through N is

$$g_0 = \min_{1 \leq n \leq N} g_0^n.$$

Let $D_0(t)$ be the end-to-end delay incurred in the network by a connection-0 bit that arrives at time t . Furthermore, let D_0^* denote the maximum end-to-end delay of connection 0 over all time and all feasible arrival processes of all streams sharing a server with connection 0, formally:

$$D_0^* = \sup_{\cup_{1 \leq n \leq N} \mathcal{A}_n} \max_{t \geq 0} D_0(t).$$

A key result of [10] is the following deterministic bound on the maximum end-to-end delay for connection 0: if $g_0 \geq \rho_0$ then

$$D_0^* \leq \frac{\sigma_0}{g_0}.$$

We note that this bound does not require the independence of the served traffic streams. The independence of the traffic streams, however, is a prerequisite for our bound on the loss probability. Given a specific delay bound d_j , finding the corresponding weights of the general GPS policy is a very tedious procedure. This procedure is greatly simplified by setting $w_j = \rho_j$ for all traffic streams. GPS with this special assignment of weights is referred to as Rate Proportional Processor Sharing (RPPS). With RPPS the connection j backlog clearing rate at server n is given by

$$g_j^n = \frac{\rho_j^n}{\sum_{i \in I(n)} \rho_i^n} C_n.$$

For ease of comparison with our smoothing/bufferless multiplexing scheme we make the same simplifying assumptions we made at the end of Section 3. In particular, we set $\rho_j = \rho$, $\sigma_j = \sigma$ and $d_j = d$ for all connections. We assume that all servers 1 through N serve J connections and have a capacity of C . With these simplifications, the minimum back log clearing rate of connection 0 along its route from node 1 to N is $g_0 = C/J$. The end-to-end delay bound of connection 0 is

$$D_0^* \leq J\sigma/C, \quad (7)$$

provided the stability condition $C/J \geq \rho$ is satisfied. We are interested in the maximum number of connections each server along the route of connection 0 can serve without violating the delay limit of connection 0 or any other connection. Let J^* denote this maximum number of connections. From (7) and the the stability condition we have:

$$J^* = \lfloor \min\left\{\frac{Cd}{\sigma}, \frac{C}{\rho}\right\} \rfloor.$$

Note that J^* does not depend on N , the number of nodes connection 0 traverses. We remark that for the example at hand, consisting of homogeneous connections with homogeneous delay bounds, J^* is the absolute maximum number of connections a deterministic service discipline can support; no matter what deterministic service discipline (GPS, RCS, etc.) is employed.

4.2 Numerical Results

In this section we compare the performance of the smoothing/bufferless multiplexing scheme with that of GPS in multihop networks numerically. We have chosen the parameters $\sigma = 11,925$ Bytes, $\rho = 150$ Kbit/sec

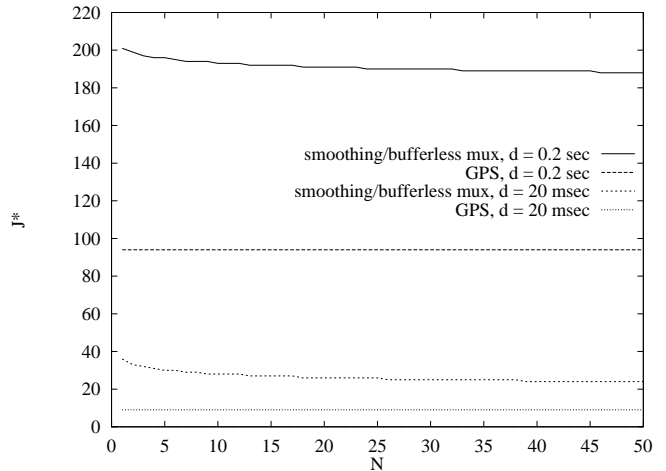


Figure 4: Maximum number of connections as a function of the number of hops for smoothing/bufferless multiplexing and GPS.




Figure 5: Maximum number of connections as a function of the delay bound for smoothing/bufferless multiplexing and GPS. The GPS performance is independent of the number of traversed hops.

and $C = 45$ Mbit/sec. For our smoothing/bufferless multiplexing scheme we set the loss bound to $\epsilon = 10^{-7}$. (These parameters are also used for some of the numerical examples in [5].) In Figure 4 we plot the maximum number of connections J^* that can be supported by the nodes 1 through N without violating any QoS requirements as a function of the number of hops, N . We do this for two delay bounds, $d = 20$ msec and $d = 0.2$ seconds. The maximum number of connections that can be supported by GPS is independent of N ; $J^* = 9$ for $d = 20$ msec and $J^* = 94$ for $d = 0.2$ seconds.

Figure 5 depicts J^* as a function of the delay bound d for $N = 5$ hops and $N = 50$ hops. Again, note that the GPS performance is independent of the number of hops.

Two points are especially noteworthy about the plots. First, with our smoothing/bufferless multiplexing scheme the number of allowable connections, J^* , drops off only slowly as the number of traversed hops, N , increases. Secondly, our smoothing/bufferless multiplexing scheme dramatically increases the connection-carrying capacity of the network. We observe from Figure 4, for instance, that for a delay bound of $d = 20$ msec and $N = 15$ hops our scheme can support more than three times the number of connections that GPS — or any other deterministic service discipline — can support. We achieve this remarkable performance by first smoothing the traffic at the network edges and then statistically multiplexing the smoothed traffic streams with miniscule loss probabilities within the network. The miniscule losses of the order of 10^{-7} can be effectively hidden by applying error concealment techniques to the multimedia streams [8]. The losses will therefore not be noticed by the viewers/listeners.

Our results are particularly relevant in light of the current debate on service disciplines for the Internet. Our results indicate that an Internet offering exclusively Guaranteed Service based on the RCS service discipline will be severely underutilized. An Internet service allowing for small losses — such as the Predictive Service framework proposed in [1] — would be able to make efficient use of the Internet resources and still provide the receivers with an enjoyable multimedia experience. Such a statistical Internet service could be based on our smoothing/bufferless multiplexing traffic management scheme.

Acknowledgements: Dorgham Sisalem's help in including the figures in this document is gratefully acknowledged.

References

- [1] D. Clark, S. Shenker, and L. Zhang. Supporting real-time applications in an integrated services packet network: Architecture and mechanism. In *ACM SIGCOMM '92*, pages 14–26, 1992.
- [2] R. Cruz. A calculus for network delay, part I: Network elements in isolation. *IEEE Transactions on Information Theory*, 37(1):114–121, January 1991.
- [3] R. Cruz. A calculus for network delay, part II: Network analysis. *IEEE Transactions on Information Theory*, 37(1):122–141, January 1991.
- [4] A. Elwalid, D. Heyman, T. Lakshman, D. Mitra, and A. Weiss. Fundamental bounds and approximations for ATM multiplexers with application to video teleconferencing. *IEEE Journal on Selected Areas in Communications*, 13(6):1004–1016, August 1995.
- [5] A. Elwalid, D. Mitra, and R. H. Wentworth. A new approach for allocating buffers and bandwidth to heterogeneous regulated traffic in an ATM node. *IEEE Journal on Selected Areas in Communications*, 13(6):1115–1127, August 1995.
- [6] L. Georgiadis, R. Guerin, V. Peris, and K.N. Sivarajan. Efficient network QoS provisioning based on per node traffic shaping. *IEEE/ACM Transactions on Networking*, 4(4):482–501, August 1996.
- [7] T. Lee, K. Lai, and S. Duann. Design of a real-time call admission controller for ATM. *IEEE/ACM Transactions on Networking*, 4(5):758–765, October 1995.
- [8] W. Luo and M. El Zarki. Analysis of error concealment schemes for MPEG-2 video transmission over ATM based networks. In *Proceedings of SPIE Visual Communications and Image Processing 1995*, Taiwan, May 1995.
- [9] A. Parekh and R. Gallager. A generalized processor sharing approach to flow control in integrated services networks: The single node case. *IEEE/ACM Transactions on Networking*, 1(3):344–357, June 1993.
- [10] A. Parekh and R. Gallager. A generalized processor sharing approach to flow control in integrated services networks: The multiple node case. *IEEE/ACM Transactions on Networking*, 2(2):137–150, April 1994.
- [11] M. Reisslein and K. W. Ross. Call admission for prerecorded sources with packet loss. *IEEE Journal on Selected Areas in Communications*, 15(6):1167–1180, August 1997.
- [12] M. Reisslein, K. W. Ross, and S. Rajagopal. Guaranteeing statistical QoS to regulated traffic: The single node case. In *submitted to Infocom '99*. available at <http://www.eurecom.fr/~ross/>.
- [13] J. Roberts, U. Mocchi, and J. Virtamo (Eds.). *Broadband Network Traffic: Performance Evaluation and Design of Broadband Multiservice Networks, Final Report of Action COST 242, (Lecture Notes in Computer Science, Vol. 1155)*. Springer Verlag, 1996.
- [14] K. W. Ross. *Multiservice Loss Models for Broadband Telecommunication Networks*. Springer-Verlag, 1995.
- [15] S. Shenker and C. Partridge. Specification of guaranteed quality of service. internet draft draft-ietf-intserv-guaranteed-svc-03.txt. November 1995.
- [16] H. Zhang. Providing end-to-end performance guarantees using non-workconserving disciplines. *Computer Communications: Special Issue on System Support for Multimedia Computing*, 18(10), October 1995.
- [17] H. Zhang and D. Ferrari. Rate-controlled service disciplines. *Journal of High Speed Networks*, 3(4):389–412, 1994.