In Proceedings of Performance '99, Istanbul, Turkey, August 1999.

Buffering vs. smoothing for end-to-end QoS: Fundamental issues and comparison

Tao Wu and Edward W. Knightly

Department of Electrical and Computer Engineering
Rice University
6100 South Main, Houston, Texas 77005
{*tao,knightly*}*@ece.rice.edu*

Abstract

Smoothing traffic flows at the network edge to reduce their burstiness has been shown to have significant benefits for video-on-demand systems and deterministic services. In this paper, we investigate the relative abilities of smoothing and buffering to improve a network's admissible region for end-to-end delay-bounded statistical services. In single multiplexer systems, we show that buffering outperforms smoothing for any delay bound and loss probability. We find that this behavior is due not only to statistical buffer sharing, but also to heterogeneity of the traffic flows' time scales. In multi-node scenarios, key issues for buffering and smoothing are user QoS requirements, traffic characteristics, and route length. For example, we find that as the number of hops traversed increases, the advantages of buffering diminish due to node-to-node buffer partitioning; and while smoothing is asymptotically superior, we find that in practice, the "critical route length" required to realize a smoothing gain is so large that buffering results in larger admissible regions, even in many multi-node scenarios.

Keywords: traffic smoothing, traffic shaping, end-to-end QoS, statistical service

1 Introduction

In guaranteed quality-of-service communication, one expects that the greater a source's burstiness in terms of peak-to-average rate ratio, temporal correlation, etc., the greater its network resource requirements. This intuitive observation motivates *traffic smoothing*, in which a flow's burstiness is reduced at the network edge to achieve a variety of goals, including the reduction of network resource demands.

Indeed, in the literature, smoothing has been shown to be beneficial in a number of scenarios. For example, for deterministic QoS guarantees, smoothing can have significant advantages in multi-hop rate-controlled networks [8,12]. Moreover, smoothing can also have significant benefits in video-on-demand systems in which traffic patterns are known in advance, clients may "work ahead" and prefetch video frames, and delay requirements are not strict [5,14,16,19,20]. Finally, networks in which traffic flows are smoothed but not buffered are more tractable than buffered networks, and a number of studies have considered such scenarios [4,9,15]. However, despite such potential advantages of smoothing, it is not yet clear what the services and scenarios are in which smoothing can improve a network's admissible region.

In this paper, we investigate the effectiveness of traffic smoothing in end-to-end delay-bounded statistical and deterministic services. Towards this end, we conduct a comparative study of two systems: (1) a *smoothing system* in which traffic flows are smoothed at the network edge with a maximum delay D and then are serviced by a network of bufferless multiplexers, and (2) a *buffering system* in which the same traffic flows are *not* smoothed,

This research is supported by Nokia Corporation, NSF CAREER Award ANI-9733610, and NSF Grant ANI-9730104. The authors can be reached via http://www.ece.rice.edu/networks.

and are serviced by a network of buffered multiplexers, which also provide a maximum end-to-end delay D. For a given end-to-end delay requirement D and loss probability P_l (which may be 0 for the special case of deterministic service), we compare the admissible regions of the two systems and identify the key factors that influence their relative performance.

A fundamental issue for smoothing and buffering is the relative extent to which network resources are partitioned vs. shared. First, in the smoothing system, each flow's smoothing buffers are individually partitioned, while in the buffering system, flows share a common buffer. Thus, an obvious property of the buffering system is that it can attain a statistical multiplexing gain from statistically sharing a common resource. However, we will show that there is a further advantage of the buffering system that is not immediately evident, namely, a gain due to sources with different *critical time scales* sharing a resource. A key observation is that this gain derives solely from time scale heterogeneity and is therefore available in both deterministic and statistical buffering systems, i.e., it is not an artifact of statistical sharing. We illustrate this aspect of buffer sharing using deterministic delay calculus [2] and show how heterogeneous sources can obtain a buffering gain under deterministic services in which statistical multiplexing gains are not available.

Consequently, due to the aforementioned advantages of resource sharing, we show using sample path analysis that in the case of a single multiplexer, the buffering system has a greater admissible region than the smoothing system for any end-to-end delay bound and loss probability, including a loss probability of zero for deterministic service. We experimentally quantify the buffering system's advantages using simulations and admission control experiments with both periodic on-off sources and long traces of compressed video. As an illustrative example with on-off sources and a delay of 70 msec, we find that the buffering system achieves an admissible region 30% larger than the smoothing system's, while the advantage is 10% for video sources.

We next turn to multiple node networks. Here again at issue is partitioning vs. sharing of network resources: in multi-node scenarios, a flow's end-to-end delay budget and hence buffering is *partitioned* among network nodes. Consequently, the significant advantages of buffering found in the single-node case are lessened in multi-node scenarios. We formally establish this property by showing that under certain conditions, there exists a *critical route length* H^* such that if the number of hops traversed is less than H^*, the buffering system's admissible region is larger than that of the smoothing system, whereas at H^* hops and beyond, either the admissible regions are equivalent or smoothing is superior. We experimentally investigate this result and find that the flows' traffic characteristics strongly influence H^*: for periodic on-off sources, H^* tends to be moderate, on the order of 6 to 9 hops in typical examples of Section 4; in contrast, for more bursty video sources, we find H^* to be so large that smoothing is unable to improve the admissible region for end-to-end statistical services.

Finally, in addition to traffic characteristics, we find that user QoS requirements also play a key role in the relative merits of smoothing and buffering. We formalize this by using envelope-based admission control tests [11] to show that the critical route length H^* is a non-decreasing function of the loss probability P_l. Indeed, with the most stringent QoS requirement of $P_l = 0$ for deterministic service, H^* can be one. This concurs with previous studies of *deterministic* services which demonstrated that smoothing traffic at the network edge can produce significant utilization gains in many multi-hop scenarios [8,12].

Thus, we study the relative merits of the two systems from the perspective of partitioning vs. sharing of network resources. By employing a number of analytical techniques, including deterministic delay calculus, sample path analysis, and statistical traffic envelopes, we show that statistical resource sharing, heterogeneity of time scales, and node-to-node buffer partitioning play key roles in these systems' admissible regions. Moreover, we explore the impact of several important system parameters such as route length, QoS requirements, and traffic characteristics on the smoothing/buffering systems. We find that in stark contrast to deterministic services and video-on-demand systems, smoothing for delay-bounded statistical services is of limited utility, and in many cases is detrimental towards improving a network's admissible region.

The remainder of this paper is organized as follows. In Section 2 we describe the smoothing and buffering systems. Next, in Section 3 we consider partitioning and sharing issues in the case of a single network multiplexer while in Section 4 we consider multi-node networks. In both cases, we perform experimental investigations. Finally, in Section 5, we conclude.

2 System Description

In this paper, we compare the relative merits of smoothing and buffering for end-to-end QoS by studying the performance of two related systems: a *smoothing system* $\mathcal{S}$, and a *buffering system* $\mathcal{B}$. Denoting a traffic flow's maximum allowable end-to-end delay as D, the smoothing system allocates the delay budget to traffic shapers at the network edge, while the buffering system allocates all of the delay budget to buffers inside the network.

2.1 The Smoothing System

In the smoothing system $\mathcal{S}$, each traffic flow is smoothed or shaped at the network edge and is serviced by a network of bufferless multiplexers. The delay incurred in the smoothing element can be bounded as follows. Denoting the arrivals of traffic flow j in the interval $[s, s+t]$ as $A_j[s, s+t]$, a non-decreasing subadditive function $b_j(t)$ is said to be a deterministic envelope of flow j [2] if $A_j[s, s+t] \leq b_j(t) \;\; \forall s, t > 0$.

The smoothing element can be characterized by an envelope $\hat{b}_j(t)$ such that by delaying packets as required, traffic flow j's arrivals are bounded by $\hat{b}_j(t)$ at the output of the smoother. The delay incurred by smoothing a traffic flow with envelope $b(t)$ to one with envelope $\hat{b}(t)$ is bounded by $D = \max_{s \geq 0}\{(\hat{b}^{-1}(b(s)) - s)^+\}$ which can be interpreted as the maximum horizontal distance between the two envelopes b and $\hat{b}$ [3,8,12]. In this paper (and in [12,15]), a traffic flow is smoothed with a buffered first-come first-serve server with rate

$$c_j = \max_{t \geq 0} \frac{b_j(t)}{t + D} \tag{1}$$

which is the minimum smoothing rate such that the smoothing delay is no larger than D.

Observe that with bufferless multiplexers inside the network, the maximum end-to-end delay is also bounded by D. Moreover, without network buffers, loss occurs in a multiplexer whenever the aggregate input rate exceeds the multiplexer's link capacity. Throughout this paper, we will study the loss probability and end-to-end delay behavior of this system.

2.2 The Buffering System

In the buffering system $\mathcal{B}$, traffic is transmitted into the network without incurring any delays due to smoothing (or one can view that the traffic smoother has an envelope with $\hat{b}_j(t) = b_j(t)$ for all t, and hence the traffic is not delayed by the smoother). In this case, the user's end-to-end delay budget D is allocated to queueing delays inside the network's buffers.

In this system, backlogged traffic is serviced in first-come-first-serve order, and each network node employs delay-jitter control [6]. A delay-jitter controller at the h^{th} hop holds packet k of connection j for $D_j^{h-1} - \delta_{j,k}^{h-1}$ seconds before queueing it, where D_j^{h-1} is connection j's delay bound at node $h-1$ and $\delta_{j,k}^{h-1}$ is the actual delay incurred by packet k of connection j at node $h-1$. Consequently, if traffic flow j's arrivals in $[s, t]$ are $A_j[s, t]$ at the entrance of the network, they are $A_j[s - \sum_{h=1}^{H} D_j^h, t - \sum_{h=1}^{H} D_j^h]$ at the entrance of the H^{th} queue. Because the arrival sequence at the H^{th} queue is a constant-delayed version of the original sequence, we can analyze networks of buffered multiplexers using the same properties of A at each network node.

While consideration of buffered networks without delay-jitter control is beyond the scope of this paper, our techniques can be extended to *rate*-controlled servers [18] using techniques such as in [11], or to more general classes of networks using other techniques for end-to-end performance evaluation, e.g., [1,13].

2.3 Experimental Workload

Throughout this paper we use two sources for admission control and simulation experiments: a periodic on-off source and a 30 minute trace of an MPEG-compressed video of an action movie. The periodic on-off source can be characterized by three parameters, i.e., the on period T_{on}, the off period T_{off}, and the peak rate R. The parameters that we use are $T_{\text{on}} = 83$ msec, $T_{\text{off}} = 750$ msec and $R = 5.87$ Mbps. The MPEG video trace exhibits rate variations over multiple time scales and has an average rate of 583 Kbps and a peak rate of 5.87 Mbps. Finally, we consider networks of FCFS servers with 45 Mbps link capacity in all simulations and admission control tests.

3 Smoothing vs. Buffering: The Single Node Case

Here, we show analytically and demonstrate experimentally that for the single node case, the buffering system attains a higher (or same) admissible region than the smoothing system for any end-to-end delay bound D and loss probability P_l. Our analysis is based on sample path behavior and addresses both deterministic and statistical services within the same framework. We find that heterogeneity of the sources' time scales and statistical multiplexing gains account for buffering's relative advantage to smoothing. We quantify these results using simulation and admission control experiments.

3.1 Loss in Delay-Bounded System

We next use sample path analysis to show that for any arrival sequence the loss in the buffering system is less than that in the smoothing system. By demonstrating this for any sample path, the result is quite general and applies to both deterministic and statistical services.

To show this, we first note that the busy period of a finite buffer FCFS server is smaller than that of an infinite buffer FCFS server when loss occurs, and the duration of this busy period is dependent on the buffer size B. We refer to such a busy period in a finite buffer FCFS server as a *finite buffer busy period* and denote it by F. We are interested only in the buffer dynamics within finite buffer busy periods since this is where loss occurs. Without loss of generality, we assume a finite buffer busy period of interest, F, starts at time 0. The aggregate arrival from the beginning of F up to time t is denoted by $A(t)$, and the link capacity of the server is C.

Lemma 1 *In a single node buffering system $\mathcal{B}$, the loss in any finite buffer busy period F, $L_{\mathcal{B}}(F)$, is* $L_{\mathcal{B}}(F) = \max(\sup_{t \in F}(A(t) - Ct) - B, 0)$.□

The proof can be found in [17]. Roughly, this lemma states that if there is a loss in a finite buffer FCFS server, the size of the loss is determined by maximizing $A(t) - Ct - B$ for t in the corresponding busy period. Figure 1(a) illustrates Lemma 1 by depicting a sample path of a finite buffer and an infinite buffer queue. Note that while there is loss of size L_1 at time t_1, $L_{\mathcal{B}}(F)$ is achieved at $t^* > t_1$ where L_1 is also accounted for. This result provides an analytical tool to obtain Theorem 1 as follows.

Theorem 1 *In a single node system in which all flows have delay bound D, the admissible region of the buffering system $\mathcal{B}$ is larger than or identical to that of the smoothing system $\mathcal{S}$, for both deterministic and statistical services.*

Proof: We prove the theorem on a sample path basis. Since the link capacity is C, the buffer size of $\mathcal{B}$ is $B = CD$. For each finite buffer busy period F where loss occurs, let t^* be the maximizing t in Lemma 1. The ending time of F, which we denote by t_e, satisfies $t_e \geq t^* + D$ since the buffer is full at t^*.

Note that $A(t^*)$ in $\mathcal{B}$ is smoothed to $A_s(t_s)$ in $\mathcal{S}$, i.e., $A_s(t_s) = A(t^*)$. We also have $t^* \leq t_s \leq t^* + D$ since A_s is deterministically smoothed and thus $t_s \leq t_e$.

Denote the loss during $[0, t_s]$ in $\mathcal{S}$ by $L_{\mathcal{S}}(t_s)$. The service $\mathcal{S}$ provides during $[0, t_s]$ is upper bounded by Ct_s. Then we have $L_{\mathcal{S}}(t_s) \geq A_s(t_s) - C \cdot t_s$ since $\mathcal{S}$ has a bufferless multiplexer. Furthermore, $L_{\mathcal{S}}(F) = L_{\mathcal{S}}(t_e) \geq L_{\mathcal{S}}(t_s) \geq A_s(t_s) - C \cdot t_s = A(t^*) - C \cdot t_s$.

On the other hand, according to Lemma 1, the loss in $\mathcal{B}$ in F is $L_{\mathcal{B}}(F) = A(t^*) - C \cdot t^* - B$. We thus have $L_{\mathcal{S}}(F) - L_{\mathcal{B}}(F) \geq B + C \cdot t^* - C \cdot t_s \geq B - C \cdot D = 0$.

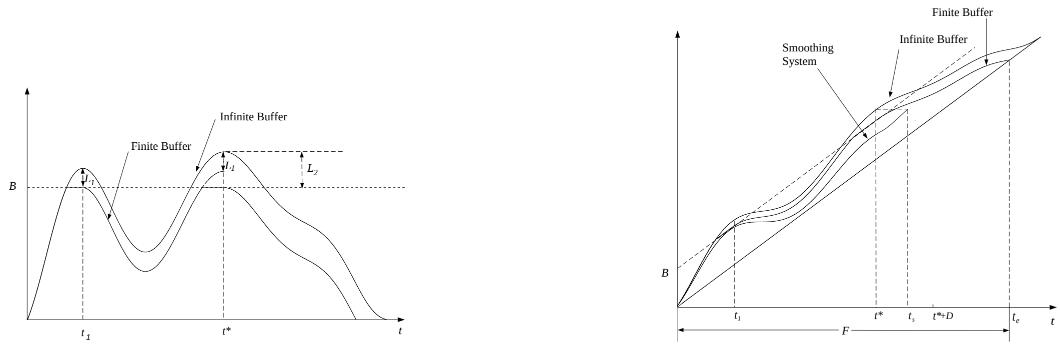

(a) Buffer dynamics (b) Service vs. arrival within a finite buffer busy period F

Fig. 1. Illustration of Lemma 1 and Theorem 1

We have shown that for each finite buffer busy period where loss occurs in $\mathcal{B}$, there is a loss of equal or larger size in $\mathcal{S}$ in the same period. Since this is true for each sample path, $\mathcal{B}$ is capable of admitting more (or the same number of) flows than $\mathcal{S}$ for both deterministic and statistical services. □

3.2 Heterogeneity-of-Time-Scales Gain of Buffering

An apparent explanation for Theorem 1 is that the statistical multiplexing gain in the buffering system $\mathcal{B}$ outweighs any advantages of smoothing. However, this explanation fails for deterministic services, in which resources are allocated according to the *worst case* scenario, and statistical sharing cannot be exploited since no loss can occur for deterministic service. Here, we show using deterministic delay calculus [2] that heterogeneity of the traffic flows' time scales partially accounts for the superiority of the buffering system.

Consider a single node buffering system $\mathcal{B}$ and smoothing system $\mathcal{S}$ and a deterministic service with delay bound D. Suppose there are N flows, each with traffic envelope $b_j(t)$, $j = 1, 2, \ldots, N$. From Equation (1), the required link capacity of the bufferless multiplexer in $\mathcal{S}$, $C_{\mathcal{S}}$, is

$$C_{\mathcal{S}} = \sum_{j=1}^{N} c_j = \sum_{j=1}^{N} \left\{ \max_{t \geq 0} \frac{b_j(t)}{t+D} \right\} \tag{2}$$

We denote the maximizing t's in Equation (2) by t_j^*, $j = 1, 2, \ldots, N$, and refer to t_j^* as source j's "critical time scale".

In the buffering system, the minimum link capacity needed to support the same set of sources is $C_{\mathcal{B}} = \max_{t \geq 0} \left\{ \frac{\sum_{j=1}^{N} b_j(t)}{t+D} \right\}$. Observe that $C_{\mathcal{B}} \leq C_{\mathcal{S}}$ since $\max_{t \geq 0} \sum_{j=1}^{N} \frac{b_j(t)}{t+D} \leq \sum_{j=1}^{N} \max_{t \geq 0} \frac{b_j(t)}{t+D}$. Thus, to support the same set of sources, smoothing requires higher bandwidth. Note that *equality* holds only when the critical time scales t_j^* of all sources are the same (homogeneous traffic satisfies this condition).

We now provide a simple example to illustrate the heterogeneity gain of buffering. Suppose there are two dual leaky bucket flows with delay requirement $D = 1$ and traffic envelopes $b_1(t) = \min(5t, 4+t)$ and $b_2(t) = \min(3t, 4+t)$. From Equation (2), $c_1 = 2.5$ and $c_2 = 2$. On the other hand, if we multiplex the two sources, the envelope for the aggregate traffic is $b(t) = \min(8t, 4+4t, 8+2t)$, and the capacity required for the buffering system for the same delay bound is $C_{\mathcal{B}} = 4$, while $C_{\mathcal{S}} = c_1 + c_2 = 4.5$.

3.3 Experiments

Here, we perform a set of simulations and admission control experiments to quantify the extent to which the buffering system outperforms the smoothing system in the single node case.

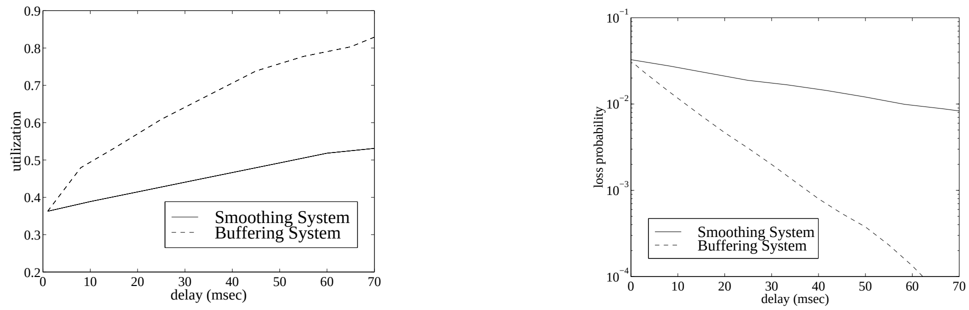

(a) Admissible Region ($P_l = 10^{-3}$) (b) Loss Rate (67.4% utilization)

Fig. 2. Simulation Results for Periodic On-off Sources

Figure 2 shows the results for periodic on-off sources. For Figure 2(a), we use simulations to experimentally determine these systems' admissible regions by finding the maximum number of traffic flows that can be supported for a given QoS requirement. The figure depicts this number of flows (scaled to average utilization) vs. delay bound for a loss probability of 10^{-3}. As expected from Theorem 1, the figure shows that the buffering system achieves a larger admissible region than the smoothing system, with the two curves converging at low delays, since with $D \approx 0$, both systems behave as a single bufferless multiplexer. We also note that for larger delay bounds, the buffering system's utilization is significantly higher than the smoothing system's; for example, the difference is approximately 30% when the delay bound is 70 msec.

For Figure 2(b), we fix the number of flows in both systems such that the utilization is 67.4% and depict loss probability vs. delay bound. Here, the loss probability of the buffering system is one to two orders of magnitude below that of the smoothing system for delay bounds above 40 msec.

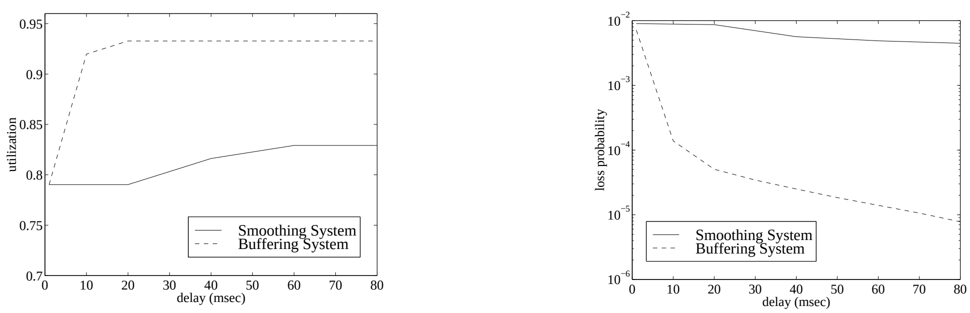

(a) Admissible Region ($P_l = 10^{-3}$) (b) Loss Rate (89.5% utilization)

Fig. 3. Simulation Results for Video Traces

Figure 3 depicts the results of analogous experiments using traces of MPEG-compressed video. For Figure 3(a), observe that the buffering system again has a larger admissible region, although the shapes of these curves differ from those of Figure 2(a). In particular, here the admissible region for the buffering system increases sharply for delays of up to 10 msec and then flattens considerably. This behavior indicates that while short time scale frame-to-frame rate variations are easily absorbed by network buffers, buffering is ineffective at absorbing longer time scale scene-to-scene rate variations and hence the admissible region flattens [10]. Comparing the buffering and smoothing curves, the difference between the two admissible regions is approximately 10% utilization.

In Figure 3(b), we fix the utilization to 89.5% and depict the experimental loss probability. As was the case for on-off sources, the buffering system's loss probability is significantly lower than the smoothing system's.

Thus, the above experiments quantify the advantages of the buffering system outlined in Sections 3.1 and 3.2 for single multiplexer networks, and indicate that in practice, the buffering system can achieve utilizations of 10% to 30% greater than the smoothing system depending on the characteristics of the traffic flows.

4 Buffering, Smoothing, and Multi-Hop Networks

In single node networks, we showed that buffering systems always achieve larger admissible regions than smoothing systems. However, this is not necessarily the case in multi-node scenarios. Here, we show that due to node-to-node partitioning of a flow's delay budget, the advantages of buffering over smoothing are reduced as an increased number of hops are traversed. In particular, we demonstrate the existence of a *critical route length* H^* such that for networks of at least H^* hops, *smoothing* achieves a larger admissible region than buffering. We then investigate the impact of the quality of service parameters and traffic characteristics on the critical route length.

4.1 Critical Route Length H^*

The proposition below establishes whether a larger admissible region is achieved by allocating an end-to-end delay budget entirely to traffic smoothing at the network edge or by equally partitioning the delay budget to queueing delays in the network's multiplexers.

Proposition 1 *For identical traffic flows with delay bound D traversing H multiplexers with capacity C, there exists a* critical route length H^* *such that for $H < H^*$, $\mathcal{B}$ has a larger admissible region than $\mathcal{S}$; for all $H \geq H^*$ the admissible region of $\mathcal{B}$ is smaller than (or the same as) that of $\mathcal{S}$.*

Proof: Since the traffic is reshaped by a delay jitter controller at each node, loss along the path that a flow traverses occurs independently. The loss and the end-to-end loss and delay-bound violation probability P_l is given by $P_l = 1 - (1 - P_{l,n})^H$, where $P_{l,n}$ is the loss probability at a single node [7]. Expanding the expression and neglecting all higher order terms yields $P_l \approx H \cdot P_{l,n}$. Hence if P_l and D are fixed, the per-node admissible regions for both $\mathcal{B}$ and $\mathcal{S}$ will decrease with increasing H. Furthermore, the buffer size at each node in $\mathcal{B}$ will also decrease since $B = \frac{D}{H}C$. Thus the per-node admissible region of $\mathcal{B}$ would asymptotically be that of the smoothing system if the sources of both systems were the same. However, $\mathcal{S}$ admits smoothed streams, and thus asymptotically outperforms $\mathcal{B}$. On the other hand, according to Theorem 1, $\mathcal{B}$ is superior to $\mathcal{S}$ in single node, thus there exists an H^* where the two systems' admissible regions cross. $\Box$

With the existence of H^* established, the key issues for smoothing and buffering in multi-node networks are (1) what is the expected range of H^* in practice, and (2) how do user QoS requirements and traffic characteristics impact H^*? We address these issues below.

4.2 H^* and Loss Probability

Here, we show that the critical route length is a non-decreasing function of loss probability, so that as user QoS constraints become more restrictive, the smoothing system obtains a relative advantage. In the extreme case of $P_l = 0$ for deterministic service, H^* can be one such that for two or more nodes, smoothing is superior to buffering. However, for statistical services with $P_l > 0$, we find that H^* can be quite large.

To explore these issues, we first introduce background on admission control for statistical services and buffered multiplexers. We employ an algorithm that determines the per-node delay-bound-violation probability using rate-variance traffic envelopes [10], and use $P_l \approx H \cdot P_{l,n}$ as discussed above for end-to-end calculations. In particular, we characterize a traffic flow by the stochastic envelope

$$\sigma_j^2(t) = Var\left(\frac{A_j[s, s+t]}{t}\right) \tag{3}$$

and approximate the loss probability in a single node as

$$P_{l,n} \approx \max_t \Psi\left(\frac{C(t+D) - t\sum_j m_j}{t\sqrt{\sum_j \sigma_j^2(t)}}\right) \tag{4}$$

where m_j is source j's mean rate, and $\Psi(x)$ is the Gaussian tail probability $\Psi(x) = \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-t^2/2} dt$.

Theorem 2 *For identical traffic flows with maximum end-to-end delay bound D, the critical route length H^* is a non-decreasing function of P_l for any traffic envelope $\sigma^2(t)$.*

Proof: Suppose the critical route length is H when the end-to-end QoS requirements are (D, P_1). Consider a flow W traversing exactly H hops, with QoS requirements (D, P_1). Then the admissible regions for a single node in both $\mathcal{B}$ and $\mathcal{S}$ along the path W are the same and denoted by N_1. Without loss of generality, we assume $\mathcal{B}$ and $\mathcal{S}$ provide the same node loss probability, i.e., $P_{1,\mathcal{B},n} = P_{1,\mathcal{S},n}$.

Now suppose the end-to-end QoS requirements change to (D, P_2), and $P_2 < P_1$. The admissible region of $\mathcal{S}$ has to change to N_2 to satisfy the QoS change, and we have $N_2 \leq N_1$. To prove the theorem, it is equivalent to prove that $\mathcal{S}$ provides a lower node loss probability than $\mathcal{B}$ for W when admitting N_2 flows, i.e., $P_{2,\mathcal{S},n} \leq P_{2,\mathcal{B},n}$.

From Equation (4), we have $P_{1,\mathcal{S},n} \approx \Psi\left(\frac{C - mN_1}{\sqrt{N_1}\sigma_{\mathcal{S}}}\right)$, where m is the mean rate of the flow and $\sigma^2_{\mathcal{S}}$ is the variance of the smoothed trace. Similarly, $P_{2,\mathcal{S},n} \approx \Psi\left(\frac{C - mN_2}{\sqrt{N_2}\sigma_{\mathcal{S}}}\right)$ and $P_{1,\mathcal{B},n} \approx \Psi\left(\frac{C + \frac{B}{t^*} - mN_1}{\sqrt{N_1}\sigma_{\mathcal{B}}(t^*)}\right)$, where $B = \frac{D}{H}C$, and $\sigma^2_{\mathcal{B}}(t)$ is the rate-variance function of interval length t in Equation (3), and t^* is the maximizing t in Equation (4).

With these relationships established, $P_{1,\mathcal{B},n} = P_{1,\mathcal{S},n}$ is equivalent to

$$\frac{C - mN_1}{\sqrt{N_1}\sigma_{\mathcal{S}}} = \frac{C + \frac{B}{t^*} - mN_1}{\sqrt{N_1}\sigma_{\mathcal{B}}(t^*)} \tag{5}$$

Since $N_2 \leq N_1$, we have $\frac{C - mN_2}{C - mN_1} \geq \frac{C - mN_2 + B/t^*}{C - mN_1 + B/t^*}$. And from Equation (5), we have $\frac{C - mN_2}{\sqrt{N_2}\sigma_{\mathcal{S}}} \geq \frac{C + \frac{B}{t^*} - mN_2}{\sqrt{N_2}\sigma_{\mathcal{B}}(t^*)}$.

Since t^* may not be the maximizing t for N_2 flows, t', $P_{2,\mathcal{B},n} \approx \Psi\left(\frac{C + \frac{B}{t'} - mN_2}{\sqrt{N_2}\sigma_{\mathcal{B}}(t')}\right) \geq \Psi\left(\frac{C + \frac{B}{t^*} - mN_2}{\sqrt{N_2}\sigma_{\mathcal{B}}(t^*)}\right) \geq \Psi\left(\frac{C - mN_2}{\sqrt{N_2}\sigma_{\mathcal{S}}}\right) \approx P_{2,\mathcal{S},n}$. Hence the theorem is established. □

We explore this result further by experimentally investigating the extent to which loss probability impacts the critical route length. Figure 4 depicts H^* vs. P_l for on-off and video sources using the admission control algorithm described above. Observe that both curves have the critical route length increasing with P_l supporting Theorem 2. Moreover, note that even for stringent loss probability requirements, the critical route length for the video trace is quite large (greater than 30), indicating that buffering is preferable in a wide range of scenarios.

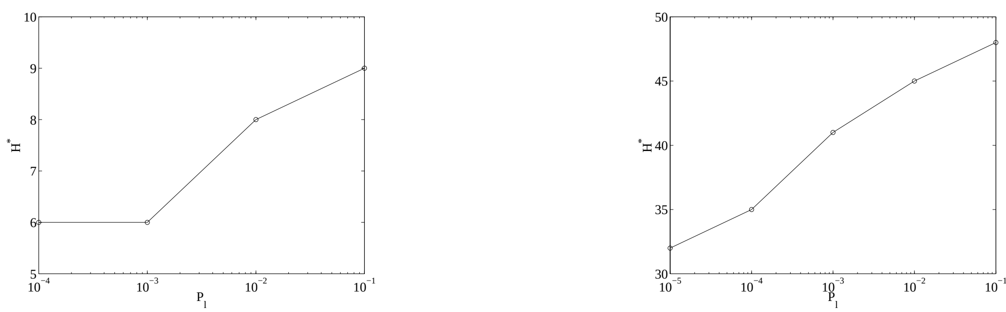

(a) Periodic On-off Sources ($D = 125$ msec) (b) Video Traces ($D = 100$ msec)

Fig. 4. H^* vs. P_l

4.3 The Impact of Traffic Characteristics on H^*

In this section, we further evaluate the impact of the traffic characteristics on H^*. Figure 5(a) depicts node utilization vs. path length for periodic on-off sources with a loss requirement of 10^{-3} and an end-to-end delay budget of 125 msec. While the buffering system's utilization is significantly higher for a single hop, the difference decays quickly with H, so that H^*, the number of hops beyond which smoothing is equivalent or superior, is 6 hops.

Results for similar experiments with video traces with a loss probability requirement of 10^{-3} and a delay budget of 100 msec are shown in Figure 5(b). We see the same trend as for on-off sources, but H^* is significantly larger, between 40 to 50 hops.

Thus, while both Figure 5(a) and (b) support Proposition 1, it is noteworthy how widely H^* varies for these two sources due to the different nature of their traffic characteristics.

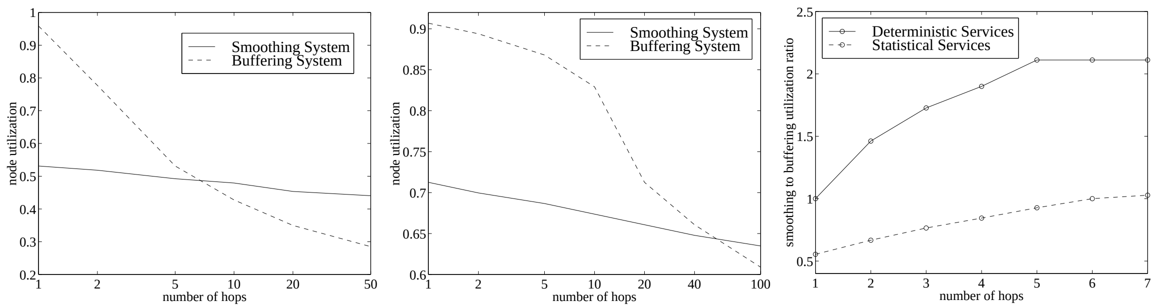

(a) Node Utilization: Periodic On-off Sources

(b) Node Utilization: Video Traces

(c) Smoothing to Buffering Utilization Ratio

Fig. 5. Admissible Region and Utilization Ratio

4.4 Utilization Ratio of Smoothing and Buffering Systems

In the experiments depicted in Figure 5(c), we further compare the smoothing and buffering systems by investigating the *ratio* of their admissible regions as a function of the number of hops traversed. The two curves represent the respective admissible regions of deterministic service, computed using [2], and statistical service, as in the experiments above. Notice that the point at which the ratio becomes greater than or equal to 1 is H^*, the route length at which the smoothing system becomes equivalent or superior. We make the following observations about the figure.

First, note that for one hop, the smoothing-to-buffering utilization ratio does not exceed 1 for both deterministic and statistical services, in agreement with Theorem 1. Moreover, for deterministic service the two systems attain the same utilization with one hop. The reason for this is that this experiment considers homogeneous sources which have identical critical time scales so that buffering's heterogeneity-of-time-scales gain (Section 3.2) is not available.

Second, notice that the curves for both deterministic and statistical services in Figure 5(c) have positive slopes. This is in agreement with Proposition 1 which states that buffering's advantages over smoothing diminish as the number of hops traversed increases.

Finally, observe that the critical route lengths for deterministic and statistical services are quite different. As Theorem 2 points out, H^* is a non-decreasing function of the loss probability and in these experiments $H^* = 1$ for deterministic service and $H^* = 6$ for statistical service.

5 Conclusions

In this paper, we conducted a comparative study of smoothing and buffering for end-to-end QoS. For a single multiplexer network, we demonstrated that buffering is superior to smoothing for both deterministic and statistical services due to the flows' heterogeneity of time scales, and for statistical services, a further gain from statistical buffer sharing. For multi-node networks, buffering's superiority diminishes as the number of hops traversed increases, and there is a "critical route length" beyond which smoothing is preferable. We further explored the impact of other system parameters, including QoS requirements and traffic characteristics on the relative merits of smoothing and buffering. We found that in contrast to video-on-demand systems and deterministic services, for delay-bounded statistical services, a traffic flow's end-to-end delay budget is often better spent in network buffers than in traffic smoothers at the network edge.

References

[1] C. Chang. Sample path large deviations and intree networks. *Queueing Systems, Theory and Applications*, 20(1-2):7–36, 1995.

[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. Quality of service guarantees in virtual circuit switched networks. *IEEE Journal on Selected Areas in Communications*, 13(6):1048–1056, August 1995.

[4] A. Elwalid and D. Mitra. Traffic shaping at a network node: theory, optimum design, admission control. In *Proceedings of IEEE INFOCOM '97*, Kobe, Japan, April 1997.

[5] W. Feng and J. Rexford. A comparison of bandwidth smoothing techniques for the transmission of prerecorded compressed video. In *Proceedings of IEEE INFOCOM '97*, Kobe, Japan, April 1997.

[6] D. Ferrari. Design and applications of a delay jitter control scheme for packet-switching internetworks. *Computer Communications*, 15(6):367–373, July 1992.

[7] D. Ferrari and D. Verma. A scheme for real-time channel establishment in wide-area networks. *IEEE Journal on Selected Areas in Communications*, 8(3):368–379, April 1990.

[8] L. Georgiadis, R. Guérin, V. Peris, and K. Sivarajan. Efficient network QoS provisioning based on per node traffic shaping. *IEEE/ACM Transactions on Networking*, 4(4):482–501, August 1996.

[9] M. Grossglauser, S. Keshav, and D. Tse. RCBR: A simple and efficient service for multiple time-scale traffic. *IEEE/ACM Transactions on Networking*, 5(6):741–755, December 1997.

[10] E. Knightly. Second moment resource allocation in multi-service networks. In *Proceedings of ACM SIGMETRICS '97*, pages 181–191, Seattle, WA, June 1997.

[11] E. Knightly. Enforceable quality of service guarantees for bursty traffic streams. In *Proceedings of IEEE INFOCOM '98*, San Francisco, CA, March 1998.

[12] E. Knightly and P. Rossaro. On the effects of smoothing for deterministic QoS. *Distributed Systems Engineering Journal: Special Issue on Quality of Service*, 4(1):3–15, March 1997.

[13] J. Kurose. On computing per-session performance bounds in high-speed multi-hop computer networks. In *Proceedings of ACM SIGMETRICS '92*, pages 128–139, Newport, RI, June 1992.

[14] J. McManus and K. Ross. Video-on-demand over ATM: constant-rate transmission and transport. *IEEE Journal on Selected Areas in Communications*, 14(6):1087–1098, August 1996.

[15] M. Reisslein, K. Ross, and S. Rajagopal. Guaranteeing statistical QoS to regulated traffic: the multiple node case. In *Proceedings of 37th IEEE Conference on Decision and Control*, Tampa, FL, December 1998.

[16] J. Salehi, Z. Zhang, J. Kurose, and D. Towsley. Supporting stored video: Reducing rate variability and end-to-end resource requirements through optimal smoothing. *IEEE/ACM Transactions on Networking*, 6(4):397–410, August 1998.

[17] T. Wu. *Efficient Statistical Service Provisioning in Broadband Networks*. M.S. Thesis, Rice University, May 1999.

[18] H. Zhang and D. Ferrari. Rate-controlled service disciplines. *Journal of High Speed Networks*, 3(4):389–412, 1994.

[19] J. Zhang and J. Hui. Applying traffic smoothing techniques for quality of service control in VBR video transmissions. *Computer Communications*, 21(4):375–89, April 1998.

[20] Z. Zhang, J. Kurose, J. Salehi, and D. Towsley. Smoothing, statistical multiplexing, and call admission control for stored video. *IEEE Journal on Selected Areas in Communications*, 15(6):1148–66, August 1997.