

Scheduling CBR Flows: FIFO or Per-flow Queuing?*

Jasleen Sahni, Pawan Goyal[†] and Harrick M. Vin

Distributed Multimedia Computing Laboratory
Department of Computer Sciences
University of Texas at Austin

‡Ensim Corporation
Mountain View, California

Abstract

In this paper, we study the effect of using FIFO or fair queuing on the end-to-end delay and jitter observed by CBR traffic in large-scale networks, where: (1) the bandwidth requirement and the packet sizes vary considerably across the CBR flows; (2) the class of CBR flows occupy different fractions of the total link bandwidth; and (3) the class of CBR flows share each network link with several other flows with different packet arrival patterns. Our results provide an empirical basis to evaluate the effectiveness of FIFO and per flow scheduling for CBR flows, as well as guidelines for deploying CBR services in the Internet.

1 Introduction

The Internet has traditionally supported the *best-effort* service model in which the network offers no assurance about when, or even if, packets will be delivered. This service model has proved to be adequate for elastic applications (e.g., ftp, telnet, and http) that tolerate packet delays and losses rather gracefully. With the commercialization of the Internet and the deployment of inelastic continuous media applications, however, the best-effort service model is increasingly becoming inadequate. For example, to meet the timeliness requirements of digital audio and video playback, most multimedia applications require greater predictability with respect to end-to-end delay and bandwidth than that offered by the current best-effort networks. To facilitate the co-existence of these emerging applications with conventional elastic applications, there is an increasing need for designing networks that *differentiate* between the services provided to different customers and applications.

One architecture for achieving service differentiation requires (1) networks to employ per-flow scheduling algorithms [3, 16] and (2) sources and receivers to exchange signaling messages that establish packet classification and forwarding state on each router along the path [14]. In this ar-

chitecture, however, the amount of state information required to be maintained at each router scales in proportion to the number of concurrent flows, which can be potentially large on high-speed links. To address this scalability requirement, the *differentiated services* architecture has been proposed [12]. This architecture achieves scalability by implementing complex classification and conditioning functions only at network boundary routers (which process lower volumes of traffic and lesser numbers of flows), and providing service differentiation inside the network for aggregated traffic rather than on a per-flow basis [12].

An example of this philosophy is the Virtual Leased Line service model [8, 12]. This service model guarantees a reserved rate to a flow. Flows requesting this service are shaped to *constant bit rate (CBR)* flows, and the packets of the flow are then marked as belonging to a particular service class (*expedited forwarding* service class in [8]) by appropriately setting the Type-of-Service (ToS) byte in the IP header of the packet [11]. The routers transmit packets belonging to this service class in the first-in-first-out (FIFO) order. To ensure that the rate guarantees of the flows can be met, the routers employ appropriate scheduling algorithm to ensure that on any given link, the rate available for this service class is at least the aggregate of the reserved rates of flows that request the Virtual Leased Line service. Note that this service does not guarantee an upper bound on delay; however, it is expected to closely resemble the service offered by a leased line.

The design of the Virtual Leased Line and other such service models are based on the following conjecture: If all the flows belonging to that service class are smoothed to CBR at the source or the ingress routers and if the bandwidth available to the service class is at least as large as the aggregate rate of all the flows, then FIFO scheduling of packets belonging to the aggregate is sufficient. The validity of this conjecture is essential for the deployment of many proposed services; yet, very little is known about: (1) the end-to-end delay and jitter observed by packets belonging to the CBR flows as they travel through many routers, and (2) the sizes of build-out buffer required at the end-point to restore the CBR nature of the flow prior to delivering the packets to the applications. This problem was investigated in [6]; however, that study assumes equal packet size for all flows. Hence, the conclusions of that study may not apply to the Internet, which allows variable-size network layer packets. We discuss in de-

*This research was supported in part by an AT&T Foundation Award, IBM Faculty Development Award, Intel, the National Science Foundation (CAREER award CCR-9624757, and Research Infrastructure Award CDA-9624082), Lucent Bell Laboratories, NASA, Mitsubishi Electric Research Laboratories (MERL), and Sun Microsystems Inc.

[†]This work was carried out when the author was with AT&T Labs—Research, Floram Park, New Jersey.

tail the related work in Section 6.

In this paper, we study—through simulations—the effect of using FIFO or fair queuing on the end-to-end delay and jitter observed by CBR traffic in large-scale networks, where: (1) the bandwidth requirement and the packet sizes vary considerably across the CBR flows; (2) the class of CBR flows occupies different fractions of the total link bandwidth; and (3) the class of CBR flows shares each network link with other flows with different packet arrival patterns. We find that the difference between the end-to-end performance of CBR flows yielded by FIFO and fair queuing is significant in networks that service flows with substantially different packet sizes. Since the current Internet allows flows to use different packet sizes, our results provide engineering guidelines for deploying CBR services in the Internet.

The rest of the paper is organized as follows. In Section 2, we provide a context for comparing the performance of FIFO and fair queuing for CBR flows. In Section 3, we describe our experimental setup, selection of experimental parameters, and the metrics for our evaluation. Sections 4 and 5 discuss results of our experiments. Section 6 describes the related work, and finally, Section 7 summarizes our conclusions.

2 Problem Formulation

FIFO is the simplest known discipline for scheduling flows in a network; it transmits packets in the order of their arrival. The simplicity of FIFO, however, comes at a cost—it is well-known that FIFO does not protect flows from each other; a burst of packet arrival from a flow affects the performance of all the other flows sharing the network. Thus, in a network with bursty traffic, FIFO does not provide end-to-end delay and bandwidth guarantees to flows. Furthermore, it has been shown that FIFO preserves or increases bunching of packets within a flow [4], thus worsening the end-to-end delay and jitter as the size of the network increases.

To address the limitations of FIFO, several different packet scheduling algorithms have been proposed [3, 16]. These algorithms maintain a per-flow state and provide bounded end-to-end delay guarantee to a flow regardless of the behavior of other flows in the network. In particular, it has been shown that for a leaky bucket flow, a network employing the Guaranteed Rate (GR) scheduling algorithms [5] guarantees that the maximum end-to-end delay is bounded by

$$\frac{\sigma}{R_f} + \frac{(K-1)L}{R_f} + \sum_{i=1}^{i=K} \beta^i \quad (1)$$

where σ , L , and R_f , respectively, denote the burstiness, the maximum packet size, and the bit rate requirement of the flow; K is the number of hops the flow is traversing; and β^i is a constant that depends on the scheduling algorithm at switch i . Though this result is applicable for a large class of

algorithms, since it has been shown that fairness is a desirable property of algorithms [3], we restrict our attention to the class of fair queuing algorithms in GR.

Now, consider a network where all the flows are smoothed out to CBR at the edge of the network (as is the case in the Virtual Leased Line service model). Since a CBR flow has a burst size of one packet, from (1), we conclude that the end-to-end delay guarantee for a CBR flow in case of fair queuing is:

$$K \frac{L}{R_f} + \sum_{i=1}^{i=K} \beta^i = K * I_f + \sum_{i=1}^{i=K} \beta^i \quad (2)$$

where $I_f = \frac{L}{R_f}$ is the inter-packet spacing for the CBR flows.

FIFO does not provide such a bound. However, it is conjectured that due to the elimination of traffic bursts at the edge of the network, the queues at network switches may never grow to a significant size; therefore, FIFO may provide worst-case delays that are comparable to the delay bounds guaranteed by fair queuing algorithms. The main objective of this work is to evaluate—through simulations—the validity of this conjecture in the context of large-scale networks.

3 Experimental Methodology

We have developed a network simulator using CSIM [1] to study the effect of using FIFO or fair queuing on the end-to-end performance of CBR traffic in large-scale networks. In this section, we describe our simulation environment, the design of our experiments, and the metrics for the performance evaluation.

3.1 Simulation Environment

3.1.1 Network Topology

For our experiments, we consider a linear, multi-hop network topology (see Figure 1(a)). This network model is fairly general and has been used in literature [7, 9, 10, 15]. Let M_n denote a linear, multi-hop network topology with n routers, and let R_i ($i \in [1, n]$) denote the i th router in the topology. Given such a topology, we are interested in the end-to-end performance of *tagged traffic*, which refers to the set of CBR flows that enters the network topology at router R_1 and traverses the multi-hop network topology M_n . Specifically, we measure how the characteristics of the individual CBR flows aggregated in the tagged traffic are altered as they interact with other CBR aggregates (referred to as the *cross traffic*) that enter and depart the network at each router along the path.

We model each router in this network as having three input ports (I_1, I_2 , and I_3) and three output ports (O_1, O_2 , and O_3). Each input port de-multiplexes 1/3 of its flows to each of the three output ports (see Figure 1(b)). Using these routers, we construct the linear, multi-hop network topology M_n as follows: M_n consists of n routers such that, for all i




Figure 1: Simulation environment: network topology and router architecture

($1 \leq i \leq n - 1$), the output port O_2 of router R_i is connected to the input port I_2 of router R_{i+1} . Through each router port, an aggregation of CBR flows enters the network. We refer to $1/3$ of the flows entering port I_2 of router R_1 as the tagged traffic. For each router, the traffic routed to the output port O_2 consists of: (1) The tagged traffic (entering the router from port I_2); (2) $1/3$ of the flows entering from input port I_1 ; and (3) $1/3$ of the flows entering from input port I_3 . We refer to flows entering from ports I_1 and I_3 that are routed to port O_2 as cross traffic. All of the remaining traffic entering each router is routed to ports O_1 and O_3 .

The above topology ensures that: (1) the tagged traffic that enters the network at router R_1 is routed all the way through the multi-hop network M_n , and (2) the cross traffic entering the network at router R_i ($i \in [1, n]$) interferes with the transmission of the tagged traffic for a single hop, and leaves the network at router R_{i+1} . This topology facilitates experimentation with different compositions of the cross traffic and different network depths.

3.1.2 Modeling Cross Traffic

The extent to which the cross traffic entering each router affects the characteristics of the CBR flows in the tagged traffic depends on the *burstiness* of the cross traffic. Note that, although we have assumed that each flow entering the network is shaped to CBR at the source, the aggregate cross traffic entering each router may be bursty. This burstiness results from: (1) the inherent differences in the bit rate requirements and the phases of the CBR flows, and (2) the traffic distortions in the network.

To reasonably approximate this burstiness, we model the cross traffic entering at each router in the network as consisting of two types of flows: (1) flows that are at the beginning of their routes or have traversed through a *small* number of routers, and (2) flows that are at the end of their routes or have traversed through a *large* number of hops. This model closely approximates the current Internet—each backbone router is a small number of hops away from some set of hosts while being far away from some others. Given this model, the interesting question in modeling the cross traffic is: what are the

reasonable values for *small* and *large*?

To address this question, we conducted an experiment with a seven-hop network topology (M_7) with 40Mbps link bandwidth and fixed size (512 Bytes) packets. The CBR flows entering the network were selected from two classes: (1) flows with bit rates in the range 1.5-5 Mbps, and (2) flows with bit rates in the range 32-95 Kbps. The bit rate for each flow was selected randomly from the respective ranges. Flows from each class were assigned equal share of the link bandwidth. The link utilization was set to 97%. For this experimental setup, we measured the burstiness of the tagged traffic emanating from port O_2 of each router in the M_7 topology. Since all the flows have the same packet size, we measure burstiness in terms of *index of dispersion for intervals (IDI)*; IDI measures the variability of packet arrivals at various time scales¹. Figure 2 plots the IDI values observed for the tagged traffic, assuming that packets are scheduled at the routers using FIFO and WF²Q+. The results indicate that: (1) although the variability in the inter-packet arrival times for individual flows may increase as CBR flows traverse through a multi-hop network, the IDI values for the aggregate tagged traffic reduces with increase in the hop count; and (2) IDI values for the aggregate tagged traffic does not change appreciably after 4 hops. Hence, for the rest of this paper, we will assume that the cross traffic entering port I_1 of each router in the network has traveled through 1 or 2 network hops, and the traffic entering port I_3 has traveled through 4 hops².

¹The IDI for an inter-arrival process $\{X_i\}$ is formally defined as [6]:

$$J_i(n) = \frac{n * var[\sum_{k=1}^n X_{i+k}]}{E^2[\sum_{k=1}^n X_{i+k}]}$$

If the process is wide-sense stationary, then $J_i(n) = J(n)$ for all i . The function $J(n)$ describes the variation in the inter-packet arrival process at different time scales. Note that for CBR flows, $J(n) = 0$, while for a Poisson arrival process, $J(n) = 1$ for all n .

²Note that our experiment to measure the burstiness of the aggregate flows was conducted for a network with fixed packet size. The conclusions presented here may change slightly if the above experiment is repeated in a different network environment.

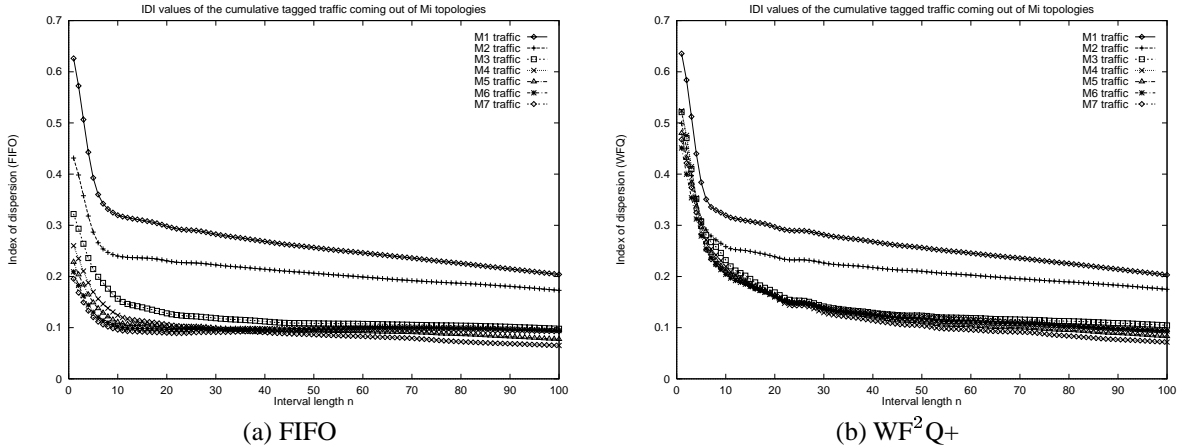


Figure 2: IDI values for traffic emanating from router R_i ($i \in [1, 7]$) in an M_7 network topology.

3.2 Design of Experiments

We conduct two sets of experiments; one for the network environment that supports a single class of service (namely the CBR service class), and the other for the network environment that simultaneously supports multiple service classes.

3.2.1 CBR Networks

In this set of experiments, we compare the end-to-end performance of CBR flows in networks where each router employs either FIFO or Worst-case Fair Weighted Fair Queuing (WF^2Q+) [2] to arbitrate access to link bandwidth, and (1) all flows are shaped to CBR at the source or at the ingress routers, and (2) all the flows request the same class of service. We experiment with different compositions of CBR flows and different network configurations.

Experiment 1: A CBR flow f_i is characterized by specifying its bit rate requirement r_i and its packet size l_i . In a large-scale network, such as the Internet, both r_i and l_i may vary considerably across the CBR flows. In this experiment, we systematically evaluate the effect of heterogeneity along each of these two dimensions on the end-to-end performance of individual CBR flows. We simulate the following three network environments.

1. A network environment in which all CBR flows have different bandwidth requirements, but the packet size for all the flows is the same. ATM networks are good examples of such environments.
2. A network environment in which all flows have the same bandwidth requirement, but the packet size for the flows are different.
3. A network environment in which the bit rate requirement as well as the packet sizes vary considerably across flows. The current Internet is an excellent example of such a network.

Experiment 2: In this experiment, we evaluate the effect of increasing the number of hops that the tagged traffic traverses through on the end-to-end performance of individual CBR flows. This allows us to draw conclusions about large networks.

Experiment 3: In this experiment, we study the impact of different network utilization levels on the results of Experiment 1 and 2.

We present the results of these experiments in Section 4.

3.2.2 Multi-class Networks

The differentiated services architecture is expected to support many different types of end-to-end services. A network can enable the co-existence of these services by employing packet scheduling algorithms to protect application classes from one another. Figure 3 depicts such a scheduling framework with k service classes (denoted by C_1, C_2, \dots, C_k). The link bandwidth available to a service class is proportional to its weight (denoted by w_1, w_2, \dots, w_k). The proportionate fair allocation of link bandwidth among classes is achieved by WF^2Q+ ; the transmission order of packets belonging to a service class is determined by a class-specific scheduler.

The objective of this set of experiments is to compare the end-to-end performance of CBR flows in network environments that: (1) supports multiple service classes, (2) uses WF^2Q+ to allocate link bandwidth among the service classes, and (3) uses either FIFO or WF^2Q+ to determine the order for transmitting packets belonging to the CBR service class. We conduct the following experiments.

Experiment 1: Flows that belong to other service classes may follow different packet arrival processes. In this experiment, we study the effect of two different arrival processes—back-logged and on-off—for flows belong-

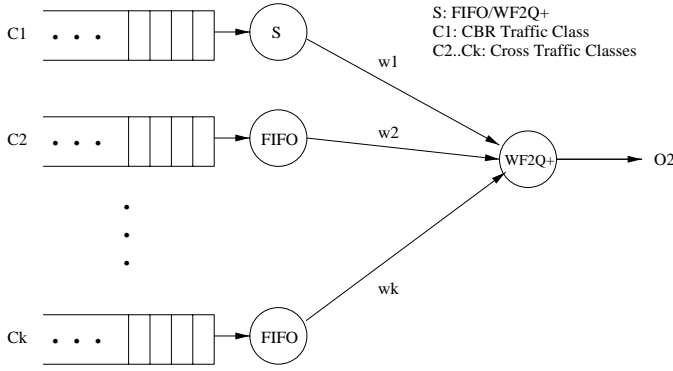


Figure 3: Scheduling framework for each output port

ing to other service classes on the end-to-end performance of CBR flows.

Experiment 2: In this experiment, we evaluate the effect of supporting different number of service classes on the end-to-end performance of CBR flows.

Experiment 3: In environments that support multiple service classes, different fractions of the total link bandwidth may be available to the CBR service class. In this experiment, we determine the effect of allocating different fractions of the total link bandwidth to the CBR class on the end-to-end performance of individual CBR flows.

We present the results of these experiments in Section 5.

3.3 Performance Measures

We use the following three metrics for our evaluation:

1. *End-to-end queuing delay:* For each flow f_i , we measure the distribution of the end-to-end queuing delay (d_i)—the difference between the end-to-end delay and the propagation latency—suffered by its packets.
2. *Normalized inter-packet separation:* For each flow, we measure the distribution and the variance of the normalized inter-packet separation at the destination. The normalization is performed with respect to the inter-packet spacing of the CBR flow at the source. This measures the extent to which a CBR flow is distorted while it travels through multiple routers.
3. *Build-out buffer size:* This is the size of the buffer required at the destination node to restore the spacing between packets to its initial value at the source. We estimate the build-out buffer size as follows.

Let d_i^{min} and d_i^{max} , respectively, refer to the minimum and the maximum delay suffered by packets of flow f_i . Then, all packets that experience the minimum delay d_i^{min} have to be buffered at the destination for an interval

of length $(d_i^{max} - d_i^{min})$ to reconstruct the nominal delay d_i^{max} . Since r_i denotes the bandwidth requirement for flow f_i , the total build-out buffer (in terms of bytes) can be estimated as: $B = [(d_i^{max} - d_i^{min}) * r_i]$. We determine the build-out buffer required for each flow, and then derive the average build-out buffer requirement.

4 Evaluation of CBR Network

In this set of experiments, we compare the end-to-end performance of CBR flows in network environments where each router employs either FIFO or WF²Q+ to arbitrate access to link bandwidth, and: (1) the CBR flows have heterogeneous requirements; (2) CBR flows traverse different number of hops; and (3) the average network utilization levels are different. For each of these experiments, we simulate a linear, 20-hop network topology (M_{20}), with 40Mbps link bandwidth.

4.1 Heterogeneity in CBR Flows

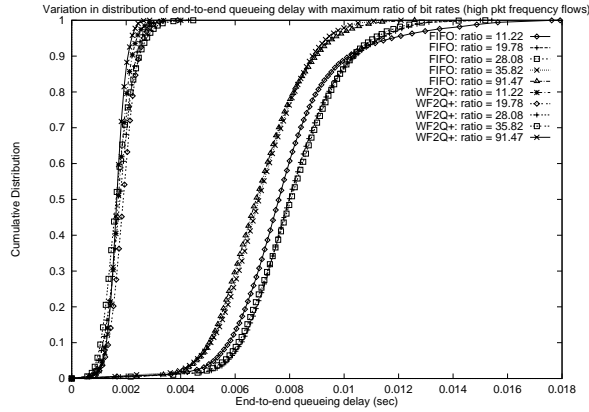
We experiment with tagged traffic consisting of flows selected from two classes; we vary the ratio of the average packet arrival rates of the flows belonging to the two classes from 10 to 100. Note that, for a CBR flow, the rate of packet arrival is given by r_i/l_i , where r_i and l_i , respectively, refer to the bandwidth requirement and packet size for flow f_i . Hence, we construct flows that belong to these classes either by appropriately controlling r_i , l_i , or both. We assume that flows belonging to each class get an equal share of the link bandwidth. Moreover, we select flows from each of these two classes such that the overall network utilization is 97%.

4.1.1 Heterogeneity in Bandwidth Requirements


We simulate a network environment where all CBR sources have packets of size 512 bytes, but their bandwidth requirements are selected from intervals such that the ratio of average bandwidth requirement of flows belonging to the two classes varies from 10 to 100. Figures 4, 5, and 6 depict the results of our simulations. The following conclusions can be derived from these figures.

1. Unlike FIFO, which schedules packets in the order of their arrival, WF²Q+ schedules packets in the increasing order of their finish tags [2]. Since finish tag computation is governed by the inverse of the packet arrival rate, WF²Q+ provides lower end-to-end queuing delay to flows with high packet arrival rate while increasing the end-to-end queuing delay for flows with low packet arrival rates. Figure 4 illustrates this behavior.

Perhaps a more interesting observation is that the end-to-end delay distributions do not change significantly with increase in the ratio of average packet arrival rates of the two classes of flows. Since this experiment assumes that




(a) High packet frequency flows




(b) Low packet frequency flows

Figure 4: Cumulative distribution of absolute end-to-end queuing delay for different ratio of bit rates




(a) High packet frequency flows




(b) Low packet frequency flows

Figure 5: Cumulative distribution of normalized inter-packet spacing for different ratio of bit rates



(a) High packet frequency flows



(b) Low packet frequency flows

Figure 6: Average buffer space required at destination to rebuild a CBR flow

packet sizes are identical for all flows, Figure 4 suggests that FIFO may be adequate for supporting CBR services in ATM networks (similar result was also noted in [6]).

2. Figure 5 illustrates that for flows with higher rates of packet arrival, the variance in the normalized inter-packet spacing increases with increase in the heterogeneity of CBR flows. For flows with low rates of packet arrival, on the other hand, the variance for both FIFO and WF²Q+ does not change appreciably. For FIFO, the maximum inter-packet spacing at the destination is about 3 times the corresponding value at the source; for WF²Q+, the maximum is about 1.5 times.
3. Figure 6 indicates that the average build-out buffer requirements for flows with high arrival rates increase with increase in the ratio of average packet arrival rates of the two classes of flows. Furthermore, FIFO requires about 4 times larger build-out buffer size than WF²Q+. For flows with low arrival rates, on the other hand, the build-out buffer size does not depend on the ratio.

Note that, for flows with high rates of packet arrivals, Figure 4 indicate that the maximum and the minimum delay observed by packets do not appreciably change with increase in heterogeneity. Hence, the increase in the build-out buffer size can be attributed predominantly to the increase in the bandwidth requirement of flows.

4.1.2 Heterogeneity in Packet Sizes

We simulate a network environment where all CBR flows have the same bandwidth requirement (namely, 113 Kbps), but their packet sizes are selected from intervals such that the ratio of average packet sizes of flows belonging to the two classes varies from 10 to 100. To obtain this range, we kept the interval of smaller packet sizes fixed, and changed the interval from which larger packet sizes are chosen. Note that the larger the packet size, the smaller is the rate of packet arrivals. Figures 7, 8, 9, and 10 plot results of our experiments.

1. Figure 7(a) indicates that under WF²Q+ scheduling, the end-to-end delay suffered by flows with high packet arrival rates does not change appreciably; however, with FIFO, the end-to-end delay increases substantially. Figure 7(b) indicates that, with increase in packet size, and hence an increase in the inter-arrival time for packets, flows suffer a larger end-to-end queuing delay with WF²Q+. With FIFO, however, the end-to-end queuing delay does not change appreciably.

Note that since the packet size interval for smaller packets is kept fixed, the packet arrival rate for the corresponding flows remains unchanged. On the contrary, for the flows with smaller packet arrival rates, increasing the packet size reduces the arrival rate even further (and thereby increases the inter-arrival time for packets).

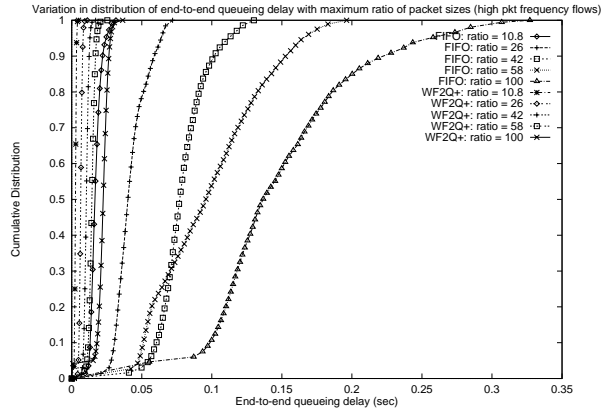
Hence, to put the results presented in Figure 7 in context, consider Figure 8, which plots 99.9%-percentile of the normalized end-to-end queuing delay for each flow; the normalization is done with respect to the inter-arrival time at the source of packets for each flow.

Figure 8(a) indicates that the 99.9%-percentile normalized end-to-end queuing delay suffered by flows in a class increases with the increase in the packet sizes for flows in the other class. Furthermore, the flows suffer almost *an order of magnitude higher end-to-end queuing delay with FIFO than with WF²Q+*. Figure 8(b), on the other hand, indicates that, although the absolute end-to-end delay for flows with low packet arrival rates increases with increase in the heterogeneity in packet sizes, the normalized delay does not change appreciably. This illustrates that the increase in absolute end-to-end queuing delays for these flows is in proportion to the increase in their packet sizes.


2. Figure 9 indicates that the variance of the normalized inter-packet separation at the destination increases with increase in ratio of packet sizes of the flows belonging to the two classes. For flows with a higher frequency of packet arrivals, the variance increases from 0.05 to 0.63 for FIFO, and from 0.004 to 0.21 for WF²Q+ when the ratio of average packet sizes for the two classes increases from 10 to 100. The variance in the normalized inter-packet separation at the destination is not significant for flows with low packet arrival frequency. For FIFO, the maximum inter-packet separation at the destination is about 10 times the corresponding value at the source; for WF²Q+ the maximum is about 3 times.
3. Figure 10 indicates that the average build-out buffer requirement increases for flows belonging to both classes, which is consistent with the increase in the end-to-end queuing delay. For flows with higher packet arrival rates, the build-out buffer requirement for FIFO is about an order of magnitude larger than WF²Q+; for the other class of flows, WF²Q+ requires twice as much build-out buffer as FIFO. It is interesting to note that although the number of packets that need to be buffered at the destination is dramatically higher for flows with high packet arrival rates than its counterpart, due to the heterogeneity in packet sizes, the maximum buffer space requirements for the high frequency flows under FIFO is similar to that of low frequency flows under WF²Q+.

4.1.3 Heterogeneity in Bandwidth Requirements and Packet Sizes

In this section, we study the effect of heterogeneity in both the bandwidth requirements or the packet sizes on the end-to-end performance of CBR flows. To conduct this experiment, we simulate a series of network environments with different




(a) High packet frequency flows

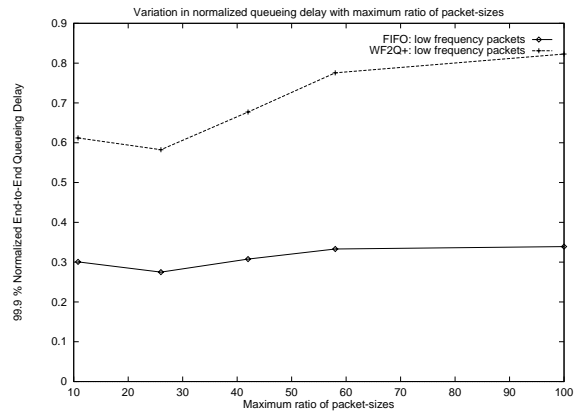


(b) Low packet frequency flows

Figure 7: Cumulative distribution of absolute end-to-end queuing delay for different ratio of packet sizes

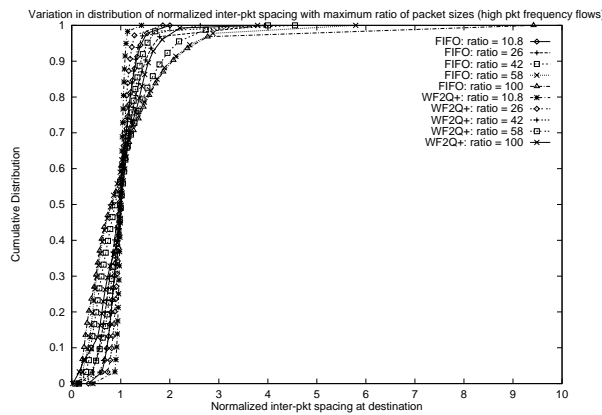


(a) High packet frequency flows




(b) Low packet frequency flows

Figure 8: 99.9%-percentile of the normalized end-to-end queuing delay for different ratio of packet sizes




(a) High packet frequency flows




(b) Low packet frequency flows

Figure 9: Cumulative distribution of normalized inter-packet spacing for different ratio of packet sizes



(a) High packet frequency flows



(b) Low packet frequency flows

Figure 10: Average buffer space required at destination to rebuild a CBR flow

ratios of packet sizes (namely, 5, 30, 60, and 120) for the two classes of flows. For each environment, we selected the bandwidth requirement of flows such that the ratio of packet arrival rates (or inter-packet separation) varied from 15 to 130. Figure 11 depicts the results of these experiments.

Figure 11(a) plots the 99.9%-percentile of the distributions of the end-to-end queuing delay. It indicates that for networks that support CBR flows with small heterogeneity in packet sizes, increasing the ratio of packet arrival rates (by changing the ratio of bandwidth requirements of flow classes) does not yield any noticeable increase the end-to-end queuing delay. However, for network environments that support larger heterogeneity in packet sizes (e.g., packet size ratios of 60 and 120), increasing the ratio of packet arrival rates increases the maximum end-to-end delay suffered by flows with high arrival rates. Figure 11(b) supports that same conclusion with respect to the normalized inter-packet separation.

4.2 Effect of Network Depth

In this section, we evaluate the effect of network depth (i.e., the number of hops that a flow traverses through) on the end-to-end performance of CBR flows. For this experiment, we select CBR flows from two classes: (1) flows with bit rates in the range 1.5-5 Mbps, and (2) flows with bit rates in the range 32-95 Kbps. The bit rate for each flow was selected randomly from the respective ranges. For these flows, we consider three possible scenarios for packet sizes: (1) high-bandwidth flows use large packet sizes, and vice versa; (2) all flows have the same packet size; and (3) high-bandwidth flows use small packet sizes, and vice versa. These selections yield CBR flow classes with ratios of packet arrival rates of 2, 36, and 540, respectively. Figures 12, 13, and 14 plot the results of these experiments.

Figures 12 and 14, respectively, indicate that the 99.9%-percentile of queuing delay suffered by flows and the average build-out buffer requirement at the destination increases lin-


early with increase in the number of hops traversed. The normalized inter-packet spacing, on the other hand, increases at first with increase in the hop count, but saturates after a few hops (see Figure 13). This illustrates that even with FIFO, the traffic distortion observed by CBR flows is bounded.

For the case where high-bandwidth flows use small packet sizes and vice versa, Figures 12 and 14 also illustrate that for high-bandwidth flows: (1) the end-to-end queuing delay observed by high-bandwidth flows is almost two orders of magnitude larger with FIFO than WF²Q+; and (2) the build-out buffer is about 7 times as large for FIFO than WF²Q+.

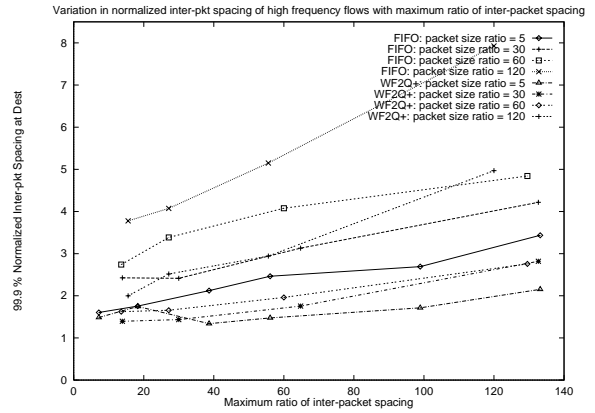
4.3 Effect of Network Utilization

In the previous sections, we had conducted all the experiments assuming a network environment at 97% utilization. In this section, we evaluate the effect of different network utilization levels (ranging from 40% to 97%) on the results presented in previous sections. Figures 15, 16, and 17 plot the results of this experiment.

Figures 15 and 17 indicate that for flows with high packet arrival rates, for FIFO scheduling, the end-to-end queuing delay and build-out buffer size reduce by a factor of 10 as the network utilization decreases from 97% to 40%. The performance under WF²Q+ remains mostly unaffected. Interestingly, in the presence of heterogeneous CBR flows, FIFO yields higher end-to-end queuing delays than WF²Q+ even at 40-50% utilization. For flows with low packet arrival rates, on the other hand, the end-to-end queuing delay and buffer size reduces for both FIFO and WF²Q+. Figure 16 illustrates that normalized inter-packet spacing at the destination is independent of the network utilization levels.

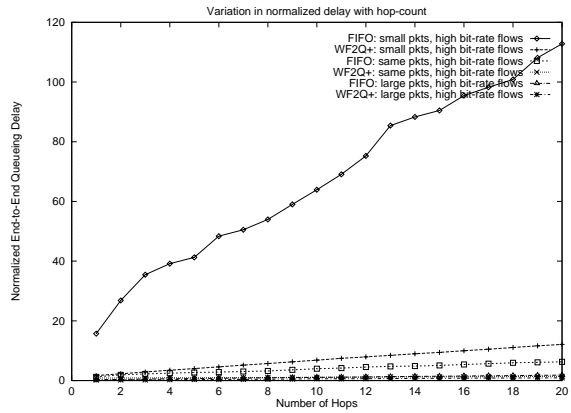


(a)

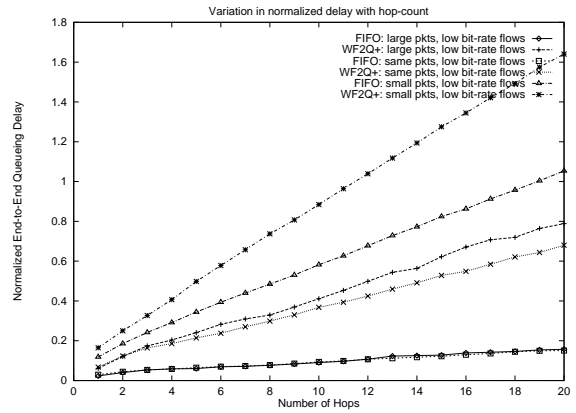


(b)

Figure 11: 99.9%-percentile of (a) end-to-end queuing delay and (b) normalized inter-packet spacing (for flows with high packet frequency)




(a) High packet frequency flows




(b) Low packet frequency flows

Figure 12: 99.9%-percentile of normalized end-to-end queuing delay




(a) High packet frequency flows

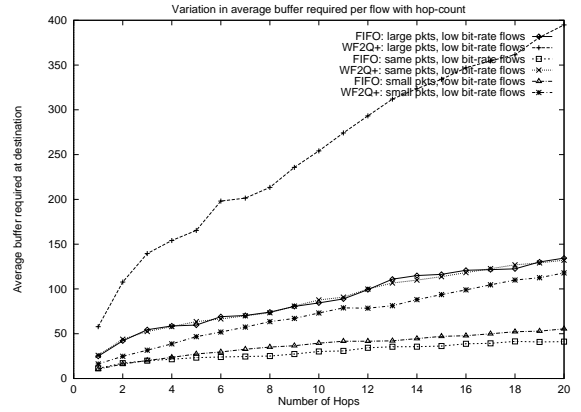


(b) Low packet frequency flows

Figure 13: 99.9%-percentile of normalized inter-packet spacing




(a) High packet frequency flows

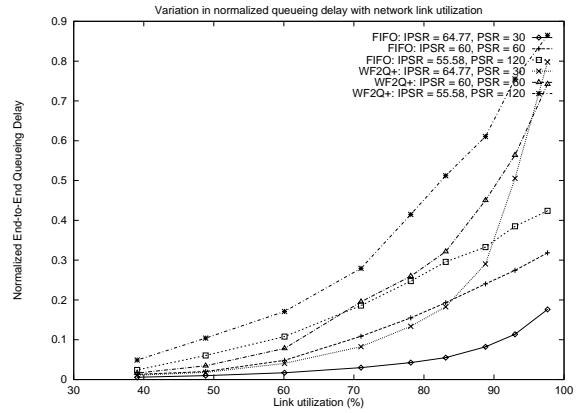


(b) Low packet frequency flows

Figure 14: Buffer requirement at destination

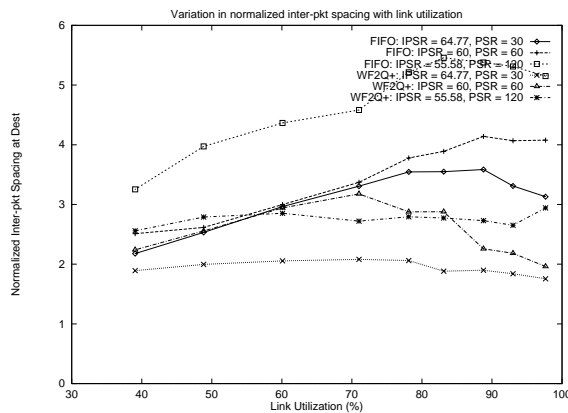


(a) High packet frequency flows




(b) Low packet frequency flows

Figure 15: 99.9%-percentile of normalized end-to-end queuing delay

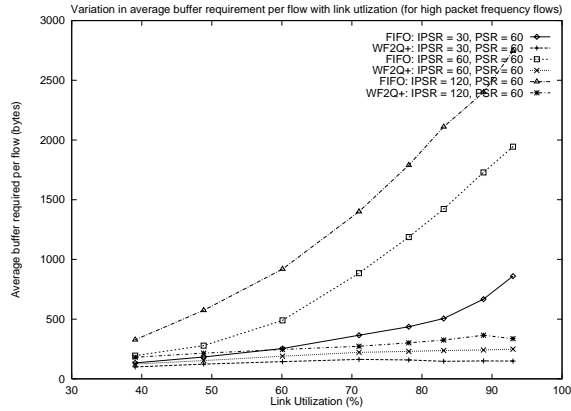


(a) High packet frequency flows

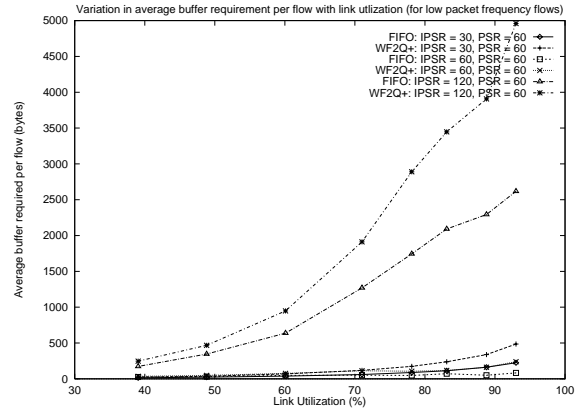


(b) Low packet frequency flows

Figure 16: 99.9%-percentile of normalized inter-packet spacing



(a) High packet frequency flows



(b) Low packet frequency flows

Figure 17: Buffer space required to rebuild a CBR stream at the destination

5 Evaluation in Multi-Class Networks

The objective of this set of experiments is to compare the end-to-end characteristics of CBR flows in networks that support multiple service classes, each receiving a fair-share of the link bandwidth according to its weight. We consider network scenarios where FIFO or WF²Q+ is used to arbitrate access to link bandwidth among flows within the CBR class and (1) flows in the other service classes have different arrival patterns; (2) different numbers of service classes are supported by the network; and (3) different fractions of the total link bandwidth are available to the CBR service class;

For each of these experiments, the CBR class is composed of flows with bandwidth selected from two intervals of 1.5-5 Mbps and 32-96 Kbps. The packets of all flows are 512 bytes in length. The network utilization level is set to 97%. The link bandwidth is scaled to an appropriate value so as to ensure that the CBR service class receives 40 Mbps link bandwidth.

5.1 Arrival Pattern of Competing Traffic

We simulate network environments with 15 classes of either *permanently back-logged* or *on-off* traffic sources, competing for 50 % of the link bandwidth with the CBR class. The on-off sources have an average on-time and off-time of 10 ms and 100 ms respectively. Figures 18 and 19 plot the 99.9%-percentile of the normalized end-to-end queuing delay and inter-packet separation observed at the destination as the CBR flows traverse a different number of hops.

These figures indicate that the end-to-end performance of the CBR flows is independent of the arrival pattern of the cross-traffic classes. This is because fair allocation of link bandwidth among the different classes of service masks any adverse effects of bursty arrival of packets for other service classes on the performance of CBR flows. Hence, for the remainder of this section, we will assume that all the cross-

traffic classes are on-off sources.

5.2 Number of Other Traffic Classes


We simulate network environments that support 2 through 16 different service classes. Independent of the number of service classes, we assume that the CBR service class receives 50% share of the link bandwidth. Figures 20 and 21 plot the results of these experiments (for an M_{20} topology). They indicate that, due to fair allocation of link bandwidth among the service classes, the end-to-end performance of flows in the CBR service class is affected only marginally with increase in the number of service classes.

5.3 Percentage link bandwidth to the CBR class


In this experiment, we simulate network environments with 15 other service classes sharing the link bandwidth with the CBR service class, and vary the percent share of link bandwidth available to the CBR service class from 10% to 100%. Figures 22 and 23 plot the results of this experiment (for an M_{20} topology). The figures indicate that with WF²Q+, the end-to-end performance of CBR flows does not change appreciably for different fractions of link bandwidth availability. However, with FIFO, the end-to-end queuing delay increases with reduction in the percentage of the link bandwidth available to the CBR class.

6 Related Work

There have been several analytical studies for evaluating the end-to-end performance of CBR flows. Roberts and Virtamo [13] derive bounds on the queue size distribution for superposed CBR streams with heterogeneity in their periods. Matragi et al. [10] provide techniques for estimating the end-to-end jitter incurred to CBR traffic in an ATM network.

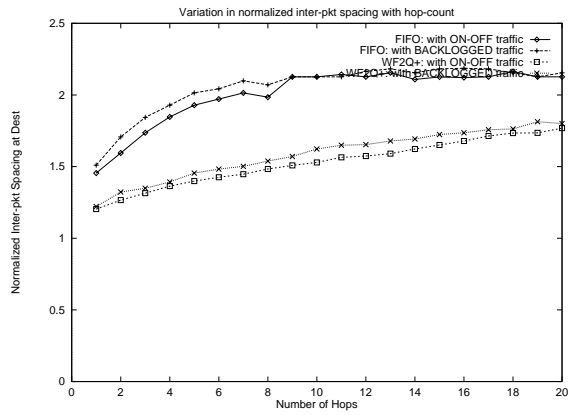


(a) High packet frequency flows




(b) Low packet frequency flows

Figure 18: 99.9 percentile of normalized end-to-end queuing delay - for different characterizations of other classes of traffic




(a) High packet frequency flows

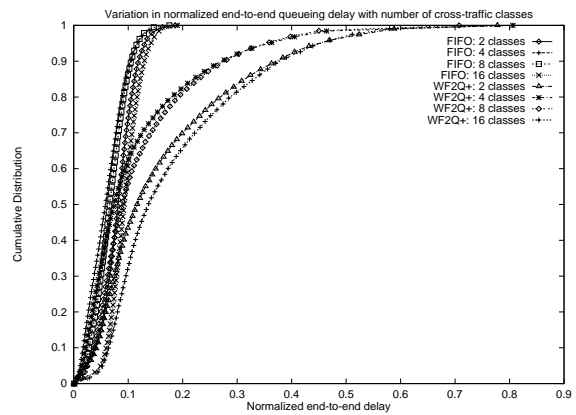


(b) Low packet frequency flows

Figure 19: 99.9 percentile of normalized inter-packets spacing - for different characterizations of other classes of traffic




(a) High packet frequency flows




(b) Low packet frequency flows

Figure 20: Cumulative distribution of normalized end-to-end queuing delay - for different number of other classes of traffic




(a) High packet frequency flows




(b) Low packet frequency flows

Figure 21: Cumulative distribution of normalized inter-packet spacing - for different number of other classes of traffic




(a) High packet frequency flows




(b) Low packet frequency flows

Figure 22: Cumulative distribution of normalized end-to-end queuing delay - for different percentage of link bandwidth available to the CBR class



(a) High packet frequency flows



(b) Low packet frequency flows

Figure 23: Cumulative distribution of normalized inter-packet spacing - for different percentage of link bandwidth available to the CBR class

Their study is thus limited to flows with fixed packet sizes, and they do not talk about the actual end-to-end delay incurred, or the build-out buffer requirement at the destination. Gruenenfelder [7] studies a connection going through multiple hops and observes that the end-to-end delay distribution depends on the auto-covariance of the cross-traffic sharing the network links. DeSimone [4] shows that the inter-arrival time of packets becomes more bursty in a traffic stream on passing through networks that carry packets of highly variable size. However, this study focuses on non-CBR traffic, and it is not obvious if the same result would hold for CBR networks. Furthermore, the analysis is carried out assuming a single cross traffic stream and a small network.

Grossglauser and Keshav [6] present a simulation study of the end-to-end delay and jitter observed by CBR flows as they traverse large networks while interacting with several cross traffic flows. However, their study assumes fixed packet sizes, making the study applicable to ATM-type networks.

7 Concluding Remarks

In this paper, we studied—through simulations—the effect of using FIFO or fair queuing on the performance of CBR traffic in the context of large-scale networks, in which: (1) the bandwidth requirement and the packet sizes vary considerably across the CBR flows; (2) the class of CBR flows occupy different fractions of the total link bandwidth; and (3) the class of CBR flows share each network link with several other flows with different packet arrival patterns. Our experiments demonstrated that:

1. For networks where all flows have the same packet size, end-to-end queuing delays do not change significantly with increase in heterogeneity of bandwidth requirements of CBR flows. Hence, FIFO is adequate to support CBR services in ATM networks.
2. End-to-end delay and build-out buffer size increases linearly with increase in heterogeneity of packet sizes. Flows can suffer almost an order of magnitude higher end-to-end queuing delay with FIFO than with WF^2Q+ . This result demonstrates the invalidity of the conjecture reported in Section 2, and thereby argues that in networks that service flows with substantially different packet sizes, per-flow scheduling techniques will be desirable.
3. The end-to-end performance for CBR flows worsens linearly with increase in the depth of the network.
4. At low-levels of network utilization, the difference between the end-to-end performance of CBR flows yielded by FIFO and WF^2Q+ reduces.
5. The above results are mostly unaffected by the number of other service classes that network may support.

References

- [1] CSIM18 - The Simulation Engine. <http://www.mesquite.com>.
- [2] J.C.R. Bennett and H. Zhang. Hierarchical Packet Fair Queueing Algorithms. *IEEE/ACM Transactions on Networking*, 5(5):675–689, October 1997.
- [3] A. Demers, S. Keshav, and S. Shenker. Analysis and Simulation of a Fair Queueing Algorithm. In *Proceedings of ACM SIGCOMM*, pages 1–12, September 1989.
- [4] A. DeSimone. Generating Burstiness in Networks: A simulation Study of Correlation Effects in Networks of Queues. *ACM Computer Communication Review*, pages 24–31, 1991.
- [5] P. Goyal and H.M. Vin. Generalized Guaranteed Rate Scheduling Algorithms: A Framework. *IEEE/ACM Transactions on Networking*, 5(4):561–571, August 1997.
- [6] M. Grossglauser and S. Keshav. On CBR Service. In *Proceedings of INFOCOM'96*, pages 129–137, March 1996.
- [7] R. Gruenenfelder. A Correlation Based End-to-End Cell Queueing Delay Characterization in an ATM Network. In *Proc. Thirteenth International Teletraffic Congress (ITC-13)*, volume 15, pages 59–64, June 1991.
- [8] V. Jacobson, K. Nichols, and K. Poduri. An Expedited Forwarding PHB. August 1998. available as IETF draft(draft-ietf-diffserv-phb-ef-00.txt).
- [9] J. Kurose. Open Issues and Challenges in Providing Quality of Service Guarantees in High-Speed Networks. *ACM Computer Communication Review*, 23:6–15, January 1993.
- [10] W. Matragi, C. Bisdikian, and K. Sohraby. Jitter Calculus in ATM Networks: Multiple Node Case. *IEEE INFOCOM '94*, June 1994.
- [11] K. Nichols, S. Blake, F. Baker, and D. Black. Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers. December 1998. Internet RFC 2474.
- [12] K. Nichols, V. Jacobson, and L. Zhang. A Two-bit Differentiated Services Architecture for the Internet. November 1997. <ftp://ftp.ee.lbl.gov/papers/dsarch.pdf>.
- [13] J.W. Roberts and J.T. Virtamo. The Superposition of Periodic Cell Arrival Streams in an ATM Multiplexer. *IEEE Transactions on Communications*, 39(2):298–303, February 1991.
- [14] S. Shenker and C. Partridge. Specification of Guaranteed Quality of Service. Available via anonymous ftp from <ftp://ftp.ietf.cnri.reston.va.us/internet-drafts/draft-ietf-intserv-guaranteed-svc-03.txt>, November 1995.
- [15] D. Yates, J.F. Kurose, D. Towsley, and M.G. Hluchyj. On per-session end-to-end delay distributions and the call admission problem for real-time applications with QOS requirements. In *Proceedings of ACM SIGCOMM*, pages 160–166, October 1993.
- [16] H. Zhang and S. Keshav. Comparison of Rate-Based Service Disciplines. In *Proceedings of ACM SIGCOMM*, pages 113–121, August 1991.