Providing Deterministic Guarantees For Bursty Traffic (Extended Abstract)

Hui Zhang and Domenico Ferrari
hzhang,ferrari@tenet.ICSI.Berkeley.EDU
The Tenet Group
Computer Science Division
University of California at Berkeley
and
International Computer Science Institute

1 Introduction

High speed networking has introduced opportunities for new applications such as video conferencing, scientific visualization and medical imaging. These applications have stringent performance requirements in terms of throughput, delay, delay jitter and loss rate. The best-effort service provided by the current packet-switching networks is not adequate. Services that support performance guarantees are needed.

To provide guaranteed service, resources have to be reserved within the network. One important consideration about any resource reservation scheme is whether it can achieve a high utilization of the network. In general, resource reservation schemes works well for smooth traffic. However, many clients requiring performance guarantees have bursty sources, for example, compressed video. If reservation is based on the peak rate of each connection, new requests will be rejected when the sum of the peak rates of all the connections reaches link speed. In this case, the network will be under-utilized by guaranteed service traffic when the peak-average rate ratio is high.

Therefore, it is important to derive solutions that can provide performance guarantees even when the sum of the peak rates of all the connections is greater than link speed, and to understand the relationship between the server utilization and traffic burstiness. In this paper, we discuss these issues within the framework of the Tenet real-time channel scheme, to be summarized in the sequel.

The reminder of the paper is organized as follows. Section 2 briefly reviews the real-time channel scheme and shows the limitations of the admission control tests for the deterministic service in the current form. Section 3 presents the analysis which provides delay bounds for the First-Come-First-Served (FCFS) discipline in both cases when the sum of the peak rates of all the connections is either larger or smaller than 1. Section 4 gives numerical examples to illustrate the results in Section 3. Section 5 discusses the implications of the results. Section 6 concludes the paper with a summary of the results.

2 Background

In this section, we give a brief overview of the current version of the Tenet resource management algorithms [3].

The Tenet algorithms are based on a communication abstraction called a *real-time channel* [4]. A real-time channel is a network channel associated with traffic and performance parameters. The parameters are provided by the clients to specify their traffic characteristics and performance requirements. The traffic specification consists of parameters $(Xmin, Xave, I, Smax)$, where $Xmin$ is the minimum packet inter-arrival time, $Xave$ is the worst-case average packet inter-arrival time over an an averaging interval, I is the averaging interval, and $Smax$ is the maximum packet size. In such a specification, it is easy to see that $Xave/Xmin$ is the peak-average-rate-ratio, which is an indicator of the traffic burstiness. It should be also noticed that I also affects the traffic burstiness. For the given values of $Xmin$ and $Xave$, I determines how long the source can continuously send packets at the peak rate in the worst case. The larger $Xave/Xmin$ and I, the burstier the traffic. Tenet algorithms provide both deterministic and statistical guarantees. For the purpose of this paper, we only consider deterministic guarantees. For a channel with deterministic delay bound $\overline{D}$, the network guarantees that the delays of all the packets on that channel will be less than $\overline{D}$ as long as the channel does not violate its traffic specification.

A channel needs to be *established* before data can be transferred. This channel establishment is achieved in the following manner: a real-time client specifies its traffic characteristics and end-to-end performance requirements to the network; the network determines the most suitable route for a channel with these traffic characteristics and performance requirements; it then translates the end-to-end parameters into local parameters at each node, and attempts to reserve resources at these nodes accordingly. This is done in a distributed manner during a round-trip communication.

In order to provide performance guarantees, two levels of controls are needed: at the channel level, channel admission control algorithms reserve resources for each of the channels and limit the maximum utilization of the network by real-time traffic; at the packet level, the service discipline at each of the switches determines the multiplexing policy and allocates resources to different channels according to the reservations.

As shown in [2, 10], many service disciplines can be used to provide real-time service. However, different service disciplines require different admission control algorithms. In [4], conditions are given for a variation of the Earliest-Due-Date discipline. Three tests need be satisfied before a new channel request can be accepted. The three tests are: (a) *deterministic test*, which limits the peak utilization of the server to be less than 1; (b) *delay bound or schedulability test*, which avoids the case of scheduling saturation; and (c) *buffer test*, which ensures enough buffer space is available. For the purpose of the discussion in this paper, we will focus on the deterministic test, which requires

$$\sum_{j=1}^{n} \frac{1}{Xmin_j} \times \frac{Smax_j}{l} < 1 \tag{1}$$

where l is the link speed [1] and n is the number of channels traversing the link.

Although, the tests guarantee that deterministic delay bounds can be provided, they are rather restrictive in the sense the peak utilization of the link by real-time traffic is limited to be less than 1. If the peak-average-rate-ratio is high, the average link utilization by real-time traffic will be low.

3 Delay Analysis

In this section, we show that deterministic delay bounds can be obtained even when the peak utilization of the link is greater than 1 (the average utilization of the link has to be less than 1). We show that the result holds even for simple discipline like FCFS. We use the techniques developed by [1] to obtain delay bounds for an FCFS server. We use the discrete traffic model $(Xmin, Xave, I, Smax)$ rather than the fluid model in [1]. A discrete model has the advantage that it better characterizes multimedia sources than fluid approximations.

Figure 1(a) illustrates some of the concepts used in the analysis. The horizontal axis is the time. The upper curve in the figure is the sum of the bits that have arrived. Each arrival of a packet causes a jump in this curve. The lower curve is the number of bits that have been transmitted. The difference between the two curves is the *backlog* function. A work-conserving server always transmits packets at a constant rate when there is a backlog. The points where the arrival curve and the service curve meet, or where the backlog is zero, divide the time axis into *busy periods* and *idle periods*.

One important observation is that the maximum length of a busy period is an upper bound on delay for any work-conserving discipline. Also, the longest backlog divided by link speed is an upper bound on delay for an FCFS server.

Central to the analysis is the concept of *traffic constraint function* b(.). $b_j(u)$ is defined to be the maximum number of bits that can arrive on connection j during *any* interval of length u. For a traffic source that obeys the model $(Xmin_j, Xave_j, I_j, Smax_j)$, it is easy to show that

$$b_j(u) = (min(\lceil \frac{u \ mod \ I_j}{Xmin_j} \rceil, \lceil \frac{I_j}{Xave_j} \rceil) + \lceil \frac{u}{I_j} \rceil \times \lceil \frac{I_j}{Xave_j} \rceil) \times Smax \tag{2}$$

[1] We consider a network with non-blocking switches, where queueing happens at the output link of each switch. The processing time of each packet is thus equal to the transmission time, and is bounded by $\frac{Smax_j}{l}$

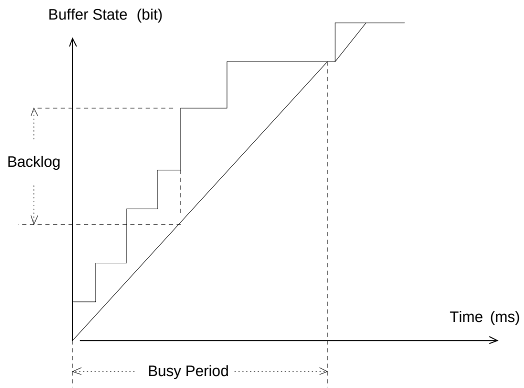

Concepts: Delay, Backlog and Busy Period
(a)

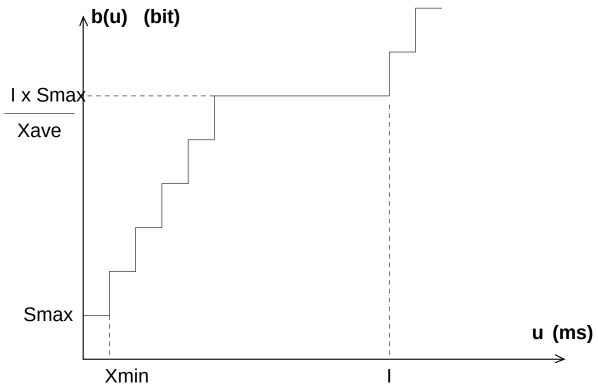

Traffic Constraint Function for (Xmin, Xave, I, Smax) specification
(b)

Figure 1: Graphical Illustration of the Analysis

Figure 1(b) illustrates the function *b(.)* for a channel with traffic specification of $(Xmin, Xave, I, Smax)$.

The following theorem bounds the delay for an FCFS server.

Theorem 1 *Let there be n channels multiplexed on a link with the FCFS service discipline and link speed l. If for $j = 1, \cdots, n$, the traffic on channel j is bounded by $b_j(.)$, then the delays of packets on all the channels are bounded by d, where d is defined by*

$$d = \frac{1}{l} \max_{u \geq 0} \left(\sum_{j=1}^{n} b_j(u) - l \times u \right) \tag{3}$$

The theorem for the case $n = 2$ has been stated and proven in [1]. The case of an arbitrary n can be proven by applying induction on n. Intuitively, equation 3 gives the maximum value of backlog in the server, which is an upper bound on delays for packets in an FCFS server.

Since $b_j(.)$ as defined in (2) is a piecewise linear function, evaluating (3) is equivalent to optimizing a piecewise linear function. This can be computed efficiently. The details of the algorithm are presented in [8]. In the following two corollaries, we give closed form solutions for two special cases. Corollary 1.1 considers the case when $\sum_{j=1}^{n} \frac{1}{Xmin_j} \times \frac{Smax_j}{l} \leq 1$, Corollary 1.2 considers the case of homogeneous sources when $\sum_{j=1}^{n} \frac{1}{Xmin_j} \times \frac{Smax_j}{l} > 1$. The proof of the two Corollaries are given in [8].

Corollary 1.1 *Let there be n channels multiplexed on a link with the FCFS service discipline and link speed l. Assume channel j has the traffic specification $(Xmin_j, Xave_j, I_j, Smax_j)$, $(j = 1, \cdots, n)$. If $\sum_{j=1}^{n} \frac{1}{Xmin_j} \times \frac{Smax_j}{l} \leq 1$, then the delays of packets on all the channels are bounded by d, where d is given by*

$$d = \frac{1}{l} \sum_{j=1}^{n} Smax_j \tag{4}$$

Corollary 1.2 *Let there be n homogeneous channels with traffic specification (Xmin, Xave, I, Smax) multiplexed on a link with the FCFS service discipline and link speed l. If $n \times \frac{1}{Xmin} \times \frac{Smax}{l} > 1$, the delays of packets from all the channels are bounded by d, where d is given by*

$$d = Xmin + (\rho_{ave} - \frac{1}{burstRatio})I \tag{5}$$

where $\rho_{ave} = n \times \frac{Xave}{I}$, and $burstRatio = \frac{Xave}{Xmin}$.

In (5), ρ_{ave} is the average utilization of the link, and $\frac{Xave}{Xmin}$ and I represent the burstiness of channels. Intuitively, the more loaded the link , and the burstier the channel traffic , the larger the delay bound. This is exactly what is shown by (5).

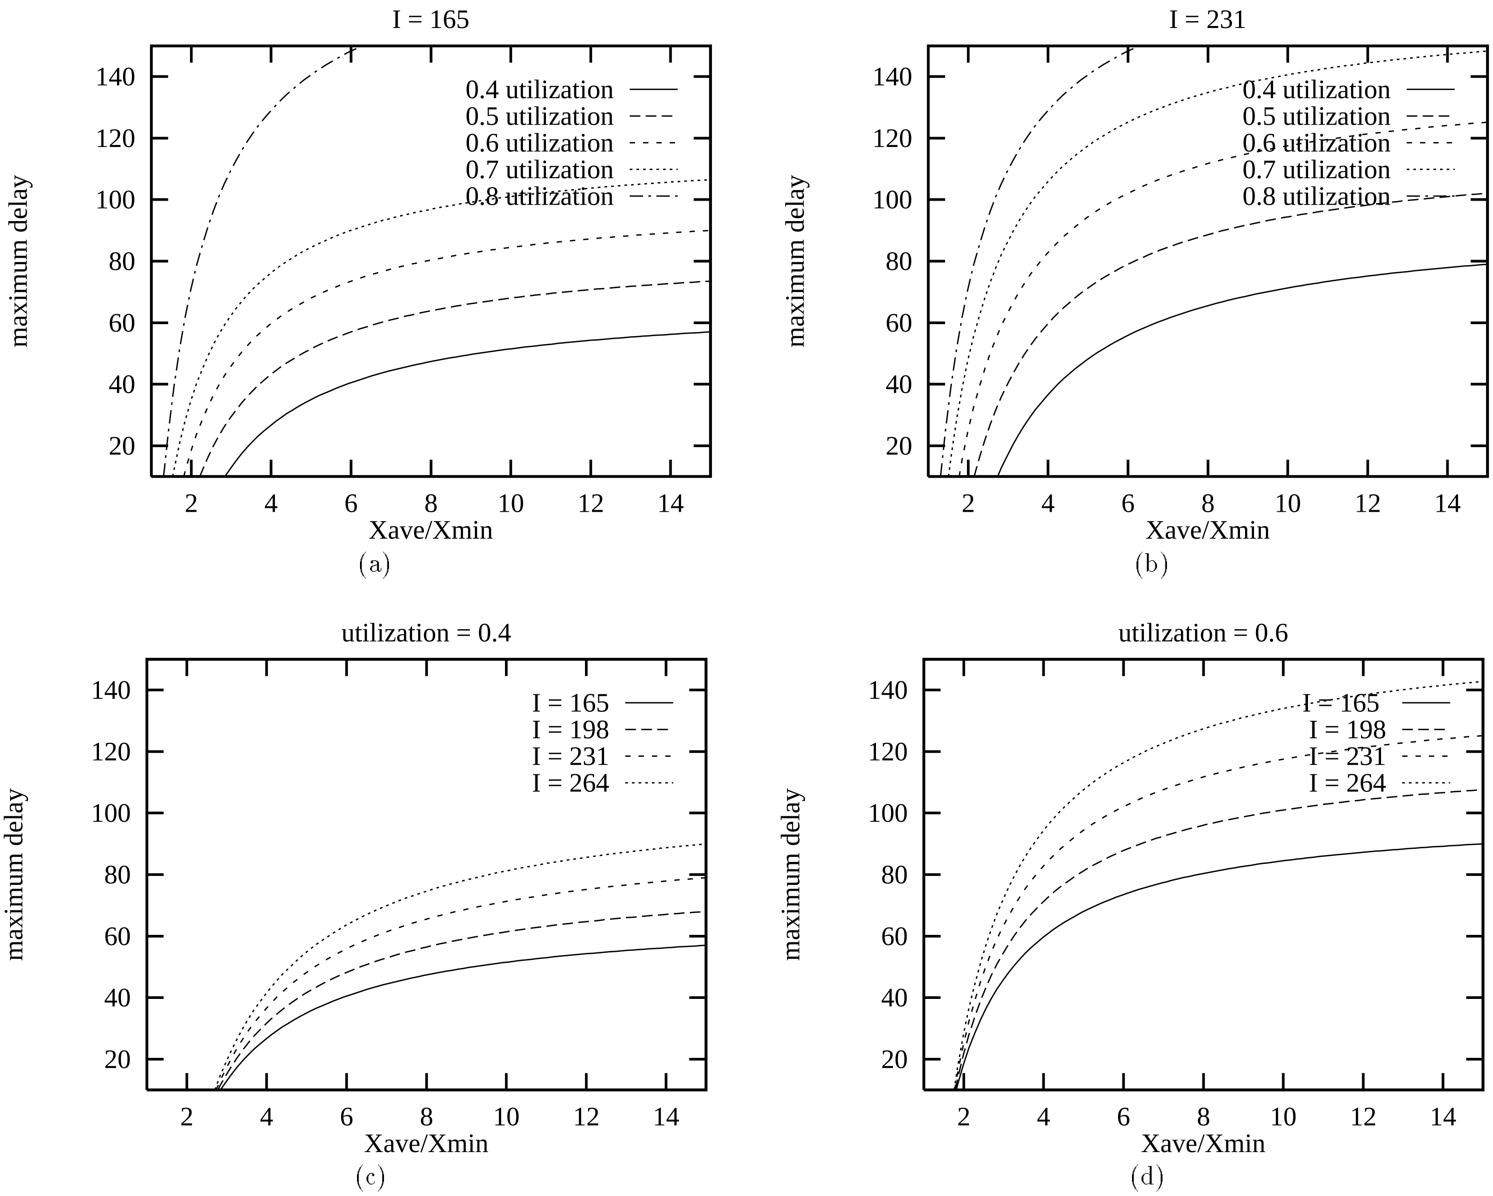

Figure 2: Effects of Burstiness and Utilization On Delay Bound

4 Numerical Examples

In this section, we show some numerical examples to illustrate the results presented in the previous section.

Figure (3) shows the effects of the traffic burstiness and the average utilization on delay bounds. The horizontal axis is the peak-average-rate ratio of the channels, the vertical axis is the delay bound. The figures only show the case when the sum of the peak rate on all the channels is greater than 1. Figures 2(a) and 2(b) show the cases with fixed averaging interval I but different average utilizations. Figures 2(c) and 2(d) show the cases with fixed utilization but different I's. The values of I are chosen to be multiples of 33 ms, which is the time interval between two frames in a 30 frames/sec video.

Figure (4) shows how much average or peak utilization can be achieved under certain delay bound constraints.

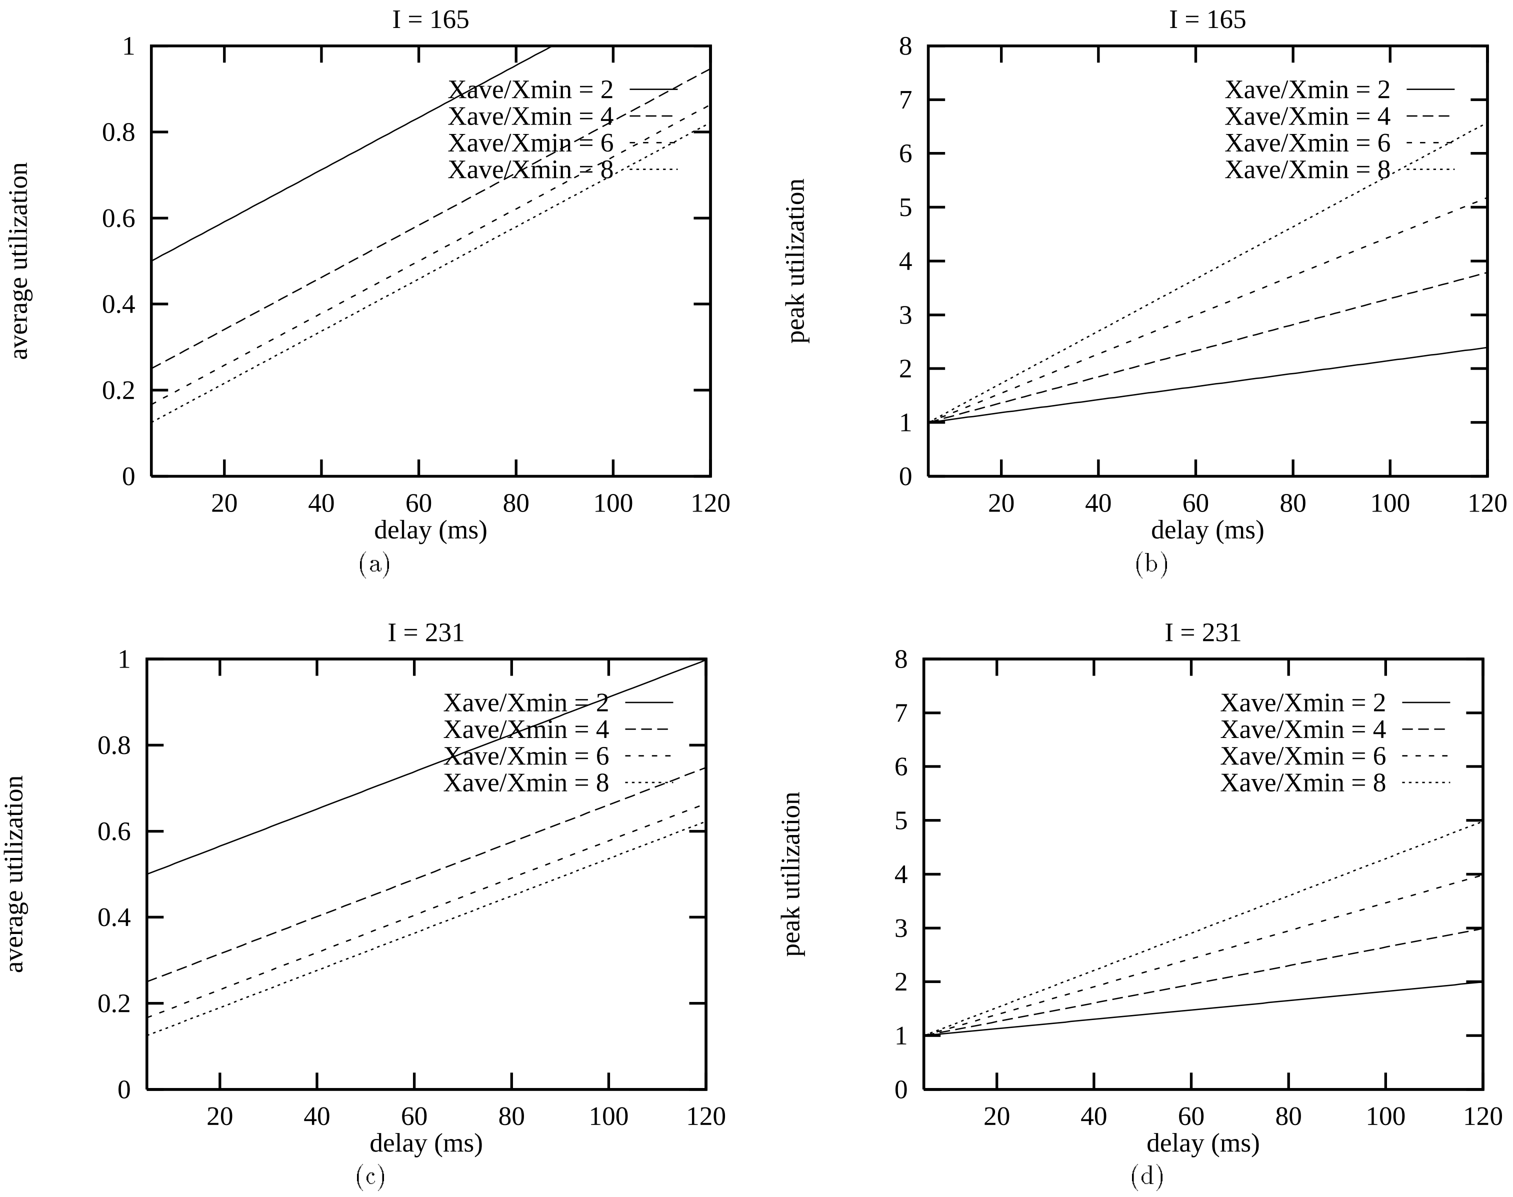

Figure 3: Average and Peak Utilizations vs. Delay Bound

5 Discussion

In the previous sections, we showed that even for simple service disciplines like FCFS, deterministic delay bounds can be obtained even when the sum of the peak rates of all the channels is greater than the link bandwidth. The question naturally arises: why does one need more complex service disciplines than FCFS? There are several reasons that FCFS alone is not enough.

First, guaranteed service requires that the network *protect* clients from two sources of variability: misbehaving users and network load fluctuations. Misbehaving users can send packets at a higher rate than the bandwidth allocated to them, and network load fluctuations may cause a higher instantaneous arrival rate from a channel at some switch, even though the channel satisfies the bandwidth allocation constraint at the entrance to the network. The results obtained in the previous section hold only when the traffic obeys the characterization at the entrance to the switch. The FCFS discipline does not offer such protections. Although putting traffic policing function at *all* network access points will prevent misbehaving users from affecting other users, traffic distortions due to network fluctuations suggest that protection should be implemented within the network using a rate-based service discipline [10].

In previous work [7, 8, 9], we have shown that, by having a server with two components, a schedule that can provide a local delay bound can be combined with a rate-controller, and end-to-end delay bounds can obtained from arbitrary topology network with these types of servers.

The second reason that FCFS is not enough is that an FCFS server can only offer a single value of delay bound for all the channels. However, the performance requirements for integrated services networks will be diverse. It is important to support multiple classes of Quality of Service. More advanced service disciplines can achieve that [5, 6, 9].

The third reason that FCFS may not be enough is that more sophisticated service disciplines can provide better bounds than the FCFS discipline.

6 Conclusion

In this paper, we showed that it is possible to provide deterministic delay bounds even when the sum of the peak rates of all the channels is greater than the link speed. We show that, even for a simple discipline like FCFS, a reasonable average utilization can be achieved for deterministic service when the burstiness of traffic is bounded.

For the traffic characterization $(Xmin, Xave, I, Smax)$, it is important to notice that the burstiness of the traffic does not only depend on peak-average-rate-ratio $\frac{Xave}{Xmin}$, but also depends on the averaging interval I.

The result for FCFS only holds in a single node case and when none of the channels is sending more packets than declared in the traffic specification. Thus, the result does not imply that more advanced service disciplines are not necessary.

References

[1] Rene L. Cruz. A calculus for network delay, part I : Network elements in isolation. *IEEE Transaction of Information Theory*, 37(1):114–121, 1991.

[2] Domenico Ferrari. Real-time communication in an internetwork. *Journal of High Speed Networks*, 1(1):79–103, 1992.

[3] Domenico Ferrari, Anindo Banerjea, and Hui Zhang. Network support for multimedia: a discussion of the Tenet approach. Technical Report TR-92-072, International Computer Science Institute, Berkeley, California, October 1992. Also to appear in *Computer Networks and ISDN Systems*.

[4] Domenico Ferrari and Dinesh 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.

[5] S. Jamaloddin Golestani. A stop-and-go queueing framework for congestion management. In *Proceedings of ACM SIGCOMM'90*, pages 8–18, Philadelphia Pennsylvania, September 1990.

[6] Abhay Kumar J. Parekh. *A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks*. PhD dissertation, Massachusetts Institute of Technology, February 1992.

[7] Dinesh Verma, Hui Zhang, and Domenico Ferrari. Guaranteeing delay jitter bounds in packet switching networks. In *Proceedings of Tricomm'91*, pages 35–46, Chapel Hill, North Carolina, April 1991.

[8] Hui Zhang. *Service Disciplines for Integrated Services Packet-Switching Networks*. PhD dissertation, University of Califoria at Berkeley, November 1993.

[9] Hui Zhang and Domenico Ferrari. Rate-controlled static priority queueing. In *Proceedings of IEEE INFOCOM'93*, San Francisco, California, April 1993.

[10] Hui Zhang and Srinivasan Keshav. Comparison of rate-based service disciplines. In *Proceedings of ACM SIGCOMM'91*, pages 113–122, Zurich, Switzerland, September 1991.