

Real-Time Transport of MPEG Video with a Statistically Guaranteed Loss Ratio in ATM Networks

Seok-Kyu Kweon, *Member, IEEE*, and Kang G. Shin, *Fellow, IEEE*

Abstract—Unlike deterministic real-time communication in which excessive resources may be required for “absolute” performance guarantees, statistical real-time communication seeks to achieve *both* probabilistic performance guarantees and efficient resource sharing. This paper presents a framework for statistical real-time communication in ATM networks, providing delay-guaranteed transport of MPEG-coded video traffic with a statistically-guaranteed cell-loss ratio. Delay-guaranteed communication is achieved with a modified version of Traffic-Controlled Rate-Monotonic Priority Scheduling (TCRM). A set of statistical real-time channels that share similar traffic characteristics are multiplexed into a common *macrochannel*. Those statistical real-time channels which are multiplexed together share the resources of a macrochannel, and individual statistical real-time channels are given timeliness and probabilistic cell-loss guarantees. A macrochannel is serviced by the modified TCRM which improves link utilization and makes channel management simpler. Based on the analysis of an $M/D/1/N$ queueing system, we propose a procedure for determining the transmission capacity of a macrochannel necessary to statistically guarantee a cell-loss ratio bound. Our extensive trace-driven simulation has shown the superiority of the proposed framework to the other approaches. The overall cell-loss ratios for *multihop* statistical real-time channels are shown to be smaller than the predetermined bounds, thus verifying our analytical results.

Index Terms—ATM, B-ISDN, statistical real-time communication, video conferencing, statistical multiplexing, MPEG.

1 INTRODUCTION

IN order to realize its potential as a B-ISDN, an ATM network must be able to transport a wide variety of traffic. Moreover, providing integrated services in high-speed store-and-forward networks like ATM is a difficult task because of the wide range of traffic patterns and Quality-of-Service (QoS) requirements to support. Real-time networking applications such as video and audio conferencing, video-on-demand, and remote medical services in an integrated network pose serious challenges, especially due to their stringent QoS requirements—such as bounded cell-delivery delay and cell-loss ratio—and their highly-bursty traffic.

Unlike best-effort delivery services, real-time communication requires *each* cell delivery delay, not the average delay, as a QoS requirement. If a cell arrives at the receiver after its deadline, its value to the application may be greatly diminished or even worthless. In some circumstances, a cell (actually the segment containing it) that missed its deadline is considered lost. For example, if a cell which contains part of a frame for real-time video conferencing arrives at the receiver later than the frame’s playback time, the receiver will not be able to play the video frame in a complete form

and the late cell will be of no use to the application. Thus, the cell-delivery delay must be bounded and predictable for real-time applications. It is impossible, however, to meet various QoS requirements for different real-time applications using the best-effort delivery service of conventional packet-switched networks, because their packet multiplexers do not differentiate between real-time and nonreal-time traffic, nor among real-time connections themselves. Several packet-multiplexing techniques have been proposed to provide different QoS for different real-time applications [1], [2], [3], [4], [5], [6], [7], [8], [9]. In particular, we proposed Traffic-Controlled Rate-Monotonic Priority Scheduling (TCRM) [9] as a cell-multiplexing scheme to achieve deterministic real-time communication in ATM networks. The prime attraction of the TCRM lies in its simplicity and, hence, makes it suitable for highspeed networks such as ATM and, at the same time, provides high channel admissibility.

Real-time communication can be classified further into two categories according to QoS requirements [10]: *deterministic* and *statistical*. In deterministic real-time communication, QoS requirements are specified in absolute terms and no cell losses or deadline misses are allowed. In order to satisfy its absolute QoS requirements, each deterministic real-time connection must reserve resources based on the worst-case source traffic-generation behavior, thus resulting in severe underutilization of network resources, especially when source traffic is bursty. By contrast, to make more efficient use of network resources, statistical real-time communication specifies QoS requirements in probabilistic (instead of deterministic) terms and tolerates a certain percentage of cell losses and/or deadline misses. Such a

- S.-K. Kweon is with Cisco systems, 170 West Tasman Dr., San Jose, CA 95134-1706. Email: skkweon@cisco.com.
- K.G. Shin is with Real-Time Computing Laboratory, Department of Electrical Engineering and Computer Science, University of Michigan, 1301 Beal Avenue, Ann Arbor, MI 48109-2122. E-mail: kgshin@eecs.umich.edu.

Manuscript received 20 Sept 1999; revised 28 July 2000; accepted 28 Sept. 2000.

For information on obtaining reprints of this article, please send e-mail to: tpd@computer.org, and reference IEEECS Log Number 110668.

specification allows for overbooking network resources and, at the same time, enhancing the multiplexing gain. Statistical real-time communication is useful especially to those applications: 1) That can tolerate a portion of cell losses and deadline misses and 2) Whose traffic is bursty. The statistical multiplexing gain is known to be substantial, especially in Variable-Bit-Rate (VBR) applications such as MPEG-coded video.

While much work has been done to realize deterministic real-time communication over multihop networks [9], [11], [12], much less has been done to provide statistical real-time guarantees over multihop networks. Ferrari and Verma [4] introduced the concept of statistical performance guarantees and proposed a method for statistical real-time communication in WANs using the Earliest-Due-Date (EDD) packet-scheduling policy. They specified a channel-establishment condition for statistical performance guarantees. However, since it requires the calculation of the probability of deadline overflows that depends on other channels' activities, their channel-establishment process is quite complex. Kurose [13] proposed to model a traffic source with a family of random variables that bound the number of packets sent during various time intervals. In his approach, the delay of each packet in the network can be upper-bounded, but the bound may be quite loose. Zhang and Knightly applied Kurose's input traffic model to Rate-Controlled Static-Priority (RCSP) and provided statistical performance guarantees for individual real-time connections in general packet-switched networks [14]. They proposed to use traffic shaping at the intermediate nodes in order to prevent the burstiness of VBR traffic as packets travel through the network. However, each packet is required to carry its timing information for traffic shaping. Since an ATM cell cannot carry any timing information, it is impossible to implement such a traffic shaper and, thus, this scheme cannot be employed in ATM networks. Without traffic shaping, the increased burstiness will lower network utilization. Moreover, it is not clear how to choose appropriate bounding random variables of real applications, like MPEG-coded video, which have great effects on network utilization. Zhang et al. [15] also derived a statistical bound on the end-to-end delay by applying the Exponentially-Bounded Burstiness (EBB) process model [16] to Generalized Processor Sharing (GPS) networks in a similar way as Parekh and Gallager [17], [18] derived a deterministic bound using the leaky bucket model [19], [20]. Although their work is theoretically attractive, it is not clear whether it can be applied to real life systems since their scheme assumes an infinite buffer at each node. Moreover, the implementation complexity of PGPS must be resolved before it can be used for highspeed networks like ATM [9]. Reisslein et al. [21], [22] also showed that statistical delay guarantee for multiplexed real-time streams can be achieved in multihop networks by employing cascaded leaky-bucket regulators, which adds to the complexity of switches.

Effective bandwidth has been investigated in order to provide statistically-guaranteed QoS in ATM networks [23], [24], [25]. This approach is based on the large deviation

theory and often employs an on-off process as a source traffic model. In particular, Elwalid et al. [25] derived the worst-case traffic parameters to achieve lossless multiplexing and used them to extract multiplexing gains from the statistical independence of traffic processes subject to the constraint of a small buffer-overflow probability. They employed the leaky-bucket-regulated periodic on-off process as their input traffic model to this end. In order to calculate the overflow probability in the buffered multiplexing system like an ATM multiplexer, they developed a virtual buffer/trunk system. This model enabled them to transform the two-resource (buffer and link bandwidths) reservation problem into a single-resource reservation problem. By using this model, they were able to use the Chernoff bound as a buffer-overflow probability estimate. Although their approach is mathematically elegant, the estimate based on the extremal on-off process is quite pessimistic as we shall see in Section 4.

In this paper, we propose a framework to provide statistical real-time communication services for MPEG videos in ATM networks equipped with the TCRM [9]. In addition to its simplicity, the TCRM achieves a good channel accommodability. The TCRM, however, was intended for providing deterministic real-time communication services and does not allow statistical multiplexing among real-time connections. We make a slight modification to the TCRM so that it may allow statistical multiplexing among a set of real-time connections. The modified TCRM retains the property of providing a CBR (Constant-Bit-Rate) pipe to each individual virtual channel. So, every cell of each virtual channel is guaranteed to be delivered within a certain bound as long as it is not lost due to buffer overrun. By employing the histogram-based model [26] as the input traffic specification for video traffic data along with the modified TCRM, we analytically derive a statistical bound for the average cell-loss ratio of each statistical real-time channel. The TCRM's ability to provide CBR pipes is crucial to our analysis. Simulation results are shown to support our analysis.

Our approach differs in several aspects from the effective bandwidth approach [23], [24], [25] in providing statistical real-time communication services. First, our approach can provide a framework that can control the capacity of a trunk over which statistical real-time channels are multiplexed using the TCRM. Therefore, it can be used not only for a large ATM network but also for a small system that multiplexes only a dozen of real-time channels. Second, as compared to the effective bandwidth approach, ours provides much tighter cell-loss estimates that can be used for channel-admission control. Finally, we can reduce the complexity of channel admission control by adjusting the number of bins in the histogram while the Chernoff bound approach requires to solve nonlinear equations in calculating cell-loss estimates.

The remainder of this paper is organized as follows: Section 2 defines a real-time connection with statistical performance guarantees (i.e., a statistical real-time channel) and reviews the characteristics of the TCRM. Section 3 describes the MPEG video source model and analytically derives the cell-loss ratios of a set of real-time channels for

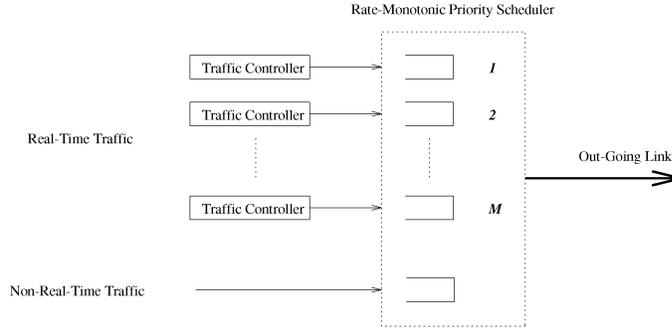


Fig. 1. Structure of TCRM.

both single-hop and multihop cases. Section 4 presents a simulation study with real MPEG video data and compares our approach with the effective bandwidth approach. The paper concludes with Section 5.

2 BACKGROUND

Providing statistical real-time communication service requires source-traffic modeling, resource reservation, and an appropriate scheme for cell-multiplexing and buffer-management. This section discusses the cell-scheduling scheme to be used in our approach.

We first define a *statistical real-time channel* as a unidirectional virtual circuit that guarantees the probability of losing a cell of this channel to be less than a given number Z :

$$Pr(\text{end-to-end cell loss}) \leq Z. \quad (2.1)$$

Although a statistical real-time channel can also be defined in terms of delay as was done in [27], we consider only cell losses due to buffer overrun, because the modified TCRM used in our scheme can guarantee the delivery deadline of every cell as long as it is not lost due to buffer overrun. (More on this will be discussed in Section 3.4.2.)

2.1 TCRM

The TCRM [9] is a cell-scheduling scheme for output-queueing ATM switches [28] in order to provide a guaranteed throughput to individual deterministic real-time channels sharing a common outgoing link. It emulates circuit switching during a period longer than the cell interarrival time of each real-time channel. Fig. 1 shows the structure of the TCRM consisting of a set of *traffic controllers* and a *rate-monotonic priority scheduler*. A traffic controller is assigned to each individual real-time channel and the rate-monotonic priority scheduler is shared by all the real-time channels running over the link.

2.1.1 Traffic Controller

Consider a real-time channel i whose throughput requirement is ρ_i (in bits/sec). For real-time channel i , the function of the traffic controller is to keep the cell-arrival rate at the scheduler below ρ_i . It holds the incoming cells until their supposed arrival times known as *logical arrival times* and then moves them to the scheduler. The logical arrival time of an incoming cell is calculated based on that of the previous cell of real-time channel i . Thus, the

logical arrival time of the k^{th} cell of real-time channel i , $X_i(k)$, is calculated as:

$$X_i(k) = \begin{cases} A_i(k) & k = 1 \\ \max\{X_i(k-1) + \frac{L}{\rho_i}, A_i(k)\} & k \geq 2, \end{cases} \quad (2.2)$$

where $A_i(k)$ is the actual arrival time of the k^{th} cell of real-time channel i , and L denotes the length of an ATM cell, 53 bytes. The traffic controller ensures that the inter-logical arrival time of the incoming cells at the scheduler is at least L/ρ_i . Assuming that the cell-arrival rate at the traffic controller from the previous link is lower than ρ_i/L , we can ensure that there will be *at most* one cell of real-time channel i in its traffic controller, as the traffic controller is allowed to transfer one cell every L/ρ_i seconds. Consequently, the traffic controller requires a buffer to store only one cell.

2.1.2 Nonpreemptive Rate-Monotonic Priority Scheduling

After the cell stream of real-time channel i passes through the traffic controller, the cell-arrival rate at the scheduler of every switch is bounded by ρ_i/L . If the minimum cell-drain rate at the scheduler is ρ_i/L , there will never be an unbounded accumulation of cells at the scheduler. The minimum cell interarrival time and the cell-delay bound at the scheduler will be given by L/ρ_i . As a result, we set the cell-delivery deadline to L/ρ_i .

Liu and Layland [29] proved that rate-monotonic priority scheduling is optimal among all fixed-priority scheduling policies when the deadline of each task is the same as the task period. If we view a cell as a "task" and a link as a "processor," then the rate-monotonic priority cell scheduling is optimal among fixed-priority scheduling policies that achieve the guaranteed throughput, because the cell interarrival time is the same as the cell-delivery deadline ($= L/\rho_i$). Thus, we adopt the rate-monotonic priority scheduler for cell transmissions. This scheduling policy assigns higher priority to real-time channels with higher request rates, i.e., higher ρ_i .

Note that the analysis in [29] is based on a preemptive scheduling policy. However, preemptive scheduling is not desirable for cell transmissions, since if in-progress transmission of a cell is interrupted, the cell will be lost and has to be retransmitted, thus wasting network resources. So, we have to use the nonpreemptive rate-monotonic priority scheduling policy as the cell scheduler. The nonpreemptive

rate monotonic priority scheduler assigns priority to each real-time channel according to its required throughput and an in-progress cell transmission will not be preempted. With an appropriate channel-admission control [9], the scheduler provides the minimum throughput ρ_i to real-time channel i .

Since the cell interarrival time is larger than, or equal to, L/ρ_i and one cell is permitted to be transmitted every L/ρ_i , at most one cell can stay in the scheduler at any time. Thus, the scheduler needs a buffer of only one cell for each real-time channel.

In summary, the rate-monotonic priority scheduler works as follows: In Fig. 1, the scheduler transmits cells according to their priorities, i.e., values of ρ_i . If there are no real-time channel cells to transmit, cells from the nonreal-time traffic queue are transmitted. The cells held at the traffic controllers are not allowed to be transmitted in any case. Although there is no statistical multiplexing among real-time channels under the TCRM, it doesn't necessarily waste bandwidth, because the bandwidth left unused by real-time channels is used to transmit nonreal-time traffic.

2.1.3 Emulating Circuit-Switching at the Cell-Level

One characteristic of the TCRM is that it emulates circuit-switching at the cell level if a cell arrives during every inter-logical arrival time from the source and if every node along the path employs the TCRM as its link scheduler. Let $T_i(t, s)$ be the amount of channel i 's traffic transmitted over a link during a time interval $[t, s)$ for any t, s such that $t \leq s$, then

$$L \lfloor \frac{s-t}{L/\rho_i} \rfloor \leq T_i(t, s) \leq L \lceil \frac{s-t}{L/\rho_i} \rceil. \quad (2.3)$$

The lower bound is derived from the fact that a cell of channel i is guaranteed to be transmitted every L/ρ_i by the scheduler of the current node, and the upper bound comes from the fact that the worst-case cell-arrival rate from the link scheduler of the previous node is ρ_i/L . Thus, the average traffic-service rate of channel i , $R_i(t, s)$, during the interval $[t, s)$ is given as:

$$R_i(t, s) = \rho_i, \quad (2.4)$$

where $s - t = k(L/\rho_i)$ and $k = 1, 2, \dots$. In other words, the throughput of channel i is guaranteed to be ρ_i during any time interval of length L/ρ_i , implying that the TCRM emulates circuit-switching at the cell level.

Since it employs rate-monotonic priority scheduling, the TCRM is simple enough to operate in highspeed networks like ATM compared to other packet-scheduling schemes which require timestamp-based sorting [1], [4], [7], [8] (see [9] for detailed implementation issues), while yielding network utilization close to that of a deadline-based scheduling which is known to be optimal in terms of network utilization [9]. However, the deterministic real-time communication service provided by the TCRM has two disadvantages due to its strict "separation" of real-time channels. The first is its inefficient use of buffer and link bandwidth when only the resource usage by real-time traffic is considered. The second is its complex channel management due to the need to monitor each individual

real-time channel *separately*. By allowing real-time channels to be multiplexed statistically, the proposed new framework for statistical real-time communication utilizes resources more efficiently and, at the same time, provides simpler channel management by monitoring a *set* of statistical real-time channels together.

3 A FRAMEWORK FOR STATISTICAL REAL-TIME COMMUNICATION

Using the guaranteed-throughput service provided by the TCRM, we now build a framework for statistical real-time communication on ATM networks. We first need to describe the model of MPEG video traffic sources.

3.1 The Histogram-Based Model for MPEG Video Traffic Sources

In order to reduce the large amount of multimedia traffic such as video, audio, and graphical data, a number of data compression techniques have been proposed and used. Compression attempts to keep the quality of played-back data at the receiving end constant at the expense of changing the bit rate. Consider an MPEG-coded movie sequence, *Starwars*,¹ in Fig. 2a. The sequence shows extremely high burstiness as *I* (Intracoded), *P* (Predictive), and *B* (Bidirectional) frames alternate. Accurate characterization of these compressed data streams is essential for real-time transport of such data over ATM networks.

Numerous models have previously been proposed for VBR video under various compression schemes [26], [31], [32], [33], [34], [35], [36], [37], [38], [39], [40], [41], [42], [43], [44], [45]. Since the VBR behavior of a video stream strongly depends on the compression technique used, many of these models do not characterize MPEG-coded video which is now widely-accepted as a standard for transmission and storage of video data in many multimedia systems. How to characterize MPEG-coded video streams has been investigated by several researchers [36], [37], [41], [43], [44], [45]. They addressed the modeling of MPEG streams using a model fitting or an analytic approach. Especially, in [43], [44], Krunz et al. characterized a video stream using its frame-size histogram and generated synthetic streams possessing the same characteristics as the original stream. These synthetic streams are then used for a simulation study of bandwidth-allocation and buffer-dimensioning problems. Compared to the analytic approach considered in this paper, their approach doesn't scale well, i.e., has a limitation in dealing with a large number of multiplexed video streams. In [45], Krunz and Tripathi solved the bandwidth-allocation problem for multiplexed video streams using a time-varying traffic envelope. To reduce the bandwidth requirement, they proposed to arrange the starting times of component video streams in such a way that *I*, *P*, and *B* frames evenly coincide with one another. Such an arrangement may be implementable for a single-hop network, but it is impractical for a multihop network since frame boundaries change as the traffic advances

1. The original sequence was generated by Garrett and Vetterli [30].

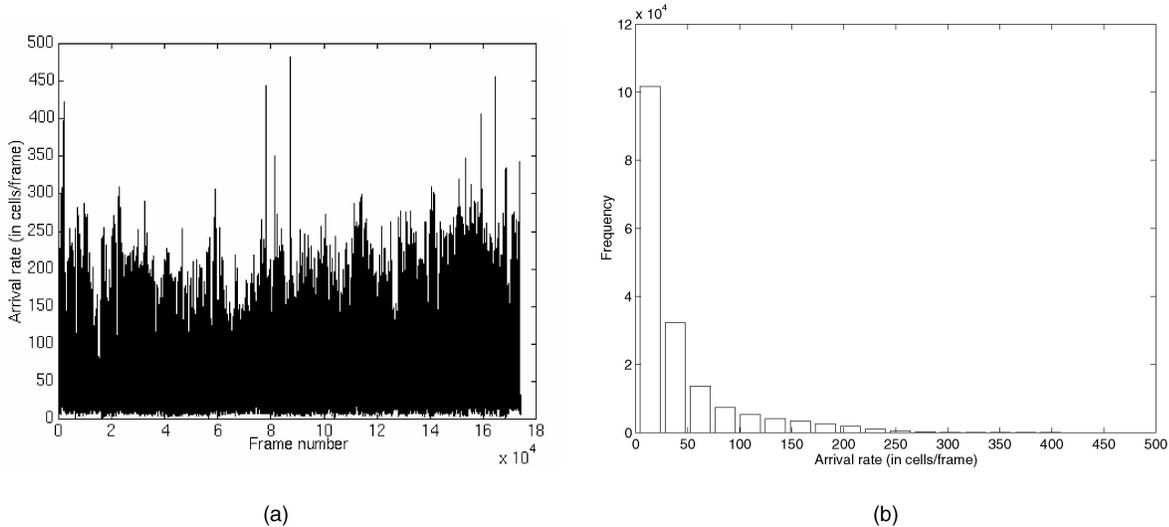


Fig. 2. MPEG-coded video sequence: *Starwars*. (a) Arrival rate sequence. (b) Arrival rate histogram.

inside the network. In contrast, our approach does not require such an arrangement.

Skelly et al. [26] proposed a histogram-based model in order to describe a slowly-varying VBR video traffic source like a motion JPEG video. In their model, the traffic-generation rate from a video source is assumed to be constant during a fairly long period since the bit rate of a motion JPEG video changes very slowly. Using this assumption, Shroff and Schwartz [46] derived an analytic solution to the bandwidth-allocation problem for multiplexed video streams which were deterministically smoothed at their sources.

The bit rate changes more rapidly in an MPEG video than in a motion JPEG video, but this change can still be considered slow compared to a cell's worst-case link delay in an ATM network, if the ratio of buffer size to the output rate is sufficiently small. If video frames are decomposed into ATM cells for transmission and cells are injected into the network after being deterministically smoothed within each frame period, then the cell-arrival rate remains constant in each frame period. For example, if a video buffer can hold 100 cells and the output rate of the buffer is 2,000 cells/frame, the worst-case cell delay at the buffer is one twentieth of a frame period² and, thus, the bit-rate change is slow relative to the worst-case cell link delay. This small ratio assumption is valid in a variety of applications since most real-time video communications require a very small link delay bound.

Under this assumption, we can derive the cell-loss behavior of aggregated video streams using an approach similar to that in [46]. To this end, we extend the histogram-based video model for an MPEG-coded video as follows: Let's assume that: 1) All the video streams have the same frame period, T , and 2) Transmissions of ATM cells are randomly scattered within a frame period, i.e., random smoothing at the source. The cell-arrival rate is measured in

each frame period. We can then think of the arrival process formed by a video stream as a modulated Poisson process whose modulating process is the cell-arrival rate sequence of the entire stream. Since the cell-arrival rate changes frame-by-frame, the modulating process keeps it constant during a certain frame period. The probability mass of a certain cell-arrival rate can be obtained from the histogram of frames' cell-arrival rates. When multiple video streams are multiplexed, the input process of the aggregate traffic can also be modeled as a modulated Poisson process if all the component streams are synchronized frame-by-frame, i.e., frame-transition times of all the component streams are synchronized. The frame-transition time kT where $k = 0, 1, 2, \dots$ is defined as the time at which the transmission of a new frame starts. During a time interval, $(kT, (k+1)T]$, each component stream generates Poisson traffic. Since the superposition of Poisson processes forms another Poisson process, the aggregate traffic becomes a Poisson process during this interval. Considering full-length videos, the aggregate traffic becomes another modulated Poisson process. In this case, the modulating process has the same form as that of a single video stream, and the probability mass function (pmf) of the cell-arrival rate of the aggregate traffic—called the rate pmf—determines the probability masses of cell-arrival rates in the modulated Poisson process. The rate pmf of the aggregate traffic is obtained by taking the convolution of the rate pmfs of all component streams.

The modulated Poisson process model described above may appear unrealistic due mainly to the condition that all the component streams are synchronized frame-by-frame. However, it can be shown that the synchronized traffic-arrival scenario is the worst case of cell losses for multiplexed video streams, which will be discussed in detail in Section 3.3. One can therefore obtain an upper bound of cell-loss ratio for *any* frame synchronization scenario using the modulated Poisson process model.

When a large number of video streams are multiplexed, the assumption of using random smoothing at the source can be relaxed since a large number of similar and

2. Note that this is a *cell* link delay bound. To calculate the end-to-end frame delay bound, we should consider additional delays such as source/destination processing delays and smoothing delay at the source since we assumed source smoothing.

independent sources can be considered as a Poisson process [47]. Thus, as long as the cells of a frame do not arrive in burst as a result of some form of smoothing—whether it is random or deterministic—at the source, we can model the cell arrivals of the aggregate video as a modulated Poisson process.

Fig. 2b shows the histogram of the traffic-generation rate of the sequence in Fig. 2a. The MPEG sequence is *IBBPBBPBBPBBIBB...*. Since *I* frames appear once every 12 frames, the frequency of large frames in the histogram is very low compared to that of small frames.

3.2 Macrochannel

In our approach, a QoS guarantee is made to a *set* of statistical real-time channels, rather than to a single real-time channel as in the deterministic approach. Specifically, we use a statistical real-time channel to transport a video stream. The set of statistical real-time channels are multiplexed onto a common “macro” real-time channel which is guaranteed to receive the minimum throughput provided by the TCRM. A macrochannel is defined as a single-hop real-time channel with parameters (ρ, N) over a link where ρ is the bandwidth (in bits/sec) guaranteed to this channel by the TCRM and N is the size of buffer needed at the traffic controller of this channel. Recall that the buffer space for only one cell is reserved for each real-time channel at its traffic controller in the original TCRM. In this paper, we change the TCRM’s buffer size from 1 to N in order to reduce the cell-loss probability when multiple cells arrive at the macrochannel in a very short time. The cell-drain rate from the buffer is ensured to be ρ/L using the cell logical arrival times. Since the admission control in [48] requires only ρ as a parameter, the change of the buffer size does not require any other modifications to the structure of the TCRM.

Within a macrochannel, we do not differentiate statistical real-time channels from one another. All the cells arriving at this macrochannel are transmitted on a FIFO (First-In-First-Out) basis. This policy simplifies significantly channel management within a macrochannel as compared to the case of treating individual statistical real-time channels separately. Note, however, that the cells of a macrochannel are serviced separately from those of the other macrochannels, deterministic real-time channels, and best-effort traffic. Since all the statistical real-time channels sharing a common macrochannel are treated *equally* on a FIFO basis, all cells in the macrochannel are given the *same* loss probability irrespective of the cell’s channel membership. This implies that individual statistical real-time channels sharing a common macrochannel have the same cell-loss ratio of the macrochannel. That is, the statistical loss guarantee of a macrochannel implies that of each of its component channels, hence allowing us to focus on the macrochannel (or a “bundle” of statistical real-time channels).

Given input traffic specifications of all of its component statistical real-time channels we can derive the parameters (ρ, N) of a macrochannel based on its QoS requirement, or its cell-loss ratio bound Z .

Fig. 3 shows a scenario in which various statistical real-time communication services are provided. Macrochannels

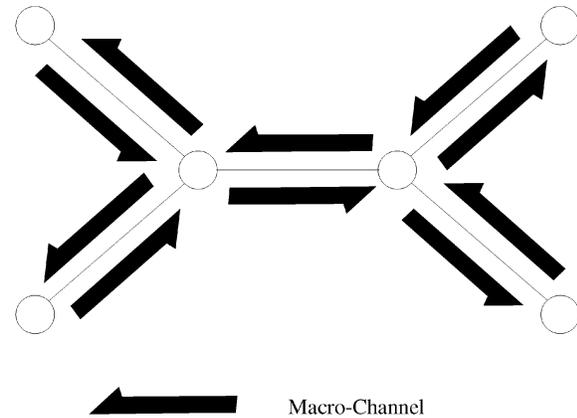


Fig. 3. Macrochannel.

with different parameters (ρ, N) are established in order to provide different QoS guarantees in a single ATM network. Although Fig. 3 shows one macrochannel per link, one can establish an arbitrary number of macrochannels with different cell-loss ratios over a link as long as there are sufficient resources.

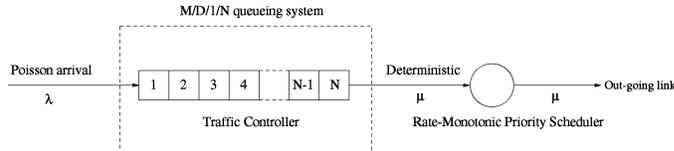
A statistical real-time channel C may run through a (fixed) multihop path between its source and destination. In such a case, since a macrochannel is established over each hop, we need to concatenate a series of “appropriate” macrochannels each of which is selected from the macrochannels established over each link along the path, and multiplex the statistical real-time channel C into them. By an “appropriate” macrochannel, we mean that it must guarantee the cell-loss ratio required for this statistical real-time channel C . We will discuss how to choose macrochannels when we consider admission control later in this section.

Given the above setting, the problem is how to determine the bandwidth ρ and the buffer space N needed to meet the given delay and cell-loss requirements. Since the TCRM bounds the delay over each link when the buffer size N and the minimum throughput ρ are fixed, we will first concentrate on the cell-loss ratio.

3.3 Cell-Loss Ratio within a Macrochannel

We want to derive the cell-loss ratio of a macrochannel using the histogram-based model for aggregate video sources. For now, we will consider only the case in which all statistical real-time channels are established over only a single hop, so that all the cell streams are fed into a macrochannel directly from external sources. We will in Section 3.4 relax this assumption.

In order to determine the cell-loss ratio of a macrochannel, we need the input traffic specification of the aggregate of statistical real-time channels multiplexed over the macrochannel. Since the histogram-based model is chosen for source traffic, we need the rate pmf of the aggregated statistical real-time channels. It can be obtained by taking the convolution of the rate pmfs of component sources, as discussed in Section 3.1. We also assume that all the statistical real-time channels multiplexed are synchronized frame-by-frame.


 Fig. 4. An $M/D/1/N$ queueing system.

According to the MPEG video source model in Section 3.1, an aggregate of multiple video streams has a constant cell-arrival rate during each frame period, and the cell-arrival process during this period forms a Poisson process. During the next frame period, the cell-arrival process forms another Poisson process. Suppose a macrochannel is fed with such a modulated Poisson process. Then, the buffer of the macrochannel can be considered as a queueing system with a modulated Poisson process input. Once a new frame period is started, after some transient period, the queueing system reaches the steady-state during which the arrival process is a Poisson process with a certain rate. During the transient period, the queueing system either serves those cells which arrived during the previous frame period or spends time reaching a steady-state queue level. Assuming the small ratio of buffer to output rate discussed in the previous section, we can consider the transient period negligible compared to the duration of the steady-state and, thus, the cell-loss behavior of the queueing system can be approximated with the steady-state analysis. Under this assumption, the cell-loss behavior of a macrochannel can be analyzed in two steps as follows: First, we fix the cell-arrival rate to a constant, say, λ (cells/frame), and analyze the cell-loss behavior of the system with a Poisson arrival input whose rate is λ . Then, using the rate pmf, we calculate the weighted sum of cell-loss ratios in order to obtain the cell-loss statistics of the overall system.

Assume that the throughput guaranteed to macrochannel j is ρ and that the buffer space reserved at the traffic controller for macrochannel j is N . Then, the buffer of macrochannel j can be seen as an $M/D/1/N$ queueing system since the cell intertransfer time from the traffic controller's buffer to the rate-monotonic priority scheduler is constant i.e., L/ρ . Then, $\mu = \rho/L$ is the service rate (in cells/frame) of the queueing system as illustrated in Fig. 4. Although there can be multiple queueing systems when multiple macrochannels run through a link, we specify only a single macrochannel since different macrochannels are virtually isolated from one another, thanks to the TCRM's *Firewall* property [48].

When the cell-arrival rate is λ in an $M/D/1/N$ system, the cell-blocking probability $P_b(\lambda)$ can be calculated using the following $O(N^2)$ algorithm [49]:

$$\phi'_0 = 1, \quad (3.1)$$

$$\phi'_{k+1} = a_0^{-1} \left(\phi'_k - \sum_{j=1}^k \phi'_j a_{k-j+1} - a_k \right), \quad 0 \leq k \leq N-1, \quad (3.2)$$

$$\phi_{0,n} = \left(\sum_{k=0}^n \phi'_k \right)^{-1}, \quad n = 0, \dots, N \quad (3.3)$$

$$P_{b,n}(\lambda) = 1 - \frac{1}{\phi_{0,n} + \lambda/\mu}, \quad n = 0, \dots, N \quad (3.4)$$

$$P_b(\lambda) = P_{b,N}(\lambda), \quad (3.5)$$

where

$$a_k = \frac{(\lambda/\mu)^k}{k!} e^{-\lambda/\mu}. \quad (3.6)$$

$\phi_{0,n}$ represents the probability that a departing customer finds an $M/D/1/n$ system empty [49]. The computation time of this recursion increases with the buffer size, N . For fast admission control, it is critical to reduce this time. Fortunately, $P_{b,n}(\lambda)$ converges to $P_b(\lambda)$ rapidly as n increases toward N except when λ/μ is close to 1. For example, when $\lambda/\mu = 2$ and $N = 100$, $\phi_{0,0}$, $\phi_{0,1}$, $\phi_{0,2}$, $\phi_{0,3}$, and $\phi_{0,4}$ are, respectively, 0.5317, 0.5062, 0.5012, 0.5003, and 0.5001, showing a rapid convergence of $\phi_{0,n}$ to $\phi_{0,100} = 0.5000$. Intuitively, when $\lambda/\mu \ll 1$, the cell-blocking probability of an $M/D/1/N$ approaches zero rapidly as the buffer size increases. When $\lambda/\mu \gg 1$, $\phi_{0,n}$ approaches zero rapidly as the buffer size increases. So, $P_b(\lambda)$ is given by $1 - \mu/\lambda$. In addition, since ϕ'_k is nonnegative [49], $\phi_{0,n}$ is a nondecreasing sequence for any λ/μ . Therefore, $P_{b,n}(\lambda)$ upper-bounds $P_b(\lambda)$ for $n = 0, \dots, N$. Using these properties, one can obtain an upper bound of $P_b(\lambda)$ which is very close to $P_b(\lambda)$ within a reasonable amount of time except when λ/μ is close to 1.

The cell-loss ratio of the macrochannel, P_{macro} , is given as the weighted sum of cell-loss ratios where weights are given by the rate pmf of the aggregated video sources. Thus,

$$P_{macro} = \frac{\sum_{i=1}^M P_b(\lambda_i) f_i \lambda_i}{\sum_{i=1}^M f_i \lambda_i}, \quad (3.7)$$

where M is the number of intervals (bins) in the histogram and f_i is the probability mass of arrival rate λ_i which is obtained from the histogram of the cell-arrival rate of the aggregate traffic.

Until now, we assumed that all the video streams to be multiplexed are synchronized. We now want to relax this assumption in order to handle a more general traffic-arrival scenario. To this end, we consider the following *semi-synchronized* traffic-arrival scenario. First, we divide a frame period, $(kT, (k+1)T]$ into J equal subframe periods, $(kT + jT/J, kT + (j+1)T/J]$, where $k = 0, 1, 2, \dots$ and $j = 0, 1, 2, \dots, J-1$. Let's assume that any component video stream is allowed to start transmitting a new video

frame only at one of the lattice points, $kT + jT/J, \forall k, j$. Now, consider a single video stream and calculate its rate-histogram. This time, we measure the cell-arrival rate in a subframe period. Since, however, the cell-arrival rate is constant during each frame period, we have J subframe periods with the same cell-arrival rate. Therefore, the rate-histogram calculated for a subframe period has the same shape as that calculated for a whole frame period. Therefore, they have exactly the same rate pmf, although their rate-histograms have different frequency sets (the rate-histogram calculated for a subframe period has J times larger frequencies than that calculated for a whole frame period). When video streams are multiplexed, the rate pmf of the aggregate traffic measured over a subframe period is calculated by taking the convolution of the rate pmfs of component video streams which are measured over a subframe period. Since the rate pmf of each component video stream measured over a subframe period is the same as that of each component video stream measured over a whole frame period, the rate pmf of the aggregate traffic measured over a subframe period is the same as that of the aggregate traffic measured over a whole frame period. The cell-arrival rates of the aggregate traffic measured over a whole frame period and a subframe period are denoted by Λ and Γ , respectively. Then,

$$Pr\{\Lambda > \alpha\} = Pr\{\Gamma > \alpha\}, \quad \forall \alpha \geq 0. \quad (3.8)$$

Since, in the synchronized case, the cell-arrival rate is measured over a whole frame period for deriving the cell-loss ratio, we also consider one frame period for the semisynchronized case for fair comparison. Let Λ_{semi} and Γ_j denote, respectively, the cell-arrival rate measured over a frame period, $(kT, (k+1)T], \forall k$, and the one measured over a subframe period, $(kT + jT/J, kT + (j+1)T/J]$, where $j = 0, 1, 2, \dots, J-1$, in the semisynchronized case. Then,

$$Pr\{\Lambda_{semi} > \alpha\} = Pr\{\Gamma_0 > \alpha, \Gamma_1 > \alpha, \dots, \Gamma_{J-1} > \alpha\}, \quad \forall \alpha \geq 0. \quad (3.9)$$

Since

$$Pr\{\Gamma_0 > \alpha, \Gamma_1 > \alpha, \dots, \Gamma_{J-1} > \alpha\} = Pr\{\Gamma_0 > \alpha\} Pr\{\Gamma_1 > \alpha, \dots, \Gamma_{J-1} > \alpha | \Gamma_0 > \alpha\} \quad (3.10)$$

and

$$Pr\{\Gamma_1 > \alpha, \dots, \Gamma_{J-1} > \alpha | \Gamma_0 > \alpha\} \leq 1, \quad (3.11)$$

$$Pr\{\Lambda_{semi} > \alpha\} \leq Pr\{\Gamma_0 > \alpha\}. \quad (3.12)$$

From (3.8),

$$Pr\{\Lambda_{semi} > \alpha\} \leq Pr\{\Lambda > \alpha\}. \quad (3.13)$$

Rewriting this relation using the probability masses of cell arrival rates,

$$\sum_{i=k}^M g_i \leq \sum_{i=k}^M f_i, \quad \text{for } k = 1, 2, \dots, M, \quad (3.14)$$

where f_i and g_i are the probability masses of cell arrival rate λ_i measured in the synchronized case and in the semisynchronized case, respectively. This inequality implies

$$f_M - g_M \geq 0, \quad (3.15)$$

and

$$f_k - g_k \geq - \sum_{i=k+1}^M (f_i - g_i), \quad \text{for } k = 1, 2, \dots, M-1. \quad (3.16)$$

Now, we calculate the cell-loss ratio of the macrochannel in the semisynchronized case, P_{macro}^{semi} , by replacing f_i by g_i in (3.7). Thus,

$$P_{macro}^{semi} = \frac{\sum_{i=1}^M P_b(\lambda_i) g_i \lambda_i}{\sum_{i=1}^M g_i \lambda_i}. \quad (3.17)$$

Since changing the way of multiplexing video streams does not affect the long-term average cell-arrival rate,

$$\sum_{i=1}^M g_i \lambda_i = \sum_{i=1}^M f_i \lambda_i. \quad (3.18)$$

As a result, in order to compare P_{macro}^{semi} and P_{macro} , we only need to compare $\sum_{i=1}^M P_b(\lambda_i) \lambda_i g_i$ and $\sum_{i=1}^M P_b(\lambda_i) \lambda_i f_i$. Now, let's assume that $\lambda_i < \lambda_j$ if $i < j$. Then, $\{\lambda_i\}$ is a positive monotonically-increasing sequence. In addition, since the cell loss of an $M/D/1/N$ system does not decrease with arrival rate λ_i , $\{P_b(\lambda_i)\}$ is a nonnegative nondecreasing sequence. Thus, $\{\lambda_i P_b(\lambda_i)\}$ is a nonnegative nondecreasing sequence. Let β_i denote $\lambda_i P_b(\lambda_i)$, then

$$\sum_{i=1}^M P_b(\lambda_i) \lambda_i f_i - \sum_{i=1}^M P_b(\lambda_i) \lambda_i g_i = \sum_{i=1}^M \beta_i (f_i - g_i). \quad (3.19)$$

By separating the first term,

$$\sum_{i=1}^M \beta_i (f_i - g_i) = \beta_1 (f_1 - g_1) + \sum_{i=2}^M \beta_i (f_i - g_i). \quad (3.20)$$

From (3.16)

$$\begin{aligned} \beta_1 (f_1 - g_1) + \sum_{i=2}^M \beta_i (f_i - g_i) &\geq \\ &- \beta_1 \sum_{i=2}^M (f_i - g_i) + \sum_{i=2}^M \beta_i (f_i - g_i). \end{aligned} \quad (3.21)$$

Arranging terms, we obtain

$$\beta_1 (f_1 - g_1) + \sum_{i=2}^M \beta_i (f_i - g_i) \geq \sum_{i=2}^M (\beta_i - \beta_1) (f_i - g_i). \quad (3.22)$$

Now, let's write the first term of the r.h.s. of (3.22) separately as:

$$\begin{aligned} \sum_{i=2}^M (\beta_i - \beta_1) (f_i - g_i) &= \\ (\beta_2 - \beta_1) (f_2 - g_2) + \sum_{i=3}^M (\beta_i - \beta_1) (f_i - g_i). \end{aligned} \quad (3.23)$$

Since β_i is a nonnegative nondecreasing sequence,

$$\beta_2 - \beta_1 > 0. \quad (3.24)$$

Therefore, from (3.16),

$$(\beta_2 - \beta_1)(f_2 - g_2) \geq -(\beta_2 - \beta_1) \sum_{i=3}^M (f_i - g_i). \quad (3.25)$$

Then, (3.23) becomes

$$\begin{aligned} \sum_{i=2}^M (\beta_i - \beta_1)(f_i - g_i) &\geq -(\beta_2 - \beta_1) \sum_{i=3}^M (f_i - g_i) \\ &+ \sum_{i=3}^M (\beta_i - \beta_1)(f_i - g_i). \end{aligned} \quad (3.26)$$

Arranging terms, we obtain

$$\sum_{i=2}^M (\beta_i - \beta_1)(f_i - g_i) \geq \sum_{i=3}^M (\beta_i - \beta_2)(f_i - g_i). \quad (3.27)$$

From (3.22) and (3.27),

$$\sum_{i=1}^M \beta_i (f_i - g_i) \geq \sum_{i=3}^M (\beta_i - \beta_2)(f_i - g_i). \quad (3.28)$$

Continuing with similar operations, we obtain the following inequality:

$$\sum_{i=1}^M \beta_i (f_i - g_i) \geq (\beta_M - \beta_{M-1})(f_M - g_M). \quad (3.29)$$

Finally, since $\beta_M \geq \beta_{M-1}$, from (3.15),

$$\sum_{i=1}^M \beta_i (f_i - g_i) \geq 0. \quad (3.30)$$

Therefore, we can conclude that

$$P_{macro}^{semi} \leq P_{macro}. \quad (3.31)$$

By increasing J towards ∞ , we can extend the semisynchronized case into the scenario that frames of component video streams start at arbitrary time points. So, the synchronized case represents the worst-case cell-loss behavior among all traffic-arrival scenarios for multiplexed video streams.

3.4 Cell Losses in an End-to-End Connection

In Section 3.3, we considered only the single-hop case in which the Poisson-arrival approximation holds. However, in general point-to-point networks, cell streams take multiple hops before arriving at their destination. In our framework, a statistical real-time channel is multiplexed over a *chain* of macrochannels on the links along the channel path. We first describe our assumptions on the traffic switched and routed via multiple hops and then derive the cell-loss ratio bound for the end-to-end connection.

3.4.1 Effects of Switching

As cells are switched and routed from one macrochannel to another, the traffic pattern may change depending on the traffic condition at each macrochannel. This raises two questions on our assumptions about the traffic from aggregate sources in a single hop. One is the accuracy of

the Poisson-arrival assumption on the traffic from aggregate sources. The other is the derivation of the new cell-arrival rate histograms at intermediate nodes, because the histogram defined at the source may change depending on the conditions of the intermediate nodes. That is, if other statistical real-time channels sharing the same macrochannel at the upstream nodes have large amounts of traffic, a statistical real-time channel may lose a large portion of its cells at those nodes and the rate pmf at the downstream links may change.

For the time being, let's assume that the histogram defined at the source node remains unchanged at all intermediate nodes. In general, the output process of an $M/D/1/N$ queue is not a Poisson process and cell interdeparture times are correlated [50]. This poses difficulty in analyzing a multihop statistical real-time channel. This is also the case in an $M/M/1$ system analysis. In the $M/M/1$ system, the packet interarrival and service times are correlated. To handle this difficulty, Kleinrock proposed to use "Independence Approximation" in analyzing a communication network using a general queueing network like the Jackson network [47], [51]. It asserts that, in an $M/M/1$ system, merging several cell streams on a transmission link has an effect akin to restoring the independence of interarrival times and service times. In particular, he emphasized the independence of service times of a packet at different nodes, which is not true in real communication networks. Since the length of a cell is fixed in ATM networks, the correlation of cell-service time is not important in our problem. What matters in the $M/D/1/N$ analysis is the independence of cell interarrival times. As with Kleinrock's independence approximation, we assume that the cell interarrival time at a macrochannel at an intermediate link is exponentially-distributed if multiple cell streams routed from different macrochannels on different links and externally-fed cell streams are merged into this macrochannel. Then, the new aggregate traffic arriving at this new macrochannel can be approximated as a Poisson process, thus enabling the application of the $M/D/1/N$ analysis result at any macrochannel as long as the number of multiplexed statistical real-time channels is large enough and the routing processes are uncorrelated. Our simulation study in Section 4 confirms the validity of this assumption.

Next, let's consider the rate-histogram of a video stream switched and routed inside the network. As the traffic traverses downstream nodes, the original traffic pattern at the source node will change and may, in general, become burstier. However, in the rate-histogram model, we assumed that the cell-arrival process is a Poisson with a certain rate, say, λ , during a single frame period. During the same period, this cell stream is multiplexed with streams of other statistical real-time channels onto a macrochannel, M_a . After departing from macrochannel M_a , the cell stream is separated from the other statistical real-time channels and then multiplexed onto a new macrochannel, M_b . While being multiplexed at M_a , some cells of this stream may be lost due to buffer overrun. Therefore, the number of cells of the stream at M_b cannot be larger than that at M_a . Over the frame period considered, the arrival rate of this stream at

M_b , denoted by $I(\lambda)$, cannot be larger than that at M_a , λ . That is,

$$I(\lambda) \leq \lambda. \quad (3.32)$$

Now, let's consider the entire stream which was modeled as a modulated Poisson process in Section 3.1. Let the rate pmf of the process be given by $\{\lambda_i, f_i\}_{i=1, \dots, M}$, where M is the number of bins and f_i is the probability mass of arrival rate λ_i . Let Λ_a and Λ_b denote the arrival rates of the stream at M_a and M_b , respectively. Then,

$$\begin{aligned} Pr\{\Lambda_b \geq I(\lambda_k)\} &= \sum_{i=k}^M I(\lambda_i) \cdot f_i \\ &\leq \sum_{i=k}^M \lambda_i \cdot f_i \\ &= Pr\{\Lambda_a \geq \lambda_k\}. \end{aligned} \quad (3.33)$$

Since $I(\lambda_k) \leq \lambda_k$,

$$Pr\{\Lambda_b \geq \lambda_k\} \leq Pr\{\Lambda_b \geq I(\lambda_k)\}. \quad (3.34)$$

Thus,

$$Pr\{\Lambda_b \geq \lambda_k\} \leq Pr\{\Lambda_a \geq \lambda_k\}. \quad (3.35)$$

This relation shows that the rate pmf of a video stream at intermediate nodes is *probabilistically bounded* from above by the rate-histogram at the source node. That is,

$$\begin{aligned} Pr(\text{cell-arrival rate at the source node} \geq x) &\geq \\ Pr(\text{cell-arrival rate at the intermediate nodes} \geq x). \end{aligned} \quad (3.36)$$

In terms of QoS guarantees, it is still effective to use the traffic characteristics calculated at the source nodes in order to calculate the convolution of the rate pmfs of component video streams at intermediate nodes since the cell-loss probability can still be bounded by using the same traffic characteristics. It allows for simple run-time channel-establishment at the expense of slightly conservative resource reservation. The amount of over-reservation of resources at intermediate nodes is negligible when the cell-loss probability is small, which is the case of most statistical real-time applications, as will be discussed in Section 4.

3.4.2 Cell-Loss Ratio Bound in an End-to-End Connection

Based on the above arguments, the end-to-end cell-loss probability of a statistical real-time channel is given by

$$Pr(\text{end-to-end cell loss}) \leq 1 - \prod_{j=1}^K (1 - P_{macro,j}), \quad (3.37)$$

where K is the number of hops the statistical real-time channel takes and $P_{macro,j}$ is the cell-loss probability of the macrochannel at the j^{th} hop. Notice that although $P_{macro,j}$ is the cell-loss ratio of the macrochannel at the j^{th} hop, it is also the cell-loss ratio of individual statistical real-time channels multiplexed onto the macrochannel.

Although we focused on deriving a cell-loss ratio bound, it must be stressed that our approach also guarantees

statistically each real-time cell's delivery delay. That is, the probability that a cell is delivered to its destination before its deadline is larger than $\prod_{j=1}^K (1 - P_{macro,j})$. This is because a cell which has "survived" buffer overruns on its way to the receiver is always guaranteed to be delivered within a bounded time because buffer size is fixed and the minimum buffer drain rate is guaranteed at each link by the TCRM. The end-to-end delay bound of the statistical real-time channel is given as:

$$D_{end-to-end} = \sum_{j=1}^K (N_j + 1)L/\rho_j, \quad (3.38)$$

where N_j and ρ_j are the buffer size (measured in number of cells) and the bandwidth of the macrochannel at the j^{th} hop, respectively, and N_jL is the maximum backlog at the macrochannel upon arrival of a cell. The reason for adding 1 to the buffer size is to account for the delay at the rate-monotonic priority scheduler, as can be seen in Fig. 4.

3.5 Admission Control for Channel Establishment Requests

When the establishment of a statistical real-time channel is requested, the network service provider must execute channel-admission control in order to guarantee the QoS promised to a new channel as well as existing real-time channels. One approach to channel admission control is to use a set of preestablished macrochannels. Each local link has its own set of preestablished macrochannels. Each macrochannel's buffer size and bandwidth are fixed, and its cell-loss ratio bound is also fixed. When a new channel request arrives, the network service provider selects a macrochannel for each local link from the preestablished macrochannels such that the end-to-end delay and cell-loss bound given by (3.38) and (3.37), respectively, are smaller than the user-requested bounds. Then, at each local link, the rate pmf of a new aggregate stream consisting of existing channels multiplexed onto the chosen macrochannel and the requested channel is derived using convolution. Using (3.7), one can calculate the maximum cell-loss ratio of the aggregate stream. If the maximum cell-loss ratio is less than, or equal to, the prespecified cell loss ratio bound of the macrochannel, the requested channel is accepted. Otherwise, the request is denied. For a multihop statistical real-time channel, such an admission test must be executed at every node along the path.

3.6 Extension of Macrochannel to Virtual Path

Thus far, we have described the macrochannel as a single-hop real-time channel. The macrochannel can be easily extended to the concept of Virtual Path (VP) by establishing a macrochannel not as a single-hop real-time channel, but as a VP which takes multiple hops. As long as each VP is guaranteed by the TCRM to have a constant throughput at intermediate links, it can be considered as a CBR pipe. That is, a VP can be considered as a single-hop real-time channel even if it passes through multiple hops. The only change in the VP extension is the end-to-end delay bound given by (3.38), which must be modified to:

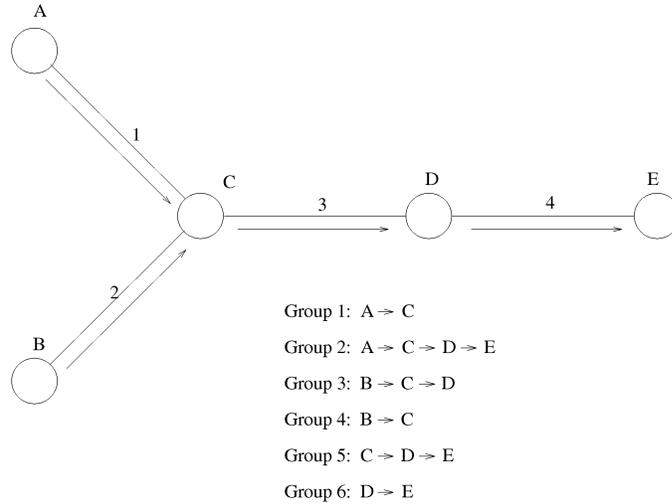
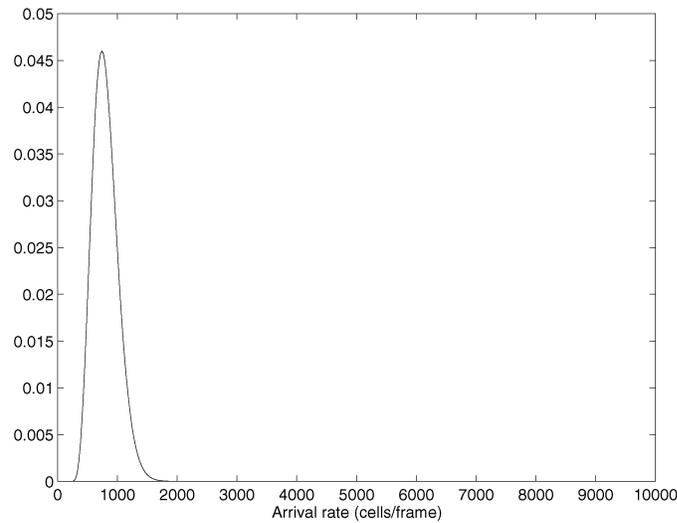


Fig. 5. The network topology for simulation.

Fig. 6. Probability mass function of arrival rate of an aggregate of 20 statistical real-time channels of *Starwars*.

$$D_{end-to-end} = \sum_{i=1}^K (N_j + l_j) L / \rho_j, \quad (3.39)$$

where l_j is the number of hops the j^{th} VP takes. The l_j is added to represent the delay at the rate-monotonic priority scheduler at each hop. Except for this additional delay, a statistical real-time channel using VPs is exactly the same as the single-hop version.

4 SIMULATION AND DISCUSSION

In order to demonstrate the usefulness of the histogram-based model for statistical real-time communication, we have conducted an extensive simulation study using MPEG-coded video traces. Since every cell which has “survived” buffer overruns is delivered in time by the TCRM, we will consider only the cell-loss ratio as the QoS parameter.

4.1 Simulation Model

Fig. 5 shows the topology of an ATM network used for the simulation study which consists of 5 nodes and 4 links. All

the links are simplex and, thus, cells are transmitted only in the direction of arrows shown in the figure. Also, for the sake of simplicity, we assume that there exists only one macrochannel over each link. That is, there is no deterministic real-time traffic, and other statistical real-time traffic or nonreal-time traffic except for the statistical real-time traffic being multiplexed over the macrochannel on each link. Since the TCRM provides a virtual circuit with a guaranteed throughput over an ATM link, a macrochannel can be considered as a CBR pipe with throughput ρ and the input buffer of size N .

In this network, we multiplexed 20 statistical real-time channels on each link. The starting frames of each statistical real-time channel are randomly selected from the clips of movie *Starwars* in Fig. 2, and 17 different MPEG-coded video clips.³ The length of each stream is 1,000 frames, and each run lasts about 50 seconds since we set one frame interval to 1/24 second. First, we conducted an experiment using streams only from the *Starwars* sequence in order to investigate cell-loss ratios in a homogeneous-traffic

3. These sequences were generated by Rose [52].

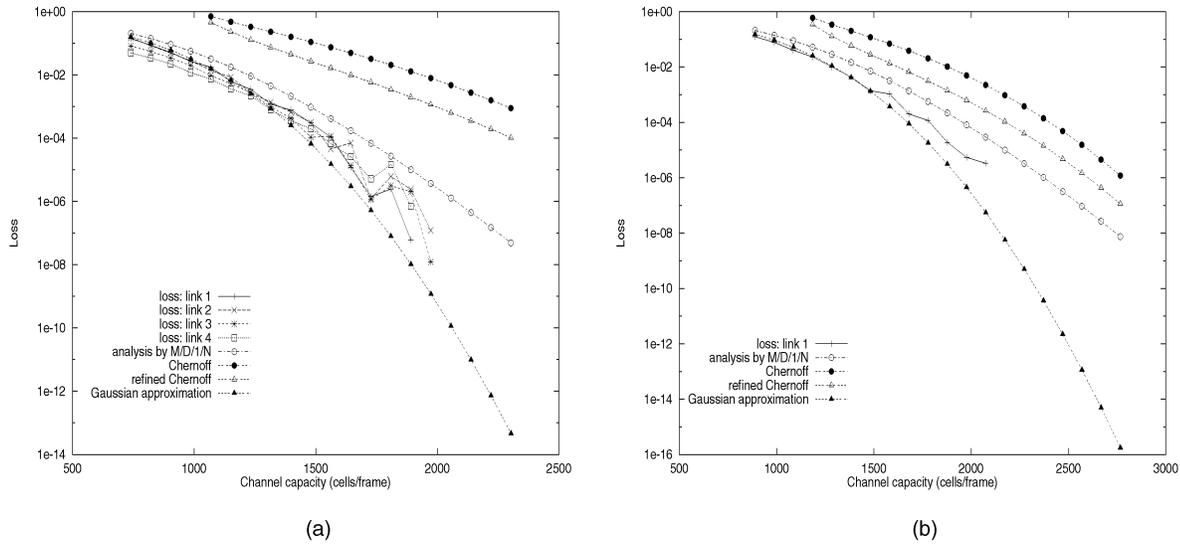


Fig. 7. Cell-loss in an external input case. (a) Homogeneous traffic. (b) Heterogeneous traffic.

environment. Using convolution, we derived the pmf of the arrival rate of the aggregate of 20 streams in Fig. 6. We derived a 20-bin histogram from the sequence, which requires simple operations for the convolution. We tried different numbers of bins such as 10 and 50. When we used a 10 bin histogram, the predicted cell loss ratio was somewhat inaccurate. But when we used a 50-bin histogram which results in a long convolution calculation time, the predicted cell loss ratios were almost similar to those obtained using a 20-bin histogram. The average cell-generation rate of the aggregate traffic is about 822 cells/frame and the maximum cell-generation rate is about 9,000 cells/frame.

Next, in order to investigate the heterogeneous-traffic case, we have conducted a similar experiment using 17 different sequences. We selected as many streams as needed for the simulation from these sequences. In particular, we chose 14 streams once and the other streams twice in order to feed 20 channels which are multiplexed over link 1. In this case, the average cell-arrival rate of the aggregate of 20 streams was 988 cells/frame and the maximum was about 12,000 cells/frame.

To investigate the various cases, the multiplexed streams are grouped according to their paths: six groups are shown in Fig. 5. For example, group 2 consists of channels whose sources (destinations) are node A (node E), and that pass through node C and D. Only the channels of group 1 and 2 traverse link 1. As a result, through link 1, no routed cells are transmitted, but only external input traffic from node A is transmitted. On link 3, group 2 and 3 are routed from link 1 and 2, respectively, and group 5 is directly fed from node C.

During the simulation, the cell transmission from each source is randomly distributed over one frame duration, and all the cells belonging to a frame must be transmitted from the source within one frame duration. At intermediate nodes, cells are transmitted on a FIFO basis regardless of their channel identities.

4.2 Simulation Results

In order to investigate the validity of our assumption on the Poisson arrival process at intermediate nodes, we have considered a case in which some links, in addition to the routed traffic from upstream links, are fed with external inputs. We have assigned 13 channels to group 1, 6 to group 2, 7 to group 3, 13 to group 4, 6 to group 5, and 7 to group 6 so that 20 streams traverse each link. Note that link 1 and link 2 are not fed with any routed traffic.

First, we consider a homogeneous traffic environment in which we multiplex only streams from the Starwars sequence. We have varied the bandwidth assigned to a macrochannel established over each link from 700 cells/frame to 2,300 cells/frame. The buffer size N is 50 (cells) and, thus, the worst-case cell delay in a single hop is $1/20$ frame period, i.e., 2.1 msec if the throughput guaranteed to a macrochannel is 1,000 cells/frame. This is small enough to satisfy the steady-state condition presented in Section 2. The cell-loss ratios are compared to the analysis of an $M/D/1/N$ system in Fig. 7a. We only show the average cell-loss ratios because the loss guarantees provided to a macrochannel and individual statistical real-time channels are the same. When utilizations of macrochannels are low, the cell-loss ratios of all links for the two systems do not show any notable difference. All the cases simulated show that the cell-loss ratios are bounded by the $M/D/1/N$ result.⁴ On the other hand, when utilizations of macrochannels are high, the cell-loss ratios of link 3 and link 4 are smaller than the bounds while those of link 1 and link 2 match the bound almost exactly. As we mentioned in Section 3.4, the tail distributions of the aggregate traffic at link 3 and link 4 decrease because of the cell losses at the upstream links, despite the fact that the decrease is negligible when the cell-loss probability is small. This explains the smaller cell-loss ratios at link 3 and link 4 which have the routed cell streams from links 1 and 2, when the cell-loss probability is large. From this observation, we conclude that if we use a

4. In all the cases in this simulation, 99 percent confidence-level intervals are 10^{-4} , so any value below 10^{-4} is considered to be noisy.

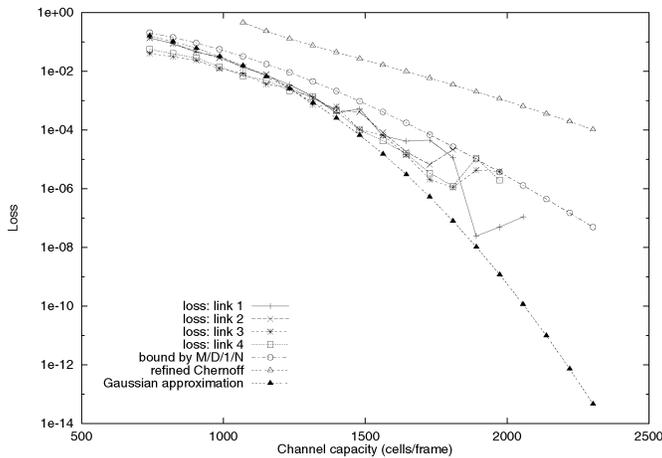


Fig. 8. Cell-loss ratio in no external input case—homogeneous traffic.

macrochannel with a high cell-loss ratio bound, our scheme will result in overreservation of network resources. However, for a macrochannel with a very small cell-loss ratio bound (e.g., 10^{-4}), our scheme provides accurate cell-loss estimates and, thus, enables efficient use of network resources

In Fig. 7a, we also show the Chernoff bound estimate of the cell-loss ratio which is calculated using the derived periodic on-off random process suggested by Elwalid et al. [25]. This approach was chosen since it allows us to analyze the cell-loss behavior of a buffered multiplexing system. Although the Chernoff bound estimate can be derived using more detailed information, e.g., the rate histogram as used in our approach, it only considers an unbuffered multiplexing system. We also show the approximation by a Gaussian distribution, which is based on Central Limit Theorem (CLT) [23]. Both approaches employed the buffer-overflow probability as a QoS parameter while ours uses the cell-loss ratio. For the purpose of comparison, we derived the cell-loss estimates from the buffer-overflow probability obtained by both methods [53]. The parameters of the on-off process derived from the *Starwars* sequence are as follows. The peak rate is 230 cells/frame,⁵ the mean rate is 41 cells/frame, and the bucket size of the leaky-bucket regulator is 462,858 cells. The on and off periods derived from the parameters are 2,450 frame intervals and 11,256 frame intervals, respectively. By substituting these parameters into the Chernoff bound estimate according to the step suggested in [25], we plotted the result in Fig. 7a. In addition to the Chernoff bound, Fig. 7a shows a more accurate refined Chernoff bound by Bahadur and Rao [54]. Compared to our analysis result based on the the $M/D/1/N$ system, both the Chernoff bound and the refined Chernoff bound estimates are too pessimistic. Considering the fact that Elwalid's approach is based on the extremal traffic description, one can anticipate the pessimistic result in Fig. 7a. In contrast, the $M/D/1/N$ analysis based on the rate histogram provides a very accurate cell-loss estimate

5. Originally, the peak rate for achieving lossless multiplexing was 483 cells/frame, but it resulted in too pessimistic a cell-loss ratio estimate. So, we choose the 99.9 percentile from the cell-arrival rate histogram as a peak rate.

with only a 20-bin histogram for which it is not difficult to compute convolutions. Specifically, when the cell-loss ratio bound is set to 10^{-4} , our scheme requires reservation of 1,712 cells/frame while Bahadur and Rao's approach requires reservation of 2,650 cells/frame.

In the CLT approximation, buffer size is ignored and only bandwidth is considered as a reservable resource. Ignoring buffer size may result in pessimistic cell-loss estimates. However, as argued in [23], the CLT approximation is shown to be too optimistic in estimating cell losses for very bursty traffic like MPEG since it tends to ignore the long tail of the rate distribution of a bursty source. By contrast, the $M/D/1/N$ analysis provides a reasonable cell-loss ratio bound that lies between the Chernoff bound estimate and the Gaussian approximation.

We have conducted the same experiment using 17 different video clips in order to study the validity of our model in a heterogeneous-traffic environment. We followed the same procedure as before and plotted the result in Fig. 7b. The only difference is choice of the peak rate of the on-off process. Instead of 99.9 percentile, we used the average cell-generation rate of I frames as a peak rate in order to favor Elwalid et al.'s approach, but it is not justifiable in a strict sense since the original peak cell-generation rate is necessary to obtain the parameters for a lossless multiplexing system [25]. Fig. 7b shows the simulation result on link 1 only, since each link has a different trunk capacity depending on the characteristics of the aggregate traffic. However, we obtained similar results on the other links. In the figure, the $M/D/1/N$ analysis provides a good estimate of cell-loss ratios as in the homogeneous-traffic case, as compared to the Chernoff bound estimate and the CLT approximation. The Chernoff bound estimate is a little closer to the simulation result than the homogeneous case due to the choice of a smaller peak rate.

Next, we considered the case in which there exists only routed traffic without any external input traffic at intermediate nodes: we disabled group 5 in Fig. 5 and changed the number of channels in each group accordingly. We assign 10 channels to each of groups 1, 2, 3, 4, and 6. Note that the number of channels multiplexed over each macrochannel on each link is kept at 20. In this case, there is no external input traffic at the macrochannel on link 3. In Fig. 8, we only show the homogeneous-traffic case using the *Starwars* sequence. The loss at the macrochannel on link 3 does not make any difference from that on links 1, 2 and 4 when the cell-loss probabilities are small. When the cell-loss probability is large (i.e., the reserved bandwidth of the macrochannel approaches the average cell-generation rate of the aggregate channel), the cell-loss ratio of the macrochannel on link 3 is smaller than others. However, the trend is clear that the cell-loss probability is bounded by the analysis result and that the difference between the simulation and analysis results is small when the cell-loss ratio is small. Thus, the Poisson-arrival assumption can be applied even when there is no external input traffic at the intermediate nodes.

Next, we considered the validity of our framework when the number of statistical channels multiplexed over

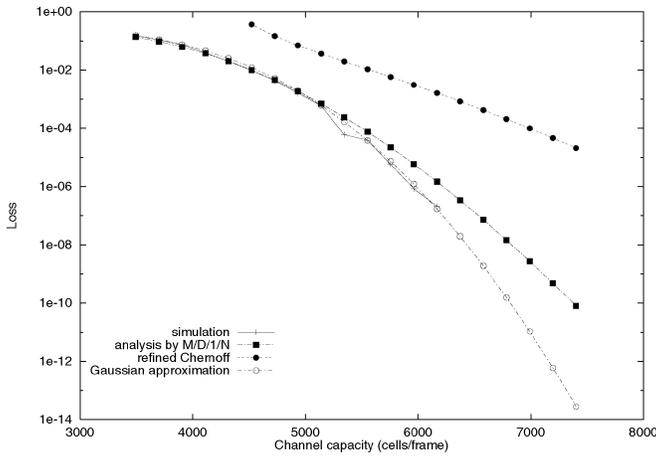


Fig. 9. Cell-loss ratio of 100 channels.

a macrochannel is quite large. We have set the number of channels to 100 and the buffer size N to 100 cells. Fig. 9 shows the simulation result along with the $M/D/1/N$ analysis, the Chernoff bound approximation and the CLT approximation. The Chernoff bound estimates remain too pessimistic, again due to the extremal traffic description. The $M/D/1/N$ analysis provides a good cell-loss ratio bound. Interestingly, the CLT approximation shows a pretty accurate cell-loss estimate in this figure unlike the case when 20 channels are multiplexed. This result shows that the cell-loss behavior resulting from the ignored longer tail of the rate distribution is mitigated as the number of channels multiplexed gets very large. Although the CLT approximation provides a good cell-loss estimate for the case of a large number of channels, our approach provides good cell-loss bounds regardless of the number of channels multiplexed.

The statistical multiplexing gain achieved by increasing the number of channels multiplexed is shown in Fig. 10, in which the cell-loss ratios are plotted against link utilization when five, 10 and 20 channels are multiplexed. The link utilization is normalized against the average cell-arrival rate of the aggregate sources. We show only the $M/D/1/N$ bound. One can see that the loss ratios are bounded for all

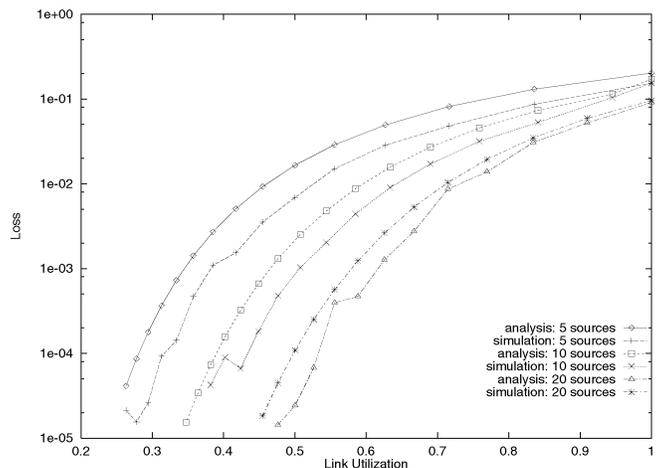


Fig. 10. Statistical multiplexing gain.

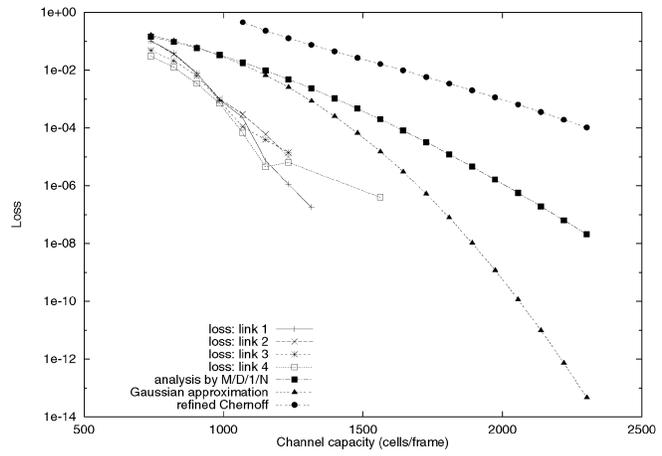


Fig. 11. Effect of large bufer size.

three cases and, thus, the histogram-based model satisfies our requirements regardless of the number of statistical real-time channels multiplexed. Moreover, the statistical multiplexing gain is shown to increase with the number of channels multiplexed. However, in order to establish a macrochannel with the cell-loss probability of 10^{-4} , we need to reserve the bandwidth which is twice the average cell-generation rate. This results from the high burstiness of MPEG data and is inevitable in order to satisfy the QoS requirement. Although the macrochannel's utilization is about 0.5, it does not necessarily mean the waste of bandwidth since the unused bandwidth by the macrochannel can be used for transmission of best-effort traffic, as in the case of real-time channels [9].

Finally, we investigated the effects of buffer size on the cell-loss ratio although a small buffer system is our main target. We set the buffer size to 500 cells. The cell-loss ratios obtained from the simulation are much smaller than the bound as seen in Fig. 11. This indicates that the assumption that a macrochannel with the aggregate input reaches the steady-state quickly is violated in a large-buffer system. The queue trajectory in Fig. 12 confirms this reasoning. The

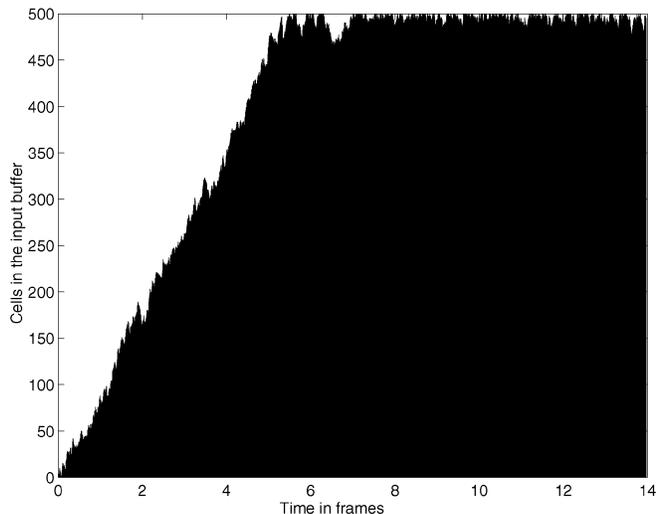


Fig. 12. Queue trajectory in a large buffer system.

buffer size N is 500 cells, the cell-arrival rate is 2,000 cells/frame, and the service rate is 1,300 cells/frame. Therefore, the system is overloaded and cells are lost in steady-state. Starting from the "empty" state, it takes several frame periods for the queue to reach its steady-state. Considering the fact that the cell-generation rate remains constant for at most one frame period, the queueing system cannot reach the steady state under most overloaded conditions obtained from the histogram-based approach, and thus, the cell-loss ratio is quite low as compared to the estimate. This, in turn, indicates that our framework results in overreservation of network resources when the buffer size is very large, but it still satisfies the basic requirements of statistical real-time communication since the cell-loss probabilities are bounded. On the other hand, the CLT approximation seems to predict the cell-loss ratio of a large buffer system comparatively well, because the longer tail of the rate distribution ignored by the CLT approximation is compensated by a large buffer. Therefore, one cannot argue that the better performance of CLT approximation is due to its accurate modeling of a large buffer system.

5 CONCLUSION

In this paper, we have proposed a framework for statistical real-time communication in ATM networks. The framework uses the TCRM [9], which was originally proposed for deterministic real-time communication in high-speed networks like ATM. While only one real-time channel is assigned in the deterministic version, multiple statistical real-time channels are aggregated as a macrochannel over each ATM link and cells from different component channels of the macrochannel are serviced on a FIFO basis. To quantify the cell-loss ratio of a macrochannel, we have proposed to use a histogram-based model for the input traffic specification of MPEG video sources. The histogram-based model specifies an MPEG video source with the histogram of time-sampled traffic-generation rates. Also, in the histogram-based model, input traffic specification for aggregate sources is derived by taking the convolution of the histograms of component sources, making it simple to apply the model to a set of statistical real-time channels. Assuming a constant cell-arrival rate, we have modeled the arriving traffic at the macrochannel as a Poisson process and derived the cell-loss ratio from the $M/D/1/N$ analysis. Then, the cell-loss ratio of the macrochannel is given as the weighted sum of cell-loss ratios, and the weight was given by the rate pmf. For a faster channel-establishment process, we presented an approximation of the cell-blocking probability of an $M/D/1/N$ system. Based on this analysis of a macrochannel, we have extended a statistical real-time channel to the multiple-hop case. In order to consider the traffic change in macrochannels at intermediate links, we have made an assumption of independence of cell interarrival times at the intermediate links. Lastly, our scheme was extended in the context of VPs.

We have conducted an extensive simulation study using real MPEG-coded video data in order to evaluate the effectiveness of our framework for statistical real-time communication and validated the assumptions used. The simulation results have reasonably well matched the

analysis that is based on the assumptions including the histogram-based modeling and Poisson arrival at each link, although, in some cases, over-reservation of network resources has been observed.

We have applied the histogram-based model to MPEG video data in order to achieve statistical multiplexing gain. However, it is not clear whether the histogram-based model can be applied to general VBR real-time traffic. Hence, our current research effort focuses on the applicability of the histogram-based model to more general VBR real-time traffic for statistical multiplexing.

ACKNOWLEDGMENTS

The authors would like to thank Mark Garret (Bellcore, Morristown, New Jersey) and O. Rose (Institute of Computer Science, University of Würzburg, Germany) for providing MPEG data used in our simulations.

The work reported in this paper was supported in part by the National Science Foundation under grant MIP-9203895 and the Office of Naval Research under grant N00014-99-1-0465. The opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the view of the funding agencies. A subset of the materials in this paper was presented at the IEEE RTSS99.

REFERENCES

- [1] L. Zhang, "Virtual Clock: A New Traffic Control Algorithm for Packet Switching Networks," *Proc. ACM SIGCOMM*, pp. 19-29, 1990.
- [2] C. Kalmanek, H. Kanakia, and S. Keshav, "Rate Controlled Servers for Very High-Speed Networks," *Proc. IEEE GLOBECOM 90*, pp. 1220, 1990.
- [3] S.J. Golestani, "A Stop-and-Go Queueing Framework for Congestion Management," *Proc. ACM SIGCOMM*, pp. 8-18, 1990.
- [4] D. Ferrari and D.C. Verma, "A Scheme for Real-Time Channel Establishment in Wide-Area Networks," *IEEE J. Selected Areas Comm.*, vol. 8, no. 3, pp. 368-379, Apr. 1990.
- [5] H. Zhang and D. Ferrari, "Rate-Controlled Static-Priority Queueing," *Proc. IEEE INFOCOM*, pp. 227-236, 1993.
- [6] A. Demers, S. Keshav, and S. Shenker, "Analysis and Simulation of a Fair Queueing Algorithm," *Proc. ACM SIGCOMM*, pp. 1-12, 1989.
- [7] A.K. Parekh, "A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks," PhD thesis, Dept. of Electrical Eng. and Computer Science, MIT, Feb. 1992.
- [8] D.D. Kandlur, K.G. Shin, and D. Ferrari, "Real-Time Communication in Multi-Hop Networks," *Proc. 11th Int'l Conf. Distributed Computing Systems*, pp. 300-307, May 1991.
- [9] S.-K. Kweon and K.G. Shin, "Providing Deterministic Delay Guarantees in ATM Networks," *IEEE/ACM Trans. Networking*, vol. 6, no. 6, pp. 838-850, Dec. 1998.
- [10] C.M. Aras, J. Kurose, D.S. Reeves, and H. Schulzrinne, "Real-Time Communication in Packet-Switched Networks," *Proc. IEEE*, vol. 82, no. 1, pp. 122-139, Jan. 1994.
- [11] L. Georgiadis, R. Guérin, V. Peris, and K.N. Sivarajan, "Efficient Network QoS Provisioning Based on per Node Traffic Shaping," *Proc. IEEE INFOCOM*, pp. 102-110, 1996.
- [12] J. LeBoudec, "Applications of Network Calculus to Guaranteed Service Networks," *IEEE Trans. Information Theory*, vol. 44, no. 3, pp. 1087-1096, May 1998.
- [13] J. Kurose, "On Computing Per-Session Performance Bounds in High-Speed Multi-Hop Computer Networks," *Proc. ACM SIGMETRICS*, pp. 128-139, 1992.
- [14] H. Zhang and E. Knightly, "Providing End-to-End Statistical Guarantees Using Bounding Interval Dependent Stochastic Models," *Proc. ACM SIGMETRICS*, pp. 211-220, 1994.

- [15] Z. Zhang, D. Towsley, and J. Kurose, "Statistical Analysis of the Generalized Processor Sharing Scheduling Discipline," *IEEE J. Selected Areas Comm.*, vol. 13, no. 6, pp. 1071-1080, Aug. 1995.
- [16] O. Yaron and M. Sidi, "Performance and Stability of Communication Networks via Robust Exponential Bounds," *IEEE/ACM Trans. Networking* vol. 1, no. 3, pp. 372-385, June 1993.
- [17] A.K. Parekh and R.G. Gallager, "A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks: The Single Node Case," *IEEE/ACM Trans. Networking*, vol. 1, no. 3, pp. 344-357, June 1993.
- [18] A.K. Parekh and R.G. Gallager, "A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks: The Multiple Node Case," *Proc. of IEEE INFOCOM*, pp. 521-530, 1993.
- [19] R.L. Cruz, "A Calculus for Network Delay, Part I: Network Elements in Isolation," *IEEE Trans. Information Theory*, vol. 37, no. 1, pp. 114-131, Jan. 1991.
- [20] R.L. Cruz, "A Calculus for Network Delay, Part II: Network Analysis," *IEEE Trans. Information Theory*, vol. 37, no. 1, pp. 132-142, Jan. 1991.
- [21] M. Reisslein, K.W. Ross, and S. Rajagopal, "Guaranteeing Statistical (QoS) to Regulated Traffic: The Multiple Node Case," *Proc. 37th IEEE Conf. Decision and Control*, pp. 531-538, Dec. 1998.
- [22] M. Reisslein, K.W. Ross, and S. Rajagopal, "Guaranteeing Statistical (QoS) to Regulated Traffic: The Single Node Case," *Proc. IEEE Infocom '99*, pp. 1060-1071, Mar. 1999.
- [23] R. Guérin, H. Ahmadi, and M. Naghshineh, "Equivalent Capacity and its Application to Bandwidth Allocation in High-Speed Networks," *IEEE J. Selected Areas Comm.*, vol. 9, no. 7, pp. 968-981, Sept. 1991.
- [24] G. Kesidis, J. Walrand, and C.-S. Chang, "Effective Bandwidths for Multiclass Markov Fluids and Other ATM Sources," *IEEE/ACM Trans. Networking*, vol. 1, no. 4, pp. 424-428 Aug. 1993.
- [25] A. Elwalid, D. Mitra, and R. Wentworth, "A New Approach for Allocating Buffers and Bandwidth to Heterogeneous, Regulated Traffic in an ATM Node," *IEEE J. Selected Areas Comm.*, vol. 13, no. 6, pp. 1115-1127 Aug. 1995.
- [26] P. Skelly, S. Dixit, and M. Schwartz, "A Histogram-Based for Video Traffic Behavior in an ATM Network Node with an Application to Congestion Control," *Proc. IEEE INFOCOM*, pp. 95-104, 1992.
- [27] C.C. Chou and K.G. Shin, "Statistical Real-Time Video Channels over a Multiaccess Network," *Proc. High-Speed Networking and Multimedia Computing*, pp. 86-96, Feb. 1994.
- [28] M. dePrycker, *Asynchronous Transfer Mode: Solution for Broadband SDN*. London: Ellis Horwood, 1991.
- [29] C. Liu and J.W. Layland, "Scheduling Algorithms for Multiprogramming in a Hard Real-Time Environment," *J. ACM*, vol. 20, no. 1, pp. 44-61, Jan. 1973.
- [30] M.W. Garrett and M. Vetterli, "Joint Source/Channel Coding of Statistically Multiplexed Real-Time Services on Packet Networks," *IEEE/ACM Trans. Networking*, vol. 1, no. 1, pp. 71-80, Feb. 1993.
- [31] J. Beran, R. Sherman, M.S. Taqqu, and W. Willinger, "Long-Range Dependence in Variable Bit-Rate Video Traffic," *IEEE Trans. Comm.*, vol. 43, pp. 1566-1579, 1995.
- [32] M.R. Frater, J.F. Arnold, and P. Tan, "A New Statistical Model for Traffic Generated by VBR Coders for Television on the Broadband ISDN," *IEEE Trans. Circuits and Systems for Video Technology*, vol. 4, no. 6, pp. 521-526, Dec. 1994.
- [33] M.W. Garrett and W. Willinger, "Analysis, Modeling, and Generation of Self-Similar VBR Video Traffic," *Proc. ACM SIGCOMM*, pp. 269-280, Sept. 1994.
- [34] D.P. Heyman, A. Tabatabai, and T. Lakshman, "Statistical Analysis and Simulation Study of Video Teleconferencing in ATM Networks," *IEEE Trans. Circuits and Systems for Video Technology*, vol. 2, no. 1, pp. 49-59, Mar. 1992.
- [35] D. Heyman, A. Tabatabai, and T. Lakshman, "Statistical Analysis of MPEG2-Coded VBR Video Traffic," *Proc. Sixth Int'l Workshop Packet Video*, 1994.
- [36] C. Huang, M. Devetsikiotis, L. Lambadaris, and A. Kaye, "Modeling and Simulation of Self-Similar Variable Bit Rate Compressed Video: A Unified Approach," *Proc. ACM SIGCOMM*, 1995.
- [37] B. Jabbari, F. Yegengolu, Y. Kuo, S. Zafar, and Y.-Q. Zhang, "Statistical Characterization and Block-Based Modeling of Motion-Adaptive Coded Video," *IEEE Trans. Circuits and Systems for Video Technology*, vol. 3, no. 3, pp. 199-207, June 1993.
- [38] B. Maglaris, D. Anastassiou, P. Sen, G. Karlsson, and J.D. Robbins, "Performance Models of Statistical Multiplexing in Packet Video Communications," *IEEE Trans. Comm.*, vol. 36, no. 7, pp. 834-844, July 1988.
- [39] B. Melamed and D. Pendarakis, "A TES-Based Model for Compressed Star Wars Video," *Proc. IEEE GLOBECOM*, pp. 120-126, 1994.
- [40] B. Melamed, D. Raychaudhuri, B. Sengupta, and J. Zdepski, "TES-Based Video Source Modeling for Performance Evaluation of Integrated Networks," *IEEE Trans. Comm.*, vol. 42, no. 10, pp. 2773-2777, Oct. 1994.
- [41] P. Pancha and M.E. Zarki, "Bandwidth-Allocation Schemes for Variable-Bit-Rate MPEG Sources in ATM Networks," *IEEE Trans. Circuits and Systems for Video Technology*, vol. 2, no. 1, pp. 49-59, Mar. 1992.
- [42] G. Ramamurthy and B. Sengupta, "Modeling and Analysis of a Variable Bit Rate Video Multiplexer," *Proc. IEEE INFOCOM*, pp. 812-827, 1992.
- [43] M. Krunz, R. Sass, and H. Hughes, "Statistical Characteristics and Multiplexing of MPEG Streams," *Proc. IEEE INFOCOM*, pp. 455-462, 1995.
- [44] M. Krunz and H. Tripathi, "On the Characterization of VBR MPEG Streams," *Proc. ACM SIGMETRICS*, pp. 192-202, June 1997.
- [45] M. Krunz and H. Tripathi, "Impact of Video Scheduling on Bandwidth Allocation for Multiplexed MPEG Streams," *Multimedia Systems J.*, vol. 5, no. 6, pp. 47-57, Dec. 1997.
- [46] N. Shroff and M. Schwartz, "Video Modeling within Networks Using Deterministic Smoothing at the Source," *Proc. IEEE INFOCOM*, pp. 342-349, 1994.
- [47] D. Bertsekas and R. Gallager, *Data Networks*, second ed. Englewood Cliffs, New Jersey: Prentice-Hall Int'l, 1992.
- [48] S.-K. Kweon and K.G. Shin, "Traffic-Controlled Rate-Monotonic Priority Scheduling," *Proc. IEEE INFOCOM*, pp. 655-662, 1996. 'Also an expanded version appeared in: *IEEE/ACM Trans. Networking*, vol. 6, no. 6, pp. 838-850, Dec. 1998.'
- [49] R.B. Cooper, *Introduction to Queueing Theory*, second ed. New York: North-Holland Publishing, 1981.
- [50] E. Gelenbe and G. Pujolle, *Introduction to Queueing Networks*. New York: John Wiley and Sons, 1987.
- [51] L. Kleinrock, *Communication Nets: Stochastic Message Flow and Delay*. New York: McGraw-Hill, 1964.
- [52] O. Rose, "Statistical Properties of MPEG Video Traffic and Their Impact on Traffic Modeling in ATM Systems," Technical Report 101, Univ. of Wuerzburg Inst. of Computer Science, Feb. 1995. (many MPEG-1 traces are available via FTP from ftp-info3.informatik.uni-wuerzburg.de in pub/MP).
- [53] M. Reisslein and K.W. Ross, "Call Admission for Pre-recorded Sources with Packet Loss," *IEEE J. Selected Areas Comm.*, vol. 15, no. 6, pp. 1167-1180, Aug. 1997.
- [54] R.R. Bahadur and R. Rao, "On Deviations of the Sample Mean," *Ann. Math. Statis.*, vol. 31, pp. 1015-1027, 1960.



Seok-Kyu Kweon received the BS and MS degree in electronics from Seoul National University, Korea, and the PhD degree in electrical engineering and computer science from the University of Michigan, Ann Arbor in 1998. Currently, he is working for Cisco systems in San Jose, California, and working on IP QoS. Before joining Cisco systems, he worked for General Motors. His research interests are traffic management, scheduling algorithms for

highspeed networks, statistical QoS provisioning, QoS routing, and real-time communication for manufacturing automation. He is a member of the IEEE.



Kang G. Shin received the BS degree in electronics engineering from Seoul National University, Korea in 1970, and both the MS and PhD degrees in electrical engineering from Cornell University, Ithaca, New York in 1976 and 1978, respectively. From 1978 to 1982, he was on the faculty of Rensselaer Polytechnic Institute, Troy, New York. He has held visiting positions at the United States Airforce Flight Dynamics Laboratory, AT&T Bell Laboratories, Computer Science Division within the Department of Electrical Engineering and Computer Science at UC Berkeley, and International Computer Science Institute, Berkeley, California, IBM T.J. Watson Research Center, and Software Engineering Institute at Carnegie Mellon University. He also chaired the Computer Science and Engineering Division, EECS Department, the University of Michigan for three years, beginning January 1991. Dr. Shin is professor and founding director of the Real-Time Computing Laboratory, Department of Electrical Engineering and Computer Science, the University of Michigan, Ann Arbor. He has supervised the completion of 40 PhD theses, and authored/coauthored more than 600 technical papers and numerous book chapters in the areas of distributed real-time computing and control, computer networking, fault-tolerant computing, and intelligent manufacturing. He has co-authored (jointly with C.M. Krishna) a textbook *Real-Time Systems*, McGraw Hill, 1997. In 1987, he received the Outstanding IEEE Transactions on Automatic Control Paper Award, and Research Excellence Award in 1989, Outstanding Achievement Award in 1999, and Service Excellence Award in 2000 from The University of Michigan. In 1985, he founded the Real-Time Computing Laboratory, where he and his colleagues are investigating various issues related to real-time and fault-tolerant computing. His current research focuses on Quality of Service (QoS) sensitive computing and networking with emphases on timeliness and dependability. He has also been applying the basic research results to telecommunication and multimedia systems, embedded systems, and manufacturing applications. He is an IEEE fellow and member of the Korean Academy of Engineering, was the general chair of the 2000 IEEE Real-Time Technology and Applications Symposium, the program chairman of the 1986 IEEE Real-Time Systems Symposium (RTSS), the general chairman of the 1987 RTSS, the guest editor of the 1987 August special issue of *IEEE Transactions on Computers* on Real-Time Systems, a Program co-chair for the 1992 *International Conference on Parallel Processing*, and served numerous technical program committees. He also chaired the IEEE Technical Committee on Real-Time Systems during 1991-93, was a distinguished visitor of the Computer Society of the IEEE, an editor of *IEEE Transactions on Parallel and Distributed Computing*, and an area editor of *International Journal of Time-Critical Computing Systems* and *Computer Networks*.