Fundamental Limits and Tradeoffs of Providing Deterministic Guarantees to VBR Video Traffic

E. Knightly[*], D. Wrege[†], J. Liebeherr[†], and H. Zhang[‡]

[*]EECS Department, U. C. Berkeley

[†]Department of Computer Science, U. of Virginia

[‡‡]Lawrence Berkeley Laboratory

TR-94-067

Abstract

Compressed digital video is one of the most important traffic types in future integrated services networks. However, a network service that supports delay-sensitive video imposes many problems since compressed video sources are variable bit rate (VBR) with a high degree of burstiness. In this paper, we consider a network service that can provide deterministic guarantees on the minimum throughput and the maximum delay of VBR video traffic. A common belief is that due to the burstiness of VBR traffic, such a service will not be efficient and will necessarily result in low network utilization. We investigate the fundamental limits and tradeoffs in providing deterministic performance guarantees to video and use a set of 10 to 90 minute long MPEG-compressed video traces for evaluation. Contrary to conventional wisdom, we are able to show that a deterministic service can be provided to video traffic even while maintaining a high level of network utilization. We first consider an ideal network environment that employs the most accurate video traffic characterizations,

Earliest-Deadline-First packet schedulers, and exact admission control conditions. The utilization achievable in this situation provides the fundamental limits of a deterministic service. We then investigate the utilization limits in a network environment that takes into account practical constraints, such as the need for fast policing mechanisms, simple packet scheduling algorithms, and efficient admission control tests. Even when considering these practical tradeoffs, we demonstrate that a considerably high network utilization is achievable by a deterministic service.

1 Introduction

Emerging Broadband Integrated Services Digital Networks (B-ISDN) must support applications with diverse traffic characteristics and performance requirements. Of the many traffic classes in future B-ISDN networks, delay-sensitive Variable Bit Rate (VBR) video traffic poses a unique challenge. Since the service is delay-sensitive, the network must support a resource reservation scheme [4] that allocates network resources for each VBR video stream. However, the burstiness of VBR video traffic makes it difficult to determine the amount of resources required. On the one hand, if resources are reserved at the average rate of the VBR video source, unacceptable delays may result if the source is transmitting at its peak rate. On the other hand, if resource reservations are based on the peak rate, the network may be under-utilized most of the time.

This study investigates utilization limits achievable in B-ISDN networks with service guarantees to delay-sensitive VBR video traffic. Throughout this study we assume that service guarantees are *deterministic*, i.e., no guarantees are ever violated. Any network that offers service guarantees, especially guarantees on the maximum delay, requires admission control tests for new connection and policing mechanisms for existing connections [1, 4, 8, 9]. *Admission control tests* determine if the network has sufficient resources to accommodate a new connection without degrading the service of existing connections. The admission tests for a new connection are based on a characterization of the expected traffic for the new connection. If accepting or admitting the new connection may result in violations of service guarantees for any existing connection, the connection is not established. For all connections that are admitted, the network performs *traffic policing*, that is, it ensures that all connections adhere to the traffic characterization given to the network for the admission control tests. Traffic from a connection that is in excess of its characterization is not permitted to enter the network.

A key challenge of incorporating VBR video traffic into networks with service guarantees lies in the difficulty of finding an appropriate traffic characterization that captures the dynamics of the traffic. A rich set of literature exists on characterizing VBR video traffic by stochastic processes such as Markov-Modulated Poisson Processes [13, 15, 16, 19]. Other approaches include Auto Regressive Moving Average models [7], histogram-based models [17], and self-similar models [6]. While stochastic models of a source have the advantage that they may be used to achieve higher network utilization via statistical multiplexing, they have several important disadvantages. First, it is difficult to implement a policing mechanism that enforces a stochastic traffic characterization. Moreover, most stochastic models for characterizing video traffic are either not powerful enough to capture the burstiness and the timely correlations of a video source, or they are too complex for practical implementation [14]. Finally, since stochastic approaches to VBR traffic characterization provide only statistical guarantees, they cannot be applied to deterministic network services [10, 11, 12, 21].

For these reasons, we use in our work *deterministic* traffic models that characterize VBR traffic in terms of a worst-case description. While it is a common belief that such a worst-case approach necessitates an allocation of network resources according to a source's peak rate, leading to low network utilization, we show that this conventional wisdom is not necessarily correct. Recently, several methods have been proposed that express the burstiness of VBR traffic sources in terms of a worst-case description. In particular, the (x_{min}, x_{ave}, I, s) model [4] and the Deterministic Bounding Interval Dependent or D-BIND model [10] have been shown to significantly increase the number of admitted connections when the traffic is bursty.

In this study, we analyze the fundamental limits and practical tradeoffs in providing deterministic performance guarantees to VBR video. Such a study has become feasible with a set of recently developed tight, i.e., necessary and sufficient, admission control tests for networks that employ Earliest-Deadline-First and Static Priority scheduling disciplines [11, 12]. We present an optimal characterization of VBR video traffic for a deterministic service by so-called *empirical envelopes*. Using the optimal characterization, the Earliest-Deadline-First algorithm for packet scheduling, and the recently presented exact admission control tests from [11, 12], we provide insights into the highest achievable utilization for actual VBR video sources.

We also investigate the fundamental tradeoffs that must be considered for practical networking environments. For example, a realistic traffic characterization must be easily policable by the network. Furthermore, the traffic characterization should be concise so that the traffic of each source can be described with a small number of parameters. We study the impact of selecting easily policable and concise traffic characterizations on the utilization limits of the deterministic service. We also investigate the impact of selecting simpler packet schedulers, such as Static Priority (SP) and First-Come-First-Served (FCFS), on the network utilization. Finally, for the Static Priority packet scheduler we demonstrate the degree to which the accuracy of the admission control tests influences the utilization of a deterministic service.

The remainder of the paper is structured as follows: In Section 2 we review the network components necessary to offer a service with deterministic delay bounds. In Section 3 we devise a method to generate an optimal traffic characterization, the *empirical envelope*, for VBR video sources. We present an algorithm that allows us to approximate the empirical envelope by an easily policable deterministic traffic model. In Section 4 we present an empirical evaluation of the deterministic approach by considering the abovementioned tradeoffs. The evaluation is performed using traces of MPEG-encoded VBR video traffic. We summarize the paper in Section 5.

2 Components of a Deterministic Network Service

For a network to provide deterministic performance guarantees on throughput, delay, delay-jitter, and loss, it must be able to tightly control all available resources. As mentioned before, admission control tests and

policing mechanisms are effective means to control the number of connections as well as the traffic from each connection; they thus determine the degree to which network resources are utilized. Both admission control tests and policing mechanisms are influenced by three components: (1) the traffic model used for the characterization of the worst-case traffic from a connection, (2) the scheduling discipline at the network switches, and (3) the accuracy of the admission control functions. In this section, we review the tradeoffs involved in the selection of these components and their impact on a deterministic network service.

2.1 Deterministic Traffic Models

In a deterministic network service, a deterministic traffic model of all sources receiving the service is required. Such a model has several fundamental requirements. First, the model must be a *worst-case* characterization of the source to provide an absolute upper bound on a source's packet arrivals. Second, the model must be *parameterized* so that a source can efficiently specify its traffic characterization to the network. Next, the model should characterize the traffic as *accurately* as possible so that the admission control algorithms do not over-estimate the resources required by the connection. Finally, the model must be *policable* so that the network can enforce a source's traffic characterization.

A worst-case representation of a traffic source may be described as follows. If the actual traffic of a connection is given by a function A such that $A[t, t+\tau]$ denotes the traffic arrivals in the time interval $[t, t+\tau]$, an upper bound on A can be given by a function A^* if for all times $t \geq 0$ and all $\tau \geq 0$ the following holds [2, 3]:

$$A[t, t+\tau] \leq A^*(\tau) \tag{1}$$

We refer to a function $A^*(\tau)$ that satisfies the property in (1) as a *traffic constraint function*. Note that a traffic constraint function provides a *time-invariant* bound of A so that a source is bounded for every interval of length τ.

In practice, a source specifies its traffic characterization with a parameterized model. The parameterized deterministic traffic model then defines a traffic constraint function that bounds the source. For example, in the (x_{min}, x_{ave}, I, s) model [4], x_{min} is the minimum packet inter-arrival time, x_{ave} is the maximum average packet inter-arrival time over any time interval of length I, and s is the maximum packet size. This worst-case parameterization of a source has an associated traffic constraint function given in Table 1. The (σ, ρ) model [3] describes traffic in terms of a rate factor ρ and a burstiness factor σ such that during any time interval of length t, the traffic from a connection is less than $\sigma + \rho t$. In this study, we consider a $(\vec{\sigma}, \vec{\rho})$ model, which can be viewed as an extension to the (σ, ρ) model. The $(\vec{\sigma}, \vec{\rho})$ model maintains a number n of (σ, ρ) pairs, where the amount of traffic in a time interval t is restricted to $\min_{1 \leq i \leq n}\{\sigma_i + \rho_i t\}$. A more general deterministic traffic model called the *Deterministic Bounding Interval Dependent* or *D-BIND* model [10] characterizes traffic by a family of rate-interval tuples $\{R_i, T_i\}_{i=1}^{n}$. The D-BIND model defines an n

Traffic Model	Traffic Constraint Function
(σ, ρ)	$A^*(t) = \sigma + \rho t$
$(\vec{\sigma}, \vec{\rho})$	$A^*(t) = \min_i\{\sigma_i + \rho_i t\}$
(x_{min}, x_{ave}, I, s)	$A^*(t) = \lfloor \frac{t}{I} \rfloor \cdot \frac{I \cdot s}{x_{ave}} + \min\left\{ \lceil (\frac{t}{I} - \lfloor \frac{t}{I} \rfloor) \cdot \frac{I}{x_{min}} \rceil, \frac{I}{x_{ave}} \right\} \cdot s$
D-BIND	$A^*(t) = \begin{cases} tR_1 & \text{for all } t \leq T_1. \\ \frac{T_{j-1}T_j(R_{j-1} - R_j) + t(R_jT_j - R_{j-1}T_{j-1})}{T_j - T_{j-1}} & \text{for all } T_{j-1} < t \leq T_j. \end{cases}$

Table 1: Traffic Models and Traffic Constraint Functions.

segment piece-wise linear traffic constraint function as given in Table 1. The $(\vec{\sigma}, \vec{\rho})$ model may be viewed as a special case of the D-BIND model because it defines an n segment piece-wise linear *concave* traffic constraint function.

Table 1 depicts the traffic constraint functions A^* for all aforementioned deterministic traffic models. Since deterministic traffic models have associated traffic constraint functions, the admission control tests may be expressed in terms of A^*. A requirement on A^* for the admission control tests is subadditivity, i.e., for all $t \geq 0$ and all $\tau \geq 0$ we have $A^*(t_1 + t_2) \leq A^*(t_1) + A^*(t_2)$. Thus, A^* can be used as its own worst-case characterization, that is, $A^*[t, t+\tau] \leq A^*(\tau)$ $(t \geq 0,\ \tau \geq 0)$ [12].

Clearly, the shape of a traffic constraint function depends on the selection of the traffic model. From the perspective of achievable network utilization, the model should have a traffic constraint function that is as tight as possible so that the admission control algorithms do not over-estimate the resources required by the connection. While in general, a model with more parameters can achieve a more accurate or tight traffic constraint function, the additional parameterization causes an increase in the complexity of *policing* the traffic model. Thus, the selection of an appropriate traffic model for a deterministic service must find a compromise between the high complexity preferred by the admission control tests and the simplicity required for the implementation of traffic policing.

Traffic policing mechanisms must verify in real-time whether the traffic transmitted on an established connection adheres to a specified set of parameters of a deterministic traffic model. To ensure that the policing mechanisms can monitor and control traffic at high data rates, the complexity of the traffic model is strictly limited. In [3] it was shown that a traffic model with a piece-wise linear concave traffic constraint function can be policed by a fixed number of leaky buckets. Since a leaky bucket can be implemented with a

counter and a single timer [18], concavity of the traffic constraint functions ensures a simple implementation of the traffic policer. Therefore, by restricting ourselves to deterministic traffic models which are guaranteed to yield concave piece-wise linear traffic constraint functions, such as the (σ, ρ) model and the $(\vec{\sigma}, \vec{\rho})$ model, we can always ensure a simple implementation of traffic policing mechanisms.

2.2 Packet Scheduling

In a connection-oriented packet-switched network, packets from a particular connection traverse the network on a fixed path of switches and links. Each switch has a packet scheduler for each outgoing link. Since the packet scheduler can transmit only one packet at a time, it maintains a queue containing all packets waiting for transmission. Here we consider the following well-known scheduling disciplines at a packet scheduler for a set $\mathcal{N}$ of connections: First-Come-First-Served (FCFS), Static Priority (SP), and Earliest-Deadline-First (EDF). Each of these disciplines has been investigated for use in bounded delay services [3, 4, 12, 20].

First-Come-First-Served (FCFS): FCFS schedulers transmit all packets in the order of their arrival. Since the maximum delay in a FCFS scheduler is the same for all connections $j \in \mathcal{N}$, all connections have an identical delay bound d.

Static Priority (SP): In an SP packet scheduler, each connection $j \in \mathcal{N}$ is assigned a priority p with $1 \leq p \leq P$, where a lower priority index indicates a higher priority. $\mathcal{C}_p$ is the set of connections with priority p, and all connections in $\mathcal{C}_p$ have the same delay bound d_p, with $d_p < d_q$ for $p < q$. SP schedulers maintain one FCFS queue for each priority level, always selecting the first packet in the highest-priority FCFS queue for transmission.

Earliest-Deadline-First (EDF): With EDF scheduling, each connection $j \in \mathcal{N}$ is assigned a delay bound d_j, where the delay bound may be different for each connection. An EDF scheduler selects packets for transmission in increasing order of packet deadlines, where packet deadlines are calculated as the sum of the arrival time and the delay bound of a packet.

The selection of a particular scheduling discipline for a packet scheduler involves a tradeoff between the need to support a large number of connections with diverse delay requirements and the need for simplicity in the scheduling operations. For example, while a FCFS scheduler can be easily implemented, it can effectively support only one delay bound for all connections. On the other extreme, while an EDF scheduler can support a different delay bound for each connection, the scheduling operations of EDF are complex since they involve a search operation for the packet with the shortest deadline.

Delay Bound Test	Condition
FCFS Exact	$d \geq \sum_{j \in \mathcal{N}} A_j^*(t) - t + \max_{k \in \mathcal{N}} s_k$ for all $t \geq 0$.
SP Exact	$(\exists \tau \leq d_p)\ t + \tau \geq \sum_{j \in \mathcal{C}_p} A_j^*(t) + \sum_{q=1}^{p-1} \sum_{j \in \mathcal{C}_q} A_j^*(t+\tau) + \max_{r>p} s_r$ for all $p,\ t \geq 0$.
EDF Exact	$\begin{cases} t \geq \sum_{j \in \mathcal{N}} A_j^*(t - d_j) & \text{for all } t \geq 0. \\ t \geq \sum_{j \in \mathcal{N}} A_j^*(t - d_j) + \max_{d_k > t} s_k & \text{for all } d_1 \leq t < d_{\lvert\mathcal{N}\rvert}. \end{cases}$

Table 2: Exact delay bound tests for FCFS, SP, and EDF packet schedulers.

Delay Bound Test	Condition
SP Sufficient 1	$t \geq \sum_{j \in \mathcal{C}_p} A_j^*(t - d_p) + \sum_{q=1}^{p-1} \sum_{j \in \mathcal{C}_q} A_j^*(t) + \max_{r>p} s_r$ for all $p,\ t \geq d_p$.
SP Sufficient 2	$d_p \geq \sum_{q=1}^{p} \sum_{j \in \mathcal{C}_q} A_j^*(d_p) + \max_{r>p} s_r$ for all p.

Table 3: Sufficient delay bound tests for an SP packet scheduler.

2.3 Admission Control Tests

While the number of admission control tests can be large, including tests for availability of transmission capacity, CPU power, buffer space, etc., the most crucial admission test in a network with a deterministic service is the *delay bound test*. The delay bound test verifies that, for all connections, the delay of each packet is less than its required delay bound. Most other admission control tests, including those that verify throughput and delay-jitter guarantees, can be directly derived from the delay bound test.

Note that the complexity of a delay bound test depends on both the scheduling algorithm and the desired accuracy of the test. For certain packet schedulers the calculations required for accurate delay bound tests can be computationally expensive. Thus, less accurate but fast delay bound tests are often desirable.

In Tables 2 and 3 we present conditions that must be satisfied to pass delay bound tests for FCFS, SP, and EDF packet schedulers.[1] Table 2 shows the delay bound tests with the highest degree of accuracy (*"Exact Tests"*) for each scheduler. Formal derivations of these tests can be found in [3, 12]. Since the exact admission test for an SP scheduler is computationally very complex, Table 3 presents less accurate tests

[1] In the tables, s_k denotes the maximum transmission time of any packet on connection k. Also, A_j^* denotes the traffic constraint function for a connection $j \in \mathcal{N}$.

(*"Sufficient Tests"*) that require less computation [12, 21]. Note that each test is formulated so that it can be applied to the traffic models discussed in Section 2.1.

As an example, we present an informal derivation of the delay bound test for the FCFS packet scheduler. Recall that the FCFS packet scheduler can offer only one delay bound d to all connections. Thus, a delay bound test for FCFS simply verifies that the maximum waiting time of any packet never exceeds d. If the scheduler is never idle in the time interval $[0, t]$, then the maximum waiting time of a packet that arrives at time t is determined by $\sum_{j \in \mathcal{N}} A_j[0, t] - t$, the maximum backlog in the transmission queue, and $\max_{k \in \mathcal{N}} s_k$, the largest remaining transmission time of any packet that is in transmission at time t. Since $A_j[0, t] \leq A_j^*(t)$ by equation (1), we are guaranteed not to have a deadline violation at time $t + d$ if and only if the following holds at time t:

$$d \geq \sum_{j \in \mathcal{N}} A_j^*(t) - t + \max_{k \in \mathcal{N}} s_k$$

Applying the argument to all times $t \geq 0$ yields the delay bound test given in Table 2.

3 Deterministic Characterizations of VBR Video Traffic

As described in Section 2.1, each deterministic traffic model has an associated traffic constraint function $A^*(t)$ that provides an upper bound on the source's arrivals in any interval of length t. In order to evaluate the maximum achievable utilization for a deterministic service to VBR video, we first present the so-called *empirical envelope* that provides the most accurate traffic constraint function for a given video trace. For any scheduling discipline, using the exact admission control tests with the empirical envelope results in the highest network utilization achievable in a deterministic service. Although the empirical envelope can be seen as the best traffic constraint function from the perspective of the admission control test, it cannot be efficiently specified or policed since the number of required parameters is the same as the number of frames in the entire video sequence. Despite its lack of practical properties, the empirical envelope provides an optimal benchmark for evaluating deterministic traffic models. Next, we describe how to obtain the empirical envelope from a given video trace. Additionally, we show how parameters for the $(\vec{\sigma}, \vec{\rho})$ model can be obtained from the empirical envelope.

In this paper, we investigate VBR video traffic that is generated according to the MPEG compression algorithm [5]. An MPEG coder generates three types of frames: I frames that use *intraframe* compression, and P and B frames that are transmitted between I frames that use *interframe* compression. While P frames (predicted frames) are coded based on only past frames, B frames (bidirectional frames) are coded based on both past and future frames. Thus, I frames exploit spatial locality, while P and B frames exploit temporal locality. More details of the MPEG algorithm may be found in [5]. Figure 1 shows an MPEG-compressed video trace of advertisement sequences with the frame pattern *IBBPBB*. The behavior of the coder is clearly

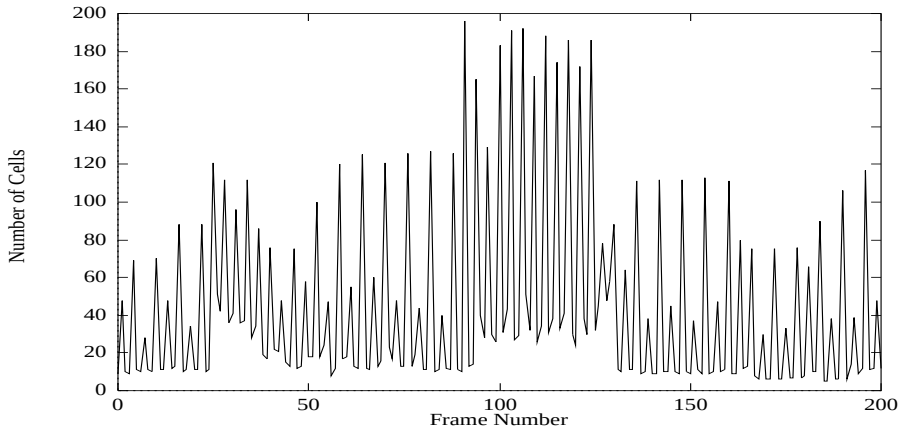

Figure 1: Traffic of an MPEG Video Stream.

visible from the dynamics of the trace. Specifically, the frame pattern can be seen since the I frames tend to be the largest, B the smallest, and P in between.

3.1 The Empirical Envelope

In order to evaluate a deterministic traffic model, we need to compare how accurately a parameterized model's traffic constraint function can represent an actual traffic source. Toward this end, we define the empirical envelope $\mathcal{E}$ as the most accurate traffic constraint function for an arrival function A. $\mathcal{E}$ is given as:

$$\mathcal{E}(\tau) = \max_t A[t, t+\tau] \tag{2}$$

Clearly, from the definition, $\mathcal{E}(\tau)$ is a time-invariant, tight bound on arrivals in any interval of length τ. Note that the tightness of $\mathcal{E}$ implies that any traffic constraint function A^* for an arbitrary traffic model satisfies $A^*(t) \geq \mathcal{E}(t)$ for all t. Since the admission control tests in Tables 2 and 3 are based on the traffic constraint functions A^*, using the most accurate constraint function $\mathcal{E}$ in these tests results in the highest achievable network utilization. That is, from the tests of Tables 2 and 3, it follows that the tightest delay conditions are obtained if A^* is selected such that $A^* \equiv \mathcal{E}$. Thus, the empirical envelope can be used as a benchmark for comparing deterministic traffic models.

Next we describe how to find the empirical envelope for a given video trace. For the traces analyzed here, the inter-frame time is fixed at $r = \frac{1}{30}$ seconds and each video fame is fragmented into 48 byte ATM cells which are transmitted at equally-spaced intervals over the frame time r. If the sequence of frame sizes of video source j is given by $\{f_1, f_2, \ldots, f_N\}$, and if the transmission of this sequence starts at time 0, [2] the

[2] Note that this assumption is made without loss of generality.

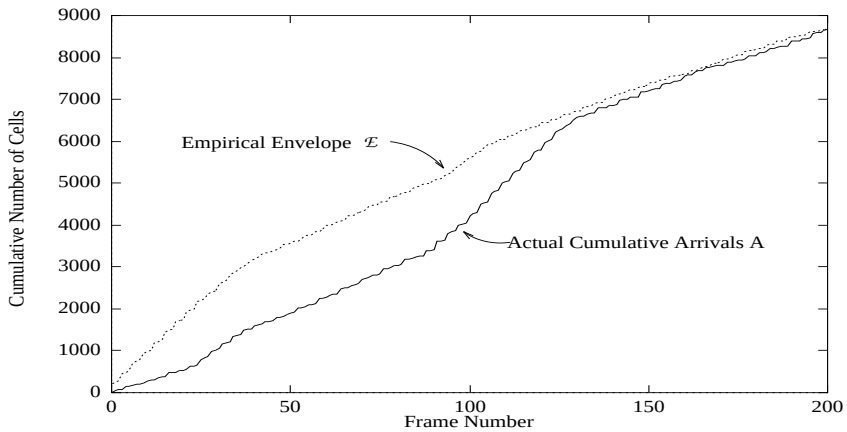

Figure 2: The accumulated traffic $A[0,t]$ of the MPEG stream from Figure 1 and its empirical envelope $\mathcal{E}(t)$.

arrival function A which describes the cumulative number of cell arrivals up to time t is given as follows:

$$A[0,t] = \sum_{i=1}^{\lfloor \frac{t}{r} \rfloor} f_i + \left(\frac{t}{r} - \left\lfloor \frac{t}{r} \right\rfloor \right) f_{\lceil \frac{t}{r} \rceil} \qquad \text{for } 0 \leq t \leq N \cdot r \tag{3}$$

The first term on the right-hand side of equation (3) denotes the number of cells from frames that are fully transmitted at time t, and the second term gives the number of transmitted cells from the frame that is being transmitted at time t. By combining equations (2) and (3), we construct the empirical envelope.

Alternatively, $\mathcal{E}$ can be obtained directly from the sequence of frames $\{f_1, f_2, \ldots, f_N\}$ by first calculating:

$$\mathcal{E}(i \cdot r) = \max_{0<k<N-i} \sum_{j=k}^{k+i} f_j \qquad \text{for } i = 1, 2, \ldots, N \tag{4}$$

Then the values of the empirical envelope at times that are not multiples of the frame time are obtained by spacing the cells in $\mathcal{E}((i+1) \cdot r) - \mathcal{E}(i \cdot r)$ evenly over the frame time $[i \cdot r, (i+1) \cdot r]$. In Figure 2 we illustrate the arrival function A and the empirical envelope $\mathcal{E}$ for the MPEG video trace shown in Figure 1.

3.2 Obtaining Model Parameters from the Empirical Envelope

Here we present an algorithm that uses the empirical envelope to construct a traffic constraint function $A^*(t)$ that conforms to the $(\vec{\sigma}, \vec{\rho})$ traffic model. Recall that the $(\vec{\sigma}, \vec{\rho})$ traffic model is defined by a set of parameter pairs $\{ (\sigma_i, \rho_i) \mid i = 1, 2, \ldots n \}$ with a traffic constraint function given by $A^*(t) = \min_i \{\sigma_i + \rho_i t\}$. Note that this traffic constraint function $A^*(t)$ is guaranteed to be concave and can thus be policed with an N-level leaky bucket [3, 10].

We now present an algorithm that approximates the empirical envelope $\mathcal{E}$ by the $(\vec{\sigma}, \vec{\rho})$ model. For the algorithm we assume that the values of the empirical envelope are available up to some time limit τ

Input: An empirical envelope $\mathcal{E}$ and a time τ.

Output: A set of leaky bucket parameters.

1. **Procedure** Find_Parameters ($\mathcal{E}$, τ)
2. $n = 0$
3. **While** $\tau > 0$ **Do**
4. $n = n + 1$
5. $\sigma_n = \max_{0 \le t < \tau} \left\{ \frac{1}{\tau - t}(\tau \mathcal{E}(t) - t\mathcal{E}(\tau)) \right\}$
6. $\rho_n = \frac{\mathcal{E}(\tau) - \sigma_n}{\tau}$
7. **Output** (σ_n, ρ_n)
8. $\tau = \min\{\, t \mid \sigma_n + \rho_n t = \mathcal{E}(t)\}$
9. **End While**
10. **End Procedure**

Figure 3: An algorithm to determine the parameters of the $(\vec{\sigma}, \vec{\rho})$ traffic model.

where τ is greater than the maximum length of the busy period. The complete algorithm is shown in Figure 3. It determines the number of (σ_i, ρ_i) pairs that are needed to approximate the envelope and selects the parameter values for each (σ_i, ρ_i) pair. The parameter selection by the algorithm proceeds as follows. For the given time τ, a parameter pair (σ_i, ρ_i) and a new time value $0 \le \tau' < \tau$ are selected such that $\mathcal{E}(\tau) = \sigma_i + \rho_i \tau, \mathcal{E}(t) \le \sigma_i + \rho_i t$ for all $0 \le t \le \tau$, and $\mathcal{E}(\tau') = \sigma_i + \rho_i \tau'$. This procedure is repeated with a newly calculated time value τ' as long as τ' is positive.

We illustrate the operations of the algorithm from Figure 3 with a simple example that results in the calculation of two (σ_i, ρ_i) pairs. As shown in Figure 4(a), we assume that the empirical envelope $\mathcal{E}$ for the sequence and some interval $[0, t_1]$ are given. Figure 4(b) depicts the selection of parameters for the pair (σ_1, ρ_1) as calculated in the first iteration of the algorithm. The parameters σ_1 and ρ_1 and the new time value t_2 are chosen such that the curve $\sigma_1 + \rho_1 t$ is never below the empirical envelope and identical to the envelope at times $t = t_1$ and $t = t_2$. Also shown in Figure 4(b) is the selection of time t_2 which determines the starting time for the next iteration of the algorithm.

Figure 4(c) illustrates the determination of the pair (σ_2, ρ_2) which is done analogously to the selection of (σ_1, ρ_1) in Figure 4(b). Note that the algorithm terminates after the calculation of the second pair of parameters, i.e., $t_3 = 0$. The final result is shown in Figure 4(d) where we show the traffic constraint function produced by the algorithm as a bold curve. In this example, two parameter pairs (σ_1, ρ_1) and (σ_2, ρ_2) determine the traffic constraint function, $A^*(t) = \min\{\sigma_1 + \rho_1 t, \sigma_2 + \rho_2 t\}$.

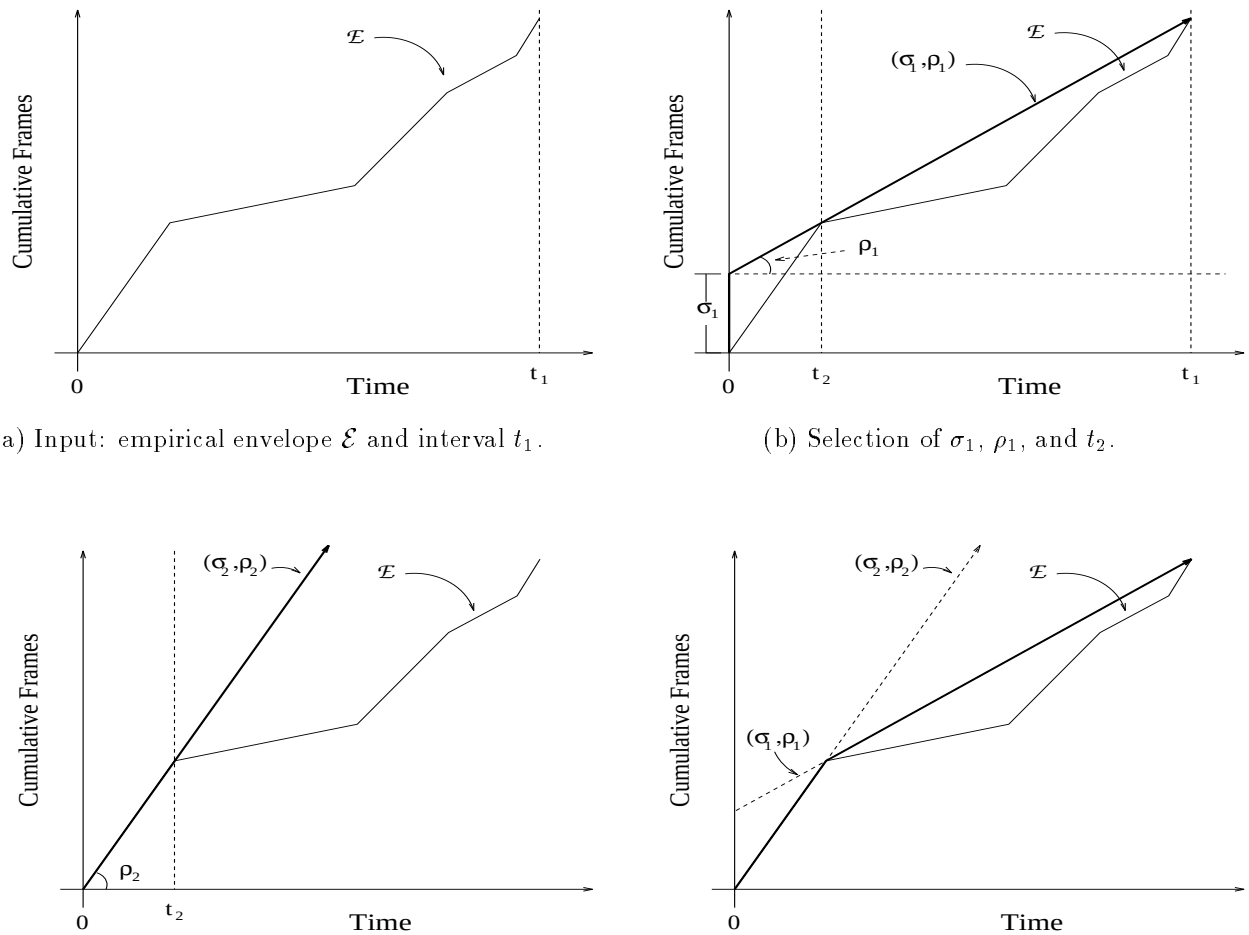

(a) Input: empirical envelope $\mathcal{E}$ and interval t_1.

(b) Selection of σ_1, ρ_1, and t_2.

(c) Selection of σ_2 and ρ_2. (Note that $\sigma_2 = 0$.)

(d) Traffic constraint function.

Figure 4: Parameter selection by the algorithm in Figure 3.

4 Evaluation of Design Choices for Deterministic Services

In this section, we use several actual traces of MPEG-compressed video to empirically evaluate the fundamental limits and tradeoffs of deterministic services as described in the previous sections. We present experiments that show the impact of the traffic constraint function, packet service discipline, and admission control test on the achievable network utilization of a deterministic service.

First, the theoretical limits on the efficiency of a deterministic network service are illustrated by combining the tightest traffic constraint function with an exact admission control test for the best packet scheduler. This is accomplished by using the empirical envelope $\mathcal{E}$ as the traffic constraint function, and an EDF scheduler with exact admission control tests. With these utilization limits as a benchmark we explore the achievable utilization of networks that use more practical traffic models, packet schedulers, and admission control tests.

In all experiments, we assume a multiplexer that operates at 45 Mbps, corresponding to the transmission rate of a T3 line. MPEG-compressed video sequences are used as traffic sources, where each video frame is fragmented into a number of 53-byte ATM cells with a payload of 48 bytes each. The cells for a frame are assumed to be spaced evenly over the frame time in accordance with the cell arrival function as given in equation (3).

Two video sequences are used in the experiments: the entertainment film *The Princess Bride* ("*Bride*") and a series of advertisements ("*Advertisements*"). The two sequences were captured and encoded by different MPEG coders. *Bride* is a 90 minute video sequence with a frame size of 320x240 pixels and a frame pattern *IBBPBBPBBPBBPBB*. *Advertisements* is about 10 minutes long with a smaller frame size of 160x120 pixels and a frame pattern *IBBPBB*. These video sequences were selected because they have very different characteristics in terms of resolution, frame pattern, and the frequency of scene changes.

4.1 Maximum Achievable Utilization of a Deterministic Service

In the first experiment, we show the maximum link utilization that can be achieved for a deterministic service. We consider the exact admission control test for an EDF scheduler from Table 2 and empirical envelopes $\mathcal{E}$ as traffic constraint functions. Note that this scenario corresponds to using the tightest traffic constraint function in combination with the most accurate admission control test for the best packet scheduler. Here we assume that all sources are homogeneous and have identical delay requirements.[3]

Figures 5(a) and 5(b) illustrate the maximum number of *Bride* and *Advertisements* connections, respectively, that can be admitted in a deterministic service. In each graph, the number of connections that can be simultaneously supported is plotted as a function of the delay bound of those connections. The general

[3] Note that all packet schedulers produce the same schedule if sources are homogeneous. We consider heterogeneous sources in Sections 4.3–4.5.

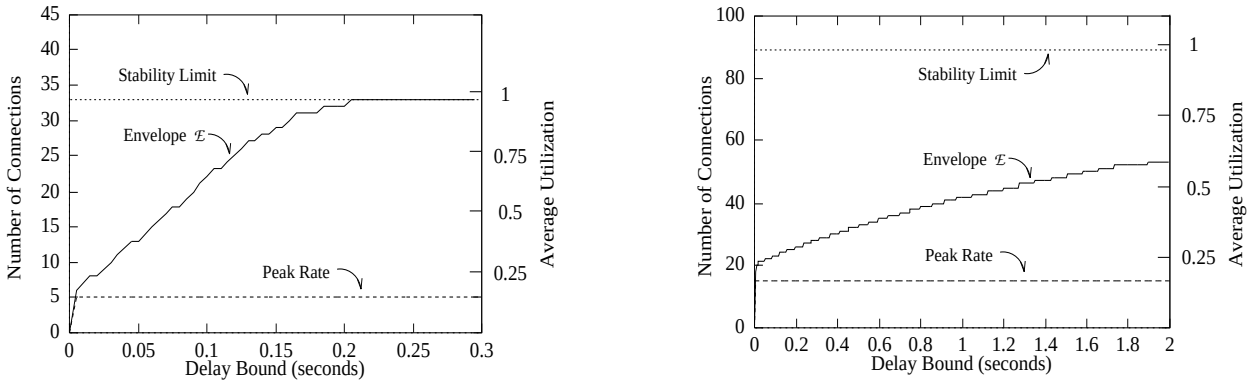

(a) *Bride*.

(b) *Advertisements*.

Figure 5: Fundamental limits of a deterministic network service.

trend of these curves is that as the delay bound increases, more connections can be established. In both figures we also include two reference cases: the curve labeled *Peak Rate*[4] indicates the maximum number of connections that can be accepted if admission control tests are based on peak rate allocation; the curve labeled *Stability Limit* shows the maximum number of connections that can be accepted if packets do not have deadlines, i.e., when the delay bounds are set to infinity. Note that in the case where packets do not have deadlines, the number of connections accepted is limited only by the average traffic rate over the entire video sequence.

One important observation is that reasonably high average utilization for a deterministic service can be achieved even when the delay bound is fairly small. In the case of *Bride* (Figure 5(a)), an average utilization of 38.3% can be achieved with a delay bound of 50 msec. With a delay bound of 200 msec, the link can be utilized up to its stability limit of 33 connections which results in an average utilization of 97.3%. These utilizations are remarkably high when compared to the 14.7% utilization achieved with a peak rate allocation scheme.

In Figure 5(b), notice that the achievable utilization for a given delay bound is considerably lower for the *Advertisements* sequence than for the *Bride* sequence. This difference in utilization can be explained by both the different frame pattern of the *Advertisements* sequence and its additional burstiness caused most likely by its fast scene changes and use of complex graphics. Regardless, admission control tests that use the empirical envelope still provide a significant utilization improvement over the 18% utilization achievable with admission control tests that use a peak-rate allocation scheme. For example, for a delay bound of 500 msec, an average utilization of 37% is achieved.

[4] In this case, the peak rate is the maximum frame size times the frame rate.

4.2 Tradeoffs in Number of Traffic Parameters

Here we explore the utilization tradeoffs in a deterministic service when practical deterministic traffic models are used rather than the empirical envelope. Using the exact admission control test for an EDF scheduler, we use the $(\vec{\sigma}, \vec{\rho})$ model to characterize the traffic and the empirical envelope $\mathcal{E}$ as a benchmark for comparison. Obtaining the values of the (σ_i, ρ_i) pairs from the algorithm of Section 3.2, we compare the utilization achieved with the $(\vec{\sigma}, \vec{\rho})$ model with that of the benchmark for various numbers of (σ_i, ρ_i) pairs. We again consider homogeneous sources with identical delay bounds.

Figures 6(a) and 6(b) depict the respective number of *Bride* and *Advertisements* connections that can be supported using different traffic constraint functions for different delay bounds. When a number j of (σ_i, ρ_i) pairs is specified, the j pairs used are the those with the smallest values of σ_i, e.g., the single pair (σ_1, ρ_1) would be the one with the smallest value of σ_i.

The first observation from Figure 6 is that as more (σ_i, ρ_i) pairs are used, more connections can be accepted. The reason for this is that a traffic constraint function A^* defined with more parameters can better approximate the empirical envelope $\mathcal{E}$, allowing the acceptance of more connections. However, it should also be noted that the benefit of adding additional (σ_i, ρ_i) pairs decays quickly. From Figure 6(a), we see that a characterization of *Bride* that uses two (σ_i, ρ_i) pairs results in the same network utilization as the empirical envelope function when the delay bound is less than 130 msec. With five pairs, the same network utilization as for the empirical envelope can be achieved for all values of the delay bound, i.e., a sixth pair introduces no additional benefit. In contrast, as shown in Figure 6(b), a larger number of (σ_i, ρ_i) pairs is required to approximate the empirical envelope for the *Advertisements* sequence. The larger number of (σ_i, ρ_i) pairs required by the *Advertisements* sequence can be attributed to its additional burstiness and complexity as compared to the *Bride* sequence.

4.3 Tradeoffs of Packet Scheduling

In this section we compare the performance of the three different packet schedulers discussed in Section 2.2, namely EDF, SP and FCFS. We consider the exact admission control tests from Table 2 and use the empirical envelopes as traffic constraint functions. The experiments in this section assume two types of connections: those carrying the *Bride* sequence and those carrying the *Advertisements* sequence. All connections of the same type have identical delay requirements.

For each of these schedulers, Figure 7 shows the maximum number of *Advertisements* and *Bride* connections that can be simultaneously supported by the network. The number of *Advertisements* connections is plotted as a function of the number of *Bride* connections. As indicated in the legend, each curve in the figures corresponds to a pair of delay bounds for the two connection types. As in the previous experiment,

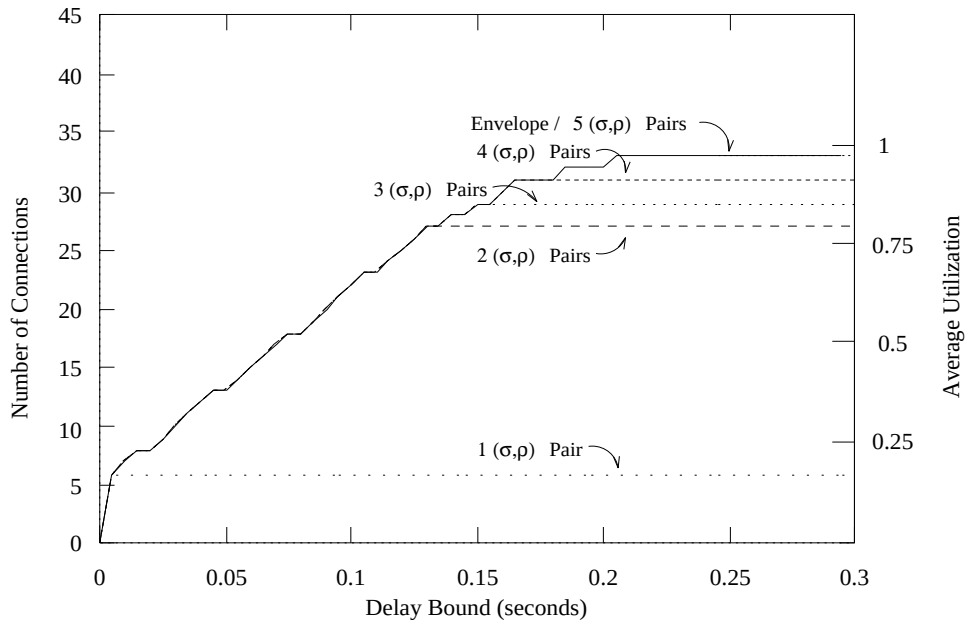

(a) *Bride.*

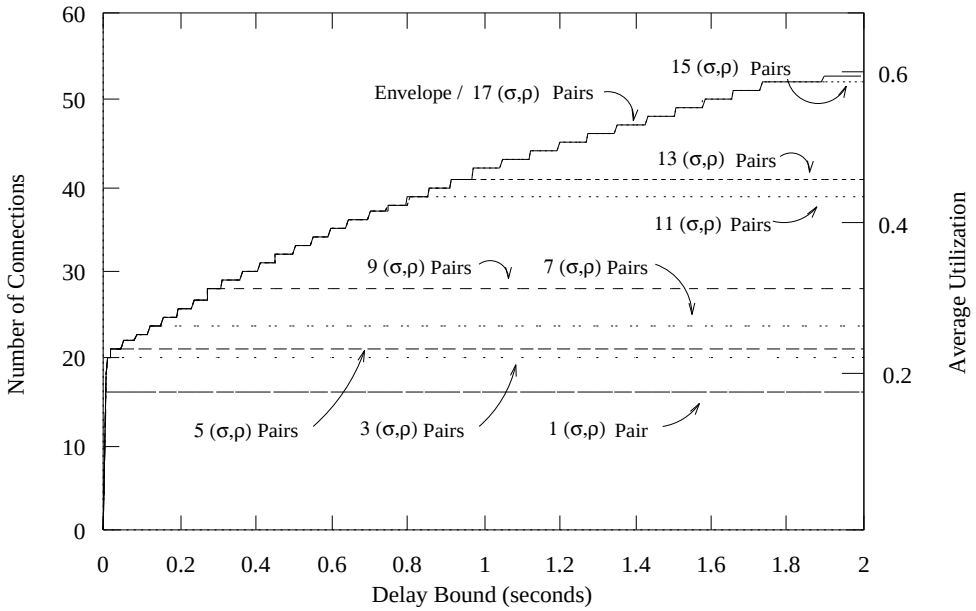

(b) *Advertisements.*

Figure 6: A comparison of different traffic constraint functions.

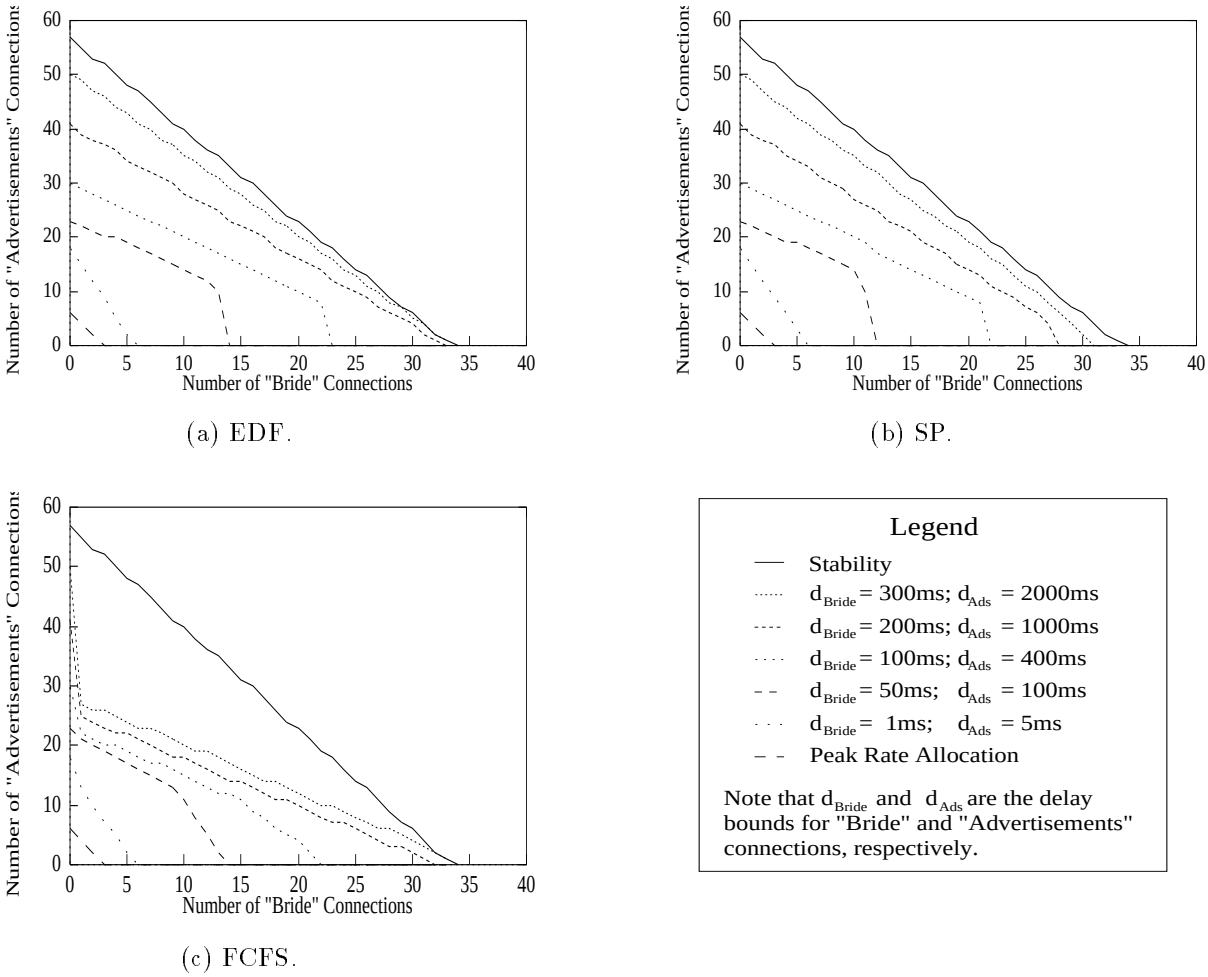

(a) EDF.

(b) SP.

(c) FCFS.

Figure 7: A comparison of packet schedulers.

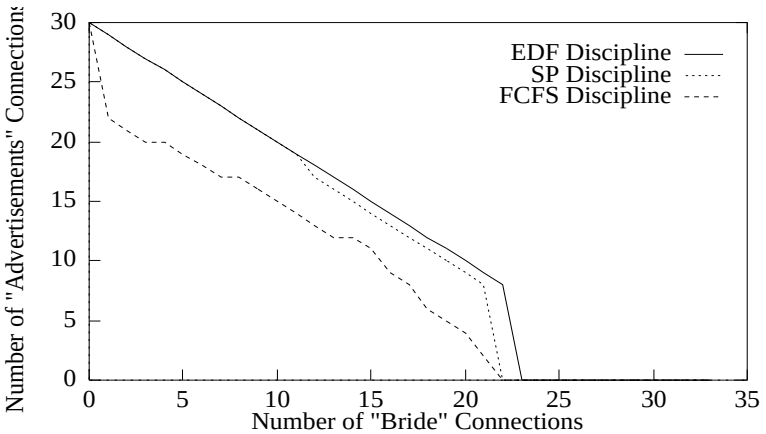

Figure 8: A comparison of packet schedulers (Here $d_{\text{Bride}} = 100$ msec and $d_{\text{Ads}} = 400$ msec).

the *Stability* curve and *Peak Rate Allocation* curve are used as reference cases. The *Stability* curve indicates the maximum number of connections that can be admitted if the number of connections accepted is limited only by the average traffic rate, and the *Peak Rate Allocation* curve depicts the number of connections that can be admitted if admission control tests are based on peak rate allocation. In each experiment, the same delay bound combinations are used so that the relative performance of the disciplines can be evaluated by comparing the contours of each plot.

By comparing Figure 7(a) and Figure 7(b) it can be seen that the contour plots are similar. Thus, for our experiment that requires only a limited number of different delay bounds, an SP scheduler can support connections almost as well as an EDF scheduler. Note from Figure 7(c) the advantages of EDF and SP disciplines over the FCFS discipline: since FCFS disciplines can in effect only guarantee a single delay bound, the number of *Advertisement* connections is limited by the tight delay bound of the *Bride* traffic.

In order to compare the three packet schedulers directly, we extract a single curve, i.e., the curve corresponding to $d_{\text{Bride}} = 100$ msec and $d_{\text{Ads}} = 400$ msec, from each of Figures 7(a), 7(b), and 7(c) and plot them together in Figure 8. For the case considered here, the SP discipline results in nearly as many admissible connections as the EDF discipline, while the FCFS discipline admits considerably fewer connections than both EDF and SP. Thus, in this experiment where only a limited number of different delay bounds are supported, the SP scheduler shows utilization limits similar to those of the EDF scheduler.

4.4 Tradeoffs of Admission Control Tests

This experiment illustrates for the SP packet scheduler how the accuracy of admission control tests impacts the utilization limits. As before we assume the empirical envelope $\mathcal{E}$ as the traffic constraint functions. We consider the following admission control tests for the SP scheduler: the exact admission control test *SP Exact*

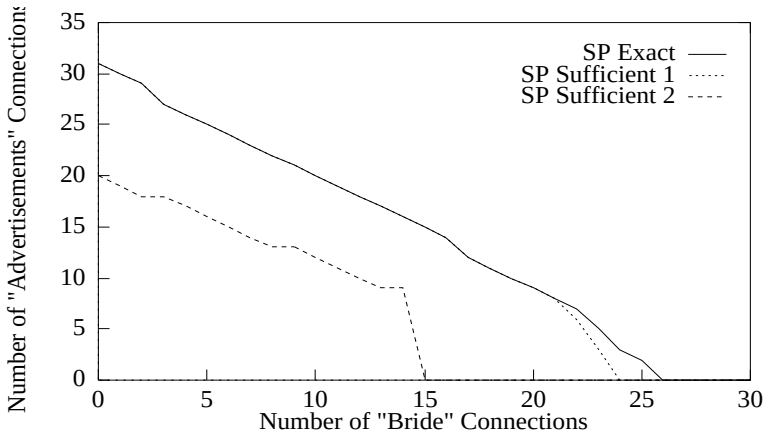

Figure 9: A comparison of admission control tests (Here, $d_{\text{Bride}} = 150$ msec and $d_{\text{Ads}} = 500$ msec).

from Table 2, and the two sufficient admission control tests from Table 3, *SP Sufficient 1* and *SP Sufficient 2*.

These admission control tests vary in their computational complexity. If there are P priority levels and the frame sequence is of length N, then the computational complexities of the admission control tests are as follows: the admission control test *SP Exact* has complexity $O(PNd_p)$; the admission control test *SP Sufficient 1* has complexity $O(PN)$; the admission control test *SP Sufficient 2* has complexity $O(P)$.

Figure 9 depicts the number of *Advertisements* and *Bride* connections that can be simultaneously supported for each of the three admission control tests. Here we only consider the case in which $d_{\text{Bride}} = 150$ msec and $d_{\text{Ads}} = 500$ msec. For this experiment, the admission control test *SP Sufficient 1* achieves a utilization that is almost as high as the slightly more complex admission control test *SP Exact*. However, with a very simple admission control test such as *SP Sufficient 2*, the achievable utilization is seen to be substantially lower. Since the admission control test *SP Sufficient 1* achieves a utilization that is very close to that of the admission control test *SP Exact*, it may be worthwhile to explore sufficient conditions. However the achievable utilization of extremely simple admission control tests such as *SP Sufficient 2* can be low.

4.5 Summary of Tradeoffs

In the previous three experiments, we analyzed the tradeoffs involved in selecting network components in isolation. Here we evaluate the utilization of a system that *simultaneously* considers practical choices for the traffic constraint functions, packet schedulers, and admission control tests, and compare this to the benchmark case.

Figure 10 depicts the maximum number of admissible connections for these two cases. Specifically, for *Advertisement* connections requiring a delay bound of 30 msec and *Bride* connections requiring a delay bound

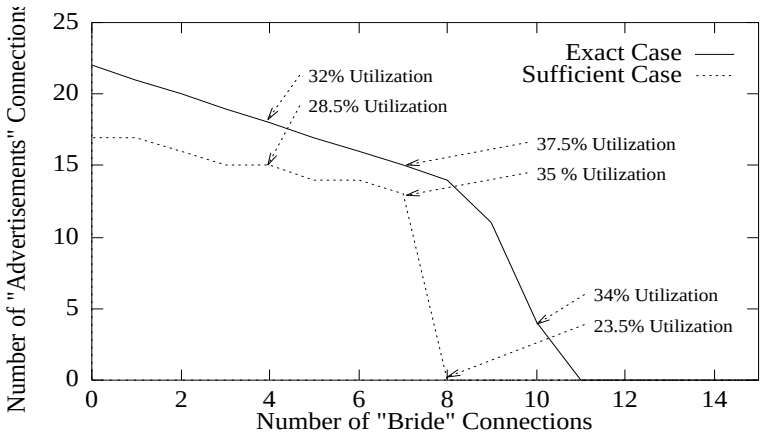

Figure 10: Achievable utilization considering all tradeoffs simultaneously (Here, $d_{\text{Bride}} = 30$ msec and $d_{\text{Ads}} = 50$ msec).

of 50 msec, admission control tests are performed for the following two cases. The first case, shown by the upper curve, represents the case of exact analysis, EDF scheduling, and the empirical envelope for the source characterization. The lower curve is subjected to the practical constraints of SP scheduling, SP Sufficient 1 admission control tests, and a traffic characterization consisting of 3 (σ, ρ) pairs. Several points on this schedulable region are labeled with their corresponding average utilization.

The figure demonstrates that reasonably high network utilizations of 20% to 35% can be achieved for deterministic service to VBR video even when taking into account many practical considerations such as the need for simple schedulers, policing mechanisms, traffic models, and admission control tests.

5 Conclusions

The conventional wisdom has been that providing deterministic performance guarantees to "bursty" video sources results in a low utilization of network resources. The reason for such pessimism is that deterministic services ensure that *all* packets of *every* connection are guaranteed to meet their performance requirements, even in worst case situations. We have studied the fundamental limits and tradeoffs of using deterministic services to support delay-sensitive VBR video in packet networks. We have identified the main factors that influence the utilization of a network with deterministic services as (1) the deterministic traffic model used to characterize the source traffic, (2) the packet scheduling disciplines at the network switches, and (3) the accuracy of the admission control tests. We have shown the degree to which these factors impact the efficiency of a deterministic VBR video service. Our results can be summarized as follows.

- We used MPEG video traces to investigate the best possible utilization achievable in a deterministic VBR video service. To this goal, we employed the most accurate traffic characterization, the tightest admission control tests, and the best possible packet scheduler. One MPEG compressed video trace containing an entertainment movie showed that an average link utilization of 38.3% can be achieved by selecting a deterministic delay bound of 50 msec. A link utilization of 97.3% was achieved if the delay bound was set to 200 msec. However, a different MPEG compressed video sequence (*Advertisements*) resulted in a significantly lower link utilization, e.g., 18% for small delay bounds and 25% for a delay bound of 100 msec.

We considered tradeoffs that must be considered in a practical network environment, such as, constraints on the complexity of the policing functions, the sophistication of the scheduling disciplines at the network switches, and the accuracy of the admission control tests.

- We considered a traffic model that can be enforced with multiple leaky buckets and varied the number of model parameters to see the effect on achievable network utilization. Our experiments indicate that within a practical range of small delay bounds, considering models that are policed with more than several leaky buckets has little advantage.

- We considered three packet schedulers: A FCFS scheduler which can enforce only a single deterministic delay bound for all connections, a Static Priority (SP) scheduler which enforces one delay bound for each priority level, and an EDF scheduler which can enforce a different delay bound for each connection. Our experimental data indicates that in networks which support only a limited number of different delay bounds, Static Priority schedulers show utilization limits that are similar to those of EDF schedulers. Thus, if the number of supported delay requirements is small, the complexity involved in the implementation of sophisticated scheduling methods may not result in a significant utilization gain.

- The calculation of exact admission control tests may consume a considerable amount of time, causing additional delays during the connection establishment process. As an example, we considered admission control tests of Static Priority schedulers. Here sufficient tests for admission control are considerably simpler to calculate than the necessary and sufficient tests. In an empirical evaluation we demonstrated that the sufficient admission control test of [21] resulted in significantly lower utilization than for the necessary and sufficient test of [3, 12]; however, a different sufficient admission control test from [12] resulted in a utilization that closely approximated that of the ideal case.

References

[1] K. Bala, I. Cidon, and K. Sohraby. Congestion Control for High Speed Packet Switched Networks. In *Proc. IEEE Infocom '90*, pages 520–526, May 1990.

[2] C.-S. Chang. Stability, Queue Length and Delay, Part I: Deterministic Queueing Networks. Technical Report RC 17708, IBM Research Division, Yorktown Heights, February 1992.

[3] R. L. Cruz. A Calculus for Network Delay, Part I: Network Elements in Isolation. *IEEE Transactions on Information Theory*, 37(1):114–131, January 1991.

[4] D. Ferrari and D. C. 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] D. Le Gall. MPEG: A Video Compression Standard for Multimedia Applications. *Communications of the ACM*, 34(4):305–313, April 1991.

[6] M. W. Garrett and W. Willinger. Analysis, modeling and generation of self-similar VBR video traffic. In *ACM Sigcomm 1994*, London, UK, August 1994.

[7] R. Grunenfelder, J.P. Cosmas, S. Manthorpe, and A. Odinma-Okafor. Characterization of Video Codecs as Autoregressive Moving Average Processes abd Related Queueing System Performance. *IEEE Journal on Selected Areas in Communications*, 9(3):284–293, June 1991.

[8] D.P. Hong and T. Suda. Congestion Control and Prevention in ATM Networks. *IEEE Network*, 5(7):10–16, July 1991.

[9] J. M. Hyman, A. A. Lazar, and G. Pacifici. Joint Scheduling and Admission Control for ATS-based Switching Nodes. In *Proc. Sigcomm '92*, pages 223–234, August 1992.

[10] E. Knightly and H. Zhang. Traffic characterization and switch utilization using deterministic bounding interval dependent traffic models. Technical Report TR-94-047, International Computer Science Institute, Berkeley, California, August 1994.

[11] J. Liebeherr and D. E. Wrege. Design and Analysis of a High-Performance Packet Multiplexer for Multiservice Networks with Delay Guarantees. Technical Report CS-94-30, University of Virginia, July 1994.

[12] J. Liebeherr, D. E. Wrege, and Domenico Ferrari. Exact Admission Control in Networks with Bounded Delay Services. Technical Report CS-94-29, University of Virginia, July 1994.

[13] B. Maglaris, D. Anastassiou, P. Sen, G. Karlsson, and J.D. Robbins. Performance Models of Statistical Multiplexing in Packet Video Communications. *IEEE Transactions on Communications*, 36(7):834–844, July 1988.

[14] I. Nikolaidis and I.F. Akyildiz. Source Characterization and Statistical Multiplexing in ATM Networks, September 1992.

[15] G. Ramamurthy and B. Sengupta. Modeling and Analysis of a Variable Bit Rate Video Multiplexer. In *Proc. IEEE Infocom '92*, pages 817–823, 1992.

[16] P. Sen, B. Maglaris, N. Rikli, and D. Anastassiou. Models for Packet Switching of Variable-Bit-Rate Video Sources. *IEEE Journal on Selected Areas in Communications*, 7(5):865–869, June 1989.

[17] P. Skelly, S. Dixit, and M. Schwartz. A Histogram-Based Model for Video Traffic Behavior in an ATM Network Node with an Application to Congestion Control. In *Proc. IEEE Infocom '92*, pages 95–104, 1992.

[18] J. S. Turner. New Directions in Communications (or Which Way to the Information Age?). *IEEE Communications Magazine*, 25(8):8–15, October 1986.

[19] F. Yegenoglu, B. Jabbari, and Y.Q. Zhang. Modeling of Motion Classified VBR Video Codecs. In *Proc. IEEE Infocom '92*, pages 105–109, 1992.

[20] H. Zhang and D. Ferrari. Rate-Controlled Static-Priority Queueing. In *Proc. IEEE Infocom '93*, pages 227–236, April 1993.

[21] H. Zhang and D. Ferrari. Improving Utilization for Deterministic Service in Multimedia Communication. In *1994 International Conference on Multimedia Computing and Systems*, pages 295–304, Boston, MA, May 1994.