

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

Edward W. Knightly* Dallas E. Wrege[†] Jörg Liebeherr[†] Hui Zhang[‡]

* EECS Department, U.C. Berkeley
and Sandia National Laboratories

[†] Department of Computer Science
University of Virginia
Charlottesville, VA 22903

[‡] School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

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, in many cases, a deterministic service can be provided to video traffic while maintaining a reasonable level of network utilization. We first consider an ideal network environment that employs the most accurate deterministic, time-invariant 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.

¹The work of D. Wrege and J. Liebeherr was supported in part by NSF under Grant No. NCR-9309224.

1 Introduction

Future packet-switching integrated services networks must support applications with diverse traffic characteristics and performance requirements. Of the many traffic classes in integrated services networks, delay-sensitive Variable Bit Rate (VBR) video traffic poses a unique challenge. Since the required service is delay-sensitive, the network must support a resource reservation scheme [3] 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 one hand, if resources are reserved according to the average rates of the VBR video sources, unacceptable delays or packet losses may result when the sources are transmitting near their peak rates. On the other hand, if resource reservations are based on sources' peak rates, the network may be under-utilized most of the time.

In this study, we investigate utilization limits achievable in packet-switching networks with *deterministic* service guarantees to delay-sensitive VBR video traffic. A deterministic service ensures that no packets are dropped or delayed beyond their guaranteed delay bound. Such a service requires *admission control* tests to determine if the network has sufficient resources to support a new connection without degrading the service of existing connections [3, 6]. The 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 admitted connections, the network performs *traffic policing*, that is, it ensures that the admitted connections adhere to the traffic characterization given to the network for the admission control tests. If a connection's traffic is policed by the network, the traffic in excess of its specified characterization may not be 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 cap-

tures the dynamics of the source. A rich set of literature exists on characterizing VBR video traffic by a *stochastic* process using models such as Markov-modulated, Auto Regressive, TES, and Self-similar (see [4, 11] and the references therein). While such stochastic models of a source have the advantage that they may be used to potentially 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 [11]. Finally, since stochastic approaches to VBR traffic characterization provide only statistical guarantees, they cannot be applied to deterministic network services [7, 8, 9, 15].

For such reasons, in this work we consider only *deterministic* traffic models that characterize VBR traffic in terms of a worst-case description. It is a common belief that such a worst-case approach necessitates allocation of network resources according to sources' peak rates, necessarily leading to low network utilization for bursty traffic. However, several recent works show that this conventional wisdom is not correct and that peak-rate allocation is not required even for providing deterministic service [1, 2, 7, 9, 15].

In this study, we analyze the fundamental limits 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 [8, 9]. We present an optimal characterization of VBR video traffic for a deterministic service using what we term *empirical envelopes*. Empirical envelopes represent the most accurate, time-invariant, deterministic characterization of a traffic source. Using this optimal characterization, together with Earliest-Deadline-First packet scheduling and the exact admission control tests from [8, 9], we provide insights into the highest achievable utilization for 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 and First-Come-First-Served, on the achievable 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 §2, we describe the network components necessary to offer

a service with deterministic delay bounds. In §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 §4 we use traces of MPEG-encoded VBR video traffic to empirically evaluate the limits of the deterministic approach as well as the abovementioned tradeoffs.

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 certain 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 multiplexers, and (3) the accuracy of the admission control tests. 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

A traffic model for a deterministic service 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 $A[\tau, \tau + t]$ denotes the traffic arrivals in the time interval $[\tau, \tau + t]$, an upper bound on A can be given by a function A^* if for all times $\tau \geq 0$ and all $t \geq 0$ the following holds [1, 2]:

$$A[\tau, \tau + t] \leq A^*(t) \quad (1)$$

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

In practice, a source specifies its traffic characterization with a parameterized model. The parameterized deterministic traffic model then defines a traffic constraint

¹In [2, 15], the traffic constraint function is denoted by $b(\cdot)$.

Traffic Model	Traffic Constraint Function
(σ, ρ)	$A^*(t) = \sigma + \rho t$
$(\vec{\sigma}, \vec{\rho})$	$A^*(t) = \min_{1 \leq i \leq n} \{\sigma_i + \rho_i t\}$
(x_{min}, x_{ave}, I, s)	$A^*(t) = \lfloor \frac{t}{I} \rfloor \cdot \frac{I \cdot s}{x_{ave}} + \min \left\{ \left[\left(\frac{t}{I} - \lfloor \frac{t}{I} \rfloor \right) \cdot \frac{I}{x_{min}} \right], \frac{I}{x_{ave}} \right\} \cdot s$
D-BIND	$A^*(t) = \begin{cases} tR_1 & \text{for all } t \leq I_1. \\ \frac{I_{j-1}I_j(R_{j-1} - R_j) + t(R_jI_j - R_{j-1}I_{j-1})}{I_j - I_{j-1}} & \text{for all } I_{j-1} < t \leq I_j. \end{cases}$

Table 1: Traffic Models and Traffic Constraint Functions.

function that bounds the source. For example, in the (x_{min}, x_{ave}, I, s) model [3], 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 [2] 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 $\min_i \{\sigma_i + \rho_i t\}$. This is similar to the ‘‘burstiness curve’’ characterization in [10] that considers the entire $\sigma(\rho)$ curve. A more general deterministic traffic model called the *Deterministic Bounding Interval Dependent* or *D-BIND* model [7] characterizes traffic by a family of rate-interval pairs $\{(R_j, I_j) | j = 1, 2, \dots, n\}$. The D-BIND model defines an n 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 the 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, $A^*[t, t + \tau] \leq A^*(\tau)$ ($t \geq 0, \tau \geq 0$) [9].

Since the selection of the traffic model determines the shape of a traffic constraint function, 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 accuracy preferred by the

admission control tests and the simplicity required for the implementation of traffic policers. The 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 limited. In [2], 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 [13], 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, a simple implementation of traffic policing mechanisms is ensured.

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 [2, 3, 9, 14].

First-Come-First-Served 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 .

In a Static Priority packet scheduler such as in [14], each connection $j \in \mathcal{N}$ is assigned a priority p with $1 \leq p \leq P$, where a lower 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.

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

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

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

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

With Earliest-Deadline-First scheduling [3], 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.

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.

Tables 2 and 3 present conditions that must be satisfied to pass delay bound tests for FCFS, SP, and EDF packet schedulers. In the tables, the link speed is normalized to 1 and s_k denotes the maximum packet size for connection k . Table 2 shows the necessary and sufficient delay bound tests (“*Exact Tests*”) for each scheduler. Formal derivations of these tests can be found in [2, 9]. Since the exact admission test for an SP scheduler is computationally very complex, Table 3 presents less accurate tests (“*Sufficient Tests*”) that require less computation [9, 15]. Note that each test is formulated in terms of A^* so that it can be applied to the traffic models discussed in §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 the maximum backlog in the transmission queue ($\sum_{j \in \mathcal{N}} A_j[0, t] - t$) and the largest remaining transmission time of any packet that is in transmission at time t ($\max_{k \in \mathcal{N}} s_k$). 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

Here we present the optimal characterization of VBR video traffic for a deterministic service using *empirical envelopes*. We then show how traffic models such as the $(\bar{\sigma}, \bar{\rho})$ model can be viewed in terms of the empirical envelope and we present an algorithm for deriving the $(\bar{\sigma}, \bar{\rho})$ parameters from the envelope.

In this paper, we use traces of video compressed with the MPEG compression algorithm as examples of VBR video traffic. Figure 1 shows a typical segment of a trace of MPEG compressed video obtained from a recorded sequence of advertisements. 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 a past and a future reference frame. 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]. The behavior of the coder is clearly visible from the dynamics of the trace. Specifically, the frame pattern (in this case *IBBPBB*) can be seen since the *I* frames tend to be the largest, *B* the smallest, and *P* in between.

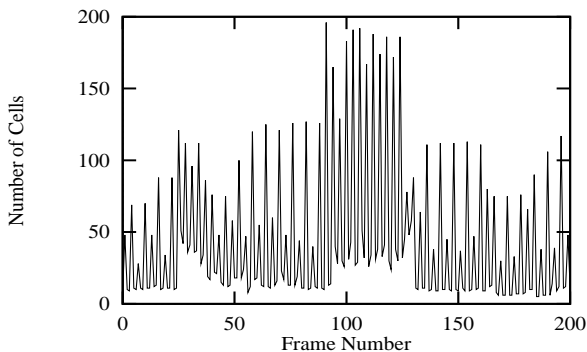


Figure 1: Traffic of an MPEG video stream.

3.1 The Empirical Envelope

As described in §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 . The empirical envelope presented next provides the most accurate traffic constraint function for a given video trace which allows evaluation of the maximum achievable utilization for a deterministic service to VBR video. 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 will be shown to lack important practical properties such as the ability to be efficiently specified or policed, it does provide an optimal benchmark for evaluating deterministic traffic models. Toward this end, we define the empirical envelope \mathcal{E} as the most accurate traffic constraint function for an arrival function A as:

$$\mathcal{E}(t) = \max_{\tau > 0} A[\tau, \tau + t] \quad \forall t > 0 \quad (2)$$

Equation (2) indicates that $\mathcal{E}(t)$ is a time-invariant, tight bound on arrivals in any interval of length t . 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.

The empirical envelope for a given video trace may be obtained in the following manner. For the traces analyzed in this paper, the inter-frame time is fixed at $r = \frac{1}{30}$ seconds and each video frame 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 $\{f_1, f_2, \dots, f_N\}$, and if the transmission of this sequence starts at time 0, the arrival function A which describes the cumulative number of cell arrivals up to time t is given as follows:

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

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 . The empirical envelope may then be constructed by combining Equations (2) and (3).

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

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

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

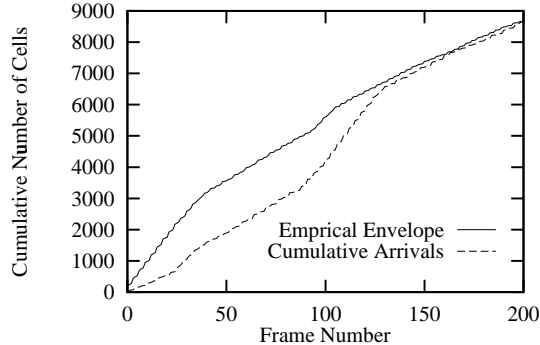


Figure 2: The accumulated traffic $A[0, t]$ of the MPEG stream from Figure 1 and its empirical envelope $\mathcal{E}(t)$.

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 $\{(\sigma_i, \rho_i) \mid i = 1, 2, \dots, n\}$ 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 [2, 7]. The algorithm, shown in Figure 3, assumes that the values of the empirical envelope are available up to some time limit τ where τ is greater than the maximum length of the busy period. 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 $0 \leq \tau' < \tau$ are selected such that $\mathcal{E}(\tau) = \sigma_i + \rho_i \tau$, $\mathcal{E}(t) \leq \sigma_i + \rho_i t$ for all $0 \leq t \leq \tau$, 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. From Equation (2) and the discussions above, it can be seen that the algorithm constructs A^* that is a piece-wise linear concave upper approximation to \mathcal{E} .

Input:	An empirical envelope \mathcal{E} and a time τ .
Output:	A set of leaky bucket parameters.
<ol style="list-style-type: none"> 1. Procedure Find_Parameters (\mathcal{E}, τ) 2. $n = 0$ 3. While $\tau > 0$ Do 4. $n = n + 1$ 5. $\sigma_n = \max_{0 \leq t < \tau} \left\{ \frac{1}{\tau - t} (\tau \mathcal{E}(t) - t \mathcal{E}(\tau)) \right\}$ 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 $(\vec{\sigma}, \vec{\rho})$ parameters.

Figure 4 illustrates the operations of the algorithm with a simple example that results in the calculation of two (σ_i, ρ_i) pairs. Figure 4(a) depicts the selection of parameters for the pair (σ_1, ρ_1) as calculated in the first iteration of the algorithm with the given empirical envelope \mathcal{E} and interval $[0, t_1]$. 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 is identical to the envelope at times t_1 and t_2 . Also shown in Figure 4(a) is the selection of time t_2 which determines the starting time for the next iteration of the algorithm.

Figure 4(b) 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(a). 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(c) 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, $A^*(t) = \min \{\sigma_1 + \rho_1 t, \sigma_2 + \rho_2 t\}$.

Finally, as an illustrative example to emphasize the potential simplicity of approximating \mathcal{E} for VBR video, we note as in [7] that for MPEG video sources, one possible concise and simple-to-obtain parameterization of A^* is to use knowledge of the frame pattern (in this case *IBBPBB*) along with a parameterization of the largest sized *I* frame, *B* frame, and *P* frame. With this alternative “worst-case” characterization, a pessimistic approximation to \mathcal{E} is easily obtained by constructing A^* as a transmission of the largest *I* frame, followed by 2 transmissions of the largest *B* frame and so on. In essence, any A^* that is a sub-additive upper approximation to \mathcal{E} can be used as a traffic constraint function, while, as described in §2.1, the shape of A^* affects its policability, and, as explored in the next section, the tightness of A^* affects the achievable network utilization.

4 Evaluation of Design Choices for Deterministic Services

In this section, we use several traces of MPEG-compressed video to empirically evaluate the fundamental limits and tradeoffs of deterministic services as described in the previous sections. First, the theoretical limits on the efficiency of a deterministic network service are illustrated by combining the tightest traffic constraint function \mathcal{E} , with an exact admission control test for the best packet scheduler, EDF. 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 consider a single multiplexer that operates at 45 Mbps, corresponding to the transmission rate of a T3 line. Note that the end-to-end delay bound is an accumulation of the local delay bounds investigated below. While there is not a consensus on the exact range of acceptable delay bounds for interactive applications, some discussion of this is found in [12].

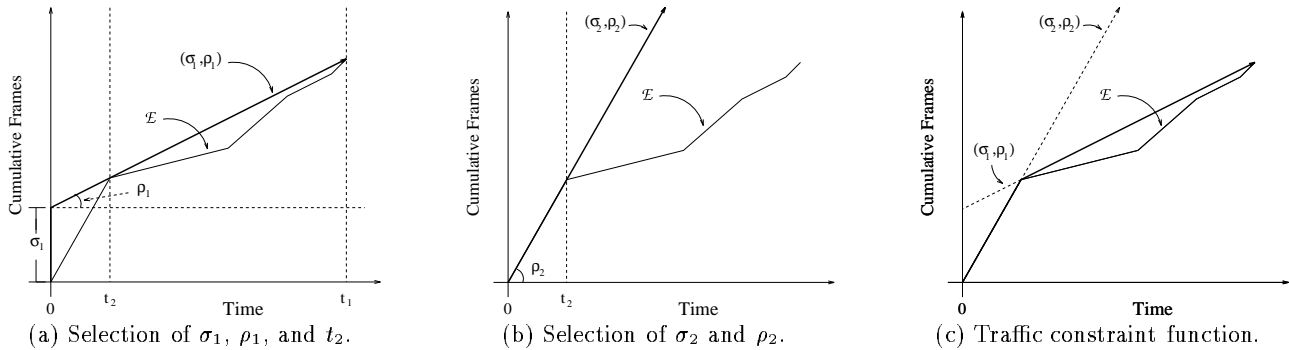


Figure 4: Parameter selection by the algorithm in Figure 3.

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 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 using 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 consider the case of homogeneous sources with identical delay requirements.² The heterogeneous case is considered in §4.3–4.5.

Figures 5(a) and 5(b) illustrate the maximum number of *Bride* and *Advertisements* connections, respectively, that can be admitted for 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. Both figures also include two reference cases: the curve labeled *Peak Rate* (where peak rate is the maximum frame size times the frame rate) 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.

The general trend of these curves is that as the delay bound increases, more connections can be established. One important observation is that a reasonable network utilization for 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 considerably higher than the 14.7% utilization achieved with a peak rate allocation scheme.

Figure 5(b) shows 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 utilization improvement over the 18% utilization achievable with admission control tests that use a peak-rate allocation scheme.

4.2 Tradeoffs in Number of Traffic Parameters

Here we explore the utilization tradeoffs in a deterministic service when using a practical deterministic traffic model 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 to compare the utilization achieved for various numbers of (σ_i, ρ_i) pairs.

Figures 6(a) and 6(b) depict the respective number of homogeneous *Bride* and *Advertisements* connections that can be supported as a function of delay bound for various numbers of (σ_i, ρ_i) pairs. 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 .

²Note that EDF, SP, and FCFS produce the same schedule when all connections have the same local delay bound.

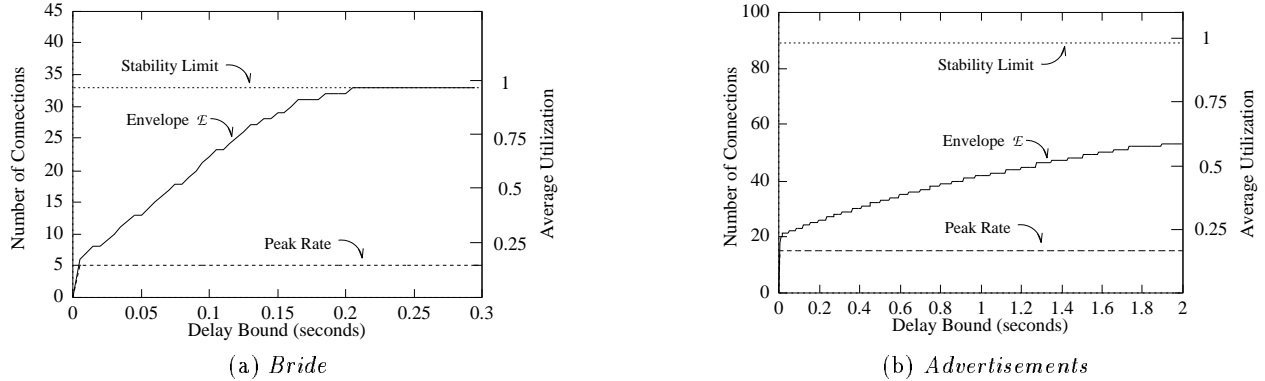


Figure 5: Fundamental limits of a deterministic network service.

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 \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. Figure 6(a) shows 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 can be achieved for all values of the delay bound, i.e., a sixth pair introduces no additional benefit. While it may seem counter-intuitive that full utilization may be achieved considering only the concave hull of the empirical envelope, note that the advantage of considering non-concave constraint functions is only apparent in the heterogeneous case, when other sources can take advantage of the difference between the non-concave constraint function and the concave hull (we defer further discussion of this issue for future work). In contrast to Figure 6(a), 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 in Packet Scheduling

In this section, we compare the performance of the three different packet schedulers discussed in §2.2, namely EDF, SP and FCFS. We use the exact admission control tests from Table 2 with empirical envelopes as traffic constraint functions, and consider 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. Each curve in the figures corresponds to a pair of de-

lay bounds for the *Advertisements* and *Bride* connection types. Thus, the relative performance of the scheduling disciplines can be evaluated by comparing the contours of each plot. The stability and peak rate allocation curves are again provided as reference cases.

By comparing Figure 7(a) and (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. Alternatively, Figure 7(c) shows 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 tighter delay bound of the *Bride* traffic.

4.4 Tradeoffs in Admission Tests

This experiment illustrates the impact of the accuracy of admission control tests on the utilization limits for the SP scheduler. Using the empirical envelope \mathcal{E} for traffic constraint functions, we consider the following admission control tests for the SP scheduler: “*SP Exact*” from Table 2, and “*SP Sufficient 1*” and “*SP Sufficient 2*” from Table 3. These admission control tests vary in their computational complexity. If there are P priority levels and the constraint function is \mathcal{E} as derived from a frame sequence of length N , then the computational complexities of the admission control tests are as follows: the test *SP Exact* has complexity $O(PNd_P)$; the test *SP Sufficient 1* has complexity $O(PN)$; and the test *SP Sufficient 2* has complexity $O(P)$. If the constraint function is based on a parameterized traffic model rather than \mathcal{E} , these complexities may be considerably reduced [9].

Figure 8 depicts the number of *Advertisements* and *Bride* connections that can be simultaneously supported for each of the three admission control tests for 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 simpler admission control test such as *SP Sufficient*

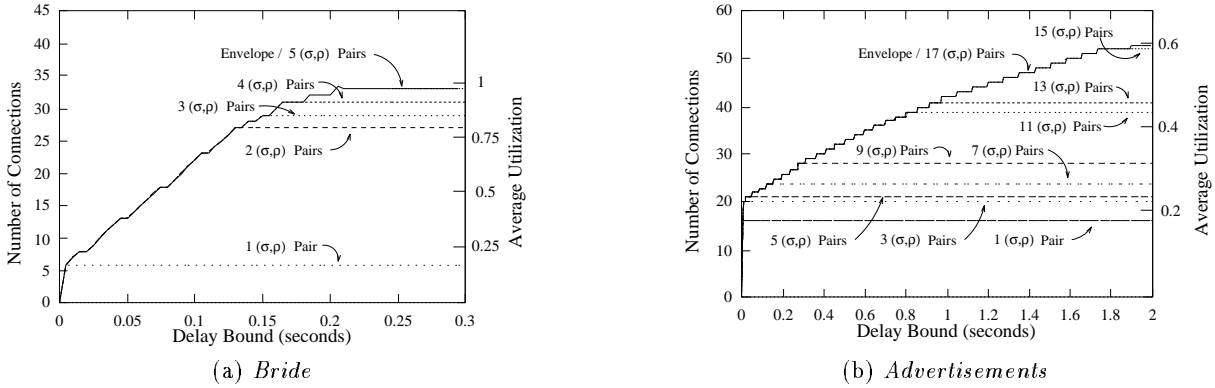


Figure 6: A comparison of different traffic constraint functions.

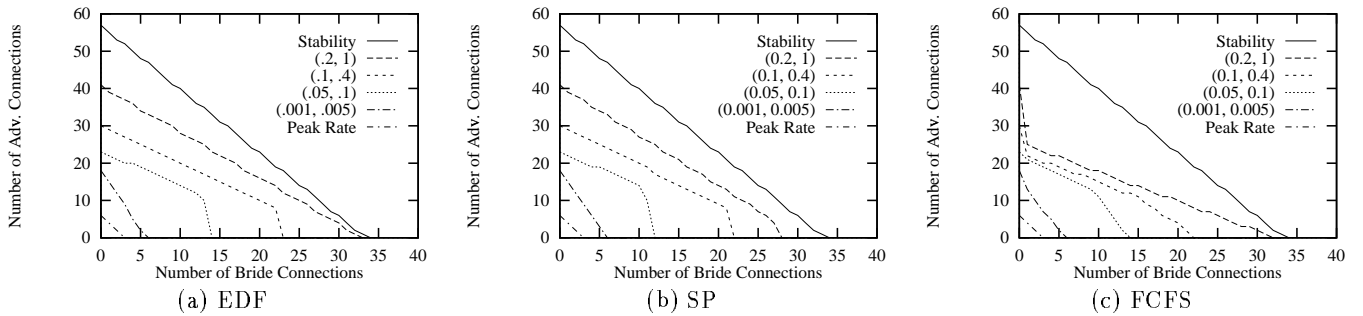


Figure 7: A comparison of packet schedulers.

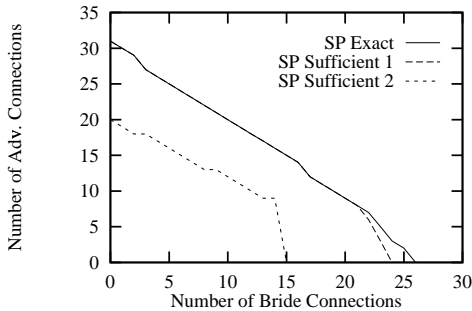


Figure 8: A comparison of admission control tests.

2, the achievable utilization is substantially lower. Since the admission test *SP Sufficient 1* achieves a utilization very close to that of *SP Exact*, it may be worthwhile to explore sufficient conditions for admissions tests.

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 9 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 of 50 msec, admission control tests are performed for the following two cases. The first benchmark 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 reasonable 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. Moreover, these utilizations are considerably above the 15% utilization achievable with a peak rate allocation scheme.

5 Conclusions

Providing deterministic guarantees requires the network to allocate resources according to the worst-case scenario. A common belief is that this requires a peak-rate resource allocation scheme that would significantly under-utilize the network when sources have high peak-to-average-rate ra-

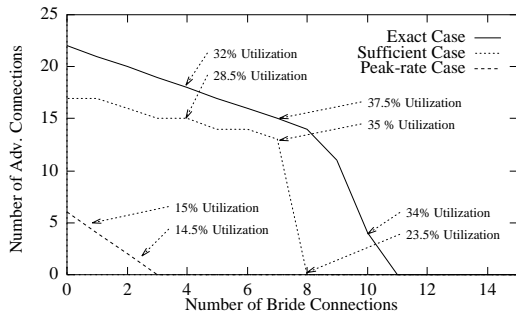


Figure 9: Achievable utilization with all tradeoffs.

tios. However, previous work has shown that a deterministic service does not require a peak-rate allocation scheme even when deterministic performance guarantees are provided to bursty traffic [9, 15]. Yet, to assess the applicability of the deterministic approach to VBR video, several important questions must be answered. For example, how high is the achievable utilization for real VBR video traffic when deterministic services are provided? What are the factors that affect the network utilization when deterministic services are provided? What are the tradeoffs between the achievable network utilization and other network design goals?

In order to address the above issues, we have studied the fundamental limits in achievable network utilization for deterministic service. We 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. The highest possible utilization is achieved for deterministic service in the “benchmark” case with (1) empirical envelopes to characterize sources, (2) EDF scheduling disciplines at the network multiplexers, and (3) necessary and sufficient conditions for admission control tests. We quantified the limits of a deterministic service to VBR video by using traces of MPEG compressed video in the scenario above and found that considerable gains over a peak rate allocation scheme are possible.

Because this benchmark scenario has limitations in its applicability to practical networks, we investigated fundamental tradeoffs for deterministic service. These tradeoffs include (1) considering a simple-to-police n level leaky bucket $(\bar{\sigma}, \bar{\rho})$ traffic model with various values of n ; (2) considering scheduling disciplines such as Static Priority and FCFS that do not require the sorted priority queue that EDF requires; (3) considering fast admission control tests that employ sufficient schedulability conditions rather than *necessary* and sufficient. Using the MPEG traces, we have shown the degree to which these factors impact the efficiency of a deterministic service to VBR video. Our experiments indicate that even in environments that consider the above tradeoffs, achievable utilizations are considerably above those obtained with a peak rate allocation scheme with average utilizations typically in the range of

25% to 35% for multiplexer delay bounds in the range of 30 to 50 msec.

As a final note, although it is tempting to draw conclusions such as “a concave constraint function with several linear segments is accurate enough to characterize VBR video” or “SP schedulers perform almost as well as EDF”, we would like to caution readers to notice that our experiments are based on only two video traces. Connections with more diverse traffic characteristics and performance requirements are needed to further explore these issues. We leave further investigation of this to future work.

References

- [1] C. Chang. Stability, queue length, and delay of deterministic and stochastic queueing networks. *IEEE Transactions on Automatic Control*, 39(5):913–931, May 1994.
- [2] R. Cruz. A Calculus for Network Delay, Part I: Network Elements in Isolation. *IEEE Transactions on Information Theory*, 37(1):114–131, January 1991.
- [3] D. Ferrari and D. Verma. A Scheme for Real-Time Channel Establishment in Wide-Area Networks. *IEEE Journal on Selected Areas in Communications*, 8(3):368–379, April 1990.
- [4] V. Frost and B. Melamed. Traffic modeling for telecommunications networks. *IEEE Communication Magazine*, 32(3):70–81, March 1994.
- [5] D. Le Gall. MPEG: A Video Compression Standard for Multimedia Applications. *Communications of the ACM*, 34(4):305–313, April 1991.
- [6] J. Hyman, A. Lazar, and G. Pacifici. Joint Scheduling and Admission Control for ATS-based Switching Nodes. In *Proc. SIGCOMM '92*, pages 223–234, August 1992.
- [7] E. Knightly and H. Zhang. Traffic characterization and switch utilization using deterministic bounding interval dependent traffic models. In *Proc. IEEE INFOCOM'95*, April 1995.
- [8] 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.
- [9] J. Liebeherr, D. E. Wrege, and D. Ferrari. Exact Admission Control in Networks with Bounded Delay Services. Technical Report CS-94-29, University of Virginia, July 1994.
- [10] S. Low and P. Varaiya. Burstiness bounds for some burst reducing servers. In *Proc. IEEE INFOCOM'93*, pages 2–9, March 1993.
- [11] I. Nikolaidis and R. Onvural. A Bibliography on Performance Issues in ATM Networks. *Computer Communications Review*, 22(5):8–23, October 1992.
- [12] C. Partridge. *Gigabit Networking*. Addison-Wesley Publishing Company, 1994.
- [13] J. Turner. New Directions in Communications (or Which Way to the Information Age?). *IEEE Communications Magazine*, 25(8):8–15, October 1986.
- [14] H. Zhang and D. Ferrari. Rate-Controlled Static-Priority Queueing. In *Proc. IEEE INFOCOM '93*, pages 227–236, April 1993.
- [15] 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, May 1994.