In Proceedings of IEEE INFOCOM'95

Traffic Characterization and Switch Utilization using a Deterministic Bounding Interval Dependent Traffic Model

Edward W. Knightly
EECS Department
University of California, Berkeley
knightly@eecs.berkeley.edu

Hui Zhang
School of Computer Science
Carnegie Mellon University
hzhang@cs.cmu.edu

Abstract

Compressed digital video is one of the most important types of traffic in future integrated services networks. It is difficult to support this class of traffic since, on one hand, compressed video is bursty, while on the other hand, it requires performance guarantees from the network. The common belief is that we are unlikely to achieve a high network utilization while providing performance guarantees to such bursty sources. In this paper, we introduce a new Deterministic Bounding Interval-Dependent (D-BIND) traffic model, together with tight analysis techniques, to explore the possibility of providing deterministic performance guarantees to VBR traffic while still achieving a reasonable network utilization. The D-BIND model consists of a family of rate-interval pairs where the rate is a bounding rate over the interval length. The model captures the intuitive property that over longer interval lengths, a source may be bounded by a rate lower than its peak rate and closer to its long-term average rate. While the D-BIND model is a general deterministic model that can be used to characterize a wide variety of sources, in this study, we focus on MPEG-compressed video. Using two 10 minute traces, we demonstrate the effectiveness of the new model and show that, contrary to common belief, reasonable network utilization can be achieved for compressed video, even when deterministic guarantees are provided.

1 Introduction

Future integrated services networks will have to support applications with diverse traffic characteristics and performance requirements. There are three important types of traffic for future integrated services networks: delay sensitive constant bit rate or CBR traffic, delay sensitive variable bit rate or VBR traffic, and best-effort or available bit rate ABR traffic. Among these, delay sensitive VBR traffic poses a unique challenge. While resource reservation schemes work best for CBR traffic [6], and there are many congestion control algorithms based on feedback and re-transmission for best-effort traffic [9, 17], there is no consensus on which strategy should be used for VBR traffic. On one hand, since VBR traffic is delay-sensitive, a resource reservation scheme seems to be the choice. On the other hand, VBR traffic is also bursty: if resources are reserved according to peak rates, the network may be under-utilized if the peak-to-average rate ratios are high. Because of this concern, feedback control algorithms have also been proposed for VBR video [2, 8, 10]. However, these algorithms cannot provide the performance guarantees desired by the application.

Thus, the key question to be answered is the following: can performance guarantees be provided to VBR traffic without significantly under-utilizing the network? In [6], two types of performance guarantees are proposed: statistical and deterministic. While statistical guarantees provide probabilistic bounds on delay and throughput, deterministic guarantees provide an *absolute* bound on delay and throughput so that performance bounds are met for *all* packets of a connection, even in the worst case. While deterministic service provides a better performance guarantee, statistical service allows the network to enhance its utilization by achieving a statistical multiplexing gain.

In this study, we propose a new traffic model and tight analysis techniques to explore the possibility of providing deterministic performance guarantees to VBR traffic while still achieving a reasonable network utilization. To better characterize the important properties of the source, we propose a Deterministic Bounding Interval-Dependent (D-BIND) model which consists of a family of rate-interval pairs where the rate is a bounding rate over the interval length. The model captures the intuitive property that over longer interval lengths, a source may be bounded by a rate lower than its peak rate and closer to its long-term average rate. We then analyze the achievable network utilization for real sources by characterizing several MPEG compressed video sequences using the D-BIND model and applying the tighter bounding techniques used in [20]. We show that, contrary to common belief, reasonable network utilization can be achieved for compressed video, even when deterministic performance guarantees are provided. Since sources may be multiplexed beyond a peak-rate-allocation scheme even while providing deterministic delay and loss and delay guarantees, we define the Deterministic Multiplexing Gain (DMG) as the gain in utilization above a peak-rate-allocation scheme that is achieved. The DMG is used to further quantify the improvements of the new model.

While the D-BIND model is a general deterministic model that can be used to characterize a wide variety of sources, in this study, we use traces of MPEG-compressed video as an example of VBR traffic to evaluate the new approach. Two observations are important in this context. The first is that although compressed video is bursty, it is much more "regular" and "structured" than data traffic. While compressed video is bursty because the size of a compressed frame varies from one frame to the next, there is an underlying structure in that a new frame is generated every 33 msec. More importantly, for an MPEG source, the largest local variation between frame sizes is due to the alternation of inter-frame coded frames with intra-frame coded frames. That is, a larger I-frame is immediately followed by a smaller B-frame so that the micro-level burst does not persist for very long. The second observation is that as long as there is adequate buffering, bursty traffic can always be serviced at a lower rate than its peak rate; the major concern is that this buffering introduces delay. The D-BIND model attempts to capture such intuitive observations in a well-defined, deterministic manner.

The remainder of this paper is organized as follows. In Section 2, we describe the underlying requirements of a deterministic source model and review previously proposed models. We then define and analyze the new model and present the new admission control tests. In Section 3 we investigate the performance of the new model using parameters derived from actual MPEG traces. Finally, in Section 4, we discuss some of the practical issues for the model including policing and parameter specification.

2 Deterministic-BIND Model

While a deterministic service cannot, by definition, employ statistical multiplexing, it does provide a better service in that no packets are dropped and none violate their guaranteed delay bound. For the network to deliver such a service, it needs a deterministic upper bound on all sources receiving the service. This approach has the added advantage that a source's traffic specification can be enforced. For example, if a source promises that its minimum packet inter-arrival time is $Xmin$, this may be easily verified and enforced by the network. Alternatively, statistical models of the source are inherently much more difficult to enforce.

2.1 Deterministic Traffic Models

In the $(Xmin, Xave, I, Smax)$ model of [6] (we will refer to this as the $Xmin$ model), a source is constrained so that its minimum packet spacing is $Xmin$, its maximum packet size is $Smax$, and that in every interval of length I, it may send no more than $I/Xave$ packets. In [4], a source is said to satisfy a (σ, ρ) leaky-bucket model if during any interval of length t, the number of bits that the source transmits is less than $\sigma + \rho t$. The (σ, ρ) model was extended in [13, 14] with a proposed burstiness curve to characterize traffic. Rather than use a single (σ, ρ) pair, this work considers the function $\sigma(\rho)$ (the burstiness curve) so that the above relationship holds for all points of the $\sigma(\rho)$ curve.

As required, all of the above characterizations provide a deterministic upper bound on each source's arrivals, and allow a worst-case analysis that provides absolute bounds on delay and throughput. Specifically, each deterministic traffic model defines a traffic constraint function $b(t)$ which constrains or bounds the source over every interval of length t. Denoting $A[t_1, t_2]$ a connection's arrivals in the interval $[t_1, t_2]$, the traffic constraint function $b(t)$ requires that $A[s, s+t] \leq b(t)$, $\forall s, t > 0$. Note that $b(t)$ is a time invariant deterministic bound since it constrains the traffic source over every interval of length t.

An important observation about the traffic constraint function is that for a given arrival process A, the tightest time invariant deterministic bound on arrivals in any interval of length t is by definition

$$\mathcal{E}(t) = \sup_{s \geq 0} A[s, s+t] \tag{1}$$

$\mathcal{E}(t)$ is called the empirical envelope in [11], and the minimum envelope process in [3]. Thus, in order for a traffic model's constraint function $b(t)$ to be a time invariant upper bound on A, it must be an upper approximation to $\mathcal{E}(t)$. A desirable property of a traffic model is therefore that it provides a constraint function that can closely bound $\mathcal{E}(t)$ for a wide variety of sources. Examples of constraint functions for the $Xmin$ and (σ, ρ) traffic models are shown in Figure 1. Conceptually, both models allow a limited-sized burst and have an additional longer-term rate constraint.

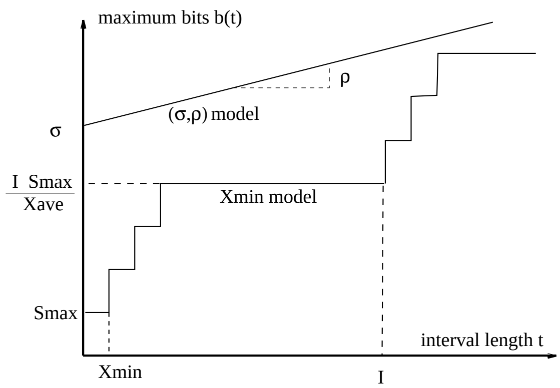

Figure 1: Traffic Constraint Curves

Before defining the D-BIND model, we further motivate it by describing the analysis techniques used to derive connection admission control conditions for deterministic guarantees.

2.2 Delay Analysis

Deterministic admission control conditions rely on the delay analysis techniques of [4, 20] which may be described as follows. Figure 2 illustrates the different components of the analysis. The horizontal axis is time and the vertical axis is bits. The upper curve represents the total number of bits that have arrived in the queue by time t and the lower curve represents the total number of bits transmitted by time t. The difference between the two curves is the number of bits currently in the queue, or the *backlog* function. When the backlog function returns to zero (the two curves meet) there are no bits in the queue and thus a busy period has ended. The key to this analysis is that if the upper curve is a deterministic bounding curve, then the maximum delay can be expressed as a function of the two curves. For example, the following two observations hold: the maximum busy period provides an upper bound on delay for any work-conserving server; the maximum backlog divided by the link speed provides an upper bound on delay for a FCFS server. Delay bounds for other policies can also be expressed [1, 4, 12, 15].

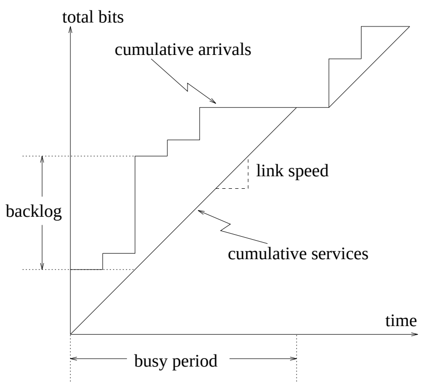

Figure 2: Concepts: Delay, Backlog and Busy Period

The constraint function provides the required bound on arrivals in any interval of length t, so that with the aggregate of individual source's respective b(t) constraint functions forming the upper curve of Figure 2, admission control conditions for deterministic delay and throughput bounds may be derived. For example, for a FCFS scheduler with $j = 1, ..., n$ multiplexed connections constrained by their respective constraints $b_j(t)$, and with a link speed l, and a maximum packet size $\overline{s}$, a deterministic upper bound on delay for all connections is given by

$$d = \frac{1}{l} \max_{t \geq 0} \{ \sum_{j=1}^{n} b_j(t) - lt + \overline{s} \}. \tag{2}$$

The proof is given in Theorem 1 of [20].

The equation indicates that even better bounds are possible with new traffic models. That is, if a given traffic source can be more tightly bounded by a different constraint function than those of previous traffic models, the resulting maximum delay bound of Equation (2) will be lower. Thus, the goal of the D-BIND model is a more accurate source characterization that results in a tighter (lower) traffic constraint function $b(t)$. The effect is thus a higher network utilization and a higher DMG for a given deterministic delay and throughput constraint.

2.3 D-BIND Model Definition

In [21], we proposed a Stochastic Bounding Interval Dependent (S-BIND) model for providing statistical performance guarantees. In the S-BIND model, a source is stochastically bounded in intervals of different length to capture the intuitive property that over longer interval lengths, a source's rate may be bounded by a random variable that is weighted nearer to its long term average rate. In this study, we consider the deterministic case by characterizing a source with a *deterministic* bound on a source's rate over intervals of different length. The key to the BIND models is that they are *bounding*, needed for admission control, and *interval-dependent*, to characterize the important interval-length dependent behavior of sources.

The D-BIND model may be defined as follows. Source j, may be described by the curve $R_j(I)$ where $R_j(I)$ is the bounding rate over every interval of length I so that

$$A_j[t, t+I]/I \leq R_j(I) \quad \forall\, t, I > 0. \tag{3}$$

Thus, the source is deterministically constrained to transmit no more than $b_j(t) = t \cdot R_j(t)$ bits during any interval of length t.

There are several points to note about this characterization. First, for a given source, the general trend of the $R(I)$ curve is that for small I, $R(I)$ will approach the source's peak rate while for larger interval lengths, $R(I)$ approaches the source's long term average rate from above, where the long term average rate is defined empirically as $\lim_{t \to \infty} A[0, t]/t$. Although a tight $R(I)$ is not necessarily convex or monotonically decreasing, its associated constraint curve, $b(t) = t \cdot R(t)$, is subadditive. While an $R(I)$ curve that is not strictly decreasing may seem unusual, we demonstrate later with real traffic traces that this may indeed be the case. Regardless, the predominant trend is that $R(I)$ decreases with increasing interval length: this is the interval-length dependent property of the BIND model — that over longer intervals, the bounding rate decreases. By explicitly characterizing the source's different bounding rates over different interval-lengths, we will show analytically and demonstrate empirically that higher network utilizations are achievable.

In practice, a traffic source must be able to specify its traffic with a small number of parameters. For this reason, the D-BIND *model* consists of N rate-interval pairs, i.e., $\{(R_n, I_n) | n = 1, 2, \cdots, N\}$. We distinguish the D-BIND *model* from the D-BIND *characterization* which consists of the entire $R(I)$ curve. In the practical case with the parameterized D-BIND model, the traffic constraint function requires interpolation between the rate-interval pairs. A practical interpolation function is the linear one shown in Figure 3. In this figure, the lower curve represents the tightest bound on the number of arrivals in any interval of length t, $\mathcal{E}(t)$ (Equation (1)). The D-BIND model provides a piece-wise linear upper approximation to this tightest bound. For example, in the figure the source is constrained (over every interval of length t) tightly by the lower curve $\sup_{s\geq 0} A[s, s+t]$, and approximately by the D-BIND model's constraint function with several D-BIND pairs $(b_n/t_n, t_n)$, and with linear interpolation between the points on the constraint curve. Thus given rate-interval pairs (R_n, I_n),

$$b(t) = \frac{R_k I_k - R_{k-1} I_{k-1}}{I_k - I_{k-1}}(t - I_k) + R_k I_k, \quad I_{k-1} \leq t \leq I_k \tag{4}$$

assuming $R_0 = I_0 = 0$.

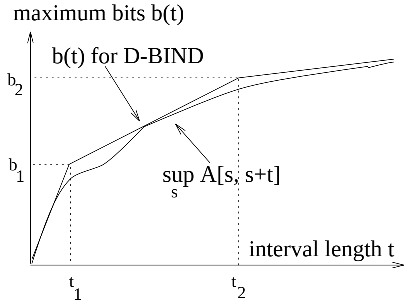

Figure 3: Traffic Constraint Function for D-BIND Model

Note that if the D-BIND curve is tight, i.e., $R(I) = \sup_{s\geq 0} A[s, s+I]/I$, then this represents the tightest deterministic time-invariant characterization of a source. (Equation (3) shows that $R(I)$ is a bound that is not *required* to be tight.) With more and more rate-interval pairs, the D-BIND model approaches this tight constraint.

2.4 D-BIND's Relationship to Other Traffic Models

Note that other deterministic traffic models may be expressed in terms of the D-BIND model. For example, a traffic model based on multiple (σ, ρ) pairs $\{(\sigma_n, \rho_n), n = 1, 2\cdots, N\}$ (a parameterized version of [13, 14]) is a special case of the D-BIND model in which the constraint function is piece-wise linear *concave*. That is, a multiple (σ, ρ) model has a constraint curve $b(t) = \min_n\{\sigma_n + \rho_n t\}$ which is necessarily concave (see Section 4). As well, the $Xmin$ model can be expressed in terms of the D-BIND model by using a different interpolation function.

2.5 Admission Control

Two important components to providing performance guarantees in a connection-oriented network are connection admission control and the packet service discipline. The RCSP service discipline [19] provides the mechanisms needed to provide integrated services to heterogeneous sources. The scheduler is based on a number of FCFS priority queues where queue k has an associated delay bound d_k. A number of heterogeneous connections can be multiplexed at a given priority level k. When a new connection is requested with a delay bound d_k, the admission control algorithm must first check that sufficient resources are available at the node so that all existing connections and the new connection will meet their respective throughput and delay bounds. As described in Section 2.2, this may be achieved by calculating a bound on the maximum backlog. The connection admission control test of Equation (2) may be extended from FCFS to the RCSP scheduler with the following theorem.

Theorem 1 *Assume a Static Priority scheduler has n priority levels. Let C_q be the set of connections at level q, and the j^{th} connection in C_q satisfies the traffic constraint function $b_{q,j}(\cdot)$. With a link speed l, and maximum packet size of $\overline{s}$, the maximum delay of any packet at priority level k is bounded above by d_k, where*

$$d_k = \max\{t \geq 0 \mid b'_k(t) \geq lt\} \tag{5}$$

and $b'_k(\alpha)$ is defined for all α by

$$b'_k(\alpha) = \max_{\beta\geq 0}\{\overline{s} + \sum_{j\in C_k} b_{k,j}(\beta) + \sum_{q=1}^{k-1}\sum_{j\in C_q} b_{q,j}(\alpha+\beta) - l\beta\}. \tag{6}$$

The proof is by extension of the results of [4]. Details may also be found in [18].

3 Traffic Characterization and Network Utilization

This section evaluates the effectiveness of the D-BIND model by analyzing the switch or multiplexer utilization achieved with the new model and comparing the results to utilizations obtained with the $Xmin$ model using the bounds in [20]. Also for comparison, utilizations obtained with peak-rate reservation are also investigated. (By peak-rate reservation, we mean an admission control scheme in which the sum of the source's peak rates are constrained to be less than the link speed.) Two traces of MPEG compressed video are analyzed as traffic sources. One video is a series of advertisements and the other is a lecture. The advertisement video is quite fast moving and has a wide variety of scenes with varying complexity. Alternatively, the lecture video does not have much action other than the speaker's movements and changes of scene from the speaker to the transparencies and back. The nature of these two video streams will be shown to have a remarkable effect on the achievable network utilization.

3.1 Deterministic Characterization

Figure 4 shows a typical four-second segment of each of the ten-minute traces. The vertical axis is rate in Mbps and the horizontal axis is time. It is assumed that the entire frame is transmitted per frame-time (as opposed

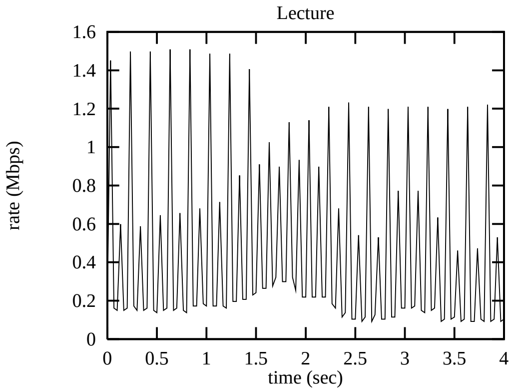

(a)

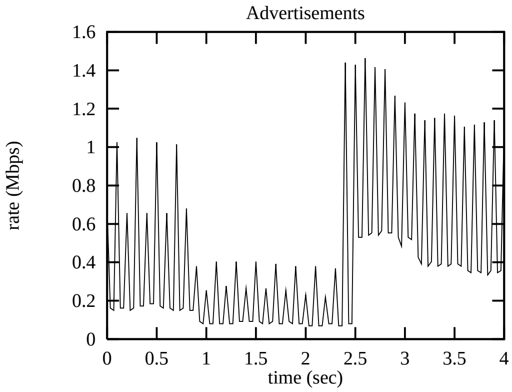

(b)

Figure 4: Segments of MPEG Compressed Video Traces

to introducing additional delay by smoothing over several frames) so that Figure 4 depicts the frame size multiplied by the frame rate (30 fps). Additionally, it is assumed that each frame is fragmented into 48 byte ATM cells with the cells being transmitted at equally spaced intervals over the frame-time.

The general shape of the traces may be explained in terms of the mechanisms used in the MPEG standard. The coder generates three types of frames: I frames that use only *Intraframe* compression, and P and B frames that are transmitted between I frames and 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 a future frame. With P and B frames, higher compression ratios can be achieved since the interframe coding makes use of motion compensation techniques. More details of the MPEG algorithm may be found in [7]. The frame pattern for Figure 4 is IBBPBB; which frames are which is apparent since the I frames tend to be the largest, B the smallest, and P in between.

Figure 5 shows the D-BIND $R(I)$ curves for the advertisement and lecture sequences. The vertical axis is the bounding rate over an interval of length I, where I is on the horizontal axis. The bounding rate may be converted to $Xave$ by inverting and multiplying by $Smax$ (the unit of $Xave$ is seconds/packet). Thus, for the $Xmin$ characterization, any $(I, Smax/Xave)$ pair from the curve along with $Xmin$ (obtained from $Smax/Xave$ for $I = 1$ frame-time on the curve) results in a valid deterministic characterization of the source.

There are several things to note about the figure. First, the bounding rate for $I = 1$ frame-time is the peak rate since the bounding rate over a single frame-time is caused by transmission of the largest frame in the sequence. Since this transmission must occur within the fixed frame-time of $\frac{1}{30}^{th}$ of a second, the cell spacing will be the minimum $(Xmin)$ and the rate will be the maximum. Second, the general trend of the curves is that the bounding rate decreases with increasing interval length. This is the intuitive property described in Section 2, that the bounding rate of a source is smaller over larger interval lengths, decreasing from the peak rate to the long-term-average rate.

Next, and of most importance for the analyses presented here, is how quickly $R(I)$ approaches the long-term-average rate (shown as a dashed line) as I increases. For example, the curve for the the lecture video in Figure 5(a) shows a rapid decrease of the bounding rate, whereas the curve for the advertisement video in Figure 5(b) decreases much more slowly. Intuitively, a slowly decreasing $R(I)$ curve indicates that bursts of high rate persist over relatively long interval-lengths which in turn implies that it will be extremely difficult to multiplex such a source (i.e., it should not be expected that we can do much better than peak-rate reservation). We conjecture that video compressed with a coder that does only *intra-frame* compression (e.g., JPEG), will have this undesirable property of a slowly decreasing $R(I)$ curve.

3.2 Constraint Functions

Figure 6 shows the traffic constraint function of the lecture sequence for both the $Xmin$ model the D-BIND model. As explained in Section 2.1, the horizontal axis is interval length and the vertical axis is the maximum number of bits that deterministically constrain the source. The lower this curve is, the more tightly the model represents the sources (i.e., it may be bounded with fewer bits for a given interval length). As shown, regardless of the choice of I in the $Xmin$ model, the D-BIND model more tightly represents the traffic source.

Note that the temporal properties of the MPEG source are evident in the D-BIND model's constraint function: both the D-BIND and $Xmin$ constraint curves begin with an initial slope which represents the source sending at its peak rate, i.e., transmitting its largest I frame. At 33 msec (1 frame-time) the slope of the D-BIND constraint function sharply decreases indicating that even in the worst case, a large I frame is followed by a typically smaller B frame. At 100 msec, after sending two B frames, D-BIND's constraint function breaks up again indicating the transmission of a

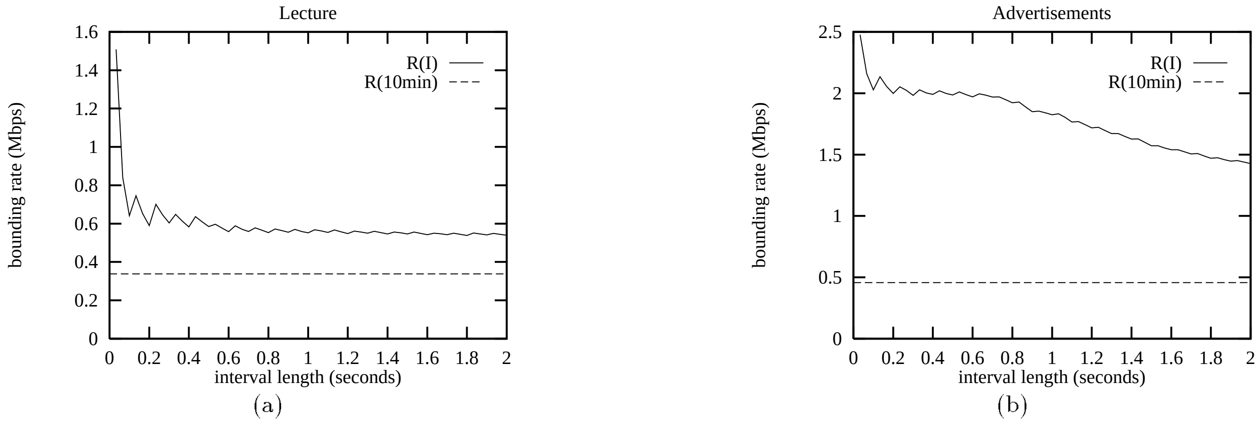

Figure 5: D-BIND $R(I)$ Bounding Rate Curve

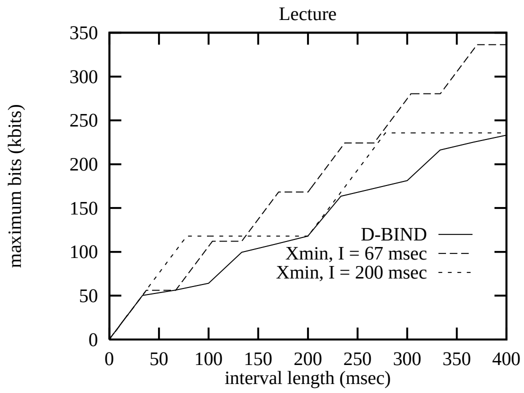

Figure 6: Traffic Constraint Function

P frame. In essence, the D-BIND model is capturing the temporal nature of the MPEG source. Finally, note that this constraint curve is not concave. Section 4 addresses the implications of this observation that regard policing.

3.3 Switch Utilizations and DMG

A key issue for evaluation of the D-BIND model is the achievable network utilization compared to other deterministic traffic models.

In the experiments below, we calculate the maximum number of homogeneous connections that can be multiplexed at a link so that all connections receive a deterministic guarantee on delay and throughput (see [11] for investigations of the heterogeneous case and scheduling disciplines other than RCSP). In this case, the maximum number of channels with constraint function $b(t)$ that may be given a deterministic delay bound of d_k at an RCSP scheduler served at link speed l is given by

$$N(d_k) = max\{n \mid \frac{1}{l} \max_{t \geq 0}\{nb(t) - lt + \overline{s}\} \leq d_k\} \qquad (7)$$

where $b(t)$ is given by Equation (4) and shown in Figure 6 for the lecture sequence. Within this context, for a deterministic delay bound d_k, the average utilization for deterministic traffic is given by $N(d_k) \cdot R(I = 10min)/l$ where $R(I = 10min)$ is the long-term-average rate. Since sources may be multiplexed beyond a peak-rate-allocation scheme even while providing deterministic delay and loss and delay guarantees, we define the Deterministic Multiplexing Gain as the gain in utilization above a peak-rate-allocation scheme that is achieved with the new model. Thus, the DMG is the sum of the peak rates of all connections with deterministic guarantees divided by the link speed, which is $N(d_k) \cdot R(I_1)/l$ in the homogeneous case. Thus, a peak-rate-allocation scheme has a DMG of 1.

For the two video streams, Figures 7(a) and 7(b) show the number of channels accepted, the average utilization, and the DMG, all as a function of the deterministic delay bound d_k. The link speed is T3 or 45 Mbps and video frames are fragmented into ATM cells and transmitted as described previously. Figure 7(a) shows the data for the lecture sequence while Figure 7(b) is for the advertisements. As expected from Equation (7), the utilizations are functions of delay, increasing with the delay bound until queue lengths cannot be bounded if more connections are accepted.

There are several noteworthy points about Figure 7. First, it is immediately apparent that the D-BIND model performs better than the $Xmin$ model for any choice of I. For example, for the lecture sequence of Figure 7(a) and a guaranteed delay bound of 58 msec, the D-BIND model is able to utilize the network to 60%. Alternatively, depending on the choice of I, the $Xmin$ model results in utilizations of approximately 40% so that, in this case, the D-BIND model results in a 50% improvement in network utilization. It should also be noted that both the D-BIND and $Xmin$ models do significantly better than peak-rate reservation which results in an average utilization of 23%. With the improved analysis techniques of [4, 20], even for small delay bounds, DMG's significantly greater than 1 are achievable. For example, for a delay bound of 9 msec, 38 channels may be multiplexed for a DMG of 1.3. For a 40

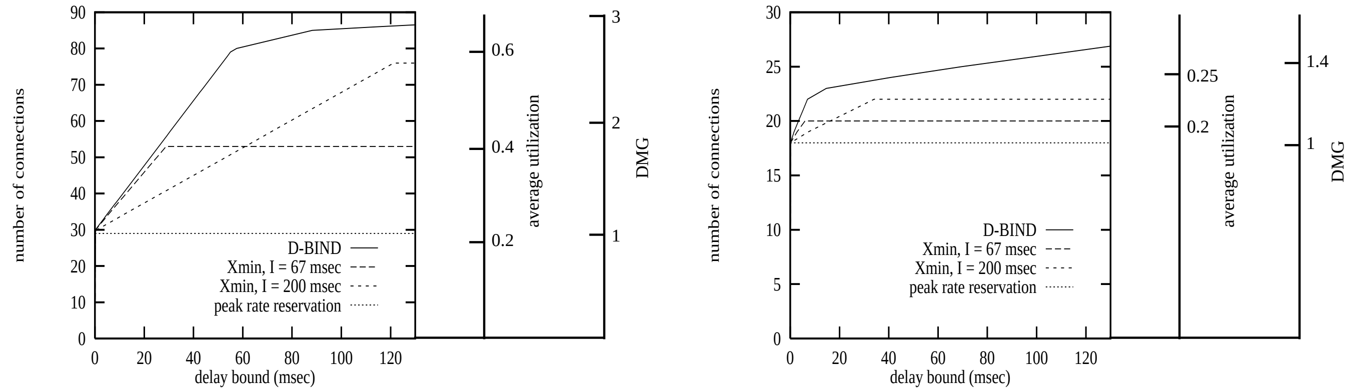

(a) lecture (b) advertisements

Figure 7: Utilization and DMG for Lecture and Advertisements

msec delay bound, the DMG is 2.2, and for a 48 msec delay bound, it is 2.7. For the $Xmin$ model, the DMG's achieved are high, but depending on the choice of I, substantially below the DMG's for the D-BIND model.

Second, note that the results of Figure 7(b) are not as pronounced. This is due to the shape of the $R(I)$ curve in Figure 5(b). Though the curve does obey the interval-dependent property that sources may be bounded by lower rates over longer interval lengths, this property is obeyed in a very lethargic manner. That is, compared to Figure 5(a), the $R(I)$ curve of Figure 5(b) decreases more slowly to its long-term-average rate. In this case, for a delay bound of 69 msec, the improvement is from an average network utilization of 18% for a peak-rate-allocation scheme to approximately 21% for the $Xmin$ model and to 25% for the D-BIND model.

Next, the figures demonstrate that for a given I, there may be a small range of delays such that the $Xmin$ model performs nearly as well as the D-BIND model. However, note that the D-BIND model still has a significant advantage with respect to practical issues of establishment of real-time connections in a network. For example, if the required end-to-end delay of a connection is 200 msec and the connection traverses several switches, these switches will have different loads. Depending on the load, each switch may wish to allocate a different local delay bound to the connection. Thus, it may easily happen that the local delay bounds are 120, 20, and 60 msec at the respective three nodes. Therefore, regardless of how cleverly the user chooses I for the $Xmin$ model, some of the nodes will be forced to allocate resources inefficiently since choosing one I tends to yield a decent bound for some delays and a poor bound for others.

Finally, note that the utilizations shown are for deterministic real-time traffic only. The remaining network resources may be used by statistical or best-effort traffic.

4 Discussion

In proposing a new source model there are several issues regarding the practicality of the model. For a deterministic model, these issues include: a) source specification - how difficult is it for a source to come up with its characterization? b) parameterization - can the model be represented in a concise manner? c) policing - can the model be effectively and efficiently enforced?

4.1 Source Specification

A problem with live sources such as live video is that the source's parameterization is not known a priori. Of course, this problem is not limited to the analysis presented here. Two factors can alleviate this problem. First, with the Dynamic Connection Management Scheme [16], a real-time channel can change its traffic specification or performance requirements during the duration of the connection. Thus, a source can adapt its (R_n, I_n) rate-interval pairs to the least upper bound. Second, although the general form of the D-BIND model consists of the entire $R(I)$ curve, in practice, specifying a small number of points (investigated below) will likely be sufficient.

4.2 Parameterization

In the example of Section 3, a (R_n, I_n) rate-interval pair was used for each frame time up to an interval length of several seconds. In the following experiment, we use four rate-interval pairs to characterize the traffic and calculate the maximum number of acceptable connections as in the previous sections. Figure 8 shows the result. While the homogeneous case does not explore all of the facets of using different constraint functions, this experiment indicates that a smaller number of D-BIND rate-interval pairs may result in utilizations close to those achieved with a large number of pairs.

For MPEG video sources, an alternative concise parameterization is to use knowledge of the frame pattern (in this

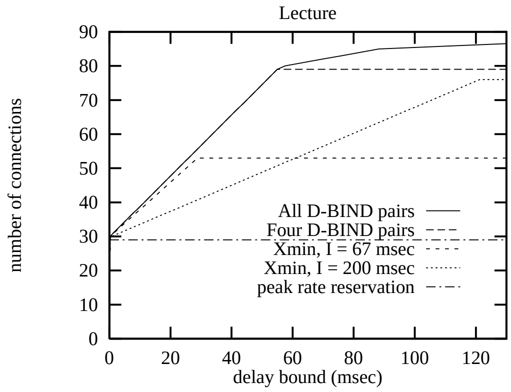

Figure 8: Utilization with Four Rate-Interval Pairs

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 the D-BIND $R(I)$ curve can be obtained by constructing the constraint function function as a transmission of the largest I frame, followed by 2 transmissions of the largest B frame and so on. In essence, any $b(t)$ that is a piece-wise linear upper approximation to $\mathcal{E}(t) = sup_{s \geq 0} A[s, s+t]$ can be used within the D-BIND framework.

4.3 Policing

Since the network must protect clients from malicious users, it needs to monitor the traffic from each source to ensure that it satisfies its traffic specification. Such an access control function at the network, called policing, is shown in Figure 9. The input to the policer comes from the source and the output goes to the network. The function of the policer is to ensure that the the traffic it outputs to the network satisfies the traffic constraint function $b(t)$ that is specified by the source's model parameters. To achieve this, the policer may need to buffer or drop packets when the input stream exceeds the limit defined by $b(t)$. If the input stream to the source policer satisfies the traffic constraint function, no buffering or delay will incur in the policer.

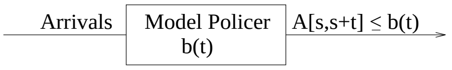

Figure 9: Traffic Constraint Function $b(t)$

As noted in Section 4.2 and shown with the traces of Section 3, a piece-wise linear function may be used to represent the D-BIND model's constraint function. Section 3 demonstrated that because of the temporal properties of MPEG sources, the sources considered here had neither monotonically decreasing $R(I)$ curves nor concave constraint functions. As addressed by the propositions below, a concave constraint function has implications for policing.

Lemma 1 *If the constraint function $b(t)$ is concave, then $R(I)$ is strictly decreasing.*

Proof: A function $b(t)$ is concave if for any $t_1 < t_2$ and $0 \leq \alpha \leq 1$, $\alpha b(t_1) + (1-\alpha) b(t_2) \leq b(\alpha t_1 + (1-\alpha) t_2)$. We need to show that for any $u_1 < u_2$, $R(u_1) \geq R(u_2)$ or $\frac{b(u_1)}{u_1} \geq \frac{b(u_2)}{u_2}$. Since $b(0) = 0$, in the inequality above, let $t_1 = 0$, $t_2 = u_2$, and $\alpha = 1 - u_1/u_2$. Thus, we have $b(u_1) \geq u_1/u_2 \times b(u_2)$.

Lemma 2 *If a piece-wise linear constraint function $b(t)$ with N linear segments is* concave*, then the source may be fully policed (i.e., Equation (3) holds) with a cascade of N leaky buckets.*

The proof is given in Theorem 5.1 of [5]. Note however, as shown in Section 3, a source does not necessarily have a concave constraint function $b(t)$. In this case, a piece-wise linear *non-concave* constraint function may be policed with a cascade of leaky buckets with *state-dependent* token-generation rates. That is, the leaky bucket's token rate is a function of the number of cells transmitted over the previous interval. Thus, for simplicity, one may opt to approximate a source's constraint curve by its concave hull so that it may be policed with a cascade of one or more leaky buckets. A concave constraint function (not necessarily piece-wise linear) also corresponds to the $\sigma(\rho)$ characterization of [13, 14] since the corresponding constraint function, $b(t) = \min_\rho \{\sigma(\rho) + \rho t\}$, is by definition concave. The experiment of Figure 8 used a concave constraint function. However, the possible utilization gain of using non-concave constraint functions is not apparent in the homogeneous case. Thus, we defer further discussion of concavity to future work.

5 Conclusion

The analysis techniques of [4, 18, 20] have shown that a peak-rate resource allocation scheme is not required in order to provide deterministic performance guarantees.

In this paper, we demonstrated several things. First, noting that a better deterministic traffic model will result in a tighter traffic constraint function, which in turn results in higher network utilizations, we proposed a new traffic model called the Deterministic Bounding Interval Dependent (D-BIND) model to better capture the property that over longer interval lengths, sources may be bounded with smaller rates. Using MPEG-compressed video traces, we demonstrated that the D-BIND model can achieve a higher network utilization for a given performance requirement than previous models.

Second, we showed that reasonable network utilization can be achieved even while providing deterministic performance guarantees to bursty traffic. Using MPEG compressed video traces, we showed that network utilizations of over 60% are achievable for "well-behaved" sources such as a lecture sequence and network utilizations of over 25% are achievable for more "ill-behaved" sources such as an advertisement sequence. The ability to efficiently multiplex these sources was demonstrated to be due largely to the shape of the D-BIND $R(I)$ curve rather than the more commonly used characterization of peak-to-average rate ratio. If the $R(I)$ curve decreases too slowly, high network utilizations are difficult to achieve since, intuitively, if a source can send at near its peak rate for a long length of time, the network cannot absorb these bursts without excessively large buffers and introducing excessively large delays.

Finally, we quantify the advantages of the new model's improvements over a peak-rate-allocation scheme by exploring the achievable DMG. As shown, with the D-BIND model's effective way of capturing the interval-dependent properties of sources as well as their temporal characteristics, high DMG's of over 2.7 were achieved for the lecture sequence and 1.4 for the advertisement sequence.

In conclusion, although only MPEG traffic was analyzed here, we believe that the D-BIND model is general enough to capture the essential properties of a wide variety of sources and that gains in utilization provided by the D-BIND model will be largely determined by the rate at which the bounding rate decreases with interval length.

6 Acknowledgements

The authors are grateful to Domenico Ferrari, Jörg Liebeherr, Dallas Wrege, and to the members of the Tenet Group for their helpful discussions, to Jörg for pointing out subadditivity of the empirical envelope, and to Sandia Labs for help acquiring the MPEG traces.

References

[1] A. Banerjea and S. Keshav. Queueing delays in rate controlled networks. In *Proceedings of IEEE INFOCOM'93*, pages 547–556, San Francisco, CA, April 1993.

[2] J.-C Bolot, T. Turletti, and I. Wakeman. Scalable feedback control for multicast video distribution in the Internet. In *Proceedings of ACM SIGCOMM'94*, pages 58–67, University College London, London, U.K., September 1994.

[3] C. Chang. Stability, queue length, and delay of deterministic and stochastic queueing networks. *IEEE Transactions on Automatic Control*, 39(5):913–931, May 1994.

[4] R. Cruz. A calculus for network delay, part I : Network elements in isolation. *IEEE Transactions on Information Theory*, 37(1):114–121, January 1991.

[5] R. Cruz. A calculus for network delay, part II : Network analysis. *IEEE Transactions on Information Theory*, 37(1):121–141, January 1991.

[6] 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.

[7] D. Le Gall. MPEG: A video compression standard for multimedia applications. *Communications of the ACM*, 34(4):46–58, April 1991.

[8] M. Gilge and R. Gusella. Motion video coding for packet switching networks – an integrated approach. In *Proceedings of SPIE Visual Communications and Image Processing '91*, pages 592–603, Boston, MA, November 1991.

[9] V. Jacobson. Congestion avoidance and control. In *Proceedings of ACM SIGCOMM'88*, pages 314–329, Stanford, CA, August 1988.

[10] H. Kanakia, P. Mishra, and A. Reibman. An adaptive congestion control scheme for real-time packet video transport. In *Proceedings of ACM SIGCOMM'94*, pages 20–31, San Francisco, CA, September 1993.

[11] E. Knightly, D. Wrege, J. Liebeherr, and H. Zhang. Fundamental limits and tradeoffs for providing deterministic guarantees to VBR video traffic. Technical Report TR-94-067, International Computer Science Institute, Berkeley, CA, December 1994.

[12] J. Liebeherr, D. Wrege, and D.Ferrari. Exact admission control for networks with bounded delay services. Technical Report CS-94-29, University of Virginia, Department of Computer Science, July 1994.

[13] S. Low and P. Varaiya. A simple theory of traffic and resource allocation in ATM. In *Proceedings of IEEE GLOBECOM'91*, pages 1633–1637, Phoenix, AZ, December 1991.

[14] S. Low and P. Varaiya. Burstiness bounds for some burst reducing servers. In *Proceedings of IEEE INFOCOM'93*, pages 2–9, San Francisco, CA, March 1993.

[15] A. Parekh and R. Gallager. A generalized processor sharing approach to flow control - the single node case. In *Proceedings of IEEE INFOCOM'92*, pages 521–530, Firenze, Italy, May 1992.

[16] C. Parris, H. Zhang, and D. Ferrari. Dynamic management scheme for real-time connections. In *Proceedings of IEEE INFOCOM'94*, pages 698–707, Toronto, Ontario, June 1994.

[17] K. Ramakrishnan, D. Chiu, and R. Jain. Congestion avoidance in computer networks with a connectionless network layer. In *Proceedings of ACM SIGCOMM'88*, pages 303–313, Stanford, CA, August 1988.

[18] H. Zhang. Service disciplines for integrated services packet-switching networks. Ph.D. Dissertation. UCB/CSD-94-788, U. C. Berkeley, November 1993.

[19] H. Zhang and D. Ferrari. Rate-controlled static priority queueing. In *Proceedings of IEEE INFOCOM'93*, pages 227–236, San Francisco, CA, March 1993.

[20] H. Zhang and D. Ferrari. Improving utilization for deterministic service in multimedia communication. In *Proceedings of 1994 International Conference on Multimedia Computing and Systems*, pages 295–304, Boston, MA, May 1994.

[21] H. Zhang and E. Knightly. Providing end-to-end statistical performance guarantees with bounding interval dependent stochastic models. In *Proceedings of ACM SIGMETRICS'94*, pages 211–220, Nashville, TN, May 1994.