to Appear 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 at Berkeley
and Sandia National Laboratories
knightly@tenet.berkeley.edu

Hui Zhang
Lawrence Berkeley Laboratory
MailStop: 50B-229
Berkeley, CA 94720
hzhang@george.lbl.gov

Abstract

Compressed digital video is one of the most important types of traffic in the 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 bursty traffic. While this is certainly true for traditional data traffic, compressed video is much more "regular" and "smooth" than data traffic. In this paper, we propose a deterministic bounding interval-dependent (BIND) model to capture the source's characteristics. We use the BIND model together with a tighter analysis technique to show that, contrary to common belief, reasonable network utilization can be achieved for compressed video even when deterministic guarantees are provided. In the study, we used several 10 minutes long MPEG compressed video sequences to demonstrate the effectiveness of the new model. Since even if all packets are deterministically guaranteed to meet their loss and delay bounds, sources may be multiplexed beyond their peak rate, we define the Deterministic Multiplexing Gain (DMG) as the fraction above a peak-rate allocation scheme that is achieved while still providing a deterministic performance guarantee. We show that with the new BIND model, network utilizations as high as 60% and DMG's of up to 2.8 are achievable for MPEG video.

keywords: quality of service, video traffic characterization, deterministic multiplexing gain.

1 Introduction

Future integrated services networks will have to support applications with diverse traffic characteristics and performance requirements. There are three important classes 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 ABR traffic. Examples for these classes are uncompressed video, compressed video and data traffic respectively. Among them, delay sensitive VBR traffic poses unique challenge. While resource reservation schemes work best for CBR traffic [5], and there are also many congestion control algorithms based on feedback and re-transmission strategy proposed for best-effort traffic [7, 12], there is no consensus on which strategy should be used for VBR traffic. On the one hand, since VBR traffic is delay-sensitive, a resource reservation scheme seems to be the choice. One the other hand, VBR traffic is also bursty: if resources are reserved according to the peak rate, the network may be under-utilized if the peak-to-average rate ratio is high. Because of this concern, various feedback control algorithms have also been proposed [8]. However, these algorithms can not 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 [5], two types of performance guarantees are proposed: statistical guarantees and deterministic guarantees. 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 the deterministic performance guarantee is a better service, statistical service allows the network to enhance its utilization by achieving a statistical multiplexing gain.

In this paper, we explore the possibility of providing deterministic performance guarantees to VBR traffic while at the same time achieving a reasonable network utilization. In this study, we use MPEG-compressed video as an example of VBR traffic. Two observations are important in the study. The first one is that although compressed video is bursty, it is much more "regular" and "smooth" compared to data traffic. Compressed video is bursty because the frame size varies from one frame to another. A closer look at compressed full motion video shows that there is more structure than a new frame every 33 msec: for an MPEG source, the largest variation between frame sizes is due to the alternation of inter-frame coding frames with intra-frame coding frames. I.e., a larger sized I-frame is immediately followed by a smaller B-frames 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 the peak rate. The major concern is that buffering introduces delay.

In this study, we address the issue of achieving a high network utilization with use of a new traffic model and tighter analysis techniques. To better characterize the important properties of the source, we propose a a deterministic bounding interval-dependent (BIND) model which captures the source's different characteristics in intervals of different length. We analyze the achievable network utilization by characterizing several MPEG compressed video sequences using the BIND model and applying the tighter bounding techniques used in [15]. We show that, contrary to common belief, reasonable network utilization can be achieved for compressed video even when deterministic performance guarantees are provided. To further quantify the improvements of the new techniques, we define the Deterministic Multiplexing Gain (DMG): even if all packets are deterministically guaranteed to meet their loss and delay bounds, sources may be multiplexed beyond their peak rate. The DMG may thus be defined as the fraction above a peak-rate allocation scheme that is achieved while still providing deterministic delay and throughput bounds.

The remainder of this paper is organized as follows. In Section 2, we discuss the background of the problem and describe the underlying requirements of a source model. In Section 3, we define and analyze the new model present the new admission control tests. In Section 4 we investigate the performance of the new model using parameters derived from actual MPEG traces. Finally, in Section 5, we discuss some of the practical issues for the model, and in Section 6, we conclude.

2 Background

In the analysis presented here, we investigate deterministic performance guarantees. While it is often assumed that the "worst-case" approach will result in low network utilization, we show that this is not necessarily so.

In order to provide deterministic performance guarantees, it is required that the source have a deterministic constraint. In [5], the constraints are that source j's packet spacing must not be less than its specified $Xmin_j$, in an interval of length I_j, it may send no more than $I_j/Xave_j$ packets, and that the maximum packet size is $Smax_j$. With such a constraint on each source, a worst-case analysis can provide absolute bounds on delay and throughput. In [2], a source is said to satisfy (σ, ρ) if during any time interval of length u, the number of bits that the source transmits is less than $\sigma + \rho u$. This leaky-bucket model is also sufficient for providing deterministic guarantees.

In [16], we propose a statistical bounding interval dependent model. In this 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 much nearer its long term average rate. In this study, we extend this model to the deterministic case by characterizing a source by a $B(I)$ BIND curve that provides a *deterministic* bound on a source's rate for intervals of length I. Analogous to the stochastic BIND model, the deterministic BIND model captures the property that for small interval lengths, a source's bounding rate approaches its peak rate, whereas for longer interval lengths, the source's bounding rate approaches its long term average rate from above. By capturing this important property of sources, higher network utilization is achievable.

There are two components to the approach described in [5] to providing performance guarantees in an integrated services network: connection admission control to limit the maximum real-time traffic load in the network, and a packet service discipline that provides the required performance during data transmission. In this paper, we assume a static priority scheduler such as RCSP [14] in the network and focus on the traffic model and admission control algorithms.

3 Deterministic BIND Model

As motivated in Section 2, the Deterministic BIND (D-BIND) model captures the interval length dependent behavior of sources as in [16] while still providing a deterministic traffic specification that is needed to provide deterministic performance guarantees.

3.1 Model Definition

More formally, the D-BIND model may be defined as follows. Source j, may be described by the curve $B_j(I)$ where $B_j(I)$ is the bounding rate over an interval of length I. Dropping the source j subscript, if $r(\tau)$ is the source's instantaneous rate at time τ (in bits or cells per second), then

$$\frac{1}{I}\int_{t}^{t+I} r(\tau)d\tau \leq B(I) \quad \forall\ t, I > 0. \tag{1}$$

Thus, the source is constrained to transmit no more than $I \cdot B(I)$ bits during any interval of length I. Equation (1) is thus a deterministic constraint on the source and therefore provides a worst-case arrival pattern for a given $B(I)$ characterization (this will be derived below).

There are several points to note about this curve. First, for a given source, the general trend of the characterization curve is that for small I, $B(I)$ will approach the source's peak rate while for larger interval lengths, $B(I)$ approaches the source's long term average rate from above, where the long term average rate is defined empirically as $\lim_{t\to\infty} \frac{1}{t}\int_0^t r(\tau)d\tau$. This curve however, is not necessarily convex or even monotonically decreasing. (Although a curve that is not strictly decreasing may seem like an odd case, we demonstrate later with real traffic traces that this may indeed be the case.) Despite the lack of desirable mathematical properties, the general trend is that $B(I)$ 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.

3.2 Delay Analysis

The admission control criteria of Section 3.3 relies on the delay analysis techniques of [2] which may be described as follows. Figure 1, illustrates the different components of the analysis. The horizontal axis is time and the vertical axis is bits. The upper curves 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 the maximum delay of packet can be expressed as a function of the two curves. For example, the following two observations hold:

1. the maximum busy period provides an upper bound on any work-conserving server;
2. the maximum backlog divided by link speed provides an upper bound on delay for a FCFS server.

Delay bounds for other policies can also be expressed [2, 10, 1, 9].

In order to derive delay bound, a deterministic constraint is needed on each traffic source so that the total number of bits that arrive at the queue by a given time also has a deterministic bound. In [5], traffic is constrained with the $(Xmin, Xave, I, Smax)$ model described in Section 2. With this model, a worst case arrival pattern may be constructed in which a source sends at its peak rate $(1/Xmin)$ until it is forced to stop transmitting because of its $I/Xave$ constraint. This is illustrated in Figure 2 with the $b(u)$ traffic constraint

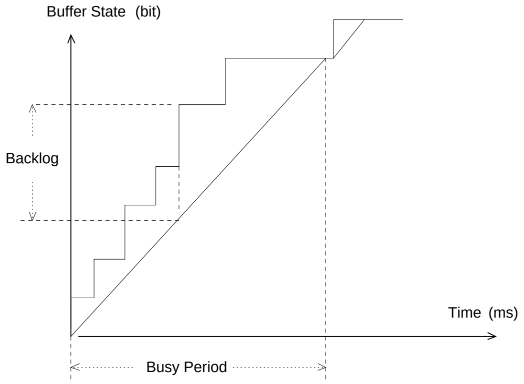

Figure 1: Concepts: Delay, Backlog and Busy Period

curve which represents the most bits that a source can send by time u. Clearly, the aggregate of the $b(u)$ curves from individual sources will determine the maximum backlog of the multiplexed stream of Figure 1.

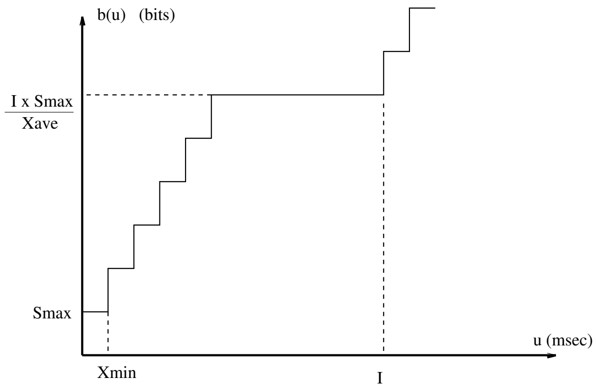

Figure 2: Traffic Constraint Curve for $(Xmin, Xave, I, Smax)$ Specification

For the $(Xmin, Xave, I, Smax)$ model, the constraint curve for connection j is given by

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

Theorem 1 of [15] shows that for a FCFS scheduler with $j = 1, ..., n$ multiplexed connections constrained by their respective constraint $b_j(u)$, a link speed l, and a maximum packet size $\overline{Smax}$, the resulting delay bound for all connections is given by

$$d = \frac{1}{l} \max_{\forall u \geq 0} \{ \sum_{j=1}^{n} b_j(u) - l \times u + \overline{Smax} \}. \tag{3}$$

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 than that in Equation (2), the resulting maximum delay bound of Equation (3) will also 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 $b(u)$. The effect is thus a higher DMG and a higher network utilization for a given deterministic delay and throughput constraint.

3.3 Admission Control

Two prerequisites to providing performance guarantees in a connection-oriented network are connection admission control and priority scheduling. An RCSP scheduler [14] provides the needed priority mechanisms. It 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. Of course, there is a limit to the number of connections that can be multiplexed at priority level k so that all packets from all connections have a delay bound of d_k. Therefore, when a new connection is requested with a delay bound d_k, the admission control algorithm must first check to 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 3.2, this may be achieved by calculating a bound on the maximum backlog.

The D-BIND model provides a deterministic constraint on the traffic source so that if source j is described by $B(I)$, the constraint function is

$$b(u) = B(u) \cdot u. \tag{4}$$

For a FCFS queue served at link speed l, Equation (3) gives a bound on the delay for all of the multiplexed connections. This result may be extended to the RCSP scheduler with the following theorem.

Theorem 1 *Assume an SP scheduler has n priority levels. Let C_q be the set of the connections at level q, and the j^{th} connection in C_q satisfies the traffic constraint function $b_{q,j}(.)$. Also assume that the link speed is l, and the size of the largest packet that can be transmitted onto the link is $\overline{Smax}$. The maximum delay of any packet at priority level m is bounded above by d^m, where*

$$d^m = max\{u : u \geq 0, b'_m(u) \geq l \times u\}$$

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

$$b'_m(\alpha) = \max_{\beta \geq 0}\{\overline{Smax} + \sum_{j \in C_m} b_{m,j}(\beta) + \sum_{q=1}^{m-1} \sum_{j \in C_q} b_{q,j}(\alpha + \beta) - l \times \beta\} \tag{5}$$

The proof is by extension of the results of [2]. Details may also be found in [13].

4 Traffic Characterization and Switch Utilization

This section evaluates the effectiveness of the D-BIND model. This is achieved by analyzing the switch utilization achieved with the new model and comparing the results to utilizations obtained with the $(Xmin, Xave, I, Smax)$ model (henceforth referred to as the $Xmin$ model) using the bounds in [15]. 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 speakers 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.

4.1 Deterministic Characterization

Figure 3 shows a typical segment of the traces of the two video streams. The vertical axis is rate in Mbps and the horizontal axis is frame number. It is assumed that the entire frame is transmitted per frame time (as opposed to introducing additional delay by smoothing over several frames) so that Figure 3 is simply the frame size multipled 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 ($\frac{1}{30}^{th}$ of a second). While neither of these assumptions is necessary, these transmission policies must be specified.

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 [6].

The frame pattern for Figure 3 is IBBPBB. Which frames are which becomes immediately apparent since the I frames tend to be the largest, B the smallest, and P in between. Further, the advertisement sequence,

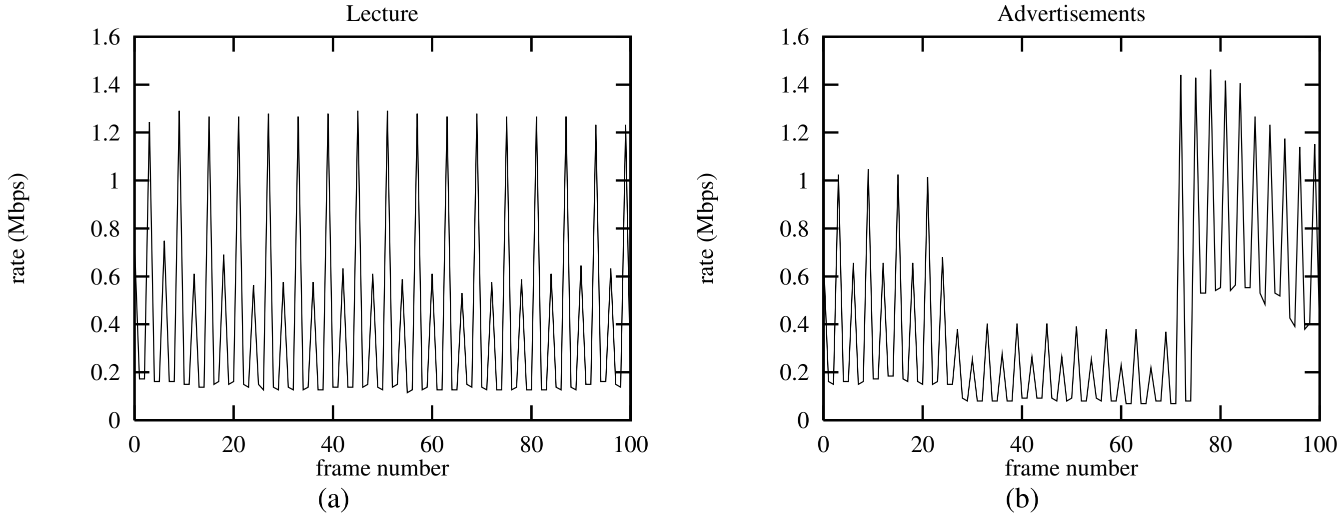

Figure 3: Segment of Trace

shown in Figure 3(b) tends to be much less regular than the lecture scene. This is as expected since the advertisement sequence has many different quickly changing scenes. The entire sequences used for the analysis below are 10 minutes of length. The segments of Figure 3 are chosen for illustration because of their representativeness of their respective sequences (i.e., they were chosen at random and are not 'worst case' sequences).

Figure 4 shows the B(I) curves for the advertisement and lecture sequences. The vertical axis is the worst case average rate averaged over an interval of length I where I is on the horizontal axis with units of frame times (again, the frame rate is 30 fps). This is a deterministic bound on the average rate since it is an upper bound on the average rate over *any* interval of length I. The bounding average 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 maximum average rate for $I = 1$ frame time is the peak rate ($Smax/Xmin$). The reason for this is that maximum average 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 maximum average rate decreases with increasing interval length. This is the intuitive property described in Section 3 that the bounding rate of a