In Proceedings of IEEE INFOCOM'96

H-BIND: A New Approach to Providing Statistical Performance Guarantees to VBR Traffic

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

Abstract

Current solutions to providing statistical performance guarantees to bursty traffic such as compressed video encounter several problems: 1) source traffic descriptors are often too simple to capture the burstiness and important time-correlations of VBR sources or too complex to be used for admission control algorithms; 2) stochastic descriptions of a source are inherently difficult for the network to enforce or police; 3) multiplexing inside the network's queues may change the stochastic properties of the source in an intractable way, precluding the provision of end-to-end QoS guarantees to heterogeneous sources with different performance requirements. In this paper, we present a new approach to providing end-to-end statistical performance guarantees that overcomes these limitations. We term the approach Hybrid Bounding Interval Dependent (H-BIND) because it uses the Deterministic-BIND traffic model to capture the correlation structure and burstiness properties of a stream; but unlike a deterministic performance guarantee, it achieves a Statistical Multiplexing Gain (SMG) by exploiting the statistical *properties of deterministically-bounded streams. Using traces of MPEG-compressed video, we show that the H-BIND scheme can achieve average network utilizations of up to 86% in a realistic scenario.*

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- and loss-sensitive Variable Bit Rate (VBR) traffic poses a unique challenge, since providing good Quality of Service (QoS) to such bursty traffic sources is at odds with achieving high network utilization. In the scheme of [8], two types of guarantees are proposed to support VBR traffic: *deterministic* guarantees, in which loss, delay and throughput are guaranteed even in the worst case, and *statistical* guarantees, in which the network provides probabilistic guarantees on loss and delay.

This research was supported in part by Sandia National Laboratories.

A statistical guarantee utilizes statistical multiplexing to achieve higher network utilization at the expense of occasionally dropping or excessively delaying packets. Previous approaches to providing such guarantees utilize explicit stochastic properties of the arrival processes and calculate or approximate quantities such as the steady state buffer overflow probability (see [6, 9, 11, 13, 16, 24] for examples). For an approach to be used as an admission control algorithm, it must be able to efficiently calculate the loss and delay-bound-violation probability for a new source, given the traffic specifications and performance requirements of all the currently established connections. Because of a multiplexer's nonlinearity, solving for these admission control conditions often becomes quickly intractable and leads to several problems. First, stochastic models of the source are often too complex to be used for admission control or too simple to capture the burstiness and important time-correlations of VBR sources such as compressed video. Second, a stochastic description of a source is inherently more difficult for the network to enforce than a deterministic description such as a peak rate. That is, the network would like to police the connection to ensure that it adheres to its specified traffic description. Third, providing *end-to-end* performance guarantees is especially difficult since the original stochastic model of the source is no longer valid after the stream traverses even a single multiplexer. Finally, because of analytical difficulties, many solutions cannot deal with both heterogeneous sources and heterogeneous QoS requirements (i.e., priority schedulers), both of which are needed to provide integrated services.

Several recent works [14, 15] indicate that in many cases, reasonable network utilizations can be achieved even while providing *deterministic* performance guarantees to VBR video. Specifically, using traces of MPEG compressed video, [15] showed that using the Deterministic Bounding Interval Dependent (D-BIND) traffic model, utilizations on the order of 15% to 30% are achievable for queueing delay bounds in the range of 30 msec. While such utilizations are considerably higher than those achieved with a peak-rate-allocation scheme, there are fundamental limitations to the utilizations achievable when providing deterministic service [14], and the numbers cited above are quite close to that limit. To achieve higher utilization, statistical multiplexing must be used.

In this paper, we propose a new approach to providing statistical performance guarantees that overcomes many of the limitations of previous approaches. As in [15], the approach uses the D-BIND traffic model to characterize a stream's burstiness properties and temporal correlation structure. However, unlike a deterministic approach, we do not consider the aggregate process as a worst case aggregation of individual sources. That is, a deterministic performance guarantee requires that even if sources multiplex so that all sources are *dependent* and exactly synchronize to each other in the worst possible way, the guarantees are still met. Alternatively, in this work we provide a *statistical* performance guarantee by taking advantage of statistical properties of deterministically constrained sources, including the maximum variance of a D-BIND-policed stream, and statistical independence of the sources' *phases*. The goal is to provide a statistical service with many of the advantages of the deterministic approach in [15], without most of the disadvantages of the statistical approaches in [9]. We term the new approach Hybrid Bounding Interval Dependent (H-BIND) because it combines the deterministic modeling techniques of the D-BIND model with the stochastic aggregation approaches of [16, 24]. The approach will be shown to be fully enforceable by the network, and to provide end-to-end statistical performance guarantees to heterogeneous sources with heterogeneous QoS requirements. We provide an admission control test for a static priority scheduler that supports both deterministic and statistical service with multiple delay bounds for each service. Hence, the H-BIND statistical test takes into account the presence of higher-priority traffic, including streams utilizing a deterministic service. This is possible because the same traffic descriptor is used by all streams that obtain a QoS guarantee. Finally, the admission control calculations can be performed relatively quickly and avoid the convolutions of [16, 24].

We evaluate the performance of the H-BIND admission control algorithm using a 28-minute trace of MPEG compressed video. We show that utilizations of up to 86% are achievable in a realistic scenario, and that the admission control test, while slightly conservative, closely approximates the performance observed in trace-driven simulations.

A related problem was studied in [3, 4, 12, 20] via analysis of the nD/D/1 queue, considering periodic sources with uniformly distributed independent phases. Using various techniques, queue length distributions were derived for the case of sources with identical periods in [3, 4, 12] and *bounds* for the queue length distribution for the case of heterogeneous periods in [20]. The nD/D/1 queue is unlike the current approach in that, rather than assuming sources are deterministic, we are using deterministic bounds on sources as provided by the D-BIND model. Moreover, in our analysis, the sources are more general than the periodic model of the nD/D/1 queue. Finally, in this work we consider QoS bounds for priority service disciplines such as RCSP [21] rather than for FCFS. Such a service discipline is important to providing integrated services to heterogeneous sources.

The recent works of [5, 19] also consider stochastic bounds for deterministically constrained sources. For example, in [19] it was shown that for processes constrained by a multi-level leaky-bucket, the Chernoff estimate of the loss-probability is maximized by a periodic on-off process with a random phase. The techniques of [5, 19] share many of the properties of the H-BIND approach presented here. For example, both techniques investigate stochastic properties of deterministically constrained streams, and both techniques find a deterministic on-off process with random phase to be extremal: for [19], an on-off process maximizes the Chernoff loss estimate, and for H-BIND, an on-off process yields the maximal variance over every interval length for D-BIND constrained sources. Our approach differs from [19] in that the D-BIND model provides a more accurate bound on a traffic stream than multiple leaky buckets (indeed, multi-level leaky buckets are a special case of D-BIND [15]), and in our use of the Central Limit Theorem rather than Chernoff's Theorem. Finally, the approach of [5, 19] has not yet been validated with realistic workloads.

The remainder of this paper is organized as follows. In Section 2, we review the D-BIND model and its application to providing deterministic performance guarantees in an integrated services environment. In Section 3, we present the analytical foundations of the H-BIND approach, and in Section 4, we evaluate the approach using a 28 minute trace of MPEG compressed video.

2 A Deterministic Approach

A deterministic service provides the best QoS with its no-drop, no-delay-bound-violation guarantee. In [1, 2, 17, 22], it was shown that peak-rate-allocation is not required to provide deterministic guarantees to VBR traffic, and in [15], it was demonstrated that the achievable utilization in a deterministic setting is largely determined by the choice of the deterministic traffic model.

2.1 D-BIND Model

As shown in [15], previous deterministic traffic models such as the (σ, ρ) model [2] and the $(Xmin, Xave, I, Smax)$ model [8] cannot capture the property that sources exhibit burstiness over a wide variety of interval lengths. The D-BIND traffic model was introduced to address this issue. The key components of the D-BIND model are that it is *bounding*, required to provide deterministic QoS guarantees, and *interval-dependent*, needed to capture important burstiness properties of sources. This more accurate traffic characterization then allows for a higher network utilization for a given delay bound (see [15]).

Each deterministic traffic model uses parameters to define a traffic constraint function $b(t)$, which constrains or bounds the source over every interval of length t. Denoting by $A[t_1, t_2]$ the number of 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 timeinvariant deterministic bound since it constrains the traffic source over every interval of length t. For example, the (σ, ρ) model is defined so that $A[s, s+t] \leq \sigma + \rho t$ for all $s, t > 0$.

The D-BIND model is defined via multiple rate-interval pairs $\{(R_k, I_k) | k = 1, 2, \cdots, P\}$, so that the constraint function is given by a piece-wise linear function:

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

with $b(0) = 0$. Thus, the rate R_k can be viewed as an upper bound on the stream's rate over every interval of length I_k, so that $A[t, t+I_k]/I_k \leq R_k \quad \forall\, t > 0, k = 1, 2, \cdots, P$.

Figure 1 shows a plot of the D-BIND bounding rate vs. interval length for a 28 minute trace of an MPEG-compressed action movie. The trace is taken from a "James Bond" movie, and was digitized to 384 by 288 pixels and compressed at 24 frames per second with frame pattern IBBPBBPBBPBB. The compression algorithm is constant-quality MPEG 1 compression performed in software (see [10] for further details of MPEG 1). Plotting the bounding rate vs. interval length, the figure shows that the model captures the source's burstiness over multiple interval lengths. For example, for small interval lengths, R_k approaches the source's peak rate, $R_1 = 5.87$ Mbps. For longer interval lengths, R_k approaches the long term average rate of 583 kbps, which is the total number of bits in the MPEG sequence divided by the length of the sequence – 28 minutes in this case. In practice, we expect that a source will specify a small number of rate-interval pairs (e.g., four or eight) for connection admission control.

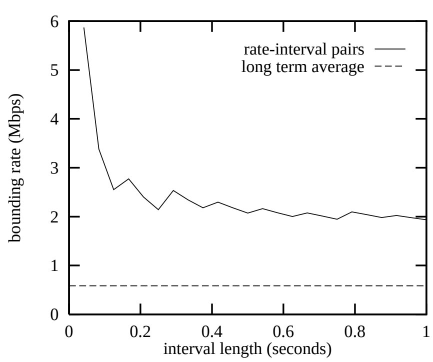

Figure 1: D-BIND Curve for Action Movie

2.2 Connection Admission Control

With the constraint function's bound on arrivals in any interval of length t, the sum of the individual sources' respective $b(t)$ constraint functions will yield an upper bound on aggregate arrivals to the queue. With this bound on aggregate arrivals, admission control conditions for deterministic delay and throughput bounds may be derived.

Below, we demonstrate how heterogeneous, deterministic, QoS guarantees can be provided to heterogeneous sources by using the Rate-Controlled Static Priority (RCSP) Service Discipline [21]. The scheduler is based on a number of prioritized FCFS queues, where queue p has an associated delay bound d_p. A number of heterogeneous connections can be multiplexed at a given priority level p. The maximum number of schedulable connections is determined by the following theorem.

Theorem 1 *Assume a Static Priority scheduler has n priority levels, where priority level p has an associated delay bound d_p. Let C_q be the set of connections at level q, and the j^{th} connection in C_q satisfies the traffic constraint function $b_{q,j}(\cdot)$. For a link speed l, the set of connections is schedulable if the following condition holds for all priorities p and for all $0 \leq t \leq \beta_p$:*

$$l(t + d_p) \geq \sum_{j \in C_p} b_{q,j}(t) + \sum_{q=1}^{p-1} \sum_{j \in C_q} b_{q,j}(t + d_p) \tag{2}$$

where β_p is a bound on the priority-p busy period.

The proof is given in [17]. The statistical bounds of Section 3 will be shown to have a form similar to that of Equation (2).

3 The H-BIND Approach

Compared to a deterministic approach, a statistical approach achieves higher network utilization at the expense of dropping or excessively delaying some small fraction of packets. In the H-BIND statistical approach described below, we achieve a statistical multiplexing gain by exploiting statistical independence of sources' phases, by deriving statistical properties of streams that are policed or deterministically constrained, and by using the D-BIND model to capture a stream's correlation structure.

In Section 3.1, we derive a stochastic bound on delay for streams multiplexed at a static priority scheduler. To make this bound practical as a connection admission control algorithm, we proceed in several steps. First, in Section 3.2 we show how the D-BIND model captures a stream's correlation structure. Next, since the framework of bounding *distributions* has practical limitations, in Section 3.3 we bound the *variance* of a process that is policed or constrained by D-BIND parameters. Finally, in Section 3.4, we utilize the Central Limit Theorem (CLT) to provide a more practical admission control condition.

3.1 Bounding Delay in a Static Priority Scheduler

A static priority scheduler as in [21] consists of a number of prioritized FIFO queues. Each connection is assigned a priority p at connection setup time. The assigned priority level is based on the connection's requested QoS, including whether it has requested deterministic or statistical service, and on the requested delay bound. As shown in Figure 2, such a scheduler can support both deterministic and statistical performance guarantees. Priority levels det_1 through det_m provide m deterministic delay bounds from d_{det_1} up to d_{det_m}. Priority levels $stat\_1$ through $stat\_m$ provide m statistically guaranteed delay bounds from d_{stat_1} up to d_{stat_m}. Connections utilizing the statistical service obtain guarantees on the loss and delay-bound-violation *probability* whereas connections utilizing the deterministic service obtain absolute bounds on delay and loss. For connections utilizing deterministic service, the admission control test utilizes Equation (2).

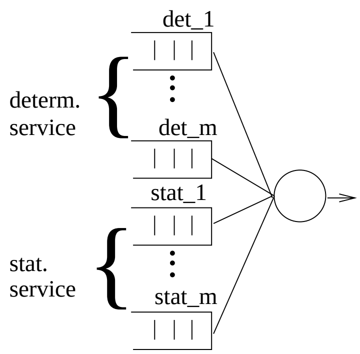

Figure 2: Static Priority Scheduler

In [16], Kurose introduced a framework for providing statistical network performance bounds based on sources characterized by a family of bounding random variables. In that work, stream j is characterized by a family of random variables $\{(B_j(t_1), B_j(t_2), ...\}$ that stochastically bound the source over the respective interval lengths t_k. A random variable B is said to be stochastically larger than a random variable S (denoted $B \succeq_{st} S$) if and only if $Prob(B > x) \geq Prob(S > x)$ for all x. A stochastic bound on delay for a FCFS multiplexer with link speed l serving N connections characterized by their respective families of bounding random variables is shown in [16] to be:

$$Prob\{D > d\} \leq \max_{0 \leq t_k \leq \beta} Prob\{\sum_{j=1}^{N} B_j(t_k) - lt_k \geq ld\} \quad (3)$$

where D is a random variable representing the delay of a randomly selected cell and β is an upper bound on the busy period (calculation of this is discussed below).

The above result can be extended to a static priority service discipline by combining the techniques of [2, 16, 17]. The following theorem provides a stochastic bound on the probability that a random cell of connection j in priority level p ($p \in stat\_1, \cdots, stat\_m$) is delayed beyond its bound d_p.

Theorem 2 *Assume a Static Priority scheduler has n priority levels, where priority level p has an associated delay bound d_p. Let C_q be the set of connections at level q, and let the distribution of the number of cells transmitted by the j^{th} connection in C_q over intervals of length t be stochastically bounded by the random variable $B_{q,j}(t)$. For a link speed l, the delay violation probability for a random cell in level p is bounded by:*

$$Prob\{D_p > d_p\} \leq \max_{0 \leq t_k \leq \beta_p} Prob\{\sum_{j \in C_p} B_{p,j}(t_k) + \sum_{q=1}^{p-1} \sum_{j \in C_q} B_{q,j}(t_k + d_p) - lt_k \geq ld_p\} \quad (4)$$

where β_p is a bound on the priority-p busy period.

Proof: Consider a random cell from level p that arrives in a given busy period. Both cells from priority level-p that arrive before the random cell, and cells that arrive in a higher priority level will contribute to the cell's delay. For the previous t_k seconds, the distribution of the number of level-p cells that arrive (and hence contribute to the random cell's delay) is stochastically bounded by $\sum_{j \in C_p} B_{p,j}(t_k)$, even in the worst case that all of these cells arrive within the random cell's busy period. This was demonstrated in [16] for the FCFS result of Equation (3). For the previous t_k seconds, cells from higher priority levels may also arrive and contribute to the random cell's delay. Moreover, higher priority cells may contribute to a delay bound violation for the random cell even if they arrive *after* the random cell by up to d_p seconds. The distribution of this total number of higher-priority cells is bounded by $\sum_{q=1}^{p-1} \sum_{j \in C_q} B_{q,j}(t_k + d_p)$. Since lt_k cells are served in this interval, that many fewer are queued. Thus, a stochastic bound on the total number of cells that arrived no more than t_k seconds ago, have priority over the random cell, and are queued is given by $\sum_{j \in C_p} B_{p,j}(t_k) + \sum_{q=1}^{p-1} \sum_{j \in C_q} B_{q,j}(t_k + d_p) - lt_k$. Since cells that arrive in other busy periods will not contribute to the random cell's delay, maximizing over all interval lengths less than the maximal busy period yields the result of Equation (4). $\Box$

Unfortunately, Equation (4) of Theorem 2 has several problems that preclude its practical use as a connection admission control condition. First, Equation (4) requires knowledge of a family of bounding distributions for each traffic stream. While these distributions can be parameterized using techniques such as those in [24], even in [24], the sum of random variables in Equation (4) must be calculated via a convolution or a Fast Fourier Transform, which may be too costly. A further problem (shared, but not usually addressed by many statistical approaches) is that to provide heterogeneous services, i.e., to provide both deterministic and statistical services with multiple delay bounds, the streams in Equation (4) will include some that are utilizing a deterministic service. Since connections receiving a deterministic service have priority over those utilizing a statistical service, to calculate the delay-violation probability for a statistical connection requires that even streams receiving a deterministic service would have to specify some statistical properties to the network.

This is potentially a problem since streams utilizing deterministic service typically specify only worst-case traffic parameters such as (σ, ρ) or D-BIND. Finally, as is the case for most statistical techniques, a source's promised traffic specification (of bounding distributions) cannot be efficiently enforced by the network.

The H-BIND approach described below uses the D-BIND traffic descriptor for both deterministic and statistical performance guarantees to alleviate these latter problems. The computational complexities of Equation (4) will also be addressed with the introduction of a simplified admission control test based on the Central Limit Theorem.

3.2 D-BIND and Source Correlation Structure

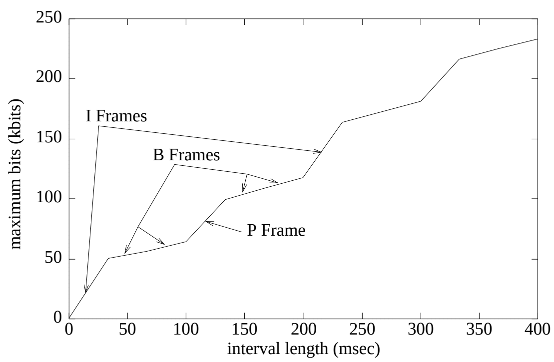

Figure 3: D-BIND Constraint Function

To motivate the use of the D-BIND characterization, even for streams receiving a statistical performance guarantee, Figure 3 shows how the D-BIND model captures a stream's correlation structure. The figure shows the D-BIND constraint function $b(t)$ (described by Equation (1)) for a 10 minute trace of a lecture sequence. As shown in the figure, the temporal properties of the MPEG source are captured by the D-BIND model's constraint function. The figure shows the maximum number of *kbits* that the source transmits over any interval of length t (shown in seconds). The initial slope (at $t = 0$) 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 always followed by a B frame, which is typically much smaller than an I frame due to MPEG's exploitation of temporal redundancies. At 100 msec, after sending the I frame and two B frames, the constraint function breaks up again, indicating the transmission of a P frame. Thus, the coder's alternation between I, P, and B frames in the pattern IBBPBB is evident from the D-BIND $b(t)$.

3.3 Bounding the Second Moment of D-BIND-constrained Streams

Since bounding the *distribution* of a D-BIND-constrained stream has practical limitations for admission control, below we bound the maximum *variance* of the D-BIND-constrained stream, maximizing the variance over every interval length.

Consider a stream that is bounded by rate-interval pairs $\{(R_k, I_k) | k = 1, 2, \cdots, P\}$. Hence, the stream's arrivals over every interval of length t, $A[s, s+t]$ is bounded by the associated constraint function $b(t)$ as given by Equation (1). While the constraint function provides a *deterministic* bound on the stream's arrivals, to achieve a statistical multiplexing gain using a statistical service, we are interested in the stochastic properties of such a deterministically constrained source. Note that there are an infinite number of arrival processes $A[0, t]$ that conform to a given traffic specification $\{(R_k, I_k) | k = 1, 2, \cdots, P\}$. To calculate the stochastic bound in Equation (4) for D-BIND constrained sources, we would want to know the stochastically greatest bounding distributions $B(t_k)$ that conform to the deterministic constraints. However, once these distributions are determined, such a formulation would still have the problem that convolutions would be required in the calculation of Equation (4). For this reason, we derive a bound on the *variances* of the $B(t_k)$ random variables subject to the $b(t)$ constraint. We then utilize these variances in the CLT-based admission control presented in the next section.

For simplicity, consider time to be slotted, and denote a stream's arrivals as a sequence of positive real numbers $\{X_{t_1}, X_{t_2}, \cdots, X_{t_N}\}$. For a video sequence, time can be considered to be slotted to the frame time, e.g., 1/24 seconds, so that the X_{t_i}'s represent the sizes of the compressed video frames. Otherwise, time can be considered slotted by the link's cell transmission time. The X_{t_i}'s are constrained by the rate-interval pairs $\{(R_k, I_k) | k = 1, 2, \cdots, P\}$. The rate-interval pairs define the constraint function $b(t)$, which we will denote for slotted time as the sequence $\{b_{t_1}, b_{t_2}, \cdots, b_{t_{P'}}\}$. Note that $t_{P'}$ can be much larger than I_P, or equivalently, I_P can be much larger than P time slots, since the rate-interval pairs simply define a piece-wise linear function and have no underlying slotted structure. Indeed, as described in Section 4, I_P should be as large as possible to capture the long-time-scale behavior of the process. But we expect P, the number of rate-interval pairs, to be on the order of four or eight.

Thus, we are interested in the sequence $\{X_{t_1}, X_{t_2}, \cdots, X_{t_N}\}$ that maximizes the sequence's empirical variances over every interval length, subject to the constraints $\{b_{t_1}, b_{t_2}, \cdots, b_{t_{P'}}\}$. Over an interval of length one time slot, the rate-variance is given by

$$\sigma^2(t_1) = \frac{1}{N-1} \sum_{i=1}^{N} \left(\frac{X_{t_i} - \mu}{t_1}\right)^2 \tag{5}$$

where μ is the empirical mean,

$$\mu = (\frac{1}{N}) \sum_{i=1}^{N} X_{t_i}. \tag{6}$$

The variance of the sequence's arrival rate over intervals of length k is given by:

$$\sigma^2(t_k) = (\frac{1}{N-k}) \sum_{i=1}^{N-k+1} \left(\frac{\sum_{m=0}^{k-1} X_{t_i+t_m} - k\mu}{t_k} \right)^2 \tag{7}$$

Theorem 3 *The sequence of positive real numbers* $\{X_{t_1}, X_{t_2}, \cdots, X_{t_N}\}$ *that maximizes the empirical variances over every interval length*

$$\max_{1 \leq k \leq P'} \sigma^2(t_k) \tag{8}$$

subject to the constraints:

$$X_{t_i} \leq b_{t_1} \tag{9}$$

$$X_{t_i} + X_{t_{i+1}} \leq b_{t_2} \tag{10}$$

$$\vdots$$

$$\sum_{k=0}^{P'-1} X_{t_{i+k}} \leq b_{t_{P'}} \tag{11}$$

consists of $X_{t_i} = 0$ *or* $X_{t_i} = b_{t_1}$*, with the* b_{t_1}*'s spaced as closely as possible subject to the* b_{t_k} *constraints of Equation (10) - (11).*

The proof, omitted because of space constraints, proceeds in two steps. First, one can show that $\sigma^2(t_k)$ is maximized when $k = 1$. Thus, maximizing Equation (8) is equivalent to maximizing Equation (5) subject to the constraints. Second, since Equation (5) is separable and convex, its maximization occurs at the maximal constraint boundaries, in this case 0 or b_{t_1}.

Theorem 8 therefore states that the process that maximizes the empirical variances is a type of on-off process that either transmits at its peak rate R_1 or transmits nothing. For example, the worst-case sequence could be $\{b_{t_1}, 0, 0, b_{t_1}, 0, 0, 0, b_{t_1}, 0, 0, b_{t_1}, 0, 0, 0, 0, \cdots\}$ where the number of interspaced 0's are determined by the D-BIND parameters.

3.4 H-BIND Admission Control

In this section, we combine the results of Sections 3.1-3.3 with the Central Limit Theorem to provide the H-BIND connection admission control test.

The Central Limit Theorem states that the sum of N independent random variables converges in distribution to a random variable that has a Normal distribution ([18], page 287). The bound of Equation (4) states that the delay violation probability at priority p is bounded by $\max_{0 \leq t_k \leq \beta_p} Prob\{\sum_{j \in C_p} B_{p,j}(t_k) + \sum_{q=1}^{p-1} \sum_{j \in C_q} B_{q,j}(t_k + d_p) - lt_k \geq ld_p\}$. Hence, the CLT states that the random variable $\sum_{j \in C_p} B_{p,j}(t_k)$ converges to $\hat{B}_p(t_k)$, which is Normally distributed with mean $\sum_{j \in C_p} k\mu_j$ and variance $\sum_{j \in C_p} t_k^2 \sigma_j^2(t_k)$. As well, the random variable $\sum_{q=1}^{p-1} \sum_{j \in C_q} B_{q,j}(t_k + d_p)$ converges to Normally Distributed $\hat{B}_{q<p}(t_k + d_p)$ with mean $\sum_{q=1}^{p-1} \sum_{j \in C_q} (k + \frac{d_p}{t_1}) \mu_{q,j}$ and variance $\sum_{q=1}^{p-1} \sum_{j \in C_q} (t_k + d_p)^2 \sigma_j^2(t_k + d_p)$. The means and variances are as calculated by Equations (6) and (7) respectively, with the X_{t_i}'s as given by Theorem 3. β_p can be calculated as in [1] as $\beta_p = min\{k \geq 1 : \sum_{j \in C_p} b_j(t_k) \leq lt_k\}$.

The H-BIND connection admission control test is summarized in the following proposition.

Proposition 1 *For a static priority scheduler multiplexing D-BIND constrained sources, the delay violation probability is approximated by*

$$Prob\{D_p > d_p\} \approx \max_{0 \leq t_k \leq \beta_p} Prob\{\hat{B}_p(t_k) + \hat{B}_{q<p}(t_k + d_p) - lt_k \geq ld_p\} \tag{12}$$

where the $\hat{B}$*'s are Normally distributed random variables with upper bounds on the means and variances determined as above from the D-BIND parameters.*

3.4.1 Discussion

While the admission control test of Equation (12) is approximate because of the CLT approximation, it tends to be conservative because the variances of $\hat{B}_p$ and $\hat{B}_{q<p}$ are calculated with upper bounds on the variances of D-BIND constrained sources as in Theorem 3. To efficiently bound the tail *probability* (rather than bounding the variances and using the CLT), one could use, for example, Chebychev's Inequality together with the variance bounds of Theorem 3. However, the Chebychev bound is too loose to be useful in practice. The validity of the approximation in Proposition 1 is further explored via trace-driven simulation in Section 4.

We also note that Proposition 1 allows all connections, including those receiving both deterministic and statistical service, to use the same D-BIND traffic descriptor. For a static priority scheduler, the D-BIND admission control test of Equation (2) will be used for connections requesting a deterministic service, and the H-BIND test of Equation (12) for connections requesting a statistical service. This provides a unified framework for providing heterogeneous services to heterogeneous types of traffic since all sources can utilize a single traffic descriptor. The D-BIND descriptor was chosen for its accuracy in its ability to capture both a stream's correlation structure (Figure 3) as well as its burstiness over multiple interval lengths (Figure 1). As well, the techniques presented here could be applied to other deterministic traffic models such as the multi-level leaky bucket, which is a special case of D-BIND [15].

By utilizing the CLT, statistical independence of streams' phases is exploited to achieve a statistical multiplexing gain. Further, computationally expensive matrix calculations or convolutions are avoided. Instead, a bound on the mean and variance of a stream's rate distribution over multiple interval lengths are obtained from the source's specified rate-interval pairs. We also note that the probability calculation of Equation (12) will in general include streams that are utilizing a deterministic service. As mentioned above, this does not pose a problem since all streams specify only deterministic constraints. A further benefit of using only deterministic traffic specifications is that such constraints can be enforced by the network, unlike traditional stochastic traffic descriptors.

3.4.2 End-to-End Guarantees

The above single hop delay-violation probability can be extended to an end-to-end probability using the delay-jitter control techniques of [7, 24]. Delay-jitter control decouples the network nodes by reconstructing the sources' original traffic pattern at each hop, effectively absorbing the delay-jitter introduced by the previous hop. For example, [24] showed that if a connection traverses a path of servers with delay-jitter controlling regulators, the traffic pattern at the input of each scheduler is *exactly the same* as that at the entrance to the network. This allows us to analyze each scheduler using the *same* traffic characterization. Thus, if a delay bound d_h is provided at hop h with violation probability ζ_h, then the end-to-end delay-bound-violation probability is $Prob\{D_{e2e} > \sum d_h\} \leq \prod \zeta_h$ (see also [8]).

Additionally, for H-BIND statistical guarantees and D-BIND deterministic guarantees, the simpler *rate*-jitter control mechanisms for D-BIND (e.g., [15, 21]) may be sufficient in practice. That is, a D-BIND rate-jitter-controller at each hop will ensure that each stream conforms to its original D-BIND characterization at each hop along the connection's path. For H-BIND's statistical guarantees, the rate-controllers will not strictly preserve statistical independence of the streams' phases as assumed by the admission control test. However, rate-jitter-control is considerably simpler than delay-jitter control (for a special case of D-BIND, rate-jitter-control may be implemented with buffered multi-level leaky buckets), hence it may be a practical approximate solution to the problem of providing end-to-end performance guarantees.

4 Evaluation of H-BIND from MPEG Traces

In the experiments below, we use trace-driven simulation to evaluate the H-BIND scheme. As described in Section 2, the trace is 28 minutes long and consists of an MPEG-compressed action movie. We consider a 45 Mbps link with packets served according to the RCSP service discipline. With homogeneous sources with the same QoS requirements, the service discipline transmits packets in the same order as the FCFS policy, but with rate controllers to enforce the traffic specification. In the first part of the experiment, we parameterize the traces using eight D-BIND rate-interval pairs. For a "live" sequence, where the trace is not available in advance, [23] presents an on-line algorithm for obtaining (and adapting) the rate-interval pairs. These rate-interval pairs are then used in the admission control condition of Equation (12) to calculate the maximum number of admissible connections for a given delay bound and delay-violation probability.

The second part of the experiment utilizes trace-driven simulation to compare the empirical delay bound violation probabilities obtained by multiplexing randomly offset traces, with those given by the admission control conditions. Figure 4 illustrates the simulation scenario. For a given simulation, a number of sources are multiplexed on the 45 Mbps link, with each stream's arrival pattern given by the movie trace, and its start time τ_j chosen uniformly over the length of the trace (28 minutes). The simulation runs until all sources have transmitted their entire trace twice, with the traces wrapped around to the beginning when they reach the end. Multiple simulations are performed with independent start times and average results are reported.

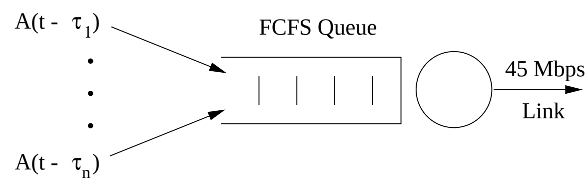

Figure 4: Scenario for Trace-Driven Simulation

For the experiments, we use average utilization of the multiplexer as a performance metric to evaluate the H-BIND scheme. This average utilization may be calculated as

$$\frac{\sum_{j=1}^{N} R_{j,28min}}{l} \tag{13}$$

where $R_{j,28min}$ is the long term average rate of source j, l is the link speed, and N is either the number of admissible connections or the number of simulated connections for the respective admission control and simulation experiments . For the simulation, the average utilization is also the total number of bits transmitted by the sources in the simulation, divided by the total number of bits that the server can transmit during the simulation (the link speed multiplied by the simulation time).

Figure 5 shows the results of the experiments. The curves show average utilization versus delay bound for the following scenarios. As a base-line, the figure shows the achievable utilization of 9% for a peak-rate allocation scheme, which corresponds to a resource allocation scheme that requires the sum of the sources' peak rates to be less than the link speed. The next curve in the figure is the achievable utilization for the D-BIND model (as in Section 2) with deterministic guarantees. For a deterministic delay bound as given on the horizontal axis, the vertical axis represents the resulting average multiplexer utilization so that the admission control condition of Equation (2) is satisfied. For example, for a deterministic queueing delay bound of 40 msec, 15 connections are admissible, which results in an average multiplexer utilization of 19%.

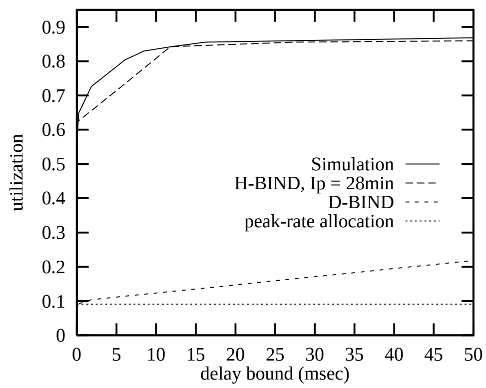

Figure 5: Achievable Utilization for H-BIND

The remaining two curves respectively depict the results for the H-BIND admission control tests and for the trace-driven simulations. For a delay bound violation probability of 10^{-6}, the curve labeled "H-BIND, $I_P = 28$min" depicts average utilization as a function of delay bound d_p such that $Prob\{D_p > d_p\} \approx 10^{-6}$ as calculated by the H-BIND admission control test of Equation (12). The curve was generated using eight D-BIND rate-interval pairs with the final interval length, I_P or I_8, the same as the length of the trace, 28 minutes. As shown below, this parameter has an important effect on the achievable utilization. The final curve of Figure 5, labeled "Simulation", shows the results of the trace driven simulations. Multiple simulations are performed aggregating different numbers of connections which correspond to different average utilizations as given by Equation (13). These simulations are then repeated and the curve depicts the average results. Specifically, for a given number of connections, the curve depicts the delay bound d_p such that an average fraction of 10^{-6} cells violated their delay bound.

From Figure 5, we note that the H-BIND admission control algorithm matches fairly well with the empirical results from the trace driven simulation. The difference between the two curves depends on the delay bound, with the admission control being most conservative over delay bounds less than 10 msec. Moreover, the utilizations achieved by the scheme are considerably high, reaching 86% for delay bounds under 40 msec. This utilization is well above that achievable with a deterministic guarantee. For example, at a 40 msec delay bound, the H-BIND scheme has 66 admissible connections for an 86% average utilization, compared to 15 admissible connections and a 19% average utilization for a deterministic service. Compared to the simulations, we conclude that the H-BIND test may be slightly conservative since it is based on upper bounds of streams' variances, but it can still achieve a considerable statistical multiplexing gain, yielding utilizations close to those obtained by the simulations.

Finally, Figure 6 explores the issue of a source's choice of I_P, the largest interval length for the source's specified rate-interval pairs (or I_8 for the 8 pairs used here). As described in Section 2.1, the worst case rate R tends to decrease with increasing interval length I. Hence, when a source chooses its rate-interval pairs to specify to the network, it should make I_P as large as possible to capture the long-time-scale properties of the stream. By choosing a large I_P, more connections can be accepted by the admission control test. For example, as a stability condition, we require that $\sum_{j \in \mathcal{A}} R_{j,8} < l$, so that if a source parameterizes its traffic with a larger I_8, it will have a smaller R_8. Hence, possibly a larger set of connections $\mathcal{A}$ can satisfy this stability condition. Indeed, Figure 6 shows that parameterizing the source with the longest possible I_P of 28 minutes provides the highest utilizations. For example, $I_8 = 28$ minutes results in an 86% utilization at a delay bound of 40 msec. A smaller I_8 of 6.9 minutes or 42 seconds reduces the achievable utilization to 71% or 40% respectively. The reason for this is that the longer-time-scale information is not as well reflected with these smaller choices of I_P or I_8.

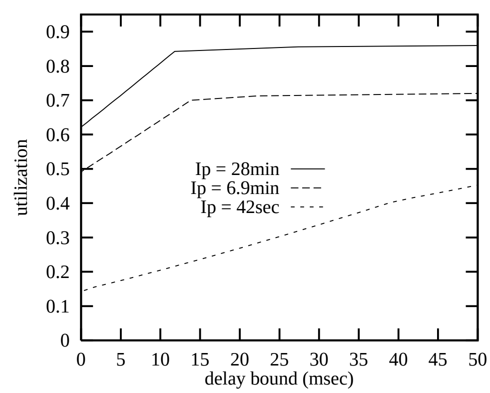

Figure 6: Effect of I_P on Achievable Utilization

5 Conclusion

In this paper, we have presented a new approach, termed H-BIND, to providing statistical performance guarantees in integrated services networks. We used the D-BIND deterministic traffic model to capture a stream's correlation structure and burstiness properties over multiple interval lengths. However, unlike a deterministic performance guarantee as in [15], the H-BIND approach exploits *statistical* properties of deterministically constrained sources. In this way, our approach is able to achieve considerable multiplexing gains while avoiding many of the disadvantages of traditional approaches to providing statistical performance guarantees. For example, with H-BIND, the traffic specification of multiple rate-interval pairs is fully enforceable by the network. Moreover, this same D-BIND traffic specification can be used for streams utilizing both a deterministic and a statistical guaranteed service. We presented a connection admission control algorithm that supports heterogeneous sources and heterogeneous QoS requirements by utilizing a simple static priority scheduler that can support both deterministic and statistical performance guarantees at multiple delay bounds. By using a single traffic descriptor for all guaranteed connections, the H-BIND statistical test is able to take into account the existence of higher priority traffic that utilizes a deterministic service, as well the traffic of other streams that utilize a statistical service.

Through trace-driven simulation using a 28 minute trace of MPEG compressed video, we showed that the H-BIND admission control test is slightly conservative, but yields considerably high utilizations close to those obtained by the trace driven simulation. For example, for a 40 msec delay bound guaranteed with violation probability 10^{-6}, an 86% average utilization was obtained.

As a part of future work, we plan to perform simulations with heterogeneous traffic mixes and to validate the performance of H-BIND under a wider range of scenarios.

6 Acknowledgements

The author is grateful to Domenico Ferrari and Hui Zhang for their helpful discussions, comments, and criticisms, and to the University of Wuerzburg for providing the trace of MPEG video.

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):114–121, January 1991.

[3] L. Dron, G. Ramamurty, and B. Sengupta. Delay analysis of continous bit rate traffic over an ATM network. *IEEE Journal on Selected Areas in Communications*, 9(3):402–407, April 1991.

[4] A. Eckberg. The single server queue with periodic arrival process and deterministic service times. *IEEE Transactions on Communications*, 27(3):556–562, March 1979.

[5] A. Elwalid, D. Mitra, and R. Wentworth. A new approach for allocating buffers and bandwidth to heterogeneous, regulated traffic in an ATM node. *IEEE Journal on Selected Areas in Communications*, 13(6):1115–1127, August 1995.

[6] A. I. Elwalid and D. Mitra. Effective bandwidth of general markovian traffic sources and admission control of high speed networks. *IEEE/ACM Transactions on Networking*, 1(3):329–43, June 1993.

[7] D. Ferrari. Design and applications of a delay jitter control scheme for packet-switching internetworks. *Computer Communications*, 15(6):367–373, July 1992.

[8] 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.

[9] V. Frost and B. Melamed. Traffic modeling for telecommunications networks. *IEEE Communication Magazine*, 32(3):70–81, March 1994.

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

[11] R. Guerin, H. Ahmadi, and M. Naghshineh. Equivalent capacity and its application to bandwidth allocation in high-speed networks. *IEEE Journal on Selected Areas in Communications*, 9(7):968–981, September 1991.

[12] B. Hajek. A queue with periodic arrivals and constant service rate. In *Chapter 10 of Probability, Statistics and Optimization —a Tribute to Peter Whittle*. F.P. Kelly ed.; John Wiley and Sons, 1994.

[13] G. Kesidis, J. Walrand, and C. Chang. Effective bandwidths for multiclass Markov fluids and other ATM sources. *IEEE/ACM Transactions on Networking*, 1(4):424–428, August 1993.

[14] E. Knightly, D. Wrege, J. Liebeherr, and H. Zhang. Fundamental limits and tradeoffs for providing deterministic guarantees to VBR video traffic. In *Proceedings of ACM SIGMETRICS'95*, Ottowa, Ontario, May 1995.

[15] E. Knightly and H. Zhang. Traffic characterization and switch utilization using deterministic bounding interval dependent traffic models. In *Proceedings of IEEE INFOCOM'95*, pages 1137–1145, Boston, MA, April 1995.

[16] J. Kurose. On computing per-session performance bounds in high-speed multi-hop computer networks. In *Proceedings of ACM SIGMETRICS'92*, pages 128–139, Newport, Rhode Island, June 1992.

[17] 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.

[18] M. Loève. *Probability Theory I*. Springer-Verlag, 4th edition, 1977.

[19] D. Mitra and J. Morrison. Multiple time scale regulation and worst case processes for ATM network control. In *Proceedings of the 34th IEEE Conference on Decision and Control*, pages 353–358, October 1995.

[20] J. Roberts and J. Virtamo. The superposition of periodic cell arrival streams in an ATM multiplexer. *IEEE Transactions on Communications*, 39(2):298–303, February 1991.

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

[22] 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.

[23] H. Zhang and E. Knightly. RED-VBR: A renegotiation-based approach to support delay-sensitive VBR video. *ACM/Springer-Verlag Multimedia Systems Journal*. To appear.

[24] 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.