

D-BIND: An Accurate Traffic Model for Providing QoS Guarantees to VBR Traffic

Edward W. Knightly and Hui Zhang

Abstract— Variable bit rate traffic that requires a bounded-delay network service is one of the most important types of traffic in future integrated services networks. In this paper, we introduce a new deterministic traffic model called Deterministic Bounding INterval-length Dependent (D-BIND) to capture the important multiplexing properties of bursty streams. With the D-BIND model, clients specify their traffic to the network via multiple rate-interval pairs, (R_k, I_k) , where a rate R_k is a bounding or worst-case rate over every interval of length I_k . 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 analyze the new model in the context of a deterministic service, and we quantify its performance benefits using a set of experiments with traces of MPEG-compressed video. We show that D-BIND's more accurate characterization of traffic streams leads to substantial improvements in network utilization as compared to previous traffic models.

1 Introduction

Future integrated services networks will have to support applications with both diverse traffic characteristics and diverse performance requirements. Of the many traffic classes in integrated services networks, delay-sensitive Variable Bit Rate (VBR) traffic poses a unique challenge. Since the required service is delay-sensitive, the network must reserve network resources for each connection. The resources are reserved based on both a source's traffic characteristics that are specified via a parameterized traffic model and the source's Quality of Service (QoS) requirements in terms of end-to-end delay bounds, delay-jitter bounds, and the maximum loss rate. However, it is unclear how large the reserved resources must be for VBR traffic sources due to their burstiness. In the literature, most traffic models are based on stochastic processes such as Markov-modulated [1] or Self-similar models [10]. In general, most stochastic models for characterizing bursty VBR traffic sources such as compressed video are either not powerful enough to capture the important burstiness and time correlations of realistic sources, or they are too complex for practical implementation for Connection Admission Control (CAC) [23]. In addition, while it is important for the network to verify and enforce, usually via a policing mechanism, that a source sends data according to its traffic specification, it is very difficult to verify whether a traffic stream satisfies

a specified stochastic characterization. Furthermore, with stochastic traffic models, it is often impossible to provide clients with *end-to-end* performance guarantees in networks with general topologies due to analytical difficulties in extending single-node results to networking environments.

A number of *deterministic* models have also been proposed to characterize traffic in integrated services networks: for example, the $(Xmin, Xave, I, Smax)$ model [7], the leaky bucket or (σ, ρ) model [4], [26], and the (peak-rate, burst length, average rate) model [8]. However, none of these models can accurately capture the burstiness of realistic sources such as compressed video. As we will show, a less accurate traffic model will result in an unnecessary over-allocation of resources, and hence, in lower average utilization of the network.

To overcome such limitations, we propose a new parameterized traffic model called D-BIND, or Deterministic Bounding INterval-length Dependent. The D-BIND model characterizes sources via multiple rate-interval pairs, (R_k, I_k) , where a rate R_k is a bounding or worst-case rate over every interval of length I_k . There are two key components to the D-BIND model. First, it *bounds* sources rather than attempting to fit their exact arrival distributions. These bounds are crucial for allowing the network to police sources, i.e., to verify that sources transmit traffic within their specified limits. Moreover, upper bounds on traffic allow the network to provide *deterministic* performance guarantees, which is not possible with typical stochastic models. Second, D-BIND parameterizes sources with bounding rates over *multiple interval lengths* to capture the most important multiplexing properties of sources. We will motivate this latter point via fundamental properties of queues and show that the new model captures the key properties of sources that *cause* the worst-case scenario to occur.

As an example of the D-BIND characterization, consider an MPEG-compressed video stream that alternates among transmission of I , P , and B frames. A source is sending at its peak rate when it is transmitting its largest-sized I frame. However, even in the worst case, the large I frame is immediately followed by a typically smaller B frame so that the micro-level burst does not persist for more than one frame time. The D-BIND model captures such burstiness characteristics with its family of rate-interval pairs. The model captures the intuitive property that over longer intervals, a source may be bounded by a rate lower than its peak rate and closer to its long-term average rate.

E. Knightly is with the Department of Electrical and Computer Engineering, Rice University, Houston, TX 77005.

H. Zhang is with the School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213.

By taking advantage of the more accurate traffic characterization offered by the D-BIND model, the CAC algorithm can allocate resources more efficiently and thus achieve a higher network utilization for a given QoS. We quantify the utilization improvements achieved with D-BIND as compared to previous traffic models by performing a set of experiments with several long traces of MPEG-compressed video. We focus on a *deterministic* service or a service without losses or delay-bound violations and show that significantly higher utilization can be achieved with the D-BIND model than with previous traffic models. Since sources may be multiplexed beyond a peak-rate-allocation scheme even while providing deterministic loss and delay-bound 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.

The accuracy of the D-BIND model can also be employed for network services other than a deterministic service. For example, a statistical service provides probabilistic guarantees on packet loss and delay-bound violation, exploiting statistical properties of multiplexed streams to obtain higher utilization than is possible for a deterministic service. In [16], statistical properties of streams are extracted from their *deterministic* constraints in order to deliver a statistical service. By basing a statistical service on D-BIND's accurate deterministic model, a substantial statistical multiplexing gain can be achieved, even while retaining the enforceability of the traffic specification. Although the focus here is on a deterministic service, D-BIND's relationship to other services, including a renegotiated service as in [30], is discussed in Section 4.4.

Lastly, we note that *shaping* traffic streams before their transmission into the network or *re-shaping* them at each network hop can impact the amount of resources that need to be reserved for each stream. For example, [11, 12] and [18] show how traffic shaping at appropriate points in the network can be used to improve end-to-end delay bounds or average network utilization in certain cases. In this paper, our focus is to characterize a given traffic stream, shaped or not, as accurately as possible in order to enable efficient resource allocation by a *network element* (switch or router) for each of the services the network offers. The joint effect of proper reshaping and accurate traffic characterization is beyond the scope of this paper.

The remainder of this paper is organized as follows. In Section 2, we describe the underlying requirements of deterministic traffic models and show how such models can be mapped to constraint functions. We then show, in Section 3, how bounds on individual streams can be applied to CAC to provide a deterministic service. In Section 4, we define the D-BIND model with motivations from both application traces as well as fundamental properties of queues. The performance of the new model is then compared to that of previous models using parameters derived from actual MPEG traces in Section 5. Finally, in Section 6, we discuss practical implementation issues for the model including policing and parameter specification.

2 Deterministic Traffic Models

A deterministic traffic model is one that parametrically describes the worst-case behavior of a traffic stream. Such a traffic characterization has the advantage that it can be policed by the network. 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.

In the $(Xmin, Xave, I, Smax)$ model of [7] (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. The *(peak rate, burst length, average rate)* model proposed by the ATM Forum [8] is similar to the $Xmin$ model, with $Smax$ fixed to 53 bytes. 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 can also be viewed in terms of its policing mechanism. In this policing mechanism, a source must have a "token" or "credit" to transmit a packet to the network. With (σ, ρ) , a source obtains credits at rate ρ and can collect up to σ credits. Hence, a source can send a burst of size σ bits or packets into the network, but, over the long term, there is an upper-average rate constraint ρ , similar to the $Smax/Xave$ rate constraint of the $Xmin$ model.

As required, all of the above traffic models provide a deterministic upper bound on each source's arrivals and allow a worst-case analysis that upper bounds delay and throughput. Specifically, a deterministic traffic model defines a deterministic traffic constraint function. A monotonic increasing function $b_j(\cdot)$ is called a *deterministic traffic constraint function* of connection j if during *any* interval of length t , the number of bits arriving on j during the interval is no greater than $b_j(t)$. Formally, let $A_j[t_1, t_2]$ be the total number of bits arrived on connection j in the interval $[t_1, t_2]$; $b_j(\cdot)$ is a traffic constraint function of connection j if $A_j[s, s+t] \leq b_j(t), \forall s, t > 0$. Notice that $b_j(\cdot)$ is a time-invariant deterministic bound since it constrains the traffic stream over every interval of length t .

For a given traffic stream, there are an infinite number of possible traffic constraint functions that can bound the source, out of which a deterministic traffic model defines a parameterized family. Since all deterministic models have an associated constraint function that is defined via the model's parameters, we can compare the accuracy or tightness of different models by comparing their constraint functions.

An important observation about the traffic constraint function is that for a given arrival process $A[0, t]$, the tightest time invariant deterministic bound on arrivals in any interval of length t is by definition

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

$\mathcal{E}(t)$ is called the empirical envelope in [27], and the min-

imum envelope process in [3]. In other words, $\mathcal{E}(t)$ is the tightest or most accurate deterministic time-invariant characterization of an arrival sequence $A[0, t]$. Thus, in order for a traffic model's constraint function $b(t)$ to be a time-invariant upper bound on the arrivals $A[s, s + t]$, it must upper bound $\mathcal{E}(t)$, viz., $b(t) \geq \mathcal{E}(t)$ for all t . Since the CAC algorithm reserves resources for the connection based on $b(t)$, the more tightly $b(t)$ bounds the actual traffic, the more efficiently resources can be allocated. A desirable property of a traffic model is therefore that it parameterizes a constraint function that can closely bound $\mathcal{E}(t)$ for a wide variety of sources. The linkage of the traffic model and the CAC algorithm, as well as the importance of the time-invariance property, will be further explored in the next section.

The *Xmin* model's constraint function is given by:

$$b(t) = (\min(\lceil \frac{t \bmod I}{Xmin} \rceil, \lceil \frac{I}{Xave} \rceil) + \lceil \frac{t}{I} \rceil \lceil \frac{I}{Xave} \rceil) Smax \quad (2)$$

and the (σ, ρ) model's is given by $b(t) = \sigma + \rho t$.

Conceptually, both the *Xmin* and (σ, ρ) models allow a limited-size burst and have an additional longer-term rate constraint. Experiments with constraint functions can be found in Section 5.2.

Before defining the D-BIND model, we further motivate it by describing the analysis techniques used to derive CAC algorithms for deterministic guarantees.

3 Deterministic CAC

A deterministic service ensures that no packets are dropped or delayed beyond their guaranteed delay bound. For the network to deliver such a service, it must reserve resources according to a worst-case scenario, so that in essence, resources cannot be "over booked". Thus, connection admission control for a deterministic service requires a worst-case bound on individual sources, and must also be able to determine how a collection of sources will interact when multiplexed inside the network, again considering the worst-case scenario.

For a deterministic service, the advantage of policing is extended to the network clients: a client can easily check that it has not had any packets dropped or delayed beyond its guaranteed bound. Alternatively, a client will have difficulty verifying a statistical guarantee that is defined over an infinite time horizon.

Deterministic admission control conditions rely on the worst-case delay analysis techniques such as in [3, 4, 19, 28]. Conceptually, an upper bound on delay can be derived beginning with an expression for the queue length at time τ expressed in terms of the actual arrivals A_j of each traffic stream and the link speed l . For example, for FCFS, such an expression is given by:

$$q(\tau) = \max_{s \leq \tau} \left\{ \sum_{j=1}^N A_j[s, \tau] - l(\tau - s) \right\}, \quad (3)$$

which is a direct consequence of the Lindley recursion [20].

By using the fact that $A[s, s + t] \leq b(t)$ together with other manipulations, it follows that for FCFS, delay is upper bounded by:

$$d = \frac{1}{l} \max_{t \geq 0} \left\{ \sum_{j=1}^N b_j(t) - lt + \bar{s} \right\}. \quad (4)$$

Delay bounds for other policies such as Hierarchical-Round-Robin, Static Priority, Earliest-Due-Date-First, and Packetized Generalized Processor Sharing, can also be expressed as functions of traffic constraint functions [2, 4, 19, 24, 29]. For example, the following theorem is given and proven in [29] to derive delay bounds in a Static Priority scheduler.

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 \bar{s} , the maximum delay of any packet at priority level k is bounded above by d_k , where*

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

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

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

Equations (4) and (5) can be used as CAC tests for FCFS and SP schedulers in that they can test whether or not a set of connections can be multiplexed so that each packet of each connection can be delivered within the delay bound that is guaranteed to the network client. For CAC, the theorems may be used to test if a *new* connection can be admitted so that all connections, including the new one, obtain their respective delay guarantees. Hence, Equations (4) and (6) can also be viewed in terms of the maximum number of admissible connections for a given QoS. For example, for FCFS, Equation (4) can be rewritten to express the maximum number of admissible connections as a function of the delay bound:

$$N(d) = \max\{n \mid \frac{1}{l} \max_{t \geq 0} \left\{ \sum_{j=1}^n b_j(t) - lt \right\} \leq d\} \quad (7)$$

Two things should be noticed about the above admission control condition. First, in both Equations (4) and (5), delay bound calculations are reduced to a problem of maximizing a linear combination of traffic constraint functions. If the traffic constraint function is piece-wise linear, as the case for *Xmin*, (σ, ρ) , and the D-BIND model defined in Section 4, this computation can be very efficient and fast. In addition, the algorithms work equally well for both homogeneous and heterogeneous sources. In contrast, accommodating heterogeneous sources with stochastic traffic characterizations often severely complicates the admission control analysis.

Second, the equations indicate that even better bounds are possible with a new traffic model. 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 (4) will be lower. Thus, a 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.

4 D-BIND Model Definition

In this section, we define the D-BIND traffic model. We present it with motivations based on two scenarios. First, we investigate a 10-minute trace of MPEG compressed video and show how the D-BIND model captures the burstiness properties of such a stream. Second, we revisit the CAC problem and show how the D-BIND model captures the properties of sources that are important to admission control and hence to achieving the highest possible network utilization.

4.1 D-BIND and Burstiness Properties of Streams

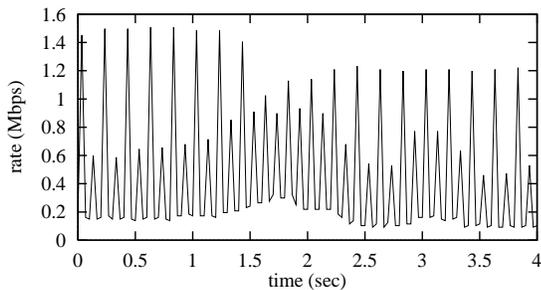


Figure 1: Segment of Lecture Trace

Figure 1 shows a 10-second segment of a lecture sequence. The lecture is digitized to 160x120 pels and compressed at 30 frames per second using the MPEG compression algorithm [9]. The Group of Pictures pattern is IBBPBB such that M is 3 and N is 6, and the quantizer scales are 8 (I), 10 (P), and 25 (B). The sequence was digitized in hardware and compressed in software with the Berkeley MPEG-1 encoder [25]. The figure shows time on the horizontal axis and rate on the vertical axis. The bit rate is calculated as the frame size multiplied by the frame rate of 30 frames per second.

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, which use only *intraframe* compression, and P and B frames, which are transmitted between I frames and use *interframe* compression. While P frames or Predicted frames are coded based on only past frames, B frames or 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. In Figure 1, which frames are which is apparent since the I frames tend to be the largest, B the smallest, and P in between.

This traffic source is *bursty* in that its transmission rate varies substantially over time. The important observation about this stream for the purpose of designing a traffic model is that the source transmits at different rates for different interval lengths. For example, while transmission of large I frames can cause bursts of high rate, these bursts will only last for the duration of one frame time, even in the worst case. That is, larger I frames are always followed by smaller B frames. Hence, the key questions about the source that we are trying to answer with the model are “what are the rates of the bursts and how long do they last?” Moreover, since we wish to provide deterministic service, these characteristics must be captured with a deterministic traffic model. Thus, the question we turn to now, is how to capture such burstiness properties of sources with a worst-case model.

The traffic model that we propose is called Deterministic Bounding INterval-length Dependent, or D-BIND. With the D-BIND model, sources characterize their traffic with P rate-interval pairs, $\{(R_k, I_k) | k = 1, \dots, P\}$, which are specified to the network at connection-setup time. The network then performs CAC tests based on both the D-BIND traffic parameters and the requested QoS parameters. The “Deterministic Bounding” part of D-BIND provides the worst-case bound on sources that is required to provide a deterministic service. The “INterval-length Dependent” part of D-BIND captures sources’ different burstiness properties over different interval lengths, the property noted in the trace of Figure 1.

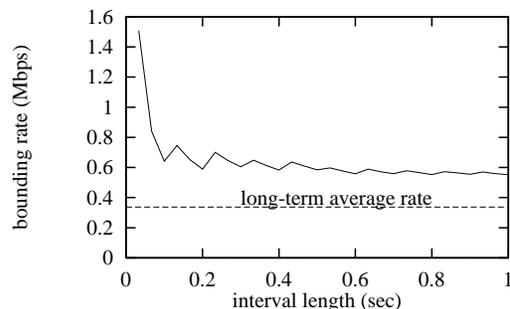


Figure 2: D-BIND Rate-Interval Pairs for Lecture

Figure 2 shows a plot of the D-BIND rate-interval pairs for the 10-minute trace of the lecture. The horizontal axis shows interval length I_k in seconds, and the vertical axis shows the bounding rate R_k in Mbps, over intervals of length I_k . For example, for an interval of length $I_1 = 1/30$ seconds, the bounding rate is determined by the largest sized frame of the entire sequence, which is the largest sized I frame. For this lecture trace, the largest sized I frame is 50.3 kbits, for a peak rate of $R_1 = 1.51$ Mbps.

As is evident from Figure 1, even in the worst case, the large intra-coded I frames are followed by smaller inter-

coded B frames. Hence, the burst at rate 1.51 Mbps lasts for only 1/30 seconds, and the worst case rate over 2/30 seconds or two frame times is considerably lower. In this case, the bounding rate over any interval of length 2/30 seconds is 840 kbps. Equivalently, the largest size of any two consecutive frames is 840 kbps times 2/30 seconds or 56 kbits. Hence, the D-BIND model captures the source’s burstiness over multiple interval lengths. For small interval lengths, R_k approaches the source’s peak rate, 1.51 Mbps. For longer interval lengths, R_k approaches the long term average rate of 337 kbps, which is total number of bits in the MPEG sequence divided by the length of the sequence.

From the initial peak rate R_1 , the bounding rate tends to decrease over longer interval lengths. It does not decrease monotonically though because of the quasi-periodic nature of the MPEG stream in which sources alternates between large and small frames over time. Regardless, the general trend of the curves is that the bounding rate decreases with increasing interval length, decreasing from the peak rate to the long-term-average rate. 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.

From all possible D-BIND pairs as shown in Figure 2, P rate-interval pairs are specified to the network as the source’s traffic specification for admission control. Typically, we view P as being small, on the order of four. Discussion of the impact of P and of other implementation issues is found in Section 6.

Lastly, we note that we do not require the special framing structure found in MPEG video in order to use the D-BIND model. Indeed, as we describe in the next section, the D-BIND model uses multiple rate-interval pairs to accurately describe a general source’s worst-case behavior for more efficient resource allocation. D-BIND characterizations and the resulting performance for applications other than MPEG-compressed video can be found in [17].

4.2 D-BIND and CAC

In Section 2, we showed that the tightest possible constraint function for an arrival sequence A is defined as the empirical envelope $\mathcal{E}(t) = \sup_s A[s, s + t]$. Hence, every traffic model parameterizes a constraint function that upper bounds $\mathcal{E}(t)$, or $b(t) \geq \mathcal{E}(t) \forall t$ and for all deterministic traffic models. To relate this property to CAC, we note that for the FCFS and SP delay bounds derived above, the minimum delay bound that can be provided to a set of connections which are bounded by their respective constraint functions $b_j(t)$ is achieved when $b_j(t) = \mathcal{E}_j(t)$ for all t and for all connections j . Equivalently, characterizing all traffic streams according to their empirical envelopes will result in the maximum number of admissible connections for any deterministic traffic model. This means that the empirical envelope is *optimal* in the sense that of all deterministic time-invariant bounds on traffic streams, the empirical envelope is the most accurate and results in the most admissible connections for a given delay bound (see

[27] for further discussion of the empirical envelope).

Unfortunately, the empirical envelope lacks practical properties for use in CAC. In particular, sources cannot efficiently specify the function $\mathcal{E}(t)$ to the network, and the network cannot efficiently police or enforce a general envelope $\mathcal{E}(t)$ without any constraints on its shape. Thus, the goal of the D-BIND model is to provide a parameterized traffic model that is suitable for CAC, that can be policed by the network, and that characterizes the traffic as tightly or accurately as possible, in order to achieve the highest possible utilization for a deterministic service. We can therefore define the D-BIND model as it relates to the empirical envelope.

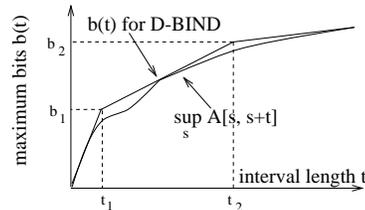


Figure 3: Traffic Constraint Function for D-BIND Model

Figure 3 shows an example of an empirical envelope. The horizontal axis represents interval length, and the vertical axis represents the maximum number of bits the source transmits over the interval length. The lower curve represents the tightest bound on the number of arrivals in any interval of length t , $\mathcal{E}(t) = \sup_s A[s, s + t]$ (Equation (1)). As described above, all traffic models define constraint functions $b(t)$ that upper bound $\mathcal{E}(t)$. The constraint function defined by the D-BIND model provides a piece-wise linear upper approximation to this tightest bound $\mathcal{E}(t)$. 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 rate-interval pairs $(b_k/t_k, t_k)$, and with linear interpolation between the points on the D-BIND curve.

Thus, given P rate-interval pairs, i.e., (R_k, I_k) , $k = 1, 2, \dots, P$, we define the D-BIND constraint function as

$$b(t) = \frac{R_k I_k - R_{k-1} I_{k-1}}{I_k - I_{k-1}} (t - I_k) + R_k I_k, \quad I_{k-1} \leq t \leq I_k \quad (8)$$

with $b(0) = 0$ and $b(\cdot)$ repeating for $t > I_P$, such that $b(t) = b(t - \lfloor t/I_P \rfloor I_P)$ for $t > I_P$.

The purpose of defining the D-BIND model’s constraint function as a piece-wise linear upper-bound on $\mathcal{E}(t)$ is multi-fold.

- **Parameterization:** In order to integrate a traffic model with a signaling protocol for connection-establishment, it must be specifiable with a small number of parameters. The D-BIND model’s rate-interval pairs provide such an interface for network clients.
- **Policing:** The network must be able to protect the network and other network clients from malicious or misbehaving users that violate their promised traffic

specifications. The D-BIND characterization can be policed or enforced by the network using mechanisms described in Section 6.3.

- Utilization - Equation (3) shows that queues build up from more arrivals over intervals than can be serviced. These arrivals over intervals are *rates*, and the D-BIND model captures these rates with multiple rate-interval pairs.

Equations (4), (3), and (5) indicate that queue lengths are determined by arrivals over intervals, and deterministic connection admission control tests rely on upper *bounds* on arrivals over intervals, or constraint functions $b(t)$. Equation (7) shows that the key to achieving high utilization (or admitting as many connections as possible) is to use a constraint function $b(t)$ that more accurately describes arrivals over intervals. The D-BIND model directly characterizes these arrivals over intervals via rates in the most accurate manner possible subject to the practical constraints above. Thus, the D-BIND model is *bounding* for CAC and policing, and *interval-dependent* to characterize the sources' different worst-case rates over different interval lengths.

4.3 D-BIND's Relationship to Other Traffic Models

We 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 $\{(\sigma_k, \rho_k), k = 1, 2, \dots, P\}$, or $(\vec{\sigma}, \vec{\rho})$, such as analyzed in [27], is a special case of the D-BIND model in which the constraint function is piece-wise linear *concave*. That is, the $(\vec{\sigma}, \vec{\rho})$ model has a constraint curve $b(t) = \min_k \{\sigma_k + \rho_k t\}$ which is necessarily concave. Implications of concavity are discussed in Section 6. As well, the *Xmin* model can be expressed in terms of the D-BIND model by using a different interpolation function.

4.4 D-BIND's Relationship to Other Network Services

The main issues that we are addressing in this work are the parameterized traffic model and the resource allocation scheme or CAC algorithm. The resource allocation scheme in turn determines the type of service that is offered. Here, we introduced the *D-BIND* traffic model, used a worst-case resource allocation scheme, and provided a scheme for delivering a *deterministic* service. As summarized in Table 1, other combinations of traffic models, resource allocation schemes, and services are possible.

In the first two rows, we have examples of schemes for providing a *deterministic service*, in which connections are guaranteed a service that avoids loss or delay-bound violations. Both [7] and this work utilize worst-case resource allocation schemes in order to provide this service, but they use different traffic models (and also use different service disciplines, traffic streams, etc.).

The second two rows of Table 1 list two recent examples of schemes for providing *statistical service*, which can

Traffic Model	Resource Allocation	Service
D-BIND	Worst Case	Deterministic
$(Xmin, Xave)$	Worst Case [7]	No Loss
D-BIND	H-BIND [16]	Statistical
$(\vec{\sigma}, \vec{\rho})$	Large Deviations [22]	$Prob\{Loss\}$
D-BIND	RED-VBR [30]	Renegotiated
Peak-Rate	R-CBR [14]	$Prob\{Block\}$

Table 1: Example Combinations of Traffic Model, Resource Allocation, and Service

achieve higher resource utilizations by exploiting statistical properties of the sources. Most previous algorithms of providing statistical services require applications to specify the statistical properties of the sources explicitly via a stochastic traffic model such as a Markov Modulated Fluid Source as in [1]. While stochastic traffic models provide the statistical properties of a source needed for the CAC algorithm, they introduce great difficulties for the policing algorithm as it is often impossible to verify whether a traffic source satisfies a specified stochastic characterization. In [16, 22], techniques are developed to extract *statistical* properties of traffic streams from their *deterministic characterizations*. With these techniques, statistical services can be provided with only deterministic traffic models. By basing a statistical service on a deterministic model, a statistical multiplexing gain can be achieved, even while retaining the enforceability of the traffic specification. With the more accurate D-BIND model, the derived statistical properties are also more accurate, which means that an even higher utilization can be achieved for the statistical service.

The last two rows of Table 1 depict two proposed schemes for providing *renegotiated services*, RED-VBR (Renegotiated Deterministic VBR) [30] and R-CBR (Renegotiated CBR) [14]. In both of these works, services are proposed in which network clients can renegotiate their traffic parameters and performance requirements with the network in order to adapt to long-time-scale dynamics of the clients' traffic characteristics.

For example, consider a movie that is digitized, compressed, and transmitted over a network. Roughly speaking, if this movie alternates between high-action, high-bandwidth scenes and low-action, lower-bandwidth scenes, a renegotiated service allows a network client to release network resources back to the network during low-action scenes by signaling its new traffic specification at the beginning of such scenes. When the client re-enters a high-action scene, it signals its request to the network to renegotiate for its new higher bandwidth requirements.

Compared to a deterministic service, a renegotiated service achieves higher utilization of network resources. Specifically, when one network client releases some of its resources back to the network, other network clients can utilize these resources. Once again however, this utilization improvement is not for free in that a renegotiated service has a profound difference as compared to a determinis-

tic service: when a renegotiating client enters a higher-bandwidth scene, its request for more bandwidth may be denied or *blocked* by the network, in which case the client will have to make do with the previously reserved bandwidth. Hence, in order to be useful, such services must provide this blocking probability $Prob\{Block\}$ to network clients as a performance or QoS parameter.

For the purposes of the discussion here, the schemes of [30] and [14] differ in that RED-VBR builds the renegotiation service on top of a deterministic variable bit rate service with the D-BIND traffic model (as proposed in this paper), while R-CBR builds the renegotiation service on top of a constant bit rate service. By using a more accurate characterization of traffic streams, the RED-VBR scheme can potentially allocate resources more efficiently than R-CBR, i.e., require fewer renegotiations for the same level of resource utilization. On the other hand, a CBR service is easier to implement than a D-VBR service.

Different network services such as renegotiated service and statistical service can be viewed as complimentary to the D-BIND model and deterministic guarantees that we are describing in this paper. The reason for this is that the D-BIND model provides a foundation for delivering different network services: D-BIND is a traffic model that upper bounds the traffic, and the resource allocation scheme determines which network service is offered. Different resource allocation schemes can be used with D-BIND to provide different services [16, 30].

5 Empirical Investigations

In this section, we evaluate the effectiveness of the D-BIND model by analyzing the link or multiplexer utilization achieved with the new model and comparing the results to utilizations obtained with the $Xmin$ and the (σ, ρ) models. Because deterministic traffic models may be viewed in terms of the constraint function $b(t)$, they are comparable via the admission control condition in Section 3. 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, or $\sum_j R_{1,j} < l$.

Two ten-minute traces of MPEG compressed video are analyzed as traffic sources. The first video trace is the lecture sequence described in Section 4.1 and the second video trace consists of a sequence of advertisements for graphics products. The advertisement sequence has the same MPEG parameters as the lecture sequence, but quite different content: the advertisement video is 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.

For the experiments, we assume that each video frame is

transmitted per frame-time as opposed to introducing additional delay by smoothing over several frames: effects of such smoothing can be found in [18]. Additionally, we assume that each frame is segmented into 48 byte ATM cells with the cells being transmitted at equally spaced intervals over the frame-time.

5.1 D-BIND Traffic Characterization

Figure 4 shows the possible D-BIND rate-interval pairs for the advertisement and lecture sequences. The vertical axis depicts R_k/R_1 or the bounding rate R_k normalized to the peak rate R_1 for the respective traces. The rate R_k is the worst-case rate over intervals of length I_k , where I_k is on the horizontal axis. The long-term average rates of the streams are shown with the arrows: for the lecture sequence, the average rate is .224 times the peak rate R_1 , and for the advertisements, it is .185 times the peak rate.

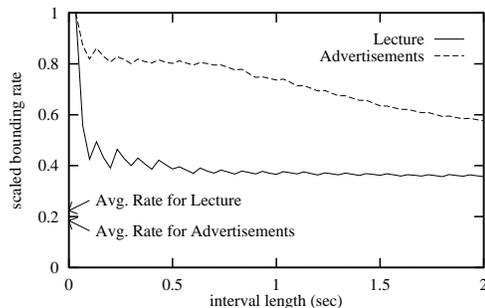


Figure 4: D-BIND Rate-Interval Pairs

Of most importance for the admission control experiments below, is how quickly R_k approaches the long-term-average rate (shown with the arrows) as I_k increases. For example, the rate-interval curve for the lecture video shows a rapid decrease of the bounding rate, whereas the curve for the advertisement video decreases much more slowly. Intuitively, a slowly decreasing rate-interval curve indicates that bursts of high rate persist over relatively long intervals, which in turn implies that it will be extremely difficult to achieve a high multiplexing gain and simultaneously provide good QoS. This is because high-rate bursts of long duration cannot be effectively absorbed by network buffers. Without network buffers, peak-rate reservation is required to provide a deterministic QoS guarantee.

5.2 Constraint Functions

Figure 5 shows the traffic constraint function of the lecture sequence for the D-BIND model with four rate-interval pairs and compares it to the constraint functions for the $Xmin$ and (σ, ρ) models. As in Figure 3, 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 or accurately the model characterizes the source, and the fewer the resources that the network must reserve for the connection.

As shown, the D-BIND model most tightly characterizes this traffic stream.

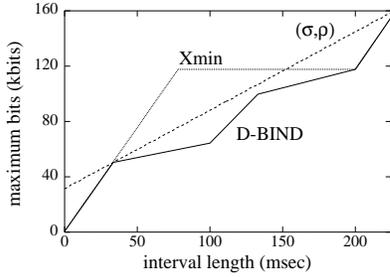


Figure 5: Traffic Constraint Function

Moreover, the temporal properties of the MPEG stream are evident only 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 smaller B frame. At 100 msec, after sending two B frames, the D-BIND constraint function breaks up again indicating the transmission of a 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. The lack of concavity is an inherent property of the MPEG stream and is due to its quasi-periodic nature, where the coder alternates between I, B, and P frames. Section 6 addresses the implications of this observation regarding policing.

5.3 Link Utilizations and DMG

A key issue for evaluation of the D-BIND model is the achievable network utilization compared to other deterministic traffic models. As described in Section 1, the deterministic approach uses upper bounds on sources to provide service without packet losses or delay-bound violations.

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. Heterogeneous cases are explored in Sections 5.5 and 6.3. In this case, the maximum number of connections 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 Equation (7), with $b(t)$ as given by Equation (8) and shown in Figure 5 for the lecture sequence. Within this context, for a deterministic delay bound d_k , the average utilization for deterministic traffic is given by $N(d_k) \cdot R_\infty / l$ where R_∞ is the stream’s 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 admissible connections with deterministic guarantees divided by the link speed,

$$DMG \triangleq \sum_{j=1}^N \frac{R_{1,j}}{l} \quad (9)$$

which is $N(d_k) \cdot R_1 / l$ in the homogeneous case. Thus, a peak-rate-allocation scheme has a DMG of at most 1.

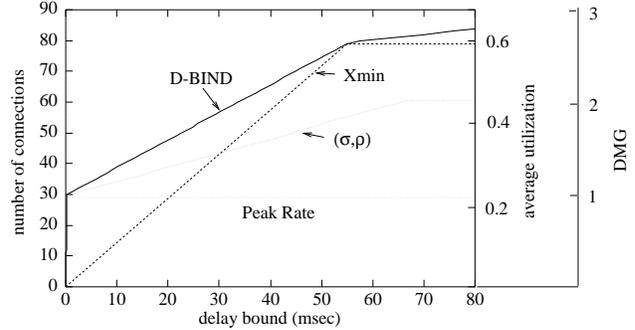


Figure 6: Utilization and DMG for Lecture

For the two traces, Figures 6 and 7 show the number of admissible connections, the corresponding 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. In other words, for the delay bound d_k reported on the horizontal axis, the vertical axis shows the maximum number of admissible video connections so that all connections obtain a worst-case delay bound d_k . Figure 6 shows the data for the lecture sequence, while Figure 7 is for the advertisements. As indicated by Equation (7), the average utilization increases with increasing delay bound. By “utilization”, we are referring to that achieved only by deterministically guaranteed real-time traffic. The remaining network resources may be used by statistical or best-effort traffic.

There are several noteworthy points about Figures 6 and 7. First, it is immediately apparent that the D-BIND model performs better than the *Xmin* and (σ, ρ) models, i.e., that it achieves higher utilization for the same QoS guarantee via a more accurate traffic characterization. For example, for the lecture sequence of Figure 6 and a guaranteed delay bound of 40 msec, the D-BIND model is able to utilize the network up to 50%, whereas the *Xmin* model results in a utilization of 36% and the (σ, ρ) model results in a 43% utilization. Thus, in this case, the D-BIND model results in 39% and 16% improvement in network utilization over the respective *Xmin* and (σ, ρ) models.

An alternative metric for the improvement may be obtained by comparisons with a peak-rate allocation scheme. Figure 6 shows that the D-BIND model performs significantly better than the 23% utilization achieved with peak-rate reservation. With the accurate D-BIND model together with the analysis techniques of [3, 4, 28], even for small delay bounds, DMG’s significantly greater than 1 are achievable. For example, for a delay bound of 9 msec, 38

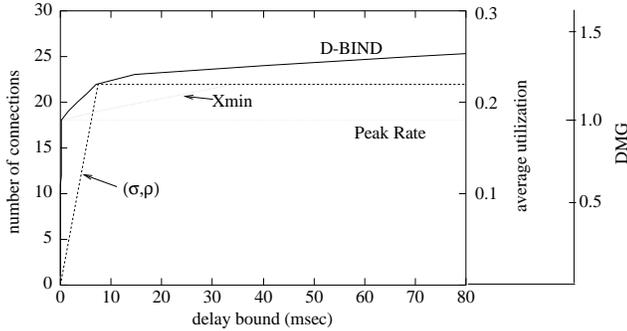


Figure 7: Utilization and DMG for Advertisements

connections may be multiplexed for a DMG of 1.3. For a 40 msec delay bound, the DMG is 2.2, and for a 48 msec delay bound it is 2.7.

Finally, note that the improvements shown in Figure 7 are not as pronounced. This is due to the different shape of the (R_k, I_k) rate-interval curve for the advertisements as shown in Figure 4. Although the curve does conform to the interval-dependent property that sources may be bounded by lower rates over longer interval lengths, this property is obeyed in a lethargic manner. That is, compared to the lecture's rate-interval curve, the advertisement sequence's rate-interval curve decreases more slowly to its long-term-average rate. Hence, when multiplexing the advertisement sequence, for a delay bound of 69 msec, the improvement is from an average network utilization of 18% for a peak-rate-allocation scheme to 25% for the D-BIND model. The reason for the limited utilization improvement is that bursts of high rate and high duration cannot be effectively absorbed by network buffers. The fact that the compressed advertisement sequence has bursts with rates near the peak rate that last for longer intervals is captured by the D-BIND rate-interval pairs.

5.4 Heterogeneous Node Loads

With the $Xmin$ and (σ, ρ) models, a source has a range of bounding parameters that it can choose from. That is, a source can choose $Smax/Xave$ or ρ anywhere between the peak rate and the average rate. The choice of this rate will in turn affect the respective values of I and σ . The effects of this choice on the utilization curves are shown in Figure 8.

The figure demonstrates that for a given $Smax/Xave$ or ρ there may be a small range of delays such that the $Xmin$ or (σ, ρ) model performs nearly as well as the D-BIND model. However, the D-BIND model still has a significant advantage with respect to practical issues of establishment of real-time connections in a network: for instance, 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 $Smax/Xave$ or ρ , some of the nodes will be forced to allocate resources inefficiently since choosing one value tends to be efficient for some delay bounds and inefficient for others. Since D-BIND captures the streams burstiness properties over multiple interval lengths, it does not have this problem.

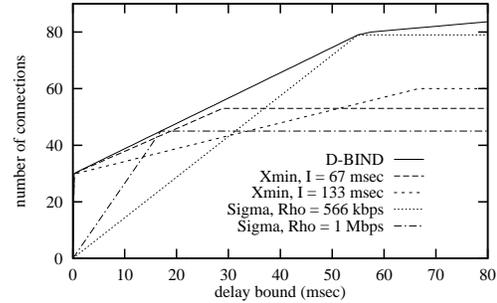


Figure 8: Effect of $Xave$ and ρ Parameter Selection

5.5 Heterogeneous Traffic and QoS

In providing *statistical* guarantees with use of traditional stochastic models such as those mentioned in the introduction, two difficulties in addition to those mentioned in Section 1 are often encountered. First, accommodating sources with heterogeneous traffic characterizations severely complicates the admission control analysis. Second, it is often difficult or impossible to provide a service in which heterogeneous users can receive heterogeneous services from the network. With the approach presented here, heterogeneity in the traffic specification and QoS are easily accommodated. It is not even required that all sources use the D-BIND model. The reason for this is that, as shown in Equation (5), admission control calculations are performed via the traffic constraint function $b(t)$. Thus, sources can choose any parameter values for whatever traffic model is provided to the user. This is then converted to a constraint function and used in the admission control test. Of course, there will be a utilization penalty for network clients that use less accurate traffic models.

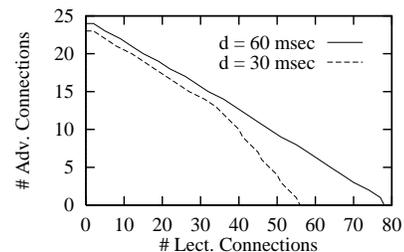


Figure 9: Admission Control for Combinations of Advertisements and Lecture

Figure 9 shows the admission control results of multiplexing different combinations of lecture and advertisement sequences by depicting a schedulable region in the manner

of [15]. For the lower curve representing a 30 msec delay bound, a point on the curve represents the maximum number of respective lecture and advertisement sequences that can be multiplexed so that all packets of all connections obtain a delay bound of 30 msec and are not dropped due to buffer overflows. Thus, any point below the curve is also schedulable for that same QoS. Similarly, the upper curve depicts the schedulable region for a 60 msec delay bound. Finally, we note that, unlike [15], since we provide a general service in which traffic sources are not restricted to classes, these CAC calculations would be made with Equation (5) rather than by table lookup.

While heterogeneity in sources is easily accommodated with D-BIND, the richness of the services that can be provided by the network is a function of the service disciplines at the switching nodes. For example, if FCFS is used, only a single local delay bound can be provided to all connections at each node. With Earliest Deadline First (EDF) scheduling [7], a continuum of delay bounds can be provided.

The choice of the service discipline also impacts the multiplexer’s utilization, its implementation complexity, and, of particular relevance here, its ability to exploit D-BIND’s more accurate traffic specification. For example, as shown in [31], the Stop-and-Go service discipline [13], must use a busy-period bound to guarantee delay because of Stop-and-Go’s framing structure. Consequentially, its delay bounds are looser than those presented in Section 3 and D-BIND’s richer traffic characterization cannot be exploited. The GPS service discipline [24] is able to exploit the D-BIND characterization, but because of its isolated treatment of streams, only to the point of a smoothing scheme, which is not always as efficient as a network’s deterministic service [18]. Alternatively, service disciplines such as EDF and Static Priority are able to use exact CAC conditions [19] and obtain all of the possible deterministic multiplexing gain. The impact of the service discipline on deterministic service is further discussed in [27] and [29].

6 Implementation Issues

In proposing a new source model, there are several issues regarding its practicality. For a deterministic model, these issues include: (1) concise parameterization - can the model be represented in a concise manner? (2) parameter specification - how difficult is it for a source to come up with its characterization? (3) policing - can the model be effectively and efficiently enforced?

6.1 Concise Parameterization

In the example of Section 5, an (R_k, I_k) 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 10 shows the result. While the homogeneous case does not explore all of the facets of us-

ing 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.

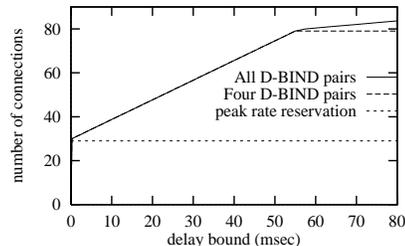


Figure 10: Utilization with Four Rate-Interval Pairs

For MPEG video sources, an alternative concise parameterization is to use knowledge of the frame pattern (in this case, IBBPBB) along with a parameterization of the largest sized I frame, B frame, and P frame. With this alternative “worst-case” characterization, a pessimistic approximation to the D-BIND constraint function can be obtained by constructing $b(t)$ 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.

There are certainly tradeoffs involved with the number of parameters used to describe $b(t)$. Our goal with D-BIND was not to simply add *more* parameters to a deterministic characterization, but rather to be sure that for a certain number of specified parameters, the most important information (for determining queue lengths and QoS) is being conveyed to the network’s resource allocation system. While additional parameters can improve resource utilization, it also increases the complexity of the policing mechanism and requires more fields in the signaling messages. Further experiments on the achievable utilization improvement for each additional parameter can be found in [27]. We note however, that a restriction to three parameters, as is the case for the current ATM and IETF standards, may overly restrict network utilizations, as seen in Section 5. Alternatively, with three or four rate-interval pairs, which could be expressed with five and seven parameters respectively if the last interval extends indefinitely, most of the achievable gain for deterministic service is realized. Moreover, in the scheme of [16], these same rate-interval pairs could be used for statistical performance guarantees, with longer interval lengths used to capture longer time-scale properties of streams that are important in obtaining a statistical multiplexing gain [16]. Alternatively, a (*peak rate, burst length, average rate*) model is not accurate enough to capture the streams’ important burstiness properties over the important time scales: characteristics that are essential for obtaining a statistical multiplexing gain (see also [22]).

6.2 Parameter Specification

Here we address the issue of how a source can determine its parameters at the connection setup time. Determining a stream’s traffic model parameters can be divided into two cases: off-line and on-line. The off-line case is for applications such as stored video where the stream’s arrival sequence is known in advance. The on-line case is for applications such as live video where little information about a stream’s characteristics is known in advance. In both of these cases, determining D-BIND parameters does not have more complexity than finding, for example, $(\vec{\sigma}, \vec{\rho})$ parameters [27].

6.2.1 Off-line Parameter Specification

For the off-line case, the following example illustrates the process of obtaining D-BIND parameters. Consider a video sequence where the size of the i^{th} frame is denoted by f_i . For P rate-interval pairs, we can set interval length I_k to be kT where T is the inter-frame time. For a trace consisting of F frames, the bounding rates can be easily calculated as

$$R_k = \frac{\max_{k \leq i \leq F} \sum_{j=0}^{k-1} f_{i-j}}{I_k}. \quad (10)$$

Each rate-interval pair can be determined with a single pass through the trace so that the computational complexity of determining P rate-interval pairs for a trace of length F is FP . Equivalently, all rate-interval pairs can be determined in a single pass through the trace using P counters.

This computational complexity of determining parameters for the $(\vec{\sigma}, \vec{\rho})$ multi-level leaky bucket model with P (σ_k, ρ_k) pairs is the same as that of the D-BIND model using an algorithm such as in [27].

6.2.2 On-line Parameter Specification

For the on-line case, the arrival sequence is not known in advance and hence, parameter values for *any* traffic model are more difficult to obtain. We therefore address the problem of how a network client that knows little or nothing about its traffic specification can best utilize a deterministic service.

Note that this problem is inherent to any resource allocation system. For example, a resource allocation system such as [6] uses the notion of a *contract* in which the network client specifies its traffic and promises not to send more. If the connection is admitted, then the network promises to deliver the client’s requested QoS. If one half of the contract cannot be specified, particularly the traffic specification, then the contract can only be honored when the client sends within what it does specify.

Our general approach for approximating a deterministic service in the on-line case is to use a renegotiated service with properly chosen parameters for the adaptation algorithm. Specifically, the renegotiation scheme of [30] has a tunable parameter that determines how aggressively a user wishes to renegotiate. A larger value of this parameter results in less-frequent renegotiations and decreases

the probability of renegotiation failure. While we do need to *adapt* the traffic parameters for a source with unknown traffic bounds, the goal of the use of renegotiation here is quite different from that of a renegotiated service per se. Indeed, a renegotiated service attempts to achieve a statistical multiplexing gain by exploiting long-time scale dynamics of streams. Consequentially, a renegotiated service provides a type of *statistical* service by targeting a non-zero probability of renegotiation failure. While there is a measure of uncertainty in providing a deterministic service to a source with unknown parameters, we attempt to deliver the closest possible service to a deterministic service by adapting the values of the traffic specification. Hence, we are trying to avoid renegotiation failures and are willing to forfeit the statistical multiplexing gain in order to obtain a better service, viz., a deterministic service.

6.2.3 Comparison to (σ, ρ) and $Xmin$

Lastly, we compare the difficulty of obtaining parameters for the (σ, ρ) and $(Xmin, Xave, I)$ models with that of the D-BIND model. It may be ostensible that specification of such two- or three-parameter models may be easier than specifying (say) four rate-interval pairs. However, compared to the D-BIND model, there are some additional difficulties in parameter selection for the (σ, ρ) and $(Xmin, Xave, I)$ models. Consider the (σ, ρ) model: the maximum burst size σ should be viewed as a function of ρ , with ρ being a rate between the source’s peak and average rate, and σ being a bound on the burst size. Clearly, a smaller ρ requires a larger σ for a given arrival sequence. As alluded to in Section 5.4, proper selection of ρ depends on the network load, and may be difficult or impossible to choose appropriately in the case of multiple hops with varying loads. Roughly, if the network is loaded such that bandwidth is plentiful and buffers are scarce, the source should choose a larger ρ and a smaller σ , and vice versa; further discussions of such bandwidth-buffer tradeoffs can be found in [21]. Unfortunately, such information about the network’s state is dynamic and may not even be available at connection setup time. Moreover, a poor choice of ρ could cause the connection to be unnecessarily rejected. Contrastly, the D-BIND model alleviates this problem with its more expressive traffic characterization.

6.3 Policing

Since the network must protect guaranteed-service 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’s edge is called policing and is shown in Figure 11. 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 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 is incurred in the policer.

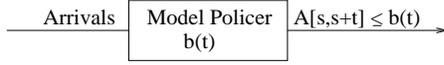


Figure 11: Policing of the Traffic Constraint Function $b(t)$

As described in Section 4.2, a piece-wise linear function may be used to represent the D-BIND model’s constraint function. Section 5 demonstrated that, because of the temporal properties of MPEG sources, the sources considered here had neither monotonically decreasing rate-interval curves nor concave constraint functions. As addressed by the propositions below, a concave constraint function has implications for policing.

Lemma 1 *If $b(t)$ is piece-wise linear concave, then R_k is strictly decreasing with increasing I_k .*

Proof: A function $b(t)$ is concave if for any $t_1 < t_2$ and $0 \leq \alpha \leq 1$, $\alpha b(t_1) + (1 - \alpha)b(t_2) \leq b(\alpha t_1 + (1 - \alpha)t_2)$. Denoting by $R(I)$ the bounding rate over the interval I , 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 $b(u_1) \geq u_1/u_2 \times b(u_2)$. \square

Lemma 2 *If a piece-wise linear constraint function $b(t)$ with P linear segments is concave, then the source may be fully policed, i.e., Equation (8) holds, with a cascade of P leaky buckets.*

The proof is given in Theorem 5.1 of [5]. Note that, as shown in Figure 5, 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.

However, considering the concave restriction of the D-BIND model results in a less accurate constraint function and potentially lower network utilization. An example of the potential utilization differences is shown in Figure 13 for the sources shown in Figure 12. Figure 12 shows the D-BIND constraint function for the lecture sequence as well as the concave hull of this function. We will refer to this as Source #1 and Source #1/Concave. Also shown is a second source, Source #2, with a peak rate of 1.5 Mbps and a bounding rate of 300 kbps over a 500 msec interval.

Figure 13 shows, in the manner of Figure 9, the maximum number of respective Source #1 (lecture) connections and Source #2 connections that can be multiplexed so that all connections have a deterministic delay bound of 60 msec. The figure shows that restricting the D-BIND

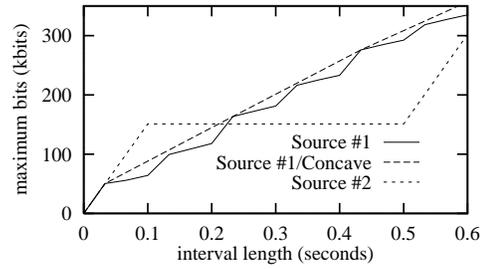


Figure 12: $b(t)$ for Concavity Experiments

model to have a concave constraint function has a utilization penalty. For example, if 25 Source #2 connections are multiplexed, the number of admissible lecture connections can be increased by 29% (from 38 to 49) by avoiding the concavity restriction.

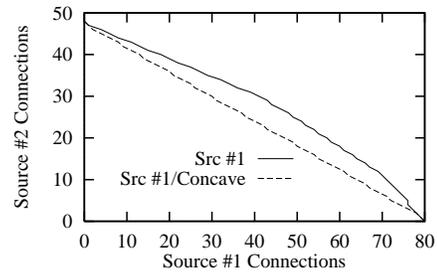


Figure 13: Utilization Cost of Concave Approximation

Finally, Figure 13 indicates that the utilization advantages of considering non-concave constraint functions are only apparent in the heterogeneous cases. This is evident from the figure in that both of the curves join at the axes intercepts showing that there is no utilization penalty when considering homogeneous connections. This can also be seen in terms of the constraint function as illustrated in Figure 14. The figure shows the calculation of the delay bound as in the admission control test illustrated in Equation 4. For homogeneous sources, the sum of the constraint curves may look like the upper curve of the figure. When calculating the maximum backlog for curves such as in Figure 14, the maximum backlog will be B_1 or B_3 depending on the shape of the original constraint curve, but never B_2 . Thus, the potential advantage of having non-concave traffic constraint functions can only be seen when there are heterogeneous sources.

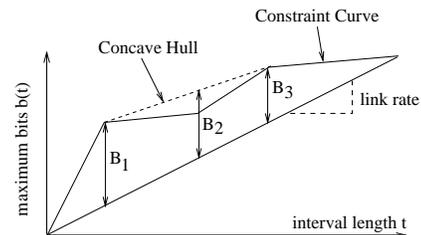


Figure 14: Illustration of Delay Calculation

7 Conclusion

In this paper, we introduced a new traffic model for providing performance guarantees in integrated services networks. The model, termed D-BIND or Deterministic Bounding INterval-length Dependent, consists of a number of rate-interval pairs that are specified to the network at connection setup time. The model captures the key properties of streams needed for resource allocation, namely, the streams' bounding or worst case rates, and the corresponding durations or interval lengths of these rates. The increased accuracy of the new model as compared to previous models allows for more efficient allocation of network resources, or equivalently, for higher network utilization for a given QoS guarantee.

Focusing on a deterministic network service, we quantified the utilization that the new model can achieve by performing a set of experiments with traces of MPEG-compressed video. For a lecture sequence with bounding rates that quickly decreased with increasing interval length, we found that network utilizations of over 60% are achievable. For burstier sources such as an advertisements sequence which has more slowly decreasing bounding rates, network utilizations of approximately 25% are achievable. We consider these utilizations to be considerably high given that the provided QoS avoids any packet losses or delay-bound violations. Moreover, these utilizations for D-BIND are substantially above those for other deterministic traffic models.

In order to achieve higher utilizations than those reported above, a statistical multiplexing gain must be exploited. Of course, such a utilization gain is not for free in that with statistical sharing of network resources, clients obtain a probabilistic guarantee on loss and delay-bound rather than an absolute guarantee as in the deterministic case. To provide a statistical service, the network needs to know the stochastic properties of the traffic streams. However, if sources specify a *stochastic* traffic characterization to the network, the network may not be able to enforce the source's specification. In [16], a solution to this problem is offered that extracts statistical properties of streams from the D-BIND model's worst-case characterization, retaining the enforcibility of the traffic specification while simultaneously achieving a statistical multiplexing gain. The accuracy of the D-BIND model translates into accuracy of the extracted statistics so that such an approach can achieve a significant statistical multiplexing gain.

With the D-BIND model, a full range of network services can be provided to network clients, including the deterministic service investigated here, as well as a renegotiated service [30] and a statistical service [16]. Which services are used by future network clients will depend on both the performance requirements of the client as well as the "cost" of delivering the service. The role of the D-BIND model is to create an efficient foundation for building and integrating these different network services by accurately characterizing traffic streams so that each of these services achieves a high level of resource utilization for the provided QoS.

References

- [1] D. Anick, D. Mitra, and M. Sondhi. Stochastic theory of a data-handling system with multiple sources. *Bell System Technical Journal*, 61:1971–1894, 1982.
- [2] A. Banerjee and S. Keshav. Queueing delays in rate controlled networks. In *Proceedings of IEEE INFOCOM'93*, pages 547–556, San Francisco, CA, April 1993.
- [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. Client requirements for real-time communication services. *IEEE Communications Magazine*, 28(11):65–72, November 1990.
- [7] 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.
- [8] ATM Forum. ATM user-network interface specification. version 3.1, September 1994. ATM Forum document.
- [9] D. Le Gall. MPEG: A video compression standard for multimedia applications. *Communications of the ACM*, 34(4):46–58, April 1991.
- [10] M. Garrett and W. Willinger. Analysis, modeling and generation of self-similar VBR video traffic. In *Proceedings of ACM SIGCOMM'94*, London, UK, August 1994.
- [11] L. Georgiadis, R. Guérin, and V. Peris. The effect of traffic shaping in efficiently providing end-to-end performance guarantees. Technical Report RC 20014, IBM Research Division, Yorktown Heights, NY, April 1995.
- [12] L. Georgiadis, R. Guérin, V. Peris, and K. Sivaraman. Efficient network QoS provisioning based on per node traffic shaping. *IEEE/ACM Transactions on Networking*, 4(4):482–501, August 1996.
- [13] S. Golestani. A framing strategy for congestion management. *IEEE Journal on Selected Areas in Communications*, 9(7):1064–1077, September 1991.
- [14] M. Grossglauser, S. Keshav, and D. Tse. RCBR: A simple and efficient service for multiple time-scale traffic. In *ACM SIGCOMM'95*, Boston, MA, September 1995.
- [15] J. Hyman, A. Lazar, and G. Pacifici. Real-time scheduling with quality of service constraints. *IEEE Journal on Selected Areas of Communications*, 9(7):1052–1063, September 1991.
- [16] E. Knightly. H-BIND: A new approach to providing statistical performance guarantees to VBR traffic. In *Proceedings of IEEE INFOCOM'96*, pages 1091–1099, San Francisco, CA, March 1996.
- [17] E. Knightly, R. Mines, and H. Zhang. Deterministic characterization and network utilizations for several distributed real-time applications. In *Proceedings of IEEE WORDS'94*, Dana Point, CA, October 1994.
- [18] E. Knightly and P. Rossaro. On the effects of smoothing for deterministic QoS. *Distributed Systems Engineering Journal: Special Issue on Quality of Service*. To appear in 1996.
- [19] J. Liebeherr, D. Wrege, and D. Ferrari. Exact admission control for networks with bounded delay services. *IEEE/ACM Transactions on Networking*, 4(6), December 1996.
- [20] D. Lindley. On the theory of queues with a single server. *Proceedings of the Cambridge Philosophical Society*, 48:277–289, 1952.
- [21] S. Low and P. Varaiya. Burst reducing servers in ATM networks. *Queueing Systems, Theory and Applications*, 20(1-2):61–84, 1995.
- [22] 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.
- [23] I. Nikolaidis and R. Onvural. A Bibliography on Performance Issues in ATM Networks. *Computer Communications Review*, 22(5):8–23, October 1992.
- [24] A. Parekh and R. Gallager. A generalized processor sharing approach to flow control in integrated services networks: the multiple node case. *IEEE/ACM Transactions on Networking*,

2(2):137–150, April 1994.

- [25] K. Patel, B. Smith, and L. Rowe. Performance of a software MPEG video decoder. In *Proceedings of the 1st ACM International Conference on Multimedia*, pages 75–82, Anaheim, CA, August 1993.
- [26] J. Turner. New directions in communications (or which way to the information age?). *IEEE Communications Magazine*, 24(10), October 1986.
- [27] D. Wrege, E. Knightly, H. Zhang, and J. Liebeherr. Deterministic delay bounds for VBR video in packet-switching networks: Fundamental limits and practical tradeoffs. *IEEE/ACM Transactions on Networking*, 4(3):352–362, June 1996.
- [28] 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.
- [29] H. Zhang and D. Ferrari. Rate-controlled service disciplines. *Journal of High Speed Networks*, 3(4):389–412, 1994.
- [30] H. Zhang and E. Knightly. RED-VBR: A new approach to support VBR video in packet-switching networks. In *Proceedings of IEEE NOSSDAV'95*, pages 275–286, Durham, NH, April 1995. Extended version to appear in *ACM Multimedia Systems Journal*, May 1997.
- [31] H. Zhang and E. Knightly. RCSP and Stop-and-Go: A comparison of two non-work-conserving disciplines for supporting multimedia communication. *ACM Multimedia Systems Journal*, 4(6), December 1996.