In Proceedings of IEEE INFOCOM '98

Enforceable Quality of Service Guarantees for Bursty Traffic Streams

Edward W. Knightly
ECE Department, Rice University
knightly@ece.rice.edu

Abstract

Providing statistical quality-of-service guarantees introduces the conflicting requirements for both deterministic traffic models to isolate and police users and statistical multiplexing to efficiently utilize and share network resources. We address this issue by introducing two schemes for providing statistical services to deterministically policed sources: (1) adversarial mode *resource allocation in which we* bound *the stochastic envelopes of policed streams and provide a statistical service for adversarial or worst case sources and (2)* non-adversarial mode *allocation in which we* approximate *the stochastic envelopes of policed, but non-worst-case streams in order to exploit a further statistical multiplexing gain in the typical case. Our key technique is to study the problem within the domain of deterministic and stochastic traffic envelopes, which allows us to explicitly consider sources with rate variations over multiple time scales, obtain results for any deterministic traffic model, and apply accurate admission control tests for buffered priority schedulers. We evaluate the scheme's performance with experiments using traces of compressed video and show that substantial statistical multiplexing gains are achieved.*

1 Introduction

Future integrated services networks will multiplex bursty multimedia traffic streams while simultaneously providing them with Quality of Service (QoS) guarantees in terms of throughput and end-to-end delay. To provide such guarantees, network resources will be reserved according to both the applications' specified traffic parameters as well as their QoS requirements.

Deterministic QoS guarantees were studied in [1-4] to support applications that have stringent performance requirements for a service without packet drops or delay bound violations. In addition to this *absolute* guarantee, a deterministic service also has the advantage of *enforceability*: when the network guarantees QoS based on the clients' worst-case descriptions of their traffic, the network can easily verify that these traffic specifications are satisfied. On the other hand, the most important drawback of a deterministic service is that, by its very nature, it must reserve resources according to a worst-case scenario, and hence it cannot achieve a statistical multiplexing gain.

To overcome the utilization limits of a deterministic service, statistical multiplexing must be introduced to exploit the fact that the worst-case scenario will occur quite rarely. To account for such statistical resource sharing and provide a statistical QoS guarantee, the traffic streams' rate fluctuations and time correlations must be characterized. In the literature, such properties are often represented via stochastic traffic models, including Markov Modulated, Self-Similar, and others [5-8]. However, in a shared public network with misbehaving or malfunctioning users, resources should not be allocated according to such stochastic source characterizations since they are inherently difficult for the network to enforce or police. This work therefore addresses the fundamental conflicting requirements for both deterministic traffic models to isolate and police users, and statistical multiplexing to efficiently utilize network resources.

In this paper, we describe a new envelope-based approach to providing statistical QoS guarantees to deterministically policed streams. Our key technique is development of new methods to determine a stream's stochastic envelope directly from its deterministic envelope, or equivalently, from the parameters of the policer, without requiring explicit determination of a particular worst-case arrival pattern. We show that with an accurate deterministic traffic model such as the D-BIND model of [2], the correlation structure and relevant time scales of the traffic which are characterized by the deterministic model are preserved in the inferred stochastic envelope. Thus, we explicitly consider the case of highly bursty traffic streams with rate variations over multiple time scales, so that if the source's deterministic envelope indicates that the stream exhibits multiple time scale behavior, so will the inferred stochastic envelope.

We focus on developing two methods for obtaining a stream's stochastic envelope from its network-enforceable parameters. First, we show how a stream's stochastic envelope can be *bounded* based on the parameters of the policer, so that regardless of how adversarial a stream may be, the network's stochastic characterization of the stream is not violated. Second, we show how the stochastic envelope of a deterministically constrained stream may be *approximated* for a non-adversarial policed stream, which, while bursty, does not always transmit its traffic in the statistically worst possible way. Compared to the former approach, the latter approach is able to achieve an increased statistical multiplexing gain, since the extracted stochas-

tic envelopes are based on properties of non-adversarial policed streams rather than worst-case ones. We term these two modes of inferring stochastic properties "adversarial mode" and "non-adversarial mode" respectively. Finally, once the stochastic envelopes of the deterministically constrained streams are obtained, we apply the admission control test of [9], which we experimentally found to be highly accurate, even for traffic streams with rate variations over multiple time scales.

This work differs from previous studies of statistical resource allocation for policed traffic [10-16] in several aspects. First, using envelope-based techniques, we find that policed sources can exhibit multiple times scale rate variations and have characteristics quite different from the periodic on-off sources studied previously. Indeed, we find that characterizing bursty traffic streams such as compressed video with a single time scale model can result in significant inaccuracies in the resource allocation algorithm. Second, by determining a policed stream's maximal envelope rather than its worst-case arrival sequence *per se*, we obtain computationally simple admission control tests, unlike [17], for example, which requires the solution to an optimization problem to find the worst-case arrival pattern subject to the policing constraints, or [16], which requires an optimization over possible buffer and bandwidth allocations. Third, in contrast to [10-16] as well as our previous work of [17], we investigate resource allocation for *non-adversarial* policed sources (in addition to adversarial ones) to exploit higher statistical multiplexing gains than are possible in a completely adversarial scenario. Finally, our approach applies to any deterministic traffic model and provides admission control tests for buffered priority schedulers.

To illustrate the scheme's performance, from trace-driven-simulation experiments with MPEG-compressed video traces and a 45 Mbps link with a buffer size corresponding to 20 msec delay, the measured maximum achievable utilization is 86% for a loss probability of 10^{-6}. For this same scenario, our adversarial-mode admission control test utilizes resources to 41%, necessarily lower than that of the trace-driven simulation since the scheme assumes that each stream is independently adversarial, which is not the case for these video streams. This represents a better estimate of the admissible region than [12] which obtains 14% utilization in this case. Finally, our non-adversarial-mode admission control test utilizes resources to 79%, achieving most of the statistical multiplexing gain by considering policed, but non-worst-case streams. Indeed, with non-adversarial-mode allocation, we find that once traffic streams are aggregated and economies-of-scale are present, even simple approximate mappings of deterministic to stochastic envelopes can lead to considerably accurate admission control tests.

The remainder of this paper is organized as follows. In Section 2, we describe the important components of both deterministic and statistical network services. In Section 3, we present the scheme for extracting stochastic envelopes of traffic streams from their enforceable parameters, which we apply to admission control in Section 4. Finally, we evaluate the scheme experimentally in Section 5.

2 Guaranteed Services

A network service that guarantees QoS may be classified as either a deterministic service, which provides an absolute guarantee, or a statistical service, which provides a statistical performance guarantee.

2.1 Deterministic Service

A deterministic service supports applications requiring that no packets are dropped due to buffer overflows and that no packets violate their guaranteed end-to-end delay bounds. The primary components of a deterministic service are the parameterized traffic model, which provides the network with a worst-case description of a source's arrivals, and the admission control test, which determines whether each source's QoS requirement can be met, even in a worst-case scenario, e.g., if all streams simultaneously send a burst of traffic.

When utilizing a deterministic service, network clients specify their traffic characteristics to the network via a deterministic traffic model which upper bounds the streams' arrivals. Specifically, a deterministic traffic model uses parameters to define a traffic constraint function $b(t)$, which constrains or bounds the number of bits that a source can transmit over any interval of length t. Denoting $A_j[s_1, s_2]$ as the number connection j arrivals in the interval $[s_1, s_2]$, a traffic constraint function (and deterministic envelope) $b_j(t)$ bounds an arrival sequence A_j if

$$A_j[s, s+t] \leq b_j(t), \ \forall \ s, t > 0. \tag{1}$$

Different traffic models parameterize different constraint functions $b(t)$. For example, the (σ, ρ) or leaky-bucket traffic model [1] defines a constraint function $b(t) = \sigma + \rho t$ so that a source is allowed to send a burst of size σ bits in an arbitrarily small interval, but over longer interval lengths, the source is constrained to an upper-average rate of ρ bits-per-second.

We introduced a more accurate traffic model, termed D-BIND, in [2] to better characterize the burstiness properties of realistic traffic streams. With the D-BIND model, sources characterize 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. With P rate-interval pairs, the model parameterizes a piece-wise linear constraint function with P linear segments given by

$$b(t) = \frac{R_k I_k - R_{k-1} I_{k-1}}{I_k - I_{k-1}}(t - I_k) + R_k I_k, \quad I_{k-1} < t \leq I_k \tag{2}$$

with $I_0 = 0$. We also showed how this source characterization captures a stream's burstiness properties and temporal correlation structure, even over long time scales [2] . For example, with an MPEG-compressed video source, the stream's pattern of alternation between large intra-coded frames and smaller inter-coded frames is evident from the values of the rate-interval pairs.

In [4], a $(\vec{\sigma}, \vec{\rho})$ model is considered along with the above traffic models. This model consists of P (σ_k, ρ_k) leaky buckets in parallel such that the resulting constraint function is piece-wise linear *concave* with P linear segments:

$$b(t) = \min_{1 \leq k \leq P} (\sigma_k + \rho_k \ t). \tag{3}$$

The $(\vec{\sigma}, \vec{\rho})$ model is therefore a special case of the D-BIND model.

All of the above deterministic traffic models have the property that they are enforceable by the network so that when a client specifies its traffic parameters to the network, the network can *verify* that these parameters are satisfied via policing elements such as multi-level leaky buckets [1]. As illustrated in Figure 1, regardless of the traffic stream's arrival pattern at the entrance of the policer, by delaying or dropping packets that violate the traffic parameters specified by the client, the network is assured that Equation (1) is satisfied at the output of the policer.

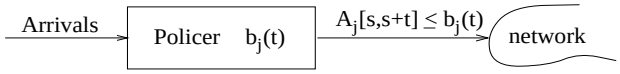

Figure 1: Policing of the Traffic Constraint Function $b_j(t)$

While a deterministic service has important advantages in terms of the strength of the guarantee itself as well as the enforceability of the traffic specification, it can have a significant limitation in terms of the achievable utilization of the network's resources. We showed in [4], that when MPEG-compressed action movie videos with multiple time scale rate variations obtain a deterministic QoS guarantee, the utilization of the multiplexer is limited to 30% for delay bounds less than 100 msec.

2.2 Statistical Service

A statistical network service provides a probabilistic performance guarantee; it achieves a statistical multiplexing gain by exploiting stochastic properties of individual traffic streams as well as statistical independence among streams. However, strong assumptions on the stochastic properties of traffic streams are inherently difficult for the network to enforce or police. For example, consider a Markovian source: in real time, it is impractical to determine whether a stream is following a certain transition matrix, is close enough to its implied marginal distribution, or has the required autocorrelation structure. Consequently, if a particular application does not conform to the chosen stochastic model, no guarantees can be made. Moreover, if admitted to the network, such a stream could adversely affect the performance of other applications if it is statistically multiplexed with them.

In this paper, we focus on providing statistical services using enforceable deterministic traffic models.

3 Enforceable QoS Guarantees

In this section, we introduce techniques for obtaining stochastic envelopes of traffic streams from their policeable parameters in order to achieve a statistical multiplexing gain and provide a network-enforceable statistical service.

3.1 Extracted Stochastic Envelopes

The stochastic envelope that we consider in this paper is the rate-variance envelope $RV_j(t)$ which describes the variance of a stream's arrival rate over intervals of length t [9]:

$$RV_j(t) = Var\left(\frac{A_j[s, s+t]}{t}\right). \tag{4}$$

This characterization captures the second moment correlation structure of an arrival process in the same way as an autocorrelation function or a variance-time characterization used in [18]. We use the $RV_j(t)$ characterization simply because it relates more directly to admission control. Additionally, we restrict ourselves to a second moment envelope rather than considering, for example, envelopes of distributions or moment generating functions, in order to devise an admission control algorithm that is computationally simple as well as highly accurate.

We now present two methods for obtaining $RV_j(t)$ from the deterministic envelope $b_j(t)$: *adversarial* mode in which the rate-variance envelope is upper bounded, and *non-adversarial* mode in which it is approximated for policed but non-worst-case traffic. We refer to the respective envelopes as $RV_j^*(t)$ (adversarial) and $\widehat{RV}_j(t)$ (non-adversarial).

3.2 Adversarial Mode

Here, we derive a bound on the rate-variance envelope of a deterministically policed stream. We refer to this as "adversarial mode" for mapping $b_j(t)$ to $RV_j^*(t)$, since by bounding the rate-variance envelope, the network is assured that regardless of the behavior of the original source, the stochastic envelope at the policer's output is upper bounded. Together with the appropriate admission control algorithm, the enforceable $RV_j^*(t)$ characterization provides a mechanism for delivering a statistical service that is able to extract a statistical multiplexing gain even if all sources are independently adversarial, i.e., if sources are adversarial, but not collusive.

The following proposition shows how a stream's stochastic envelope can be upper bounded from the parameters of the policing elements.

Proposition 1 *If stream j is stationary and its arrivals are upper bounded such that $A_j[s, s+t] \leq b_j(t)$ for all s, $t > 0$, then its rate-variance envelope is upper bounded by:*

$$RV_j^*(t) \leq \frac{\phi_j b_j(t)}{t} - \phi_j^2 \tag{5}$$

where ϕ_j is defined as:[1]

$$\phi_j = \lim_{t \to \infty} \frac{b_j(t)}{t}. \tag{6}$$

Proof: Let the random variable $r_j(s)$ represent source j's instantaneous rate at time s and let $a_j(t)$ represent the total arrivals in an interval of length t,

$$a_j(t) = \int_s^{s+t} r_j(s)\, ds \tag{7}$$

which depends only on t for stationary sources.

[1] For example, for a source parameterized by multiple (σ_k, ρ_k) pairs as in Equation (3), ϕ_j is simply the minimum of the ρ_k's.

Denoting $f_{t,j}(x)$ as the distribution of $a_j(t)$, we show that for any t, the maximal value of $RV_j(t) = Var(a_j(t)/t)$ subject to the constraints of the policing elements

$$\int_s^{s+t} r_j(s)\, ds \le b_j(t) \quad \forall s, t \ge 0 \tag{8}$$

is given by Equation (5) and is attained when the distribution of $a_j(t)$ is given by

$$f^*_{t,j}(x) = \left(\frac{b_j(t) - \phi_j t}{b_j(t)}\right) \delta(x) + \frac{\phi_j t}{b_j(t)}\, \delta(x - b_j(t)) \tag{9}$$

such that for an interval length t, Equation (9) describes a binomial distribution.

According to (8), $f_{t,j}(x) = 0$ for $x > b(t)$ and $x < 0$ so that the rate-variance envelope of a policed stream is given by

$$\begin{aligned} RV_j(t) &= \frac{Ea_j(t)^2 - (Ea_j(t))^2}{t^2} \qquad (10)\\ &= \frac{1}{t^2}\int_0^{b_j(t)} x^2 dF_{t,j}(x) - \frac{1}{t^2}\left(\int_0^{b_j(t)} x dF_{t,j}(x)\right)^2 \end{aligned}$$

for some distribution $f_{t,j}(x)$ satisfying (8). For the distribution $f^*_{t,j}(x)$ of Equation (9), $RV_j^*(t)$ is given by Equation (5). To show that $RV_j^*(t) \ge RV_j(t)$ for all t and for all distributions $f_{t,j}(x)$ satisfying (8), observe that

$$\begin{aligned} RV_j^*(t) - RV_j(t) &= \frac{b_j(t)\phi_j}{t} - \frac{1}{t^2}\int_0^{b_j(t)} x^2 dF_{t,j}(x)\\ &= \frac{b_j(t)}{t^2}\int_0^{b_j(t)} x\, dF_{t,j}(x) - \frac{1}{t^2}\int_0^{b_j(t)} x^2 dF_{t,j}(x) \end{aligned} \tag{11}$$

since the mean rate $Ea_j(t)/t$ is given by

$$\frac{1}{t}\int_0^{b_j(t)} x\, dF_{t,j}(x) = \phi_j.$$

Rewriting Equation (11),

$$RV_j^*(t) - RV_j(t) = \frac{1}{t}\int_0^{b_j(t)} \frac{x b_j(t)}{t}\left(1 - \frac{x}{b_j(t)}\right) dF_{t,j}(x)$$

which is clearly non-negative. □

Proposition 1 applies to any deterministic traffic model since each deterministic traffic model parameterizes a constraint function $b_j(t)$ as described in Section 2.1; the more accurately the model characterizes the traffic stream, the tighter the corresponding bound on $RV_j^*(t)$.

We also note that for an adversarial source to realize the variance bound at a time-scale T, it would first transmit its maximal burst such that $A_j[0,t] = b_j(t)$ for $t \le T$. Next, the source would remain idle in order to obtain enough credits or tokens from the policer to send this same burst of size $b_j(T)$ again. This is different than a "greedy" source defined in [19] which always transmits a packet when allowed to do so by the policer and never remains idle to collect tokens; for a greedy source, $A_j[0,t] = b_j(t)$ for all t. For example, consider a (σ, ρ) source with $b_j(t) = \sigma_j + \rho_j t$. A greedy source would send a burst of size σ_j bits at $t = 0$ and then send traffic at constant rate ρ_j for the remainder of the connection's lifetime. In contrast, a source that alternately sends bursts of size σ_j and remains idle for a time σ_j/ρ_j has the same mean but greater variance and hence is more adversarial for statistical multiplexing. The admission control test of Section 4 shows how rate variation at different time scales leads to an increased loss probability and delay-bound violation probability.

3.3 Non-Adversarial Mode

By upper bounding the stochastic envelope of policed traffic streams as in Proposition 1, a statistical service can be provided even in the case that all traffic streams are independently adversarial. Below, we show how this same rate-variance envelope can be *approximated* for policed arrival streams. In other words, we derive an expression for $\widehat{RV}_j(t)$ to describe the envelope of a non-adversarial policed stream that satisfies Inequality (1).

Proposition 2 *If stream j is stationary and its arrivals are upper bounded such that $A_j[s, s+t] \le b_j(t)$ for all s, $t > 0$, then its rate-variance envelope is approximately:*

$$\widehat{RV}_j(t) \approx \frac{\phi_j b_j(t) - t\phi_j^2}{12t} \tag{12}$$

where ϕ_j is defined by Equation (6).

Proof: A deterministic traffic constraint function $b_j(t)$ bounds the worst-case arrivals of connection j. On a time scale T, a burst of size $b_j(T)$ is the largest-sized burst allowed by the policer. But what is the probability or fraction of time that the source transmits such bursts? In Proposition 1, the rate-variance envelope bound $RV_j^*(T)$ is realized when a source sends the worst-case burst at time scale T as often as possible subject to the policer constraints, namely, when the source achieves

$$Prob\{a_j(t) = b_j(t)\} \le \frac{\phi_j t}{b_j(t)}. \tag{13}$$

Contrastly, we define a "non-adversarial" policed traffic stream as one in which $\hat{a}_j(t)$ can take on values across its entire policeable range $[0, b_j(t)]$ rather than only its extreme values, 0 and $b_j(t)$, that is,

$$\hat{f}_{t,j}(x) > 0, \quad 0 \le x \le b_j(t). \tag{14}$$

In particular, denoting a uniform distribution on $[x_1, x_2]$ as $U[x_1, x_2]$ we define the distribution of $\hat{a}_j(t)$ as

$$\hat{a}_j(t) \sim \begin{cases} U[0, \phi_j t] & w.p.\ \frac{\phi_j t}{b_j(t)} \\ U[\phi_j t, b_j(t)] & w.p.\ 1 - \frac{\phi_j t}{b_j(t)} \end{cases} \tag{15}$$

so that $\widehat{RV}_j(t)$ is given by Equation (12). □

In other words, with probability $\phi_j t/b_j(t)$, $\hat{a}_j(t)$ is distributed uniformly between 0 and the mean number of bits in intervals of length t, $\phi_j t$; and with probability $1 - \phi_j t/b_j(t)$ it is uniformly distributed between the mean and the maximum, $b_j(t)$.

Importantly, we note that while Equation (15) is that of a weighted uniform distribution, this arrival characterization has no

relationship to a "uniform" source in the traditional sense, i.e., a source with *iid* uniform interarrival times. In contrast, the rate-variance envelopes of Propositions 1 and 2 allow for an arbitrary autocorrelation structure over any time scales. The exact form of $RV(t)$ will depend on the traffic model used to bound the arrival stream and the parameter values for that particular stream.

Finally, by comparing Equations (5) and (12) we notice the relationship between $RV_j^*(t)$ and $\widehat{RV}_j(t)$. This is due to the relationship between an extremal distribution which takes on values of 0 and $b_j(t)$ and the weighted uniform distribution.

3.4 Example Envelopes

Figure 2 illustrates the rate-variance envelopes obtained from Propositions 1 and 2 for the MPEG-compressed video trace described in Section 5. The curve labeled "Actual $RV(t)$" is the true rate-variance envelope as directly computed from the trace as in [9]. To obtain the "Adversarial Mode" and "Non-adversarial Mode" envelopes, we first calculate the deterministic parameters of the source. In particular, we characterize the source with 6 rate-interval pairs using the D-BIND traffic model [2]. These rate-interval pairs, which are policeable by the network, parameterize a traffic constraint function as given by Equation (2), from which $RV_j^*(t)$ is calculated using Proposition 1 and $\widehat{RV}_j(t)$ using Proposition 2.

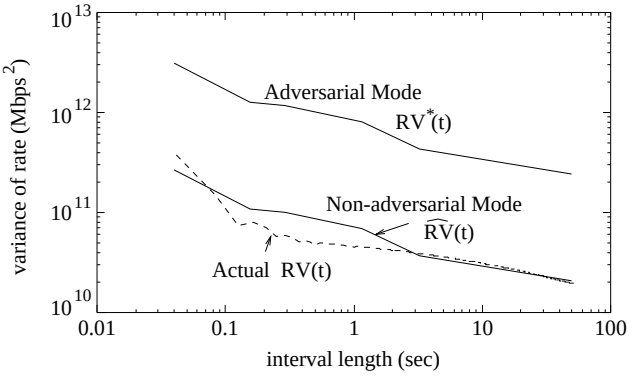

Figure 2: Envelopes from Video Trace

We make the following observations about the figure. First, the trace itself exhibits a non-trivial autocorrelation structure, even over relatively long time scales. This can be seen from the slope of the $RV(t)$ curve as depicted on the figure's log-log scale: if arrivals in successive intervals of length t are uncorrelated, then the slope of this curve would be -1 at t. However, the curve for the actual source has a slope considerably greater than -1 even for large t. Second, we observe that both of the inferred envelopes, $RV^*(t)$ and $\widehat{RV}(t)$ exhibit this same behavior, i.e., they reflect the long-time-scale characteristics of the source. This indicates that even deterministic traffic models are capable of capturing the stochastic properties of sources that exhibit rate variations over multiple time scales. Finally, we observe that the non-adversarial mode rate-variance envelope $\widehat{RV}(t)$ is quite close to the stream's actual envelope $RV(t)$. Thus, even Proposition 2's simple mapping from policeable deterministic parameters to a stochastic envelope is able to approximately characterize the complex dynamics and autocorrelation structure of this highly bursty trace.

4 Admission Control for Policed Streams

Here, we describe an admission control test for policed streams multiplexed at a Static Priority scheduler. In particular, we show how to determine packet loss or delay bound violation probability as a function of the streams' rate-variance envelopes, which in turn can be calculated from the parameters of the policing elements using Propositions 1 and 2. While we focus on a single multiplexer, our approach is applicable across multiple network nodes using techniques from [17]. For example, if traffic streams are reshaped at each network node as in [20, 21], a stream's deterministic envelope $b_j(t)$ is reconstructed, and hence so is its inferred rate-variance envelope.

A rate-controlled static priority scheduler [21] consists of per-stream rate controllers and a number of prioritized FCFS queues (Figure 3). Each stream is assigned a priority p at connection setup time based on its requested QoS, including whether it requires deterministic or statistical service, and on the requested delay bound. Additionally, packets are rate controlled (and hence policed) before being queued to ensure that each stream conforms to its specified deterministic parameters, namely, its envelope $b_j(t)$.

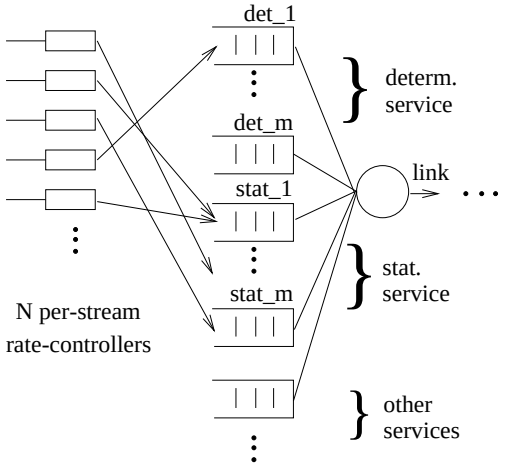

Figure 3: Static Priority Scheduler

As shown in Figure 3, such a scheduler can provide both deterministic and statistical performance guarantees as well as other services. 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 *probabilities* whereas connections utilizing the deterministic service obtain absolute bounds on delay and loss. Thus, our

approach provides uniform semantics for delivering QoS guarantees, allowing all traffic streams to use the same policeable traffic specification, regardless of the service they obtain. Admission control tests that support connections utilizing a deterministic service can be found in [4]. At the lower priority levels, other services can be provided, including measurement-based services and best-effort service.

Denoting $RV_{p,j}(t)$ as the rate-variance envelope of source j at priority level p, the delay bound violation probability at level p for a static priority scheduler with capacity C is approximately

$$P\{D_p > d_p\} \approx \max_{0 \le t \le \beta_p} \frac{1}{\sqrt{2\pi}} exp\left(\frac{-(C(t+d_p) - \mu_{t,p})^2}{2\sigma_{t,p}^2}\right) \quad (16)$$

where

$$\mu_{t,p} = \sum_j t\phi_{p,j} + \sum_{q=1}^{p-1}\sum_j (t+d_p)\phi_{q,j},$$

$$\sigma_{t,p}^2 = \sum_j t^2 RV_{p,j}(t) + \sum_{q=1}^{p-1}\sum_j (t+d_p)^2 RV_{q,j}(t+d_p),$$

and β_p is the busy period bound at priority level p [1]

$$\beta_p = \min\{t > 0 \mid \sum_{q=1}^{p}\sum_j b_{q,j}(t) \le Ct\}. \quad (17)$$

Equation (16) uses [9] together with a well-known approximation for the tail of a Gaussian distribution. The key technique in its derivation is the analysis of the traffic and buffer dynamics within a busy period, which has a duration bounded by Equation (17). Moreover, with a heavy traffic approximation, $\sum_j a_j(t)$ approaches a Gaussian distribution with mean $\sum_j \phi_j t$ and variance $\sum_j t^2 RV_j(t)$. The test then considers that buffer overflows and delay-bound violations can occur at any time-scale up to the maximal busy period. In [9], we experimentally found Equation (16) to be highly accurate in predicting the performance of a buffered priority multiplexer.

Notice that the probability of delay-bound violation is strictly increasing with $RV_j(t)$, so that by considering the maximal rate-variance envelope $RV_j^*(t)$ of each policed source (as in Proposition 1), our estimate of this probability is also maximized. Finally, we note that for both adversarial and non-adversarial allocation, traffic streams must be statistically independent or non-collusive. If traffic streams are collusive, then a fully deterministic approach must be employed [4].

5 Experimental Investigations

In this section, we evaluate our proposed scheme for provisioning enforceable statistical QoS guarantees via a set of trace-driven experiments. With an implementation of the proposed resource reservation scheme of Sections 3 and 4, we compare the streams' performance obtained in trace-driven simulations with that predicted by the admission control tests and $RV(t)$ traffic characterizations.

5.1 Experimental Scenario

The workload consists of a 30 minute trace of MPEG-compressed video taken from an action movie. It was digitized to 384 by 288 pixels and compressed with constant-quality MPEG 1 compression at 24 frames per second with frame pattern IBBPBBPBBPBB. Further details of the trace and its characteristics may be found in [22].

For each simulation, N streams or traces are multiplexed on a simulated 45 Mbps first-come-first-serve link, with each stream's arrival pattern given by the movie trace with a start time chosen uniformly over the length of the trace (30 minutes). For a given number of connections N and buffer size $C \cdot d$ (the link capacity times the delay bound) we measure the fraction of packets ε that are dropped due to buffer overflow. Many simulations are performed with independent start times and average results are reported.

In the admission control part of the experiments, we determine the streams' rate-variance envelopes from their enforceable deterministic parameters as described in Section 3 and depicted in Figure 2. We then use the admission control test of Section 4 to determine the maximum number of admissible connections, N, subject to the QoS constraints for delay, d, and loss probability, ε.

5.2 Results

Here, we compare the results of the trace-driven simulations with the admission control tests. To further evaluate our approach, we also compare with the admission control algorithm of [12].

Figure 4 shows the results of the trace-driven simulation and admission control experiments.[2] The figure shows the average utilization of the multiplexer (which is proportional to the number of connections as $N\phi/C$) versus buffer size scaled to delay. In other words, for a given delay d depicted on the horizontal axis, the vertical axis shows shows the maximum number of connections N (scaled to utilization) that can be multiplexed such that all connections are guaranteed a probability of delay-bound violation or buffer overflow of 10^{-3} in Figure 4(a) and 10^{-6} in Figure 4(b).

In the figures, four curves are depicted (from top to bottom): (1) the results of the trace-driven simulation; (2) admission control tests based on the *Non-Adversarial Mode* $\widetilde{RV}(t)$ traffic characterization of Proposition 2 (an approximate rate-variance envelope for a non-adversarial, but policed, traffic stream); (3) admission control tests based on the *Adversarial Mode* $RV^*(t)$ traffic characterization of Proposition 1 (the worst-case rate-variance envelope of a policed stream); and (4) the admission control test of [12].

Trace-driven Simulation - For the simulation curves of Figures 4(a) and 4(b), the average utilization of the multiplexer, and hence the number of multiplexed connections, increases with increasing delay or buffer size. However, notice that increasing the buffer size beyond that of a 10 to 20 msec delay is of little benefit, i.e., larger buffers will not provide a better QoS or support more connections for a given QoS. Regardless, the utilizations are in

[2] 95% confidence intervals for the simulations are all within a single connection and are therefore not shown.

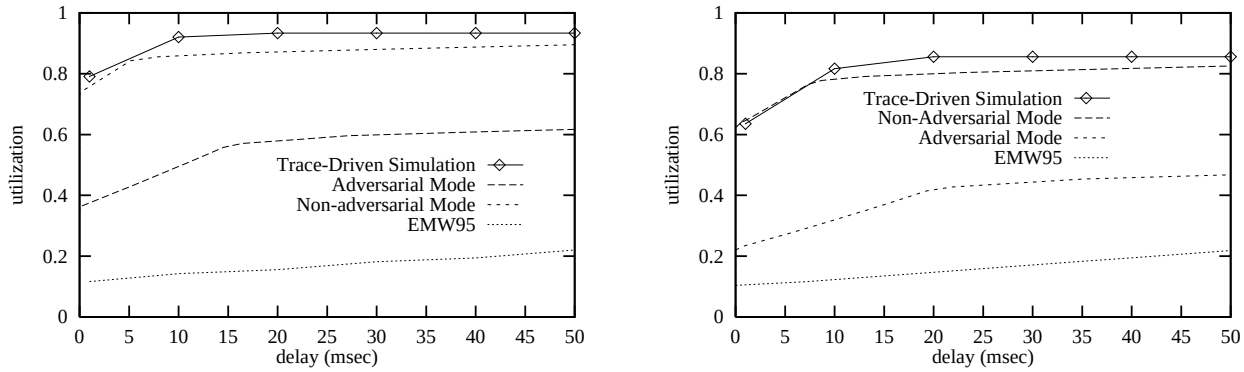

(a) $\varepsilon = 10^{-3}$ (b) $\varepsilon = 10^{-6}$

Figure 4: Utilization vs. Delay Bound

the range of 79% to 92% (61 to 72 connections on the simulated 45 Mbps link) for $\varepsilon = 10^{-3}$, and in the range of 64% to 88% (49 to 68 connections) for $\varepsilon = 10^{-6}$. Such high utilizations indicate that these MPEG streams are well suited to statistical multiplexing, despite their burstiness over multiple time-scales.

Non-Adversarial Mode Admission Control - The second curve from the top depicts the admission control experiments that use the $\widehat{RV}(t)$ characterization for non-adversarial policed streams. Notice that the non-adversarial-mode curves are quite close to those of the trace-driven simulation, indicating that with only knowledge of the streams' deterministic parameters (in this case, six worst-case rate-interval pairs), the scheme of Proposition 2 is able to deliver a statistical service that exploits nearly all of the achievable statistical multiplexing gain.

Adversarial Mode Admission Control - The third curve shows the results of the admission control experiments using the $RV^*(t)$ bound on a policed stream's rate-variance envelope. As described in Proposition 1, $RV^*(t)$ *bounds* the stochastic properties of policed streams so that statistical QoS guarantees can be provided even if all streams are independently adversarial. Consequently, the $RV^*(t)$ envelope is necessarily more pessimistic than the $\widehat{RV}(t)$ envelope for non-worst-case policed streams (cf. Figure 2) so that the adversarial-mode scheme captures some, but not all, of the possible statistical multiplexing gain. Its utilizations are 38% to 64% ($\varepsilon = 10^{-3}$) and 23% to 48% ($\varepsilon = 10^{-6}$) for delays between 1 and 50 msec, utilizations that are considerably below that of the trace-driven simulation. However, despite not capturing all of the multiplexing gain, this scheme does have a distinct advantage in terms of protection: if there were many *adversarial* sources rather than MPEG video sources (the MPEG trace is bursty, but not worst-case), then the adversarial mode service is still able to deliver a rigorous statistical QoS guarantee.

EMW95 Admission Control - The final curve depicts admission control experiments based on [12]. Here, the trace is characterized using the dual leaky bucket model with a peak rate of 5.87 Mbps, a maximum burst time (T_{on}) of 41.7 msec, and an upper average rate ρ of 1.98 Mbps. The test assumes that sources transmit traffic according to an extremal periodic on-off model with these parameters. As shown, the test is quite conservative limiting utilization of the multiplexer to 12% to 22% for delays less than 50 msec and loss probabilities less than 10^{-3}. Moreover, for the smaller loss probability of $\varepsilon = 10^{-6}$ in Figure 4(b), the admission control test of [12] deemed the situation "non-statistically-multiplexable" so that the admissible region shown in the figure is that of the deterministic (no-loss case) as computed in [1]. The primary reasons for this conservatism are two-fold. First, the dual-leaky bucket traffic model captures only a single time scale of the source, and characterizing such highly bursty traffic streams as compressed video with an on-off model is necessarily restrictive [23]. While one could capture longer time scale behavior by characterizing the source with a smaller value of ρ and a larger T_{on}, the values of T_{on} required to significantly reduce ρ are so large that performance gains are not achieved. Second, in [9], we showed that resource allocation using Equation (16) and rate-variance envelopes can be highly effective for sources with rate variations over multiple time scales; this approach eliminates the need for a number of conservative approximations in [12].

6 Conclusions

Providing statistical performance guarantees in networks encounters a conflicting requirement between the need to obtain a statistical multiplexing gain, which often engenders the use of a *statistical* traffic model, and the need to police traffic streams, which necessitates a *deterministic* traffic model. In this paper, we introduced a new approach for delivering a statistical service that extracts a traffic stream's stochastic envelope from its network-enforceable deterministic parameters. We first showed how to bound a policed stream's rate-variance envelope to provide a statistical service and achieve a statistical multiplexing gain even in the case that all traffic sources are independently adversarial. We then showed how to approximate this same rate-variance envelope for perhaps the more typical case of policed, but non-worst-case traffic streams; this latter approach allows the network to exploit a further statistical multiplexing gain when multiplexing non-adversarial sources. The key components of our approach are (1) simple-to-compute mechanisms to bound and approximate

stochastic envelopes from enforceable deterministic parameters, (2) use of an accurate deterministic model to characterize the important properties of the traffic, and (3) stochastic envelope based admission control tests for buffered, priority multiplexers. Evaluations of our approach with experiments using compressed video traces showed that the scheme is able to achieve a substantial statistical multiplexing gain.

References

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

[2] E. Knightly and H. Zhang. D-BIND: An accurate traffic model for providing QoS guarantees to VBR traffic. *IEEE/ACM Transactions on Networking*, 5(2):219–231, April 1997.

[3] J. Liebeherr, D. Wrege, and D. Ferrari. Exact admission control for networks with bounded delay services. *IEEE/ACM Transactions on Networking*, 4(6):885–901, December 1996.

[4] 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.

[5] N. Adas. Traffic models in broadband networks. *IEEE Communications*, 35(7):82–89, July 1997.

[6] M. Garret and W. Willinger. Analysis, modeling and generation of self-similar VBR video traffic. In *Proceedings of ACM SIGCOMM'94*, pages 269–280, London, UK, August 1994.

[7] M. Krunz, R. Sass, and H. Hughes. Statistical characteristics and multiplexing of MPEG streams. In *Proceedings of IEEE INFOCOM'95*, pages 455–462, Boston, MA, April 1995.

[8] A. Lazar, G. Pacifici, and D. Pendarakis. Modeling video sources for real time scheduling. *ACM Multimedia Systems Journal*, 1(6):253–266, April 1994.

[9] E. Knightly. Second moment resource allocation in multi-service networks. In *Proceedings of ACM SIGMETRICS '97*, pages 181–191, Seattle, WA, June 1997.

[10] A. Baiocchi, N. Blefari-Melazzi, F. Cuomo, and M. Listanti. Worst deterministic pattern allocation: a viable approach to attain statistical gain in ATM. In *IEEE ICC '94*, pages 106–110, New Orleans, LA, May 1994.

[11] I. Cidon, R. Guérin, and I. Kessler A. Khamisy. Analysis of a statistical multiplexer with generalized periodic sources. *Queueing Systems, Theory and Applications*, 20(1-2):139–169, 1995.

[12] 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.

[13] 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.

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

[15] P. Oechslin. Worst case arrivals of leaky bucket constrained sources: the myth of the on-off source. In *Proceedings of the 1997 International Workshop on Quality of Service*, pages 67–76, New York, NY, May 1997.

[16] F. Lo Presti, Z. Zhang, D. Towsley, and J. Kurose. Source time scale and optimal buffer/bandwidth tradeoff for regulated traffic in an ATM node. In *Proceedings of IEEE INFOCOM '97*, Kobe, Japan, April 1997.

[17] 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.

[18] W. Leland, M. Taqqu, W. Willinger, and D. Wilson. On the self-similar nature of Ethernet traffic. *IEEE/ACM Transactions on Networking*, 2(1):1–15, February 1994.

[19] A. Parekh and R. Gallager. A generalized processor sharing approach to flow control in integrated services networks: the single-node case. *IEEE/ACM Transactions on Networking*, 1(3):344–357, June 1993.

[20] L. Georgiadis, R. Guérin, V. Peris, and K. Sivarajan. Efficient network QoS provisioning based on per node traffic shaping. *IEEE/ACM Transactions on Networking*, 4(4):482–501, August 1996.

[21] H. Zhang and D. Ferrari. Rate-controlled service disciplines. *Journal of High Speed Networks*, 3(4):389–412, 1994.

[22] O. Rose. Statistical properties of MPEG video traffic and their impact on traffic modeling in ATM systems. In *Proceedings of IEEE Conference on Local Computer Networks*, pages 397–406, Minneapolis, MN, October 1995.

[23] E. Knightly. On the accuracy of admission control tests. In *Proceedings of IEEE ICNP '97*, pages 125–133, Atlanta, GA, October 1997.