To appear in ACM Multimedia Systems Journal, 1999.

Resource Allocation for Multimedia Traffic Flows using Rate-Variance Envelopes

Edward W. Knightly
Department of Electrical and Computer Engineering
Rice University

Abstract

In order for networks to support the delay and loss requirements of interactive multimedia applications, resource management algorithms are needed that efficiently allocate network resources. In this paper, we introduce a new resource allocation scheme based on rate-variance envelopes. Such envelopes capture a flow's burstiness properties and autocorrelation structure by characterizing the variance of its rate distribution over intervals of different length. From this traffic characterization, we develop a simple and efficient resource allocation algorithm for static priority schedulers by employing a Gaussian approximation over intervals and considering a maximal busy period. Our approach supports heterogeneous quality of service requirements via our consideration of prioritized service disciplines, and supports heterogeneous and bursty traffic flows via our general framework of traffic envelopes. To evaluate the scheme, we perform trace-driven simulation experiments with long traces of compressed video and show that our approach is accurate enough to capture most of the available statistical multiplexing gain, achieving average network utilizations of up to 90% for these traces and substantially outperforming alternate schemes.

1 Introduction

Bursty traffic sources that require Quality of Service (QoS) guarantees in terms of loss and delay bound are emerging as one of the most important types of traffic in future integrated services networks. Because of the variable-bit-rate nature and multiple-time-scale correlation characteristics of many realistic sources, e.g., [8, 13, 17, 20], it is difficult to determine the amount of resources that need to be allocated to individual flows such that each flow obtains the performance that it requires. This problem is exacerbated when different flows require different services, such as different throughputs, delay bounds, and loss probabilities, since resources must then be allocated in networks that use *prioritized* service disciplines such as Static Priority (SP) or Earliest Deadline First.

In this paper, we introduce a scheme for allocating resources to heterogeneous and bursty traffic sources that are multiplexed with the static priority service discipline. Our approach is based on a simple and general traffic characterization which we term a *rate-variance envelope*; this envelope describes the variances of the flows' rates as a function of interval length. In the literature, statistical QoS has been studied via envelopes of bounding moment generating functions [2], exponentially bounded envelopes [26, 28], and envelopes consisting of a family of bounding distributions [15, 30]. Statistical envelopes have also been applied to resource allocation for inter-class resource sharing [23] and video-on-demand services [16]. Moreover, *deterministic* traffic envelopes have been studied in the context of worst-case performance guarantees [2, 5, 6, 12, 19, 27]. Here, we use a rate-variance envelope as a simple way to capture the second-moment properties of temporally correlated traffic flows, and to describe how quickly the rate-distribution becomes concentrated at the mean rate with increasing interval-length, a key factor for computing delay-bound-violation probabilities. We show empirical rate-variance envelopes for several long traces of compressed video to show how it captures the burstiness properties of realistic network traffic sources.

This work was supported by NSF CAREER Award ANI-9733610, NSF Grant ANI-9730104, and Nokia Corporation. The author can be reached via http://www.ece.rice.edu/networks.

Based on this envelope, we introduce a resource allocation algorithm that provides a simple-to-compute approximation for QoS parameters such as packet-loss probability and delay-bound-violation probability. Studying the delay incurred up to the maximal busy period of a static priority scheduler, we introduce a Gaussian approximation over intervals (based on the Central Limit Theorem) that allows us to determine QoS parameters from simple computations on the rate-variance envelopes of the traffic flows. Our solution applies to a networking environment that allows for both heterogeneous traffic flows and heterogeneous performance requirements, considering the case of priority service disciplines that are well suited for delivering integrated network services.

Our approach differs from previous work on statistical resource allocation for multimedia traffic in two ways. First, to ensure our approach applies to a general class of traffic flows, we use a stochastic envelope to bound the traffic, rather than attempting to model the arrival process itself with a particular statistical model such as a Markov Modulated models, histogram-based models, and other more sophisticated video models [1, 8, 13, 14, 17, 20, 21, 25]. Second, by approaching the problem from the perspective of traffic envelopes, the resource allocation problem is significantly simplified, so that we are able to study priority service disciplines and heterogeneous traffic flows, a scenario which is quite complex using other techniques (see [31] for example).

Finally, we investigate the effectiveness of the new scheme via trace-driven simulation experiments. We utilize 30-minute traces of MPEG-compressed video and simulate a number of flows aggregating at a SP multiplexer. For a given traffic mix, we compare the QoS actually obtained by the flows in trace-driven simulations with that predicted by the algorithm. The results indicate that for these bursty variable bit rate video sources, the scheme is accurate enough to capture most of the available statistical multiplexing gain and consequently, attains utilizations considerably higher than those of effective bandwidth techniques [4, 9].

The remainder of this paper is organized as follows. In Section 2, we develop the rate-variance admission control algorithm and resource allocation scheme. Next, in Section 3, we investigate our scheme's implications for traffic models and admission control tests in multi-service networks. In Section 4, we perform a set of trace-driven simulation and admission control experiments, and in Section 5, we conclude.

2 Resource Allocation using Rate-Variance Envelopes

Here, we present our approach for using traffic flows' rate-variance envelopes to manage network resources and provide statistical quality of service guarantees. The techniques apply to priority service disciplines that may be used to support network clients that have heterogeneous QoS requirements. The traffic characterization is also quite general in that it allows for an arbitrary autocorrelation structure of individual flows and heterogeneous statistical characteristics among flows.

2.1 Background on Traffic Envelopes

A traffic envelope provides a means of bounding a traffic flow's arrivals over intervals of different length. In particular, denote the arrivals of traffic flow j (in bits for example) in the interval $[s, s+t]$ as $A_j[s, s+t]$.

Deterministic traffic envelopes were introduced by Cruz [5] in defining a constraint function $b_j(t)$ which upper bounds source j's arrivals in any interval of length t. In particular, $b_j(t)$ is a deterministic envelope of traffic flow j if

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

Since such a worst-case bound is enforceable by network policers, properties of deterministic envelopes have primarily been used to provide deterministic network services [2, 5, 6, 12, 19, 27], but also to provide enforceable statistical services [10, 11].

In [15], Kurose introduced a stochastic analogue to the above envelope such that a collection of bounding distributions $B_j(t)$ form a stochastic envelope of arrival process A_j if

$$B_j(t) \geq_{st} A_j[s, s+t] \ \ \forall, s, t > 0 \tag{2}$$

where the random variable X is said to be stochastically greater than the random variable Y, i.e., $X \geq_{st} Y$, if $P(X > x) > P(Y > x)$ for all x. In [15, 30], such stochastic envelopes were used to study end-to-end statistical QoS guarantees.

In [2], Chang studied stochastic traffic envelopes that bound a traffic flows moment generating function. In particular, the function $\hat{B}_j(\theta, t)$ is an envelope of arrival sequence A_j with respect to parameter θ if

$$\hat{B}_j(\theta, t) \geq \frac{1}{\theta} \log E e^{\theta A_j[s, s+t]} \quad \forall, s, t > 0. \tag{3}$$

From this general traffic characterization, Chang also studied end-to-end statistical QoS guarantees and showed how effective bandwidths can be computed from properties of such envelopes.

2.2 Rate-Variance Envelope

Here, we define and study a *rate-variance envelope* such that traffic flow j is bounded by rate-variance envelope $RV_j(t)$ if

$$RV_j(t) \geq \text{var}\left(\frac{A_j[s, s+t]}{t}\right) \quad \forall s, t > 0 \tag{4}$$

Our goal in defining such a traffic envelope is to employ a general yet accurate characterization of a traffic flow in designing computationally simple admission control tests for priority schedulers. Note that we place no restrictions on the autocorrelation function of the arrival sequence A_j so that we consider a general class of flows with bounded variance. Moreover, as shown below, characterizing a flow by the variance of its rate as a function of interval length provides an intuitive representation of the flow's stochastic properties.

To illustrate such an envelope, consider a stationary random process $\{X_1, X_2, X_3, \cdots\}$ which represent the frame sizes of a compressed-video source. Denote the first and second moments of the frame-size distribution by EX and EX^2 respectively so that the variance of the frame size distribution, $\text{var}(X)$, is $EX^2 - (EX)^2$. The flow's mean rate is then $m = EX/T$ where T is the duration of the time slot or frame time.

The sequence $\{X_1, X_2, X_3, \cdots\}$ may have an arbitrary autocorrelation structure and the *variance* of the distribution of the total number of arrivals over n consecutive frames is given by $\text{var}(X_i + X_{i+1} + \cdots + X_{i+n-1})$. By normalizing this variance to the length of the interval, we have the variance of the *rate* distribution over the respective interval length, which is

$$RV(nT) = \text{var}\left(\frac{X_i + X_{i+1} + \cdots + X_{i+n-1}}{nT}\right), \; n \geq 1. \tag{5}$$

which is independent of i under the stationarity assumption; if the process were non-stationary, the rate-variance envelope describes the maximum variance over any interval of length t.

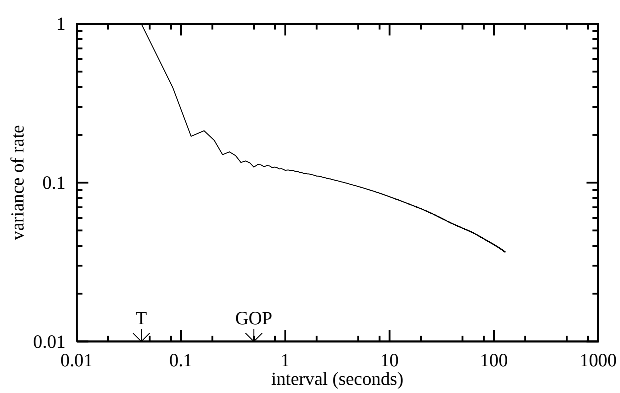

Figure 1: Rate-Variance Envelope for Video Trace

Figure 1 shows the normalized rate-variance envelope for a 30-minute trace of an MPEG-compressed movie. The trace is that of an action movie taken from [24], with the movie digitized to 384x288 pels and compressed at 24 frames per second using the MPEG compression algorithm with frame pattern *IBBPBBPBBPBB*.

For the 30-minute sequence of frame sizes, the figure shows the empirical variance of the rate as given by Equation (4) as a function of the interval length on a log-log scale. The rate-variance depicted on the vertical axis is normalized to the frame-size variance so that $RV(t_1 = 1/24) = 1$. The horizontal axis depicts the interval length in seconds, which is nT, or the interval length in frame-times, n, multiplied by the frame time T, $\frac{1}{24}^{th}$ of a second in this case. The figure depicts intervals of up to 100 seconds rather than up to the length of the trace so that there are an adequate number of sample points to determine the sample variance.

We make several observations about the figure. First, note that the $RV(t)$ characterization differs dramatically from that of an uncorrelated sequence. In other words, without dependencies among the frame sizes, the curve in Figure 1 would depict a straight line with a slope of -1. Indeed, the movie's rate-variance envelope indicates a correlation structure that is present over long interval lengths.

Second, we note that regardless of the correlation structure of the flow, a resource allocation system has the potential to exploit this second-moment characterization. For example, over wider and wider interval-lengths, the distribution of the flow's rate becomes more and more concentrated at the mean rate. How quickly this occurs has implications in the amount of resources that need to be allocated to the flow. Indeed, an uncorrelated sequence will have lower delays for a given average utilization than the one in Figure 1.

Lastly, we note that the shape of the curve in Figure 1 has implications on the relevant time-scales of the flow's correlation structure and whether the flow exhibits long-range dependence (see [8] for example): the resource allocation algorithm we develop below places no restrictions on the shapes of the flows' rate-variance envelopes.

2.3 Envelope-Based Admission Control

As shown in Figure 2, a static priority scheduler consists of a number of prioritized FCFS queues, where each queue has an associated delay bound and probability of delay bound violation. At flow setup time, each flow is assigned a priority level that is based on the requested QoS, including the requested end-to-end delay bound and probability of delay bound violation. During data transmission, each packet of the flow is serviced at its pre-established priority level. Thus, a static priority service discipline has an advantage as compared to many other service disciplines in its simplicity of implementation. For example, in the Earliest Deadline First service discipline, as well as Generalized Processor Sharing [22], when the scheduler determines which packet to service next, it must search through all packets to find the one with the smallest deadline, or for Generalized Processor Sharing, the one with the highest priority index. Alternatively, for static priority, the packet at the head of the highest-priority non-empty queue is always serviced next.

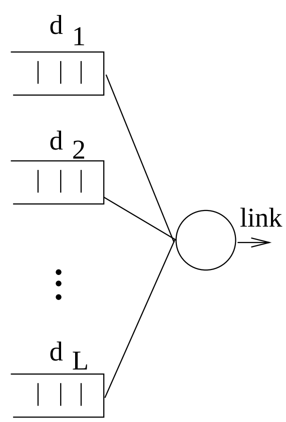

Figure 2: Static Priority Scheduler

The algorithm below describes our admission control test for SP schedulers and statistical service.

Result: *Consider a SP scheduler with link capacity C such that priority level l has an associated delay bound d_l. Let $\mathcal{L}$ be the set of flows at level l, so that the j^{th} flow in $\mathcal{L}$ has mean rate $\phi_{l,j}$, rate-variance envelope $RV_{l,j}(t)$, and deterministic envelope $b_{l,j}(t)$. With application of a Gaussian approximation over* *intervals, the delay-bound violation probability at priority level l is approximately*

$$Prob\{D_l > d_l\} \approx \max_{0\le t\le \beta} Prob\{\hat{B}_l(t) + \hat{B}_{m<l}(t+d_l) \ge C(t+d_l)\} \tag{6}$$

with

$$\hat{B}_l(t) \sim N\left(\sum_{j\in\mathcal{L}} t\phi_{l,j},\ \sum_{j\in\mathcal{L}} t^2 RV_{l,j}(t)\right) \tag{7}$$

and

$$\hat{B}_{m<l}(t+d_l) \sim N\left(\sum_{m=1}^{l-1}\sum_{j\in\mathcal{M}} (t+d_l)\phi_{m,j},\ \sum_{m=1}^{l-1}\sum_{j\in\mathcal{M}} (t+d_l)^2 RV_{m,j}(t+d_l)\right) \tag{8}$$

where $X \sim N(\mu,\sigma^2)$ denotes a Gaussian distribution such that

$$Prob\{X > x\} = \frac{1}{\sqrt{2\pi}}\int_x^\infty exp\left(\frac{-(y-\mu)^2}{2\sigma^2}\right) dy$$

and β denotes the maximum busy period

$$\beta = \min\{t > 0 \mid \sum_l \sum_{j\in\mathcal{L}} b_{l,j}(t) \le Ct\}. \tag{9}$$

Proof. Assume that all queues are initially empty at time 0 and consider a marked level-l packet arriving at time t. Let $\mathcal{P}_l(t)$ denote the packets or bits from level-l flows that have priority over the marked packet. The arrivals of all level-l packets arriving by t is stochastically bounded by

$$\mathcal{P}_l(t) = \sum_{j\in\mathcal{L}} A_{l,j}[0,t] \le_{st} \sum_{j\in\mathcal{L}} B_{l,j}(t) \tag{10}$$

by definition of B in Equation (2). All of these packets have priority over the marked packet as they share the same priority level and arrived earlier than the marked packet. Denoting $\psi_l(A[s,t])$ as the backlog induced by level l arrivals in $[s,t]$ that must be serviced before the marked packet, note that

$$\psi_l(A_{l,j}[t,u]) = 0 \tag{11}$$

for $u > t$ since packets within a priority level are serviced in first-come first-serve order.

At higher priority levels,

$$\psi_{m<l}(A_{m,j}[0,t]) \ge 0 \tag{12}$$

since packets that arrived previously and at higher priority level will be serviced before the marked packet. Moreover, higher priority packets may contribute to the delay of the marked packet, even if they arrive *after* the marked packet by up to d_l seconds, i.e.,

$$\psi_{m<l}(A_{m,j}[t,t+d_l]) \ge 0. \tag{13}$$

The distribution of this total number of higher-priority packets that may be served ahead of the marked packet is bounded by

$$\mathcal{P}_{m<l}(t) = \sum_{m=1}^{l-1}\sum_{j\in\mathcal{M}} A_{m,j}[0,t+d_l] \le_{st} \sum_{m=1}^{l-1}\sum_{j\in\mathcal{M}} B_{m,j}(t+d_l). \tag{14}$$

Further, using the definition of SP,

$$\psi_{m>l}(A_{m,j}[0,u]) = 0 \tag{15}$$

for all $u > 0$, ignoring the possible packet transmission delay from a lower priority packet already in service when the marked packet arrives.

Next, observe that the number of packets having priority over the marked packet consists of the aggregate arrivals summed over a large number of independent traffic flows. Consequently, we approximate its distribution as Gaussian since the central limit theorem states that if the random variable Y_i has mean μ_i and variance σ_i^2, and the Y_i's are mutually independent, then the distribution of the sum $Y_1 + Y_2 + \cdots Y_n$ converges to a Gaussian distribution with mean $\sum_{i=1}^{n} \mu_i$ and variance $\sum_{i=1}^{n} \sigma_i^2$. The theorem can be shown to hold under fairly general conditions on the Y_i's, including for example, the Lindeberg sufficient condition.

Finally, Equation (6) follows by observing that up to Ct bits may be serviced by time t and by considering all values of t up to the maximal busy period of Equation (9), which is the maximum busy period of any work conserving service discipline and therefore of SP [5]. □

Our approach therefore applies the central limit theorem over multiple interval lengths in order to approximate the delay-bound violation probability in static priority schedulers. We explore the validity of the Gaussian assumption in Section 4 for loss probabilities as small as 10^{-6}. Further rationale for such techniques can also be found in [3]. While smaller loss probabilities could be studied using large deviations refinements to the above scheme, it would require increased computational complexity and more detailed traffic models that specify higher-moment envelopes. Moreover, significantly smaller loss probabilities may not be meaningful to many multimedia applications; for example, a loss probability of 10^{-9} corresponds to an average of one ATM cell lost approximately every 100 hours for a 1 Mbps flow.

3 Implications for Traffic Models Admission Control

In this section, we consider three practical issues of the rate-variance resource allocation scheme in the previous section: (1) the manner in which network clients express their traffic characterization to the network for on-line admission control tests, (2) calculation of the maximum busy period, and (3) resource allocation across multiple network nodes.

3.1 Specifying a Rate-Variance Envelope

One application of the above resource allocation scheme is in network planning and design such as buffer sizing for switches and routers. In that case, the rate-variance envelopes of the expected workloads can be used directly in the calculation of the terms in Equation (6).

For an on-line resource allocation system such as an admission control algorithm, or for an adaptive resource management algorithm, Equation (6) has the advantage that it is a simple computation that does not require, for example, convolutions or large matrix computations. However, in order to integrate the scheme with a network signaling protocol, flows must be able to specify their traffic characteristics more concisely than the entire $RV(t)$ curve. Several possibilities are described below.

First, the $RV(t)$ characterization can be *inferred* or bounded based on the flows' specified parameters, which may be worst case parameters such as the multi-level leaky bucket $(\vec{\sigma}, \vec{\rho})$ model [27] or the D-BIND model's rate-interval pairs [12]. For example, in [10, 11], a scheme is presented to upper bound a flow's rate-variance envelope based on its worst-case parameters as given by the D-BIND model. Such a scheme has the advantage that the *stochastic* properties of the flow can then be policed by using the appropriate deterministic filters at the network edge.

Second, if a specific shape of the $RV(t)$ rate-variance curve is assumed, then the curve can be specified to the network concisely by specifying several points on the curve and interpolating. For example, consider Figure 3 which depicts empirical rate-variance curves for various traces of compressed video.

The figure depicts $RV(t)$ as given by Equation (4) on a log-log scale, just as in Figure 1 for the MPEG trace of the action movie. All of the traces in the figure are of MPEG-compressed video except for the one labeled "JPEG" which uses only intraframe compression in a manner similar to motion JPEG. The Star Wars traces are from Bellcore [8], the action movie (Bond) and news traces are from the University of Würzburg [24], and the advertisements and lecture traces are from Berkeley [12]. Further details of the traces may be found in the respective references. Based on the shape of the curves in Figure 3, we observe that using two or three piece-wise linear segments to bound or approximate the $RV(t)$ curve (on a log-log scale) would closely bound the rate-variance envelopes of these flows.

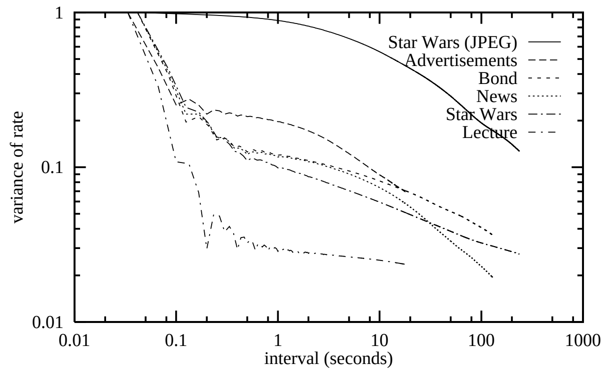

Figure 3: Variance of Rate over Multiple Interval Lengths for Video Traces

Lastly, we note that the slope at which the rate's variance decreases with interval length, as depicted on Figure 3's log-log scale, has implications on the time-scales of the source's correlation structure and whether or not the flow exhibits long-range dependence.

For example, as noted in [8] and [18], an uncorrelated process with $EX_m X_{m+n} = 0$ for $n \neq 0$ has $RV(n) = Var(X)/n$. As well, $RV(n)$ of a short-range-dependent process asymptotically decreases as $1/n$. However for long range dependent processes, $RV(n)$ asymptotically decreases as $n^{-k}Var(X)$ for $0 < k < 1$. In other words, on a log-log scale, a long-range dependent process has an $RV(n)$ curve with a slope greater than -1. While none of the curves in Figure 3 behave as strongly in this manner as the Ethernet trace of [18], whether or not long-range dependence is present can impact the parameters that one chooses to concisely describe $RV(n)$.

Our point here is not to speculate on the existence of long-range dependence in these traces, but rather to show the relationship of traffic envelopes to this topic and to point out that typical shapes of rate-variance envelopes for realistic multimedia applications will impact how one chooses to map $RV(t)$ to a concise traffic specification.

3.2 Calculation of the Busy Period Bound

Equation (6) requires the calculation of the worst-case busy period β. While this might be impossible to obtain for most traditional stochastic models, it is quite easy to obtain when traffic flows are policed at the network edge.

When the network polices a traffic flow, it ensures that the flow conforms to some client-specified parameters of a deterministic traffic model such as the leaky bucket (σ, ρ) model, the multi-level leaky bucket $(\vec{\sigma}, \vec{\rho})$ model, or the D-BIND (R_k, I_k) model. Moreover, each deterministic (or policeable) traffic model has an associated constraint function $b(t)$ that upper bounds a traffic flow over intervals of length t [27]. For example, for the D-BIND model, the constraint function is piece-wise linear; for the (σ, ρ) model, it is $\sigma + \rho t$. Hence, if flow j is policed at the network edge, then the network has the required information to determine the flow's constraint function $b_j(t)$.

Hence, Equation (9) shows how to obtain the parameter β for sets of policed flows sharing a work conserving multiplexer. For example, consider N homogeneous (σ, ρ) sources (or leaky bucket sources) of [5]. In this case the constraint function is $b(t) = \sigma + \rho t$ for a maximum burst size of σ and an upper average rate of ρ so that the maximal busy period is $\beta = N\sigma/(l - N\rho)$.

3.3 End-to-end QoS

Here, we describe several techniques for extending our single-node result to networks of general topology. The difficulty of this problem stems from the fact that the original stochastic properties of traffic flows are distorted when the flows are multiplexed inside the network.

First, in [15] Kurose showed how a general stochastic description of a traffic flow can be translated across multiple hops using the busy period bound β. While such a transformation requires no special support from the network such as sophisticated service disciplines, its primary disadvantage is that it can result in excessively low utilization of network resources when a number of hops are traversed [30]. The reason for this is that the flows' traffic characterizations become more and more bursty with each hop traversed.

Second, in [7] Ferrari showed how delay-jitter control can be used to reconstruct a flow's original traffic pattern at each network node. Consequentially, delay-jitter control reconstructs the *stochastic* properties of the traffic flows at each node, including $RV(t)$, allowing the nodes of the network to be analyzed independently (see [30] for an example). Such a scheme is able to improve the utilization of network resources, but requires more sophisticated service disciplines inside the network (e.g., time stamping, synchronized clocks, and sorted priority queues).

Finally, in contrast to delay-jitter control which reconstructs the *exact* arrival sequence at each network node, a class of service disciplines termed rate-controlled service disciplines [29] reconstruct each flow's *deterministic* parameters at each network node. In [10], we showed how stochastic properties of traffic flows can be derived from parameters of the deterministic model. Here, we note that these *derived* stochastic properties are reconstructed at each network node when rate-controlled service disciplines are used. The advantage of this approach is that rate-controlled service disciplines are simpler to implement than delay-jitter controlled service disciplines.

4 Experimental Investigations

In this section, we evaluate the resource allocation scheme described in the previous section using 30-minute traces of MPEG-compressed video. We perform a set of experiments using trace-driven simulations considering various scenarios with different loads, QoS parameters, and so on. We compare the multiplexer's performance obtained in these trace-driven simulations with that predicted by our admission control algorithm.

4.1 Trace-driven Simulation Scenario

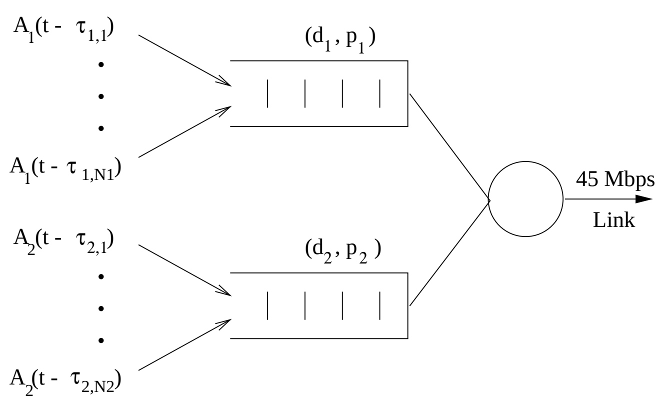

Figure 4: Trace-Driven Simulation Scenario

The trace-driven simulation scenario is depicted in Figure 4 using video traces from a newscast and an action movie as described in Section 2. For each simulation of *homogeneous* traffic flows and QoS requirements, N flows or traces are multiplexed on a simulated 45 Mbps link, with each flow's arrival pattern given by the movie trace with a start time τ_j chosen uniformly over the length of the trace (30 minutes). All flows obtain a single QoS represented by the pair (d, p) where d is the delay bound and p is the probability that a packet violates its delay bound or is dropped due to buffer overflow. The buffer size of the queue for each simulation is set to $C \cdot d$ bits. Hence, the outputs of these experiments consist of three-tuples (N, d, p). Multiple simulations are performed with independent start times and average results are reported.

We also consider *heterogeneous* traffic mixes and QoS requirements in Section 4.5. In these simulations, N_1 flows of type 1 are multiplexed with N_2 flows of type 2. In this case, type 1 flows obtain the QoS pair

(d_1, p_1) and type 2 flows (d_2, p_2). The outputs of these trace driven simulations report values of (d_i, p_i) for different (N_1, N_2) combinations. These results are then compared with the the resource allocation which determines the maximum numbers of admissible flows (N_1, N_2) such that the (d_i, p_i) QoS requirements are satisfied.

4.2 Performance Metrics

Throughout the experiments, we focus on the following two performance metrics. The first is average utilization of the link. For a given video trace that consists of F frames, we define its average rate as

$$\gamma = \frac{\sum_{i=1}^{F} x_i}{TF} \tag{16}$$

where x_i is the size of the i^{th} frame. In other words, γ is the total number of bits transmitted by the source divided by the duration of the transmission. For N multiplexed flows, the average utilization of the link is therefore

$$\frac{\sum_{n=1}^{N} \gamma_n}{C}. \tag{17}$$

In the simulations, this average utilization is also the total number of bits transmitted by all sources for the duration of a simulation, divided by the total number of bits that the server can transmit during the duration of the simulation (the link speed multiplied by the simulation time).

Our second performance metric is the total fraction of packets that either violate their delay bound or are dropped due to buffer overflow. We set the buffer size to be equal to the delay bound multiplied by the link speed, and drop packets that arrive to a full buffer. Note that if the buffer size was larger, these packets would violate their delay bounds rather than being dropped. We consider a range of delay bounds d and corresponding buffer sizes, and report the empirical $Prob\{Delay > d\}$ or $Prob\{loss\}$ as the measured fraction of packets that are dropped due to buffer overflow. We refer to this probability as p.

4.3 Admissible Region

For a given QoS (d, p) pair, we calculate the maximum number of admissible flows allowed by our resource allocation algorithm in Section 2. Specifically, we first determine the $RV(t)$ characterization for the video trace as in Equation (4) and depicted in Figures 1 and 3. Next, we determine the worst-case busy period β for the set of flows by first calculating deterministic parameters for the trace as in [12]. As described in Section 3.2, these are the parameters of the policing elements if the flow's parameters are to be enforced [10]. After parameterizing the flow with four rate-interval pairs, we determine β using Equation (9). Finally, with knowledge of the $RV(t)$ characterization and β, we calculate the maximum number of admissible flows for the QoS (d, p) pair using Equation (6).

To compare our approach with previous resource allocation schemes, we also perform admission control experiments using effective bandwidth admission control tests of [4]. An effective bandwidth scheme as in [4, 9] reserves a bandwidth for each flow according to its stochastic properties as well as the required loss probability p. Once the effective bandwidth of a flow is determined, which we denote by $E_j(p)$, the admission control test checks that $\sum_j E_j(p) < C$, where C is the link capacity. Such techniques rely on either large buffer asymptotics or large N asymptotics, and are typically based on large deviations theory or eigenvalue decomposition of Markovian flows. The admission control test that we consider here is based on large buffer asymptotics using large deviations theory [4]. In computing the admissible region, we use the same second moment statistics of the source as for the rate-variance approach.

Figure 5 shows a plot of the average utilization of the multiplexer as a function of the delay bound. Specifically, the vertical axis depicts average utilization which is directly proportional to the number of flows N as given by Equation (17).[1] This N is the maximum number of flows that can be multiplexed such that all flows obtain their required QoS. The QoS is depicted on the horizontal axis with the guaranteed delay bound d. As described above, there is also a corresponding loss or delay-bound violation probability p. Figure 5(a) depicts the case of $p = 10^{-3}$ and Figure 5(b) depicts the case of $p = 10^{-6}$.

[1] Since this trace's average rate is 583 kbps, the average utilization is $N \cdot 583$ kbps / 45 Mbps.

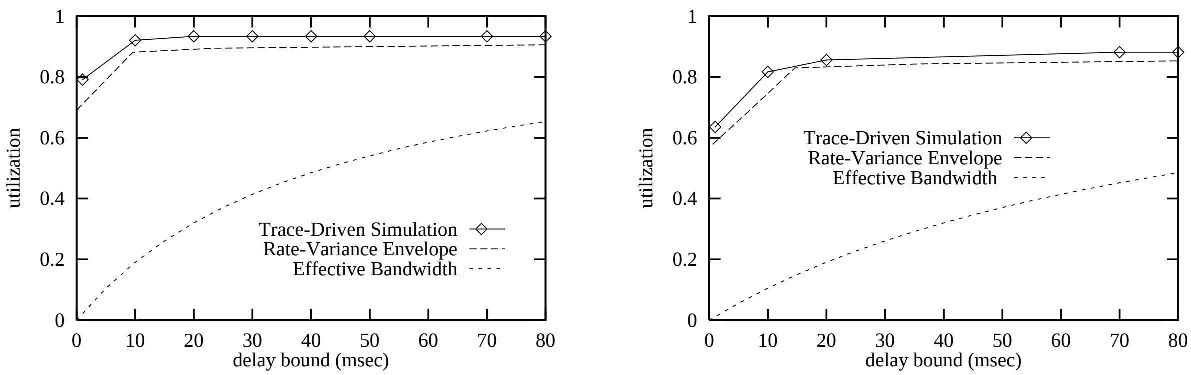

(a) $p = 10^{-3}$ (b) $p = 10^{-6}$

Figure 5: Utilization vs. Delay Bound

In both Figures 5(a) and 5(b), three curves are depicted: the trace-driven simulation, the rate-variance resource allocation scheme, and the effective bandwidth resource allocation scheme. We first focus on the two upper curves of Figure 5. For the trace-driven simulation, a point on the curve indicates the maximum number of flows that could be multiplexed so that for a given delay bound d, at most a fraction of 10^{-3} packets violated that delay bound (or a fraction 10^{-6} for Figure 5(b)). Alternatively, the rate-variance curve uses the variances-over-intervals traffic characterization, together with the resource allocation scheme of Section 2, to make an *a priori* determination of how many flows can be multiplexed such that all flows obtain the desired QoS. If Section 2's resource allocation scheme is to be used for capacity allocation, then one would desire that the *Rate-Variance Envelope* curve be as close as possible to, but not greater than, the *Trace-Driven Simulation* curve. Indeed, if the resource allocation scheme allows for more flows than can actually be supported, then violations of the promised QoS will occur, an undesirable situation for a guaranteed-services network. Hence, Figure 5 shows that our scheme is able to achieve *most* of the achievable statistical multiplexing gain, coming quite close to the results of the trace-driven simulation. Moreover, as desired, the scheme errs slightly on the conservative side rather than over-committing resources.

In contrast, we note that the effective bandwidth test unnecessarily rejects flows and under-utilizes network resources. For example, in Figure 5(b), for a 30 msec delay bound with a delay-violation probability of 10^{-6}, the simulation curve shows that 66 flows can be multiplexed for an average utilization of 86%. However, for this same QoS, the effective bandwidth test will only admit 20 flows, blocking the remainder and restricting the network utilization to 26%.

Next, we take note of the general shapes of the *Trace-Driven Simulation* and *Rate-Variance Envelope* curves of Figure 5. Since the delay bound depicted on the horizontal axis corresponds to a buffer size Cd, we can see the impact of increasing the buffer size at the network nodes. Both Figures 5(a) and 5(b) show a considerable increase in utilization for minor increases in buffer size for delay bounds in the range of up to 10 msec for Figure 5(a) and 20 msec for Figure 5(b) (10 msec of buffering corresponds to roughly 1000 ATM cells on a 45 Mbps link). However, after the respective delay bounds of 10 and 20 msec, the curves flatten considerably, indicating that increasing the buffer size further is of relatively little use. While the flattening of this curve is most likely due to the longer time-scale characteristics of the traffic flows, we note here that our variance-based resource allocation is able to follow the knee of this curve and to approximate the admissible region quite closely.

Finally, regarding the relevant time scales of these sources for resource allocation, the maximizing t of Equation (6) in these experiments is on the order of several seconds, approximately two orders of magnitude larger than the buffer size. We note that this interval length, which can be referred to as the system's *critical time scale*, is of most importance for admission control so that the traffic characteristics (mean and rate variance) must be accurately specified for this interval length.

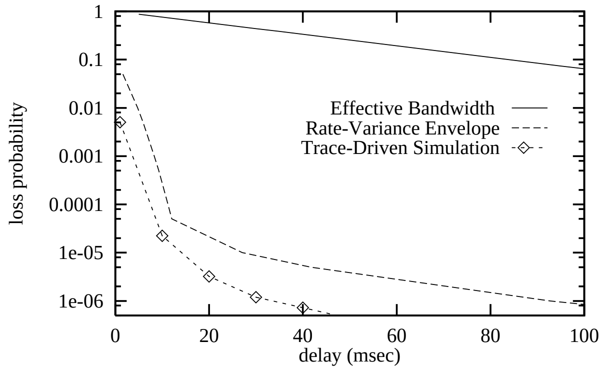

(a) 85.6% utilization

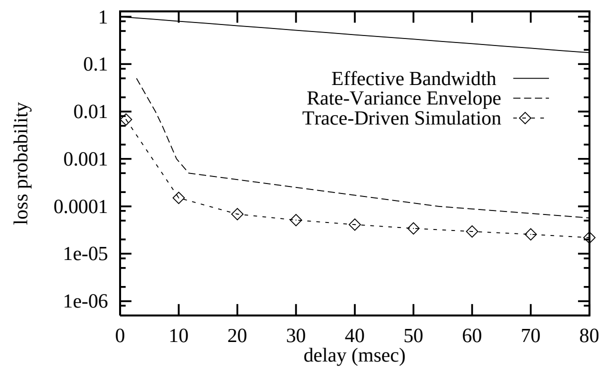

(b) 88.2% utilization

Figure 6: Probability of Loss vs. Buffer Size

4.4 Loss Probability and Buffer Size

Figure 6 further describes the results of the experiments by plotting the loss or delay-bound violation probability p vs. delay bound d for a fixed number of flows N and hence a fixed utilization. Figure 6(a) depicts the case of 66 multiplexed flows for an average utilization of 85.6%, and Figure 6(b) depicts the case of 68 multiplexed flows for an average utilization of 88.2%.

From Figure 6, we take note of the shapes of the loss probability as a function of delay bound, d, or buffer space, Cd. For the simulation curves, Figure 6(a) shows a sub-exponential relationship (note that an exponential relationship would be linear on the figure's semi-log scale). We make the following two observations. First, an effective bandwidth scheme approximates the delay-bound violation probability by a single exponential $p \approx e^{-\eta d}$. As shown by the figure, the effective bandwidth scheme considerably over-estimates this probability and does not capture its non-exponential behavior.

Second, we note that our resource allocation scheme based on application of the central limit theorem over intervals tracks this non-exponential loss-behavior quite well. Indeed, as shown in Figure 5, the scheme is able to admit *most* of the allowable flows, and as shown in both Figures 5 and 6, it is able to track the complex behavior of the multiplexer's performance across a wide range of loads and QoS parameters, including behavior expected from long-range dependent sources.

4.5 Heterogeneous Sources and QoS Requirements

To investigate the case of heterogeneous traffic mixes with different QoS requirements, we multiplex N_1 news traces with N_2 movie traces (both traces are from [24]). For the static priority scheduler, the news traces are serviced at the higher priority level with a delay bound of 10 msec and a delay-bound-violation probability of 10^{-6}. The movie traces are serviced at the lower priority level with a delay bound of 50 msec and a delay-bound-violation probability of 10^{-3}.

For the above QoS requirements, Figure 7 shows two curves. The curve labeled "Trace-Driven Simulation" depicts the actual admissible region, or the maximum (N_1, N_2) combinations such that the respective QoS requirements are met. The curve labeled "Rate Variance Envelope" depicts the maximum number of flows that the resource allocation scheme of Equation (6) will allow based on the flows' second moment statistics and their QoS requirements. As was the case in the homogeneous experiments, the figure shows that our approach extracts most of the statistical multiplexing gain, admitting nearly all of the flows that can be multiplexed for the required QoS constraints.

5 Conclusion

In this paper, we introduced a new scheme for allocating resources to multimedia traffic flows. Our approach uses simple computations on the traffic flows' rate-variance envelopes to estimate QoS parameters such as

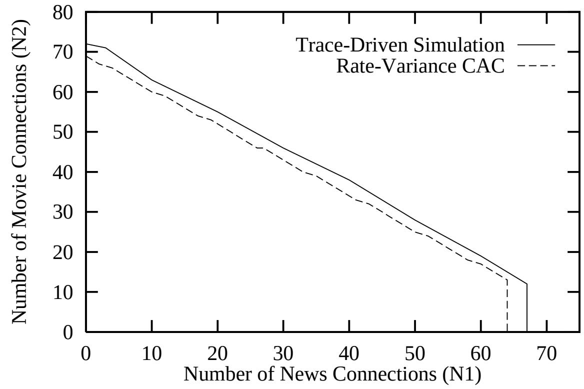

Figure 7: Admissible Region for Heterogeneous Traffic and SP Scheduler

delay-bound-violation probability for heterogeneous and bursty traffic flows. Our experiments with traces of MPEG-compressed video and trace-driven simulations indicate that our scheme is accurate enough to capture most of the achievable statistical multiplexing gain over a wide range of utilizations, buffer sizes, and loss probabilities, achieving typical average utilizations in the range of 60% to 90%.

Our results have implications for network design and planning as well as on-line capacity allocation in admission control algorithms.

References

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

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

[3] J. Choe and N. Shroff. A central limit theorem based approach to analyze queue behavior in ATM networks. *IEEE/ACM Transactions on Networking*, 6(5):659–671, October 1998.

[4] C. Courcoubetis, G. Fouskas, and R. Weber. On the performance of an effective bandwidths formula. In *Proceedings of 14th International Teletraffic Congress*, pages 201–212, June 1994.

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

[6] R. Cruz. Quality of service guarantees in virtual circuit switched networks. *IEEE Journal on Selected Areas in Communications*, 13(6):1048–1056, August 1995.

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

[8] 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.

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

[10] 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.

[11] E. Knightly. Enforceable quality of service guarantees for bursty traffic streams. In *Proceedings of IEEE INFOCOM '98*, San Francisco, CA, March 1998.

[12] 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.

[13] M. Krunz and A. Makowski. Modeling video traffic using $M/G/\infty$ input processes: A compromise between Markovian and LRD models. *IEEE Journal on Selected Areas in Communications*, 16(5):733–748, June 1998.

[14] M. Krunz and S. Tripathi. On the characterization of VBR MPEG stream. In *Proceedings of ACM SIGMETRICS '97*, pages 192–202, Seattle, WA, June 1997.

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

[16] S. Kweon and K. Shin. Video-on-demand service using a statistical traffic envelope. Technical report, University of Michigan, Ann Arbor, MI, 1998.

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

[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] 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.

[20] A. Lombardo, G. Morabito, and G. Schembra. An accurate and treatable Markov model of MPEG-video traffic. In *Proceedings of IEEE INFOCOM '98*, San Francisco, CA, March 1998.

[21] Y. Nikolaidis and I. Akyildiz. An overview of source charcterization in ATM networks. In *Modeling and Simulation of Computer and Communication Networks: Techniques, Tools and Tutorials*, pages 123–150. Gordon and Breach Publishing Co., 1997.

[22] 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.

[23] J. Qiu and E. Knightly. Inter-class resource sharing using statistical service envelopes. In *Proceedings of IEEE INFOCOM '99*, New York, NY, March 1999.

[24] 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.

[25] P. Skelly, M. Schwartz, and S. Dixit. A histogram-based model for video traffic behavior in an ATM multiplexer. *IEEE/ACM Transactions on Networking*, 1(4):446–459, August 1993.

[26] D. Starobinski and M. Sidi. Stochastically bounded burstiness for communication networks. In *Proceedings of IEEE INFOCOM '99*, New York, NY, March 1999.

[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] O. Yaron and M. Sidi. Performance and stability of communication networks via robust exponential bounds. *IEEE/ACM Transactions on Networking*, 1(3):372–385, June 1993.

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

[31] Z. Zhang, D. Towsley, and J. Kurose. Statistical analysis of generalized processor sharing scheduling discipline. *IEEE Journal on Selected Areas in Communications*, 13(6):368–379, August 1995.