In Computer Communications Journal: Special Issue
on Multimedia Communications over the Internet.

Enforceable and Efficient Service Provisioning ⋆

Tao Wu

Nokia Research Center, Burlington, MA

Edward W. Knightly

Rice University, Houston, TX

Abstract

Provisioning resources for network services introduces the conflicting requirement 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 complimentary schemes for QoS management for deterministically policed flows. The first is *adversarial mode* resource allocation: here we *bound* the stochastic envelopes of policed flows and achieve a statistically multiplexed QoS-controlled service, even in the case that all flows are independently adversarial, i.e., when all flows are non-collusively behaving in a worst-case manner at all time scales within the constraints of their policing functions. The second scheme is *non-adversarial mode, maximum-entropy* allocation: here we determine the maximum-entropy stochastic envelopes of policed (but non-worst-case) flows. Consequently, this scheme exploits a further statistical multiplexing gain via a characterization of the "most likely" behavior of policed flows. 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 design accurate admission control tests for buffered priority schedulers. We evaluate the schemes' performance with experiments using traces of compressed video and single and dual time-scale periodic sources and show that substantial statistical multiplexing gains are achieved.

Key words: deterministic traffic model, admission control, quality of service, statistical multiplexing, traffic envelopes

⋆ This research is supported by Nokia Corporation, NSF CAREER Award ANI-9733610, NSF Grant ANI-9730104, and Texas Instruments. The authors can be reached via http://www.ece.rice.edu/networks.

1 Introduction

A key challenge for future packet networks is to efficiently multiplex bursty traffic flows while simultaneously supporting Quality of Service (QoS) objectives in terms of throughput, loss probability, and end-to-end delay. At one extreme, performance can be assured even in the worst case via deterministic service [4]. In addition to its absolute guarantee, 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 has fundamental limits in its achievable utilization [29].

To overcome the utilization limits of 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, the traffic flows' rate fluctuations and temporal correlation must be characterized. In the literature, such properties are often represented via stochastic traffic models, including Markov Modulated, Self-Similar, and others [7,13,15,18,26]. However, in a shared public network with misbehaving or malfunctioning users, provisioning resources according to such stochastic source characterizations incurs a significant risk, as the underlying assumptions of the model are inherently difficult for the network to enforce or police.

The need to address the fundamental conflicting requirement for both deterministic traffic models to isolate and police users, and statistical multiplexing to efficiently utilize network resources was perhaps first recognized in [6], and remains an important problem both for per-flow and aggregate services. In [6], a policeable traffic model similar to the now standard dual leaky bucket model is used to provide a statistical network service and exploit a statistical multiplexing gain. However, in [6] as well as later studies including [2,5], network services are considered only for single time scale flows. We will show that when traffic flows have rate variations over even *two* time scales, significant inaccuracies are encountered when applying a single time scale solution. Consequently, for traffic flows more complex than periodic on-off, new techniques are needed for enforcing network services.

In this paper, we develop a general framework for enforceable network services. Our key technique is to formulate the relationship between deterministic [4] and statistical [22] traffic envelopes in order to devise general and enforceable services applicable to any deterministic traffic model. In this way, we are able to provide efficient (high utilization) services to multiple-time scale flows, exploiting statistical resource sharing, while also enabling network service providers to police traffic arrivals and ensure that their promised services are delivered.

We develop two complementary schemes for providing network services to policed flows. First, we develop an *adversarial mode* scheme. Here, we show how *statistical* network services can be assured even if each flow independently behaves in a worst-case manner at any or all

time scales. Specifically, we show how the flow's policeable deterministic envelope yields a simple *bound* on its statistical envelope so that QoS can be assured even if all flows are (non-collusively) adversarial at the worst possible time scale.

Second, we develop a *non-adversarial mode* scheme using maximum entropy techniques. Here, our goal is to better estimate the statistical properties of more typical policed but non-worst-case flows. The maximum entropy statistical envelope can be viewed as the most likely statistical characterization of the flow given only its deterministic bounds. Thus, non-adversarial mode allocation yields a controlled mechanism for network service providers to achieve an increased statistical multiplexing gain in scenarios where flows are expected to be "random" rather than strictly worst-case adversarial.

Using this framework of traffic envelopes, we also explore the relationship between source time scales, deterministic traffic models, and the admission control algorithm's ability to accurately determine the network's true admissible region. We find that if dual time scale flows are characterized by a single time scale model (such as the standard peak rate, burst length, and average rate model), network clients must either ignore their long time scale characteristics and over-state their mean rate, or ignore their short time scale rate variations and over-state their burst length and hence the impact of their temporal correlation structure on network buffers. We quantify the impact of such traffic mischaracterizations on admission control by examining subsequent errors in the computed admissible region that an inaccurate traffic model causes. We show that even with an ideal mapping of a dual time scale source to a single time scale model, significant errors in the admissible region can occur. Furthermore, such errors necessarily *under*-estimate the true admissible region as the traffic parameters must be *over*-stated: under-statement of traffic parameters would result in traffic being blocked by the network's policing elements.

Finally, we study the schemes' performance using trace-driven-simulation experiments as well as simulations with single and dual time scale periodic flows. As an illustrative example, we find through simulations that 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 scheme utilizes resources to 41%, necessarily lower than that of the trace-driven simulation since the approach assumes that each flow is independently adversarial, which is not the case for these video flows. However, this represents a significant improvement over a single time scale approach which obtains 14% utilization in this case. Alternatively, the non-adversarial-mode admission control scheme achieves an average link utilization of 79%. Indeed, with non-adversarial-mode allocation, we find that once traffic flows are aggregated and economies-of-scale are present, the maximum entropy mapping from deterministic to stochastic envelopes can lead to a considerably accurate admission control test.

In addition to the aforementioned work, our approach is related to several other schemes for resource provisioning for policed flows. In [8], we first studied statistical services for multiple time scale policed flows using the D-BIND model [12]. In that work as here, the traffic model is policeable, yet accurate enough to capture key properties more typically associated with

statistical models such as autocorrelation structure. However, [8] requires the on-line solution of an optimization problem in order to find a flow's worst-case statistical traffic envelope. Here, extending [10], we develop a simpler and more direct approach.

Further work using quite different approaches is found in [21,19,23,25], where multi-level leaky buckets (a special case of the D-BIND model) are employed to characterize the multiple time scale nature of flows. These schemes are developed in the context of smoothing and bufferless multiplexing which has the advantage of simplifying multiple node issues (as studied in [24]) as traffic traverses bufferless multiplexers undistorted. Comparatively, while our approach does have utilization advantages from exploiting buffering rather smoothing [30], the key difference lies in the simplicity and generality of the envelope-based approach, viz., it applies to any deterministic traffic model and, employing [22], to a broad class of traffic schedulers including weighted fair queueing, static priority, and earliest deadline first. Finally, our *non-adversarial* scheme is unique in its control of network services to policed but non-worst case flows.

The remainder of this paper is organized as follows. In Section 2, we describe the important aspects of deterministic traffic models for provisioning network resources. In Sections 3 and 4, we present the scheme for extracting stochastic envelopes of traffic flows from their enforceable parameters, which we apply to admission control in Section 5. We investigate the impact of the source time scales 6 and evaluate the scheme experimentally using video traces in Section 7.

2 Background on Traffic Envelopes

As described above, strict enforcement of network services requires that a flow's traffic be specified and policed via a deterministic traffic model which upper bounds its 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 can be transmitted over any interval of length t. Denoting $A_j[s, s+t]$ as the number of flow j arrivals in the interval $[s, s+t]$, 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 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 [12] to better characterize the burstiness properties of realistic traffic flows. With the D-BIND model, sources charac-

terize 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 $I_0 = 0$. In [12], we showed how this source characterization captures a flow's burstiness properties and temporal correlation structure, even over long time scales. For example, with an MPEG-compressed video source, the flow'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 [29], 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 (a user or traffic class) 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. As illustrated in Figure 1, regardless of the flow'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.

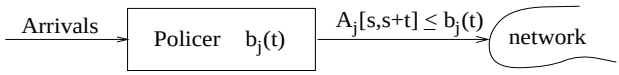

Fig. 1. Policing of the Traffic Constraint Function $b_j(t)$

Figure 2 illustrates deterministic traffic modeling with an example D-BIND source characterization. The figure depicts the bounding or worst-case rate versus interval length so that for small interval lengths, the bounding rate approaches what is commonly called the "peak-rate," and for long interval lengths, it approaches the source's long-term average rate.

Observe that the worst-case arrivals over different interval lengths characterize a flow's burstiness over different time scales in a manner analogous to the variance-time plot of [16] which describes the second moment of the arrivals over different interval lengths: both models describe the traffic in terms of dispersions from the mean rate as a function of interval length. Indeed, observe from Figure 3 that the statistically described time scales of the source char-

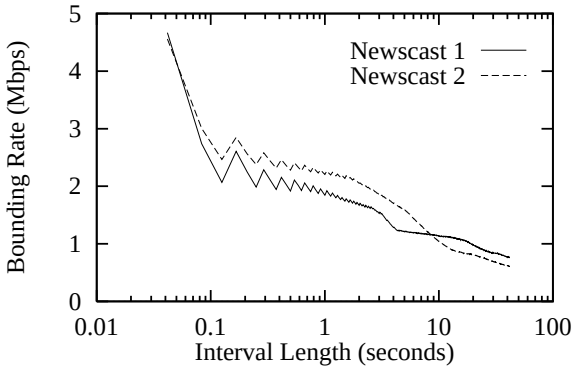

Fig. 2. Deterministic Source Characterizations for MPEG Traces

acterized by the rate-variance envelope are also evident in the deterministic parameters of Figure 2.

Regardless, the ultimate effectiveness of a deterministic model in describing the important properties of traffic flows is best determined by its effectiveness in resource allocation. We evaluate this experimentally in Section 7 with trace-driven-simulation experiments using video traces that exhibit multiple time scale behavior.

Finally, in the sections following we describe schemes for bounding and approximating a flow's stochastic envelope given its deterministic envelope. In general, the random variables $B_j(t)$ are a stochastic envelope of flow j if [14]

$$P(A_j[s, s+t] > x) \leq P(B_j(t) > x) \tag{4}$$

for all s, t, and x. In this paper, we consider a rate-variance envelope $RV_j(t)$ which describes the variance of a flow's arrival rate over intervals of length t as [9]

$$RV_j(t) = Var\left(\frac{A_j[s, s+t]}{t}\right). \tag{5}$$

While our approach is easily generalizable to higher-moment envelopes, this second moment characterization has computational advantages in admission control while maintaining a high degree of accuracy [11].

3 Adversarial Mode Allocation

In this section, we describe a technique for enforcing network services even in the case that all flows are (non-collusively) adversarial. We consider a general deterministic and statistical traffic envelope so that our solution applies to any deterministic traffic model and a broad class of service disciplines via the theory of statistical service envelopes [22].

Our key technique is to bound the stochastic properties of the flow at each time scale t using properties of the deterministic envelope and hence the policer parameters. For policed flows, the rate-variance envelope is bounded as follows.

Adversarial Envelope: *If flow 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:*

$$RV_j^*(t) = \frac{\phi_j b_j(t)}{t} - \phi_j^2 \tag{6}$$

where ϕ_j is defined as:[1]

$$\phi_j = \lim_{t \to \infty} \frac{b_j(t)}{t}. \tag{7}$$

The adversarial bound can be shown as follows. 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,

$$a_j(t) = \int_s^{s+t} r_j(s)\, ds \tag{8}$$

which depends only on t for stationary sources.

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

$$\int_s^{s+t} r_j(s)\, ds \leq b_j(t) \;\; \forall s, t \geq 0 \tag{9}$$

is given by Equation (6) and is attained when the distribution of $a_j(t)$ is given by

$$f_{t,j}^*(x) = \left(\frac{b_j(t) - \phi_j t}{b_j(t)} \right) \delta(x) + \frac{\phi_j t}{b_j(t)}\, \delta(x - b_j(t)) \tag{10}$$

such that for an interval length t, Equation (10) describes a binomial distribution.

According to (9), $f_{t,j}(x) = 0$ for $x > b(t)$ and $x < 0$ so that the rate-variance envelope of a policed flow is given by

[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.

$$RV_j(t) = \frac{Ea_j(t)^2 - (Ea_j(t))^2}{t^2} \tag{11}$$
$$= \frac{1}{t^2} \int_0^{b_j(t)} x^2 dF_{t,j}(x) - \frac{1}{t^2}\left(\int_0^{b_j(t)} x dF_{t,j}(x)\right)^2$$

for some distribution $f_{t,j}(x)$ satisfying (9). For the distribution $f^*_{t,j}(x)$ of Equation (10), $RV^*_j(t)$ is given by Equation (6). To show that $RV^*_j(t) \geq RV_j(t)$ for all t and for all distributions $f_{t,j}(x)$ satisfying (9), observe that

$$RV^*_j(t) \; - \; RV_j(t) \quad = \quad \frac{b_j(t)\phi_j}{t} - \frac{1}{t^2} \int_0^{b_j(t)} x^2 dF_{t,j}(x)$$

$$= \frac{b_j(t)}{t^2} \int_0^{b_j(t)} x \; dF_{t,j}(x) - \frac{1}{t^2} \int_0^{b_j(t)} x^2 dF_{t,j}(x) \tag{12}$$

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 (12),

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

which is clearly non-negative.

Notice that the bound applies to any deterministic traffic model since each deterministic traffic model parameterizes a constraint function $b_j(t)$ as described in Section 2; the more accurately the model characterizes the traffic flow, 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 $t \leq T$. 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 [20] 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 flow'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.

Thus, together with an envelope-based admission control algorithm, the enforceable $RV_j^*(t)$ characterization provides a mechanism for ensuring that network services which extract a statistical multiplexing gain can be provisioned and policed, even if all sources are independently adversarial, i.e., if sources are adversarial, but not collusive.

4 Non-Adversarial Mode Using Maximum Entropy

The adversarial rate-variance bound is achieved only when a flow transmits in an on-off mode at all time scales. Indeed, while it is possible for a flow to realize $RV^*(t)$ for a particular t (as in the above example), it is often not possible to simultaneously realize this bound for all t. Moreover, one may expect that in a large scale network servicing many flows, only a small fraction of the flows will be truly adversarial. Consequently, in this section we describe a scheme for *approximating* the statistical envelopes of more typical flows, again using their policed parameters. Here, the ultimate goal is to increase the network's utilization for non-adversarial flows with enforced traffic constraints.

In this scenario, the only information available is the deterministic envelope (and the mean rate derived from it). Thus, the problem is to approximate $RV_j(t)$ or more generally the distribution of $B_j(t)$ given $b_j(t)$ and ϕ_j (the maximum and mean). As this is an under-determined problem with infinitely many solutions, adversarial mode allocation can be viewed as the worst case solution with the additional constraint of bounding $RV(t)$.

Here, we propose maximizing entropy of the probability density function to approximate the rate variance. As a measure of uncertainty or randomness, entropy is defined as [3,28]

$$h(f) = -\int_S f(x) \ln f(x) dx \tag{13}$$

for a continuous random variable with probability density function $f(x)$. The maximum entropy principle states that among the many possible distributions satisfying the known constraints, we should choose the one that maximizes the entropy. The rationale for doing this is to assume the least about the distribution, and choose the distribution that is most uncertain given the available information. In other words, the maximum entropy distribution is the one which is maximally noncommittal regarding missing information.

We now apply the maximum entropy principle to approximate the distribution (and therefore second moment statistics) based on the peak-rate envelope and mean rate. The only information known about the distribution is its range (from 0 to rate envelope $v_j(t) = \frac{b_j(t)}{t}$) and mean ϕ_j. Based on the maximum entropy principle, we have the following result.

Non-adversarial Maximum Entropy Envelope: *Given a flow j's deterministic traffic rate envelope $v_j(t)$ and traffic mean rate ϕ_j, the maximum entropy estimate of rate variance $RV_j(t)$ is*

$$\widehat{RV_j}(t) = \frac{A_j}{\lambda_{j,1}^3}[((\lambda_{j,1}v_j(t) - 1)^2 + 1)e^{\lambda_{j,1}v_j(t)} - 2] - \phi_j^2 \tag{14}$$

where $\lambda_{j,1}$ is the non-zero[2] *solution of*

$$e^{\lambda_{j,1}v_j(t)} + \frac{1 + \phi_j\lambda_{j,1}}{(v_j(t) - \phi_j)\lambda_{j,1} - 1} = 0 \tag{15}$$

and A_j is

$$A_j = \frac{\lambda_{j,1}}{e^{\lambda_{j,1}v_j(t)} - 1}. \tag{16}$$

This envelope is derived as follows. We would like to maximize Equation (13) subject to the constraints

$$\int_0^{v_j(t)} f_j(x)dx = 1 \tag{17}$$

and

$$\int_0^{v_j(t)} xf_j(x)dx = \phi_j. \tag{18}$$

By using the Lagrange multiplier method, the distribution that maximizes the entropy is of the form

$$f_j(x) = \begin{cases} e^{\lambda_{j,0}+\lambda_{j,1}x} & 0 \le x \le v_j(t) \\ 0 & \text{otherwise} \end{cases} \tag{19}$$

where $\lambda_{j,0}$ and $\lambda_{j,1}$ are coefficients that must satisfy Equations (17) and (18). Let $A = e^{\lambda_{j,0}}$, and we have

$$\int_0^{v_j(t)} e^{\lambda_{j,1}x}dx = \frac{1}{A_j} \tag{20}$$

[2] If $v_j(t) = 2\phi_j$, then $\lambda_{j,1} = 0$

and

$$\int_0^{v_j(t)} xe^{\lambda_{j,1}x}dx = \frac{\phi_j}{A_j} \tag{21}$$

And these two equations are algebraically equivalent to Equations (15) and (16).

Now let A_j and $\lambda_{j,1}$ be given by Equations (15) and (16), then by definition

$$\widehat{RV}_j(t) = \int_0^{v_j(t)} x^2 e^{\lambda_{j,0}+\lambda_{j,1}x}dx - \phi_j^2. \tag{22}$$

Thus, one can easily show that Equations (22) and (14) are equivalent.

Notice that Equation (15) does not have an analytical solution in general, but we can obtain a numerical solution by the following method. First, observe the following relationships about $\lambda_{j,1}$:

$$\lambda_{j,1} < 0 \text{ if } v_j(t) > 2\phi_j$$

$$\lambda_{j,1} = 0 \text{ if } v_j(t) = 2\phi_j$$

$$\lambda_{j,1} = -\frac{1}{\phi_j} \text{ if } v_j(t) = \infty.$$

Assuming $v_j(t) > 2\phi_j$, Equation (15) can be rewritten as $g(\lambda_{j,1}) = 0$ where

$$g(\lambda_{j,1}) = e^{\lambda_{j,1}v_j(t)} + \frac{1+\phi_j\lambda_{j,1}}{(v_j(t)-\phi_j)\lambda_{j,1}-1}.$$

Furthermore, $g(x) < 0$ if $x > \lambda_{j,1}$ and $g(x) > 0$ if $x < \lambda_{j,1}$. Thus standard numerical techniques can be used to compute $\lambda_{j,1}$ efficiently. Table 1 shows the rate variances using adversarial mode and maximum entropy for mean rate $r = 1$ and different peak rates p. This table shows that adversarial mode gives an unbounded rate variance proportional to the peak rate. Consequently, mis-specifying a flow's peak rate will yield inaccurate admission control decisions. On the other hand, the maximum entropy approximation gives a bounded rate variance and provides a more "realistic" estimation when the peak rate is unknown (i.e., infinity).

p	1	2	4	8	16	32	64	∞
Adversarial Mode Variance	0	1	3	7	15	31	63	∞
Maximum Entropy Variance	0	0.333	0.778	0.988	1	1	1	1

Table 1
Rate Variances Using Adversarial Mode and Maximum Entropy

5 Integrated Admission Control for Policed Flows

Here, we develop an integrated framework for multi-service admission control for policed flows. We consider deterministic (or guaranteed) service together with statistically multiplexed services via our use of a single framework of deterministic traffic envelopes to describe traffic, regardless of the required service. In this way, we not only enforce all services, but also incorporate the impact of flows obtaining a deterministic service in the calculation of the loss probabilities for flows obtaining a statistical service. Consequently, a further statistical multiplexing gain can be obtained by flows employing statistical services due to unused capacity of deterministic flows.

5.1 Multi-Class Admission Condition

In the same way that $b_j(t)$ and $B_j(t)$ provide a general way of describing a traffic flow's *arrivals* as a function of interval length, the available *service* to a flow can be lower bounded [4] and statistically described [22] by the deterministic and statistical service envelopes $s_j(t)$ and $S_j(t)$.

Denoting superscript i as the *aggregate* arrivals or service of a class, in [22] we showed that class i with service envelope $S^i(t)$ and traffic envelope $B^i(t)$ has a delay-bound violation probability

$$P[D^i > d] \leq P[\max_{t \geq 0}\{B^i(t) - S^i(t+d)\} > 0]. \tag{23}$$

Using this general result, multiclass admission control tests can be described succinctly for a broad class of schedulers by specifying their arrival and service envelopes which are functions of the policing parameters and service discipline.

We pursue the particular example of static priority scheduling: consider an SP scheduler with N priority queues, link speed C, and the aggregate traffic in class i bounded by $B^i(t)$ and $b^i(t)$, with $i = 1, \ldots, N$ denoting the priority level from higher priority to lower priority. In this case, the available service to a flow at priority level i over intervals of length t is the available capacity less all traffic arriving at higher priority levels, i.e., the statistical service

envelope for class i is

$$S^i(t) = (Ct - \sum_{j=1}^{i-1} B^j(t))^+ \tag{24}$$

and the deterministic service envelope for class i is

$$s^i(t) = (Ct - \sum_{j=1}^{i-1} b^j(t))^+ \tag{25}$$

where $b^i(t) = \sum_{j \in C_i} b_j(t), B^i(t) = \sum_{j \in C_i} B_j(t)$, and $b_j(t)$ and $B_j(t)$ are the statistical and deterministic envelopes of the jth flow in class i, and C_i denotes the set of flows in class i.

Thus at the highest priority levels, delay bound d^i is deterministically guaranteed if [17]

$$\max_t \{b^i(t) + \sum_{k=1}^{i-1} b^k(t + d^i) - C(t + d^i)\} \leq 0 \tag{26}$$

Similarly, for statistically multiplexed services with delay bound d^i is guaranteed with probability $P^i > 0$ if

$$P[\max_t \{B^i(t) + \sum_{k=1}^{i-1} B^k(t + d^i) - C(t + d^i)\} > 0] \leq P^i. \tag{27}$$

since for statistical service classes, Equation (23) and Equation (24) indicate that

$$B^i(t) - S^i(t + d^i) \leq_{st} B^i(t) - C(t + d^i) + \sum_{k=1}^{i-1} B^k(t + d^i), \tag{28}$$

Thus, if $P[\max_t \{B^i(t) + \sum_{k=1}^{i-1} B^k(t + d^i) - C(t + d^i)\} > 0] \leq P^i$, then the statistical service in the ith service class is satisfied. Results for other service disciplines can be found in [22].

To calculate $P[\max_t \{B(t) - S(t + d)\} > 0]$ in Equation (27) we utilize the "maximum variance" approximation of [1]. Let

$$\begin{aligned} \sigma_t^2 &= \text{var}\{B(t) - S(t + d)\}, \\ \alpha_t &= \frac{0 - E\{B(t) - S(t + d)\}}{\sigma_t}, \\ \alpha &:= \inf_t \alpha_t. \end{aligned}$$

Approximating $\{B(t) - S(t+d)\}$ as Gaussian, under conditions (C1) – (C2) in [1],

$$P[\max_t\{B(t) - S(t+d)\} > 0] \geq \max_t P[B(t) - S(t+d) > 0] = \phi(\alpha) \tag{29}$$

and

$$P[\max_t\{B(t) - S(t+d)\} > 0] \leq e^{-\frac{\alpha^2}{2}} \tag{30}$$

where $\phi(\alpha) = \frac{1}{\sqrt{2\pi}}\int_\alpha^\infty e^{-\frac{x^2}{2}}dx$. Proof of these two bounds is given in [1], and we utilize the former in the experiments below.

Finally, 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, our estimate of this probability is also maximized. We note further that for both adversarial and non-adversarial allocation, traffic flows must be statistically independent or non-collusive. If traffic flows are collusive, then a fully deterministic approach must be employed [29].

5.2 Empirical Adversarial and Maximum Entropy Envelopes

Figure 3 illustrates the adversarial mode and maximum entropy (non-adversarial mode) rate-variance envelopes for the MPEG-compressed video trace described in Section 7. 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" and "Maximum Entropy" 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 [12]. 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 Equation (6) $\widehat{RV}_j(t)$ using Equation (14).

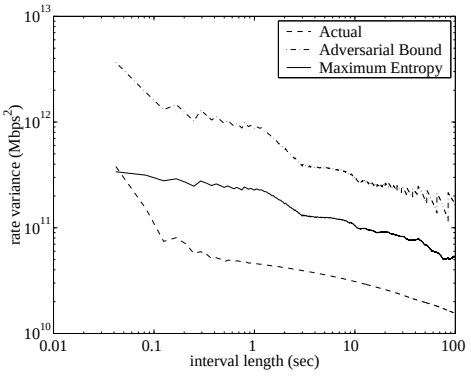

Fig. 3. Envelopes from Video Trace

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 $\widehat{RV}(t)$ better approximates the flow's actual envelope $RV(t)$, especially at smaller time scales, while it is more conservative for longer interval lengths. The reason for this is that, while the maximum entropy estimation of the marginal rate distribution is quite accurate, over longer interval lengths, the traffic (when aggregated with itself over long interval lengths) becomes more constant rate, or has lower variance: the maximum entropy envelope does not incorporate this behavior and hence is conservative for large t.

5.3 Numerical Examples

We now provide a simple example illustrating admission control using adversarial and non-adversarial mode allocations. Consider a FCFS multiplexer with link capacity C and buffer space B serving N homogeneous flows with parameters $(P = 2\rho, \sigma, \rho)$. Using adversarial mode allocation, each flow's rate variance is

$$RV^*(t) = \begin{cases} \rho^2 & \text{if } 0 < t < \frac{\sigma}{2\rho} \\ \frac{\sigma}{2t}\rho & \text{if } t > \frac{\sigma}{2\rho} \end{cases} \tag{31}$$

since adversarial mode follows the Bernoulli distribution. The admission control condition in Equation (29) can then be simplified as

$$P(D > d) \approx \phi\left(\frac{C - N\rho + 2Cd\rho/\sigma}{\sqrt{N\rho}}\right) = \phi(a). \tag{32}$$

For example, if $C = 45$Mbps, $N = 60$, $\rho = 583$Kbps, $d = 0.1$s, and $\sigma = 1$Mb, then $P(D > d) \approx 3.61 \times 10^{-4}$. For non-adversarial mode, since $P = 2\rho$, the maximum entropy distribution is uniform, and the variance is thus 1/3 of adversarial distribution's. Thus the delay bound violation probability for maximum entropy approximation is

$$P(D > d) \approx \phi(\sqrt{3}a). \tag{33}$$

In the above setting, the delay bound violation probability for maximum entropy mode allocation is 2.38×10^{-9}, about five orders of magnitude less than that of adversarial mode.

6 Time Scales of Policed Flows

In this section, we explore the impact of the traffic's time scales on the effectiveness of enforceable network services. In particular, we use single and dual time-scale traffic sources to illustrate the importance of explicitly incorporating the source's multiple-time-scale nature into both the deterministic traffic model as well as the admission control algorithm. By investigating the errors introduced with a single time scale scheme, we show that even with an ideal choice of traffic parameters, a single-time scale approach can significantly underestimate the true admissible region for dual time-scale traffic.

6.1 Single Time Scale Sources

As a baseline, we first consider periodic on-off sources as depicted in Figure 4. Such a source can be characterized by three parameters such as the dual leaky bucket's peak rate, maximum burst length, and upper average rate (P, σ, ρ) with constraint function $b(t) = \min(Pt, \sigma + \rho t)$. For notational convenience, we consider a transformation of these parameters (P, I_B, I) such that I_B is the maximal duration of a burst at rate P and I is the minimum spacing of such bursts. In other words $I_B = \sigma/(P - \rho)$ and $I = I_B + \sigma/\rho$.

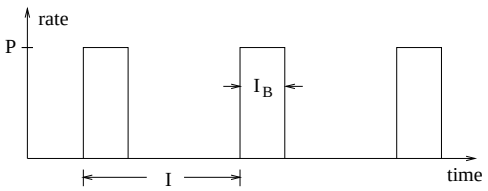

Fig. 4. Example Single Time Scale Source (P, I_B, I)

To evaluate an admission control algorithm's ability to predict the correct admissible region, we perform the following experiments.

First, we simulate a first-come first-serve multiplexer with capacity C and buffer size B which services N sources with parameters (P, I_B, I). Each source is given a random phase uniformly distributed between 0 and I such that the sources are statistically independent. For a given set of source and multiplexer parameters, we perform over 2500 simulations (each with independent source phases) and measure the empirical average loss probability P_L.

For the admission control experiments, we consider the loss probability P_L to be given as the quality of service parameter. For sources with parameters (P, I_B, I) and a multiplexer with buffer size B and link capacity C we use the scheme of Sections 3 and 5 to determine the maximum number of admissible flows N subject to the traffic and QoS constraints.

Figure 5 depicts a typical set of experimental results. In the figure, the link capacity is 45 Mbps and the sources have a peak rate of $P = 5.87$ Mbps, a burst length of I_B=.083 seconds, and a period I of 0.83 seconds. For simplicity of the discussion below, we can consider time

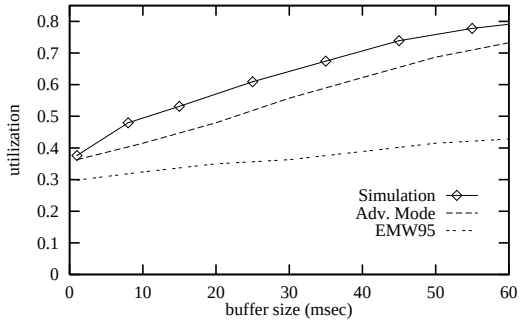

Fig. 5. Admission Control for Single Time Scale Source

to be slotted to 1/24 seconds, such that the burst length is 2 time slots and the period is 20 time slots. Thus, the source's peak rate P is 10 times its average rate ρ. The loss probability for the figure is $P_L = 10^{-3}$.

The curve in Figure 5 labeled "Simulation" depicts the maximum number of flows N (scaled to utilization as $N\rho/C$) vs. the buffer size B (scaled to delay as B/C) that could be multiplexed such that the average loss probability is less than P_L. In other words, this curve should be viewed as the true admissible region.

The curve labeled "Adv. Mode" depicts the admissible region for the adversarial mode admission control test of Sections 3 and 5, and the curve labeled "EMW95" depicts that of [5]. We make the following observations about these three curves. First, note that both analytic admission control regions are below the actual admissible regions obtained in the simulations. The reason for this is that both admission control tests employ bounds at different stages of their derivations. Next, notice that for small buffer sizes both algorithms perform well, with our proposed scheme closest to the true admissible region. This is because for a bufferless multiplexer (B approaching 0), both tests in essence calculate the probability that the aggregate rate distribution exceeds the link capacity for sources with binomial rate distributions. The discrepancy is due to our use of a Gaussian approximation for the aggregate rate distribution and [5]'s use of the Chernoff bound. Finally, we note the system's buffer scalings. As illustrated, the simulations show that the number of admissible flows increases almost linearly with increasing buffer size for delays ranging from 0 to 60 msec; over this range, the utilization doubles, increasing from 38% to nearly 80%. The buffer scalings of [5] are more conservative due to an additive approach of estimating the total buffer requirement from per source requirements. In contrast, our approach scales in a manner quite similar to the simulation results, primarily due to its use of envelope-based resource allocation [9].

6.2 Dual Time Scale Sources

Here, we consider dual time scale sources with parameters (P, I_{B1}, I_{B2}, I) as depicted in Figure 6. The sources exhibit rate-variations over two time scales in the sense that on the

slower time scale, the sources sends bursts of size $(2PI_{B1})$ bits every I seconds. However, in contrast to the single time scale source of Figure 4, each burst of duration I_{B2} consists of two separate bursts, each of duration I_{B1}.

Fig. 6. Example Dual Time Scale Source

In Figure 6, we have considered sources with peak rate $P = 5.87$ Mbps, burst lengths $I_{B1} = .083$, $I_{B2} = 0.292$, and period $I = 1.67$ seconds. Considering time to be slotted to 1/24 to seconds, I_{B1} is 2 time slots, I_{B2} 7, and I 40. The peak and average rates are therefore the same as for the single time scale example above. For the figure's "Simulation" curve, we determine the true admissible region in the same manner as for the single time scale sources.

For the admission control curve labeled "Adv. Mode", we use an approach analogous to that of the single time scale source, except that the constraint function for this dual time scale source is given by $b(t) = \int_0^t r(s)ds$ where $r(s)$ is depicted in Figure 6. Notice that such a constraint function is piece-wise linear as in the D-BIND model of [12].

Observe that to characterize a dual time scale source with a single time scale traffic model such as (P, I_B, I), either the short time scale or long time scale behavior must be ignored. To ignore the longer time scale, a dual-time-scale source with parameters (P, I_{B1}, I_{B2}, I) as in Figure 6 can be upper bounded with a single time scale model with parameters $(P, I_{B1}, I_{B2} - I_{B1})$. Note that such a characterization is necessarily conservative as the parameterized mean rate is $PI_{B1}/(I_{B2} - I_{B1})$ which is greater than the true mean rate $2PI_{B1}/I$.

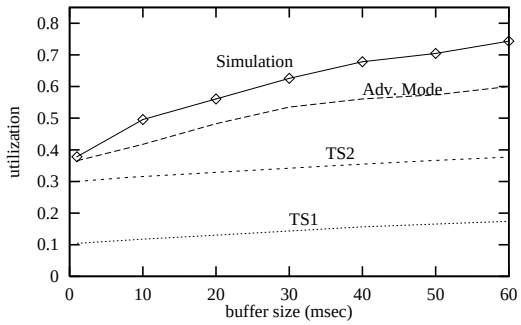

Fig. 7. Admission Control for Dual Time Scale Source

An alternative single time scale characterization of this same source of Figure 6 is given by the parameters $(P, 2I_{B1}, I)$. In this manner, the source has specified its true long term average rate but has necessarily over specified its maximum burst length.

Notice that in characterizing a dual time scale source with a single time scale model as above, the resulting traffic characterization must *upper* bound the actual traffic: otherwise,

the excess traffic will be blocked by the policing elements.

Figure 7 depicts the results of these admission control experiments. As was the case of the single time scale sources, the "Adv. Mode" admission control algorithm is somewhat conservative, but is able to approximately match the simulation's admissible region and exploit the relative benefits of buffering. From the remaining two curves, we see that in this case, it was better for the single time scale approach to ignore the source's short time scale dynamics and characterize its longer time scale behavior rather than vice versa. This is evidenced by the TS2 curve's superior performance to TS1, with the former using the $(P, 2I_{B1}, I)$ traffic parameters. However, we show in the next section that ignoring short time scales rather than long is not always beneficial.

6.3 From Modeling Errors to Admission Control Errors

Here, we explicitly consider the impact of characterizing a dual time scale traffic source with a single time scale model by determining the error in the admissible region resulting from such a bound. We calculate this error as $\epsilon = \frac{N_{STS} - N_{DTS}}{N_{DTS}}$ where N_{DTS} is the number of admissible flows calculated as in the "Adv. Mode" curve of Figure 7 using the source's true traffic parameters, and N_{STS} is the number of admissible flows using the same envelope-based admission control algorithm of Section 3 but with a single time scale traffic model (recall that these techniques apply to any deterministic traffic model). As described above, when the long time scale is ignored, the source's envelope is bounded by parameters $(P, I_{B1}, I_{B2} - I_{B1})$, and when its short time scale is ignored, it is bounded by parameters $(P, 2I_{B1}, I)$.

Figure 8 depicts the results of this experiment for a buffer size of 50 msec, a link capacity of 45 Mbps, and a loss probability of 10^{-3}. In the figure, the vertical axis depicts the error in the admissible region, ϵ, and the horizontal axis depicts the inter-burst time, $I_{B2} - 2I_{B1}$, i.e., the time between the two fast time scale bursts. In all of the experiments, I_{B1} is 2 time slots and I is 40 time slots.

First, notice that for an inter-burst time of 0 as well as an inter-burst time of 18, the source is actually a single time scale source. In other words, with an inter-burst time of 0, the source transmits a single burst for 4 time slots and then remains idle for 36 time slots. Likewise, when the inter-burst time is 18, the source's short time scale burst matches its long time scale burst such that the source is actually periodic with period 20 and burst length 2. Hence, it is not surprising that for an inter-burst time of 0, one is able to ignore the fast time scale and for an inter-burst time of 18, one is able to ignore the slow time scale.

Second, as the inter-burst time increases, the error caused by ignoring the source's fast time scale is increased. The reason for this is that as the inter-burst time increases, the source's fast time-scale bursts are more spread out and hence more flows can be admitted. However, if the fast time scale is ignored due to use of a single time scale model, this effect cannot be exploited by admission control and the same number of flows are admitted regardless of the

inter-burst time. The net effect is that the error increases.

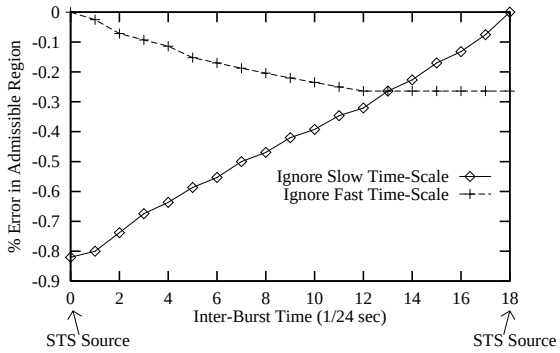

Fig. 8. Error in Admissible Region ϵ

In contrast, by ignoring the slow time scale, one has in essence exactly characterized the short term burstiness at the expense of over-stating the source's mean rate. Without the true mean rate, such an approach is quite conservative when the inter-burst time is small, or equivalently, when the average rate over I_{B2} seconds is much greater than the average rate over I seconds. However, as these two rates become closer with longer inter burst times, this approach becomes the superior of the two, as the disadvantage of mis-characterizing the mean rate is outweighed by the benefits of the source's correct bound of its maximal burst length before shutting off.

Note therefore that even with an ideal selection of traffic parameters to suit the inter-burst time, it is not possible to match the performance of a two time scale model. In this example, the maximal error under an ideal selection of traffic parameters is -26%, i.e., 26% of flows will be unnecessarily blocked for sources with an inter-burst time of 13 ($I_{B2} = 15$).

Lastly, we reiterate that the error in the admissible region due to using an overly simple traffic model is necessarily negative, as the traffic must be parameterized with an upper approximation; otherwise the source's excess traffic will be blocked by the policer.

7 Experiments with Video Traces

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 adversarial-mode and maximum-entropy resource reservation schemes, we compare the flows' performance obtained in trace-driven simulations with that predicted by the admission control tests and $RV(t)$ traffic characterizations.

7.1 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 [27].

For each simulation, N flows or traces are multiplexed on a simulated 45 Mbps first-come-first-serve link, with each flow'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 flows N and buffer size $C \cdot d$ (the link capacity times the delay bound) we measure the fraction of packets P_L that are dropped due to buffer overflow. Many simulations are performed with independent start times and average results are reported.

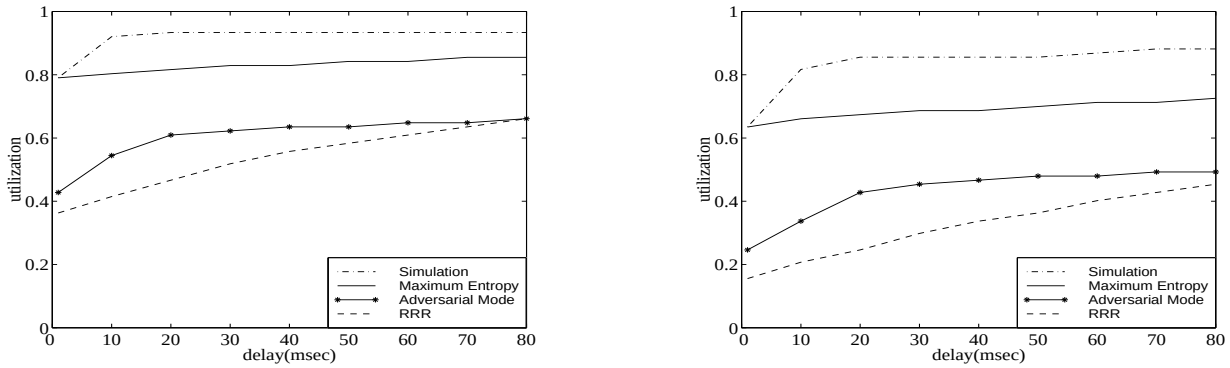

(a) $P_L = 10^{-3}$ (b) $P_L = 10^{-6}$

Fig. 9. Utilization vs. Delay Bound

In the admission control part of the experiments, we determine the flows' rate-variance envelopes from their enforceable deterministic parameters as depicted in Figure 3. We then use the admission control test of Section 5 to determine the maximum number of admissible flows, N, subject to the QoS constraints for delay, d, and loss probability, P_L.

7.2 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 [25].

Figure 9 shows the results of the trace-driven simulation and admission control experiments. [3] The figure shows the average utilization of the multiplexer (which is proportional to the

[3] 95% confidence intervals for the simulations are all within a single flow and are therefore not shown.

number of flows 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 flows N (scaled to utilization) that can be multiplexed such that all flows are guaranteed a probability of delay-bound violation or buffer overflow of 10^{-3} in Figure 9(a) and 10^{-6} in Figure 9(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* $\widehat{RV}(t)$ traffic characterization (an approximate rate-variance envelope for a non-adversarial, but policed, traffic flow); (3) admission control tests based on the *Adversarial Mode* $RV^*(t)$ traffic characterization (the worst-case rate-variance envelope of a policed flow); and (4) the admission control test of [25].

Trace-driven Simulation - For the simulation curves of Figures 9(a) and 9(b), the average utilization of the multiplexer, and hence the number of multiplexed flows, 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 flows for a given QoS. Regardless, the utilizations are in the range of 79% to 93% (61 to 72 flows on the simulated 45 Mbps link) for $P_L = 10^{-3}$, and in the range of 64% to 88% (49 to 68 flows) for $P_L = 10^{-6}$. Such high utilizations indicate that these MPEG flows 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 flows. Notice that the non-adversarial-mode curves are quite close to those of the trace-driven simulation, indicating that with only knowledge of the flows' deterministic parameters (in this case, six worst-case rate-interval pairs), the maximum entropy scheme 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 flow's rate-variance envelope. As described in Section 3, $RV^*(t)$ *bounds* the stochastic properties of policed flows so that statistical QoS guarantees can be provided even if all flows are independently adversarial. Consequently, the $RV^*(t)$ envelope is necessarily more pessimistic than the $\widehat{RV}(t)$ envelope for non-worst-case policed flows (cf. Figure 3) so that the adversarial-mode scheme captures some, but not all, of the possible statistical multiplexing gain. Its utilizations are 43% to 67% ($P_L = 10^{-3}$) and 25% to 51% ($P_L = 10^{-6}$) for delays between 1 and 80 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 are 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.

Comparison - The final curve depicts admission control experiments based on [25]. Here,

the trace is characterized in the same way as in adversarial and non-adversarial modes. The test assumes that sources transmit traffic according to an extremal periodic on-off model with these parameters. As shown, the test is achieves utilizations in the range of 36% to 67% for delays less than 80 msec and loss probabilities less than 10^{-3}. For a more stringent loss probability requirement of 10^{-6}, the admissible region of [25] is 16% to 48%. The primary reason for the performance gain of adversarial allocation over [25] is the underlying admission control algorithm (cf. Section 5 and [11]) rather than the bounds on the flows parameters, as both approaches find the most adversarial flow to be a type of on-off flow. Moreover, the non-adversarial scheme, unlike both adversarial mode and [25], is able to extract nearly the full statistical multiplexing gain via the maximum entropy technique.

8 Conclusions

Design of admission control and capacity allocation algorithms encounter 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 flows, which necessitates a *deterministic* traffic model. In this paper, we introduced two schemes for delivering a statistically multiplexed service that extracts a traffic flow's stochastic envelope from its network-enforceable deterministic parameters. We first showed how to bound a policed flow's rate-variance envelope to provide a probabilistically guaranteed 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 flows; this latter approach uses maximum entropy techniques to allow the network to exploit a further statistical multiplexing gain when multiplexing policed, but non statistically-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 for both deterministic and statistical services, 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] 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.

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

[3] T. Cover and J. Thomas. *Elements of Information Theory.* John Wiley & Sons, 1991.

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

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

[6] 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.

[7] 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.

[8] 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.

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

[11] E. Knightly and N. Shroff. Admission control for statistical QoS: Theory and practice. *IEEE Network*, 13(2):20–29, March 1999.

[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 S. Tripathi. On the characterization of VBR MPEG streams. In *Proceedings of ACM SIGMETRICS '97*, pages 192–202, Seattle, WA, June 1997.

[14] 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.

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

[16] 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.

[17] 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.

[18] 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.

[19] 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.

[20] 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.

[21] F. Lo Presti, Z. Zhang, D. Towsley, and J. Kurose. Source time scale and optimal buffer/bandwidth tradeoff for heterogeneous regulated traffic in an network node. *IEEE/ACM Transactions on Networking*, 7(4):490–501, August 1999.

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

[23] S. Rajagopal, M. Reisslein, and K. Ross. Packet multiplexers with adversarial regulated traffic. In *Proceedings of IEEE INFOCOM '98*, San Francisco, CA, March 1998.

[24] M. Reisslein, K. Ross, and S. Rajagopal. Guaranteeing statistical QoS to regulated traffic: the multiple node case. In *Proceedings of 37th IEEE Conference on Decision and Control*, Tampa, FL, December 1998.

[25] M. Reisslein, K. Ross, and S. Rajagopal. Guaranteeing statistical QoS to regulated traffic: the single node case. In *Proceedings of IEEE INFOCOM '99*, New York, NY, March 1999.

[26] V. Ribeiro, R. Riedi, M. Crouse, and R. Baraniuk. Simulation of nongaussian long-range-dependent traffic using wavelets. In *Proceedings of ACM SIGMETRICS '99*, pages 1–12, Atlanta, GA, June 1999.

[27] 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.

[28] J. Shore and R. Johnson. Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Transactions on Information Theory*, 26(1):26–37, January 1980.

[29] 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.

[30] T. Wu and E. Knightly. Buffering vs. smoothing for end-to-end QoS: Fundamental issues and comparison. In *Proceedings of Performance '99*, Istanbul, Turkey, October 1999.