Statistical Multiplexing Gain of Link Scheduling Algorithms in QoS Networks *

Technical Report: University of Virginia, CS-99-21

Robert Boorstyn * *Almut Burchard* ** *Jörg Liebeherr* † *Chaiwat Oottamakorn* *

* Department of Electrical Engineering
Polytechnic University
Brooklyn, NY 11201

** Department of Mathematics
University of Virginia
Charlottesville, VA 22903

† Department of Computer Science
University of Virginia
Charlottesville, VA 22903

July 1999

Abstract

A statistical network service which allows a certain fraction of traffic to not meet its QoS guarantees can extract additional capacity from a network by exploiting statistical properties of traffic. Here we consider a statistical service which assumes statistical independence of flows, but does not make any assumptions on the statistics of traffic sources, other than that they are regulated, e.g., by a leaky bucket. Under these conditions, we present functions, so-called *local effective envelopes* and *global effective envelopes*, which are, with high certainty, upper bounds of multiplexed traffic. We show that these envelopes can be used to obtain bounds on the amount of traffic on a link that can be provisioned with statistical QoS. A key advantage of our bounds is that they can be applied with a variety of scheduling algorithms. In fact, we show that one can reuse existing admission control functions that are available for scheduling algorithms with a deterministic service. We present numerical examples which compare the number of flows with statistical QoS guarantees that can be admitted with our effective envelope approach to those achieved with existing methods.

Key Words: Statistical Multiplexing, Statistical Service, Scheduling, Quality-of-Service.

*This work is supported in part by the National Science Foundation through grants NCR-9624106 (CAREER), ANI-9730103, and DMS-9971493, and by the New York State Center for Advanced Technology in Telecommunications (CATT).

1 Introduction

Performance guarantees in QoS networks are either deterministic or statistical. A *deterministic service* guarantees that all packets from a flow satisfy given worst-case end-to-end delay bounds and no packets are dropped in the network [4, 6, 10, 17]. A deterministic service provides the highest level of QoS guarantees, however, it leaves a significant portion of network resources on the average unused [30].

A *statistical service* makes probabilistic service guarantees, for example, of the form:

$$Pr[Delay > X] < \varepsilon \quad \text{or} \quad Pr[Loss] < \varepsilon \ ,$$

where ε is generally small, e.g., $\varepsilon = 10^{-6}$. By allowing a fraction of traffic to violate its QoS guarantees, one can improve the statistical multiplexing gain at network links and increase the achievable link utilization. The key assumption that leads to the definition of statistical services is that traffic arrivals are viewed as random processes. With this assumption a statistical service can improve upon a deterministic service by (1) taking advantage of knowledge about the statistics of traffic sources, and (2) by taking advantage of the statistical independence of flows.

Since it is often not feasible to obtain a reliable statistical characterization of traffic sources, recent research on statistical QoS has attempted to exploit statistical multiplexing without assuming a specific source model. Starting with the seminal work in [10], researchers have investigated the statistical multiplexing gain by only assuming that flows are statistically independent, and that traffic from each flow is constrained by a deterministic regulator, e.g., by a leaky bucket [8, 10, 9, 11, 12, 14, 19, 20, 23, 24, 25].

In this paper we attempt to provide new insights into the problem of determining the multiplexing gain of statistically independent, regulated, but otherwise adversarial traffic flows at a network link. We introduce the notion of *effective envelopes*, which are, with high certainty, upper bounds on the aggregate traffic of regulated flows. We use effective envelopes to devise admission control tests for a statistical service for a large class of scheduling algorithms. We show that with effective envelopes, admission control for a statistical service can be done in a similar fashion as with deterministic envelopes for a deterministic service [4, 6]. In fact, we show that one can reuse admission control conditions derived for various packet scheduling algorithms in the context of a deterministic service, e.g., [6, 17, 33]. This is encouraging, since, with few exceptions [9, 14], only few results are available on statistical multiplexing of adversarial traffic, which can consider non-trivial scheduling algorithms.

The related work for this paper is all previous work which attempts to consolidate the deterministic network calculus [6] with statistical multiplexing (e.g., [4, 12, 13, 14, 16]). In addition, of particular relevance to this paper are all previous results on statistical multiplexing gain with adversarial regulated traffic, as cited above. In Section 5 we compare our work to the existing literature.

The results derived in this paper only apply to a single node. Since traffic from multiple flows passing through the same sequence of congested nodes may become correlated, the assumption

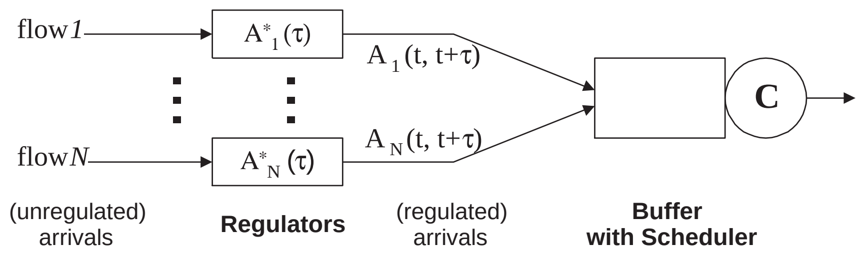

Figure 1: Regulators and scheduler at a link.

of statistical independence of flows may not hold. Only few results are currently available on end-to-end QoS guarantees for adversarial regulated traffic [9, 24, 25].

The remaining sections of this paper are structured as follows. In Section 2 we specify our assumptions on the traffic and define the effective envelopes. In Section 3 we derive sufficient schedulability conditions for a general class of packet schedulers, which can be used for a deterministic and (two types of) statistical QoS guarantees. In Section 4, we use large deviations results to derive bounds for effective envelopes. In Section 5 we contrast our results to the existing literature. In Section 6 we compare the statistical multiplexing gain attainable with the effective envelopes approach to those obtained with other methods [10, 14, 23, 25]. In Section 7 we present conclusions of our work. In Appendix A we provide additional numerical examples using MPEG video traces. In Appendix B we provide some additional comments on the simulation experiments in Section 7 and Appendix A.

2 Traffic Arrivals and Envelope Functions

We consider traffic arrivals to a single link with transmission rate C. As shown in Figure 1, the arrivals from each flow are policed by a regulator, and then inserted into a buffer. A scheduler determines the order in which traffic in the buffer is transmitted. In the following, we view traffic mainly as continuous-time fluid-flow traffic. Note, however, that our discussion applies, without restrictions, to discrete-time or discrete-size (packetized) views of traffic arrivals.

QoS guarantees for a flow j are specified in terms of a delay bound d_j. A QoS violation occurs if traffic from flow j experiences a delay exceeding d_j. (We assume that delays consist only of waiting time in the buffer and transmission time.)

2.1 Traffic Arrivals

Traffic arrivals to the link come from a set of flows which is partitioned into Q classes. We use $\mathcal{C}_q$ to denote set of flows in class q and N_q to denote the number of flows in class q. (Each flow may itself be an aggregate of the traffic from multiple sessions.)

The traffic arrivals from flow j in an interval $[t_1, t_2)$ are denoted as $A_j(t_1, t_2)$. We assume that a traffic flow is characterized by a family of nonnegative random variables $A_j(t_1, t_2)$ which is characterized as follows:

(A1) **Additivity.** For any $t_1 < t_2 < t_3$, we have $A_j(t_1, t_2) + A_j(t_2, t_3) \ = \ A_j(t_1, t_3)$.

(A2) **Subadditive Bounds.** Traffic A_j is regulated by a deterministic subadditive envelope A_j^* as

$$A_j(t, t+\tau) \leq A_j^*(\tau) \qquad \forall t \geq 0, \forall \tau \geq 0 \ . \tag{1}$$

(A3) **Stationarity.** The A_j are *stationary* random variables, i.e., $\forall t, t' > 0$

$$Pr[A_j(t, t+\tau) \leq x] = Pr[A_j(t', t'+\tau) \leq x] \ . \tag{2}$$

In other words, all time shifts of A_j are equally probable.

(A4) **Independence.** The A_i and A_j are stochastically independent for all $i \neq j$.

(A5) **Homogeneity within a Class.** Flows in the same class have identical deterministic envelopes and identical delay bounds. So, $A_i^* = A_j^*$ and $d_i = d_j$ if i and j are in the same class. Henceforth, we denote by d_q the delay bound associated with traffic from class q. By $A_{\mathcal{C}_q}$ we denote the arrivals from class q, that is, $A_{\mathcal{C}_q}(t, t+\tau) = \sum_{j \in \mathcal{C}_q} A_j(t, t+\tau)$.

Remarks:

- We want to point out that the above assumptions are quite general. The class of subadditive deterministic traffic envelopes is the most general class of traffic regulators [4, 6]. The assumptions on the randomness of flows are also quite general. Note that we do not require ergodicity.

- The traffic regulators most commonly used in practice are *leaky buckets* with a peak rate enforcer [1, 2]. Here, traffic on flow j is characterized by three parameters (P_j, σ_j, ρ_j) with a deterministic envelope given by

$$A_j^*(\tau) = \min\{P_j \tau, \sigma_j + \rho_j \tau\} \qquad \forall \tau \geq 0 \ , \tag{3}$$

  where $P_j \geq \rho_j$ is the peak traffic rate, ρ_j is the average traffic rate, and σ_j is a burst size parameter. We will use this type of regulators in our numerical examples in Section 6.

  As a generalization of the peak-rate enforced leaky bucket the traffic on flow j may be characterized by a set of parameters $\{\sigma_{jk}, \rho_{jk}\}_{k=1,\ldots K_j}$, with a deterministic envelope

$$A_j^*(\tau) = \min_{k=1,\ldots K_j} \{\sigma_{jk} + \rho_{jk} \tau\} \qquad \forall \tau \geq 0 \ . \tag{4}$$

  We will use this type of regulators in our numerical examples in Appendix A.

- A consequence of subadditivity of the A_j^* is that the limit $\rho_j := \lim_{\tau \to \infty} A_j^*(\tau)/\tau$ exists, and that it provides an upper bound for the longterm arrival rate for A_j. We assume without loss of generality, that for all t,

$$\lim_{\tau \to \infty} \frac{A_j(t, t+\tau)}{\tau} \ = \ \rho_j \ . \tag{5}$$

- Stationarity has the useful consequence that expected values can be computed as *long-time averages*. For example, for any function F,

$$E\left[F(A_j(t,t+\tau)\right] \;=\; \lim_{T\to\infty} E\left[\frac{1}{T}\int_0^T F(A_j(s,s+\tau))\,ds\right] \;. \tag{6}$$

  Similar relations hold for the joint distributions of several random variables.

2.2 Definition of Effective Envelopes

We next define *local effective envelopes* and *global effective envelopes* which are, with high certainty, upper bounds on aggregate traffic from a given class q. The envelopes are defined for a set of flows $\mathcal{C}$ with arrival functions A_j and aggregate traffic $A_{\mathcal{C}}(t,t+\tau) = \sum_{j\in\mathcal{C}} A_j(t,t+\tau)$.

Definition 1 *A* **local effective envelope** *for $A_{\mathcal{C}}(t,t+\tau)$ is a function $\mathcal{G}_{\mathcal{C}}$ that satisfies for all $\tau \geq 0$ and all t*

$$Pr\Big[A_{\mathcal{C}}(t,t+\tau) \leq \mathcal{G}_{\mathcal{C}}(\tau;\varepsilon)\Big] \geq 1-\varepsilon \;. \tag{7}$$

In other words, a *local effective envelope* provides a bound for the aggregate arrivals $A_{\mathcal{C}}(t,t+\tau)$ for *any specific* ('local') time interval of length τ. Under the stationarity assumption (A3), Eqn. (7) holds for all times t, provided that it only holds for one value $t = t_o$.

It is easy to see that there exists a smallest local effective envelope, since the minimum of two local effective envelopes is again such an envelope. Note, however, that local effective envelopes are in general not subadditive in τ, but satisfy the weaker property

$$\mathcal{G}_{\mathcal{C}}(\tau_1+\tau_2,\varepsilon_1+\varepsilon_2) \;\leq\; \mathcal{G}_{\mathcal{C}}(\tau_1,\varepsilon_1)+\mathcal{G}_{\mathcal{C}}(\tau_2,\varepsilon_2) \;. \tag{8}$$

A local effective envelope $\mathcal{G}_{\mathcal{C}}(\tau;\varepsilon)$ is a bound for the traffic arrivals in an arbitrary, but fixed interval of length τ. Global effective envelopes, to be defined next, are bounds for the arrivals in all subintervals $[t,t+\tau)$ of a larger interval.

For the definition of global effective envelopes, we take advantage of the notion of empirical envelopes, as used in [4, 30]. Consider a time interval I_β of length β. The **empirical envelope** $\mathcal{E}_{\mathcal{C}}(\,.\,;\beta)$ of a collection $\mathcal{C}$ of flows is the maximum traffic in any time interval of length $\tau \leq \beta$ in the interval I_β, as follows:

$$\mathcal{E}_{\mathcal{C}}(\tau;\beta) = \sup_{[t,t+\tau)\subseteq I_\beta} A_{\mathcal{C}}(t,t+\tau) \;. \tag{9}$$

Definition 2 *A* **global effective envelope** *for an interval I_β of length β is a subadditive function $\mathcal{H}_{\mathcal{C}}(.;\beta)$ which satisfies*

$$Pr\Big[\mathcal{E}_{\mathcal{C}}(\tau;\beta) \leq \mathcal{H}_{\mathcal{C}}(\tau;\beta,\varepsilon),\;\; \forall 0 \leq \tau \leq \beta\Big] \geq 1-\varepsilon \;. \tag{10}$$

The attribute 'global' is justified since $\mathcal{H}_{\mathcal{C}}(\tau;\beta,\varepsilon)$ is a bound for traffic for all intervals of length $\tau \leq \beta$ in I_β. Note we can always choose $\mathcal{E}_{\mathcal{C}}(\tau;\beta) \leq \sum_{j\in\mathcal{C}} \mathcal{E}_j(\tau)$ for all $\tau \leq \beta$, where $\mathcal{E}_j(\tau;\beta) = \sup_{[t,t+\tau)\subseteq I_\beta} A_j(t,t+\tau)$ is the empirical envelope of a single flow j.

Due to stationarity of the A_j, Eqn. (10) holds for *all* intervals of length β, if it holds for one specific interval I_β. When applied to scheduling, we will select β such that $\beta \geq B_{max}$, i.e., it has at least the length of the longest busy period.[1]

Assuming that one has obtained local or global effective envelopes separately for each traffic class, the following lemma helps to obtain bounds for the traffic from all classes.

Lemma 1 Given a set of flows that is partitioned into Q classes $\mathcal{C}_q$, with arrival functions $A_{\mathcal{C}_q}$. Let $\mathcal{G}_{\mathcal{C}_q}(.;\varepsilon)$ and $\mathcal{H}_{\mathcal{C}_q}(.;\beta,\varepsilon)$ be local and global effective envelopes for class q. Then the following inequalities hold.

(a) If $\sum_q \mathcal{G}_{\mathcal{C}_q}(\tau,\varepsilon) \leq x$, then, for all t, $Pr\left[\sum_p A_{\mathcal{C}_q}(t,t+\tau) > x\right] < Q\cdot\varepsilon$.

(b) If $\sum_q \mathcal{H}_{\mathcal{C}_q}(\tau,\beta;\varepsilon) \leq x(\tau)$ for all τ, then $Pr\left[\exists q\, \exists\tau :\ \sum_q \mathcal{E}_{\mathcal{C}_q}(\tau,\beta) > x(\tau)\right] < Q\cdot\varepsilon$.

Proof: We only prove part (a) of the lemma. The proof for (b) is almost identical as for (a). Assume

$$\sum_q \mathcal{G}_{\mathcal{C}_q}(\tau,\alpha) \leq x. \tag{11}$$

Consider the event

$$\sum_q A_{\mathcal{C}_q}(t,t+\tau) > x, \quad \forall t. \tag{12}$$

In this case, it follows that:

$$\sum_q A_{\mathcal{C}_q}(t,t+\tau) > x \geq \sum_q \mathcal{G}_{\mathcal{C}_q}(\tau,\alpha)\ . \tag{13}$$

This in turn implies that

$$\exists q : A_{\mathcal{C}_q}(t,t+\tau) > \mathcal{G}_{\mathcal{C}_q}(\tau,\alpha)\ . \tag{14}$$

Using the defining property of the local effective envelope, we can bound the probability of the last event by:

$$Pr\left[\exists q : A_{\mathcal{C}_q}(t,t+\tau) > \mathcal{G}_{\mathcal{C}_q}(\tau,\alpha)\right] \leq \sum_q Pr\left[A_{\mathcal{C}_q}(t,t+\tau) > \mathcal{G}_{\mathcal{C}_q}(\tau,\alpha)\right] < Q\alpha\ . \tag{15}$$

□

[1]For arrival functions A_j and regulators with deterministic envelopes A_j^*, the longest busy period in a work-conserving scheduler is given by: $B_{max} = \inf\{\tau > 0\ ;\ \sum_{j\in\mathcal{C}} A_j^*(\tau) \leq \tau\}$.

Our derivations in Section 4 will make it clear that for ε small enough, both $\mathcal{G}_{\mathcal{C}_q}(\tau;\varepsilon)$ and $\mathcal{H}_{\mathcal{C}_q}(\tau;\beta,\varepsilon)$ are not very sensitive with respect to ε, so that the bounds for ε and $Q\cdot\varepsilon$ are comparable.

In Section 3, we will use the local and global effective envelopes to derive sufficient schedulability conditions for a statistical service for a general scheduling algorithms. We will show that admission control with local and global effective envelopes can be done in an analogous fashion as admission control with deterministic envelopes is used in a deterministic service.

In Section 4 we will use well-known large deviations results to provide bounds on envelopes $\mathcal{G}_{\mathcal{C}_q}(.;\varepsilon)$ and $\mathcal{H}_{\mathcal{C}_q}(.;\beta,\varepsilon)$ for collections of flows satisfying (A1)-(A5). Our goal is to select the bounds as small as possible.

3 Deterministic and Statistical Schedulability Conditions

In this section, we present three schedulability conditions for a general class of work-conserving scheduling algorithms. The first condition, expressed in terms of deterministic envelopes, ensures deterministic guarantees. The second and third conditions, which use the local and global effective envelopes, respectively, yield statistical guarantees. All three schedulability conditions will be derived from the same expression for the delay of a traffic arrival in an arbitrary work-conserving scheduler (Eqn. (21) in Section 3.1).

In our discussions, we will not take into considerations that packet transmissions on a link cannot be preempted. This assumption is reasonable when packet transmission times are short. For the specific scheduling algorithms considered in this paper, accounting for non-preemptiveness of packets does not introduce principal difficulties, however, it requires additional notation (see [17]). Also, to keep notation minimal, we assume that the transmission rate of the link is normalized, that is $C = 1$.

3.1 Schedulability

Suppose a (tagged) arrival from a flow j in class q $(j \in \mathcal{C}_q)$ arrives to a work-conserving scheduler at time t. Without loss of generality we assume that the scheduler is empty at time 0. We will derive a condition that must hold so that the arrival does not violate its delay bound d_q.

Let us use $A^{q,t}(t_1,t_2)$ to denote the traffic arrivals in the time interval $[t_1,t_2)$ which will be served before a class q arrival at time t. Let $A^{q,t}_{\mathcal{C}_p}(t_1,t_2)$ denote the traffic arrivals from flows in $\mathcal{C}_p$ which contribute to $A^{q,t}(t_1,t_2)$.

Suppose that $t-\hat{\tau}$ is the last time before t when the scheduler does not contain traffic that will be transmitted before the tagged arrival from class q. That is,

$$\hat{\tau} = \inf\{x \geq 0 \mid A^{q,t}(t-x,t) \leq x\} \ . \tag{16}$$

So, in the time interval $[t-\hat{\tau},t)$ the scheduler is continuously transmitting traffic which will be served before the tagged arrival. (Note that $\hat{\tau}$ is a function of t and q. To keep notation simple, we do not make the dependence explicit.)

Given $\hat{\tau}$, the tagged class-q arrival at time t will leave the scheduler at time $t + \delta$ if $\delta > 0$ is such that

$$\delta = \inf \left\{ \tau_{out} \mid A^{q,t}(t - \hat{\tau}, t + \tau_{out}) \leq \hat{\tau} + \tau_{out} \right\} \ . \tag{17}$$

Hence, the tagged class-q arrival does not violate its delay bound d_q if and only if

$$\forall \hat{\tau} \ \exists \tau_{out} \leq d_q \ : \left\{ A^{q,t}(t - \hat{\tau}, t + \tau_{out}) \leq \hat{\tau} + \tau_{out} \right\} \ . \tag{18}$$

Then, the traffic arrival does not have a deadline violation if d_q is selected such that

$$\sup_{\hat{\tau}} \left\{ A^{q,t}(t - \hat{\tau}, t + d_q) - \hat{\tau} \right\} \leq d_q \ . \tag{19}$$

In general, Eqn. (19) is a sufficient condition for meeting a delay bound. For FIFO and EDF schedulers, the condition is also necessary [17]. [2]

For a specific work-conserving scheduling algorithm, let $\bar{\tau}$ (with $-\hat{\tau} \leq \bar{\tau}_p \leq d_q$) denote the smallest values for which

$$A_{\mathcal{C}_p}(t - \hat{\tau}, t + \bar{\tau}_p) \geq A^{q,t}_{\mathcal{C}_p}(t - \hat{\tau}, t + d_q) \ . \tag{20}$$

Remark: For most work-conserving schedulers one can easily find $\bar{\tau}_p$ such that equality holds in Eqn. (20). For example, for FIFO, SP,[3] and EDF schedulers, we have:

$$\begin{array}{ll} \text{FIFO:} & \bar{\tau}_p = 0 \\ \text{SP:} & \bar{\tau}_p = \begin{cases} -\hat{\tau} & , p > q \\ 0 & , p = q \\ d_q & , p < q \end{cases} \\ \text{EDF:} & \bar{\tau}_p = \max\{-\hat{\tau}, d_q - d_p\} \ . \end{array}$$

With Eqn. (20), the arrival from class q at time t does not have a violation if d_q is selected such that

$$\sup_{\hat{\tau}} \left\{ \sum_p A_{\mathcal{C}_p}(t - \hat{\tau}, t + \bar{\tau}_p) - \hat{\tau} \right\} \leq d_q \ . \tag{21}$$

Next, we show how Eqn. (21) can be used to derive schedulability conditions for deterministic and statistical services, using deterministic envelopes, local effective envelopes, and global effective envelopes. For a deterministic service, the delay bound d_q must be chosen such that Eqn. (21) is never violated. For a statistical service, d_q is chosen such that a violation of Eqn. (21) is a rare event.

[2] A FIFO scheduler transmits traffic in the order of arrival times. An EDF (Earliest-Deadline-First) scheduler tags traffic with a deadline which is set to the arrival time plus the delay bound d_q, and transmits traffic in the order of deadlines.

[3] An SP (Static Priority) scheduler assigns each class a priority level (we assume that a lower class index indicates a higher priority), and has one FIFO queue for traffic arrivals from each class. SP always transmits traffic from the highest priority FIFO queue which has a backlog.

3.2 Schedulability with Deterministic Envelopes

Exploiting the property of deterministic envelopes in Eqn. (1), we can relax Eqn. (21) to

$$\sup_{\hat{\tau}} \left\{ \sum_p \sum_{j \in \mathcal{C}_p} A_j^*(\bar{\tau}_p + \hat{\tau}) - \hat{\tau} \right\} \leq d_q \ . \tag{22}$$

Since, $\bar{\tau}_p + \hat{\tau}$ is not dependent on t, we have obtained a sufficient schedulability condition for an arbitrary traffic arrival. We refer the reader to [17] to verify that for FIFO and EDF scheduling algorithms the condition in Eqn. (22) is also necessary, in the sense that if it is violated, then there exist arrival patterns conforming with A_j^* leading to deadline violations for class q. For SP scheduling, the condition is necessary only if the deterministic envelopes are concave functions.

Next we present bounds on the likelihood of a violation of Eqn. (21), using local and global effective envelopes.

3.3 Schedulability with Local Effective Envelopes

With Eqn. (21), the probability that the tagged arrival from time t experiences a deadline violation is less than ε if d_q is selected such that

$$Pr\left[\sup_{\hat{\tau}} \left\{ \sum_p A_{\mathcal{C}_p}(t - \hat{\tau}, t + \bar{\tau}_p) - \hat{\tau} \right\} \leq d_q \right] \geq 1 - \varepsilon \ . \tag{23}$$

Let us, for the moment, make the convenient assumption that

$$Pr\left[\sup_{\hat{\tau}} \left\{ \sum_p A_{\mathcal{C}_p}(t - \hat{\tau}, t + \bar{\tau}_p) - \hat{\tau} \right\} \leq d_q \right] \approx \inf_{\hat{\tau}} Pr\left[\sum_p A_{\mathcal{C}_p}(t - \hat{\tau}, t + \bar{\tau}_p) - \hat{\tau} \leq d_q \right] \ . \tag{24}$$

Assuming that equality holds in Eqn. (24), we can re-write Eqn. (23) as

$$\inf_{\hat{\tau}} Pr\left[\sum_p A_{\mathcal{C}_p}(t - \hat{\tau}, t + \bar{\tau}_p) - \hat{\tau} \leq d_q \right] \geq 1 - \varepsilon \ . \tag{25}$$

Remark: The assumption in Eqn. (24) requires further justification, since, in general, the right hand side is larger than the left hand side. Note that standard extreme-value theory [3] is not immediately applicable to the left hand side of Eqn. (24), since the supremum in $\sup_{\hat{\tau}} \left\{ \sum_p A_{\mathcal{C}_p}(t - \hat{\tau}, t + \bar{\tau}_p) - \hat{\tau} \right\}$ is taken over a family of random variables indexed by the continuous parameter $\hat{\tau}$. Thus, one must consider the correlations between the $A_{\mathcal{C}_p}(t - \hat{\tau}, t + \bar{\tau}_p)$ for different values of $\hat{\tau}$ in order to obtain any useful estimate for the distribution of the supremum. (Even if $\hat{\tau}$ is discrete, neglecting these correlations will lead to poor estimates.) One way to provide a theoretical justification for the assumption in Eqn. (24) is to assume that arrivals follow a Gaussian processes. Also, several works on statistical QoS have used Eqn. (24) with equality [5, 13, 14, 15, 16], and, in several cases, have supported the assumption in Eqn. (24) with numerical examples.

Recall from the definition of the local effective envelope that $\mathcal{G}_{\mathcal{C}_p}(\tau,\varepsilon) \le x$ implies $Pr\left[A_{\mathcal{C}_p}(t,t+\tau) > x\right] < \varepsilon$. Then, with Lemma 1(a) and assuming that Eqn. (24) holds with equality, we have that a class-q arrival has a deadline violation with probability $< \varepsilon$ if d_q is selected such that

$$\sup_{\hat{\tau}} \left\{ \sum_p \mathcal{G}_{\mathcal{C}_p}(\bar{\tau}_p + \hat{\tau}, \varepsilon/Q) - \hat{\tau} \right\} \le d_q \ . \tag{26}$$

With Eqn. (26) we have found an expression for the probability that an arbitrary traffic arrival results in a violation of delay bounds. This condition can be viewed as a generalization of schedulability conditions for statistical QoS from [13, 14, 16].

The drawback of the condition in Eqn. (26) is its dependence on the assumption in Eqn. (24). Empirical evidence from numerical examples, including those presented in this paper, as well as numerical evidence from previous work which employed this assumption [14, 5], suggests that Eqn. (26) is not overly optimistic. However, it should be noted that the bound in Eqn. (26) is not a rigorous one.

3.4 Schedulability with Global Effective Envelopes

We next use global effective envelopes to express the probability of a deadline violation in a time interval. We will see that this bound, while more pessimistic, can be made rigorous.

Consider again the traffic arrival from class q which occurs at time t. The arrival time t lies in a busy period of the scheduler I_β of length at most β, which starts at time $\le t - \hat{\tau}$ and which ends at a time after the tagged arrival has departed.

Using the properties of the empirical envelope $\mathcal{E}_{\mathcal{C}_p}(.\,;\beta)$ as defined in Section 2 we have that, for all t and $\bar{\tau}_p + \hat{\tau} \ge 0$,

$$\mathcal{E}_{\mathcal{C}_p}(\bar{\tau}_p + \hat{\tau}; \beta) \ge A_{\mathcal{C}_p}(t - \hat{\tau}, t + \bar{\tau}_p). \tag{27}$$

Thus, we do not have any deadline violation of any class-p arrival in the time interval I_β, if

$$\sup_{\hat{\tau}} \left\{ \sum_q \mathcal{E}_{\mathcal{C}_p}(\bar{\tau}_p + \hat{\tau}; \beta) - \hat{\tau} \right\} \le d_p \ . \tag{28}$$

With Lemma 1(b), the probability that an arrival from class q experiences a deadline violation in the interval I_β is $< \varepsilon$, if d_q is selected such that

$$\sup_{\hat{\tau}} \left\{ \sum_p \mathcal{H}_{\mathcal{C}_p}(\bar{\tau}_p + \hat{\tau}; \beta, \varepsilon/Q) - \hat{\tau} \right\} \le d_q \ . \tag{29}$$

Note that the nature of the statistical guarantees derived with local effective envelopes (in Subsection 3.3) and with global effective envelopes (in Subsection 3.4) are quite different. Local effective envelopes are (under the assumption in Eqn. (24)) concerned with the probability that a deadline violation occurs at a certain time. Global effective envelopes address the probability that a deadline violation occurs in a certain time interval. Clearly, a service which guarantees the latter is more stringent, and will lead to more conservative admission control.

Lastly, we want to point to the structural similarities of the conditions in Eqs. (22), (26), and (29). Thus, schedulability conditions which have been derived for a deterministic service can be reused, without modification, for a statistical service if the respective envelopes or bounds thereof are available.

4 Construction of Effective Envelopes

In this section we will construct the local and global effective envelopes for the aggregate traffic from a set of flows as described in (A1)-(A5). Throughout this section, we will work only with flows from a single class. So, we will drop the index 'q', and denote by $\mathcal{C}$ and N, respectively, the set of flows and the number of flows. We denote by $A^*(\tau)$ the common deterministic envelope for the flows in $\mathcal{C}$ (that is, $A_j^*(\tau) = A^*(\tau)$ for all $j \in \mathcal{C}$), and by $A_{\mathcal{C}}(t, t+\tau)$ the aggregate traffic. The empirical envelope of the aggregate traffic will be denoted by $\mathcal{E}_{\mathcal{C}}$, and the local and global effective envelopes by $\mathcal{G}_{\mathcal{C}}$ and $\mathcal{H}_{\mathcal{C}}$.

Our derivations proceed in the following steps:

Step 1. We compute bounds for the moments of the individual flows $A_j(t, t+\tau)$. Since the flows are independent, this directly leads to bounds for the moments of $A_{\mathcal{C}}(t, t+\tau)$.

Step 2. We use the Chernoff bound to determine a local effective envelope $\mathcal{G}_{\mathcal{C}}$ directly from our bounds on the moments.

Step 3. We use a geometric argument to construct $\mathcal{H}_{\mathcal{C}}$ from any local effective envelope $\mathcal{G}_{\mathcal{C}}$. Specifically, we will provide bounds of the following nature:

$$\mathcal{G}_{\mathcal{C}}(\tau; \varepsilon) \le \mathcal{H}_{\mathcal{C}}(\tau; \beta, \varepsilon) \le \mathcal{G}_{\mathcal{C}}(\tau'; \varepsilon') \ . \tag{30}$$

where $\tau'/\tau > 1$ and $\varepsilon'/\varepsilon < 1$ depend on β. We claim that for ε sufficiently small and β not too large, $\tau'/\tau \approx 1$, and the resulting global effective envelope is reasonably close to the local effective envelope.

The three steps will be discussed in Subsections 4.1, 4.2, and 4.3.

4.1 Moment bounds

The moment generating functions of the distributions of $A_{\mathcal{C}}$ and the A_j are defined as follows:

$$M_{\mathcal{C}}(s, \tau) \ := \ E[e^{A_{\mathcal{C}}(t,t+\tau)s}] \ , \tag{31}$$

$$M_j(s, \tau) \ := \ E[e^{A_j(t,t+\tau)s}] \ . \tag{32}$$

Stationarity (A3) guarantees that the moment generating functions do not depend on t. Due to the stochastic independence (A4) and homogeneity (A5) of the flows, we can write:

$$M_{\mathcal{C}}(s,\tau) \;=\; \prod_{j=1}^{N} M_j(s,\tau) \;=\; \Big(M_j(s,\tau)\Big)^N \;. \tag{33}$$

Thus, to obtain a bound on $M_{\mathcal{C}}(s,\tau)$, it is sufficient to bound the moment-generating function of a single flow $A_j(t,t+\tau)$.

The k-th moments of $A_j(t,t+\tau)$ and $A_{\mathcal{C}}(t,t+\tau)$ are defined by

$$m_{\mathcal{C}}^{(k)}(\tau) \;:= E[(A_{\mathcal{C}}(t,t+\tau))^k] \;, \tag{34}$$

$$m_j^{(k)}(\tau) \;:= E[(A_j(t,t+\tau))^k] \;. \tag{35}$$

The moments are related with the moment generating functions by

$$M_j(s,\tau) \;=\; \sum_{k=0}^{\infty} m_j^{(k)}(\tau)\frac{s^k}{k!} \;, \tag{36}$$

$$M_{\mathcal{C}}(s,\tau) \;=\; \sum_{k=0}^{\infty} m_{\mathcal{C}}^{(k)}(\tau)\frac{s^k}{k!} \;. \tag{37}$$

The following lemma will be used to provide bounds on the moment generation function and the moments of the arrivals on a flow $A_j(t,t+\tau)$.

Lemma 2 *Assume that $A_j(t,t+\tau)$ satisfies Conditions (A1), (A2), and (A3). Then, for every convex increasing function F,*

$$E\left[F(A_j(t,t+\tau)\right] \;\le\; \left(\frac{\rho\tau}{A^*(\tau)}\right) F(A^*(\tau)) + \left(1-\frac{\rho\tau}{A^*(\tau)}\right) F(0) \;. \tag{38}$$

Proof: By stationarity (A3), we may write

$$E[F(A_j(t,t+\tau)] \;= \lim_{T\to\infty} E\left[\frac{1}{T}\int_0^T F(A_j(t,t+\tau))\,dt\right] \;. \tag{39}$$

Since the limit in Eqn. (39) exists, we may compute it by considering only cases where T is a large integer multiple of τ. We compute the average over $[0,T]$ by partitioning $[0,T]$ into subintervals of length τ, and then averaging over the position of the subintervals.

$$E[F(A_j(t,t+\tau)] \;=\; \lim_{T\to\infty} E\Bigg[\frac{1}{T}\sum_{i=1}^{T/\tau}\int_0^\tau F\big(A_j((i-1)\tau+\omega, i\tau+\omega)\big)\Bigg] \tag{40}$$

$$=\; \lim_{T\to\infty} E\Bigg[\frac{1}{\tau}\int_0^\tau \frac{1}{T/\tau}\sum_{i=1}^{T/\tau} F\big(A_j((i-1)\tau+\omega, i\tau+\omega)\big)d\omega\Bigg] \tag{41}$$

$$\le\; \lim_{T\to\infty} E\Bigg[\max_{0\le\omega\le\tau}\frac{1}{T/\tau}\sum_{i=1}^{T/\tau} F\big(A_j((i-1)\tau+\omega, i\tau+\omega)\big)\Bigg] \;. \tag{42}$$

For a fixed value of T and a fixed arrival pattern $\{A_j(t, t+\tau)\}_{t\geq 0}$, let ω_0 be the shift for which the maximum is assumed in Eqn. (42). Set

$$y_i = A_j((i-1)\tau + \omega_0, i\tau + \omega_0) \ . \tag{43}$$

To obtain an upper bound for the limit in Eqn. (42), we consider the following optimization problem:

$$\text{maximize} \quad \frac{1}{n}\sum_{i=1}^{n} F(y_i) \qquad (n = T/\tau) \tag{44}$$

subject to

$$0 \leq y_i \leq A^*(\tau) \qquad i = 1, \ldots, N \tag{45}$$

$$\sum_{i=1}^{N} y_i \leq A^*(T) \ . \tag{46}$$

By convexity, the maximal value is attained at some point on the boundary of the region defined by the side conditions. Moreover, since F is increasing, side condition (46) holds at this point with equality. Exploiting the symmetry of the problem under permutations of the y_i, we see immediately that the following is a maximizing solution:

$$y_i = \begin{cases} A^*(\tau) & \text{if } i \leq \lfloor \frac{A^*(T)}{A^*(\tau)} \rfloor \\ A^*(T) - i\, A^*(\tau) & \text{if } i = \lfloor \frac{A^*(T)}{A^*(\tau)} \rfloor + 1 \\ 0 & \text{otherwise} \ . \end{cases} \tag{47}$$

This assigns the maximum value $A^*(\tau)$ to as many y_i as possible, subject to the first side condition. So, the maximum of Eqn. (44), up to a rounding error of $O(1/n)$, is

$$\lambda(T)\, F\left(A^*(\tau)\right) \ + \ \left((1 - \lambda(T)\right) F(0) \ , \tag{48}$$

where

$$\lambda(T) \ = \ \frac{A^*(T)}{nA^*(\tau)} \ < \ 1 \ . \tag{49}$$

Inserting the maximum back into Eqn. (42) and recalling that $n = T/\tau$, we obtain the following bound

$$\bar{F}(\tau) \quad \leq \quad \lim_{T\to\infty} \left\{ \lambda(T)\, F\left(A^*(\tau)\right) \ + \ (1 - \lambda(T))\, F(0) \right\} \tag{50}$$

$$= \quad \left(\frac{\rho\,\tau}{A^*(\tau)} \right) F\left(A^*(\tau)\right) + \left(1 - \frac{\rho\,\tau}{A^*(\tau)} \right) F(0) \ . \tag{51}$$

In the evaluation of the limit, we have used that

$$\lim_{T\to\infty} \lambda(T) \ = \ \lim_{T\to\infty} \frac{\tau}{T} \frac{A^*(T)}{A^*(\tau)} \ = \ \frac{\rho\tau}{A^*(\tau)} \tag{52}$$

by the definition of ρ. This completes the proof. □

With Lemma 2, we can easily obtain bounds for the moment-generating function $M_j(s, \tau)$ and the k-th moments $m_j^{(k)}$. These bounds are formulated in Theorems 1 and 2.

Theorem 1 *Given a set of flows $\mathcal{C}$ from a single class which satisfy conditions (A1)–(A5). Let $A_j(t, t+\tau)$ denote the arrivals from a flow $j \in \mathcal{C}$, let $A_{\mathcal{C}}(t, t+\tau)$ denote the aggregate traffic, and let $A^*(\tau)$ denote the subadditive envelope for each flow in $\mathcal{C}$. Then,*

$$M_j(s, \tau) \leq 1 + \frac{\rho\tau}{A^*(\tau)} \left(e^{sA^*(\tau)} - 1 \right) \tag{53}$$

and

$$M_{\mathcal{C}}(s, \tau) \;\leq\; \left(1 + \frac{\rho\,\tau}{A^*(\tau)} \left(e^{sA^*(\tau)} - 1 \right) \right)^N \quad . \tag{54}$$

Proof: Eqn. (53) is obtained by setting $F(y) = e^{sy}$ in Lemma 2. Combining Eqn. (53) with Eqn. (33) yields Eqn. (54). □

The bound in Eqn (54) can be strengthened to bounds for individual moments.

Theorem 2 *Under the assumptions of Theorem 1,*

$$m_j^{(k)}(\tau) \;\leq\; \rho\tau \left(A^*(\tau) \right)^{k-1} \qquad (k \geq 1) \tag{55}$$

and

$$m_{\mathcal{C}}^{(k)}(\tau) \leq k! \cdot \left(\textit{coefficient of } s^k \textit{ in } \quad 1 + \frac{\rho\tau}{A^*(\tau)} \left(e^{sA^*(\tau)} - 1 \right) \right) \quad . \tag{56}$$

Proof: Lemma 2 with $F(y) = y^k$ yields Eqn. (55). Using the formula for the moments in Eqn. (37), we compute

$$\begin{aligned}
m_{\mathcal{C}}^{(k)}(\tau) \;&=\; k! \cdot \left(\text{coefficient of } s^k \text{ in } M_{\mathcal{C}}(s, \tau) \right) & (57)\\
&=\; k! \cdot \left(\text{coefficient of } s^k \text{ in } \left(M_j(s, \tau) \right)^N \right) & (58)\\
&=\; \sum_{k_1 + \cdots + k_n = k} \; \prod_{j=1}^{N} m_j^{(k_j)}(\tau) & (59)\\
&\leq\; k! \cdot \left(\text{coefficient of } s^k \text{ in } 1 + \frac{\rho\tau}{A^*(\tau)} \left(e^{sA^*(\tau)} - 1 \right) \right) \quad . & (60)
\end{aligned}$$

Here, the first line is from Eqn. (37), the second follows from independence, the third combines the Cauchy product formula for power series with Eqn. (36). The inequality in the last line follows from the bounds in Eqn. (56) and the positivity of the moments $m_j^{(k)}(\tau)$. □

The results in this Subsection will now be used to derive bounds on the effective envelope.

4.2 Local Effective Envelopes

4.2.1 Using the Central Limit Theorem

Combining the bound for the second moment from Theorem 2 with the assumption that $E[A_j(t,t+\tau)] = \rho\tau$ yields the bound

$$Var[A_j(t,t+\tau)] \le \underbrace{\rho\tau(A^*(\tau)-\rho\tau)}_{=:\hat{\mathbf{s}}^2} \tag{61}$$

for the variance of the individual flows. As indicated, we define

$$\hat{\mathbf{s}} = \rho\tau\sqrt{\frac{A^*(\tau)}{\rho\tau}-1} \ . \tag{62}$$

By the independence and homogeneity of the flows, it follows that

$$\begin{aligned} Var[A_{\mathcal{C}}(t,t+\tau)] &= N\, Var[A_j(t,t+\tau)] && (63)\\ &\le N\rho\tau(A^*(\tau)-\rho\tau) \ . && (64)\end{aligned}$$

Using first the Central Limit Theorem and then the bound on the variance in Eqn.(64), we see that for $x > \rho\tau$

$$\begin{aligned} Pr[A_{\mathcal{C}}(t,t+\tau) \ge Nx] &\approx 1-\Phi\left(\frac{Nx-N\rho\tau}{Var[A_{\mathcal{C}}(t,t+\tau)]}\right) && (65)\\ &\le 1-\Phi\left(\sqrt{N}\,\frac{x-\rho\tau}{\hat{\mathbf{s}}}\right) \ , && (66)\end{aligned}$$

where Φ is the cumulative normal distribution.

To find $\mathcal{G}_{\mathcal{C}}$ so that

$$Pr[A_{\mathcal{C}}(0,\tau) \ge \mathcal{G}_{\mathcal{C}}(\tau;\varepsilon)] \le \varepsilon \ , \tag{67}$$

we set $Pr[A_{\mathcal{C}}(t,t+\tau) \ge Nx] \approx \varepsilon$ in Eqn. (66) and solve for Nx. This produces an (approximate) local effective envelope of the form

$$\mathcal{G}_{\mathcal{C}}(\tau;\varepsilon) \approx N\rho\tau + z\sqrt{N}\rho\tau\sqrt{\frac{A^*(\tau)}{\rho\tau}-1} \ , \tag{68}$$

where z is defined by $1-\Phi(z) = \varepsilon$ and has the approximate value $z \approx \sqrt{|\log(2\pi\varepsilon)|}$.

We remark that our bound in Eqn. (61) is equivalent to Knightly's bound on the *rate variance* in [14]. The rate variance in [14] is defined by

$$RV[A_j(t,t+\tau)] := Var\left(\frac{A_j(t,t+\tau)}{\tau}\right). \tag{69}$$

Knightly's bound states that

$$RV[A_j(t,t+\tau)] \le \frac{A^*(\tau)}{\tau}\rho - \rho^2 \ , \tag{70}$$

which is just Eqn. (61) with both sides multiplied by τ^{-2}.

4.2.2 Using the Chernoff Bound

While the estimate in Eqn. (68) is asymptotically correct, for finite values of N it is only an approximation. To obtain a rigorous upper bound on $Pr\left[A_{\mathcal{C}}(0,\tau) \geq Nx\right]$, recall the Chernoff bound for a random variable Y (see [21]):

$$Pr[Y \geq y] \leq e^{-sy}\, E[e^{sY}] \qquad \forall s \geq 0 \; . \tag{71}$$

In particular, for $A_{\mathcal{C}}$, this gives

$$\begin{aligned} Pr[A_{\mathcal{C}}(0,\tau) \geq Nx] \;&\leq\; e^{-Nxs} M_{\mathcal{C}}(s,\tau) && (72) \\ &\leq\; \left[e^{-xs} \left(1 + \frac{\rho\,\tau}{A^*(\tau)} \left(e^{sA^*(\tau)} - 1 \right) \right) \right]^N . && (73) \end{aligned}$$

Here, Eqn. (72) simply uses the Chernoff bound, and Eqn. (73) uses Eqn. (54). We want to find a value s which makes the bound in Eqn. (73) as tight as possible. For $x < A^*(\tau)$, the right hand side of (73) is minimal when s is chosen so that

$$e^{sA^*(\tau)} \;=\; \frac{x}{\rho\tau} \, \frac{A^*(\tau) - \rho\tau}{A^*(\tau) - x} \; . \tag{74}$$

Substituting this value of s into Eqn. (73) yields

$$Pr\big[A_{\mathcal{C}}(0,\tau) \geq Nx\big] \leq \left[\left(\frac{\rho\tau}{x} \right)^{\frac{x}{A^*(\tau)}} \left(\frac{A^*(\tau) - \rho\tau}{A^*(\tau) - x} \right)^{1 - \frac{x}{A^*(\tau)}} \right]^N . \tag{75}$$

Again, our goal is to find $\mathcal{G}_{\mathcal{C}}$ satisfying Eqn. (67). Using the bound in Eqn. (75) and enforcing that $\mathcal{G}_{\mathcal{C}}(\tau;\varepsilon)$ is never larger than $NA^*(\tau)$ we may set

$$\mathcal{G}_{\mathcal{C}}(\tau;\varepsilon) \;=\; N \min(x, A^*(\tau)) \; , \tag{76}$$

where x is set to be the smallest number satisfying the inequality

$$\left(\frac{\rho\tau}{x} \right)^{\frac{x}{A^*(\tau)}} \left(\frac{A^*(\tau) - \rho\tau}{A^*(\tau) - x} \right)^{1 - \frac{x}{A^*(\tau)}} \;\leq\; \varepsilon^{1/N} \; . \tag{77}$$

It can be verified that for N sufficiently large, this bound matches closely the CLT bound of Eqn. (68).

Remark: For deterministic envelopes with a peak-rate constraint $A^*(\tau) \leq P\tau$, the expressions for $\mathcal{G}_{\mathcal{C}}$ in Eqn. (76) and Eqn. (68) describe lines, with slopes which depend on ρ, P, N, and ε. In other words, the arrivals $A_{\mathcal{C}}(t, t+\tau)$ satisfy, with probability at least $1-\varepsilon$, again a rate constraint. The new rate differs from the mean rate $N\rho$ by an error of order $\sqrt{N}$ (for fixed values of ρ, P, and ε).

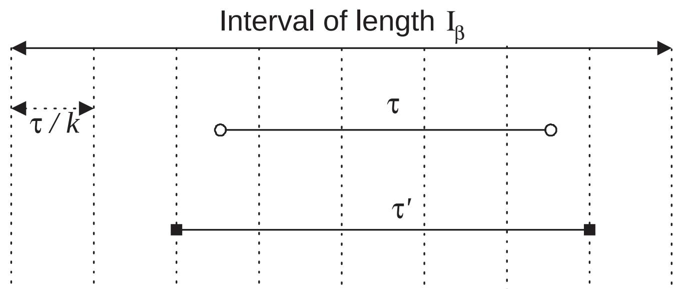

Figure 2: Embedding intervals.

4.3 From Local to Global Effective Envelopes

We use the results from the previous subsection to construct a global effective envelope $\mathcal{H}_\mathcal{C}$ for $A_\mathcal{C}$. The first step is a geometric estimate for the empirical envelope $\mathcal{E}_\mathcal{C}(\tau;\beta)$ for a particular value of τ in terms of the local effective envelope. The second step fixes the value of the global effective envelope for a finite collection of values τ_i. Finally, we obtain the entire envelope by extrapolation. The construction of the global effective envelope requires the choice of a number of parameters. We conclude with a heuristic optimization of these parameters.

Let us define two events:

$$B(x,t,\tau) = \{A_\mathcal{C}(t,t+\tau) \geq Nx\} \ . \tag{78}$$

$$B_\beta(x,\tau) = \{\mathcal{E}_\mathcal{C}(\tau;\beta) \geq Nx\} \ . \tag{79}$$

where I_β is an interval of length β. The event $B(x,t,\tau)$ occurs if the arrivals in the specific time interval $[t,t+\tau]$ exceed Nx, while $B_\beta(x,\tau)$ occurs if there is some interval of length τ in the interval I_β where the arrivals exceed Nx.

With Eqn. (75), we have a bound for the probability of events $B(x,t,\tau)$. The following bound for $B_\beta(x,\tau)$ in terms of $B(x,t,\tau)$ will be used to construct $\mathcal{H}_\mathcal{C}(\tau;\beta,\varepsilon)$ from $\mathcal{G}_\mathcal{C}(\tau;\varepsilon)$.

Lemma 3 Let $k \geq 2$ be a positive integer, I_β an interval of length β, $t \in I_\beta$, and $0 \leq \tau \leq \beta$. Then

$$Pr[B(x,t,\tau)] \ \leq \ Pr[B_\beta(x,\tau)] \ \leq \frac{\beta k}{\tau} Pr[B(x,t,\tau')] \ , \tag{80}$$

with $\tau'/\tau = (k+1)/k$.

Proof: By stationarity, we may assume that $I_\beta = [0,\beta]$ and $t = 0$. The left inequality holds by definition, since $B(x,0,\tau) \subseteq B_\beta(x,\tau)$. To see the inequality on the right, let $t_i = i\tau/k$ ($i = 0,\ldots,\lceil \beta k/\tau \rceil$), and consider the intervals $I_i = [t_i, t_{i+k+1}]$ of length $\tau' = \frac{k+1}{k}\tau$ for $i = 1,\ldots,\lceil (\beta - \tau)k/\tau \rceil$ (all but possibly the last are subintervals of $[0,\beta]$.) See Figure 2 for an illustration of this construction. Clearly, every subinterval of length τ in I_β is contained in at least one of the I_i. The claim now follows with stationarity. □

Lemma 3 provides a bound on arrivals in all subintervals of length τ in I_β. One of its implications is that for every value of τ,

$$Pr\left[\mathcal{E}_\mathcal{C}(\tau;\beta) \geq \mathcal{G}_\mathcal{C}\big(\frac{k+1}{k}\tau;\varepsilon\big)\right] \;\leq\; \frac{\beta k}{\tau}\varepsilon \;, \tag{81}$$

where $\mathcal{E}_\mathcal{C}(.;\beta)$ is the empirical envelope, and $\mathcal{G}_\mathcal{C}(\tau,\varepsilon)$ is any local effective envelope.

Constructing a finite number of values for $\mathcal{H}_\mathcal{C}^\beta$: We next assign a finite number of values for $\mathcal{H}_\mathcal{C}(.;\beta,\varepsilon)$: Pick a collection of values τ_i and k_i $(i = 1,\ldots,n)$ and define

$$\mathcal{H}_\mathcal{C}(\tau_i;\beta,\varepsilon) \;=\; \mathcal{G}_\mathcal{C}(\tau_i';\varepsilon') \;, \tag{82}$$

where

$$\tau_i' = \frac{k_i+1}{k_i}\tau_i \quad \text{and} \quad \varepsilon' = \varepsilon\left(\sum_{i=1}^{n}\frac{\beta k_i}{\tau_i}\right)^{-1} \;. \tag{83}$$

To justify this construction, note that by Eqn. (81) we have

$$Pr\Big[\exists i:\; \mathcal{E}_\mathcal{C}(\tau_i;\beta) \geq \mathcal{G}_\mathcal{C}(\tau_i',\varepsilon')\Big] \;\leq\; \sum_{i=1}^{n}\frac{\beta k_i}{\tau_i}\varepsilon' \tag{84}$$

$$\leq\; \varepsilon \;. \tag{85}$$

Interpolation: We define a first approximation to $\mathcal{H}_\mathcal{C}(\tau;\beta,\varepsilon)$ by

$$f(\tau) \;=\; \min\{NA^*(\tau),\; \mathcal{H}_\mathcal{C}(\tau_i;\beta,\varepsilon)\} \qquad \tau \in [\tau_{i-1},\tau_i) \;, \tag{86}$$

where the values $\mathcal{H}_\mathcal{C}(\tau_i;\beta,\varepsilon)$ are given by Eqn. (82). Since the empirical envelope $\mathcal{E}_\mathcal{C}(\tau;\beta)$ increases with τ and cannot exceed $NA^*(\tau)$ by assumption (A2), we see that

$$Pr\left[\exists\tau \in I_\beta:\; \mathcal{E}_\mathcal{C}(\tau;\beta) \geq f(\tau)\right] \;\leq\; \varepsilon \;. \tag{87}$$

Since $\mathcal{E}_\mathcal{C}$ is subadditive, we may take $\mathcal{H}_\mathcal{C}(\tau;\beta,\varepsilon)$ to be the largest subadditive function which does not exceed $f(\tau)$, in formulas:

$$\mathcal{H}_\mathcal{C}(\tau;\beta,\varepsilon) \;=\; \inf_{\sum\theta_i=\tau}\sum_i f(\theta_i) \;. \tag{88}$$

Heuristic optimization: Since there exists no universal "best" global effective envelope, it is clearly impossible to make an optimal choice for the values of τ_i and k_i in Eqs. (82) and (83). It is, however, possible to make good choices, which lead to global effective envelopes that approximate the given local effective envelope well, at least when ε is sufficiently small.

We will discuss only the case where the traffic regulators satisfy a peak rate constraint with peak rate P and average rate ρ. Our goal is to find an effective envelope satisfying again a rate constraint

$$\mathcal{H}_{\mathcal{C}}(\tau;\beta,\varepsilon) \;\leq\; N\alpha\tau \tag{89}$$

with $\alpha < P$ as close to ρ as possible. In this case we set

$$k_i = k \qquad \text{and} \qquad \tau_i = \gamma^i \tau_o \qquad i = 1,\ldots,n \;, \tag{90}$$

where τ_o is a small number, and

$$\gamma \;\approx\; 1 + \frac{1}{k+1} \;, \tag{91}$$

$$k \;\approx\; z\left(z + \sqrt{N}\,\frac{\rho\tau}{\hat{\mathbf{s}}}\right) = z\left(z + \frac{\sqrt{N}}{\sqrt{P/\rho - 1}}\right) \;, \tag{92}$$

where z is defined by $1 - \Phi(z) = \varepsilon$, and

$$\hat{\mathbf{s}} \;=\; \rho\tau\sqrt{P/\rho - 1} \tag{93}$$

in accordance with Eqn. (62). This choice of the τ_i and k_i is used in all our numerical simulations. We motivate the choice below by appealing to the Central Limit Theorem.

Let us for the moment accept Eqn. (90), fix τ_o and β, and optimize over the parameters k and γ. Eqn. (90) guarantees that

$$\mathcal{H}_{\mathcal{C}}(\tau;\beta,\varepsilon) \;\leq\; \mathcal{G}_{\mathcal{C}}\left(\frac{k+1}{k}\;\gamma\;\tau;\varepsilon'\right) \tag{94}$$

for all $\tau \in [\tau_o, \beta]$, where, by Eqn. (83),

$$\varepsilon' \;=\; \frac{\tau_o(\gamma - 1)}{\beta k}\cdot\varepsilon \;. \tag{95}$$

We estimate

$$Pr\left[\exists\tau \in [\tau_o,\beta] :\; \mathcal{E}_{\mathcal{C}}^{\beta}(\tau) \;\geq\; N\alpha\tau\right] \;\leq\; Pr\left[\exists i :\; \mathcal{E}_{\mathcal{C}}^{\beta}(\tau_i) \;\geq\; N\alpha\tau_i/\gamma\right] \tag{96}$$

$$\leq\; \sum_{i=1}^{n}\frac{\beta k}{\tau_o\gamma^i}Pr\left[A_{\mathcal{C}}\left(0, \frac{k+1}{k}\tau_i\right) \leq N\alpha\tau_i/\gamma\right] \tag{97}$$

$$\approx\; \frac{\beta k}{\tau_o(\gamma-1)}\left(1 - \Phi\left(\sqrt{N}\,\frac{\alpha k/(\gamma(k+1)) - \rho}{\rho\sqrt{P/\rho - 1}}\right)\right) \tag{98}$$

where the first step follows from monotonicity, the second step uses Lemma 3, and the third step invokes the Central Limit Theorem, and a simple estimate for the geometric series.

We next solve for α in the equation

$$\varepsilon \;=\; \frac{\beta k}{\tau_o(\gamma-1)}\left(1 - \Phi\left(\sqrt{N}\,\frac{\alpha k/(\gamma(k+1)) - \rho}{\rho\sqrt{P/\rho - 1}}\right)\right) \;. \tag{99}$$

For every integer of k and $\gamma > 1$, an (approximate) envelope is given by

$$\mathcal{H}_{\mathcal{C}}(\tau;\beta,\varepsilon) \approx N\alpha\tau \tag{100}$$

$$= \frac{(k+1)\gamma}{k}\left(N\rho\tau + z'\sqrt{N}\rho\tau\sqrt{P/\rho - 1}\right) , \tag{101}$$

where ε' is given by Eqn. (95), and $1-\Phi(z') = \varepsilon'$. This approximation is valid for τ in the interval $[\tau_o, \beta]$.

Our goal is to choose k and γ so that the right hand side of Eqn. (101) is as small as possible. The difficulty is that z' depends on the choice of k and γ. We can achieve our goal by minimizing instead the right hand of Eqn. (98). It is easy to see that the minimum value is achieved for some finite positive value of k and γ. Using the approximation $1-\Phi(z) \approx z^{-1}\phi(z)$, where ϕ is the density of standard normal normal distribution, differentiating with respect to k and γ, and solving approximately for the critical values, we see that the minimum is attained at a point satisfying

$$\gamma \approx 1 + \frac{1}{k+1} , \tag{102}$$

$$k \approx z'\left(z' + \frac{\sqrt{N}}{\sqrt{P/\rho - 1}}\right) , \tag{103}$$

where z' is defined by $1-\Phi(z') = \varepsilon'$. Approximating z' by z we arrive at the conditions in Eqs. (91) and (92).

We turn to the basic choice made in Eqn. (90). The fact that the right hand side of Eqn (98), and hence k and γ determined by either Eqs. (102)–(103) or Eqs. (91)–(92), does not depend on τ_o and β indicates that optimal choices of k_i and $\gamma_i := \tau_i/\tau_{i-1}$ do not depend on the size of τ_i, provided that τ_i is small enough to lie in the region where $A^*(\tau_i) = P\tau_i$.

The above heuristic method to choose parameters for the global effective envelope is for peak-rate constrained leaky buckets with parameters (P, σ, ρ). We propose the following analogous heuristic optimization for general regulators with subadditive deterministic envelope A^*.

Assuming as above that β is given, we set τ_o to be a small number, and set z such that $1-\Phi(z) = \varepsilon$. We replace Eqs. (90) – (92) by the following method to recursively determine the k_i, γ_i, and τ_i for $1 \le i \le n$, where n is the first time such that $\tau_n \ge \beta$.

$$k_i = z\left(z + \frac{\rho\tau_{i-1}}{\hat{\mathbf{s}}_{i-1}}\sqrt{N}\right) \tag{104}$$

$$= z\left(z + \sqrt{N}\left(\frac{A^*(\tau_{i-1})}{\rho\tau_{i-1}} - 1\right)^{-1/2}\right) , \tag{105}$$

where $\hat{\mathbf{s}}_i$ is as given in Eqn. (62) (the subscript i in $\hat{\mathbf{s}}_i$ corresponds to τ_i), and

$$\gamma_i = 1 + \frac{1}{k_i + 1} , \tag{106}$$

$$\tau_i = \gamma_i\tau_{i-1} \tag{107}$$

When the algorithm terminates after n steps, we obtain ε' as in Eqn. (83).

5 Related Work

The literature on statistical services and statistical multiplexing in Quality of Service networks is extensive and a full discussion is beyond the scope of this paper. Excellent discussions on the state of the art of statistical multiplexing can be found in [15, 26, 28].

Here, we will only discuss two groups of prior work on statistical multiplexing which we regard as particularly relevant to this paper. The first group studies the statistical multiplexing gain of statistically independent, regulated, adversarial traffic at a buffered multiplexer with fluid flow service. The second group of work are extensions of results on deterministic QoS [6, 7] to a probabilistic framework.

5.1 Regulated, Adversarial Traffic at a Multiplexer

Several researchers have studied the multiplexing gain of statistically independent, regulated, adversarial traffic at a buffered multiplexer, such as the one shown in Figure 1, where each flow is allocated a fixed amount of link bandwidth and buffer capacity. Such a system is sometimes referred to as a virtual buffer/virtual trunk system. In these studies, it is assumed that traffic is served in a fluid flow fashion, that is, the multiplexer can simultaneously transmit traffic from all flows at their respective allocated rates. These works generally do not consider scheduling algorithms at the multiplexer. The allocated rate and the buffer capacity for a flow is selected such that the probability of losses due to buffer overflows is smaller than some small number ε. The allocated rate is sometimes referred to as the *effective bandwidth* of a flow.

Elwalid, Mitra, and Wentworth [10] consider a virtual buffer/virtual trunk systems with flows which are regulated by peak-rate constrained leaky buckets with parameters (P, σ, ρ) and deterministic envelopes as given in Eqn. (3). The analysis of the system proceeds by reducing a two resource system (buffer and bandwidth) to a single resource system (bandwidth), and then analyze the (bufferless) single resource system. The adversarial traffic pattern used in [10] is a periodic on-off source, which is known to maximize the overflow probability in a bufferless multiplexer. In [9], the solution approach of [10] is applied to the GPS [22] scheduling algorithm.

LoPresti, Zhang, Towsley, and Kurose [19] consider the same system. They analyze the virtual buffer/virtual trunk system by transforming it into two systems, each with one resource: a bufferless trunk and a buffer with no server. Losses occur if the demand for bandwidth exceeds the capacity of the trunk, or the demand for storage space exceeds the size of the buffer. Rajagopal, Reisslein, Ross [23] study the same system as [19], and consider more general traffic regulators with deterministic envelopes as given in Eqn. (4).

The question of the adversarial traffic pattern at a buffered multiplexer with (P, σ, ρ) regulators has received a lot of interest. As suggested in [8] and by others, and supported by numerical data presented in [20], on-off traffic sources are adversarial for bufferless multiplexers, but not for buffered multiplexers. Kesisidis and Konstantopoulos [11, 12] address the problem of finding explicit expressions for the adversarial traffic patterns at a buffered multiplexers. In [11], the authors derive the adversarial traffic pattern for a single flow, and in [12] the authors analyze the queue length distribution for multiple flows. In all cases, the adversarial traffic pattern is shown to be periodic, with multiple on-phases and different rates in each 'on' phase.

Rajagopal, Reisslein, Ross [24, 25] investigate the statistical multiplexing gain for a specific type of multiplexer, which consists of one dedicated buffer for each flow and a bufferless multiplexer. All arrivals on a flow are stored in the dedicated buffer for this flow. The output rate of the buffer of a flow is set to a fixed rate, such that no traffic from the flow experiences a delay violation. The output of all buffers is multiplexed at a bufferless multiplexer. Losses eccur at the bufferless multiplexer when the aggregate output from the buffers exceeds the link capacity.

Our bounds in Section 4 have similar goals as the studies cited above, in that they investigate statistical multiplexing for statistically independent, regulated, and adversarial traffic. On the other hand, our approach deviates from the above papers in several ways. For example, different from most studies mentioned above, we do not use a particular adversarial traffic pattern in our analysis. In fact, for the scheduling algorithms considered in this paper (e.g., EDF) and for the general class of traffic regulators with subadditive deterministic envelope functions, an explicit derivation of an adversarial pattern may be a formidable, if not, impossible task. Another difference of our work is that we do not analyze a buffered multiplexer with a fluid flow server. Rather, we consider the scheduling algorithm of the multiplexer throughout the analysis.

5.2 Probabilistic Extensions of Deterministic Calculus

Our general approach can be characterized as generalizing schedulability conditions for a deterministic service to a probabilistic framework (see Section 3). Several researchers before us made probabilistic extensions to deterministic service models. In particular, Cruz's deterministic network calculus [6, 7] has inspired several researchers to define a statistical service within a probabilistic interpretation of the network calculus.

Chang [4] derives probabilistic bounds for the delays in a multiplexer with a shared buffer, as discussed in the previous subsection, i.e., each flow is served in a fluid-flow fashion at an allocated rate. The arrivals on a flow j to the buffer in time interval $[t_1, t_2)$, $A_j(t_1, t_2)$, are assumed to be bounded by a function $\hat{A}_j(s\ ;.)$ such that

$$\frac{1}{s} \ln E[e^{A_j(t,t+\tau)s}] \leq \hat{A}_j(s, \tau) \ .$$

Thus, $\hat{A}_j(s,.)$ is a bound on the moment generating function of the arrival function $A_j(t, t+\tau)$. For a certain class of bounds $\hat{A}_j(s,.)$, i.e., bounds of the form

$$\hat{A}_j(s, \tau) \ = \ \hat{\sigma_j}(s) + \hat{\rho}_j(s) \cdot \tau \ ,$$

Chang derives bounds on the loss probability. Even though the notation and the formalism in [4] bears similarity to those used in this paper, the system studied by Chang is quite different from our work. A key difference is that we do not assume that there exists an *a priori* bound on the moment generating function of the arrivals. Instead, we assume a deterministic traffic regulator for each flow (as shown in Figure 1), and, with this constraint, *derive* bounds on the moment generating functions of the arrivals.

Yaron and Sidi [31] present another probabilistic extension of Cruz's deterministic network calculus. They consider stochastic traffic sources which satisfy the constraint

$$Pr\left[A_j(t, t+\tau) \geq \rho_j(\tau) + \sigma_j\right] \leq B\ e^{-\alpha\sigma} ,$$

where α and B are constants. These source are referred to as having *exponentially bounded burstiness*. Yaron and Sidi prove that if the incoming flows to a multiplexer satisfy exponentially bounded burstiness constraints, then the output of the multiplexer has exponentially bounded burstiness. The work was recently extended in [29] to more general bounding functions.

Kurose [16] explores bounds on the distribution of the delay and buffer occupancy of a flow in a network environment. The traffic on a flow is determined by a family of random variables, where each random variable denotes the traffic over a time interval of a certain length. (The family of random variables can be thought of as a family of local effective envelopes.) For a link with FIFO scheduling and making an assumption analogous to that in Eqn. (24), [16] derives an expression for the distribution of the delays at a single node. Kurose also provides bounds on the output of a node, and, thus, can obtain bounds for networks with multiple nodes. Zhang and Knightly [34] extend this work, and, for specific types of random variables, i.e., Markov modulated fluid flow sources or sources with a binomial arrival distribution, obtain probabilities for delay bound violations at a rate-controlled SP scheduler.[4] Both [16] and [34] calculate the arrival distribution of aggregate sources directly, without resorting to large deviations results.

Knightly [13, 14] extends the approach of [16] and [34] by introducing a characterization of flow arrivals using first and second moment information on the sources. The notion of a *rate-variance envelope* $RV_j(\tau)$ is introduced as a function which describes the variance of the arrivals of a flow j over a time period of length τ, as defined in Eqn. (70). In [13], arrivals on a flow are assumed to be characterized by the rate-variance envelope and the long-term arrival rate ρ_j (see Eqn. (5)). Then, applying the CLT (Central Limit Theorem), a bound for the probability of a delay bound violation is derived for an SP scheduler. In [14], the same framework is used to address bounds on the rate-variance envelope for regulated, adversarial traffic sources. Recall from Section 4.2.1, that these bounds on $RV(\tau)$ can also be obtained by using the second moment bound from Theorem 2 from Section 4.

A common characteristic of the works in [13, 14, 16, 34] is that they consider bounds of the form shown in Eqn. (23). Hence, all results make an assumption as in Eqn. (24).[5]

Our work can be viewed as a generalized framework for the approach pursued in [13, 14, 16, 34]. The generalization is done in several directions. First, we consider more general scheduling algorithms. Also, we derive different types of bounds, including those, which allow us to not

[4] A rate-controlled SP (RCSP) scheduler is proposed in [32] performs per-flow shaping of traffic at each node, and, hence, is not workconserving.

[5] Chang [4] makes a similar assumpition in his work. The assumption occurs between Eqns. (55) and (56) of [4].

depend on assumptions as in Eqn. (24). Finally, we use a formalism which enables us to derive schedulability conditions for deterministic and statistical QoS guarantees in a uniform fashion.

6 Evaluation

In this section, we evaluate the effective envelope approach, using the schedulability conditions from Section 3 and the bounds derived in Section 4. The key criteria for evaluation is the amount of traffic which can be provisioned on a link with QoS guarantees.

As benchmarks for statistical QoS provisioning we consider the following non-statistical methods:

- **Peak Rate Allocation:** It is well-known that peak rate allocation provides deterministic QoS guarantees, but is an inefficient method for achieving QoS. The number of connections that can be supported with a peak rate allocation serve as a lower bound for any method for provisioning QoS.
- **Deterministic Allocation:** We use admission control tests for deterministic QoS from Eqn. (22). The admissible traffic is dependent on the scheduling algorithm.
- **Average Rate Allocation:** Since average rate allocation only guarantees finiteness of delays and average throughput, the number of connections that can be supported with an average rate allocation are an upper bound for any method for provisioning QoS.

We will evaluate the methods for provisioning statistical QoS which are presented in this paper.

- **Local Effective Envelope:** We use Eqn. (25) to determine admissibility. We evaluate the quality of the following two bounds, derived in Section 4:
  - **Local Effective Envelope (CB)**: Uses the bound from Eqn. (77), obtained with the Chernoff bound.
  - **Local Effective Envelope (CLT)**: Uses the bound from Eqn. (68), obtained with the Central Limit Theorem.
    Recall from our discussion in Section 4 that the *local effective envelope (CLT)* results are equivalent to the rate-variance envelope method described in [14].
- **Global Effective Envelope:** We use the procedure, including the heuristic optimization, developed in Section 4.3 to determine admissibility. We select β such that it is larger than the longest busy period (see Footnote 1). In Eqn. (82), we use the local effective envelope (CB) rather than the corresponding CLT bound, since the latter would yield only approximate bounds.

We compare our results with the effective bandwidth approach for regulated adversarial traffic from the literature:

- **Effective Bandwidth [10, 19, 23]**:[6] The effective bandwidth approach assigns to each flow a fixed capacity, the *effective bandwidth*, and assumes that each flow is serviced at a rate which corresponds to the effective bandwidth. General scheduling algorithms are not considered (with the exception of [9], which extends the results in [10] approach to GPS scheduling.). Thus, when we compare effective bandwidth results to our approach, we will assume FIFO scheduling for the effective envelopes.

  The delay bounds is indirectly derived from the buffer size. We set the delay bound d to $d = B/C$, where B is the buffer size at the scheduler and C is the transmission rate of the link.

  In our examples, we include the following results on effective bandwidth:

  - **EB-EMW**: This is the result from the classical paper by Elwalid/Mitra/Wentworth (Eqn. (39) in [10]).
  - **EB-RRR**: We use Eqn. (9) from [23] by Rajagopal/Reisslein/Ross which presents an improvement to the *EB-EMW* result.

- **Bufferless Multiplexer (Bufferless MUX) [25]:** In [25], the multiplexing gain of regulated adversarial traffic is analyzed for a particular switch architecture, where arriving traffic on a flow is shaped at a dedicated buffer for this flow. The output rate of each buffer is set to a fixed rate, such that no traffic experiences a violation of its delay bound in the buffer. The output from the buffers is multiplexed on the link, however, there are no buffers at the multiplexer (hence, the term *Bufferless multiplexer*). Thus, if the aggregate output from the buffers exceeds the link capacity, traffic is dropped.

In all our experiments, we consider traffic regulators which are obtained from peak rate controlled leaky buckets with deterministic envelopes as given in Eqn. (3).[7] In all experiments, we consider a 45 Mbps link ($C = 45$ Mbps), and we consider two traffic classes. The traffic parameters for single flows in the classes are as follows:

Class	Peak Rate P (Mbps)	Mean Rate ρ (Mbps)	Burst Size σ (bits)
1	1.5	0.15	95400
2	6.0	0.15	10345

The parameters are selected so as to match (approximately) the examples presented in [10, 23].

In this section we present four sets of examples. In the first example, we compare the deterministic envelopes with our bounds for the local and global effective envelopes for different sets of

[6] The cited works calculate effective bandwidth for regulated adversarial sources. The complete literature on effective bandwidth is much more extensive.

[7] Most of the methods listed here can work with more complex regulators. However, since peak-rate enforced leaky buckets are widely used in practice, [1, 2], they serve as good benchmarks.

sources. In the second example, we compare the maximum number of admissible flows in a FIFO scheduler for given delay bounds d and delay-violation probability ε. In the third example, we show how the statistical multiplexing gain increases at higher link capacities. In the fourth example, we investigate the case of heterogeneous traffic with different QoS requirements, and we compare the admissible regions for different scheduling algorithms (SP, EDF).

6.1 Example 1: Comparison of Envelope Functions

In the first example, we study the shape of local and global effective envelopes for homogeneous sets of flows as functions of the lengths of time intervals. The envelopes are compared to the deterministic envelope $A_j^*(\tau) = \min\{P_j, \tau, \sigma_j + \rho_j, \tau\}$, to the peak rate function $P_j\,\tau$, and to the average rate function $\rho_j\,\tau$. In our graphs, we plot the amount of traffic per flow for the various envelopes (e.g., we present $\sum_{j\in\mathcal{C}} \mathcal{G}_j(\tau;\varepsilon)/N$).

Figures 3(a) and 3(b) show the results for multiplexed flows from Class 1 and Class 2, respectively. We set $\varepsilon = 10^{-6}$ for all envelopes. By depicting the amount of traffic per flow for different numbers of flows (N denotes the number of flows), we can observe how the statistical multiplexing gain increases with the number of flows.

The first observation to be made is that local and global effective envelopes are much smaller than the deterministic envelope or the peak rate. Another observation is that, for a fixed number of flows N, the effective envelope is larger than the local effective envelopes, and the local effective envelope bound is smaller when using CLT (central limit theorem), as compared to CB (Chernoff bound). Figure 3 also shows that the difference between local and global effective envelopes narrows as the number of flows N is increased. Note that in Figure 3(b) local effective envelopes are identical for CB and CLT bounds for $N = 10000$.

In Figure 4 we depict the sensitivity of the effective envelopes to the selection of the parameter ε. We use the same parameters as before. We fix the value for the number of flows to $N = 1000$, and we show the effective envelopes for $\varepsilon = 10^{-3}, 10^{-6}$, and 10^{-9}. Figure 4 shows that the effective envelopes are not very sensitive to variations of the parameter ε.

In Figure 5 we show how the effective envelopes vary if the number of flows N is increased. We use the same parameters as before, but only consider the values of the envelopes for the time interval $\tau = 50$ ms. For this value of τ, Figures 5(a) and (b) show, for flows from Class 1 and Class 2, respectively, the values of the rates $\mathcal{G}_{\mathcal{C}}(\tau;\varepsilon)/(N\tau)$ and $\mathcal{H}_{\mathcal{C}}(\tau;\varepsilon)/(N\tau)$, as the number of flows N is varied. For comparison, we include the peak and average rates into the graph. There are three noteworthy observations to be made. First, as the number of flows N is increased, the effective envelopes, both local and global, are close to the average traffic rate. Second, the difference between the two local effective envelopes diminishes when N is large. Third, for large values of N the difference between the local effective envelopes and the global effective envelope is quite small.

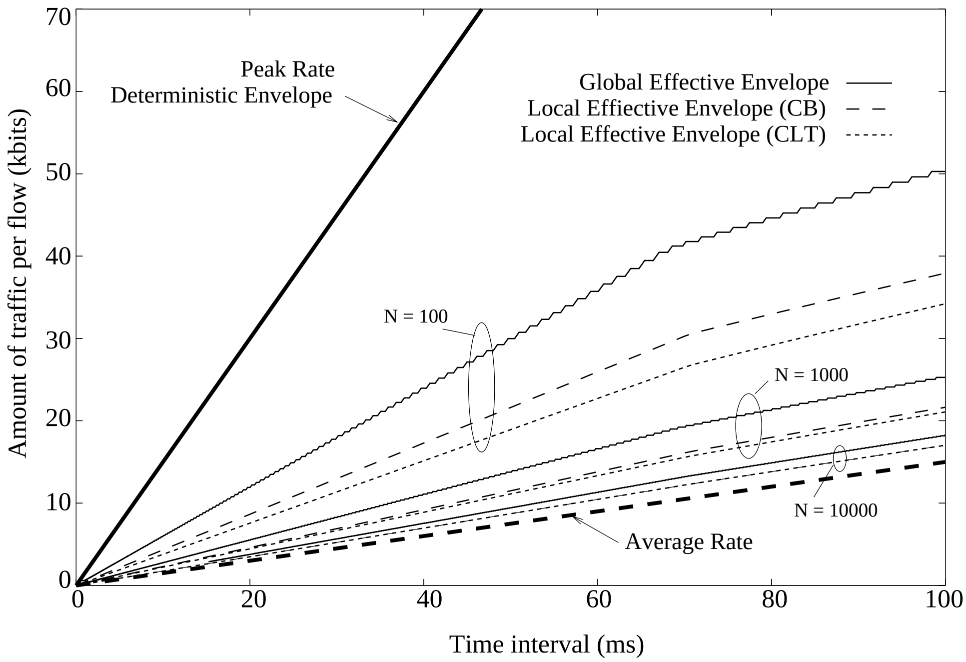

(a) Class 1.

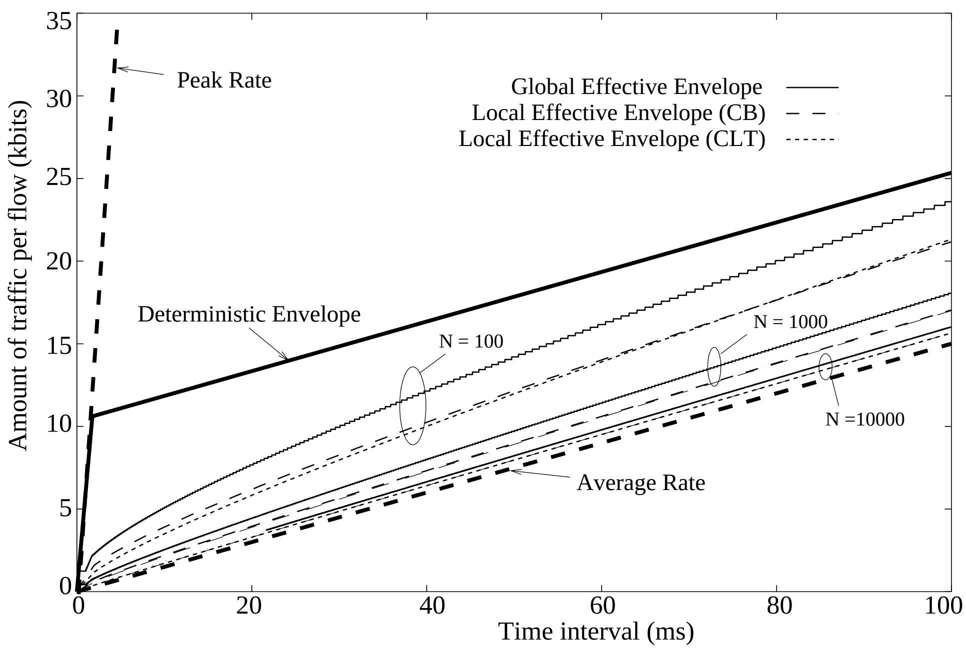

(b) Class 2.

Figure 3: Example 1: Comparison of envelopes for $\tau \leq 100\ ms$, $\varepsilon = 10^{-6}$, and for number of flows $N = 100, 1000, 10000$.

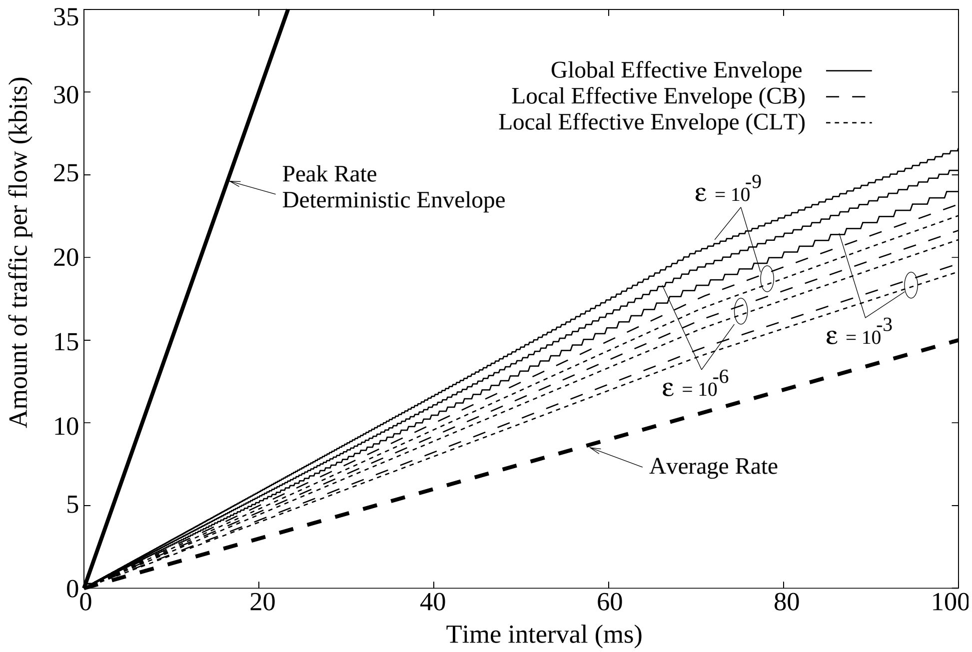

(a) Class 1.

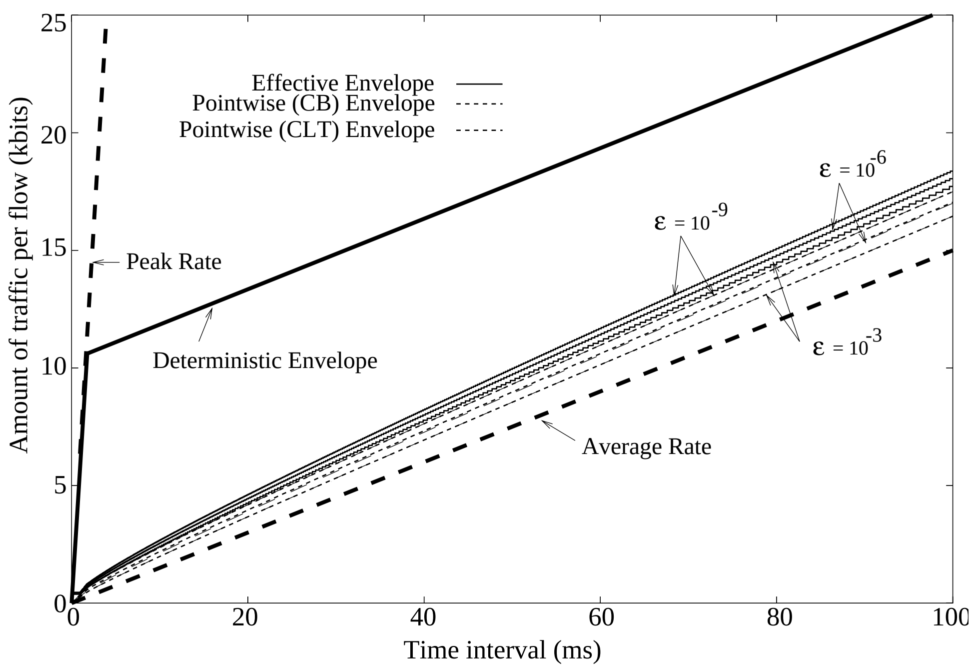

(b) Class 2.

Figure 4: Example 1: Comparison of envelopes for $\tau \leq 100\ ms$, number of flows $N = 1000$ and for different values of ε.

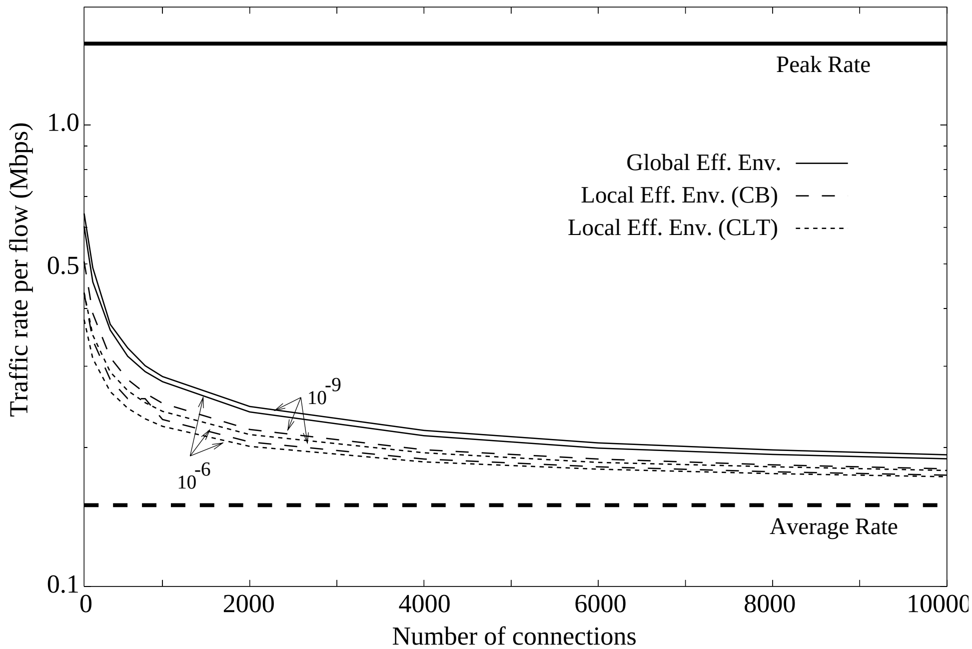

(a) Class 1.

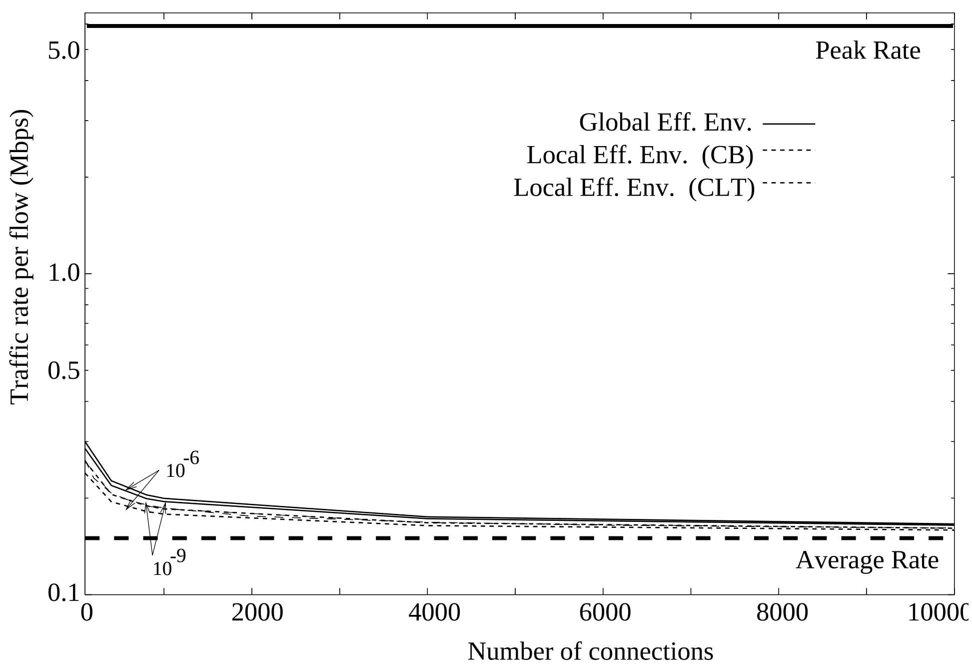

(b) Class 2.

Figure 5: Example 1: Traffic rates $\mathcal{G}_{\mathcal{C}}(\tau;\varepsilon)/(N\tau)$ and $\mathcal{H}_{\mathcal{C}}(\tau;\varepsilon)/(N\tau)$ for $\tau = 50$ ms and $\varepsilon = 10^{-6}$ or 10^{-9} .

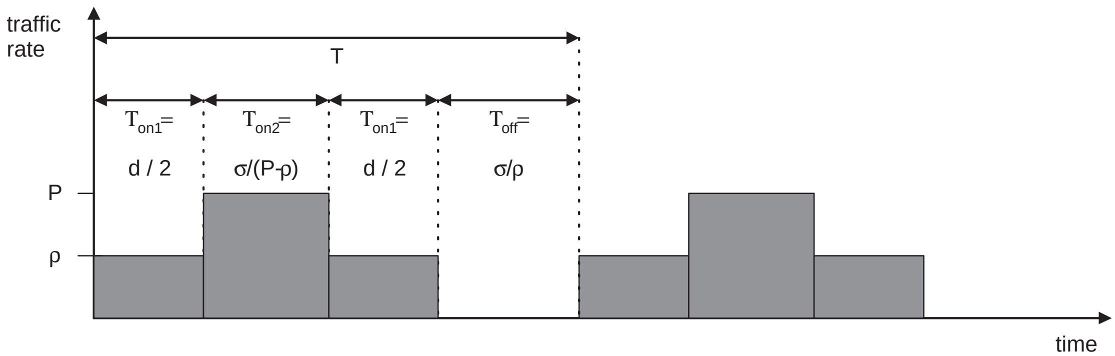

Figure 6: Traffic pattern for (P, ρ, σ) flows used in simulations.

6.2 Example 2: Admissible Region for Homogeneous Flows

In this example, we investigate the number of flows admitted by various admission control methods for guaranteeing QoS at a link with a FIFO scheduler. We assume that flows are homogeneous, that is, all flows belong to a single class. We compare the admissible regions[8] of the the local and global effective envelopes, to those of the effective bandwidth techniques (both *EB-EMW* and *EB-RRR*), the bufferless Multiplexer (*bufferless MUX*) and to deterministic QoS guarantees.

We compare the results with those obtained from a discrete event simulation. For the simulation, we take a pattern which we expect, based on the simulations in [20], to be close to an adversarial traffic pattern for peak-rate controlled leaky buckets. However, do not claim that the results from the simulation scenario are the worst possible.

In the simulations, the traffic for a flow with parameters given by (P, ρ, σ), is periodic with a pattern as shown in Figure 6. A flow transmits at the average rate ρ for a duration $T_{on1} = d/2$, where d is the delay bound. Then, the flow transmits at the peak rate P for a duration $T_{on2} = \sigma/(P - \rho)$, followed by another phase of length T_{on1} at which the flo transmits at rate ρ. Then, the source shuts off, waits for a duration $T_{off} = \sigma/\rho$ and then repeats the pattern. The starting time of a pattern of the flows are uniformly and independently chosen over the length of its period. We refer to Appendix B for additional comments of of the simulations.

Figures 7(a) and (b) depict the number of admitted flows as a function of the delay bound. Here, the probability of a violation of QoS guarantees is set to $\varepsilon = 10^{-6}$. The figures show that all methods for statistical QoS admit many more connections than a deterministic admission control test. In both Figures, the effective envelopes (both CLT and CB) are closest to the simulation results. (Once again, we point out that the results using the local effective (CLT) bounds are identical to the rate-variance results presented in [14].) Note, however, that results obtained with local effective envelopes are approximate and are not guaranteed to be upper bounds on the admissible regions.

Comparing the results from effective envelopes to the effective bandwidth results, we observe that the effective envelope methods admits more connections than the effective bandwidth methods if delay bounds are large.

[8]The *admissible region* is the range of parameters which results in a positive admission control decision.

The difference of the admissible regions in Figure 7(a) to those in Figure 7(b) illustrate the high degree to which the size of the admissible region is dependent on the selection of parameters. The lower burst sizes of flows in Class 2 lead to larger admissible regions for all methods. Specifically, notice that deterministic QoS in Figure 7(b) yields similar results to the statistical methods, if the delay bounds are large.

In Figure 8 we show the results for the same experiment as before, however, with ε set to $\varepsilon = 10^{-9}$. A comparison of Figures 7 and 8 indicates that the admissible regions are not very sensitive to variations of ε.

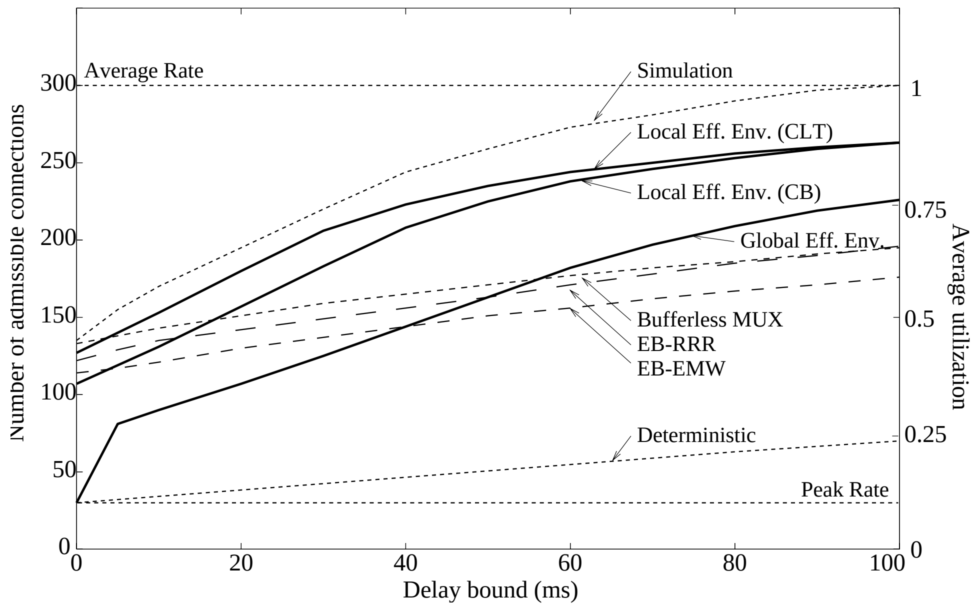

(a) Class 1.

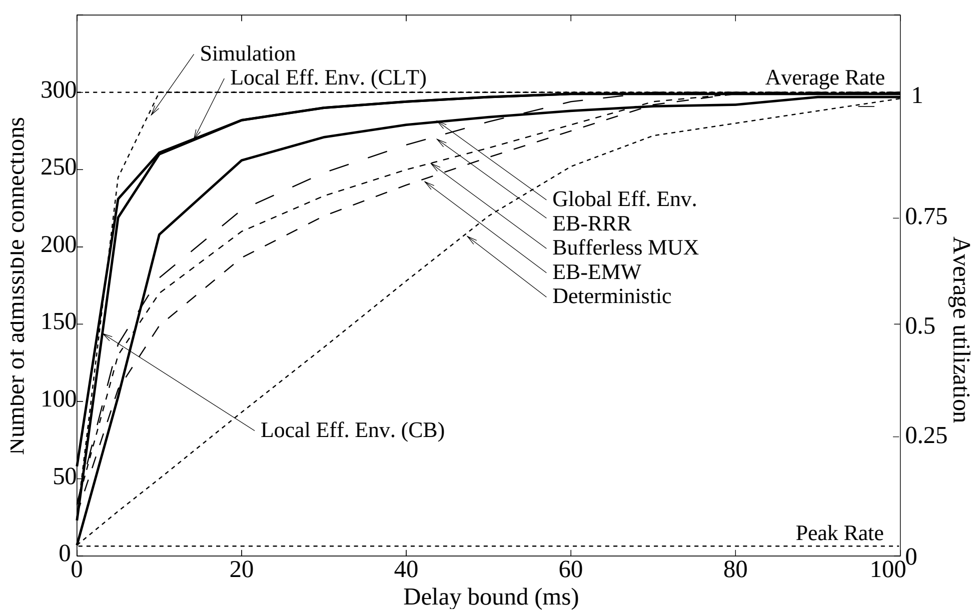

(b) Class 2.

Figure 7: Example 2: Admissible number of flows at a FIFO scheduler for homogeneous flows as a function of delay bounds ($\varepsilon = 10^{-6}$).

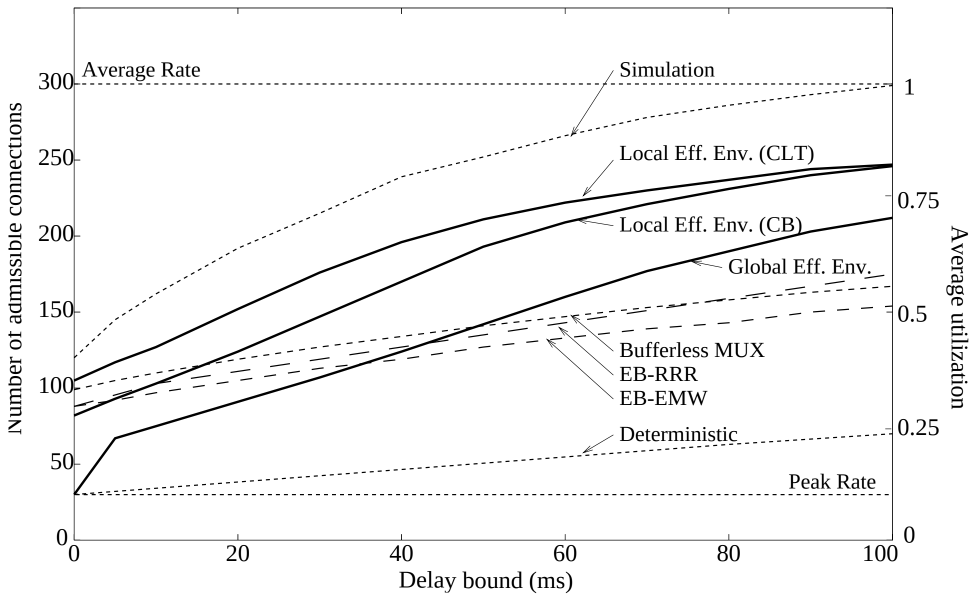

(a) Class 1.

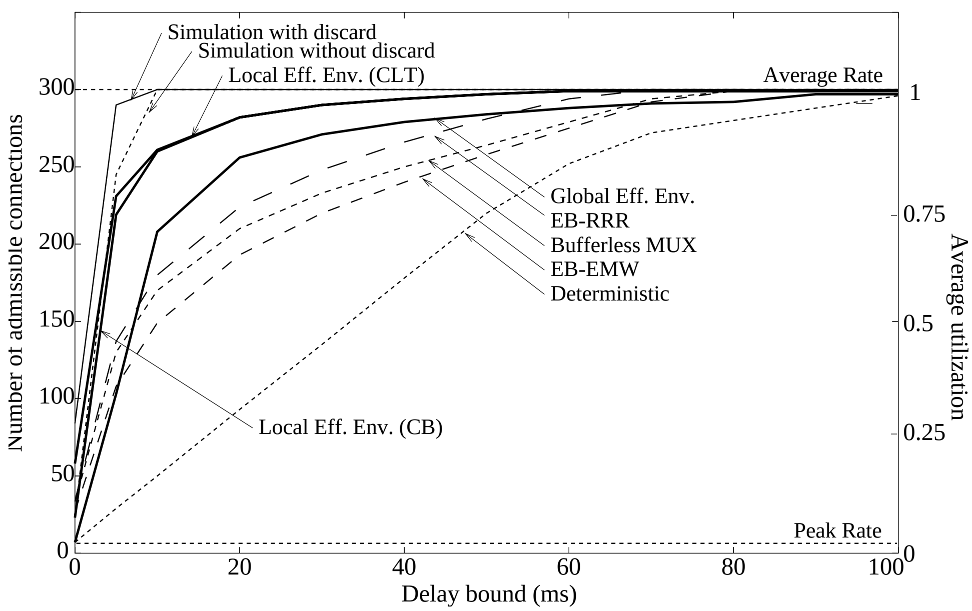

(b) Class 2.

Figure 8: Example 2: Admissible number of flows at a FIFO scheduler for homogeneous flows as a function of delay bounds ($\varepsilon = 10^{-9}$).

6.3 Example 3: Link Utilization for Homogeneous Flows

Our goal in Example 3 is to illustrate how the achievable utilization of a link with a FIFO multiplexer increases as the capacity of the link is increased. We use a similar setup as in Example 2, that is, we consider a FIFO scheduler with homogeneous traffic, where traffic is either from Class 1 or from Class 2. As before, we compare our results to the effective bandwidth methods from [10] and [23], and to the *bufferless MUX* method [25]. We fix the delay bound of traffic to $d = 50$ ms and we use $\varepsilon = 10^{-6}$.

The results of this experiment are shown in Figure 9. We depict the achievable average link utilization as a function of the link capacity. The average achievable link utilization is the sum of the average rates of flows which can be accepted according to a chosen schedulability conditions.

In Figures 9(a) and (b) we observe, for Class 1 and Class 2, respectively, that the achievable average link utilization increases with the link capacity for all considered methods. This illustrates that all methods exploit the statistical multiplexing gain well at high link capacities. Note that the achievable utilization with deterministic QoS is only about 20% in Figure 9(a), but more than 70 % in Figure 9(b). The achievable utilization for deterministic QoS is (almost) constant as a function the link capacity, which illustrates that deterministic QoS does not exploit the statistical multiplexing gain. The effective envelope method performs well over the entire parameter range. Note that the difference between the admissible regions of the local and the global effective envelopes is small at high link capacities.

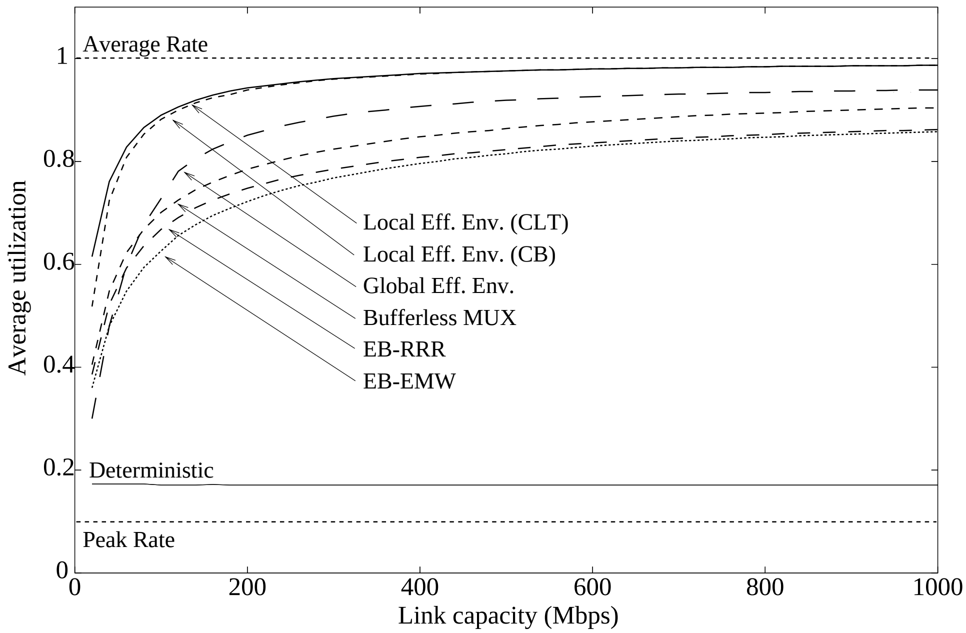

(a) Class 1.

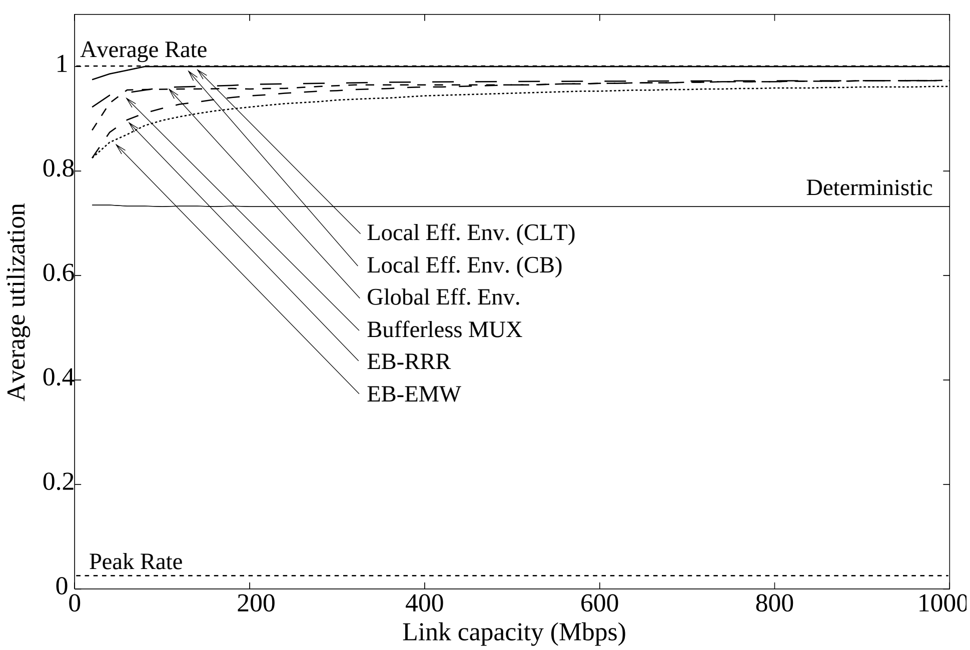

(b) Class 2.

Figure 9: Example 3: Achievable average utilization vs. link capacity, $\varepsilon = 10^{-6}$ and $d = 50$ ms.

6.4 Example 4: Admissible Region for Heterogeneous Flows

In this example we consider different scheduling algorithms with heterogeneous traffic arrivals. As scheduling algorithms, we consider Static Priority (SP) and Earliest-Deadline-First (EDF). For a deterministic service, EDF is optimal, in the sense that the admissible regions with EDF scheduling are maximal [17]. To our knowledge, results for a statistical service (with adversarial traffic) have not been reported for EDF.

In this example, we multiplex a number of flows from Class 1 and from Class 2 on 45 Mbps. We fix the delay bounds such that the delay bound for Class-1 flows is relatively long, $d_1 = 100\ ms$, and the delay bound for Class-2 flows is relatively short, $d_2 = 10\ ms$. For any particular method, we determine the maximum number of Class-1 and Class-2 flows that can be supported simultaneously on the 45 Mbps link.

The result are shown in Figure 10. The plot depicts the admissible region for SP schedulers and EDF scheduling using the results from the local effective envelope, effective envelope, and deterministic envelope approaches, respectively. We also include the admissible regions using the effective bandwidth approaches (*EB-EMW* and *EB-RRR*). Note, however, that effective bandwidth method assume a simple multiplexer (with virtual buffer partitioning (see [9]) and do not account for different scheduling algorithms. The results in Figure 10 show that the difference between SP and EDF schedulers is small in all cases. The effective envelope is, again, more conservative than the local effective envelope method. Figure 10 illustrates that with heterogeneous flows and the effective bandwidth methods (*EB-EMW*, *EB-RRR*) may not perform as well as methods which consider scheduling algorithms.

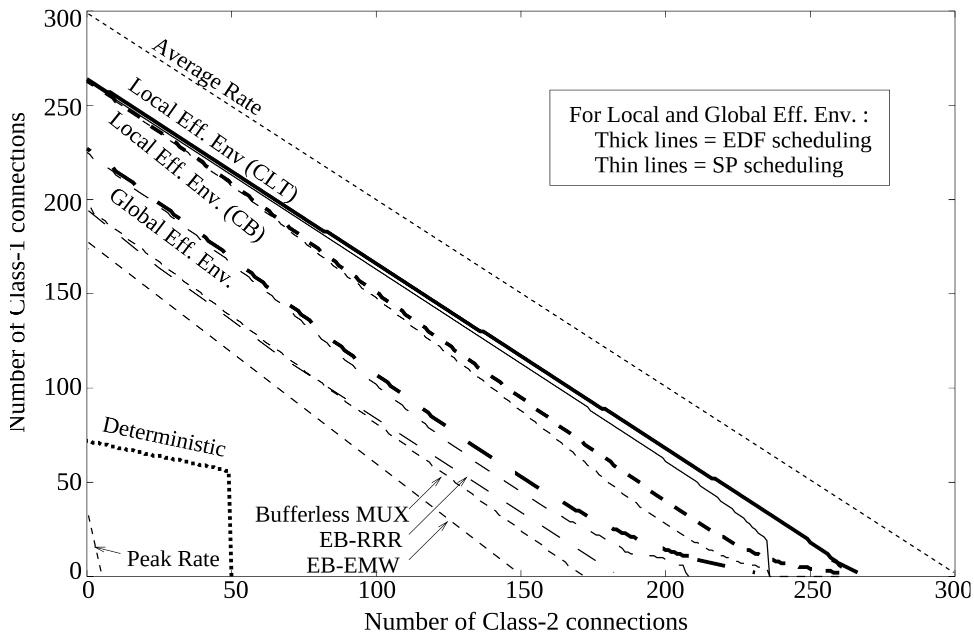

Figure 10: Example 4: Admissible region of multiplexing Class 1 and Class 2 flows with $\varepsilon = 10^{-6}$ and $d_1 = 100$ ms and $d_2 = 10$ ms.

7 Discussion

We have presented new results on evaluating the statistical multiplexing gain for packet scheduling algorithms. We have introduced the notions of local and global effective envelopes, which are, with high probability, bounds on aggregated traffic flows, and we have derived admission control tests for these bounds. We want to conclude with a brief discussion of some issues:

- We have presented two schedulability conditions for a statistical service at a single node. The first condition, which uses local effective envelopes (Section 3.3), provides a guarantee for the probability of a QoS violation for an arbitrary arrival. The second condition, which uses global effective envelopes (Section 3.4), provides a guarantee for the probability of a QoS violation in an arbitrary time interval. Our motivation to introduce a second, more conservative (and possibly less intuitive) condition is motivated by a technical problem which appears in the derivation of the first schedulability condition (see Section 3.3). The technical problem is the assumption in Eqn. (24), which, as we discussed, has been made in several previous papers. Without a verification of this assumption, however, any admission control conditions which applies this assumption may be too optimistic. The condition based on global effective envelopes does not require us to make any assumptions and provides a true lower bound on the admissible region.

  In our numerical examples, for large number of flows and small values of ε, the difference between the admissible regions for local and global effective envelopes may be small. Thus, in many cases, it may not be necessary to make the assumption given in Eqn. (24).

- We believe that our approach to separate the consideration of the service definition (deterministic, statistical), the scheduling algorithm (FIFO, SP, EDF), and the choice of the large deviations result (Central Limit Theorem, Chernoff Bound) may prove to be useful, as it simplifies the task of testing new scheduling algorithm or large deviations results.

- The schedulability conditions derived in Section 2 for deterministic service (Eqn. (22)), and for a statistical service (Eqs. (26) or (29)) have a similar structure. Thus, schedulability conditions which have been derived for a deterministic service can be reused, without modification, for a statistical service if the local or global effective envelopes are available.

- The derivations of the local and global effective envelopes in Section 4 have used novel bounds on the moment generating function and the moments for a set of regulated flows. Our numerical results have shown that our bounds are tight if the number of flows N is large and the probability of QoS violation ε is small.

- Our work does not attempt to explicitly derive an adversarial traffic pattern. Even though recently results on adversarial patterns have been obtained for buffered multiplexers [11, 12, 23], it may not be feasible to derive adversarial traffic patterns for more complex scheduling algorithms. On the other hand, our results show that tight bounds on the admissible regions are attainable even without knowledge of adversarial traffic patterns.

- A few years ago, a study has addressed the question on the fundamental limits of a deterministic service [30]. The findings of [30] were that (a) the admissible region for deterministic QoS is significantly greater than the schedulable region with a peak-rate allocation, and (b) the choice of the packet scheduling algorithm has a significant impact on the size of the admissible region. Since this paper is, to our knowledge, the first to analyze EDF scheduling (which is the optimal scheduling algorithm for deterministic QoS, in the sense that the admissible regions with EDF scheduling is maximal [17]) in the context of a statistical service for regulated, adversarial sources, and which compares the impact of scheduling algorithms in such a setting, the numerical results may shed light on the fundamental limits of a statistical service:

    1. The examples in this paper (in Section 6 and Appendix A) show that the difference of the admissible region of statistical and deterministic is significant, even if ε is selected very small, e.g., $\varepsilon = 10^{-9}$.
    2. The results from Examples 4 and A.4 suggest that the selection of the scheduling algorithm (SP vs. EDF in our case) has a noticeable, but, in relative terms, small impact on the size of the admissible region. It appears that additional numerical data is required to make more conclusive statements on the importance of scheduling algorithms in a statistical service.
    3. The examples in this paper show that, for high-capacity links, the statistical multiplexing gain is significant. The admissible region for a statistical service is sometimes close to that of an average rate allocation. Since an average rate allocation provides an upper bound on the admissible region, the additional gain achievable by improving currently available methods appears marginal.

As directions for future work, the admission control methodology presented in this paper only applies to a single node, and must be extended to a network environment. A main problem is that buffering and scheduling destroys the independence of flows at the output of a node.

References

[1] ATM Forum, ATM Forum Traffic Management Specification Version 4.0, April 1996.

[2] R. Braden, D. Clark, and S. Shenker. Integrated services in the internet architecture: an overview. IETF RFC 1633, July 1994.

[3] E. Castillo. *Extreme Value Theory in Engineering.* Academic Press, 1988.

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

[5] J. Choe and Ness B. Shroff. A central-limit-theorem-based approach for analyizing queue behvior in high-speed network. *IEEE/ACM Transactions on Networking*, 6(5):659–671, October 1998.

[6] R. Cruz. A calculus for network delay, Part I : Network elements in isolation. *IEEE Transaction of Information Theory*, 37(1):114–121, 1991.

[7] R. L. Cruz. A Calculus for Network Delay, Part II: Network Analysis. *IEEE Transactions on Information Theory*, 37(1):132–141, January 1991.

[8] B. T. Doshi. Deterministic rule based traffic descriptors for broadband ISDN: Worst case behavior and connection acceptance control. In *International Teletraffic Congress (ITC)*, pages 591–600, 1994.

[9] A. Elwalid and D. Mitra. Design of generalized processor sharing schedulers which statistically multiplex heterogeneous QoS classes. In *Proceedings of IEEE INFOCOM'99*, pages 1220–1230, New York, March 1999.

[10] 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.

[11] G. Kesidis and T. Konstantopoulos. Extremal shape-controlled traffic patterns in high-speed networks. Technical Report 97-14, ECE Technical Report, University of Waterloo, December 1997.

[12] G. Kesidis and T. Konstantopoulos. Extremal traffic and worst-case performance for queues with shaped arrivals. In *Proceedings of Workshop on Analysis and Simulation of Communication Networks*, Toronto, November 1998.

[13] 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.

[14] E. Knightly. Enforceable quality of service guarantees for bursty traffic streams. In *Proceedings of IEEE INFOCOM'98*, pages 635–642, San Francisco, March 1998.

[15] E. W. Knightly and Ness B. Shroff. Admission control for statistical QoS: Theory and practice. *IEEE Network*, 13(2):20–29, March/April 1999.

[16] J. Kurose. On computing per-session performance bounds in high-speed multi-hop computer networks. In *ACM Sigmetrics'92*, pages 128–139, 1992.

[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] J. Liebeherr and D. E. Wrege. An Efficient Solution to Traffic Characterization of VBR Vid eo in Quality-of-Service Networks. *ACM/Springer Multimedia Systems Journal*, 6(4):271–284, July 1998.

[19] F. LoPresti, Z. Zhang, D. Towsley, and J. Kurose. Source time scale and optimal buffer/bandwidth tradeoff for regulated traffic in an ATM node. In *Proceedings of IEEE INFOCOM'97*, pages 676–683, Kobe, Japan, April 1997.

[20] P. Oechslin. On-off sources and worst case arrival patterns of the leaky bucket. Technical report, University College London, September 1997.

[21] A. Papoulis. *Probability, Random Variables, and Stochastic Processes. 3rd edition.* McGraw Hill, 1991.

[22] A. K. Parekh and R. G. 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] S. Rajagopal, M. Reisslein, and K. W. Ross. Packet multiplexers with adversarial regulated traffic. In *Proceedings of IEEE INFOCOM'98*, pages 347–355, San Francisco, March 1998.

[24] M. Reisslein, K. W. Ross, and S. Rajagopal. Guaranteeing statistical QoS to regulated traffic: The multiple node case. In *Proceedings of 37th IEEE Conference on Decision and Control (CDC)*, pages 531–531, Tampa, December 1998.

[25] M. Reisslein, K. W. Ross, and S. Rajagopal. Guaranteeing statistical QoS to regulated traffic: The single node case. In *Proceedings of IEEE INFOCOM'99*, pages 1061–1062, New York, March 1999.

[26] J. Roberts, U. Mocci, and J. Virtamo (Eds.). *Broadband Network Traffic: Performance Evaluation and Desgin of Broadband Multiservice Networks. Final Report of Action. COST 242. (Lecture Notes in Computer Science. Vol. 1152).* Springer Verlag, 1996.

[27] O. Rose. Statistical properties of MPEG video traffic and their impact on traffic modeling in ATM systems. Technical Report 101, Institute of Computer Science, University of Wurzburg, February 1995. The *Talk* trace used in this paper is available via anonymous ftp from the site ftp-info3.informatik.uni-wuerzburg.de in the directory /pub/MPEG/.

[28] M. Schwartz. *Broadband Integrated Networks.* Prentice Hall, 1996.

[29] D. Starobinski and M. Sidi. Stochastically bounded burstiness for communication networks. In *Proc. IEEE Infocom '99*, pages 36–42, March 1999.

[30] 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.

[31] 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.

[32] H. Zhang and D. Ferrari. Rate-Controlled Static-Priority Queueing. In *Proc. IEEE Infocom '93*, pages 227–236, April 1993.

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

[34] 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.

Movie Trace	Average frame size (bits/frame)	Mean Rate (Mbps)	Peak Rate (Mbps)
Terminator	10,904	0.261	1.90
Lambs	7,312	0.171	3.22

Table 1: Parameters of the movie traces.

A Additional Examples: MPEG Traces

In this section, we continue our evaluation of the effective envelope method for admission control. Different from the previous sections, however, we use statistics of MPEG-compressed video as traffic sources. The evaluation with MPEG streams is analogous to that in [30] which investigated the maximum achievable utilization at a link with a deterministic service. Several works on statistical multiplexing [15, 25] have used MPEG traffic sources to evaluate methods for admission control.

In our examples, a number of MPEG-compressed video sequences are multiplexed on 45-Mbps and 622-Mbps links. We assume that the video sequences are played continuously with a randomly shifted starting time chosen uniformly over the length of the trace. We consider two traces of MPEG-compressed video from [27]. The first trace is taken from the movie "Terminator 2" (*Terminator*), and the second trace is obtained from the movie "Silence of the Lambs" (*Lambs*). Both traces are digitized to 384 by 288 pixels with 12 bit color information and compressed at 24 frames per second with frame pattern IBBPBBPBBPBB (12 frames). Each sequence consists of a total of 40,000 video frames, corresponding to approximately 30 minutes of video. The data of these traces is given in terms of frame sizes. In Table 1, we show some of the statistical parameter of the traces. We assume that the arrival of a frame is spread evenly over an interframe interval (of length 1/24 s); Hence, a (normally instantaneous) frame arrival occurs at a constant rate.

For each of the MPEG traces, we assume that a deterministic regulator is obtained using the method described in [30]: (1) the empirical envelope is obtained from the MPEG trace (using Eqn. (9) with I_β set to the length of the movie trace), (2) the convex hull of the empirical envelope is determined, yielding a piecewise linear function, and (3) the segments of the concave hull are taken as leaky bucket parameters. In Table 2 we present the leaky bucket parameters which are obtained from the two MPEG traces, yielding a large number of parameters. An algorithm which achieves an accurate characterization with a significantly lower number of leaky bucket parameters can be found in [18].

We now conduct a similar set of experiments as in Section 6. As before, we compare the results with those from simulations, and other methods to calculate the statistical multiplexing gain. Due to the complexity of implementing the method in [23], we do not include results for the *EB-RRR* method. In our simulation, the simulated time is equal to twice the length of the transmission time of the MPEG trace.

It should be noted that the MPEG traffic sources are not *random* sources. We have not verified if the assumptions of stationarity is satisfied for the MPEG sources.

Silence of the Lambs (*Lambs*)			Terminator 2 (*Terminator*)	
Rate parameter (Bits per second)	Burst parameter (Bits)		Rate parameter (Bits per second)	Burst parameter (Bits)
$\rho_1 = 3,221,376.0$	$\sigma_1 = 0.0$		$\rho_1 = 1,909,440.0$	$\sigma_1 = 0.0$
$\rho_2 = 867,008.0$	$\sigma_2 = 98,098.7$		$\rho_2 = 869,056.0$	$\sigma_2 = 43,349.3$
$\rho_3 = 759,628.8$	$\sigma_3 = 156,262.4$		$\rho_3 = 791,680.0$	$\sigma_3 = 75,589.3$
$\rho_4 = 694,336.0$	$\sigma_4 = 246,149.3$		$\rho_4 = 624,776.3$	$\sigma_4 = 165,995.4$
$\rho_5 = 656,472.0$	$\sigma_5 = 321,122.0$		$\rho_5 = 592,576.0$	$\sigma_5 = 214,296.0$
$\rho_6 = 647,850.7$	$\sigma_6 = 372,131.6$		$\rho_6 = 425,421.1$	$\sigma_6 = 485,922.6$
$\rho_7 = 563,438.9$	$\sigma_7 = 1,126,242.3$		$\rho_7 = 361,641.5$	$\sigma_7 = 679,919.0$
$\rho_8 = 502,912.0$	$\sigma_8 = 2,042,261.3$		$\rho_8 = 346,464.0$	$\sigma_8 = 961,968.0$
$\rho_9 = 448,013.1$	$\sigma_9 = 2,911,892.3$		$\rho_9 = 317,920.00$	$\sigma_9 = 1,563,770.7$
$\rho_{10} = 208,800.0$	$\sigma_{10} = 3,157,800.0$		$\rho_{10} = 304,514.7$	$\sigma_{10} = 1,853,100.7$

Table 2: Leaky bucket parameters of the movie traces using the algorithm from [30].

A.1 Example A.1: Comparison of Envelope Functions for MPEG Traces

Analogous to Example 1 in Section 6, we study the shape of local and global effective envelopes as functions of the lengths of time intervals. The envelopes are compared to the deterministic envelope $A_j^*(\tau) = \min_k\{\sigma_{jk} + \rho_{jk}\tau\}$, where the parameters $\{\sigma_{jk}, \rho_{jk}\}_{k=1,\ldots,K}$, given in Table 2, are obtained from the concave hull of the empirical envelope of the movie traces [30].

Figures 11(a) and 11(b), respectively, show the results for N multiplexed *Lambs* and *Terminator* traces, where N is set to $N = 100, 1000$, and 10000 flows. In the graphs, we plot the size of the envelopes normalized by the number of flows as functions of time. We use $\varepsilon = 10^{-6}$ for all envelopes.

We can make the same observations as in Example 1. Local and global effective envelopes are much smaller than the deterministic envelope or the peak rate. The global effective envelope is always larger than the local effective envelopes. Increasing the number of flows N increases the statistical multiplexing gain, leading to a lower traffic rate for each flow. Also, increasing N reduces the difference between the local and the global effective envelopes.

In Figure 12 we show the shape of the envelopes for a fixed number of flows, $N = 1000$, and different values of ε, namely $\varepsilon = 10^{-3}, 10^{-6}$ and 10^{-9}. Figure 12 shows that the effective envelopes are not very sensitive to variations of the parameter ε.

In Figure 13 we show how the effective envelopes vary if the number of flows N is increased. As in Example 1, we consider the values of the envelopes at time interval $\tau = 50$ ms. For comparison, we include the peak and average rates into the graph. As in Example 1, when N is large, the local and global effective envelopes are close to the average traffic rate. Also, the values for local and global envelopes are similar when N is large.

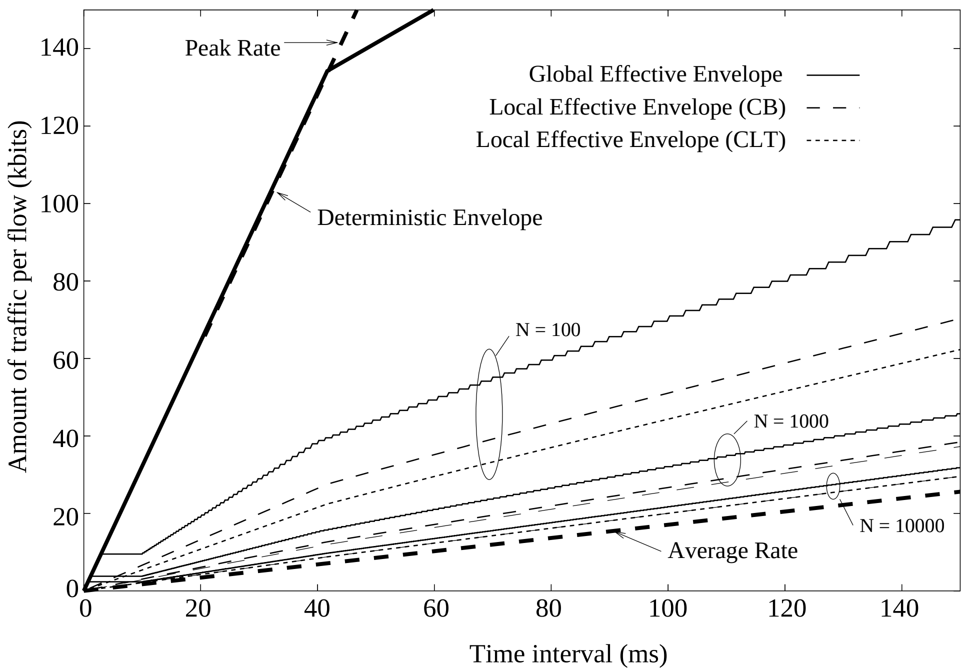

(a) *Lambs.*

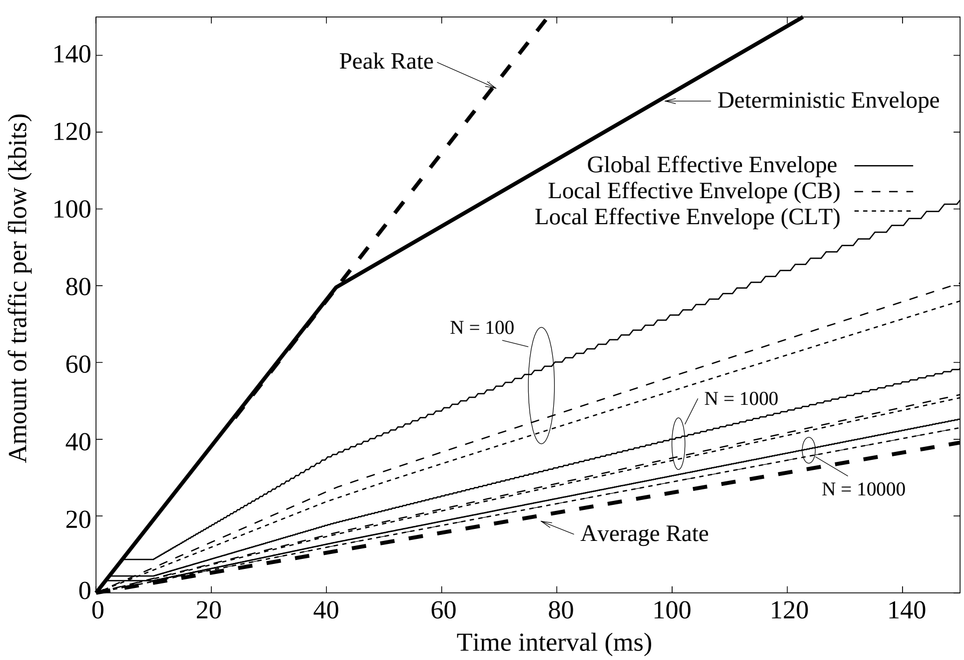

(b) *Terminator.*

Figure 11: Example A.1: Comparison of envelopes for $\tau \leq 150$ ms, $\varepsilon = 10^{-6}$, and for number of flows $N = 100, 1000, 10000$.

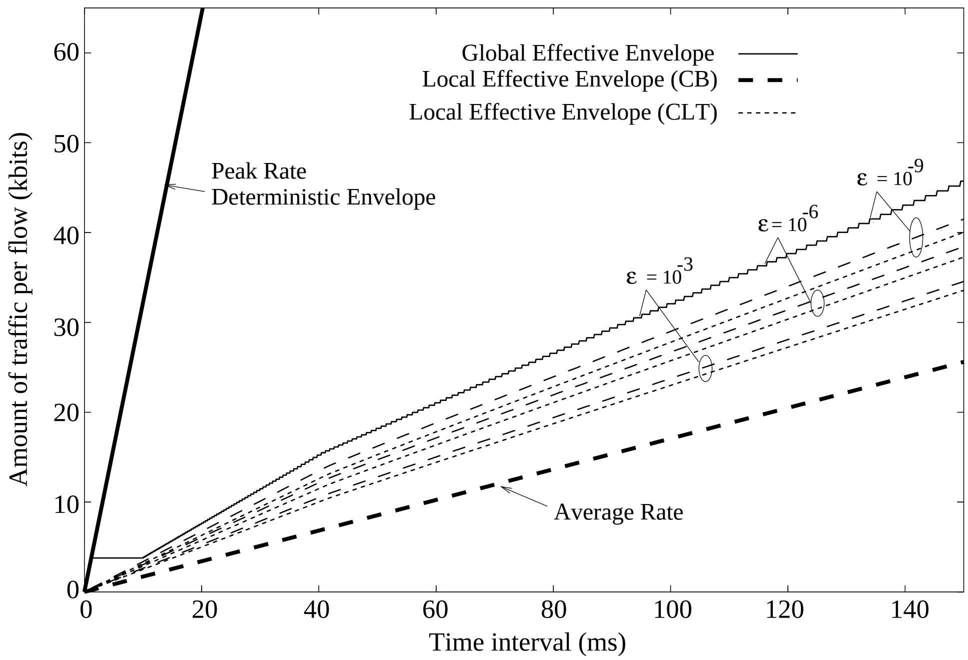

(a) *Lambs.*

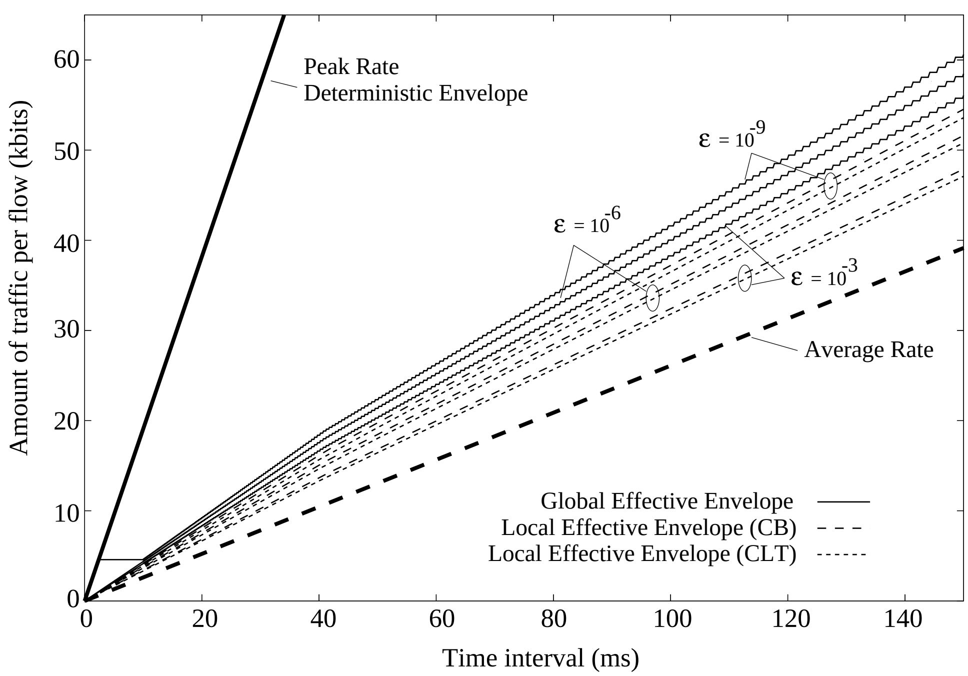

(b) *Terminator.*

Figure 12: Example A.1: Comparison of envelopes for $\tau \leq 150$ ms, number of flows $N = 1000$ and $\varepsilon = 10^{-3}, 10^{-6}, 10^{-9}$.

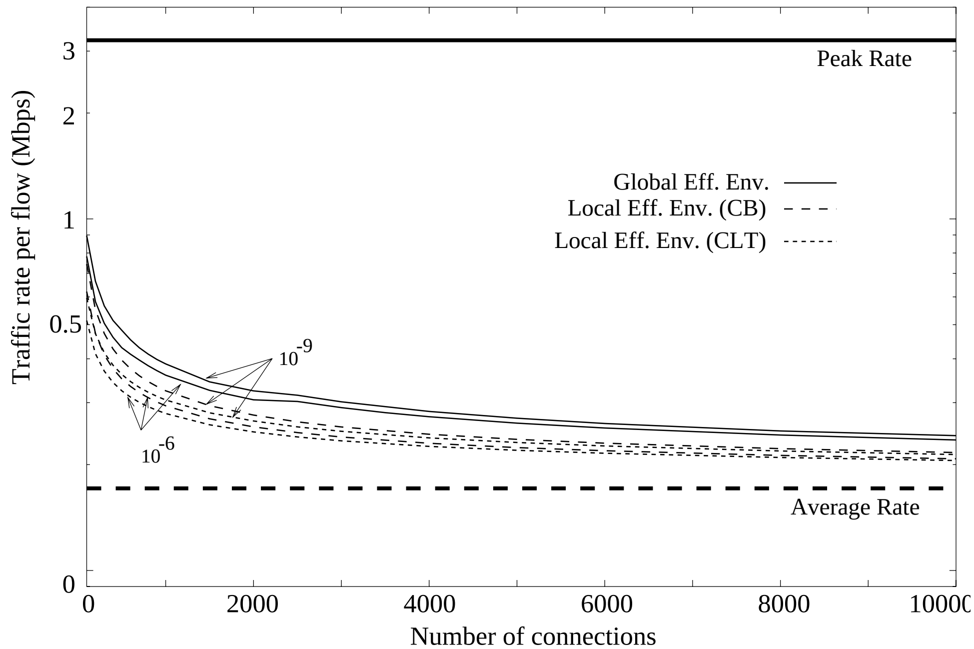

(a) *Lambs.*

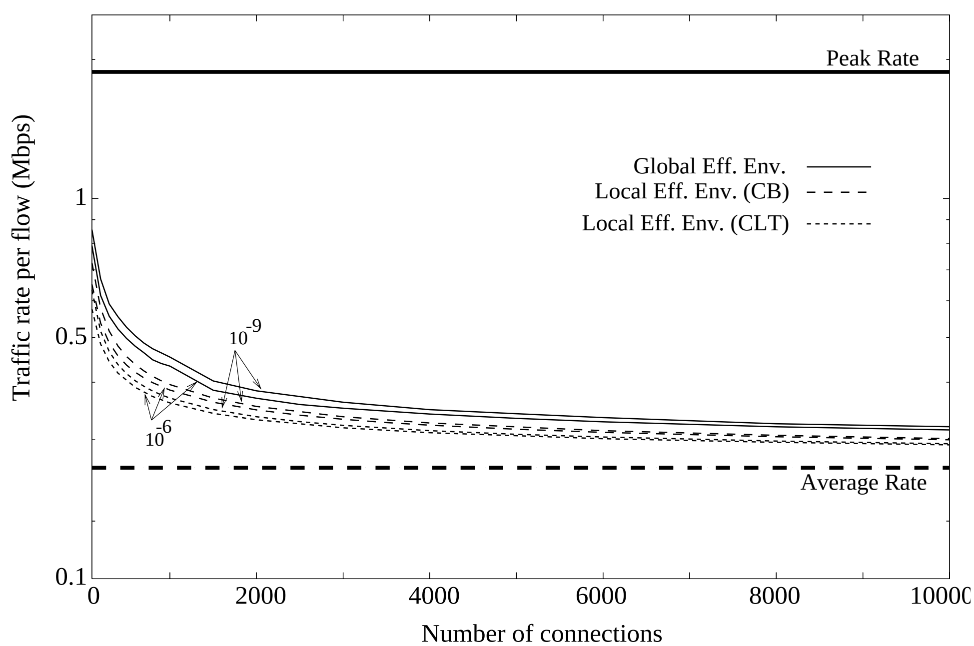

(b) *Terminator.*

Figure 13: Example A.1: Traffic rates $\mathcal{G}_{\mathcal{C}}(\tau;\varepsilon)/(N\tau)$ and $\mathcal{H}_{\mathcal{C}}(\tau;\varepsilon)/(N\tau)$ of *Lambs* and *Terminator* for $\tau = 50$ ms and $\varepsilon = 10^{-6}$ or 10^{-9} .

A.2 Example A.2: Admissible Region for Homogeneous MPEG Flows

This example is analogous to Example 2 in Section 6. We consider a single link with a FIFO scheduler and compare the admissible region of our and other admission control methods for guaranteeing QoS. The traffic sources are either all flows from the *Lamb* MPEG trace or all flows from the *Terminator* MPEG trace. We include the following methods in our comparison:

- *Peak Rate and Average Rate Allocation.*
- *Deterministic QoS.*
- *Statistical QoS with Effective Bandwidth (EB-EMW) [10].*
  Note that this method only considers traffic which is regulated by peak-rate enforced leaky buckets. Research on deterministic QoS has provided evidence that such a regulator does not lead to an accurate characterization of MPEG video traces [18, 30].

  Recent work has adapted the *EB-EMW* method to more general traffic regulators [19, 23]. However, due to the significant complexity of implementing these methods, we are unable to provide numerical data.
- *Statistical QoS with Bufferless Multiplexer (Bufferless MUX)* [25].
- *Statistical QoS using Local Effective Envelopes (CLT and CB) and the Global Effective Envelope.*
- *Simulations.*

The results for this example are shown in Figures 14–17. All Figures show the number of admitted flows as a function of the delay bound. The selection of parameters C and ε in the Figures is as follows.

Figure 14:	$C = 45$ Mbps	$\varepsilon = 10^{-6}$
Figure 15:	$C = 45$ Mbps	$\varepsilon = 10^{-9}$
Figure 16:	$C = 622$ Mbps	$\varepsilon = 10^{-6}$
Figure 17:	$C = 622$ Mbps	$\varepsilon = 10^{-9}$

The results show that the effective envelope methods perform well if compared to other methods. All statistical QoS approaches, with exception of *EB-EMW*, admit many more flows than a deterministic QoS. As mentioned earlier, the poor showing of the *EB-EMW* method is due to the traffic characterization which is not appropriate for MPEG traces. The *EB-EMW* results are included for reference purposes only. The results in Figures 16 and 17 show that the statistical multiplexing gain of all methods is significantly higher when the link capacity is high.

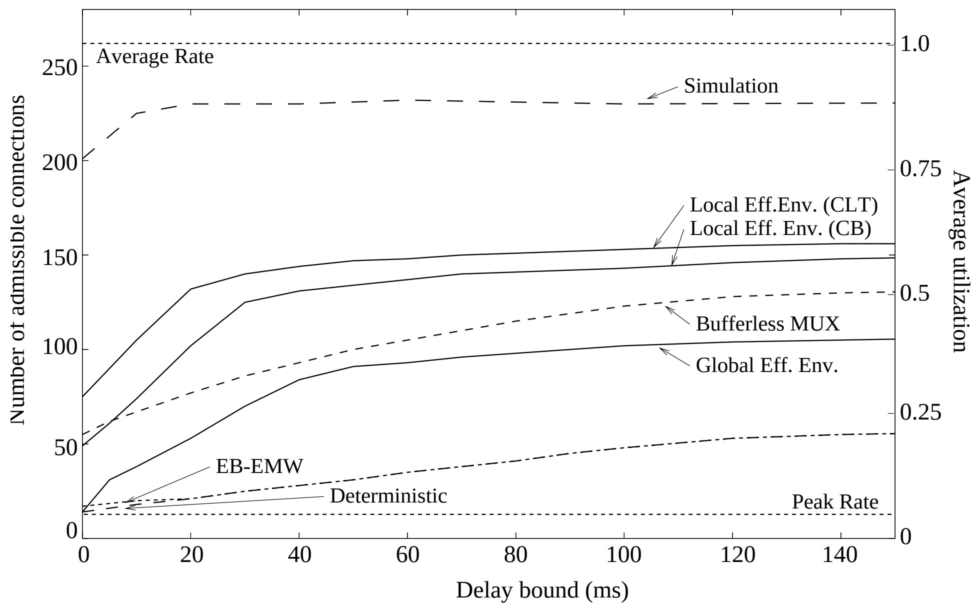

(a) *Lambs.*

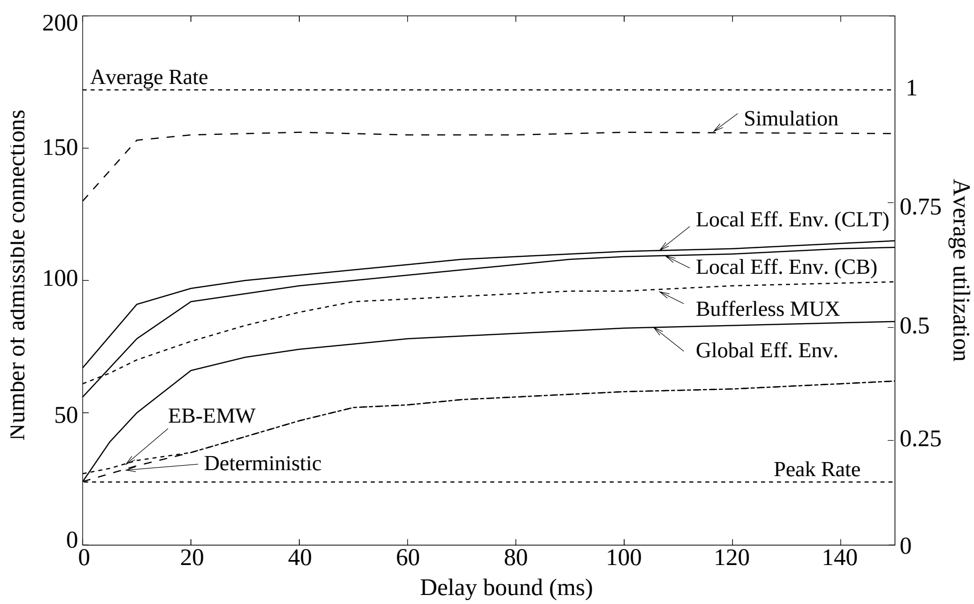

(b) *Terminator.*

Figure 14: Example A.2: Admissible number of homogeneous MPEG flows at a FIFO scheduler as a function of delay bounds, $C = 45$ Mbps, $\varepsilon = 10^{-6}$.

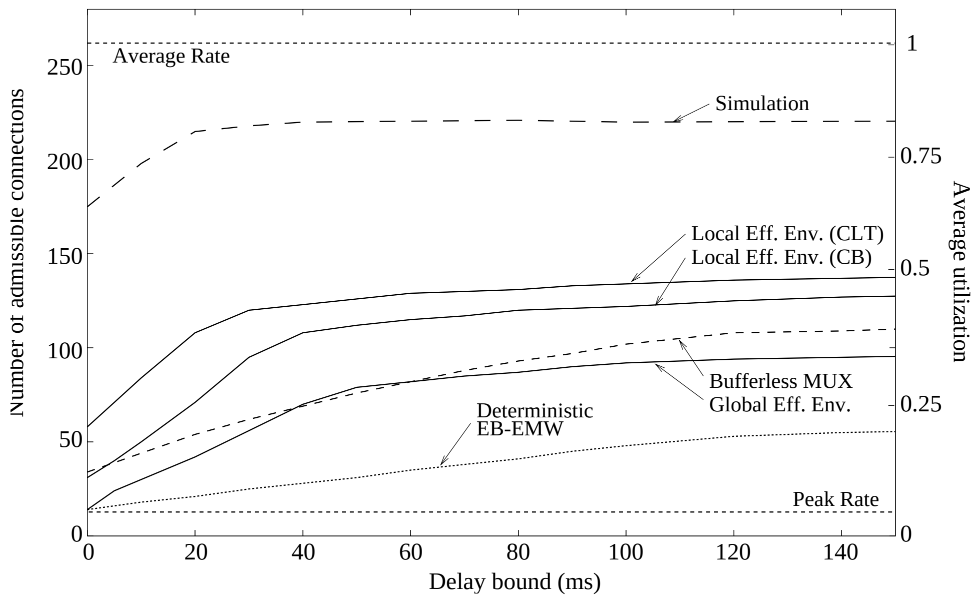

(a) *Lambs.*

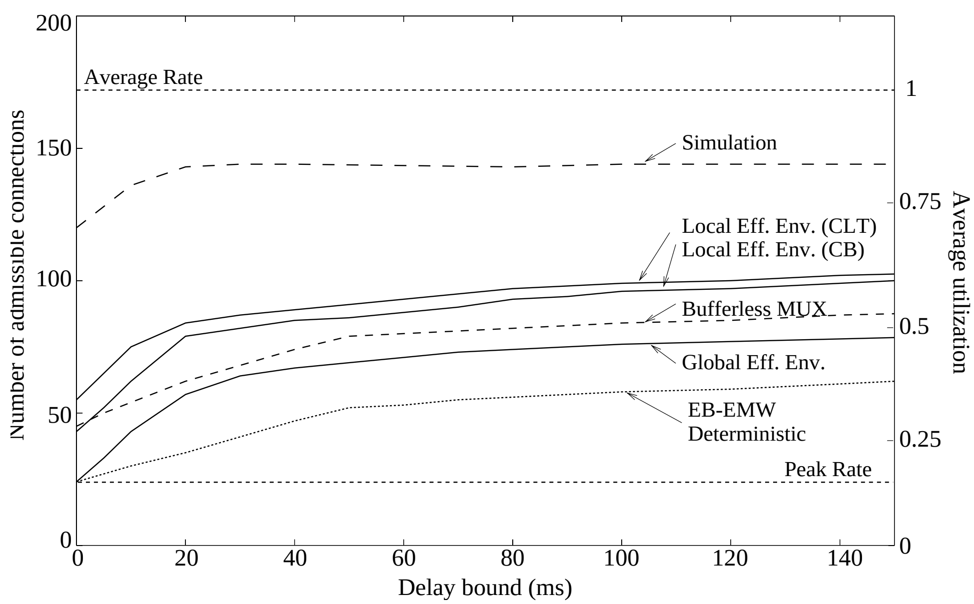

(b) *Terminator.*

Figure 15: Example A.2: Admissible number of homogeneous MPEG flows at a FIFO scheduler as a function of delay bounds, $C = 45$ Mbps, $\varepsilon = 10^{-9}$.

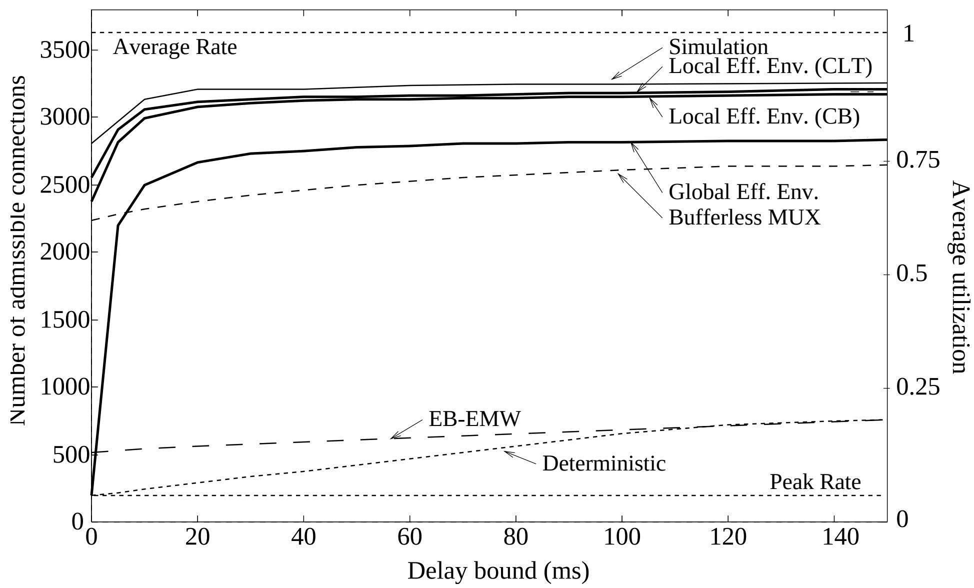

(a) *Lambs.*

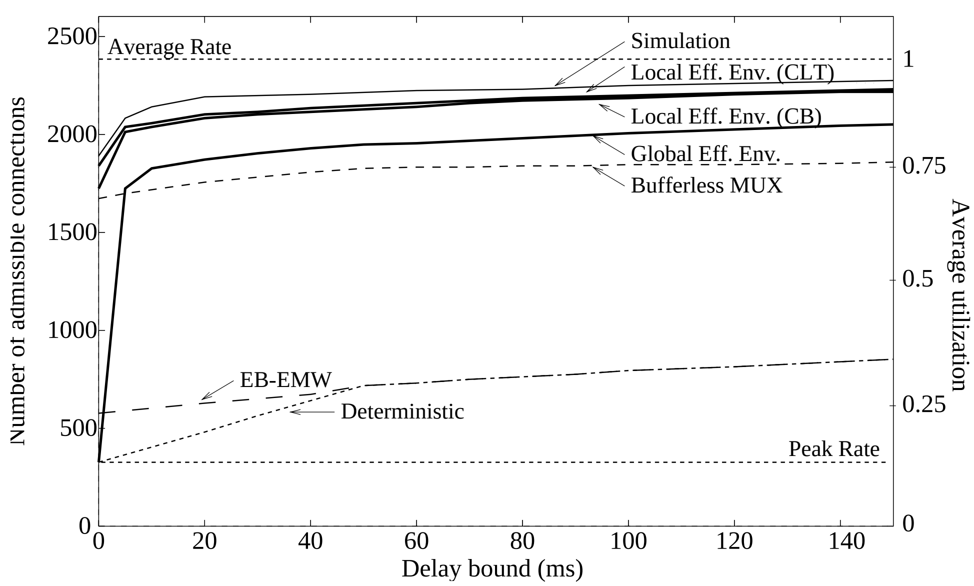

(b) *Terminator.*

Figure 16: Example A.2: Admissible number of flows at a FIFO scheduler for homogeneous flows as a function of delay bounds, $C = 622$ Mbps, $\varepsilon = 10^{-6}$.

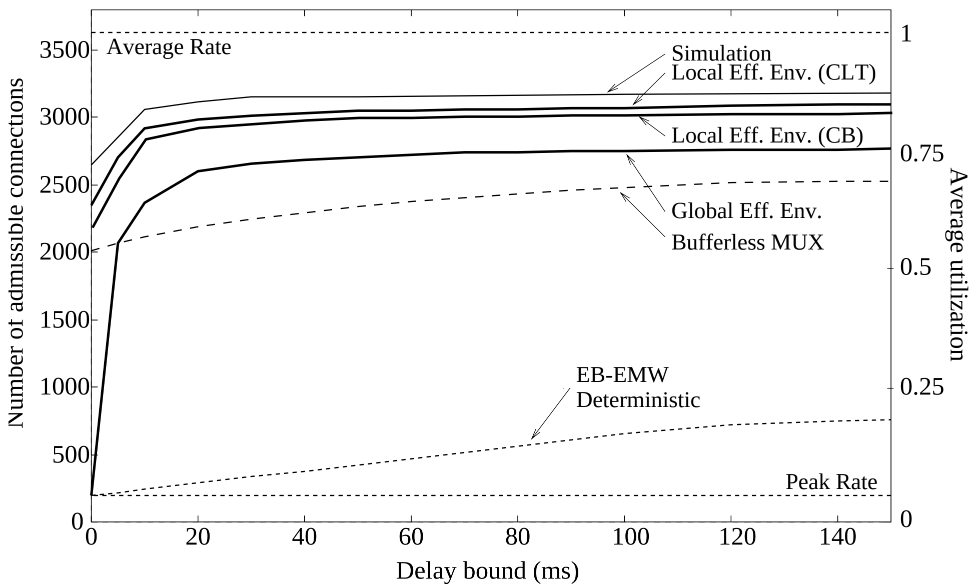

(a) *Lambs.*

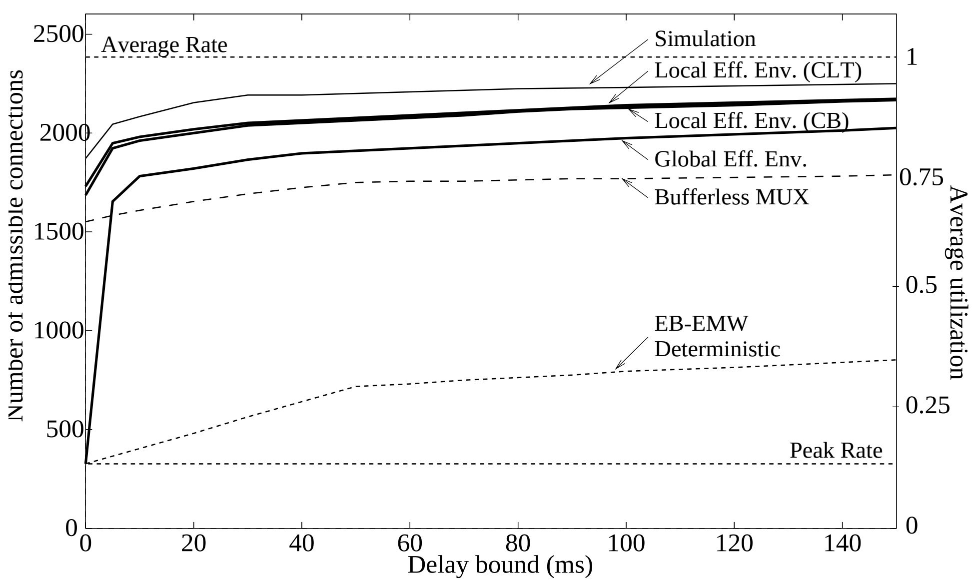

(b) *Terminator.*

Figure 17: Example A.2: Admissible number of flows at a FIFO scheduler for homogeneous flows as a function of delay bounds, $C = 622$ Mbps, $\varepsilon = 10^{-9}$.

A.3 Example A.3: Link Utilization for Homogeneous MPEG Flows

In this example we show, similarly as in Example 3 of Section 6, how the achievable average utilization of a link with a FIFO multiplexer increases as the capacity of the link is increased. We use the same MPEG traces as in the previous two examples. We fix the delay bound of traffic to $d = 50$ ms and we set $\varepsilon = 10^{-6}$.

The results of this experiment are shown in Figure 18. We depict the achievable average link utilization as a function of the link capacity. The average achievable link utilization is the sum of the average rates of flows which can be accepted according to a chosen admission control test.

The results in Figures 18(a) and (b) show that for both MPEG traces, an average utilization of 60% and higher is attainable if the link capacity is 400 Mbps or more.

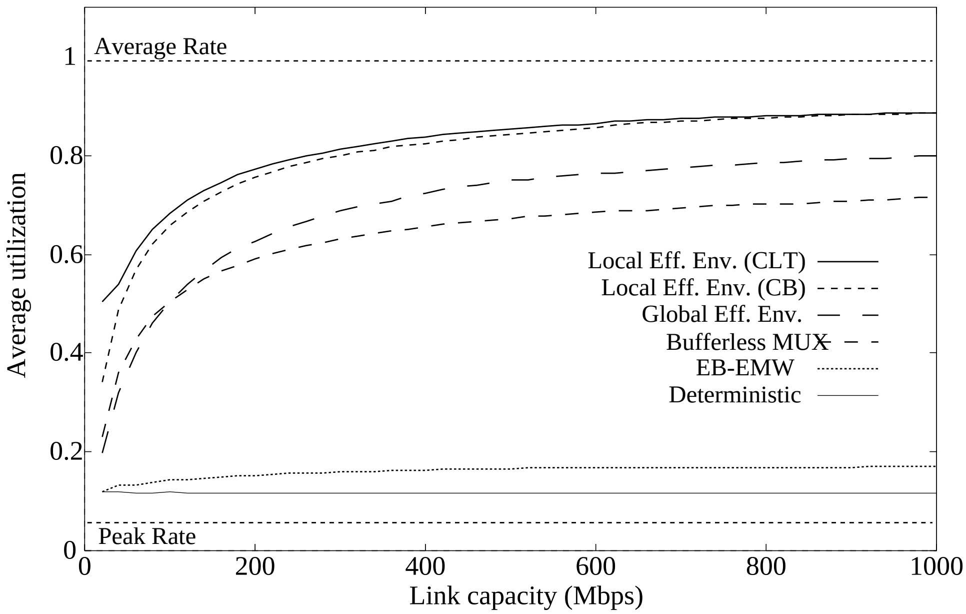

(a) *Lambs.*

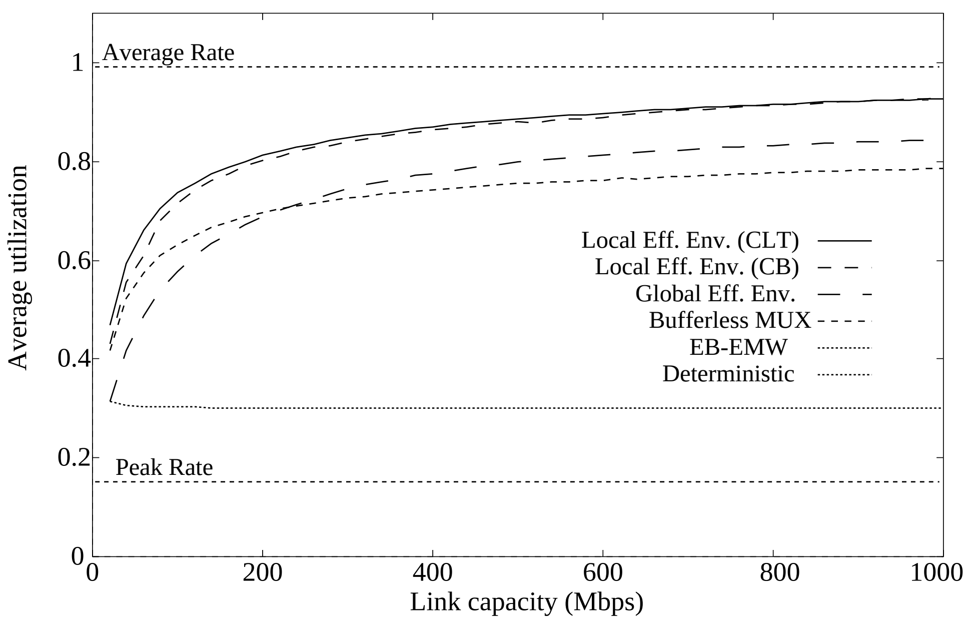

(b) *Terminator.*

Figure 18: Example A.3: Average utilization vs. link capacity, $\varepsilon = 10^{-6}$ and $d = 50$ ms.

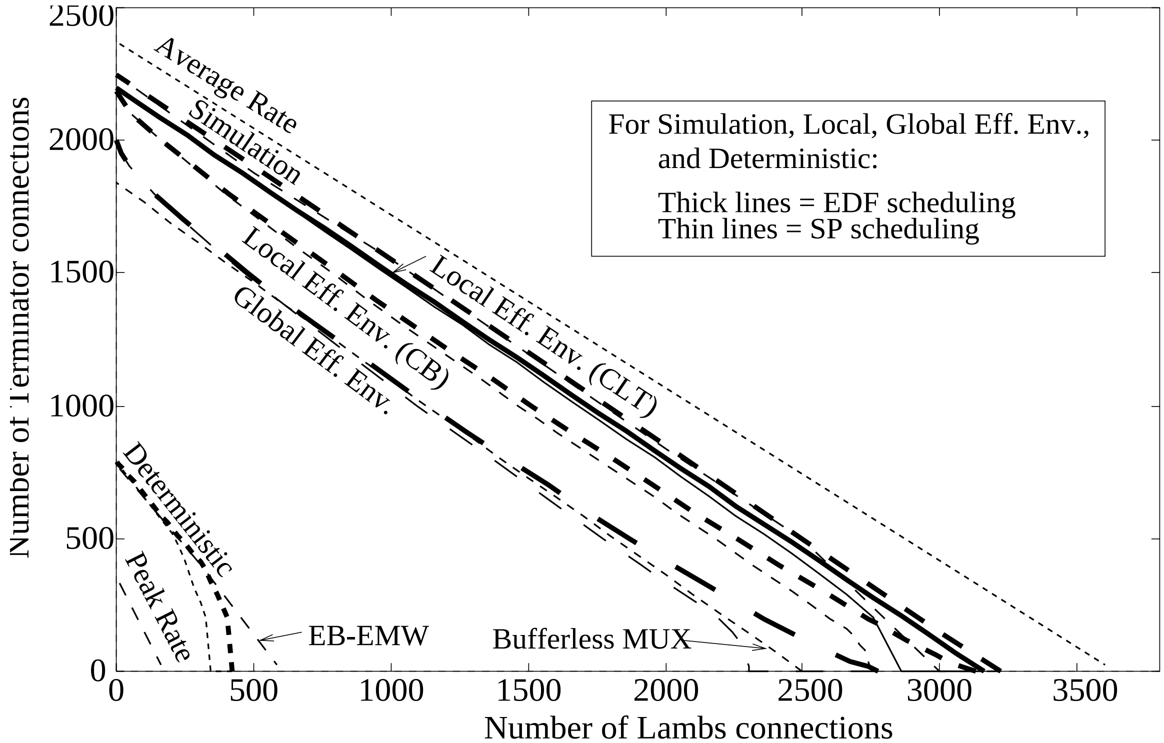

Figure 19: Example A.4: Admissible region of multiplexing *Lambs* and *Terminator* flows with $\varepsilon = 10^{-6}$ and $d_{Terminator} = 50$ ms and $d_{Lambs} = 100$ ms.

A.4 Example A.4: Admissible Region for Heterogeneous MPEG Flows

We investigate the statistical multiplexing gain at a link with two classes of traffic, flows of type *Lambs* and flows of type *Terminator*. The link has a capacity of 622 Mbps. The delay bounds are set to $d_{Terminator} = 50$ ms for flows of type *Terminator*, and to $d_{Lambs} = 100$ ms for flows of type *Lambs*.

We consider two scheduling algorithms, Static Priority (SP) and Earliest-Deadline-First (EDF). For purposes of comparison, we include the methods EB-EMW and *bufferless MUX*, which do not consider scheduling algorithms. Also, we include results for a peak rate allocation, average rate allocation, and deterministic QoS.

In Figure 19. we show the admissible region for the various methods. As in Example 7 in Section 6, the difference between SP and EDF schedulers is small in all cases. The effective envelope is, again, more conservative than the local effective envelope method. The results with the local effective envelope are quite close to those attainable with an average rate allocation.

B Simulations

In this section, we provide some details on our simulation experiments from Section 6 and Appendix A. Note that, in all simulations, the packet size is ignored and we assume that traffic arrivals and transmissions are fluid flow. This is consistent with our earlier decision to ignore the discrete size of packets in this study.

QoS Violations

In our framework, a QoS violation occurs if traffic exceeds a delay bound. Buffers are assumed to be large enough so that no buffer overflow occurs. Note, however, that in FIFO schedulers, a deadline violation occurs if the buffer length exceeds a certain length. We have exploited this feature when comparing the effective envelope approach to equivalent bandwidth approaches.

One can think of different definitions of the relative frequency of a delay bound violation. (Additional definitions of QoS violations can be found in [25].)

1. The fraction of traffic which experiences a delay bound violation.

2. The fraction of time during which deadline violations occur.

3. The fraction of busy periods which contain a delay bound violation.

We use the first notion of QoS violation, probably the most intuitive one, in our simulation. The second notion may be useful if a QoS violation is related to a certain state of the scheduler, e.g., the buffer has a maximum length and arrivals to a full buffer are lost. The third definition is of interest in the context of the global effective envelope. Recall that the global effective envelope is used to bound the likelihood of a QoS violation in a time interval.

In a simulation, which measures delay bound violations, one needs to specify how traffic which has experienced a QoS violation is handled. There are two options.

- **Option 1: No Discard.** Traffic which experiences a delay bound violation is being transmitted.

- **Option 2: With Discard.** Traffic which experiences a delay bound violation is discarded and not transmitted.

In our simulations, we implement the first option. Implementing the second option requires the availability of monitoring functions for detecting a QoS violations. Not, however, that the first type of simulations is more pessimistic than the second one.

In Figure 20 we show the difference of the *No Discard* and *With Discard* options. The figure is the same as Figure 7, but includes simulation results for both Option 1 ('Simulation without discard') and Option 2 ('Simulation with discard').

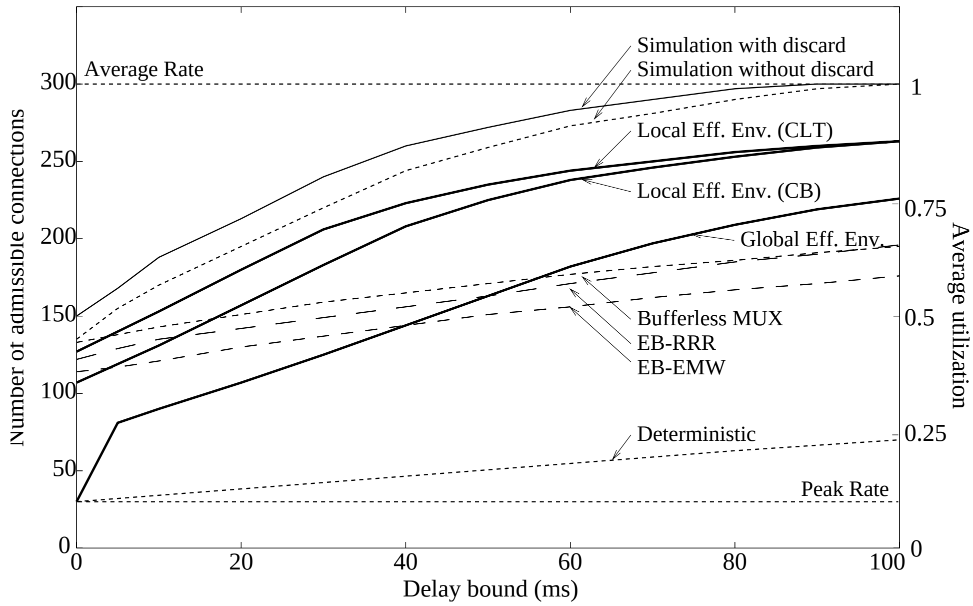

(a) Class 1.

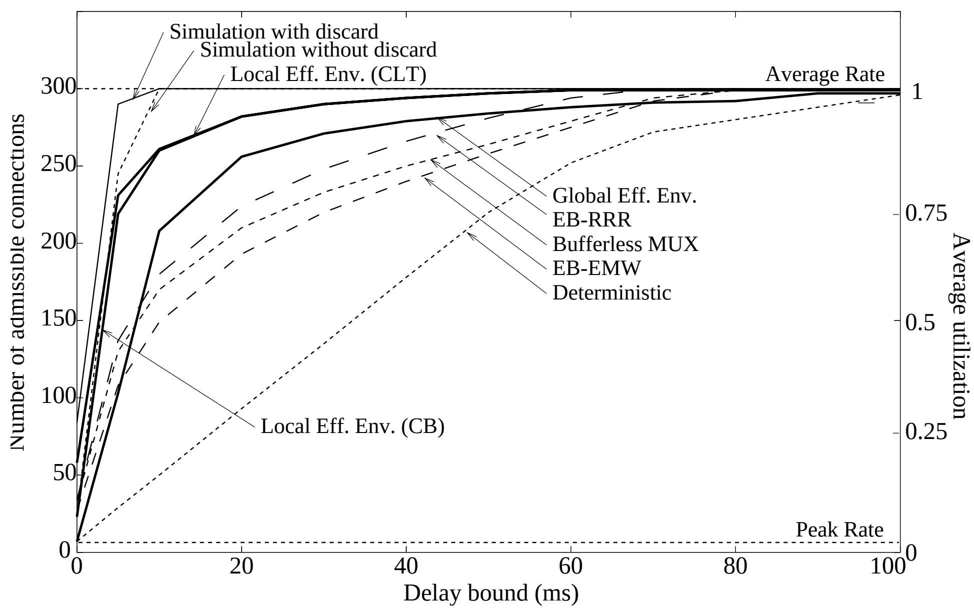

(b) Class 2.

Figure 20: This figure is similar to Figure 7. The only difference to Figure 7 is the inclusion of the 'With Discard' simulation.