Providing End-to-End Statistical Delay Guarantees with Earliest Deadline First Scheduling and Per-Hop Traffic Shaping

Vijay Sivaraman
Dept. of Computer Science
Univ. of California at Los Angeles
Los Angeles, CA 90095, USA.

Fabio Chiussi
Data Networking Systems Research
Bell Labs, Lucent Technologies
Holmdel, NJ 07733, USA.

Abstract—**This paper develops a framework for *statistically* guaranteeing end-to-end delay bounds to leaky-bucket-constrained flows transporting real-time traffic in a network of switches using Earliest Deadline First (EDF) packet scheduling and per-hop traffic shaping. We first analyze the delay-bound violation probabilities at an isolated EDF scheduler fed by fluid source processes generating extremal dual-leaky-bucket-regulated traffic. We compute a close upper bound by applying the Beneš approach to an equivalent hypothetical system derived from the real one. We then extend the analysis to the end-to-end scenario in presence of traffic reshaping at each node in the network. We compare the analytical results with simulations, and show that the match is very close. We also investigate the advantages of smoothing the traffic at the ingress to the network, and propose a simple choice of smoothing parameters, which performs very well. Using realistic traffic parameters, we compare the schedulable region of our statistical framework with that of the corresponding deterministic framework, and demonstrate that the statistical framework allows tremendous improvements in network utilization, even for very low delay-violation probabilities. The framework developed in this paper is therefore highly useful in practical packet networks to provide quality of service to real-time applications in the form of statistical, rather than deterministic, end-to-end delay bounds.**

Keywords— **Scheduling, Earliest Deadline First, Statistical Delay Bounds, Schedulable Region.**

I. INTRODUCTION

This paper addresses the issue of providing Quality of Service (QoS) guarantees in terms of end-to-end delay bounds to real-time applications operating over high-speed packet switched networks. Such guarantees are inherently required by a wide range of applications with very heterogeneous characteristics, e.g., voice and video. To support these heterogeneous QoS guarantees, and allow real-time traffic and best-effort traffic to coexist on the same network infrastructure, packet scheduling schemes more sophisticated than the First-In-First-Out (FIFO) discipline are necessary at each switch in the network; Generalized Processor Sharing (GPS) [20], [21] and Earliest Deadline First (EDF) [10], [26] constitute two popular examples of such scheduling schemes.

Call Admission Control (CAC) frameworks which employ these scheduling schemes to guarantee *deterministic* delay bounds to flows whose traffic is constrained according to a specified traffic envelope have been developed in literature (see [21] for a framework involving GPS and [14], [1] for a treatment of EDF). However, these deterministic frameworks, which provide hard guarantees on end-to-end delay, are excessively conservative in admitting flows into the network, since they have to account for the worst-case scenarios that can be encountered in the switches, even though such worst cases may occur with exceedingly low probabilities. This known over-engineering, intrinsic of the deterministic approach, coupled with the observation that real-time applications are typically resilient to infrequent violations in their delay bounds (i.e., are not unduly hindered if a small fraction, say 10^{-5}, of their packets get excessively delayed or dropped within the network), points to a compelling need for a *statistical* framework where the end-to-end delay bounds are statistically guaranteed, i.e., guaranteed with a reasonably high probability. Such statistical frameworks are expected to allow the links to operate at a much higher utilization, and still meet the QoS requirements of real-time traffic.

Frameworks for *statistical* delay guarantees can operate in the setting of either *stochastic* (e.g., markovian and exponentially bounded burstiness) or *regulated* (e.g., dual leaky bucket) source models. Numerous frameworks that consider stochastic source models have been developed in literature - for example, see [29], [31], [5] for GPS and our work in [24] for EDF. However, two considerations make frameworks based on regulated source models of greater practical relevance than those based on stochastic models: 1) conformance with stochastic models is difficult to enforce, and the problem worsens in the multi-node setting where traffic interactions at a switch distort the traffic entering the downstream nodes; on the contrary, conformance with regulated source models is easy to enforce by means of traffic shaping and policing (even at nodes internal to the network); and 2) there is a large heterogeneity in accepted stochastic models for real-time traffic (for example, markovian models may be acceptable for video conferencing [7] but are considered insufficient for modeling stored video); on the contrary, the dual-leaky-bucket-regulated source model provides a uniform basis for the description of real-time traffic, and dual-leaky-bucket traffic descriptors are generally accepted in practical networks.

Recently, statistical frameworks in the setting of *regulated* source models have been developed for priority scheduling [15] and GPS [9]. In our case, we have chosen the EDF scheduling scheme, since it is known to provide the *optimal* delay performance [12], [17] in the deterministic environment, and we expect the advantages to carry over to the statistical setting (indeed, a comparison study is performed in [25] to quantify the advantages of EDF over GPS in the statistical setting). Specifically, in this paper, we formulate a framework which guarantees end-to-end statistical delay bounds to flows offering dual leaky bucket regulated traffic in networks employing the EDF scheduling scheme in conjunction with per-node traffic shaping. (This work is drastically different from our work in [24], which centered on stochastic sources.)

Our analysis proceeds in two steps. First, we compute the *aggregate* probability of delay-bound violations at an isolated EDF scheduler fed by independent sources generating *extremal* on-off fluid traffic conformant with their dual leaky-bucket regulators. This is carried out by relating the aggregate probability of delay violations to the queue length in a hypothetical system derived from the real one, and then applying the Beneš approach [23], [19] for the characterization of the queue length to the newly derived system. Then, following what was done in the deterministic setting [14], we extend the analysis to multiple nodes by reshaping the traffic at each node. The choice of optimal shaping parameters for a flow is in general a very hard task; here, we propose a simple choice of shaper (based upon the number of hops traversed by the flow) which is argued to be reasonable and can easily be computed independent of the exogenous traffic in the network.

We validate our analytical framework against simulations, and show that the match is very close. Moreover, at very low probabilities, the schedulable region of our statistical analysis is shown to converge to that of the deterministic approach. We then quantify the enormous advantages of the statistical framework over the deterministic one; for realistic traffic parameters, the improvement in network utilization is as high as 35%, even for as low a delay violation probability as 10^{-5}. We also show that the benefits accrued from appropriate smoothing of the traffic at the ingress to the network can be quite large. Our framework is therefore of great significance for networks providing statistical end-to-end QoS guarantees. The one drawback of our approach is its computational complexity, which scales exponentially with the number of QoS classes. Our current research is addressing this issue.

Concurrently to our work, a rather general framework for statistical QoS, which handles EDF scheduling as a special case, has been developed in [22]. The approach employed there is very different from the one developed in this paper, and is based upon "statistical service curves", an extension of Cruz's "deterministic service curves" [3]. Our results (some of which are discussed in Section V) indicate that our framework in general yields more accurate bounds than [22], and the difference can be quite significant at low probability values.

The rest of this paper is organized as follows. In Section II, we provide relevant background information on EDF scheduling and the Beneš approach for characterizing the unfinished work in the system. In Section III, we present the analysis of an isolated EDF scheduler fed by independent extremal regulated sources, and in Section IV we extend it to the multi-node scenario in the presence of traffic reshaping at each node. We present numerical results in Section V. Finally, in Section VI, we offer concluding remarks and point to directions for future research.

II. Background

A. EDF Scheduling

We review some basic concepts related to EDF scheduling, and briefly describe the deterministic analysis of [14], [1] and the statistical framework of [22]. The EDF scheduling discipline [10], [26] works as follows: each flow i at switch m is associated with a *local* delay bound d_i^m; then, an incoming packet of flow i arriving to the scheduler at time t is stamped with a deadline $t + d_i^m$, and packets in the scheduler are served by increasing order of their deadline.

In the deterministic setting, EDF is known to be the optimal scheduling policy at a single switch [12]. Optimality is defined in terms of the *schedulable region* associated with the scheduling policy. Given K flows with traffic envelopes $A_i(t)$[1] ($i = 1, 2, \ldots, K$) sharing an output link, and given a vector of delay bounds $\vec{d} = (d_1, d_2, \ldots d_K)$, where d_i is an upper bound on the local scheduling delay that packets of flow i can tolerate, the schedulable region of a scheduling discipline π is defined as the set of all vectors $\vec{d}$ that are schedulable under π. The authors in [12], [17] show that EDF has the largest schedulable region of all scheduling disciplines, given by the vectors that satisfy the following constraint:

$$\sum_{i=1}^{K} A_i(t - d_i) \le Ct, \qquad \forall t > 0 \tag{1}$$

where C denotes the link rate, $A_i(t) = 0$ for $t < 0$, and it is assumed that either the packet transmission time is negligible or the scheduler is preemptive (in the case of non-premeptive scheduling with non-negligible packet sizes, replacing d_i by $d_i - L/C$ in the above constraint, where L denotes the maximum packet size at the switch, can guarantee the same delay bound to all flows). Given the traffic envelopes and the delay requirements of each flow, inequality (1) can be used directly to devise a single-node CAC mechanism. In the multinode setting, however, the traffic envelopes are no longer known at the inputs of the nodes inside the network, and the interactions that distort the traffic are not easily characterizable. To overcome this problem, Zhang and Ferrari in [30] propose a class of schemes called *Rate-Controlled Service* (RCS) disciplines which reshape the traffic at each hop within the network (EDF with per-hop reshaping is referred to as *Rate Controlled EDF*, RC-EDF). Georgiadis *et al.*, in [14] build upon this model and derive expressions for the end-to-end delay bounds in terms of the shaper envelope $E_i(t)$ and the scheduling delay at each node; they show that no advantages are gained by having non-identical shapers for a flow at each switch it traverses. The end-to-end delay bound d_i for flow i is given by

$$d_i = d_i^{sh} + \sum_{m=1}^{M} d_i^m \tag{2}$$

[1] Flow i has envelope $A_i(t)$ if the amount of flow i traffic entering the network in any interval of length t is bounded by $A_i(t)$ [2]. A typical example is the *dual leaky bucket* descriptor (p_i, σ_i, ρ_i) denoting the envelope $A_i(t) = \max\{p_i t, \sigma_i + \rho_i t\}$.

where $d_i^{sh} = D(A_i \| E_i)$ denotes the maximum shaper delay and d_i^m is the local scheduler delay bound at the m-th switch for flow i. The maximum shaper delay is incurred only *once*, and is independent of the number of nodes on the path. Equation (2), in conjunction with the single-node CAC derived from inequality (1) above readily leads to an end-to-end CAC framework [14], [1]: once an appropriate shaper envelope E_i has been chosen, the delay incurred in the shaper is computed; the remaining delay is split among the schedulers on the path of the flow, and the flow is admitted only if the single-node CAC at each switch along the path admits the flow.

The deterministic framework presented above has the inherent drawback that it accounts for the worst-case scenarios that can be encountered by the scheduler, and is therefore overly conservative in admitting flows into the network. Our object in this paper, therefore, is to devise a corresponding statistical framework which is more aggressive in admitting flows, but at the same time permits the resulting degradation in QoS to be quantified so that it can be managed at a desirable level.

Concurrently to our work, a rather general framework for guaranteeing statistical delay bounds to regulated sources under a variety of scheduling schemes, including EDF, has been developed in [22]. It is based on a stochastic extension of the idea of "service curves" developed by Cruz [3]. Let $B_i(t)$ denote the *statistical traffic envelope* for flow i, defined as a sequence of random variables such that in any interval $[u, u+t)$ the input traffic $A_i[u, u+t]$ satisfies

$$A_i[u, u+t] \leq_{st} B_i(t)$$

where $A_i[u, u+t] \leq_{st} B_i(t)$ (stochastic inequality) denotes $P[A_i[u, u+t] > z] \leq P[B_i(t) > z]$ for all z. Then, the authors in [22] show that the delay violation probability P_{vio} at an EDF scheduler serving K flows is upper bounded by

$$P_{\text{vio}} \leq P[\max_t \{\sum_{i=1}^{K} B_i(t-d_i) - Ct\} > 0] \tag{3}$$

In practice, the quantity on the right is computed by employing Gaussian approximations. For this purpose, the following quantities are defined:

$$\begin{aligned}
\sigma_t^2 &= \text{var}\{\sum_{i=1}^{K} B_i(t-d_i) - Ct\} \\
\alpha_t &= \frac{0 - E\{\sum_{i=1}^{K} B_i(t-d_i) - Ct\}}{\sigma_t} \\
\alpha &= \inf_t \alpha_t .
\end{aligned}$$

Approximating $\sum_{i=1}^{K} B_i(t-d_i) - Ct$ as Gaussian, the quantity P_{vio} can be upper bounded by

$$P_{\text{vio}} \leq e^{-\frac{\alpha^2}{2}} \tag{4}$$

The reader is referred to [22] for proofs and more detailed explanations of the above equations.

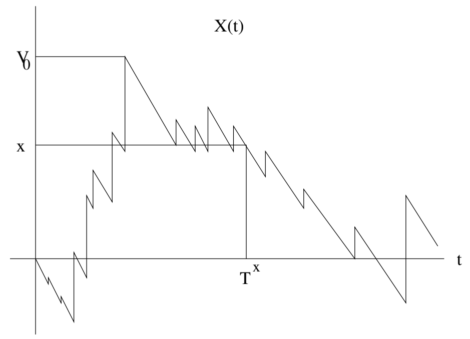

Fig. 1. A realization of the process $X(t)$ with the last exit time T^x

B. The Beneš Approach to Virtual Waiting Time

We describe in brief the essentials of the Beneš approach which allows the *virtual waiting time*, alias *unfinished work in the system* (or, in simpler terms, the "fluid" equivalent of the queue length), to be characterized. We review it only to the extent essential for our work; the reader is referred to [23, Ch. 14] for a more detailed treatment. Consider a service system with unlimited buffer capacity and service capacity of 1 unit of work per unit of time. Assume that the system is stationary, so that 0 represents an arbitrary time instant. Let $A(t), t \geq 0$, denote the amount of work arriving to the system in the interval $[-t, 0)$, and let V_t be the amount of work still in the system at time $-t$. Define $X(t) = A(t) - t, t \geq 0$, to be the *excess work* arriving in $[-t, 0)$. Then, V_t is given by Reich's formula

$$V_t = \sup_{w \geq t} \{X(w) - X(t)\} \tag{5}$$

In particular,

$$V_0 = \sup_{t \geq 0} \{X(t)\} \tag{6}$$

The virtual waiting time can thus be obtained as the maximum of a stochastic process. Further, if we let T^x denote the largest value of t such that $X(t) = x$ (see Figure 1 for an illustration), the following equivalence can be deduced:

$$\begin{aligned}
\{V_0 \geq x\} \quad \Leftrightarrow \quad & \{\exists\, unique\ T^x\ such\ that\ X(T^x) = x \\
& \text{and } X(w) < x \text{ for } w > T^x\}
\end{aligned} \tag{7}$$

The complementary distribution function of V_0, i.e.,

$$\nu(x) = P\{V_0 > x\}, \tag{8}$$

can therefore be expressed by the generic Beneš principle

$$\nu(x) = P\{T^x \in [0, \infty)\} \tag{9}$$

Employing the definition of T^x from (7) yields:

$$\begin{aligned}
\nu(x) &= \int_{u=0}^{\infty} P\{u \leq T^x < u + du\} \\
&= \int_{u=0}^{\infty} P\{X(u+du) < x \leq X(u) \\
&\qquad \text{and } X(w) < x \text{ for } w > u\}
\end{aligned} \tag{10}$$

Applying (5) at the point $t = u$ yields

$$\{V_u = 0\} \Leftrightarrow \{X(w) \leq X(u) \text{ for } w > u\}, \tag{11}$$

which, when applied to (10), leads to the Beneš formula:

$$\nu(x) = \int_{u>0} P\{X(u) \geq x > X(u + du) \text{ and } V_u = 0\} \tag{12}$$

Although the integral can seldom be calculated exactly, simple omission of the condition $\{V_u = 0\}$ yields an upper bound for $\nu(x)$ which is often both reasonably accurate and numerically calculable.

In the next section, we shall see how the above framework, originally developed for characterizing the virtual waiting time in the system, can be adapted for the computation of the delay violation probabilities at an EDF scheduler.

III. Single-Node Statistical Delay Bounds

This section develops the framework for computing the delay violation probabilities at an isolated EDF scheduler fed by dual-leaky-bucket-regulated sources[2]. We first introduce the system model and the theoretical basis for the computation of the delay violation probabilities. We then present the mathematical framework, along with the approximation techniques which allow the quantity of interest to be computed numerically.

A. System Model and Theoretical Basis

Consider an EDF scheduler multiplexing flows onto a transmission link of constant capacity C. The flows are categorized into J classes, with class j ($j = 1, \ldots, J$) having k_j sources, each offering (p_j, σ_j, ρ_j) dual-leaky-bucket-regulated traffic and requiring a delay bound d_j at the scheduler. Without loss of generality, we assume $d_1 \leq d_2 \leq \ldots \leq d_J$. We use the two-tuple (j, i) to refer to the i-th source belonging to the j-th class; $A_{ji}[0, t]$ denotes the amount of work arriving from flow (j, i) in the interval $[0, t)$. We make the following assumptions:

Assumption 1: Traffic is modeled as a fluid; hence packetization issues are ignored.

Nevertheless, we use the term "packet" loosely to refer to an infinitesimal quantity of the traffic. Also, as discussed in II-A, in the presence of non-negligible packet transmission delays, the delay requirement for every source is reduced by L/C (where L denotes the maximum packet size at the switch) and resources are reserved assuming the fluid model.

Assumption 2: Each flow generates traffic independent of all other flows.

Assumption 3: The probability of delay bound violations P_{vio} (i.e., the fraction of aggregate traffic that does not meet its delay bound at the scheduler) has a steady-state stationary value.

The above assumptions are quite similar to the ones routinely adopted in the literature for analyzing buffer occupancies. We further make the assumption that

[2]The terms *source* and *flow* are used interchangeably in this section

Assumption 4: Packets are neither dropped (due to buffer overflows) nor discarded (due to expired deadlines) at the scheduler.

Thus, buffer space is unlimited, and packet losses do not arise. (Indeed, if buffers are sufficiently large and packet losses are rare enough, the delay violation probability and the packet loss probability are orthogonal, and can be analyzed independently). Moreover, the EDF scheduler transmits packets (if any are present) in order of deadline, irrespective of whether the deadline of the packet chosen for transmission has expired or not[3]. This serves as a conservative approximation of the delay violation probability in the presence of packet discard policies.

We compute the delay violation probability, i.e., the fraction of traffic which does not meet its delay requirement at the EDF scheduler, for the *aggregate* traffic at the switch, not on an individual flow or class basis. The individual metrics are very hard to evaluate due to the strong traffic interactions at the EDF scheduler; moreover, they can be controlled to an extent by an appropriate choice of packet discard policy, which is beyond the scope of this work. Under the assumptions stated earlier, the following theorem [24] establishes the theoretical basis for our approach:

Theorem 1: Consider the EDF server at a random time 0. Construct a hypothetical system H which discards all class j ($j = 1, \ldots, J$) traffic arriving in interval $[-d_j, 0)$. Then

$$P_{\text{vio}} = \frac{1}{\rho} P\{Q^H(0) > 0\} \tag{13}$$

where P_{vio} denotes the stationary probability of delay bound violations at the EDF server, ρ the server utilization, and $Q^H(0)$ the queue length, at time 0, in the hypothetical system H.

Proof: Partition the traffic queued at the server into two (logical) queues Q^1 and Q^2 such that every flow (j, i)'s traffic arriving before time $-d_j$ enters queue Q^1 while traffic arriving in $[-d_j, 0)$ enters queue Q^2. Every packet, upon arrival, is assigned a *timestamp* representing the time by which the packet has to be served in order to satisfy its delay bound. By the definition of timestamp, all traffic queued in Q^1 has timestamp less than 0, while no traffic queued in Q^2 has timestamp less than 0. Since the EDF scheduler serves traffic in strict order of timestamp, traffic queued in Q^1 is given pure priority in service over traffic queued in Q^2. Thus, the evolution of Q^1 is as if the server is always available to serve it, i.e., as if Q^2 is completely ignored. This means that Q^1 behaves in the exact identical way as the queue length Q^H in the hypothetical system H. Having established that $Q^H \equiv Q^1$, we also note that if Q^1 is non-empty at time 0, then the traffic being served at time 0 has an expired deadline (since all timestamps in Q^1 are less than 0), while if Q^1 is empty at time 0, there is no delay violation at time 0. Thus the probability of encountering a delay bound violation at the time instant 0 is given by $P\{Q^H(0) > 0\}$,

[3]This is the corresponding assumption to the infinite buffer assumption commonly employed when analyzing packet losses in a finite buffered system.

and P_{vio}, which denotes the probability of encountering a delay bound violation at time 0 given that the system is not idle at time 0, is $P\{Q^H(0) > 0\}/\rho$. △

It should be pointed out that (13) holds *only* for EDF scheduling, since it is the only scheduling policy which serves traffic in strict order of deadline.

B. Mathematical Formulation

From Theorem 1, it follows that if the complementary distribution $\nu(x) = P\{Q^H(0) > x\}$ of the queue length $Q^H(0)$ at time $t = 0$ in the hypothetical system H is computed, the delay violation probability at the EDF scheduler is easily obtained from

$$P_{\text{vio}} = \frac{1}{\rho} \lim_{x \to 0} \nu(x) \tag{14}$$

To compute $\nu(x)$ for the system H, we require more precise knowledge of the arrival process than just its envelope. A conservative approach would be to identify, for each flow, the arrival pattern that is conformant with its regulator and yet "worst-case" in the sense of imposing the maximum stress on network resources, thereby realizing the maximum delay violation probability. Identifying such a traffic pattern, however, is non-trivial. In frameworks which employ buffer losses as the performance measure, numerous authors (e.g., [8], [18], [9]) have suggested the use of the *extremal* on-off traffic model. In this model, a (p, σ, ρ) dual-leaky-bucket-regulated source switches on whenever its token bucket is full, and keeps generating traffic at the peak rate p till the bucket empties; at this point, it switches off and waits for the token bucket to be full again, and the whole process repeats. Though such an extremal process has generally been believed to model "worst-case" behavior, the work in [6], [27] has shown that this need not always be the case. Nevertheless, in this work, we assume the extremal model as it makes the analysis simpler.

Assumption 5: Each flow generates extremal on-off traffic to the extent permitted by its dual leaky bucket regulator.

Note that the traffic pattern under this model for a (p, σ, ρ) dual-leaky-bucket-regulated source is periodic with period $T = T^o + T^f$, where the on-period $T^o = \frac{\sigma}{p-\rho}$ and the off-period $T^f = \frac{\sigma}{\rho}$. Though each source exhibits periodic behavior, the randomness in the system arises from the statistical independence of the sources, i.e., the phase of each source being randomly and uniformly distributed over its period.

We now use the Beneš approach for the computation of the complementary distribution $\nu(x)$ of the queue length $Q^H(0)$ at time $t = 0$ in the hypothetical system H introduced in Theorem 1, which discards all flow (j, i) traffic arriving in $[-d_j, 0)$. The approach here follows similar lines as the one used in [19], although our system is subject to different conditions. Developing on the notation of Section II-B, let Λ_t denote the arrival rate at time $-t$, and let $A(t) = \int_{-t}^{0} \Lambda_t \, dt, (t \geq 0)$ be the total amount of work arriving in the interval $[-t, 0)$. All sources (j, i) with $d_j > t$ do not contribute towards Λ_t or $A(t)$. Further, let V_t be the work still in the system at time $-t$, and let $X(t) = A(t) - Ct$ denote the excess work arriving in $[-t, 0)$, where C denotes the link rate. Then, we have

$$\Lambda_t = \frac{dA(t)}{dt} = C + \frac{dX(t)}{dt} \tag{15}$$

Writing $X(u + du) = X(u) + (\Lambda_u - C)du$ in (12) gives us:

$$\nu(x) = \int_{u>0} P\{x \leq X(u) < x + (C - \Lambda_u)du \text{ and } V_u = 0\} \tag{16}$$

Summing over all possible values of Λ_u, and noting that the system is certainly not empty at u if the arrival rate is greater than C, we have:

$$\begin{aligned}\nu(x) &= \int_{u>0} \sum_{0 \leq \lambda < C} P\{x \leq X(u) < x + (C - \lambda)du, \\ &\qquad \Lambda_u = \lambda \text{ and } V_u = 0\}\end{aligned}$$

which can be rewritten as:

$$\begin{aligned}\nu(x) &= \int_{u>0} \sum_{0 \leq \lambda < C} (C - \lambda) \frac{d}{dx} P\{X(u) \leq x, \\ &\qquad \Lambda_u = \lambda \text{ and } V_u = 0\} du\end{aligned} \tag{17}$$

Now, let $\phi_t(w, \lambda)$ denote the joint density of $A(t)$ and Λ_t:

$$\phi_t(w, \lambda) = \frac{d}{dw} P\{A(t) \leq w, \Lambda_t = \lambda\} \tag{18}$$

Then (17) can be rewritten as:

$$\begin{aligned}\nu(x) &= \int_{u>0} \sum_{0 \leq \lambda < C} (C - \lambda) \phi_u(x + Cu, \lambda) \\ &\qquad P\{V_u = 0 | X(u) = x, \Lambda_u = \lambda\} du\end{aligned} \tag{19}$$

We are unable to compute the conditional probability in the above integral. However, since it is less than 1, we have the upper bound:

$$\nu(x) \leq \int_{u>0} \sum_{0 \leq \lambda < C} (C - \lambda) \phi_u(x + Cu, \lambda) du \tag{20}$$

The quality of the upper bound depends on the value of $P\{V_u = 0 | X(u) = x, \Lambda_u = \lambda\}$. In practice, this quantity is very close to one, since for a high capacity link multiplexing a large number of sources the system behaves like a multiserver system for which the empty queue probability is very much closer to 1 than $1 - \rho$.

To evaluate the integral in inequality (20), we distinguish the contribution of the work $A(t)$ arriving in $[-t, 0)$ of sources which are on at $-t$ and of sources which are off at $-t$. Let $A_{ji}(t)$ be the contribution of source (j, i) and define

$$\begin{aligned}\alpha_j(t, w) &= \frac{d}{dw} P\{A_{ji}(t) \leq w \text{ and } (j, i) \text{ on at } -t\} \\ \beta_j(t, w) &= \frac{d}{dw} P\{A_{ji}(t) \leq w \text{ and } (j, i) \text{ off at } -t\}\end{aligned}$$

Thus, $A(t) = \sum_{j:d_j<t} A_{ji}(t)$ and we can express $\phi_t(w,\lambda)$ as:

$$\phi_t(w,\lambda) = \prod_{j=1}^{m} \binom{k_j}{l_j} \bigoplus_{j=1}^{m} \alpha_j(t,.)^{\oplus(l_j)} \oplus \beta_j(t,.)^{\oplus(k_j-l_j)}(.,w) \quad (21)$$

where

$$m = \max\{j : d_j < t\} \text{ and } l_1, l_2, \dots l_m : \sum_{j=1}^{m} l_j p_j = \lambda \quad (22)$$

and the symbol $\oplus$ denotes convolution (over the variable w while t is held constant).

C. Shifted Normal Approximations

The exact computation of the convolution in (21) is very difficult. To arrive at an approximated solution, we find inspiration from the approach based on the shifted (or tilted) normal approximations, used in [11, pg. 188] and [19] for the computation of a similar expression. Introduce the Laplace transforms $\alpha_j^*(s,t), \beta_j^*(s,t)$ and $\phi_t^*(s,\lambda)$ defined as $\alpha_j^*(s,t) = \int \alpha_j(t,w)e^{-sw}dw$, etc. Taking transforms in (21) gives us:

$$\phi_t^*(s,\lambda) = \prod_{j=1}^{m} \binom{k_j}{l_j} \alpha_j^*(s,t)^{l_j} \beta_j^*(s,t)^{k_j-l_j} \quad (23)$$

where m and l_j's are as given in (22). Now define the shifted density

$$\phi_t^{(s)}(w,\lambda) = \frac{\phi_t(w,\lambda)e^{-sw}}{\phi_t^*(s,\lambda)} \quad (24)$$

The shifted normal approximation technique consists of replacing $\phi_t^{(s)}$ by a central limit approximation and using (24) to provide an estimate of ϕ_t. Since the Gaussian approximation is most accurate about the mean of the distribution, s is chosen so that the mean of $\phi_t^{(s)}(w,\lambda)$ equals w:

$$w = \frac{\int u\phi_t(u,\lambda)e^{-su}du}{\phi_t^*(s,\lambda)}. \quad (25)$$

Equivalently, s is the unique fixed point of the equation

$$w = -\sum_{1\le j\le m} \left(l_j \frac{d\alpha_j^*}{ds}/\alpha_j^* + (k_j - l_j)\frac{d\beta_j^*}{ds}/\beta_j^* \right) \quad (26)$$

where m and l_j's are given by (22). The normal approximation for $\phi_t^{(s)}(w,\lambda)$ is then

$$\phi_t^{(s)}(w,\lambda) \approx \frac{1}{\sqrt{2\pi}\sigma_t} \quad (27)$$

where σ_t^2 is the variance of $\phi_t^{(s)}$ and is given by

$$\sigma_t^2 = \sum_{1\le j\le m} \left[l_j \left(\frac{d^2\alpha_j^*}{ds^2}/\alpha_j^* - \left(\frac{d\alpha_j^*}{ds}/\alpha_j^*\right)^2 \right) + (k_j - l_j) \left(\frac{d^2\beta_j^*}{ds^2}/\beta_j^* - \left(\frac{d\beta_j^*}{ds}/\beta_j^*\right)^2 \right) \right] \quad (28)$$

where m and the l_j's are again obtained from (22) and the value of s from (26).

Finally, we have the approximation which depends only on the Laplace transforms:

$$\phi_t(w,\lambda) = \frac{e^{sw}\phi_t^*(s,\lambda)}{\sqrt{2\pi}\sigma_t} \quad (29)$$

The above equations require the values of $\alpha_j^*(s,t), \beta_j^*(s,t)$ and their derivatives. These are stated in Appendix A for (p,σ,ρ) regulated sources that generate extremal traffic as per Assumption 5. These closed form expressions make it possible to solve (26) numerically for s. The shifted normal approximation in (29) above is then easily evaluated, which allows the integral in (20) to be computed numerically; further, setting the value of z in inequality (20) to 0 yields P_{vio} as per (14).

The complexity of the computation procedure is $O(\prod_{j=1}^{J}(1+k_j))$, since the the summation (over λ) within the integral in inequality (20) requires λ to go through at most $\prod_{j=1}^{J}(1+k_j)$ values. For a small number of classes (recall that a class consists of flows with identical dual leaky bucket descriptions and identical delay bounds at the scheduler), the complexity is manageable (for example, with 2 classes, the complexity is at most $O(n^2)$), but if the number of classes is large the complexity could become prohibitively high. Our current research is addressing means of reducing this computational complexity.

IV. End-to-End Statistical Delay Bounds

The extension of the above single-node analysis technique to the end-to-end case is achieved by introducing traffic reshaping at each node in the network, similar to what was done in the deterministic setting [14], [1]. Thus, for a flow with envelope $A(t)$ and end-to-end delay requirement d (excluding the packetization delays $\sum_{m=1}^{M} L^m/C^m$ where M denotes the number of hops, L^m the maximum packet length at switch m and C^m the link rate at switch m), an appropriate shaper envelope $E(t)$ is selected and the shaping delay d^{sh} computed; the remaining delay $d - d^{sh}$ is split among the schedulers on the path of the flow (the issue of how to split this delay among the schedulers on the path of the flow is an interesting and complex problem, but beyond the scope of this work; here we assume an equal split across the hops). The analytic method of the previous section is used to perform a single-node admission check at each hop, and the flow is admitted only if it is admissible at each hop, thereby guaranteeing it the statistical end-to-end delay bound. Note that the single-node admission check outlined in the previous section requires the traffic from the various flows to be *independent* (Assumption 2). Though this might not hold in general, in networks where only a small fraction of traffic at any switch also shares a downstream switch, the "nodal decomposition" assumption is valid and allows the switches to be analyzed independently (see [4], [16] for a partial justification and verification, though under a different framework).

The choice of a "good" shaper envelope $E(t)$, i.e., one which maximizes the chances of acceptance of the flow into the network, is in general a very difficult design problem. In fact

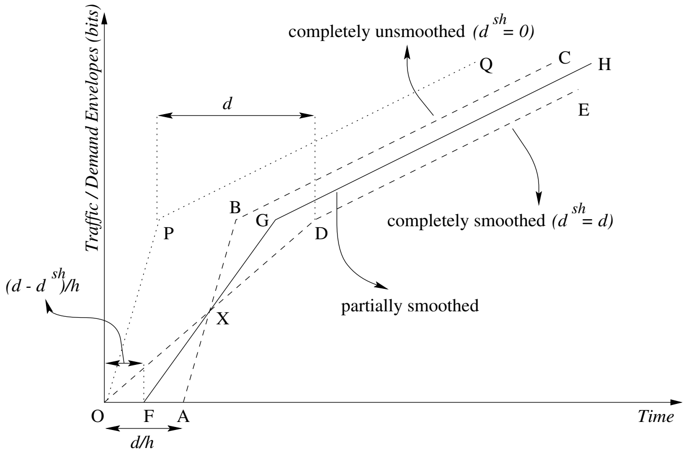

Fig. 2. Demand envelopes for various choices of smoothing parameters

the same problem is encountered in the deterministic framework [14] as well. The difficulty is established using formal arguments in [25]; here we present a simple diagrammatic illustration. Consider an incoming flow with dual-leaky-bucket arrival envelope $A(t)$, depicted by OPQ in Figure 2, and let d denote its end-to-end delay requirement and h the number of hops. Further, recall inequality (1), which requires the aggregate traffic "demand" at the EDF scheduler, denoted by the sum of the traffic envelopes of the flows (right shifted by their delay requirements), to be no more than the aggregate "capacity" Ct at the scheduler, in order for the flows to be schedulable in the deterministic setting. Obviously, the lower the "demand" offered by each flow, the larger is the number of flows that the EDF scheduler can accommodate.

For a given shaping delay budget d^{sh}, where $0 \leq d^{sh} \leq \sigma/\rho$, the smallest multiple-leaky-bucket shaper that guarantees a shaping delay of no more than d^{sh} is the (p', σ', ρ) dual-leaky-bucket shaper, where [14]

$$p' = \frac{p}{1 + d^{sh}(p-\rho)/\sigma}, \quad \sigma' = \sigma\,\frac{(1-\rho/p')}{(1-\rho/p)} \tag{30}$$

Thus, fixing the shaping delay budget d^{sh} completely defines the shaper. There is a wide range of choices for d^{sh}. On one extreme, d^{sh} could be chosen to be 0, in which case $E(t)$ is identical to $A(t)$, i.e., the traffic is left unsmoothed. In this case, the contribution of the flow to the aggregate "demand" at the scheduler is given by $A(t - d/h)$, shown by ABC in Figure 2. On the other extreme, the traffic could be "completely smoothed", i.e., the entire delay budget allocated towards shaping or $d^{sh} = d$. For this case the contribution of the flow traffic to the "demand" at the scheduler is shown by ODE in the figure (note that since all delays have been absorbed by the shaper, no delays are permitted at the scheduler).

Since neither of ABC or ODE lies entirely below the other (i.e., offers lower "demand" at all time scales), none can be claimed to be the better. The unsmoothed case (ABC) offers lower demand at small time scales (to the left of point X), and may be desirable if other traffic at the scheduler has very tight delay requirements; alternatively, the smoothed case (ODE) offers lower demand at larger time scales and is preferable if there is a larger volume of traffic with loose delay requirements. The choice of an "optimal" shaper (one which maximizes the chances of acceptance of the flow), therefore, depends very strongly on the traffic mix at the schedulers on the path of the flow. This, obviously, makes it infeasible to compute.

The infeasibility not withstanding, some general observations provide guidelines for heuristic methods for choosing "reasonably" good shapers. For example, it is known that when the flow traffic traverses only one hop, choosing any $d^{sh} > 0$ is detrimental (since it would yield a "demand" at the scheduler larger than the "unsmoothed demand" ABC at *all* time scales; moreover, the statistical multiplexing advantage at the scheduler would be traded off for deterministic delays at the shaper where there is no statistical multiplexing). On the other hand, when the number of hops traversed is very large, smoothing the traffic completely, i.e., choosing $d^{sh} = d$ is reasonable (since the "completely smoothed" case ODE offers lower demand than any other choice of d^{sh}) This observation is substantiated by formal arguments (though in a restricted setting with homogeneous traffic) by the work in [28], which establishes that when the number of hops is large, smoothing can be beneficial, but when the number of hops is below a critical value, smoothing is detrimental to network performance.

Based on these observations, we propose the following heuristic choice for the shaping delay d^{sh}:

$$d^{sh} = \min\left\{ d\left(1 - \frac{1}{h}\right), \frac{\sigma}{\rho} \right\} \tag{31}$$

The (p', σ', ρ) shaper parameters corresponding to this choice of d^{sh} are computed using (30), and the resulting envelope is depicted in Figure 2 by FGH. Such a choice of shaper has the advantage that it can be computed independent of the exogenous traffic in the network; moreover, by being dependent on the number of hops traversed by the flow, it varies gracefully between the desired limiting cases (when $h = 1$, it can be verified that $d^{sh} = 0$, whereas $h \to \infty$ yields $d^{sh} = d$).

Existing work [14], [1] on shaping in connection with EDF has suggested employing the shaping parameters induced by a rate based scheduler such as GPS. This in effect completely smoothes the traffic at the ingress, i.e., chooses $d^{sh} = d$. As argued above, such a choice is not always desirable, especially when the number of hops traversed by the flow is small. Though traffic mixes can be conceived where complete smoothing outperforms our proposed method based on (31), our numerical investigations confirm the intuition that the latter, in general, yields larger performance gains.

V. Analytical and Simulation Results

To validate the model presented in this paper and to illustrate the advantages over the deterministic analysis [14], [1],

TABLE I

FLOW PARAMETERS

type	k	p(Mbps)	σ(Kbits)	ρ(Mbps)	d(msec)	hops
type-0 (video conference)	varied	10	80	0.5	80	4
type-1 (stored video)	15	10	800	3	180	3

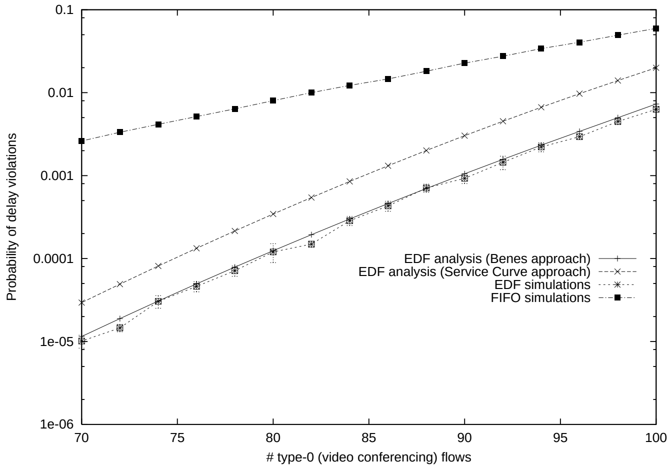

Fig. 3. P_{vio} vs. number of video conferencing flows (no smoothing)

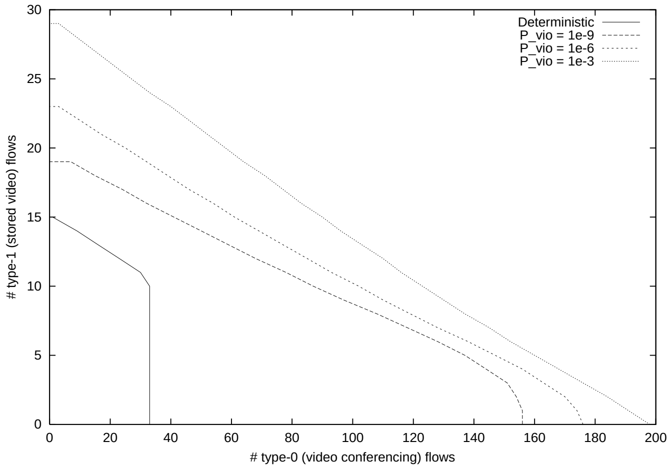

Fig. 4. Schedulable regions under deterministic and statistical QoS

we consider a switch multiplexing a traffic mix consisting of just two types of flows. The parameters for these two flow types are shown in Table I, and are consistent with the ones selected in [13]. For stored video, we use typical values from MPEG traces of a movie like Star Wars, with an average rate of about 3 Mbps and a burst size of 100 Kbytes. The video conference flows have a burst size of 10 Kbytes and an average rate of 1.5 Mbps. The peak rate for both stored video and video conferencing are limited to 10 Mbps (ethernet rate). The stored video flows are assumed to have an end-to-end delay requirement of 180 msec (propagation delays are excluded), and traverse 3 hops, while the video conferencing flows have an end-to-end delay requirement of 80 msec and traverse 4 hops. The link operates at 100 Mbps, and packet sizes are fixed at 10 Kbits (close to the ethernet MTU size). The packet transmission delays are thus 0.1 ms, and are ignored as they are negligible compared to the queueing delay requirements.

For our study, the number of stored video flows is fixed; we then study the effect of varying the number of video conferencing flows on the probability of delay violations at the chosen switch. We assume that the switch considered is the bottleneck for the video conferencing flows; hence the selected switch determines the number of video conferencing flows that can be accepted into the network. In all our experiments, the delay budget $d - d^{sh}$ is split equally among the switches on the path of the flow.

We first consider the case where there is no smoothing within the network, i.e., envelope $E(t) = A(t)$. Thus, no reshaping delays are incurred ($d^{sh} = 0$), and therefore the video conferencing and stored video flows have delay requirements of 20 and 60 msec respectively at the scheduler of interest. Figure 3 shows the delay violation probability as a function of the number of type-0 flows as obtained from 1) our analysis employing the Beneš approach, 2) the analysis using the Service Curve approach in [22], 3) simulations of EDF (95% confidence intervals are also shown), and 4) simulations of First-In-First-Out (FIFO) service (the confidence intervals are too small to be shown). For the Service Curve approach, the variance of $B_j(t)$ is computed using the moments of $\alpha_j^*(s,t) + \beta_j^*(s,t)$ (as given in the appendix) under the assumption of extremal traffic, instead of the Rate Variance (RV) method outlined in [15], since the former was found to yield tighter results. For the simulations of EDF and FIFO, the relative phases of the extremal sources was periodically resynchronized to avoid phase locking.

The plot shows that as expected, EDF can achieve up to multiple orders of magnitude lower delay violation probabilities than FIFO, hence realizing much larger schedulable regions. Moreover, the match between the simulations and our analysis of EDF employing the Beneš approach is almost perfect. The advantage of the statistical framework over the deterministic one is striking. In contrast to the deterministic framework based on (1) which admits at most 1 video conferencing flow (when the number of stored video flows is fixed at 15), thereby restricting the link utilization to around 45%, the statistical framework admits more than 70 video conferencing flows even for as low a delay violation probability as 10^{-5}. This allows link utilizations of around 80% to be realized, a utiliza-

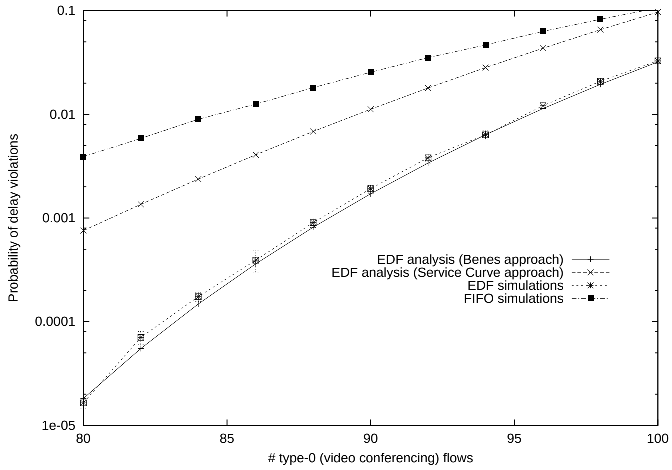

Fig. 5. P_{vio} vs. number of video conferencing flows (with smoothing)

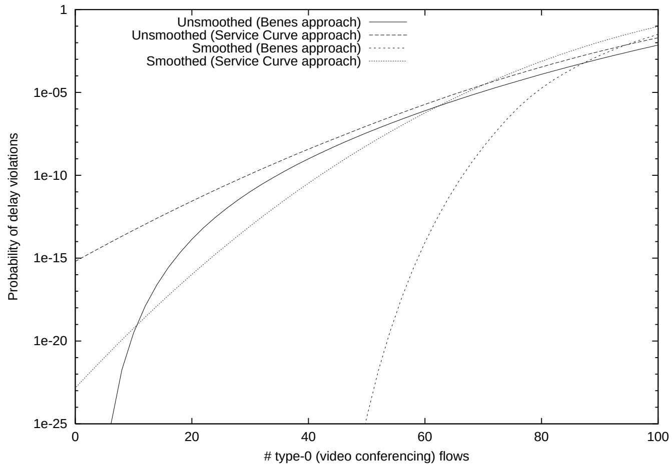

Fig. 6. Comparison between smoothed and unsmoothed traffic

tion gain of more than 35%. The increase in schedulable region is illustrated explicitly in Figure 4 for various delay violation probabilities. The advantages of the statistical framework are immediately obvious from these plots.

The plot of figure 3 also shows the analytic results obtained from the statistical service curve approach of [22]. The match with the EDF simulations is not as close, and the difference is even more pronounced at low probability values (this is evident from Figure 6), indicating that the statistical service curve approach is not as accurate as the Beneš approach. We believe that the reason for this inaccuracy lies not in the service curve framework itself, but in the Gaussian approximations employed for the computation of the probability metric in (3). The Gaussian approximations are notoriously not very accurate when employed for estimating *tail* probabilities; in such situations, the *shifted* normal approximations employed in the Beneš framework are much more accurate, since the Gaussian estimates are consistently applied at the mean rather than the tail of the distributions.

We now investigate the effect of traffic smoothing on the admissibility region of the scheduler. The shaping parameters are derived using (31) and (30). Thus, of the 80 ms delay budget for the video conferencing flows, 60 ms is allocated for smoothing, and the remaining 20 ms split equally among the 4 hops. For the stored video flows, 120 ms of the 180 ms delay budget is allocated towards shaping and the remaining 60 ms split equally among the 3 hops. In the deterministic setting, this smoothing allows 40 video conferencing flows to be admitted into the network (this is a gain of about 20% in utilization over the unsmoothed deterministic setting). For the statistical setting, Figure 5 shows the delay violation probability as a function of the number of type-0 flows, again confirming the excellent matching of our analysis with simulation, while the Service Curve approach is again in this case considerably less accurate. The effect of smoothing on P_{vio} is shown in Figure 6, which shows a comparative plot of P_{vio} vs. number of video conferencing flows for both the smoothed and unsmoothed case. For P_{vio} less than approximately 10^{-3}, the advantages of shaping is quite evident; for example, if $P_{\text{vio}} \approx 10^{-6}$, smoothing allows 76 video conferencing flows at the switch as opposed to only 62 without smoothing, an utilization gain of about 7%. From Figure 6, it can also be seen that as $P_{\text{vio}} \to 0$, the number of video conferencing flows admitted by the Beneš approach converges to the same number as from the deterministic analysis; this convergence is not observed for the statistical service curve approach of [22].

VI. Conclusions and Future Work

In this paper, we have addressed the issue of providing statistical guarantees on end-to-end delays to dual-leaky-bucket-regulated flows in networks employing the EDF scheduling discipline with per-node traffic shaping. We have developed an analytical model for an isolated EDF scheduler fed by extremal regulated sources. Using comparisons with simulations, we have shown that the analysis is very tight, and for the limiting values of the delay violation probabilities, converges to the deterministic model. We have demonstrated that the statistical framework allows to realize much larger schedulable regions than the deterministic approach; for the traffic mix we considered, the gain in utilization was as large as 35% for as low a delay violation probability as 10^{-5}. We have extended our single-node model to the end-to-end case in the presence of smoothing at the ingress and reshaping at each hop. We have proposed a simple heuristic for choosing the smoothing parameters, and shown that it leads to further increases in the schedulable region. We believe our framework could be highly useful in the design of CAC schemes for emerging networks with QoS support.

The main concern regarding the framework proposed in this paper is its computational complexity, which increases with the number of QoS classes. This is to be addressed by our future research. A more thorough comparison of our framework with those existing in literature for both EDF and GPS is also needed, and is addressed in [25].

APPENDIX

I. $\alpha_j^*(s,t)$, $\beta_j^*(s,t)$ AND MOMENTS

We present the expressions for the transforms $\alpha_j^*(s,t)$ and $\beta_j^*(s,t)$ corresponding to the periodic extremal on-off source (j,i). Recall that T_j^o and T_j^f denote respectively the on and off periods of source (j,i), and $T_j = T_j^o + T_j^f$ its period. We define the quantities:

$$x_j = t - T_j \left\lfloor \frac{t}{T_j} \right\rfloor \tag{32}$$

$$a_j = p_j T_j^o \left\lfloor \frac{t}{T_j} \right\rfloor + p_j (x_j - T_j^f)^+ \tag{33}$$

$$b_j = p_j T_j^o \left\lfloor \frac{t}{T_j} \right\rfloor + p_j \min(x_j, T_j^o) \tag{34}$$

Using these, we can express the quantities of interest as:

$$\alpha_j^*(s,t) = e^{-sa_j} \frac{(x_j - T_j^f)^+}{T_j} + e^{-sb_j} \frac{(T_j^o - x_j)^+}{T_j} + \frac{T_j^o - (x_j - T_j^f)^+ - (T_j^o - x_j)^+}{sT_j(b_j - a_j)} \left[e^{-sa_j} - e^{-sb_j} \right] \tag{35}$$

$$\beta_j^*(s,t) = e^{-sa_j} \frac{(T_j^f - x_j)^+}{T_j} + e^{-sb_j} \frac{(x_j - T_j^o)^+}{T_j} + \frac{T_j^f - (x_j - T_j^o)^+ - (T_j^f - x_j)^+}{sT_j(b_j - a_j)} \left[e^{-sa_j} - e^{-sb_j} \right] \tag{36}$$

The moments can be derived directly by differentiating the above expressions, and are not presented here due to space constraints.

REFERENCES

[1] F. M. Chiussi and V. Sivaraman. Achieving high utilization in guaranteed services networks using early-deadline-first scheduling. In *International Workshop on Quality of Service*, pages 209–217, Napa, California, May 1998.

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

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

[4] G. de Veciana, C. Courcoubetis, and J. Walrand. Decoupling bandwidths for networks: a decomposition approach to resource management. In *Proceedings of INFOCOM '94*, pages 466–473, Los Alamitos, CA, June 1994.

[5] G. de Veciana and G. Kesidis. Bandwidth allocation for multiple qualities of service using generalized processor sharing. *IEEE Transactions on Information Theory*, 42:268–272, 1996.

[6] B. T. Doshi. Deterministic rule based traffic descriptors for broadband isdn: Worst case behaviour and its impact on connection acceptance controls. *International Journal of Communication Systems*, 8(2):91–109, March-April 1995.

[7] A. Elwalid, D. Heyman, T. V. Lakshman, D. Mitra, and A. Weiss. Fundamental bounds and approximations for atm multiplexors with applications to video teleconferencing. *IEEE Journal on Selected Areas in Communications*, 13(6):1004–1016, August 1995.

[8] A. Elwalid, D. Mitra, and R. H. 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.

[9] A. I. Elwalid and D. Mitra. Design of generalized processor sharing schedulers which statistically multiplex heterogeneous qos classes. In *Proceedings of INFOCOM '99*, New York, NY, March 1999.

[10] D. Ferrari and D. Verma. A scheme for real-time channel establishment in wide-area networks. *IEEE Journal on Selected Areas in Communications*, 8(3):368–379, April 1990.

[11] R. G. Gallager. *Information Theory and Reliable Communication*. Wiley, New York, 1968.

[12] L. Georgiadis, R. Guérin, and A. Parekh. Optimal multiplexing on a single link: Delay and buffer requirements. *IEEE Transactions on Information Theory*, 43(5):1518–1535, September 1997.

[13] L. Georgiadis, R. Guérin, V. Peris, and R. Rajan. Efficient support of delay and rate guarantees in an internet. In *Proceedings of ACM SIGCOMM*, Stanford University, CA, August 1996.

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

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

[16] W.-C. Lau and S.-C. Li. Traffic analysis in large-scale high-speed integrated networks: Validation of nodal decomposition approach. In *Proceedings of INFOCOM '93*, pages 1320–1329, San Francisco, CA, March 1993.

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

[18] F. LoPresti, Z. Zhang, D. Towsley, and J. Kurose. Source time scale and optimal buffer/bandwidth trade-off for regulated traffic in an atm node. In *Proceedings of INFOCOM '97*, pages 675–682, Kobe, Japan, April 1997.

[19] I. Norros, J. W. Roberts, A. Simonian, and J. T. Virtamo. The superposition of variable bit rate sources in an atm multiplexor. *IEEE Journal on Selected Areas in Communications*, 9(3):378–387, April 1991.

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

[21] A. Parekh and R. Gallager. A generalized processor sharing approach to flow control in integrated services networks: The multiple node case. *IEEE/ACM Transactions on Networking*, 2(2):137–150, April 1994.

[22] J. Qiu and E. D. Knightly. Inter-class resource sharing using statistical service envelopes. Technical report, Rice University, January 1999. An earlier version of this paper appeared in the Proceedings of INFOCOM '99.

[23] J. Roberts, U. Mocci, and J. Virtamo. *Broadband Network Teletraffic: Performance Evaluation and Design of Broadband Multiservice Networks: Final Report of Action COST 242*. Springer, 1996.

[24] V. Sivaraman and F. M. Chiussi. Statistical analysis of delay bound violations at an earliest deadline first (edf) scheduler. *Performance Evaluation*, 36(1):457–470, 1999.

[25] Vijay Sivaraman. *End-to-end Delay Service in High Speed Networks using Earliest Deadline First Scheduling*. PhD thesis, University of California, Los Angeles. In preparation.

[26] D. Verma, H. Zhang, and D. Ferrari. Guaranteeing delay jitter bounds in packet switching networks. In *Proceedings of TRICOMM*, pages 35–46, Chapel Hill, NC, April 1991.

[27] T. Worster. Modelling deterministic queues: The leaky bucket as an arrival process. In *Proceedings of ITC 14*, pages 581–590, Nice, France, June 1994.

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

[29] O. Yaron and M. Sidi. Generalized processor sharing networks with exponentially bounded burstiness arrivals. *Journal of High Speed Networks*, 3(4):375–387, 1994.

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

[31] Z.-L. Zhang, D. Towsley, and J. Kurose. Statistical analysis of the generalized processor sharing scheduling discipline. *IEEE Journal on Selected Areas in Communications*, 13(6):1071–1080, August 1995.