A Network Calculus with Effective Bandwidth

Chengzhi Li, Almut Burchard, Jörg Liebeherr

***Abstract*— This paper establishes a link between two principal tools for the analysis of network traffic, namely, effective bandwidth and network calculus. It is shown that a general version of effective bandwidth can be expressed within the framework of a probabilistic version of the network calculus, where both arrivals and service are specified in terms of probabilistic bounds. By formulating well-known effective bandwidth expressions in terms of probabilistic envelope functions, the developed network calculus can be applied to a wide range of traffic types, including traffic that has self-similar characteristics. As applications, probabilistic lower bounds are presented on the service given by three different scheduling algorithms: Static Priority (SP), Earliest Deadline First (EDF), and Generalized Processor Sharing (GPS). Numerical examples show the impact of specific traffic models and scheduling algorithms on the multiplexing gain in a network.**

I. Introduction

To exploit statistical multiplexing gain of traffic sources in a network, service provisioning requires a framework for the stochastic analysis of network traffic and commonly-used scheduling algorithms. Probably the most influential framework for service provisioning is the *effective bandwidth* (see [14], [15] and references therein), which describes the minimum bandwidth required to provide an expected service for a given amount of traffic. The effective bandwidth of a flow determines a bandwidth somewhere between the average and peak rate of the flow. Effective bandwidth expressions have been derived for many traffic types including those with self-similarity [14].

An alternative method to determine resource requirements of traffic flows in a packet network is the *network calculus*, which takes an envelope approach to describe arrivals and services in a network. Starting with Cruz's seminal work [12] the deterministic network calculus has evolved to an elegant framework for worst-case analysis in a network. Probabilistic extensions of the network calculus, commonly referred to as *statistical network calculus*.

The contribution of this paper is the complete integration of the effective bandwidth theory into the statistical network calculus. As a result of this paper, it is feasible to analyze link scheduling algorithms that are not easily tractable with an effective bandwidth approach, for network traffic types that could previously not be analyzed in a network calculus context. The connections between network calculus and effective bandwidth were first investigated by Chang [8]. This paper continues to explore this relationship, and exploits recent advances in the statistical network calculus to analyze effective bandwidth in a multi-node network.

The network calculus in this paper provides bounds on backlog, delay, and burstiness, from very general description of arrival and service. Specific arrival and service models are inserted at a late stage in the analysis. The advantage of this approach is that it permits us to study the impact of varying scheduling algorithms and arrival models on the multiplexing gain in a network in a single framework. While an analysis that is tailored to specific arrival and service models can lead to tighter bounds, such a direct analysis generally only applies to a single node and is not easily extended to a multi-node setting. A recent paper showed that, in some cases, a network calculus analysis can reproduce bounds obtained with a direct statistical analysis [11].

Extending the deterministic network calculus to a probabilistic setting has shown to be challenging, in particular with respect to a multi-node analysis. In this paper we argue that the availability of a maximum relevant time scale, that is, a bound on the maximum time period at which system events are correlated, is an enabling factor for a statistical calculus analysis. There are numerous scenarios where such time scales can be provided. For example, sometimes it is feasible to provide *a priori* bounds on the busy period at nodes, limits on the maximum buffer lengths at links, or a maximum lifetime of traffic. The analysis in this paper exploits the availability of such time scale bounds, and discusses conditions under which time scale bounds can

The research in this paper is supported in part by the National Science Foundation under grants ANI-0085955, CNS-0435061, ANI-0085955, CNS-0435061, and DMS-0308040, and by two NSERC Discovery grants.

Chengzhi Li (cli@cse.uta.edu) is with the University of Texas at Arlington. Almut Burchard (almut@math.toronto.edu) and Jörg Liebeherr (jorg@comm.utoronto.ca) are with the University of Toronto.

be derived. Finding a multi-node calculus that dispenses with these time scale bounds remains an open question.

The remaining sections are structured as follows. In Section II, we present the statistical network calculus that is used to accommodate effective bandwidth expressions. In Section III, we explore the relationship between effective bandwidth and effective envelopes. This enables us to construct effective envelopes for all traffic models for which effective bandwidth results are available. Specifically, we consider regulated arrivals, a memoryless On-Off traffic model, and a Fractional Brownian Motion traffic model. In Section IV, we derive probabilistic lower bounds on the service offered by the scheduling algorithms SP, EDF, and GPS, in terms of effective service curves. In Section V, we apply the network calculus in a set of examples, and compare the multiplexing gain achievable with the traffic models and scheduling algorithms used in this paper. We present brief conclusions in Section VI.

II. A Network Calculus with Time Scale Bounds

In this section we derive a network calculus that exploits the availability of time scale bounds. Before motivating the need for such bounds, we first introduce necessary notation, and review results from the deterministic and statistical network calculus.

A. Deterministic and Statistical Network Calculus

We consider a discrete time model, where time slots are numbered $0, 1, 2, \ldots$. Arrivals to a network node and departures from a network node in the time interval $[0, t]$ are denoted by nonnegative, nondecreasing functions $A(t)$ and $D(t)$, respectively, with $D(t) \leq A(t)$. The backlog is given by $B(t) = A(t) - D(t)$, and the delay is given by $W(t) = \inf\{d \geq 0 \,|\, A(t-d) \leq D(t)\}$. If $A(t)$ and $D(t)$ are represented as curves, $B(t)$ and $W(t)$, respectively, are the vertical and horizontal differences between the curves.

We use subscripts to distinguish arrivals and departures from different flows or different classes of flows, e.g., $A_i(t)$ denotes the arrivals from flow i, and $A_{\mathcal{C}}(t) = \sum_{i \in \mathcal{C}} A_i(t)$ denotes the arrivals from a collection $\mathcal{C}$ of flows. We use the same convention for the departures, the backlog, and the delay. When we refer to a network with multiple nodes, we use superscripts to distinguish between different nodes, i.e., we use $A_i^h(t)$ to denote the arrivals to the h-th node on the route of flow i, and

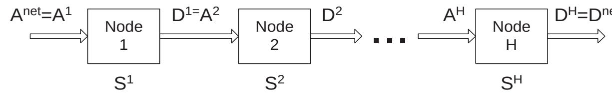

Fig. 1. Traffic of a flow through a set of H nodes. The arrivals and departures from the network are given by random processes A^{net} and D^{net}. The arrivals and departures from the h-th node are described by A^h and D^h, with $A^1 = A^{net}$, $A^h = D_{h-1}$ for $h = 2, \ldots, H$, and $D^{net} = D^H$.

$A_i^{net}(t) = A_i^1(t)$ to denote the arrivals of flow i to the first node on its route. In Figure 1 we show the route of a flow that passes through H nodes, where $A^{net} = A^1$ and $D^{net} = D^H$ denote the arrivals and departures from the network, and where $A^h = D^{h-1}$ for $h = 2, \ldots, H$. To simplify notation, we drop subscripts and superscripts whenever possible. We assume that the network is started at time 0 and that all network queues are empty at this time, i.e., $A_i(0) = D_i(0) = 0$ for all i. Under this assumption, the backlog $B(t)$ increases stochastically with t, in the sense that $Pr(B(t+1) > b) \geq Pr(B(t) > b)$ for all t and all $b \geq 0$, and converges to the steady-state backlog distribution as $t \to \infty$ (see Lemma 9.1.4 of [9]). Thus a stochastic bound on $B(t)$ that does not depend on t provides a bound on the steady-state distribution of the backlog. The corresponding statements hold for the distribution for delays and the departures over time intervals of a given length.

The min-plus algebra formulation of the network calculus [1], [5], [9], defines, for given functions f and g, the *convolution* operator $*$ and *deconvolution* operator $\oslash$ by

$$\begin{aligned} f * g(t) &= \inf_{\tau \in [0,t]} \left\{ f(t-\tau) + g(\tau) \right\} \ , \\ f \oslash g(t) &= \sup_{\tau \geq 0} \left\{ f(t+\tau) - g(\tau) \right\} \ . \end{aligned}$$

These operators are used to express service guarantees and performance guarantees.

B. Overview of the Network Calculus

In the deterministic network calculus in [1], [5], [9], service guarantees to a flow at a node are expressed in terms of *service curves*. A (minimum) service curve for a flow is a nonnegative nondecreasing function S which specifies a lower bound on the service given to the flow such that, for all $t \geq 0$,

$$D(t) \geq A * S(t) \ . \tag{1}$$

When the arrivals are bounded by an *arrival envelope* A^*, such that $A(t+\tau) - A(t) \leq A^*(\tau)$ for all $t, \tau \geq$

0, the guarantee given by the service curve in Eqn. (1) implies worst-case bounds for output burstiness, backlog and delay. According to [1], [5], [9], an envelope for the departures from a node offering a service curve S is given by $A^* \oslash S$, the backlog is bounded by $A^* \oslash S(0)$, and the delay at the node, $W(t)$, is bounded by d, if d satisfies $\sup_{\tau \geq 0}\{A^*(\tau - d) - S(\tau)\} \leq 0$.

If service curves are available at all nodes on the path of a flow through a network, these single-node bounds can be easily extended to end-to-end bounds. Suppose a flow is assigned a service curve S^h on the h-th node on its route ($h = 1, \ldots, H$). Then the service given by the network as a whole can be expressed in terms of a network service curve S^{net} as

$$S^{net} \quad = \quad S^1 * S^2 * \ldots * S^H \ . \tag{2}$$

With a network service curve, bounds for the output burstiness, backlog and delay for the entire network follow directly from the single-node results.

A drawback of the deterministic network calculus is that the consideration of worst-case scenarios ignores the effects of statistical multiplexing, and, therefore, generally leads to an overestimation of the actual resource requirements of multiplexed traffic sources. This has motivated the search for a statistical network calculus, which extends the deterministic calculus to a probabilistic setting with the goal to exploit statistical multiplexing gain. Here, traffic arrivals and departures in the interval $[0, t]$ are viewed as random processes that satisfy certain assumptions, and the arrival and departure functions $A(t)$ and $D(t)$ represent sample paths. In this paper, we assume that arrivals at the network entrance satisfy stationary bounds, in the sense that, for any $\tau > 0$, the arrivals A_i^{net} from any flow i to the network satisfy

$$\lim_{x \to \infty} \sup_{t \geq 0} Pr\Big\{A_i^{net}(t + \tau) - A_i^{net}(t) > x\Big\} = 0 \ .$$

We also assume that the arrivals A_i^{net} and A_j^{net} from different flows $i \neq j$ are stochastically independent.

The assumptions are made only at the network entrance when traffic is arriving to the first node on its route. No such assumptions are made after traffic has entered the network. The stationary bounds are needed so that we can make statements that do not depend on specific instances of time and extend to the steady-state. Assuming independence of traffic sources at the network entrance allows us to exploit statistical multiplexing gain.

We next describe the probabilistic framework used in this paper. We follow the framework for a statistical calculus presented in [4] and [7]. For traffic arrivals, we use a probabilistic measure called *effective envelopes* [4]. An effective envelope for an arrival process A is defined as a non-negative function $\mathcal{G}^\varepsilon$ such that for all t and τ

$$Pr\Big\{A(t + \tau) - A(t) \leq \mathcal{G}^\varepsilon(\tau)\Big\} > 1 - \varepsilon \ . \tag{3}$$

Simply put, an effective envelope provides a stationary bound for an arrival process. Effective envelopes can be obtained for individual flows, as well as for multiplexed arrivals (see Section III below). To characterize the available service to a flow or a collection of flows we use *effective service curves* [7] which can be seen as a probabilistic measure of the available service. Given an arrival process A, an effective service curve is a nonnegative nondecreasing function $\mathcal{S}^\varepsilon$ that satisfies for all $t \geq 0$,

$$Pr\Big\{D(t) \geq A * \mathcal{S}^\varepsilon(t)\Big\} \geq 1 - \varepsilon \ . \tag{4}$$

By letting $\varepsilon \to 0$ in Eqs. (3) and (4), we recover the arrival envelopes and service curves of the deterministic calculus with probability one.

Studies that attempt to construct a statistical network calculus using the min-plus algebra formulation with convolution and deconvolution operators are found in [2], [7]. The challenge in this approach is to construct a probabilistic network service curve that can be expressed as the convolution of per-node service curves, analogous to Eqn. (2). This was pointed out in [7] for a network as shown in Figure 1, with $H = 2$ nodes, and is repeated here. An effective service curve $\mathcal{S}^{2,\varepsilon}$ in the sense of Eqn. (4) at the second node guarantees that, for any given time t, the departures from this node are with high probability bounded below by

$$D^2(t) \geq \inf_{\tau \in [0,t]}\Big\{A^2(t - \tau) + \mathcal{S}^{2,\varepsilon}(\tau)\Big\} \ . \tag{5}$$

Suppose that the infimum in Eqn. (5) is assumed at some value $\hat{\tau} \leq t$. Since the departures from the first node are random, even if the arrivals to the first node satisfy the deterministic bound A^*, $\hat{\tau}$ is a random variable. An effective service curve $\mathcal{S}^{1,\varepsilon}$ at the first node guarantees that for any arbitrary but fixed time x, the arrivals $A^2(x) = D^1(x)$ to the second node are with high probability bounded below by

$$D^1(x) \geq A^1 * \mathcal{S}^{1,\varepsilon}(x) \ . \tag{6}$$

Since $\hat{\tau}$ is a random variable, we cannot simply evaluate Eqn. (6) for $x = t - \hat{\tau}$ and use the resulting bound in

Eqn. (5). (This is different in the deterministic calculus, where deterministic service curves make guarantees that hold for all values of x.) The problem can be resolved if a time scale bound T_{max} is available, which limits the range over which the infimum is taken as follows:

$$A^2 * \mathcal{S}^{2,\varepsilon}(t) = \inf_{\tau \in [0,T_{max}]} \left\{ A^2(t-\tau) + \mathcal{S}^{2,\varepsilon}(\tau) \right\} .$$

C. Network Calculus for Probabilistically Bounded Arrivals and Service

We now present a network calculus that exploits the availability of appropriate time scale limits. The time scale limit is introduced by assuming that service curves $\mathcal{S}^{\varepsilon_s}$ satisfy the additional requirement that there exists a time scale T such that for all $t \geq 0$,

$$Pr\Big\{ D(t) \geq \inf_{\tau \leq T} \{ A(t-\tau) + \mathcal{S}^{\varepsilon}(\tau) \} \Big\} \geq 1 - \varepsilon . \quad (7)$$

Here, T bounds the range of the convolution in Eqn. (4). Since, by assumption, $\mathcal{S}^{\varepsilon}(\tau)$ is nondecreasing in τ and $A(t-\tau) = 0$ for $\tau > t$, the infimum is always attained for some $\tau \leq t$. Hence, Eqn. (7) implies Eqn. (4). This assumption solves the problem of convolving multiple service curves, as discussed in the previous subsection. It turns out that in many networks, in particular, in networks with workconserving schedulers, a time scale bound can be established from probabilistic bounds of the busy period or from constraints of buffer sizes. This will be addressed in Subsection II-D. In general, the value of T depends on the arrival process as well as on the service curve.

The following theorem establishes statistical bounds for delay and backlog in terms of min-plus algebra operations on effective envelopes and effective service curves. Note that we distinguish two violation probabilities: ε_g is the probability that arrivals violate the effective envelope, and ε_s is the probability that the service violates the effective service curve or the condition in Eqn. (7).

Theorem 1: Assume that $\mathcal{G}^{\varepsilon_g}$ is an effective envelope for the arrivals A to a node, and that $\mathcal{S}^{\varepsilon_s}$ is an effective service curve satisfying Eqn. (7) with some $T < \infty$. Define ε to be

$$\varepsilon = \varepsilon_s + T\varepsilon_g . \quad (8)$$

Then the following hold:

1) **Output Traffic Envelope:** The function $\mathcal{G}^{\varepsilon_g} \oslash \mathcal{S}^{\varepsilon_s}$ is an effective envelope for the output traffic from the node.
2) **Backlog Bound:** $\mathcal{G}^{\varepsilon_g} \oslash \mathcal{S}^{\varepsilon_s}(0)$ is a probabilistic bound on the backlog, in the sense that, for all $t \geq 0$, $Pr\Big\{ B(t) \leq \mathcal{G}^{\varepsilon_g} \oslash \mathcal{S}^{\varepsilon_s}(0) \Big\} \geq 1 - \varepsilon$.
3) **Delay Bound:** If $d \geq 0$ satisfies $\sup_{\tau \leq T} \Big\{ \mathcal{G}^{\varepsilon_g}(\tau - d) - \mathcal{S}^{\varepsilon_s}(\tau) \Big\} \leq 0$, then d is a probabilistic delay bound, in the sense that, for all $t \geq 0$, $Pr\Big\{ W(t) \leq d \Big\} \geq 1 - \varepsilon$.

By setting $\varepsilon_s = \varepsilon_g = 0$, we recover the corresponding statements of the deterministic network calculus. Similarly, when only $\varepsilon_g = 0$, the time scale bound T disappears from Eqn. (8) and one can take $T \to \infty$. Thus, the statistical calculus from [7], which deals with deterministic arrivals (where $\varepsilon_g = 0$) and effective service curves $\mathcal{S}^{\varepsilon_s}$, is also recovered by the above theorem.

Proof. We only prove that $\mathcal{G}^{\varepsilon_g} \oslash \mathcal{S}^{\varepsilon_s}$ is an effective envelope for the output traffic. The proofs of the other parts of the theorem are similar [16]. Fix $t, \tau \geq 0$.

$$\begin{aligned}
& Pr\Big\{ D(t+\tau) - D(t) \leq \mathcal{G}^{\varepsilon_g} \oslash \mathcal{S}^{\varepsilon_s}(\tau) \Big\} \\
\geq\ & Pr\Big\{ D(t+\tau) - D(t) \leq \sup_{x \leq T} \{ \mathcal{G}^{\varepsilon_g}(\tau + x) - \mathcal{S}^{\varepsilon_s}(x) \} \Big\} \\
\geq\ & Pr\left\{ \exists x \leq T : \begin{pmatrix} A(t+\tau) - A(t-x) \leq \mathcal{G}^{\varepsilon_g}(\tau + x) \\ \text{and} \ \ D(t) \geq A(t-x) + \mathcal{S}^{\varepsilon_s}(x) \end{pmatrix} \right\} \\
\geq\ & Pr\left\{ \begin{array}{l} \forall x_1 \leq T : \ A(t+\tau) - A(t-x_1) \leq \mathcal{G}^{\varepsilon_g}(\tau + x_1) \\ \text{and} \ \ \exists x_2 \leq T : \ D(t) \geq A(t-x_2) + \mathcal{S}^{\varepsilon_s}(x_2) \end{array} \right\} \\
\geq\ & 1 - \Big(\varepsilon_s + T\varepsilon_g \Big) .
\end{aligned}$$

First, we have expanded the deconvolution operator and reduced the range of the supremum. Then, we replaced $D(t+\tau)$ by $A(t+\tau)$, and added the condition that $D(t) \geq A(t-x) + \mathcal{S}^{\varepsilon_s}(x)$. In the next step, we further restricted the event, by demanding that the first condition in the previous line holds for all values of x. Finally, we apply the assumption in Eqn. (7), and the definition of $\mathcal{G}^{\varepsilon_g}$. Adding the violation probabilities of the two events and using Boole's inequality yields the result. □

Next we derive an expression for a probabilistic version of a network service curve. Consider the path of a flow through a network, as illustrated in Figure 1. At each node $h = 1, \ldots, H$, the arrivals are allotted an effective service curve, $\mathcal{S}^{h,\varepsilon_s}$. Similar to Eqn. (7), we assume that

$$Pr\Big\{ D^h(t) \geq \inf_{\tau \leq T^h} \{ A^h(t-\tau) + \mathcal{S}^{h,\varepsilon_s}(\tau) \} \Big\} \geq 1 - \varepsilon_s \quad (9)$$

for some numbers $T^1, \ldots, T^H < \infty$. For notational convenience, we assume that the violation probabilities ε_s are identical at each node. This assumption is easily relaxed.

Theorem 2: **Effective Network Service Curve.** Assume that the service offered at each node $h = 1, \ldots, H$ on the path of a flow through a network is given by a service curve $\mathcal{S}^{h,\varepsilon_s}$ satisfying Eqn. (9). Then an effective network service curve $\mathcal{S}^{net,\varepsilon}$ for the flow is given by

$$\mathcal{S}^{net,\varepsilon} = \mathcal{S}^{1,\varepsilon_s} * \mathcal{S}^{2,\varepsilon_s} * \ldots * \mathcal{S}^{H,\varepsilon_s} \ , \qquad (10)$$

with violation probability bounded above by

$$\varepsilon = \varepsilon_s \sum_{h=1}^{H} \Big(1 + (h-1) T^h \Big) \ . \qquad (11)$$

The convolution expression in Eqn. (10) has the same form as the corresponding expression in a deterministic setting seen in Eqn. (2), and the deterministic statement is recovered with probability one by letting $\varepsilon \to 0$. On the other hand, the violation probability ε in Eqn. (11) increases at each hop by $\varepsilon_s T^h$. Clearly, it is important to control the time scale bound T^h.

Proof. We start the proof with a deterministic argument for a sample path. Fix $t \geq 0$, and suppose that, for a particular sample path, we have

$$\begin{cases} \text{If } h < H : \forall \tau \leq \sum_{k=h+1}^{H} T^k : \ D^h(t-\tau) \\ \quad \geq \inf_{x_h \leq T^h} \{ A^h(t-\tau-x_h) + \mathcal{S}^{h,\varepsilon_s}(x_h) \} \ , \\ \text{If } h = H : \\ D^H(t) \geq \inf_{x_H \leq T^H} \{ A^H(t-x_H) + \mathcal{S}^{H,\varepsilon_s}(x_H) \} \ , \end{cases} \qquad (12)$$

Since the arrivals at each node are given by the departures from the previous node, that is, $A^h = D^{h-1}$ for $h = 2, \ldots, H$, we see by repeatedly inserting the first line of Eqn. (12) into the second line of Eqn. (12) that

$$D^H(t) \geq \inf_{x_k \leq T^k, \ k=h,\ldots,H} \Big\{ A^k(t - (x_k + \ldots + x_H)) + \sum_{k=h}^{H} \mathcal{S}^{k,\varepsilon_s}(x_k) \Big\} \ . \qquad (13)$$

Setting $h = 1$ in Eqn. (13), and using the definitions of A^{net}, D^{net}, and $\mathcal{S}^{net,\varepsilon}$, we obtain

$$D^{net}(t) \geq A^{net} * \mathcal{S}^{net,\varepsilon}(t). \qquad (14)$$

We conclude proof of the theorem by

$$Pr\Big\{ D^{net}(t) \geq A^{net} * \mathcal{S}^{net,\varepsilon}(t) \Big\} \geq Pr\Big\{ \text{ Eqn. (12) holds } \Big\} \qquad (15)$$

$$\geq \ 1 - \varepsilon_s \cdot \sum_{h=1}^{H} \Big(1 + \sum_{k=h+1}^{H} T^k \Big) \ . \qquad (16)$$

In Eqn. (15) we have used that Eqn. (12) implies Eqn. (14). In Eqn. (16), we have applied Eqn. (9) and added the violation probabilities of Eqn. (12) over all possible values of $h = 1, \ldots, H$. Exchanging the order of summation completes the proof. □

D. Derivation of Time Scale Bounds

We turn to the time scale T which is required in order to apply Theorems 1 and 2. At any given node, we need to bound T from information on the capacity of the node, the properties of the scheduler, and the incoming traffic.

Consider for a moment the corresponding problem in the deterministic calculus. Suppose a node offers a service curve S to a flow, and that the arrivals from the flow are deterministically bounded by an arrival envelope A^*. If the long-term arrival rate is strictly smaller than the long-term service rate guaranteed by S, then

$$T = \sup\{ \tau \geq 0 \mid A^*(\tau) > S(\tau) \} < \infty \ . \qquad (17)$$

A short computation shows that

$$A * S(t) = \inf_{\tau \leq T} \{ A(t-\tau, t) + S(\tau) \} \ ,$$

which, together with the definition of S in Eqn. (1) yields the deterministic statement corresponding to Eqn. (7) with $\varepsilon = 0$. This argument applies to any node along the path of a flow through a network, since the long-term rate of arrivals from the flow to downstream nodes cannot exceed the long-term rate of arrivals from the flow to the ingress node.

In the statistical setting, we restrict the discussion to workconserving schedulers, where the time scale T is bounded by the length of the busy period of the scheduler at time t. To see this, let $A_{\mathcal{C}}(t)$, $D_{\mathcal{C}}(t)$, and $B_{\mathcal{C}}(t)$ denote the aggregate arrivals, the departures, and the backlog of a set $\mathcal{C}$ of flows arriving at the scheduler. By definition, the busy period for a given time $t \geq 0$ is the maximal time interval containing t during which the backlog from the flows in $\mathcal{C}$ remains positive. The beginning of the busy period of t is the last idle time before t, given by

$$\underline{t} = \max\{ \tau \leq t : \ B_{\mathcal{C}}(\tau) = 0 \} \ . \qquad (18)$$

The assumption that the queues are empty at time $t = 0$ guarantees that $0 \leq \underline{t} \leq t$. Since a workconserving scheduler that operates at a constant rate C satisfies

$$D_{\mathcal{C}}(t) \geq A_{\mathcal{C}}(\underline{t}) + C(t - \underline{t})$$

by definition, to obtain the desired time scale bound in Eqn. (7) it suffices to prove that

$$Pr\Big\{ t - \underline{t} \leq T \Big\} \geq 1 - \varepsilon \ . \qquad (19)$$

The following lemma establishes such a busy period bound for a scheduler that operates at a constant rate C. We will show in Section IV that the service available to a *single* flow at many different types of workconserving schedulers can be similarly described by a service curve satisfying Eqn. (7).

Lemma 1: Assume that the aggregate arrivals $A_{\mathcal{C}}$ to a workconserving scheduler with a constant rate C satisfy

$$\sum_{\tau=1}^{\infty} \sup_{t\geq 0} Pr\left\{A_{\mathcal{C}}(t+\tau) - A_{\mathcal{C}}(t) > C\tau\right\} < \infty \ . \quad (20)$$

For a given $\varepsilon \in (0,1)$ choose T large enough so that

$$\sum_{\tau=T+1}^{\infty} \sup_{t\geq 0} Pr\left\{A_{\mathcal{C}}(t+\tau) - A_{\mathcal{C}}(t) > C\tau\right\} \leq \varepsilon \ . \quad (21)$$

Then T is a probabilistic bound on the busy period that satisfies Eqn. (19).

Proof. Fix $t > 0$, and assume that $\underline{t} < t$. Since $B_{\mathcal{C}}(\tau) > 0$ for $\underline{t} < \tau \leq t$, we have by definition of the workconserving scheduler that $D_{\mathcal{C}}(t) - D_{\mathcal{C}}(\underline{t}) \geq C(t - \underline{t})$. Since $D_{\mathcal{C}}(t) < A_{\mathcal{C}}(t)$, and $D_{\mathcal{C}}(\underline{t}) = A_{\mathcal{C}}(\underline{t})$ by definition of $\underline{t}$, this implies $A_{\mathcal{C}}(t) - A_{\mathcal{C}}(\underline{t}) > C(t - \underline{t})$. It follows that

$$\begin{aligned} &Pr\big\{t - \underline{t} > T\big\} \\ &\quad \leq Pr\big\{\exists \tau > T: \ A_{\mathcal{C}}(t) - A_{\mathcal{C}}(t-\tau) > C\tau\big\} \\ &\quad \leq \sum_{\tau=T+1}^{\infty} Pr\big\{A_{\mathcal{C}}(t) - A_{\mathcal{C}}(t-\tau) > C\tau\big\} \quad (22) \\ &\quad \leq \varepsilon \ , \end{aligned}$$

where we have used Boole's inequality in the second line and the choice of T in the third line. □

The lemma is easily extended from constant-rate workconserving systems to output links that offer a (deterministic) *strict service curve*, which is a nonnegative function $S(\tau)$ such that for every $t_2 \geq t_1 \geq 0$ and every sample path, $D_{\mathcal{C}}(t_2) - D_{\mathcal{C}}(t_1) \geq S(t_2 - t_1)$ whenever $B_{\mathcal{C}}(t) > 0$ for $t \in [t_1, t_2]$. This includes, in particular, latency-rate service curves [22] with $S = K(t - L)$ for a rate K and a latency L.

The assumption in Eqn. (20) amounts to two requirements. First, the average rate of the incoming traffic should lie strictly below the rate C of the scheduler. This is a standard stability condition; if it is violated, stability of the backlog process is not guaranteed. Secondly, the probability that the arrivals exceed this average rate by a large amount should satisfy a suitable tail estimate. Such tail estimates hold for many commonly used traffic descriptions, including the models in [8], [21], [24]. This includes some long-range dependent processes but can fail for heavy-tailed arrival models. Examples are discussed in Section III-A.

Inserting Lemma 1 into Theorem 1 immediately provides bounds on output, delay, and backlog for a single node in terms of the arrivals and the available service at that node. Using Lemma 1 in Theorem 2 for the construction of a statistical network service curve is less straightforward. The difficulty is that Theorem 2 requires bounds on the time scales T^h at each node $h = 1, \ldots, H$ on the path of a flow. In principle, such arrival bounds can be obtained by iterating the input-output relation of Theorem 1. However, this approach leads to bounds on the violation probabilities that grow exponentially in the number of nodes.

In the numerical examples, we use the instead the following strategy. We assume that any packet whose delay at a node exceeds a certain delay threshold d^* is dropped. For the arrivals to the network, we construct a function $\overline{\mathcal{G}}_{\mathcal{C}}^{net,\varepsilon}$ for the arrivals to the network satisfying

$$Pr\Big(A_{\mathcal{C}}^{net}(t+\tau) - A_{\mathcal{C}}^{net}(t) > \overline{\mathcal{G}}_{\mathcal{C}}^{net,\varepsilon}(\tau)\Big) \leq \frac{2\varepsilon}{\pi(1+\tau^2)} \ . \quad (23)$$

This definition is analogous to the definition of the effective envelope in Eqn. (3), with ε replaced by $(2\varepsilon)/(\pi(1+\tau^2))$. We define corresponding bounds at downstream nodes by

$$\mathcal{G}_{\mathcal{C}}^{h,\varepsilon}(\tau) = \mathcal{G}_{\mathcal{C}}^{net,\varepsilon}(\tau + (h-1)d^*) \ . \quad (24)$$

At the h-th node on the path of the flow, we set

$$T^h = \sup\{\tau \geq 0 \ \mid \ \overline{\mathcal{G}}_{\mathcal{C}}^{h,\varepsilon}(\tau) > C\tau\} \ , \quad (25)$$

in analogy with Eqn. (17). This choice of T^h satisfies Eqn. (21), because

$$\begin{aligned} &\sum_{\tau=T^h+1}^{\infty} \sup_{t\geq 0} Pr\left\{A_{\mathcal{C}}(t+\tau) - A_{\mathcal{C}}(t) > C\tau\right\} \\ &\leq \quad \sum_{\tau=T^h+1}^{\infty} \sup_{t\geq 0} Pr\left\{A_{\mathcal{C}}^h(t+\tau) - A_{\mathcal{C}}^h(t) > \overline{\mathcal{G}}_{\mathcal{C}}^{h,\varepsilon}(\tau)\right\} \\ &\leq \quad \frac{2\varepsilon}{\pi}\sum_{\tau=0}^{\infty}(1+\tau^2)^{-1} \leq \varepsilon \ . \quad (26) \end{aligned}$$

By Lemma 1, T^h provides the desired time scale bound at the h-th node. Finally, we use Theorems 1 and 2 to verify that d^* is large enough so that the loss rate due to this dropping policy is a small fraction of the traffic rate.

The above assumption on an a priori delay threshold d^* is analogous to an assumption in [3] that all traffic exceeding a certain delay bound is dropped. Bounds for T^h can also be obtained from a priori bounds on the backlog, e.g., as done in [23]. Such bounds on the backlog naturally result from finite buffer sizes in a network. Alternatively, a priori bounds on delay, backlog, and the length of busy periods can be obtained from the deterministic calculus. Generally, it suffices to derive loose bounds on T^h, because the violation probabilities provided in Eqn. (8) and Eqn. (11) depend only linearly on the values of T^h, while effective envelopes $\mathcal{G}^\varepsilon$, the bound $\overline{\mathcal{G}}^{net,\varepsilon}$, and consequently the time scale bound T, typically deteriorate very slowly as $\varepsilon \to 0$.

III. Effective Envelopes and Effective Bandwidth

We now reconcile two methods for probabilistic traffic characterization, effective envelopes and effective bandwidth, and explore the relationship between them. Using the general definition from [14], the *effective bandwidth* of an arrival process A is defined as

$$\alpha(s,\tau) = \sup_{t\geq 0} \left\{ \frac{1}{s\tau} \log E\big[e^{s(A(t+\tau)-A(t))}\big] \right\} \ , \qquad (27)$$

for all $s, \tau \in (0, \infty)$. The parameter τ is called the time parameter and indicates the length of a time interval. The parameter s is called the space parameter and contains information about the distribution of the arrivals. Near $s = 0$, the effective bandwidth is dominated by the mean rate of the traffic, while near $s = \infty$, it is primarily influenced by the peak rate of the traffic. Thus, the space parameter s can be seen as relating to a violation probability ε (see Lemma 2).

A crucial result in the effective bandwidth theory concerns the large buffer asymptotics for links with FIFO scheduling, i.e., as long as the effective bandwidth of a set of flows is below the capacity of the link, the probability of a packet loss due to a buffer overflow decays exponentially fast as a function of the buffer size. This frequently cited result, however, is an asymptotic approximation for large buffer sizes, and has shown to be inaccurate if arriving traffic is bursty [10]. A network calculus approach with effective bandwidth works explicitly with finite buffer sizes. Such non-asymptotic bounds have been presented by Chang [8], [9] for a class of linear envelope processes with parameters $\big(\sigma(s), \rho(s)\big)$, characterized by

$$\frac{1}{s} \log\big(E\big[e^{sA(t,t+\tau)}\big]\big) \leq \sigma(s) + \rho(s)\tau \ . \qquad (28)$$

If $\rho(s) < C$ for these processes, Chang [8] bounds the tail probability of the backlog behavior by $Pr\big(B > x\big) \leq \beta(s) e^{-sx}$, where the constant $\beta(s)$ is explicitly given as $\beta(s) = e^{s\tau(s)}\big(1 - e^{s(\rho(s)-C)}\big)^{-1}$. Chang uses these and other results on envelope processes to draw analogies to the deterministic network calculus [12]. Chang [8] also shows that the output at a link with FIFO scheduling is again a linear envelope processes. In principle, this property can be iteratively applied to obtain delay and backlog bounds for a network with multiple nodes. In practice, however, the bounds obtained with such an iterative procedure deteriorate quickly in the number of nodes. (Closely related results, without referring to effective bandwidth, are obtained by Yaron and Sidi for traffic with exponentially bounded burstiness [24]).

We extend the results established by Chang in several directions. First, we do not assume a specific class of arrival models, but consider all arrival models for which effective bandwidth expressions are available. Second, applying the results of Section II, we can obtain an effective network service curve which yields end-to-end backlog and delay bounds over multiple nodes. Lastly, as shown in Section IV we can apply our analysis to more complex scheduling algorithms.

A. Relating Effective Bandwidth and Effective Envelopes

The choice of the term 'effective envelope' as introduced in [4] suggests a connection to the notion of effective bandwidth, but without making that connection explicit. The following lemma establishes a formal relationship between the two concepts, and thus, links the effective bandwidth theory to the statistical network calculus.

Lemma 2: Given an arrival process A with effective bandwidth $\alpha(s,\tau)$, an effective envelope is given by

$$\mathcal{G}^\varepsilon(\tau) \quad = \quad \inf_{s>0} \big\{ \tau\alpha(s,\tau) - \frac{\log\varepsilon}{s} \big\} \ . \qquad (29)$$

Conversely, if, for each $\varepsilon \in (0,1)$, the function $\mathcal{G}^\varepsilon$ is an effective envelope for the arrival process, then its effective bandwidth is bounded by

$$\alpha(s,\tau) \quad \leq \quad \frac{1}{s\tau} \log\left(\int_0^1 e^{s\mathcal{G}^\varepsilon(\tau)} d\varepsilon \right) \ . \qquad (30)$$

We emphasize that the effective envelope is a more general concept than effective bandwidth, in the sense that each effective bandwidth expression can be immediately expressed in terms of an effective envelope, whereas there may not be an effective bandwidth corresponding to a given effective envelope. Even when the effective bandwidth $\alpha(s,\tau)$ is infinite for some values of s and τ, and the corresponding construction in Lemma 2 is not applicable, it may be feasible to specify a finite effective envelope $\mathcal{G}^{\varepsilon}(\tau)$ according to Eqn. (3) for all values of ε and τ.

Proof. To prove the first statement, fix $t, \tau \geq 0$. By the Chernoff bound [19],[1] we have for any x and any $s \geq 0$

$$Pr\Big\{A(t+\tau) - A(t) \geq x\Big\} \leq e^{s(-x+\tau\alpha(s,\tau))} \ . \quad (31)$$

Setting the right hand side equal to ε and solving for x, we see that, for any choice of $s > 0$, the function

$$x^{\varepsilon,s}(\tau) = \tau\alpha(s,\tau) - \frac{\log\varepsilon}{s}$$

is an effective envelope for A, with violation probability bounded by ε. (The superscripts are added to show the dependence of x on ε and s.) Minimizing over s proves the claim.

For the second statement, fix $t, \tau \geq 0$, and let

$$F^{t,\tau}(x) = Pr\big\{A(t+\tau) - A(t) \leq x\big\} \quad (32)$$

be the probability distribution function of $A(t+\tau) - A(t)$. For any $s > 0$, we can write the moment-generating function of $A(t+\tau) - A(t)$ in the form

$$E\left[e^{s(A(t+\tau)-A(t))}\right] = \int_0^{\infty} e^{sx} dF^{t,\tau}(x) \ . \quad (33)$$

By using a suitable approximation, we may assume without loss of generality that $F^{t,\tau}$ is continuous and strictly increasing for $x \geq 0$. Let $G^{t,\tau}$ be the inverse function of $1 - F^{t,\tau}$. Since

$$Pr\left\{A(t+\tau) - A(t) > G^{t,\tau}(\varepsilon)\right\} = \varepsilon \ , \quad (34)$$

we must have $G^{t,\tau}(\varepsilon) \leq \mathcal{G}^{\varepsilon}(\tau)$ by the definition of the effective envelope. Performing the change of variables $1 - F^{t,\tau}(x) = \varepsilon$, i.e., $x = G^{t,\tau}(\varepsilon)$ in the integral and using that $G^{t,\tau}(\varepsilon) \leq \mathcal{G}^{\varepsilon}(\tau)$, we obtain

$$E\left[e^{s\left(A(t+\tau)-A(t)\right)}\right] \ \leq \ \int_0^1 e^{s\mathcal{G}^{\varepsilon}(\tau)} d\varepsilon \ . \quad (35)$$

By the definition of effective bandwidth, this implies the claim in Eqn. (30). □

We next use the lemma to obtain effective envelopes for regulated arrivals, memoryless on-off traffic, and FBM.

[1]For a random variable X, the Chernoff bound is given by $Pr\Big\{X \geq x\Big\} < e^{-sx} E\left[e^{sX}\right]$.

B. Regulated Arrivals

The regulated arrival model is a suitable description when the amount of traffic that enters the network is limited at the network ingress, e.g., by a leaky bucket. More formally, let A^* be a nondecreasing, nonnegative, subadditive function. We say that an arrival process A is *regulated by* A^* if

$$\forall t, \tau \geq 0: \quad A(t+\tau) - A(t) \leq A^*(\tau) \quad (36)$$

holds for every sample path. The peak rate and the average rate of regulated traffic, denoted by P and ρ, are defined as $P = A^*(1)$ and $\rho = \lim_{t\to\infty} \frac{A^*(t)}{t}$.

Consider a collection $\mathcal{C}$ of flows, where A_i^*, P_i and ρ_i are the arrival envelope, the peak rate, and the average rate of flow $i \in \mathcal{C}$. Clearly, the aggregate of the flows $A_{\mathcal{C}}$ is bounded by $A_{\mathcal{C}}^* = \sum_{i\in\mathcal{C}} A_i^*$, with peak and average rates of $P_{\mathcal{C}} = \sum_{i\in\mathcal{C}} P_i$ and $\rho_{\mathcal{C}} = \sum_{i\in\mathcal{C}} \rho_i$. We assume that each flow $i \in \mathcal{C}$ satisfies the stationary bound

$$E\big[A_i(t+\tau) - A_i(t)\big] \leq \rho_i\tau \ , \quad (37)$$

and that the arrivals from different flows are independent. The effective bandwidth for such a collection of flows $A_{\mathcal{C}}$ satisfies [14]

$$\alpha_{\mathcal{C}}(s,t) \leq \frac{1}{st}\sum_{i\in\mathcal{C}} \log\left(1 + \frac{\rho_i t}{A_i^*(t)}(e^{sA_i^*(t)} - 1)\right) \ . \quad (38)$$

By Lemma 2, the corresponding effective envelope is given by

$$\mathcal{G}_{\mathcal{C}}^{\varepsilon}(t) \ = \ \inf_{s>0}\Big\{\sum_{i\in\mathcal{C}} \frac{1}{s}\log\Big(1+\frac{\rho_i t}{A_i^*(t)}\big(e^{sA_i^*(t)}-1\big)\Big) - \frac{\log\varepsilon}{s}\Big\} \ . \quad (39)$$

This effective envelope satisfies $\rho_{\mathcal{C}} t \leq \mathcal{G}_{\mathcal{C}}^{\varepsilon}(t) \leq A_{\mathcal{C}}^*(t)$ for all $t \geq 0$.

C. Memoryless On-Off traffic

On-Off traffic models are frequently used to model the behavior of (unregulated) compressed voice sources. We consider a variant of On-Off traffic with independent increments. We describe an On-Off traffic source as a two-state memoryless process. In the 'On' state, traffic is produced at the peak rate P, and in the 'Off' state, no traffic is produced, with an overall average traffic rate $\rho < P$. For a collection $\mathcal{C}$ of independent flows with peak rates P_i and average rates ρ_i ($i \in \mathcal{C}$), the effective bandwidth for the aggregate traffic of the flows in $\mathcal{C}$ is given by [14]

$$\alpha_{\mathcal{C}}(s,t) = \frac{1}{s}\sum_{i\in\mathcal{C}} \log\left(1 + \frac{\rho_i}{P_i}\big(e^{P_i s} - 1\big)\right) \ . \quad (40)$$

Lemma 2 gives the corresponding effective envelope as

$$\mathcal{G}^{\varepsilon}_{\mathcal{C}}(t) \;=\; \inf_{s>0}\Big\{\frac{t}{s}\sum_{i\in\mathcal{C}}\log\Big(1+\frac{\rho_i}{P_i}\big(e^{P_i s}-1\big)\Big)-\frac{\log\varepsilon}{s}\Big\}\,. \tag{41}$$

D. Fractional Brownian Motion (FBM) traffic

As pointed out in [18], the self-similarity properties of measured traffic data can be modeled by processes of the form

$$A(t) = \rho t + \beta Z_t\ , \tag{42}$$

where Z_t is a normalized fractional Brownian motion with Hurst parameter $H > \frac{1}{2}$, $\rho > 0$ is the mean traffic rate, and β^2 is the variance of $A(1)$. By definition, $\{Z_t\}_{t\in\mathbb{R}}$ is a Gaussian process with stationary increments which is characterized by its starting point $Z_0 = 0$, expected values $E[Z_t] = 0$, and variances $E[Z_t^2] = |t|^{2H}$ for all t.

Following [18], we will refer to Eqn. (42) as the *Fractional Brownian Motion (FBM)* traffic model. Note that the sum of the arrivals from a collection $\mathcal{C}$ of independent FBM sources with common Hurst parameter is again of type FBM. where the mean traffic rate is given by $\rho_{\mathcal{C}} = \sum_{i\in\mathcal{C}}\rho_i$, and the variance β^2 is given by $\beta^2_{\mathcal{C}} = \sum_{i\in\mathcal{C}}\beta_i^2$. FBM traffic is of interest because the statistical analysis of actual network traffic has shown to be self-similar, that is, traffic exhibits long range dependence [13].

The effective bandwidth for fractional Brownian traffic has been derived as [14]

$$\alpha_{\mathcal{C}}(s,t) = \rho_{\mathcal{C}} + \frac{1}{2}\beta^2_{\mathcal{C}} s t^{2H-1}\ . \tag{43}$$

By Lemma 2, this results in an effective envelope of

$$\mathcal{G}^{\varepsilon}_{\mathcal{C}}(t) \;=\; \rho_{\mathcal{C}} t + \sqrt{-2\log\varepsilon}\,\beta_{\mathcal{C}} t^H\ . \tag{44}$$

IV. Effective Service Curves for Scheduling Algorithms

We next present probabilistic lower bounds on the service guaranteed to a class of flows in terms of effective service curves. We derive effective service curves at a node for a set of well-known scheduling algorithms.

From here on, we assume that each flow belongs to one of Q classes. We denote the arrivals from all flows in class q by A_q, and the arrivals to the collection $\mathcal{C}$ of all flows in all classes $q = 1,\dots,Q$ by $A_{\mathcal{C}}$. We make similar conventions for departures and backlogs. We use $\mathcal{G}^{\varepsilon_g}_q$ to denote an effective envelope for the arrivals from class q. We consider a workconserving link with rate C, and three scheduling algorithms: Static Priorities (SP), Earliest Deadline First (EDF), and Generalized Processor Sharing (GPS). We begin with a brief description of the three schedulers.

1) In an SP scheduler, every class is assigned a priority index, where a lower priority index indicates a higher priority. An SP scheduler selects for transmission the earliest arrival from the highest priority class with a nonzero backlog.
2) In an EDF scheduler, every class q is associated with a delay index d_q. A class-q packet arriving at t is assigned the deadline $t+d_q$, and the EDF scheduler always selects the packet with the smallest deadline for service. Note that, in a probabilistic context, actual delays may violate the delay index and departures may miss their deadlines.
3) In a GPS scheduler, every class q is assigned a weight index ϕ_q and is guaranteed to receive at least a share $\frac{\phi_q}{\sum_p \phi_p}$ of the available capacity. If any class uses less than its share, the extra bandwidth is proportionally shared by all other classes.

For these schedulers, we now present effective service curves for each traffic class q. The effective service curves consider the ‘leftover’ bandwidth which is not used by other traffic classes $p \neq q$. A similar construction was used in the *statistical service envelopes* from [20]. A major difference between statistical service envelopes and our effective service curves is that the latter are non-random functions. This makes the analysis of effective service curves more tractable. In [17] such leftover service curves were used to derive lower bounds on the service for an individual flow when the scheduling algorithms are not known ([6], Chp. 1.4 and Chp. 6.2).

Lemma 3: Consider the arrivals from Q classes to a workconserving scheduler with capacity C. For each class $q = 1,\dots,Q$, let $\mathcal{G}^{\varepsilon_g}_q$ be an effective envelope for the arrivals A_q from flows in class q. Let T be a busy period bound for the aggregate $A_{\mathcal{C}}$ that satisfies Eqn. (19) with some $\varepsilon_b < 1$. Assume the scheduling algorithm employed is either SP, EDF, or GPS. In the case of GPS, assume additionally that the functions $\mathcal{G}^{\varepsilon_p}_p$ are concave. Define functions $\mathcal{S}^{\varepsilon_s}_q$ as follows:[2]

SP: $\mathcal{S}^{\varepsilon_s}_q(t) = \Big[Ct - \sum_{p<q}\mathcal{G}^{\varepsilon_g}_p(t)\Big]_+$,
$\varepsilon_s = \varepsilon_b + (q-1)T\varepsilon_g$.

EDF: $\mathcal{S}^{\varepsilon_s}_q(t) = \Big[Ct - \sum_{p\neq q}\mathcal{G}^{\varepsilon_g}_p(t-[d_p-d_q]_+)\Big]_+$,

[2]We use the notation $[x]_+ = \max(x,0)$ to denote the positive part of x.

$\varepsilon_s = \varepsilon_b + (Q-1)T\varepsilon_g$.

GPS: $\mathcal{S}_q^{\varepsilon_s}(t) = \lambda_q\Big(Ct + \sum_{p\neq q}\Big[\lambda_p Ct - \mathcal{G}_p^{\varepsilon_g}(t)\Big]_+\Big)$, $\varepsilon_s = \varepsilon_b + (Q-1)T\varepsilon_g$, where $\lambda_p = \phi_p / \sum \phi_r$ is the guaranteed share of class p.

Then, in each case $\mathcal{S}_q^{\varepsilon_s}$ is an effective service curve for class q, satisfying

$$Pr\Big\{D_q(t) \geq \inf_{\tau \leq T}\{A_q(t-\tau) + \mathcal{S}_q^{\varepsilon_s}(\tau)\}\Big\} \geq 1 - \varepsilon_s \ . \tag{45}$$

By setting all violation probabilities $\varepsilon_b, \varepsilon_g = 0$ in Lemma 3, we can recover a deterministic (worst-case) statement on the lower bound of the service seen by a service class. The assumption that the scheduler is workconserving is used to establish that the service curves $\mathcal{S}_q^{\varepsilon_s}$ are nonnegative. The lemma easily extends to schedulers offering a strict deterministic service curve S, which need not be constant-rate (see the remark after Lemma 1). In that case, the term Ct should be replaced by $S(t)$ in the conclusions. Given a service curve S satisfying only Eqn. (1), the leftover service curve for class q in the case of an SP scheduler is given by $S(t) - \sum_{p<q}\mathcal{G}_p^{\varepsilon_g}(t)$, which is likely to be negative for small values of t. The corresponding formulas hold for EDF and GPS schedulers.

Proof. Here, we only show that Eqn. (45) holds for the SP and EDF scheduling algorithms. We refer to [16] for the proof of GPS scheduling.

1. SP scheduling: Denote the arrivals from flows of priority at least q by $A_{\leq q}$, and the arrivals from flows of priority higher than q by $A_{<q}$, and correspondingly for departures and backlogs. Fix $t \geq 0$, and let

$$\underline{t}_{\leq q} = \max\big\{x \leq t : \ B_{\leq q}(x) = 0\big\} \tag{46}$$

be the beginning of the busy period containing t from the perspective of class q. If the class-q backlog $B_q(t) = 0$, there is nothing to show. If $B_q(t) > 0$, then we have by the properties of the SP scheduler that

$$\begin{aligned} D_q(t) &= D_q(\underline{t}_{\leq q}) + \big(D_{\leq q}(t) - D_{\leq q}(\underline{t}_{\leq q})\big) \\ &\quad -\big(D_{<q}(t) - D_{<q}(\underline{t}_{\leq q})\big) &(47)\\ &\geq A_q(\underline{t}_{\leq q}) + \Big[C(t-\underline{t}_{\leq q}) - \big(A_{<q}(t) - A_{<q}(\underline{t}_{\leq q})\big)\Big]_+ &(48) \end{aligned}$$

In Eqn. (48), we have used that $D_p(\underline{t}_{\leq q}) = A_p(\underline{t}_{\leq q})$ for all $p \leq q$, that $D(t) - D(\underline{t}_{\leq q}) \geq C(t - \underline{t}_{\leq q})$ by the properties of the workconserving scheduler, and that $D_p(t) \leq A_p(t)$ for all p. It follows that

$$\begin{aligned} &Pr\Big\{D_q(t) \geq \inf_{\tau\leq T}\big(A_q(t-\tau) + \mathcal{S}_q^{\varepsilon_s}(\tau)\big)\Big\} \\ &\geq Pr\Big\{t - \underline{t}_{\leq q} \leq T \text{ and } D_q(t) \geq A_q(\underline{t}_{\leq q}) \\ &\qquad + \Big[C(t-\underline{t}_{\leq q}) - \sum_{p<q}\mathcal{G}_p^{\varepsilon_g}(t-\underline{t}_{\leq q})\Big]_+\Big\} &(49)\\ &\geq Pr\Big\{t - \underline{t}_{\leq q} \leq T \text{ and} \\ &\qquad A_{<q}(t) - A_{<q}(\underline{t}_{\leq q}) \leq \sum_{p<q}\mathcal{G}_p^{\varepsilon_g}(t-\underline{t}_{\leq q})\Big\} &(50)\\ &\geq Pr\Big\{t - \underline{t} \leq T \text{ and} \\ &\qquad \forall p<q, \ \forall\tau\leq T: \ A_p(t) - A_p(t-\tau) \leq \mathcal{G}_p^{\varepsilon_g}(\tau)\Big\} &(51)\\ &\geq 1 - \Big(\varepsilon_b + (q-1)T\varepsilon_g\Big) \ , &(52) \end{aligned}$$

where $\underline{t}$ is the beginning of the busy period of the scheduler. In Eqn. (49), we have set $\tau = t - \underline{t}_{\leq q}$ and inserted the definition of $\mathcal{S}_q^{\varepsilon_s}$, and in Eqn. (50), we have used Eqn. (48). In Eqn. (51), we have restricted the event and used that $\underline{t} \leq \underline{t}_{\leq q}$, and in the last line, we have applied the definitions of T and $\mathcal{G}_p^{\varepsilon_g}$. This proves the claim for SP.

2. EDF scheduling: Fix $t \geq 0$, and let $\underline{t}$ be the beginning of the busy period containing time t. If $B_q(t) > 0$, then according to the EDF scheduling algorithm, class-p packets which arrive after $t + d_q - d_p$ will not be served by time t. Since the system is workconserving, this implies

$$\begin{aligned} D_q(t) &= D_q(\underline{t}) + \big(D_{\mathcal{C}}(t) - D_{\mathcal{C}}(\underline{t})\big) - \sum_{p\neq q}\big(D_p(t) - D_p(\underline{t})\big) \\ &\geq A_q(\underline{t}) + \Big[C(t-\underline{t}) - \sum_{p\neq q}\big(A_p(t-(d_p-d_q)_+) - A_p(\underline{t})\big)\Big]_+ \ . \end{aligned}$$

We argue as in Eqs. (49)-(52) that

$$\begin{aligned} &Pr\Big\{D_q(t) \geq \inf_{\tau\leq T}\Big(A_q(t-\tau) + \mathcal{S}_q^{\varepsilon_s}(\tau)\Big)\Big\} \\ &\geq Pr\Big\{t - \underline{t} \leq T \text{ and } \forall p \neq q, \ \forall \tau \leq T: \\ &\qquad A_p(t) - A_p(t-\tau) \leq \mathcal{G}_p^{\varepsilon_g}(\tau)\Big\} &(53)\\ &\geq 1 - \Big(\varepsilon_b + (Q-1)T\varepsilon_g\Big) \ . \qquad \square &(54) \end{aligned}$$

V. Numerical Examples

In this section, we present numerical examples to illustrate the multiplexing gain for the different traffic models (Regulated, On-Off, Fractional Brownian Motion) and scheduling algorithms (SP, EDF, GPS) considered in this paper.

For each of the three traffic models, we consider two types of flows. The parameters are given in Table I.

	Regulated Traffic			On-Off Traffic		FBM Traffic		
Type	P (Mbps)	ρ (Mbps)	σ σ (bits)	P (Mbps)	ρ (Mbps)	ρ (Mbps)	β (Mbps)	H
1	1.5	0.15	95400	1.5	0.15	0.15	4.5	0.78
2	6.0	0.15	10345	6.0	0.15	0.15	0.94	0.78

TABLE I

Source Traffic Parameters.

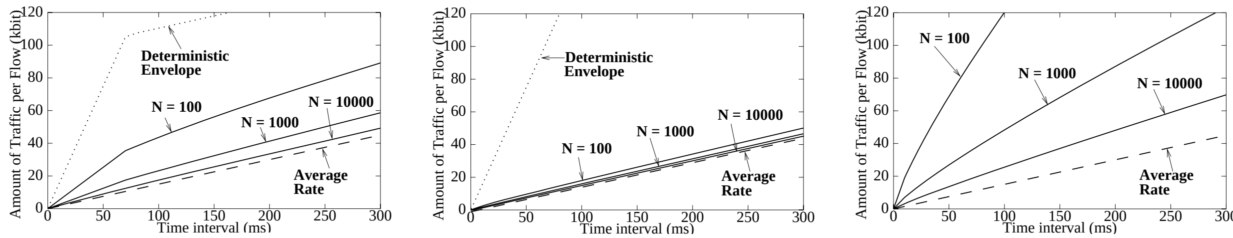

(a) (P, σ, ρ) Regulated. (b) On-Off. (c) FBM.

Fig. 2. **Example 1:** Per-flow effective envelopes $\mathcal{G}^{\varepsilon}_N(t)/N$ for Type-1 flows (with $\varepsilon = 10^{-9}$).

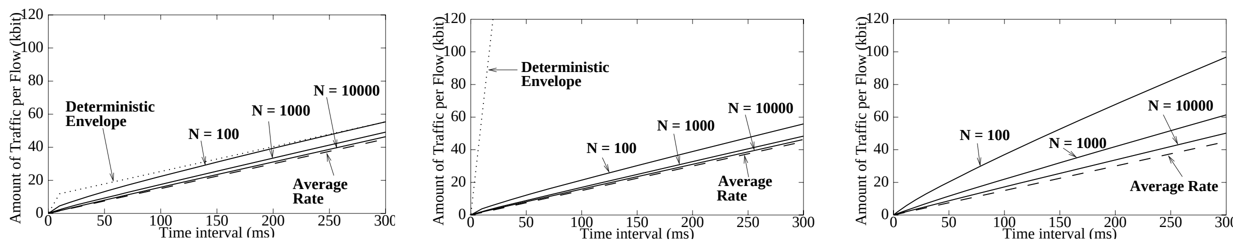

(a) (P, σ, ρ) Regulated. (b) On-Off. (c) FBM.

Fig. 3. **Example 1:** Per-flow effective envelopes $\mathcal{G}^{\varepsilon}_N(t)/N$ for Type-2 flows (with $\varepsilon = 10^{-9}$).

Since we are working in a discrete time domain, we need to select a time unit, which we set to 1 ms. For regulated traffic, we select a peak-rate constrained leaky bucket with arrival envelope $A^*(t) = \min(Pt, \sigma + \rho t)$, with parameters as in [4]. The parameters of the other traffic sources are selected to match the average rate ($\rho = 0.15$ Mbps). For FBM traffic, we set the Hurst parameter to $H = 0.78$ as suggested in [18], and select $\beta = 4.5$.

A. Example 1: Comparison of Effective Envelopes

In the first example, we evaluate the effective envelopes for Regulated traffic, On-Off traffic, and FBM traffic. We evaluate the effective envelope normalized by the number of flows as $\mathcal{G}^{\varepsilon}_N(t)/N$, where $\mathcal{G}^{\varepsilon}_N(t)$ is the effective envelope for N homogeneous flows. Figures 2 and 3 show the per flow effective envelopes with $\varepsilon = 10^{-9}$ for Type-1 and Type-2 flows, respectively. For comparison, we also include the average rate of the sources. For regulated traffic we also include the deterministic envelopes $\min(Pt, \sigma + \rho t)$, and for On-Off traffic we include the peak rate.

We make the following observations. The effective envelopes capture a significant amount of statistical multiplexing gain for each of the considered traffic types, the multiplexing gain increases sharply with the number of flows N. The effective envelope for FBM traffic is larger than for the other source models. This is due to our selection of the parameters H and β.

B. Example 2: Number of Admissible Flows

Next we consider three scheduling algorithms (SP, EDF, and GPS) and multiplex Type-1 and Type-2 flows on a link with 100 Mbps capacity. The evaluation focuses on the service given to flows from Type 1. We assume that Type-1 flows must satisfy a probabilistic delay bound of 100 ms. Given a certain number of Type-2 flows on the 100 Mbps link, we determine the maximum number of Type-1 flows that can be added to the link without violating their probabilistic delay bounds using the results from Lemma 3. Such an admission control decision is greedy, in the sense that it entirely ignores the delay requirements of other flow types. For example, using Lemma 3 for admission control of Type-1 flows ignores the delay requirements of Type-2 flows.

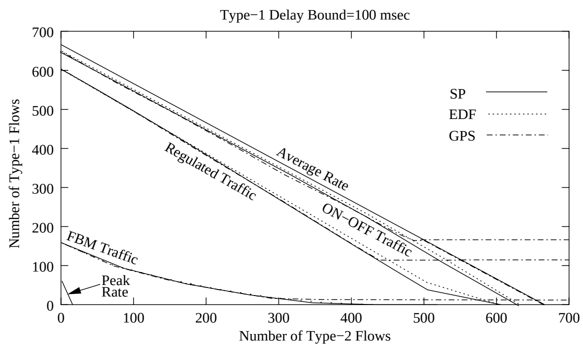

Fig. 4. **Example 2:** Number of admissible Type-1 flows as a function of the number of Type-2 flows (C = 100 Mbps) for different schedulers and traffic models with $\varepsilon = 10^{-6}$, $d_1 = 100$ ms, $\phi_1 = 0.25$, $\phi_2 = 0.75$.

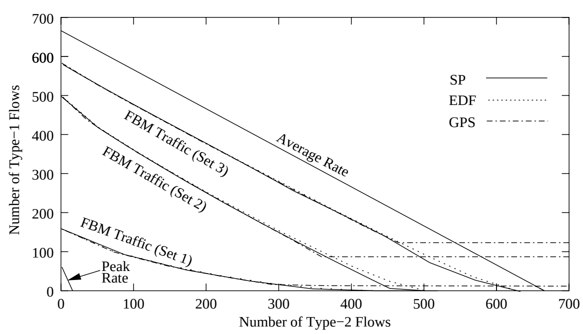

Fig. 5. **Example 2:** Number of admissible Type-1 flows as a function of the number of Type-2 flows (C = 100 Mbps) for FBM traffic with different choices of β with $\varepsilon = 10^{-6}$, $d_1 = 100$ ms, $\phi_1 = 0.25$, $\phi_2 = 0.75$.

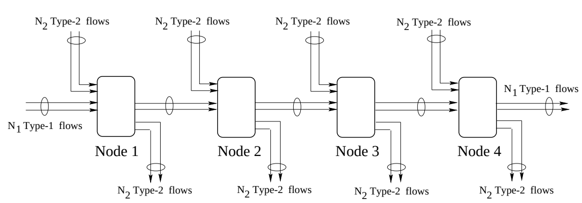

Fig. 6. **Example 3:** A network with four nodes and with cross traffic.

The parameters of the scheduling algorithms are the priority indices for SP, the delay indices for EDF, and the weights for GPS. For SP, Type-1 flows have a higher priority index, and, therefore, a lower priority, than Type-2 flows. For EDF, the delay index of Type-1 flows is $d_1 = 100$ ms and that of Type-2 flows is $d_2 = 10$ ms. For GPS, we set the weights to $\phi_1 = 0.25$ and $\phi_2 = 0.75$. As in the previous examples, we consider three traffic models: regulated traffic, On-Off traffic, and FBM traffic. The source traffic parameters are as shown in Table I. For comparison, we also include the number of flows that can be accommodated on the link with an average rate allocation and a peak rate allocation.

Figure 4 depicts the number of Type-1 flows that can be admitted without violating the probabilistic delay bounds, as a function of the number of Type-2 flows already in the system. We observe that the choice of the traffic model has a significant impact on the number of admitted Type-1 flows. The number of Type-1 flows that can be admitted with FBM traffic is much smaller than with the other traffic models. We also observe in the figure, that the selection of the scheduling algorithm has only a limited impact.

C. Example 3: Multiple Nodes with Cross Traffic.

In this example, we consider a network with four nodes, as shown in Figure 6. We assume that all links have the same capacity of $C = 100$ Mbps. There are N_1 Type-1 flows that pass through all four nodes. At each node, there is cross traffic from N_2 Type-2 flows. We assume $N_1 = N_2$.

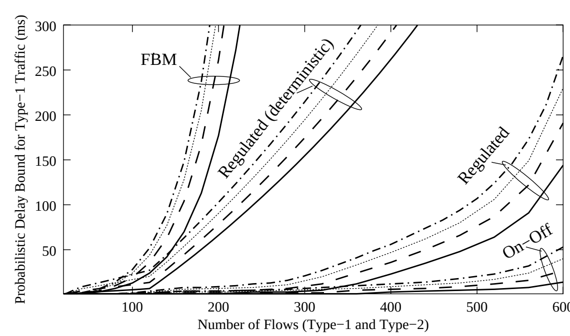

Fig. 8. **Example 3:** Probabilistic bounds for the total queueing delay experienced by Type-1 traffic when leaving Node 1 (solid line), Node 2 (dashed line), Node 3 (dotted line), and Node 4 (dotted-dashed line) with violation probability $\varepsilon = 10^{-6}$. The x-axis corresponds to $N_1 + N_2$, the number of Type-1 and Type-2 flows, where we assume $N_1 = N_2$.

First, we demonstrate how our bounds of the busy period grow as the number of flows increases and how the busy period varies at different nodes. We calculate the probabilistic busy period bounds at the first and the last node for violation probabilities $\varepsilon = 10^{-3}, 10^{-6}, 10^{-9}$ using the approach outlined in Subsection II-D. We

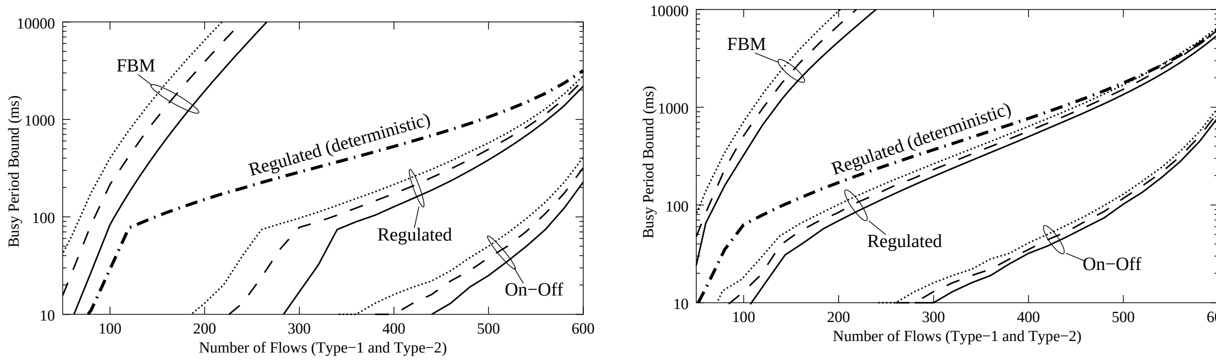

(a) Node 1.

(b) Node 4.

Fig. 7. **Example 3:** Probabilistic Busy Period Bounds for $\varepsilon = 10^{-3}$ (solid line), $\varepsilon = 10^{-6}$ (dashed line), and $\varepsilon = 10^{-9}$ (dotted line). The x-axis corresponds to $N_1 + N_2$, the number of Type-1 and Type-2 flows, where we assume $N_1 = N_2$. The thick dotted-dashed line is a deterministic busy period bound for regulated traffic.

use the formula for the effective envelope given in Eqn. (29), with ε replaced by $\varepsilon/(\pi(1+\tau^2))$ to construct for each class $q = 1, 2$ a function $\overline{\mathcal{G}}_q^{net,\varepsilon/2}$ satisfying $Pr\left\{A^{net}(t) - A^{net}(t-\tau) > \overline{\mathcal{G}}_q^{net,\varepsilon/2}(\tau)\right\} \leq \varepsilon/(\pi(1+\tau^2))$, as required in Eqn. (23). At the h-th node on the route of the through flows, we set $\overline{\mathcal{G}}_1^{h,\varepsilon/2}(\tau) = \overline{\mathcal{G}}_1^{net,\varepsilon/2}(\tau + (h-1)d^*)$, see Eqn. (24). For regulated traffic, we choose the threshold d^* comparable to the worst-case delay bound experienced by the Type-1 traffic at Node 1, as provided by the deterministic calculus. For On-Off and FBM traffic, we choose d^* comparable to the delay bound of Type-1 traffic at Node 1, as provided by Theorem 1 with $\varepsilon = 10^{-15}$. We assume that any packet experiencing a delay exceeding d^* per node is dropped before entering the next node. Since all nodes are ingress nodes for the Type-2 flows, we can use the same bound $\overline{\mathcal{G}}_2^{h,\varepsilon/2}(\tau) = \overline{\mathcal{G}}_2^{net,\varepsilon/2}(\tau)$ at each node, where $\overline{\mathcal{G}}_2^{\varepsilon/2}$ is the function computed above. We obtain bounds on the busy periods T^h by using Eqn. (25) with $\overline{\mathcal{G}}_{\mathcal{C}}^{h,\varepsilon} = \overline{\mathcal{G}}_1^{h,\varepsilon/2} + \overline{\mathcal{G}}_2^{1,\varepsilon/2}$. Finally, we use Theorems 1 and 2 to check that the loss rate due to the dropping threshold never exceeds a fraction of 10^{-15} of the traffic rate.

Figure 7 shows the probabilistic busy period bounds at each node for the three different traffic models, where the number of flows is varied from 60 to 600. Note that 600 flows corresponds to a utilization of 90%. As a reference point, we also plot the exact value for the worst-case busy period of the regulated traffic (plotted as thick dotted-dashed line). While regulated traffic permits to determine the worst-case busy period, such deterministic bounds are not available for On-Off and FBM traffic. We observe that the probabilistic busy period bounds for downstream nodes are larger than that for upstream nodes and that the probabilistic busy period bounds for FBM traffic are significantly larger than those for Regulated or On-Off traffic at each node.

Next, we exhibit the queueing delay experienced by Type-1 traffic in the network described in Figure 6. For the SP scheduling algorithm, as in Example 2, Type-1 flows have a higher priority index, and, therefore, a lower priority, than Type-2 flows. Figure 8 depicts the probabilistic bounds of the total queueing delay experienced by Type-1 traffic when leaving Node h, $h = 1, 2, 3, 4$, with the violation probability 10^{-6} in the network with SP scheduling. The total queueing delay experienced by Type-1 traffic when leaving Node h includes the queueing delay experienced by Type-1 traffic at Node h, Node $h-1$, and down to Node 1. As expected, the probabilistic bounds for the total queueing delay experienced by Type-1 traffic increase when the path traveled by Type-1 traffic increases. As a reference point, we also plot the worst case queueing delay experienced by Regulated traffic. From Figure 8, for Regulated traffic, we observe that the probabilistic bounds for the total queueing delay are dramatically smaller than the worst case queueing delay. Note that the probabilistic bounds for FBM traffic are larger than those for Regulated or On-Off traffic. For EDF and GPS scheduling algorithms, the end-to-end delay bounds experienced by Type-1 traffic in the same network with the violation probability 10^{-6} are similar to those in Figure 8 and omitted.

VI. Conclusions

We have presented a statistical network calculus for determining delays and backlog where both arrivals and service are described in terms of probabilistic bounds.

We presented bounds on the queueing behavior in terms of the min-plus algebra, and integrated the concept of effective bandwidth into the envelope-based approach of the statistical network calculus. We derived backlog and delay bounds for several traffic models (regulated, On-Off, FBM), and scheduling algorithms (SP, EDF, GPS). An important assumption for the derived calculus is the existence of a time-scale bound at each node that decorrelates arrivals and departures. For a single node, such a bound can often be obtained from an estimate on the busy period. For multiple nodes, as seen in Example 3, we require additional assumptions, e.g., that traffic exceeding a maximum delay be dropped. While such an assumption can often be justified, a goal of future work is to determine when and how to dispense with such assumptions.

References

[1] R. Agrawal, R. L. Cruz, C. Okino, and R. Rajan. Performance bounds for flow control protocols. *IEEE/ACM Transactions on Networking*, 7(3):310–323, June 1999.

[2] S. Ayyorgun and R. Cruz. A service-curve model with loss and a multiplexing problem. In *Proc. of ICDCS 2004*, pages 756–765, Tokyo, March 1994.

[3] S. Ayyorgun and W. Feng. A probabilistic definition of burstiness characterization: A systematic approach. Technical Report LA-UR-03-3668, Los Alamos National Laboratory, May 2003.

[4] R. R. Boorstyn, A. Burchard, J. Liebeherr, and C. Oottamakorn. Statistical service assurances for traffic scheduling algorithms. *IEEE Journal on Selected Areas in Communications*, 18(12):2651–2664, December 2000.

[5] J. Y. Le Boudec. Application of network calculus to guaranteed service networks. *IEEE/ACM Transactions on Information Theory*, 44(3):1087–1097, May 1998.

[6] J.-Y. Le Boudec and P. Thiran. *Network calculus*. Springer Verlag, Lecture Notes in Computer Science, LNCS 2050, 2001.

[7] A. Burchard, J. Liebeherr, and S. D. Patek. A calculus for end-to-end statistical service guarantees (revised). Technical Report CS-2001-19, University of Virginia, Computer Science Department, May 2002. Available from http://www.cs.virginia.edu/~jorg/cs-01-19.pdf.

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

[9] C. S. Chang. *Performance guarantees in comunication networks*. Springer, 2000.

[10] G. Choudhury, D. Lucantoni, and W. Whitt. Squeezing the most out of ATM. *IEEE Transactions on Communications*, 44(2):203–217, February 1996.

[11] F. Ciucu, A. Burchard, and J. Liebeherr. A network service curve approach for the stochastic analysis of networks. In *Proc. of ACM Sigmetrics'05*, pages 279–290, 2005.

[12] R. Cruz. A calculus for network delay, parts I and II. *IEEE Transactions on Information Theory*, 37(1):114–141, Jan. 1991.

[13] A. Erramilli, O. Narayan, and W. Willinger. Experimental queueing analysis with long-range dependent packet traffic. *IEEE/ACM Transactions on Networking*, 4(2):209–223, April 1996.

[14] F. Kelly. Notes on effective bandwidths. In *Stochastic Networks: Theory and Applications*. Oxford University Press, 1996.

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

[16] C. Li, A. Burchard, and J. Liebeherr. A network calculus with effective bandwidth. Technical Report CS-2003-20, University of Virginia, Computer Science Department, November 2003.

[17] J. Liebeherr, A. Burchard, and S. D. Patek. Statistical per-flow service bounds in a network with aggregate provisioning. In *Proceedings of IEEE Infocom 2003*, San Francisco, May 2003.

[18] I. Norros. On the use of fractional brownian motion in the theory of connectionless networks. *IEEE Journal on Selected Areas in Communications*, 13(6):953–962, August 1995.

[19] A. Papoulis. *Probability, Random Variables, and Stochastic Processes (3rd ed).* New York: McGraw-Hill, 1991.

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

[21] D. Starobinski and M. Sidi. Stochastically bounded burstiness for communication networks. *IEEE Transactions on Information Theory*, 46(1):206–212, January 2000.

[22] D. Stiliadis and A. Varma. Latency-rate servers: a general model for analysis of traffic scheduling algorithms. *IEEE/ACM Transactions on Networking*, 6(5):611–624, 1998.

[23] M. Vojnovic and J.-Y. Le Boudec. Bounds for independent regulated inputs multiplexed in a service curve network element. *IEEE Transactions on Communications*, 51(5):735–740, May 2003.

[24] 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.

PLACE PHOTO HERE

Chengzhi Li (S'97, M'99) received the Ph.D. degree in computer engineering from Texas A&M University, College Station, in 1999. From 1999 to 2001, he was a Postdoctoral Fellow at Rice University, Houston, TX. From 2001 to 2003, he was a Research Scientist at the University of Virginia, Charlottesville. He is currently a visiting Assistant Professor at the University of Texas at Arlington.

PLACE PHOTO HERE

Almut Burchard received the Ph.D. degree in Mathematics from the Georgia Insitute of Technology in 1994. She was on the faculty of the Department of Mathematics at Princeton University (1994-1998) and the University of Virginia (1994-2005). Since Fall 2005, she is an Associate Professor of Mathematics at the University of Toronto.

PLACE PHOTO HERE

Jörg Liebeherr (S'88, M'92, SM'03) received the Ph.D. degree in Computer Science from the Georgia Institute of Technology in 1991. He was on the faculty of the Department of Computer Science at the University of Virginia from 1992–2005. Since Fall 2005, he is with the University of Toronto as Professor of Electrical and Computer Engineering and Nortel Chair of Network Architecture and Services.