A Network Calculus with Effective Bandwidth *

Chengzhi Li[†] *Almut Burchard*[††] *Jörg Liebeherr*[†]

[†] Department of Computer Science
[††] Department of Mathematics
University of Virginia
Charlottesville, VA 22904

Abstract

We present a statistical network calculus in a setting where both arrivals and service are specified in terms of probabilistic bounds. We provide explicit bounds on delay, backlog, and output burstiness in a node or a network. By formulating well-known effective bandwidth expressions in terms of envelope functions, we are able to apply our calculus to a wide range of traffic source models, including Fractional Brownian Motion. We present probabilistic lower bounds on the service for three scheduling algorithms: Static Priority (SP), Earliest Deadline First (EDF), and Generalized Processor Sharing (GPS).

Key Words: Network calculus, effective bandwidth, Quality-of-Service, statistical multiplexing.

1 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. The most influential such framework is the *effective bandwidth* [11, 12, 13, 14] which, from a qualitative point of view, describes the minimum bandwidth required to provide an expected service for a given amount of traffic. The effective bandwidth concept was related to the theory of large deviations in [4, 25]. However, applications of the effective bandwidth approach have generally been limited to large buffer asymptotics and other asymptotic approximations [16].

In this paper, we take an envelope approach to describe arrivals and services in a network. This approach is motivated by the deterministic network calculus [8] which provides an elegant framework for worst-case analysis in a network. Several researchers have extended the network calculus to a probabilistic setting, including [5, 17, 20, 22, 24, 26, 28, 27, 29]. An advantage of an envelope approach is that it can provide finite bounds on delay and backlog in a network, as opposed to asymptotic approximations.

We present a network calculus in a fully probabilistic setting, where both arrivals and service are expressed in terms of probabilistic bounds. The principal tools of the calculus are *effective envelopes* [28], which are probabilistic upper bounds on arrivals and *effective service curves* [29], which are probabilistic lower bounds on service. By relating the concepts of effective envelopes and effective bandwidth, we obtain

*This work is supported in part by the National Science Foundation through grants ANI-9730103, ANI-9903001, and ANI-0085955, and by an Alfred P. Sloan research fellowship.

explicit bounds on delay and backlog for all traffic source characterizations for which an effective bandwidth (in the sense of [4, 15]) has been determined. Specifically, we consider traffic characterizations that relate to regulated, On-Off, and Fractional Brownian Motion traffic. The effective service curves in this paper, which are inspired by the 'statistical service envelopes' in [20], can express the service for a wide range of scheduling algorithms. In this way, the effective bandwidth theory can be easily related to scheduling algorithms used in practice. Thus, a contribution of this paper is to link two principal tools for the analysis of network traffic, i.e., effective bandwidth and network calculus.

In addition to reconciling effective bandwidth with the envelope approach used in the network calculus, we improve the state-of-the art of statistical network calculus analysis. We generalize [22] by considering a stochastic description of service. In comparison to [29], the network calculus presented here considers other traffic models than regulated traffic.

The paper is structured as follows. In Section 2, we introduce our notation and present definitions of effective envelopes and effective service curves. In Section 3, we propose a statistical network calculus for probabilistic arrivals and service guarantees. The calculus requires the availability of a statistical bound on the busy period bound at a node. Such a bound is derived in Section 4. In Section 5, we explore a duality of the effective bandwidth and the effective envelope. This enables us to construct effective envelopes for all traffic models for which effective bandwidth results are available. Specifically, we consider regulated arrivals, an On-Off traffic model, and a fractional Brownian motion traffic model. In Section 6, we derive probabilistic lower bounds on the service offered by the scheduling algorithms SP, EDF, and GPS, in terms of effective service curve. In Section 7, 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 8.

2 Notation and Definitions

We consider a discrete time model, where time slots are numbered $0, 1, 2, \ldots$. Traffic arrivals and departures in the interval $[0, t]$ are random processes which are defined over a joint probability space which is suppressed in our notation. Sample path are given by nonnegative, nondecreasing functions $A[t]$ for arrivals and $D[t]$ for departures, with $D[t] \leq A[t]$. The backlog at time t is given by $B(t) = A[t] - D[t]$, and the delay at time t is given by $W(t) = \inf\{d \geq 0 \mid A[t-d] \leq D[t]\}$. If A and D 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

$$A_{\mathcal{C}}[t] = \sum_{i \in \mathcal{C}} A_i[t] \tag{1}$$

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 are referring 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 $A_i^{net}[t] = A_i^1[t]$ to denote the arrivals of flow i to the first node on its route. To simplify notation, we drop subscripts and superscripts whenever possible

We assume that the network is started at time 0 and that network queues are empty at this time, i.e., $\forall i$, $A_i[0] = D_i[0] = 0$. Furthermore, we make two fundamental assumptions about the source traffic.

1. *Stationary Bounds*: For any $\tau > 0$, the processes A_i satisfy

$$\lim_{x \to \infty} \sup_{t \geq 0} Pr\big\{ A_i[t+\tau] - A_i[t] > x \big\} = 0 \ .$$

2. *Independence*: A_i and A_j are stochastically independent for all $i \neq j$.

The assumptions are made 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.

We introduce the operators which are used in the min-plus algebra formulation of the deterministic network calculus [1, 3, 7]. For given functions f and g, the convolution operator $*$ and deconvolution operator $\oslash$ are defined by

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

To characterize the available service to a flow or a collection flows we use *effective service curves*, which were recently proposed in [29] as a probabilistic measure of the available service.

Definition 1 Given an arrival process A, an *effective service curve* is a non-negative real-valued function $\mathcal{S}^{\varepsilon}$ that satisfies for all $t \geq 0$,

$$Pr\Big\{ D(t) \geq A * \mathcal{S}^{\varepsilon}(t) \Big\} \geq 1 - \varepsilon \ . \tag{2}$$

By letting $\varepsilon \to 0$, effective service curves recover the service curves of the deterministic calculus with probability one [2].

We also use a probabilistic measure for the traffic arrivals, called effective envelopes from [28].

Definition 2 An *effective envelope* for an arrival process A is 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, an effective envelope provides a stationary bound for an arrival process. We emphasize that both effective service curves and effective envelopes are non-random functions.

3 Statistical Network Calculus

In this section, we state results for a network where arrivals and service are described in terms of probabilistic bounds.

We first consider a single node. We assume that $\mathcal{G}^{\varepsilon_g}$ is an effective envelope curve for the arrivals A to a node, and that $\mathcal{S}^{\varepsilon_s}$ is an effective service curve. For the calculus we make the assumption that there exists a number $T^{\varepsilon_b} < \infty$ such that for all $t \geq 0$,

$$Pr\Big\{ \exists \tau \leq T^{\varepsilon_b} : \ D[t] \geq A[t-\tau] + \mathcal{S}^{\varepsilon_s}(\tau) \Big\} \geq 1 - (\varepsilon_b + \varepsilon_s) \ . \tag{4}$$

T^{ε_b} is bound on the range of the convolution in Eqn. (2) which holds with violation probability ε_b. Thus, Eqn. (4) is a probabilistic bound on the largest relevant time scale that relates arrivals and departures. In a workconserving scheduler, such a bound can be established in terms of a probabilistic bound of the busy period, or from a priori backlog or delay bounds. This will be addressed in Section 4.

The following theorem establishes statistical bounds for queueing delay and backlog in terms of min-plus algebra operations on effective envelopes and effective service curves. Note that we are dealing with three violation probabilities: ε_g is the probability that arrivals violate the effective envelope, ε_s is the probability that the service violates the effective service curve, and ε_b is the probability that the bound on the time scale T^{ε_b} is violated.

Theorem 1 *Assume that $\mathcal{G}^{\varepsilon_g}$ is an effective envelope for the arrivals A to a node, that $\mathcal{S}^{\varepsilon_s}$ is an effective service curve, and that T^{ε_b} satisfies Eqn. (4). Define ε to be*

$$\varepsilon = \varepsilon_s + \varepsilon_b + T^{\varepsilon_b}\varepsilon_g \ .$$

Then the following hold:

1. **Output Traffic Envelope:** *The function $\mathcal{G}^{out,\varepsilon} := \mathcal{G}^{\varepsilon_g} \oslash \mathcal{S}^{\varepsilon_s}$ is an effective envelope for the output traffic from the node. More precisely,*

$$Pr\Big\{\exists x \le T^{\varepsilon_b} : \ D[t+\tau] - D[t] \le \mathcal{G}^{\varepsilon_g}(\tau + x) - \mathcal{S}^{\varepsilon_s}(x)\Big\} \ge 1 - \varepsilon. \tag{5}$$

2. **Backlog Bound:** *For any time $t \ge 0$,*

$$Pr\Big\{B(t) \le \mathcal{G}^{\varepsilon_g} \oslash \mathcal{S}^{\varepsilon_s}(0)\Big\} \ge 1 - \varepsilon \ . \tag{6}$$

3. **Delay Bound:** *If $d \ge 0$ satisfies $\max_{\tau \le T^{\varepsilon_b}} \Big\{\mathcal{G}^{\varepsilon_g}(\tau) - \mathcal{S}^{\varepsilon_s}(d+\tau)\Big\} \le 0$, then for any time t,*

$$Pr\Big\{W(t) \le d\Big\} \ge 1 - \varepsilon \ . \tag{7}$$

By setting $\varepsilon_s = \varepsilon_b = \varepsilon_g = 0$, we recover the corresponding statements of the deterministic network calculus as presented in [2, 6, 1].

Proof. First, we prove that $\mathcal{G}^{\varepsilon_g} \oslash \mathcal{S}^{\varepsilon_s}$ is an effective output traffic envelope. Fix $t, \tau \ge 0$.

$$\begin{aligned}
Pr\Big\{D[t+\tau] - D[t] \le \mathcal{G}^{\varepsilon_g} \oslash \mathcal{S}^{\varepsilon_s}(\tau)\Big\} &\ge Pr\Big\{\exists x \le T^{\varepsilon_b} : \ D[t+\tau] - D[t] \le \mathcal{G}^{\varepsilon_g}(\tau + x) - \mathcal{S}^{\varepsilon_s}(x)\Big\} && (8)\\
&\ge Pr\left\{\exists x \le T^{\varepsilon_b} : \begin{pmatrix} A[t+\tau] - A[t-x] \le \mathcal{G}^{\varepsilon_g}(\tau + x) \\ \text{and } D[t] \ge A[t-x] + \mathcal{S}^{\varepsilon_s}(x) \end{pmatrix}\right\} && (9)\\
&\ge Pr\left\{\begin{array}{l} \forall x_1 \le T^{\varepsilon_b} : \ A[t+\tau] - A[t-x_1] \le \mathcal{G}^{\varepsilon_g}(\tau + x_1) \\ \text{and } \exists x_2 \le T^{\varepsilon_b} : \ D[t] \ge A[t-x_2] + \mathcal{S}^{\varepsilon_s}(x_2) \end{array}\right\} && (10)\\
&\ge 1 - \Big(\varepsilon_s + \varepsilon_b + T^{\varepsilon_b}\varepsilon_g\Big) \ . && (11)
\end{aligned}$$

In Eqn. (8), we have expanded the deconvolution operator and reduced the range of the supremum, i.e., by assuming that the supremum is achieved for a value $x \le T^{\varepsilon_b}$. In Eqn. (9), we replaced $D[t+\tau]$ by $A[t+\tau]$. Further, by adding the condition that $D[t] \ge A[t-x] + \mathcal{S}^{\varepsilon_s}(x)$ we were able to replace $D[t]$ by $A[t-x] + \mathcal{S}^{\varepsilon_s}(x)$. The inequality holds since adding the condition and the replacements restrict the event. In

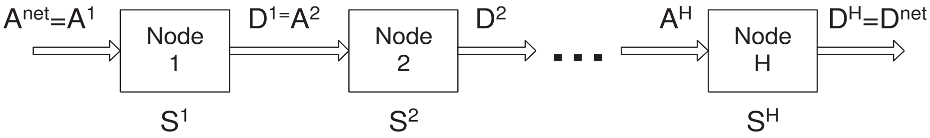

Figure 1: Traffic of a flow through a set of H nodes. Let A^h and D^h denote the arrival and departures at the h-th node, with $A^1 = A^{net}$, $A^h = D^{h-1}$ for $h = 2, \dots, H$ and $D^H = D^{net}$.

Eqn. (10) we further restricted the event, by demanding that the first condition in Eqn. (9) holds for all values of x. To obtain Eqn. (11), we applied the assumption in Eqn. (4), and used the definition of $\mathcal{G}^{\varepsilon_g}$. We added the violation probabilities of the two events using Boole's inequality ($Pr\left\{A \cap B\right\} \geq 1 - Pr\left\{\bar{A}\right\} - Pr\left\{\bar{B}\right\}$ for two events A and B). The factor T^{ε_b} in front of ε_g appears since we added the violation probabilities over all values of x_1.

The proof of the backlog bound proceeds along the same lines. We estimate

$$
\begin{aligned}
Pr\Big\{B(t) \leq \mathcal{G}^{\varepsilon_g} \oslash \mathcal{S}^{\varepsilon_s}(0)\Big\} &= Pr\Big\{A[t] \leq D[t] + \mathcal{G}^{\varepsilon_g} \oslash \mathcal{S}^{\varepsilon_s}(0)\Big\} && (12)\\
&\geq Pr\left\{\exists x \leq T^{\varepsilon_b} : \begin{pmatrix} A[t] \leq A[t-x] + \mathcal{S}^{\varepsilon_s}(x) + \mathcal{G}^{\varepsilon_g} \oslash \mathcal{S}^{\varepsilon_s}(0) \\ \text{and } D[t] \geq A[t-x] + \mathcal{S}^{\varepsilon_s}(x) \end{pmatrix}\right\} && (13)\\
&\geq Pr\left\{ \begin{array}{l} \forall x_1 \leq T^{\varepsilon_b} : \; A[t] - A[t-x_1] \leq \mathcal{G}^{\varepsilon_g}(x_1) \\ \text{and } \exists x_2 \leq T^{\varepsilon_b} : \; D[t] \geq A[t-x_2] + \mathcal{S}^{\varepsilon_s}(x_2) \end{array} \right\} && (14)\\
&\geq 1 - \Big(\varepsilon_s + \varepsilon_b + T^{\varepsilon_b}\varepsilon_g\Big) . && (15)
\end{aligned}
$$

In Eqn. (12), we have used the definition of the backlog $B(t)$. The arguments made in Eqn. (13)-(15) are analogous to those used in Eqn. (9)-(11).

Finally, we prove the delay bound in Eqn. (7). Again, the steps of the proof are similar as before. In Eqn. (16), we use the definition of the delay $W(t)$. The remaining steps apply the same arguments as the proofs of the output bound and the backlog bound.

$$
\begin{aligned}
Pr\Big\{W(t) \leq d\Big\} &= Pr\Big\{A[t-d] \leq D[t]\Big\} && (16)\\
&\geq Pr\left\{\exists x \leq T^{\varepsilon_b} : \begin{pmatrix} A[t-d] \leq A[t-x] + \mathcal{S}^{\varepsilon_s}(x) \\ \text{and } D[t] \geq A[t-x] + \mathcal{S}^{\varepsilon_s}(x) \end{pmatrix}\right\} && (17)\\
&\geq Pr\left\{ \begin{array}{l} \forall x_1 \leq T^{\varepsilon_b} : \; A[t-d] - A[t-x_1] \leq \mathcal{G}^{\varepsilon_g}([x_1 - d]_+) \\ \text{and } \exists x_2 \leq T^{\varepsilon_b} : \; D[t] \geq A[t-x_2] + \mathcal{S}^{\varepsilon_s}(x_2) \end{array} \right\} && (18)\\
&\geq 1 - \Big(\varepsilon_s + \varepsilon_b + T^{\varepsilon_b}\varepsilon_g\Big) . && (19)
\end{aligned}
$$

□

Next we consider multiple nodes. As in the deterministic calculus, the service given by the network as a whole can be expressed as a convolution of the service at each node. Suppose the arrivals $A[t]$ to a flow pass through H nodes, labeled $h = 1, \dots, H$, in series, as shown in Figure 1. Let $A^{net} = A^1$ and $D^{net} = D^H$ denote the arrivals and departures from the network, and let $A^h = D^{h-1}$ for $h = 2, \dots, H$.

At each node, the arrivals are allotted an effective service curve, where $\mathcal{S}^{h,\varepsilon_s}$ denotes the effective service curve at node h. Similar to Eqn. (4), we assume that each node satisfies

$$
Pr\Big\{\exists \tau \leq T^{h,\varepsilon_b} : \; D^h[t] \geq A^h[t-\tau] + \mathcal{S}^{h,\varepsilon_s}(\tau)\Big\} \geq 1 - \varepsilon_h \qquad (20)
$$

for some numbers $T^{h,\varepsilon}$. For notational convenience, we assume that the violation probabilities ε_s and ε_b are identical at each node. This assumption is easily relaxed.

Theorem 2 Effective Network Service Curve. *Assume that $\mathcal{S}^{h,\varepsilon_s}$ is an effective service curve for node h that satisfies Eqn. (20) for all $t \geq 0$ and $h = 1, \ldots, H$ with some numbers T^{h,ε_b}. Then, an effective service curve for the sequence of nodes is given by*

$$\mathcal{S}^{net,\varepsilon} = \mathcal{S}^{1,\varepsilon_s} * \mathcal{S}^{2,\varepsilon_s} * \ldots * \mathcal{S}^{H,\varepsilon_s} , \tag{21}$$

with violation probability bounded above by

$$\varepsilon = \varepsilon_s \sum_{h=1}^{H} \Big(1 + (h-1) T^{h,\varepsilon_b}\Big) . \tag{22}$$

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

$$\begin{cases} \forall \tau \leq \sum_{k=h+1}^{H} T^{k,\varepsilon_b} \; \exists x_h \leq T^{h,\varepsilon_b} : \; D^h[t-\tau] \geq A^h[t-\tau-x_h] + \mathcal{S}^{h,\varepsilon_s}(x_h) \, , & \text{if } h < H, \\ \exists x_H \leq T^{H,\varepsilon_b} : \; D^H[t] \geq A^H[t-x_H] + \mathcal{S}^{H,\varepsilon_s}(x_H) \, , & \text{if } h = H \, . \end{cases} \tag{23}$$

Since the departures from the $(h-1)$-th node are the arrivals at the h-th node, that is, $A^h = D^{h-1}$ for $h = 2, \ldots, H$, we see by repeatedly inserting the first line of Eqn. (23) into the second line of Eqn. (23) that there exist numbers $x_h \leq T^{h,\varepsilon_b}$ such that

$$D^H[t] \geq A^h[t - (x_h + \ldots + x_H)] + \sum_{k=h}^{H} \mathcal{S}^{k,\varepsilon_s}(x_k) \, . \tag{24}$$

Setting $h = 1$ in Eqn. (24), and using the definitions of A^{net}, D^{net}, and $\mathcal{S}^{net,\varepsilon}$, we obtain

$$\exists x_1 \leq T^{1,\varepsilon_b}, \ldots, \exists x_H \leq T^{H,\varepsilon_b} \; : \; D^{net}[t] \geq A^{net}[t - (x_1 + \cdots + x_H)] + \mathcal{S}^{net,\varepsilon}(x_1 + \cdots + x_H) \, . \tag{25}$$

Thus, we have shown that Eqn. (23) implies Eqn. (25). This argument will be used below.
To proof the claim of the theorem we will show that

$$Pr\Big\{\exists \tau \leq \sum_{h=1}^{H} T^{h,\varepsilon_b} : \; D^{net}[t] \geq A^{net}[t-\tau] + \mathcal{S}^{net,\varepsilon}(\tau)\Big\} \quad \geq \quad \varepsilon \tag{26}$$

The proof is as follows:

$$\begin{aligned} Pr\Big\{\exists \tau \leq \sum_{h=1}^{H} T^{h,\varepsilon_b} : \; D^{net}[t] \geq A^{net}[t-\tau] + \mathcal{S}^{net,\varepsilon}(\tau)\Big\} \quad &\geq \quad Pr\Big\{ \text{ Eqn. (25) holds } \Big\} && (27) \\ &\geq \; Pr\Big\{ \text{ Eqn. (23) holds } \Big\} && (28) \\ &\geq \; 1 - \varepsilon_s \cdot \sum_{h=1}^{H} \Big(1 + \sum_{k=h+1}^{H} T^{k,\varepsilon_b}\Big) \, . && (29) \end{aligned}$$

In Eqn. (27), we have set $\tau = x_1 + \ldots + x_H$. In the next step, we have used that Eqn. (23) implies Eqn. (25), as explained above. In Eqn. (29), we have applied Eqn. (20) and added the violation probabilities of Eqn. (23) over all possible values of $h = 1, \ldots, H$. Lastly, exchanging the order of summation completes the proof.

$\square$

The presented results assume that bounds T^{ε_b} are readily available. In the next section, we obtain these bounds from bounds on the busy period at a workconserving scheduler.

4 Busy Period Analysis

In the network calculus from the previous section, the time scale T^{ε_b} in Eqn. (4) played a central role. For workconserving schedulers, a bound on this time scale can be obtained by bounding the busy period of the scheduler, where a busy period for a given time t is the maximal time interval containing t during which the backlog from the flows in $\mathcal{C}$ remains positive. For $t \geq 0$, define the last idle time before t by

$$\underline{t} = \max\{\tau \leq t : \ B(\tau) = 0\} \ . \tag{30}$$

Our assumption that $B(0) = 0$ guarantees that $0 \leq \underline{t} \leq t$.

In this section, we bound the time scale T^{ε_b} in terms of a bound on the busy period. Alternatively, one can derive a bound on T^{ε_b} from bounds on the backlog, e.g., using those derived in [24].

Lemma 1 *For an arrival process A and a workconserving scheduler with a constant rate C, assume that*

$$\sum_{\tau=1}^{\infty} \sup_{t \geq 0} Pr\left\{A[t+\tau] - A[t] > C\tau\right\} < \infty \ . \tag{31}$$

For $\varepsilon \in (0,1)$, choose T^{ε} such that

$$\sum_{\tau=T^{\varepsilon}+1}^{\infty} \sup_{t \geq 0} Pr\left\{A[t+\tau] - A[t] > C\tau\right\} \leq \varepsilon \ . \tag{32}$$

Then T^{ε} is a probabilistic bound on the busy period, satisfying

$$Pr\big\{t - \underline{t} \leq T^{\varepsilon}\big\} \geq 1 - \varepsilon \ . \tag{33}$$

The lemma is easily extended to output links that offer a (deterministic) *strict service curve*, which is defined for an arrival process A as a nonnegative function such that $B(t) > 0$ for $t \in [t_1, t_2]$ implies that $D[t_2] - D[t_1] \geq S(t_2 - t_1)$ for each sample path and every $t_2 > t_1 \geq 0$, such as *latency-rate service curves* [23] with $S = K(t - \Delta)$ for a rate K and a delay Δ.

Proof. Fix $t > 0$. Since $B(\tau) > 0$ for $\underline{t} < \tau \leq t$, we have by definition of the workconserving scheduler, that $D(t) - D(\underline{t}) \geq C(t - \underline{t})$. Since $D(t) \leq A(t)$, and $D(\underline{t}) = A(\underline{t})$ by definition of $\underline{t}$, this implies $A(t) - A(\underline{t}) \geq C(t - \underline{t})$ with equality only when $B(t) = 0$, in which case $\underline{t} = t$. It follows that

$$\begin{aligned} Pr\big\{t - \underline{t} > T^{\varepsilon}\big\} &\leq Pr\big\{\exists \tau > T^{\varepsilon} : \ A[t] - A[t-\tau] > C\tau\big\} && (34) \\ &\leq \sum_{\tau=T^{\varepsilon}+1}^{\infty} Pr\big\{A[t] - A[t-\tau] > C\tau\big\} && (35) \\ &\leq \varepsilon \ , && (36) \end{aligned}$$

where we have used the definition of T^{ε} in the third line. □

Lemma 1 enables us to bound the tail distribution of a server's busy periods from bounds on the distribution function of the arrival process. This bound can be used in Theorem 1. In the next Section we use Lemma 1 to determine busy period bounds for Regulated traffic, On-Off traffic, and Fractional Brownian Motion traffic.

5 Effective Envelopes and Effective Bandwidth

In this section, we reconcile two methods for probabilistic traffic characterization, *effective envelopes* and *effective bandwidth*, and explore the relationship between them. The notion of an *effective envelope*, which was introduced in [28], is motivated by the traffic envelopes used in the deterministic calculus. The *effective bandwidth*, which has been extensively studied, is motivated by the rate functions that appear in the theory of large deviations. Effective bandwidth expressions have been derived for numerous source traffic models with applications in computer networks. We refer to [7, 15, 16] for a detailed discussion. By providing a link between effective bandwidth and effective envelopes, the results in this section make effective bandwidth results applicable to the network calculus.

The following definition was first provided in [4].

Definition 3 The *effective bandwidth* of an arrival process A that satisfies a stationary bound is given by

$$\alpha(s,\tau) = \sup_{t\geq 0} \left\{ \frac{1}{s\tau} \log E[e^{s(A[t+\tau]-A[t])}] \right\} , \qquad s,\tau \in (0,\infty) . \tag{37}$$

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.

In the following lemma, we establish an approximate duality between effective envelopes and effective bandwidth.

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

$$\mathcal{G}^{\varepsilon}(\tau) \quad = \quad \inf_{s>0} \big\{ \tau\alpha(s,\tau) - \frac{\log \varepsilon}{s} \big\} . \tag{38}$$

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

$$\alpha(s,\tau) \quad \leq \quad \frac{1}{s\tau} \log \left(\int_0^1 e^{s\mathcal{G}^{\varepsilon}(\tau)} d\varepsilon \right) . \tag{39}$$

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

$$Pr\Big\{ A[t+\tau] - A[t] \geq x \Big\} \quad \leq \quad e^{-sx} E\left[e^{s(A[t+\tau]-A[t])} \right] \tag{40}$$

$$\leq \quad e^{s(-x+\tau\alpha(s,\tau))} . \tag{41}$$

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

$$x^{\varepsilon,s}(\tau) = \tau\alpha(s,\tau) - \frac{\log \varepsilon}{s} \tag{42}$$

is an effective envelope for A, with violation probability bounded by ε. Minimizing over s proves the claim.

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

$$F^{t,\tau}(x) = Pr\big\{ A[t+\tau] - A[t] \leq x \big\} \tag{43}$$

be the 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

$$E\left[e^{s(A[t+\tau]-A[t])}\right] = \int_0^\infty e^{sx} dF^{t,\tau}(x) \ . \tag{44}$$

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

$$Pr\left\{A[t+\tau]-A[t] > G^{t\tau}(\varepsilon)\right\} = \varepsilon \ , \tag{45}$$

we must have $G^{t,\tau}(\varepsilon) \geq \mathcal{G}^\varepsilon(\tau)$ 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, we obtain

$$E\left[e^{s\left(A[t+\tau]-A[t]\right)}\right] = \int_0^1 e^{sG^{t,\tau}(\varepsilon)} d\varepsilon \ \leq \int_0^1 e^{s\mathcal{G}^\varepsilon(\tau)} d\varepsilon \ . \tag{46}$$

It follows that

$$\alpha(s,\tau) \leq \frac{1}{s\tau} \int_0^1 e^{s\mathcal{G}^\varepsilon(\tau)} d\varepsilon \ , \tag{47}$$

as claimed. □

With this lemma we can construct an effective envelope for a traffic class if its effective bandwidth is known. Since many effective bandwidth formulas have been provided in the literature (e.g., [7, 15]), Lemma 2 provides a useful tool to apply the presented network calculus to a wide range network traffic.

We emphasize that the effective envelope is a more general concept than effective bandwdith. 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 an effective envelope $\mathcal{G}^\varepsilon(\tau)$ according to Definition 2 which is finite for all values of ε and τ.

Next we derive effective envelopes and busy period estimates for three arrival models, by applying Lemmas 1 and 2.

5.1 Regulated Arrivals

Let A^* be a nondecreasing, nonnegative, subadditive function. We say that an arrival process A is *regulated by* A^* if

$$\forall t,\tau \geq 0: \quad A[t+\tau]-A[t] \leq A^*(\tau) \tag{48}$$

holds for every sample path. Using P and ρ, respectively, to denote the peak rate and the long-time average rate of regulated traffic, Eqn. (48) implies that P and ρ are bounded by

$$P = A^*(1) \ , \quad \rho = \lim_{t\to\infty} \frac{A^*(t)}{t} \ . \tag{49}$$

The regulated arrival model is a suitable description when the amount of traffic that enters the network is limited at the network edge, e.g., by a leaky bucket, and has been studied extensively [9, 10, 21].

Consider a collection $\mathcal{C}$ of flows which are regulated by subadditive functions A_i^*, and let P_i and ρ_i be the peak rate and average rate constraints implied by A_i^*. 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 longtime average rates bounded by $P_{\mathcal{C}} = \sum_{i\in\mathcal{C}} P_i$ and $\rho_{\mathcal{C}} = \sum_{i\in\mathcal{C}} \rho_i$. We assume that each flow $i \in \mathcal{C}$ satisfies the stationary bound

$$E\big(A_i[t+\tau] - A_i[t]\big) \leq \rho_i \tau \ , \tag{50}$$

and that the arrivals from different flows are independent. The effective bandwidth for such a collection of flows $A_{\mathcal{C}}$ satisfies [15]

$$\alpha_{\mathcal{C}}(s,t) \leq \frac{1}{st} \sum_{i\in\mathcal{C}} \log\left(1 + \frac{\rho_i t}{A^*_i(t)}(e^{sA^*_i(t)} - 1)\right) \ . \tag{51}$$

By Lemma 2, the corresponding effective envelope is given by

$$\mathcal{G}^{\varepsilon}_{\mathcal{C}}(t) \ = \ \inf_{s>0}\Big\{\sum_{i\in\mathcal{C}} \frac{1}{s}\log\Big(1 + \frac{\rho_i t}{A^*_i(t)}\big(e^{sA^*_i(t)} - 1\big)\Big) - \frac{\log\varepsilon}{s}\Big\} \ . \tag{52}$$

This effective envelope satisfies

$$\rho_{\mathcal{C}} t \leq \mathcal{G}^{\varepsilon}_{\mathcal{C}}(t) \leq A^*_{\mathcal{C}}(t) \tag{53}$$

for all $t \geq 0$.

We next construct a probabilistic bound for the busy period of a collection of regulated flows at a workconserving link with a fixed-rate capacity C, by verifying the assumptions of Lemma 1. We consider only the situation where $\rho_{\mathcal{C}} < C < P_{\mathcal{C}}$, since otherwise the expected backlog is either infinite or zero. Under these conditions, a deterministic bound on the busy period is given by $t - \underline{t} \leq T^0$, where $T^0 = \inf\{t > 0 : \ A^*_C(t) \leq Ct\}$. Since $t - \underline{t} \leq T^0$ for all t, the sum in Eqn. (31) contains only finitely many nonzero terms, and the first condition of Lemma 1 is satisfied. By Lemma 1, a probabilistic bound on the busy period is given by any number T^{ε} satisfying

$$\sum_{\tau=T^{\varepsilon}+1}^{T^0} \sup_{t\geq 0} Pr\big\{A_{\mathcal{C}}[t+\tau] - A_{\mathcal{C}}[\tau] > C\tau\big\} \leq \varepsilon \ . \tag{54}$$

In order to estimate the left hand side of Eqn. (54), we use the Chernoff bound

$$\begin{aligned} Pr\big\{A_{\mathcal{C}}[t+\tau] - A_{\mathcal{C}}[t] > C\tau\big\} \ &\leq \ \inf_{s\geq 0} e^{s\tau(\alpha_{\mathcal{C}}(s,\tau) - C)} && (55)\\ &\leq \ \inf_{s\geq 0}\Big\{e^{-Cs\tau}\prod_{i\in\mathcal{C}}\Big(1 + \frac{\rho_i\tau}{A^*_i(\tau)}(e^{sA^*_i(\tau)} - 1)\Big)\Big\} \ . && (56)\end{aligned}$$

In general, the infimum on the right hand side of Eqn. (56) cannot be expressed analytically. However, since the objective function is strictly convex in s, the minimum in Eqn. (56) can be found numerically by standard methods.

Ideally, we would like to find T^{ε} as small as possible. In order to obtain a rough quantitative estimate for T^{ε} and its dependence on the link rate C and the parameters of the flows, we use the peak rate constraint $A^*_i(\tau) \leq P_i\tau$ and replace the variable s by s/τ in Eqn. (56), which results in the simpler, looser bound

$$Pr\big\{A_{\mathcal{C}}[t+\tau] - A_{\mathcal{C}}[t] > C\tau\big\} = \inf_{s\geq 0}\Big\{e^{-Cs}\prod_{i\in\mathcal{C}}\Big(1 + \frac{\rho_i}{P_i}(e^{sP_i} - 1)\Big)\Big\} =: \gamma \ . \tag{57}$$

Inserting Eqn. (57) into Eqn. (54) and solving for T^{ε} yields

$$T^{\varepsilon} = T^0 - \left\lfloor \frac{\varepsilon}{\gamma} \right\rfloor \ . \tag{58}$$

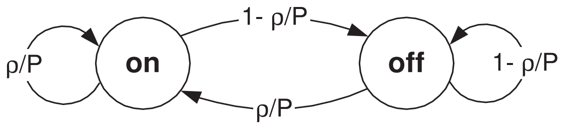

Figure 2: On-Off Transition Model.

5.2 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. As illustrated in Figure 2, 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 [7]

$$\alpha_{\mathcal{C}}(s,t) = \frac{1}{s}\sum_{i\in\mathcal{C}} \log\left(1 + \frac{\rho_i}{P_i}\left(e^{P_i s} - 1\right)\right) . \tag{59}$$

Lemma 2 gives the corresponding effective envelope as

$$\mathcal{G}^{\varepsilon}_{\mathcal{C}}(t) \;=\; \inf_{s>0}\Big\{\frac{t}{s}\sum_{i\in\mathcal{C}} \log\Big(1 + \frac{\rho_i}{P_i}\big(e^{P_i s} - 1\big)\Big) - \frac{\log\varepsilon}{s}\Big\} . \tag{60}$$

To obtain a busy period estimate at a link with a constant rate C with $\rho_{\mathcal{C}} < C < P_{\mathcal{C}}$, we use again Lemma 1. By the Chernoff bound, we have for any $t, \tau \geq 0$

$$\begin{aligned} Pr\big\{A_{\mathcal{C}}[t+\tau] - A_{\mathcal{C}}[t] > C\tau\big\} &\leq \inf_{s\geq 0} e^{s\tau(\alpha_{\mathcal{C}}(s,\tau) - C)} && (61)\\ &\leq \inf_{s>0}\Big\{e^{-Cs}\prod_{i\in\mathcal{C}}\Big(1 + \frac{\rho_i}{P_i}\big(e^{P_i s} - 1\big)\Big)\Big\}^{\tau} =: \gamma^{\tau} . && (62)\end{aligned}$$

For C with $\rho_{\mathcal{C}} < C < P_{\mathcal{C}}$ the objective function in Eqn. (62) has a unique minimum, with $\gamma < 1$. It follows that, for any $T \geq 0$,

$$\sum_{\tau=T+1}^{\infty} \sup_{t\geq 0} Pr\big\{A_{\mathcal{C}}[t+\tau] - A_{\mathcal{C}}[t] \geq C\tau\big\} \leq \frac{\gamma^{T+2}}{1-\gamma} < \infty , \tag{63}$$

verifying Eqn. (31) of Lemma 1. Setting the right hand side equal to ε and solving for T shows that

$$T^{\varepsilon} = \left\lfloor \frac{\log\left((1-\gamma)\varepsilon\right)}{\log\gamma} \right\rfloor - 1 \tag{64}$$

satisfies Eqn. (32).

5.3 Fractional Brownian Motion traffic

As pointed out in [18], the time autocorrelations of measured traffic data can sometimes be modeled by processes of the form

$$A[t] = \rho t + \beta Z_t , \tag{65}$$

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 $EZ_t = 0$, and variances $EZ_t^2 = |t|^{2H}$ for all t.

Following [18], we will refer to Eqn. (65) as the *Fractional Brownian Motion (FBM)* traffic model. Note that the sum of the arrivals from a collection $\mathcal{C}$ of independent fractional Brownian traffic sources with common Hurst parameter is again fractional Brownian traffic, where the mean traffic rate and the variance parameter β^2 are given by

$$\rho_{\mathcal{C}} = \sum_{i\in\mathcal{C}} \rho_i \ , \quad \beta_{\mathcal{C}}^2 = \sum_{i\in\mathcal{C}} \beta_i^2 \ . \tag{66}$$

We remark that the fractional Brownian traffic model is an idealization that fails to capture certain basic properties of actual traffic. Most notably, even though the average rate is positive, increments can be negative, and there is positive probability that a sample path fails to be nondecreasing, or even nonnegative. Furthermore, fractional Brownian traffic is defined for continuous time, while we consider here discrete-time arrival processes. We note that the estimates below hold for all (discrete-time) arrival processes that have nonnegative increments, and whose moment generating function is bounded by the moment generating function of fractional Brownian traffic.

The effective bandwidth for fractional Brownian traffic is given by [15]

$$\alpha_{\mathcal{C}}(s,t) = \rho_{\mathcal{C}} + \frac{1}{2}\beta_{\mathcal{C}}^2 s t^{2H-1} \ . \tag{67}$$

By Lemma 2, this results in an effective envelope of

$$\mathcal{G}_{\mathcal{C}}^{\varepsilon}(t) \quad = \quad \rho_{\mathcal{C}} t + \sqrt{-2\log\varepsilon}\, \beta_{\mathcal{C}} t^H \ . \tag{68}$$

Turning to the construction of busy period bounds, to verify the assumptions of Lemma 1, we use the Chernoff bound to obtain

$$Pr\big\{A_{\mathcal{C}}[t+\tau] - A_{\mathcal{C}}[t] \geq C\tau\big\} \leq \inf_{s\geq 0} e^{s\tau(\alpha_{\mathcal{C}}(s,\tau)-C)} = e^{-\frac{1}{2}\left(\frac{C-\rho_{\mathcal{C}}}{\beta_{\mathcal{C}}}\right)^2 \tau^{2-2H}} \ . \tag{69}$$

Since $H < 1$, the right hand side is a nonincreasing integrable function of τ. It follows that for any $T > 0$,

$$\sum_{\tau=T}^{\infty} \sup_{t\geq 0} Pr\big\{A_{\mathcal{C}}[t+\tau] - A_{\mathcal{C}}[\tau] \geq C\tau\big\} \leq \int_T^{\infty} e^{-\frac{1}{2}\left(\frac{C-\rho_{\mathcal{C}}}{\beta_{\mathcal{C}}}\right)^2 \tau^{2-2H}} \, d\tau < \infty \ , \tag{70}$$

which implies Eqn. (31). To estimate T^{ε}, we need a quantitative bound for the integral on the right hand side of Eqn. (70). Set

$$b = \frac{1}{2}\left(\frac{C-\rho_{\mathcal{C}}}{\beta_{\mathcal{C}}}\right)^2 \ , \quad \theta = 2-2H \ , \tag{71}$$

and compute, for $H \in (\frac{1}{2}, 1)$,

$$\begin{aligned}
\int_T^{\infty} e^{-b\tau^{\theta}} \, d\tau \quad &= \quad \int_{T^{\theta}}^{\infty} e^{-bx} x^{1/\theta-1} \, dx && (72)\\
&\leq \inf_{0<s<1} \Big\{ \Big(\frac{1/\theta-1}{ebs}\Big)^{1/\theta-1} \int_{T^{\theta}}^{\infty} e^{-(1-s)bx} \, dx \Big\} && (73)\\
&\leq \inf_{0<s<1} \Big\{ \Big(\frac{1/\theta-1}{ebs}\Big)^{1/\theta-1} \frac{1}{(1-s)b} e^{-b(1-s)T^{\theta}} \Big\} \ . && (74)
\end{aligned}$$

In Eqn. (72), we have changed variables to $x = \tau^\theta$. In Eqn. (73), we have used that

$$e^{-sbx} x^{1/\theta - 1} \leq \left(\frac{1/\theta - 1}{ebs}\right)^{1/\theta - 1} \tag{75}$$

for any $x \geq 0$, $s > 0$, and in Eqn. (74) we have performed the remaining integration. Setting the right hand side of Eqn. (74) equal to ε and solving for T yields the desired bound

$$T^\varepsilon = \inf_{0<s<1} \left\{ \frac{1}{(1-s)b} \Big[-\log\big((1-s)b\varepsilon\big) + (1/\theta - 1)\log\Big(\frac{1/\theta - 1}{ebs}\Big)\Big]\right\}^{1/\theta}. \tag{76}$$

6 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. With the busy period bounds from the previous section, the service curves can be used to verify probabilistic delay guarantees.

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 $q = 1, \ldots, Q$ by $A_{\mathcal{C}}$. We make similar conventions for departures and backlogs. We use $\mathcal{G}_q^{\varepsilon_g}$ to denote an effective envelope for the arrivals from class q. We consider a workconserving link with link rate C, and three scheduling disciplines: 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 deadline $t + d_q$, and the EDF scheduler always selects the packet with the smallest deadline for service. We allow that packets miss their assigned deadline. A delay index is a parameter of the scheduling algorithm, which determines the order of transmission. The delay index by itself does not provide an upper bound on delays.

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.

Lemma 3 *Consider the arrivals from Q classes to a workconserving server with capacity C. Assume the arrivals have non-negative increments. Let T^{ε_b} be a busy period bound for the aggregate $A_{\mathcal{C}}$ which satisfies*

Eqn. (33). Assume the scheduling algorithm employed at the server is either SP, EDF, or GPS. In the case of GPS, assume additionally that the functions $\mathcal{G}_p^{\varepsilon_p}$ are concave. Let $\varepsilon > 0$ and the function $\mathcal{S}_q^{\varepsilon}$ be given as follows:

1. **SP:** $$\mathcal{S}_q^{\varepsilon}(t) = \Big[Ct - \sum_{p<q} \mathcal{G}_p^{\varepsilon_g}(t)\Big]_+ , \quad \varepsilon = \varepsilon_b + (q-1)T^{\varepsilon_b}\varepsilon_g \ . \tag{77}$$
2. **EDF:** $$\mathcal{S}_q^{\varepsilon}(t) = \Big[Ct - \sum_{p\neq q} \mathcal{G}_p^{\varepsilon_g}(t - [d_p - d_q]_+)\Big]_+ , \quad \varepsilon = \varepsilon_b + (Q-1)T^{\varepsilon_b}\varepsilon_g \ . \tag{78}$$
3. **GPS:** $$\mathcal{S}_q^{\varepsilon}(t) = \lambda_q\Big(Ct + \sum_{p\neq q} \Big[\lambda_p Ct - \mathcal{G}_p^{\varepsilon_g}(t)\Big]_+\Big) , \quad \varepsilon = \varepsilon_b + (Q-1)T^{\varepsilon_b}\varepsilon_g \ , \tag{79}$$
   where $\lambda_p = \phi_p / \sum \phi_r$ is the guaranteed share of class p.

Then $\mathcal{S}_q^{\varepsilon}$ is an effective service curve for class q, satisfying

$$Pr\Big\{\exists \tau \le T^{\varepsilon_b} : \ D_q[t] \ge A[t-\tau] + \mathcal{S}_q^{\varepsilon}(\tau)\Big\} \ge 1 - \varepsilon \ . \tag{80}$$

By setting all violation probabilities $\varepsilon_b, \varepsilon_g = 0$ in Lemma 1, we can recover a deterministic (worst-case) statement on the lower bound of the service seen by a service class. Such worst-case bounds, however, are generally too pessimistic to be of practical relevance.

Proof.

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

$$\underline{t}_{\le q} = \max\big\{x \le t : \ B_{\le q}(x) = 0\big\} \tag{81}$$

be the beginning of the busy period containing t from the perspective of class q. If $B_q(t) > 0$, where B_q is the class-*q* backlog, then we have by the properties of the SP scheduler that

$$\begin{aligned} D_q[t] &= D_q[\underline{t}_{\le q}] + \big(D_{\le q}[t] - D_{\le q}[\underline{t}_{\le q}]\big) - \big(D_{<q}[t] - D_{<q}[\underline{t}_{\le q}]\big) &(82)\\ &\ge A_q[\underline{t}_{\le q}] + \Big[C(t - \underline{t}_{\le q}) - \big(A_{<q}[t] - A_{<q}[\underline{t}_{\le q}]\big)\Big]_+ \ . &(83)\end{aligned}$$

In Eqn. (83), we have used that $D_p(\underline{t}_{\le q}) = A_p(\underline{t}_{\le q})$ for all $p \le q$, that $D[t] - D[\underline{t}_{\le q}] \ge C(t - \underline{t}_{\le q})$ by the properties of the workconserving server, and that $D_p[t] \le A_p[t]$ for all *p*. It follows that

$$\begin{aligned} &Pr\Big\{\exists \tau \le T^{\varepsilon_b} : \ D_q[t] \ge A_q[t-\tau] + \Big[C(\tau) - \sum_{p<q}\mathcal{G}_p^{\varepsilon_g}(\tau)\Big]_+\Big\} \\ &\ge Pr\Big\{t - \underline{t}_{\le q} \le T^{\varepsilon_b} \text{ and } D_q[t] \ge A_q[\underline{t}_{\le q}] + \Big[C(t-\underline{t}_{\le q}) - \sum_{p<q}\mathcal{G}_p^{\varepsilon_g}(t - \underline{t}_{\le q})\Big]_+\Big\} &(84)\\ &\ge Pr\Big\{t - \underline{t}_{\le q} \le T^{\varepsilon_g} \text{ and } A_{<q}[t] - A_{<q}[\underline{t}_{\le q}] \le \sum_{p<q}\mathcal{G}_p^{\varepsilon_g}(t - \underline{t}_{\le q})\Big\} &(85)\\ &\ge Pr\Big\{t - \underline{t} \le T^{\varepsilon_g} \text{ and } \forall p < q, \ \forall \tau \le T^{\varepsilon_b} : \ A_p[t] - A_p[t-\tau] \le \mathcal{G}_p^{\varepsilon_g}(\tau)\Big\} &(86)\\ &\ge 1 - \Big(\varepsilon_b + (q-1)T^{\varepsilon_b}\varepsilon_g\Big) , &(87)\end{aligned}$$

where $\underline{t}$ is the beginning of the busy period of the server. The above proves the claim for SP. In Eqn. (84), we have set $\tau = t - \underline{t}_{\le q}$, and in Eqn. (85), we have used Eqn. (83). In Eqn. (86), we have restricted the event and used that $\underline{t} \le \underline{t}_{\le q}$, and in the last line, we have used the definitions of T^{ε_b} and $\mathcal{G}_p^{\varepsilon_g}$.

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 server is workconserving, this implies

$$\begin{aligned} D_q[t] &= D_q[\underline{t}] + \big(D_{\mathcal{C}}[t] - D_{\mathcal{C}}[\underline{t}]\big) - \sum_{p \neq q} \big(D_p[t] - D_p[\underline{t}]\big) && (88) \\ &\geq A_q[\underline{t}] + \Big[C(t - \underline{t}) - \sum_{p \neq q} \big(A_p[t - (d_p - d_q)_+] - A_p[\underline{t}]\big)\Big]_+ . && (89) \end{aligned}$$

We argue as in Eqs. (84)-(87) that

$$\begin{aligned} & Pr\Big\{\exists \tau \leq T^{\varepsilon_b} :\ D_q[t] \geq A_q[t-\tau] + \Big[C(\tau) - \sum_{p \neq q} \mathcal{G}_p^{\varepsilon_g}(\tau - (d_p - d_q)_+)\Big]_+ \text{ and } B_q(t-\tau) = 0\Big\} \\ & \geq Pr\Big\{t - \underline{t} \leq T^{\varepsilon_b} \text{ and } \forall p \neq q,\ \forall \tau \leq T^{\varepsilon_b} :\ A_p[t] - A_p[t-\tau] \leq \mathcal{G}_p^{\varepsilon_g}(\tau)\Big\} && (90) \\ & \geq 1 - \Big(\varepsilon_b + (Q-1)T^{\varepsilon_b}\varepsilon_g\Big) . && (91) \end{aligned}$$

GPS scheduling: For $t \geq 0$, let

$$\underline{t}_p = \max\big\{x \leq t :\ B_p(x) = 0\big\} \qquad (92)$$

be the beginning of the busy period of t with respect to class p. Clearly,

$$B_p(t) = A_p[t] - D_p[t] \leq A_p[t] - A_p[\underline{t}_p] - \lambda_p C(t - \underline{t}_p) \qquad (93)$$

by the properties of the GPS scheduler. For $t \geq 0$ and $p \neq q$, let

$$\underline{t}_{qp} = \max\big\{x \leq \underline{t}_q :\ B_p(x) = 0\big\} , \qquad (94)$$

then Eqn. (93) with t replaced by $\underline{t}_q$ implies that

$$\begin{aligned} D_p[t] - D_p[\underline{t}_q] &\leq A_p[t] - A_p[\underline{t}_q] + B_p(\underline{t}_q) && (95) \\ &\leq A_p[t] - A_p[\underline{t}_{qp}] - \lambda_p C(\underline{t}_q - \underline{t}_{qp}) . && (96) \end{aligned}$$

It follows that

$$\begin{aligned} D_q[t] - D_q[\underline{t}_q] &\geq \lambda_q\Big(C(t - \underline{t}_q) + \sum_{p \neq q} \big[\lambda_p C(t - \underline{t}_q) - D_p[t] + D_p[\underline{t}_q]\big]_+\Big) && (97) \\ &\geq \lambda_q\Big(C(t - \underline{t}_q) + \sum_{p \neq q} \big[\lambda_p C(t - \underline{t}_{qp}) - A_p[t] + A_p[\underline{t}_{qp}]\big]_+\Big) . && (98) \end{aligned}$$

Fix $t \geq 0$, and assume for the moment that

$$t - \underline{t} \leq T^{\varepsilon_b} \text{ and } \forall p \neq q, \forall \tau \leq T^{\varepsilon_b} :\ A_p[t] - A_p[t-\tau] \leq \mathcal{G}_p^{\varepsilon_g}(\tau) . \qquad (99)$$

Since $\underline{t} \leq \underline{t}_{qp} \leq \underline{t}_q$, it follows with by Eqn. (98) that

$$\begin{aligned} D_q[t] &\geq D_q[\underline{t}_q] + \lambda_q\Big(C(t - \underline{t}_q) + \sum_{p \neq q} \big[\lambda_p C(t - \underline{t}_{qp}) - A_p[t] + A_p[\underline{t}_{qp}]\big]_+\Big) && (100) \\ &\geq A_q[\underline{t}_q] + \lambda_q\Big(C(t - \underline{t}_q) + \sum_{p \neq q} \big[\lambda_p C(t - \underline{t}_{qp}) - \mathcal{G}_p^{\varepsilon_g}(t - \underline{t}_{pq})\big]_+\Big) . && (101) \end{aligned}$$

	Regulated Traffic			On-Off Traffic		FBM Traffic		
Type	P (Mbps)	ρ (Mbps)	σ σ (bits)	P (Mbps)	ρ (Mbps)	ρ (Mbps)	β (Mbs)	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 1: Source Traffic Parameters.

Since $\mathcal{G}_p^{\varepsilon_g}$ is concave, the function $\left[\lambda Ct - \mathcal{G}_p^{\varepsilon_g}(t)\right]_+$ is nondecreasing in t. Replacing $t - \underline{t}_{pq}$ with the smaller value $t - \underline{t}_q$ in Eqn. (101) and using the definition of $\mathcal{S}_q^{\varepsilon}$ yields

$$D_q[t] \geq A_q[\underline{t}_q] + \mathcal{S}_q^{\varepsilon}(t - \underline{t}_q) \ . \tag{102}$$

Finally, we estimate

$$\begin{aligned} Pr\Big\{\exists \tau \leq T^{\varepsilon_b} :\ D_q[t] \geq A_q[t-\tau] + \mathcal{S}_q^{\varepsilon}(\tau)\Big\} &\geq Pr\Big\{t - \underline{t}_q \leq T^{\varepsilon_b} \text{ and Eqn. (102) holds }\Big\} && (103) \\ &\geq Pr\Big\{\text{Eqn. (99) holds}\Big\} && (104) \\ &\geq 1 - \Big(\varepsilon_b + (Q-1)T^{\varepsilon_b}\varepsilon_g\Big) \ . && (105) \end{aligned}$$

This completes the proof. □

We note that the formulas in Eqs. (77)–(79) do not fully characterize the service available to class q for the three schedulers. Rather, they represent lower bounds on the leftover capacity that is left by other classes. Among the three scheduling algorithms, Eqn. (77) describes the performance of an SP scheduler rather closely. Eqn. (79) for the GPS scheduler is not the best possible description, but improves on the minimal guaranteed rate $\lambda_q C$. On the other hand, Eqn. (78) does not entirely reflect the properties of the EDF scheduler. For example, in the limit where $d_p \approx d_q$ for all classes $p \neq q$, Eqn. (78) approaches the service guarantees of an SP scheduler for the lowest priority class, while the actual EDF scheduler approaches FIFO.

7 Numerical Examples

In this section, we present numerical examples for a single network node, 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 1. The unit of time is 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 [28]. The parameters of the other traffic sources are selected to match the average rate ($\rho = 0.15\ Mbps$) and to have the same variance at $t = 1\ ms$. The Hurst parameter is set to $H = 0.78$ as in [18]. In the examples, we consider violation probabilities $\varepsilon = 10^{-3}, 10^{-6}$, and 10^{-9}.

7.1 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 $\mathcal{G}_N^{\varepsilon}(t)/N$, where $\mathcal{G}_N^{\varepsilon}(t)$

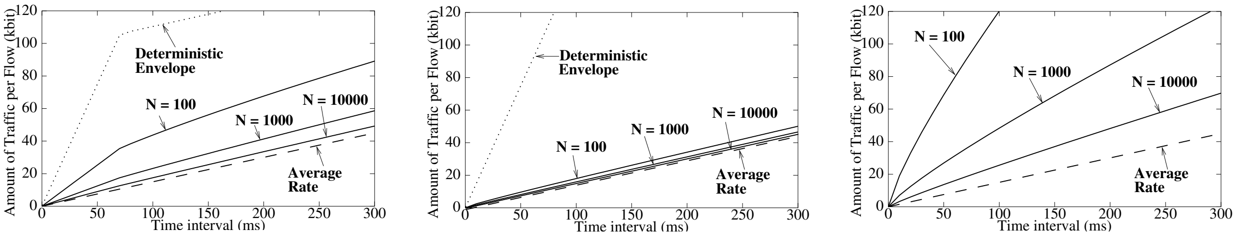

(a) (P, σ, ρ) Regulated. (b) On-Off. (c) FBM.

Figure 3: Effective envelopes for Type-1 flows ($\varepsilon = 10^{-9}$).

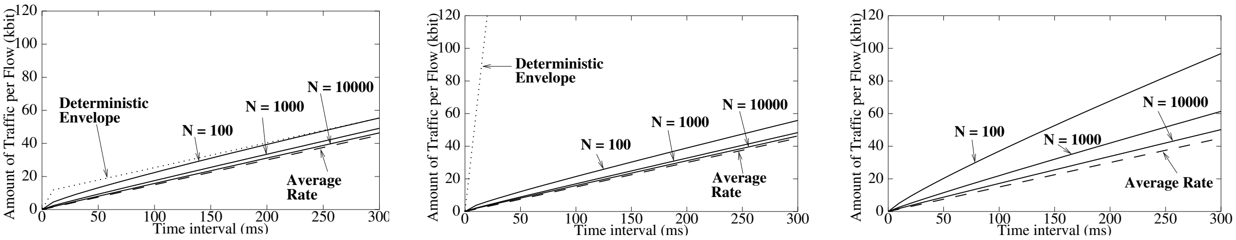

(a) (P, σ, ρ) Regulated. (b) On-Off. (c) FBM.

Figure 4: Effective envelopes for Type-2 flows ($\varepsilon = 10^{-9}$).

is the effective envelope for N homogeneous flows. Figures 3 and 4 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 and On-Off traffic, respectively, we include the deterministic envelopes $\min(Pt\,, \sigma + \rho t)$ and Pt. Note that a deterministic envelope does not exist for FBM traffic.

We make the following observations. First, the effective envelopes are able to capture a significant amount of statistical multiplexing gain for each of the considered traffic types. The multiplexing increases sharply with the number of flows N. Second, as expected, FBM traffic exhibits less multiplexing gain than the other source models.

7.2 Example 2: Probabilistic Busy Period Bounds.

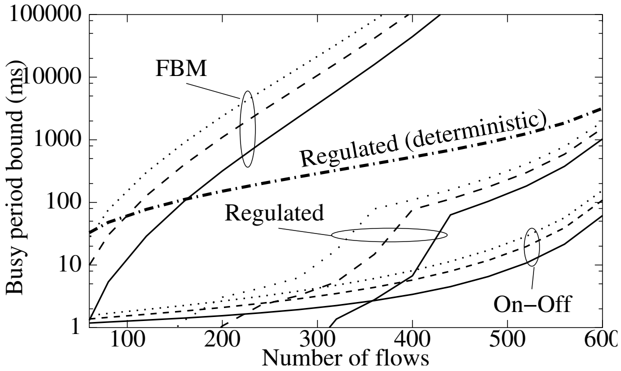

Figure 5: Example 2: Probabilistic Busy Period Bounds for $\varepsilon = 10^{-3}$ (solid line), $\varepsilon = 10^{-6}$ (dashed line), and $\varepsilon = 10^{-9}$ (dotted line). The thick dotted-dashed line is a deterministic busy period bound for regulated traffic.

Next we investigate the probabilistic bound for the busy periods. We assume a link with capacity of $C = 100\ Mbps$. We observe how the busy period grows as the number of flows increases. In this example, we assume a traffic mix of an equal number of Type-1 and Type-2 flows. We calculate the probabilistic busy period bounds for violation probabilities $\varepsilon = 10^{-3}, 10^{-6}, 10^{-9}$ using the bounds from Section 4. Figure 5 depicts the bounds for the three different traffic models, where the number of flows is varied from 60 to 600, corresponding to a utilization of 9% to 90%. As a reference point, we also plot the exact value for the worst-case busy period of the regulated traffic (plotted as thick solid line). While regulated traffic permits to determine the worst-case busy period, such deterministic bounds are not available for On-Off and FBM traffic. Note that the busy period bounds for FBM traffic are significantly larger than those for Regulated or On-Off traffic.

7.3 Example 3: Number of Admissible Flows

In this example, we consider three scheduling disciplines (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. This maximum number can be used for admission control of Type-1 flows. Note that such an admission control decision is greedy, in the sense that it entirely ignores delay requirements of other flow types. For example, the delay requirements of Type-2 flows are ignored in this example.

The admission control algorithms takes into consideration the scheduling algorithm and its parameters, as well as the the source traffic parameters of all flow types. 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 precedence, 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 1. 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 6 depicts the number of Type-1 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. This can be explained by the higher burstiness permitted by FBM traffic. We also observe in the figure, that the selection of the scheduling algorithm has only a limited impact. Given a traffic model, the number of admitted Type-1 flows is similar for all scheduling algorithm, with one notable exception. For GPS, the minimum number of Type-1 flows admitted is independent of the number of Type-2 flows. This is the result of the rate guarantee provided by GPS.

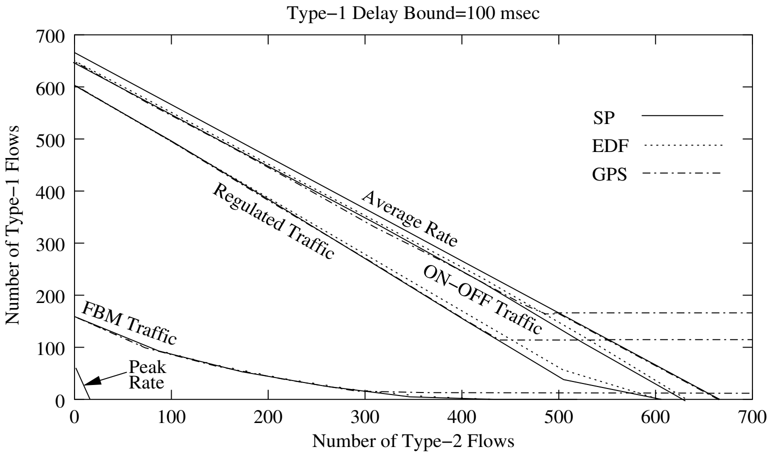

Figure 6: Admissible Regions (C = 100 Mbps) for different schedulers and traffic models with $\varepsilon = 10^{-6}$, $d_1 = 100\ ms$, $\phi_1 = 0.25$, $\phi_2 = 0.75$.

8 Conclusions

We have presented a statistical network calculus for statistical QoS provisioning where both arrivals and service are described in terms of probabilistic bounds. We have shown that it is feasible to integrate the concept of effective bandwidth into the envelope-based approach of the statistical network calculus. We have derived backlog and delay bounds for several traffic models (regulated, On-Off, FBM), and scheduling algorithms (SP, EDF, GPS), and presented bounds on the queueing behavior in terms of the min-plus algebra.

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] J.-Y. Le Boudec. Application of network calculus to guaranteed service networks. *IEEE/ACM Transactions on Information Theory*, 44(3):1087–1096, October 1998.

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

[4] C. Chang. Stability, queue length and delay part II: Stochastic queueing networks. In *Proceedings of IEEE Conference on Decision and Control*, pages 1005–1010, Tucson, AZ, December 1992.

[5] 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.

[6] C. S. Chang. On deterministic traffic regulation and service guarantees: a systematic approach by filtering. *IEEE/ACM Transactions on Information Theory*, 44(3):1097–1110, May 1998.

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

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

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

[10] V. Firoiu, J.-Y. Le Boudec, D. Towsley, and Z.-L. Zhang. Theories and models for internet quality of service. *Proceedings of the IEEE*, 2002. To appear.

[11] R. J. Gibbens and P. J. Hunt. Effective bandwidths for the multitype uas channel. *Queueing Systems Theory and Applications*, 16(9):17–27, October 1991.

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

[13] J. Hui. Resource allocation for broadband networks. *IEEE Journal on Selected Areas in Communications*, 6(6):1598–1608, August 1988.

[14] F. Kelly. Effective bandwidths at multi-class queues. *Queueing Systems, Theory and Applications*, 9(1-2):5–16, 1991.

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

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

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

[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] V. Sivaraman, F. Chiussi, and M. Gerla. End-to-end statistical delay service under GPS and EDF scheduling: A comparison study. In *Proceedings of IEEE INFOCOM 2001*, Anchorage, Alaska, April 2001.

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

[23] 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.

[24] M. Vojnovic and J.-Y. Le Boudec. Bounds for independent regulated inputs multiplexed in a service curve network element. *IEEE Transactions on Communications*, October 2002. To appear.

[25] W. Whitt. Tail probabilities with statistical multiplexing and effective bandwidths in multi-class queues. *Telecommunication Systems*, 2:71–107, 1993.

[26] 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.

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

[28] R. R. ZZBoorstyn, 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.

[29] A. ZZBurchard, 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.