IEEE/ACM Transactions on Networking

A Network Calculus with Effective Bandwidth

Journal:	*IEEE/ACM Transactions on Networking*
Manuscript ID:	TNET-00403-2003.R1
Manuscript Type:	Original Article
Date Submitted by the Author:	n/a
Complete List of Authors:	Li, Chengzhi; University of Texas, Arlington, Computer Science and Engineering Burchard, Almut; University of Toronto, Mathematics Liebeherr, Jorg; University of Toronto, Department of Electrical and Computer Engineering
Keywords:	network calculus, statistical multiplexing, Quality-of-Service

A Network Calculus with Effective Bandwidth *

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

[†] Department of Computer Science, University of Texas, Arlington
[††] Department of Mathematics, University of Toronto
[*] Department of Electrical and Computer Engineering, University of Toronto

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 formulation 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.

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. Probably the most influential framework for service provisioning is the *effective bandwidth* (see [30, 33] 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 [30].

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 [19] the deterministic network calculus has evolved to an elegant framework for worst-case analysis, which can be used to derive upper bounds for delay and backlog for a wide variety of link scheduling algorithms (e.g., [19, 40, 37]). A strength of the deterministic network calculus is that it can be used to determine delay and backlog over multiple network nodes. Since the worst-case view of the deterministic network calculus does not reap the benefits of statistical multiplexing and generally results

*This work is supported in part by the National Science Foundation through grants DMS-9971493, ANI-0085955, CNS-0435061, and by an Alfred P. Sloan research fellowship.

in an overestimation of the actual resource requirements of traffic, researchers have seeked to extend the network calculus to a probabilistic setting, e.g., [2, 7, 11, 13, 35, 41, 45, 49, 51, 53]. Probabilistic extensions of the network calculus are 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 [13] (see Subsection 3.1). 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.

In the statistical network calculus, arrivals to a network node are described in terms of probabilistic upper bounds (*effective envelopes* [7]) and the service at a node is described in terms of probabilistic lower bounds (*effective service curves* [11]). The effective service curves in this paper can express the service available to one flow in terms of the capacity unused by other flows. We do this for a wide range of scheduling algorithms. By relating the concepts of effective envelopes and effective bandwidth, we obtain explicit bounds on delay and backlog for all traffic source characterizations for which an effective bandwidth (in the sense of [13, 30]) has been determined. As a result, much of the literature on the effective bandwidth theory can now be applied in a network calculus context. This, enables the analysis of network models, such as FBM traffic at nodes with a deadline-based scheduling algorithm, which have not been analyzed before.

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 sometimes lead tighter bounds, such a direct analysis generally only applies to a single node and is not easily extended to a multi-node setting. Recent developments showed that, in some cases, a network calculus analysis can sometimes reproduce bounds obtained with a direct statistical analysis [18].

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 discuss some of the difficulties. We argue that the availability of a maximum relevant time scale, that is, a bound on the maximum time period at which events are correlated, makes a statistical calculus analysis tractable. 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, and 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 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 2, we present the statistical network calculus that is used to accommodate effective bandwidth expressions. In Section 3, 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 4, 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 5, we apply the network calculus in a set of examples, and

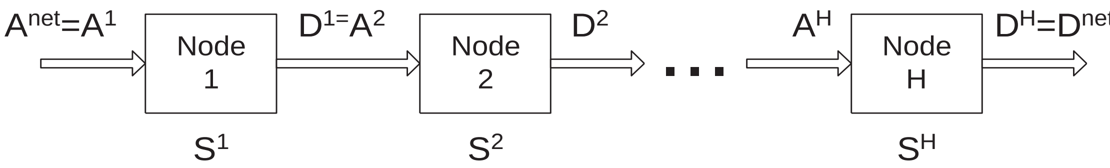

Figure 1: Traffic of a flow through a set of H nodes. The arrivals and departures from the network are given by random processes A^{net} and D^{net}. The arrivals and departures from the h-th node are described by A^h and D^h, with $A^1 = A^{net}$, $A^h = D_{h-1}$ for $h = 2, \ldots, H$, and $D^{net}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

compare the multiplexing gain achievable with the traffic models and scheduling algorithms used in this paper. We present brief conclusions in Section 6.

2 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. The main contribution of this section are the probabilistic bounds for output burstiness, backlog, delay and network service derived in Subsection 2.5 and time scale estimates for general service and arrival scenarios.

2.1 Notation and Definitions

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 are denoted by nonnegative, nondecreasing functions $A(t)$ and $D(t)$, respectively, 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 $W(t) = \inf\{d \geq 0 \,|\, A(t-d) \leq D(t)\}$. 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 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. 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 [Chang00] Thus a stochastic bound on $B(t)$ that does not depend on t provides a bound on the steady-state distribution of $B(t)$. The corresponding statements hold for the delay distribution $W(t)$, and the distribution of $D(t) - D(t-\tau)$ for any given value of τ.

The min-plus algebra formulation of the network calculus [1, 8, 15], defines, for given functions f and

g, the convolution operator $*$ and deconvolution operator $\oslash$ by

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

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

2.2 Overview of Deterministic Network Calculus

In the deterministic network calculus in [1, 8, 15], service guarantees to a flow at a node are expressed in terms of *service curves*. A (minimum) service curve for a flow is a 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, 8, 15], 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 each node 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

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

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

End-to-end delay bounds obtained with the network calculus are generally lower than the sum of the delay bounds at each node. For example, when the service curve at each node is given as a constant rate function, $S^h(\tau) = C\tau$ for all $h = 1, 2, \ldots, H$, we obtain $S^{net} = S^1 * S^2 * \ldots * S^H = C\tau$. Here, the end-to-end backlog and delay bounds are identical to the bounds at the first node.

At this time, the deterministic calculus has been extensively explored. Its results have led to the development of new scheduling algorithms [20, 40] and have been used to specify new network services [6, 10]. We refer to [9] for a comprehensive discussion of available 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.

2.3 Overview of Statistical Network Calculus

The statistical network calculus 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 make the following assumptions on arrivals:

1. *Stationary Bounds*: For any $\tau > 0$, the arrivals A_i^{net} from any flow i to the network satisfy

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

2. *Independence*: 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 that extend to the steady-state. Assuming independence of traffic sources at the network entrance allows us to exploit statistical multiplexing gain.

The literature contains a number of different approaches to devise a statistical network calculus. One group of studies investigates network traffic that, in addition to the assumptions above, satisfies certain a priori assumptions on the arrival functions, such as exponentially bounded burstiness [51], linear envelope processes [13], stochastically bounded burstiness [45], general burstiness characterization [3], or stochastic domination by a given random variable [35]. Other studies assume that arrivals of individual flows at the network ingress are regulated by (deterministic) arrival envelopes A_i^*. Then, by exploiting the independence assumption of flows, they use either the Central Limit Theorem [7, 31, 32], or large deviations tools such as the Chernoff Bound [7, 24] and the Hoeffding Bound [49, 48].

With such arrival assumptions, probabilistic backlog and delay bounds for a single node have been derived for FIFO schedulers with a fixed service rate. Some studies [13, 45, 51] also derive probabilistic bounds for the output of a node, which can then be iterated to yield end-to-end bounds. However, end-to-end bounds obtained in this fashion degrade rapidly with the number of nodes. Other studies consider more complex scheduling algorithms [7, 41, 48] for a single node. There are a few results available for end-to-end statistical guarantees, generally for special arrival or service models [42, 43, 44].

A different set of studies attempts to express a statistical network calculus using the min-plus algebra formulation with convolution and deconvolution operators [2, 11]. 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). In [11] it was shown that a network service curve in the statistical network calculus can be constructed if the service curve satisfies additional properties. In [2], a probabilistic network service curve is derived under the assumption that each node drops traffic that locally violates a given delay guarantee. The results in [11] and [2] do not make any assumptions on arrivals and hold for all sample paths of the arrivals. The current state of the statistical network calculus has shown that expressions for backlog, delay, and output bounds at a single node carry over from the deterministic network calculus to a statistical framework. However, a network service curve requires to make significant additional assumptions. At present, finding suitable assumptions that permit a formulation of a network service curve as in Eqn. (2), without restricting the applicability of the framework, is an open research problem.

We next describe the probabilistic framework used in this paper. We follow the framework for a statistical calculus presented in [7] and [11]. For traffic arrivals, we use a probabilistic measure called *effective envelopes* [7]. An effective envelope for an arrival process A is defined as a non-negative function $\mathcal{G}$ such that for all t and τ

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

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 3 below). To characterize the available service to a flow or a collection of flows we use *effective service curves* [11] which can be seen as a probabilistic measure of the available service. Given an arrival process A, an effective service curve is a non-negative 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.

2.4 What Makes Statistical Network Calculus Hard?

To illustrate that the statistical network calculus is not a straightforward extension of the deterministic network calculus, we want to mention two technical difficulties encountered when extending the calculus to a probabilistic setting. The first appears when estimating the tail distribution for the backlog or the envelope of the output traffic at a node. In the case of the backlog, the expression takes the form

$$Pr\Big\{B(t) > y\Big\} = Pr\Big\{\sup_{\tau\geq 0}\{A(t-\tau,t) - S(\tau)\} > y\Big\}\,, \tag{5}$$

where we have used $A(t-\tau,t)$ to denote $A(t) - A(t-\tau)$. The difficulty relates to the evaluation of the right hand side of the equation. Note that in Eqn. (5), the arrivals are random but service is deterministic; a probabilistic view of service causes no additional complications here. In [7] and [16], the above expression is approximated by

$$Pr\Big\{\sup_{\tau\geq 0}\{A(t-\tau,t) - S(\tau)\} > y\Big\} \approx \sup_{\tau\geq 0} Pr\Big\{A(t-\tau,t) - S(\tau) > y\Big\}\,, \tag{6}$$

by using an argument from extreme-value theory [12]. The approximation can be justified in some situations, for instance when traffic is described by a Gaussian process. However, in general the right hand side of Eqn. (5) is only a lower bound for the left hand side. In [3], the right hand side of Eqn. (5) is controlled by assuming the existence of a probabilistic bound for the entire arrival sample path. Another way to deal with Eqn. (5) is to use Boole's inequality, which yields

$$Pr\Big\{\sup_{\tau\geq 0}\{A(t-\tau,t) - S(\tau) > y\}\Big\} \leq \sum_{\tau=0}^{\infty} Pr\Big\{A(t-\tau,t) - S(\tau) > y\Big\}\,. \tag{7}$$

where the sum is replaced by an integral in a continuous time domain. This can yield a useful bound if one has available a tail estimate on the distribution of $A(t-\tau,t) - S(\tau)$, or if there exists a maximum relevant time scale, say T_{max}, such that $Pr\{A(t-\tau,t) - S(\tau) > y\} = 0$ for $\tau > T_{max}$ so that the sum contains only finitely many terms.

The second difficulty arises in the derivation of a probabilistic version of a network service curve. This issue was pointed out in [11] 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

$$D^2(t) \geq A^2 * \mathcal{S}^{2,\varepsilon}(t) = \inf_{\tau\in[0,t]}\big\{A^2(t-\tau) + \mathcal{S}^{2,\varepsilon}(\tau)\big\}\,. \tag{8}$$

Suppose that the infimum in Eqn. (8) 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

$$D^1(x) \geq A^1 * \mathcal{S}^{1,\varepsilon}(x) \; . \tag{9}$$

Since $\hat{\tau}$ is a random variable, we cannot simply evaluate Eqn. (9) for $x = t - \hat{\tau}$ and use the resulting bound in Eqn. (8). Furthermore, there is, a priori, no time-independent bound on the distribution of $\hat{\tau}$. This is different in the deterministic calculus, where deterministic service curves make guarantees that hold for all values of x. This problem can also be resolved if a time scale bound T_{max} is available, which limits the range over which the infimum is taken as follows:

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

2.5 Network Calculus for Probabilistically Bounded Arrivals and Service

As we pointed out in the previous subsection, the difficulties of the statistical network calculus can be dealt with by assuming appropriate time scale limits. A key assumption made in this paper is that the node offers a service curve $\mathcal{S}^{\varepsilon_s}$ which satisfies the additional requirement that there exists a time scale T such that for all $t \geq 0$,

$$Pr\Big\{ D(t) \geq \inf_{\tau \leq T} \left\{ A(t - \tau) + \mathcal{S}^{\varepsilon}(\tau) \right\} \Big\} \geq 1 - \varepsilon \; . \tag{10}$$

Thus T bounds the range of the convolution in Eqn. (4). This assumption solves both problems discussed in the previous subsection. In general, the value of T depends on the arrival process as well as on the service curve. 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 Subsection 2.6.

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. (10).

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. (10) with some $T < \infty$. Define ε to be*

$$\varepsilon = \varepsilon_s + T\varepsilon_g \; . \tag{11}$$

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 from Subsection 2.3 as presented in [1, 8, 14]. Similarly, when only $\varepsilon_g = 0$, the time scale bound T disappears from Eqn. (11) and one can take $T \to \infty$. Thus, the statistical calculus from [11], 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. First, we prove that $\mathcal{G}^{\varepsilon_g} \oslash \mathcal{S}^{\varepsilon_s}$ is an effective envelope for the output traffic. Fix $t, \tau \geq 0$.

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

In Eqn. (12), 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 $x \leq T$. In Eqn. (13), we replaced $D(t+\tau)$ by $A(t+\tau)$. Further, by adding the condition that $D(t) \geq A(t-x) + \mathcal{S}^{\varepsilon_s}(x)$ 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 Eqn. (14) we further restricted the event, by demanding that the first condition in Eqn. (13) holds for all values of x. To obtain Eqn. (15), we applied the assumption in Eqn. (10), and used the definition of $\mathcal{G}^{\varepsilon_g}$. We added the violation probabilities of the two events using Boole's inequality. The factor T 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

$$\begin{aligned} Pr\Big\{B(t) \leq \mathcal{G}^{\varepsilon_g} \oslash \mathcal{S}^{\varepsilon_s}(0)\Big\} &= Pr\Big\{A(t) \leq D(t) + \mathcal{G}^{\varepsilon_g} \oslash \mathcal{S}^{\varepsilon_s}(0)\Big\} && (16) \\ &\geq Pr\left\{\exists x \leq T : \left(\begin{array}{l} A(t) \leq A(t-x) + \mathcal{S}^{\varepsilon_s}(x) + \mathcal{G}^{\varepsilon_g} \oslash \mathcal{S}^{\varepsilon_s}(0) \\ \text{and } D(t) \geq A(t-x) + \mathcal{S}^{\varepsilon_s}(x) \end{array} \right)\right\} && (17) \\ &\geq Pr\left\{ \begin{array}{l} \forall x_1 \leq T: \ A(t) - A(t-x_1) \leq \mathcal{G}^{\varepsilon_g}(x_1) \\ \text{and } \exists x_2 \leq T: \ D(t) \geq A(t-x_2) + \mathcal{S}^{\varepsilon_s}(x_2) \end{array} \right\} && (18) \\ &\geq 1 - \Big(\varepsilon_s + T\varepsilon_g\Big) . && (19) \end{aligned}$$

In Eqn. (16), we have used the definition of the backlog $B(t)$. The arguments made in Eqs. (17)–(19) are analogous to those used in Eqs. (13)–(15).

Finally, we prove the delay bound. If d satisfies $\sup_{\tau \leq T}\{\mathcal{G}^{\varepsilon_g}(\tau - d) - \mathcal{S}^{\varepsilon_s}(\tau)\} \leq 0$, then

$$\begin{aligned} Pr\Big\{W(t) \leq d\Big\} &= Pr\Big\{A(t-d) \leq D(t)\Big\} && (20) \\ &\geq Pr\left\{\exists x \leq T : \left(\begin{array}{c} A(t-d) \leq A(t-x) + \mathcal{S}^{\varepsilon_s}(x) \\ \text{and } D(t) \geq A(t-x) + \mathcal{S}^{\varepsilon_s}(x) \end{array} \right)\right\} && (21) \\ &\geq Pr\left\{ \begin{array}{l} \forall x_1 \leq T: \ A(t-d) - A(t-x_1) \leq \mathcal{G}^{\varepsilon_g}([x_1 - d]_+) \\ \text{and } \exists x_2 \leq T: \ D(t) \geq A(t-x_2) + \mathcal{S}^{\varepsilon_s}(x_2) \end{array} \right\} && (22) \\ &\geq 1 - \Big(\varepsilon_s + T\varepsilon_g\Big) . && (23) \end{aligned}$$

In Eqn. (20), we have used the definition of the delay $W(t)$, and in Eqn. (21), we have used the assumption on d. The remaining steps apply the same arguments as the proofs of the output bound and the backlog bound. □

Next we derive an expression for a probabilistic version of a network service curve, which expresses the service given by the network as a whole as a convolution of the service at each node. Consider the path of a flow through a network, as illustrated in Figure 1. 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. (10), we assume that each node satisfies

$$Pr\Big\{D^h(t) \geq \inf_{\tau \leq T^h} \{A^h(t-\tau) + \mathcal{S}^{h,\varepsilon_s}(\tau)\}\Big\} \geq 1 - \varepsilon_s \tag{24}$$

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. (24). Then an effective network service curve $\mathcal{S}^{net,\varepsilon}$ for the flow is given by*

$$\mathcal{S}^{net,\varepsilon} = \mathcal{S}^{1,\varepsilon_s} * \mathcal{S}^{2,\varepsilon_s} * \ldots * \mathcal{S}^{H,\varepsilon_s} \,, \tag{25}$$

with violation probability bounded above by

$$\varepsilon = \varepsilon_s \sum_{h=1}^{H} \Big(1 + (h-1)T^h\Big) \,. \tag{26}$$

The convolution expression in Eqn. (25) 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. (26) 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

$$\begin{cases} \forall \tau \leq \sum_{k=h+1}^{H} T^k : \; D^h(t-\tau) \geq \inf_{x_h \leq T^h} \{A^h(t-\tau-x_h) + \mathcal{S}^{h,\varepsilon_s}(x_h)\} \,, & \text{if } h < H, \\ D^H(t) \geq \inf_{x_H \leq T^H} \{A^H(t-x_H) + \mathcal{S}^{H,\varepsilon_s}(x_H)\} \,, & \text{if } h = H \,. \end{cases} \tag{27}$$

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. (27) into the second line of Eqn. (27) that

$$D^H(t) \geq \inf_{x_h \leq T^h,\; h=1,\ldots,H} \Big\{A^h(t - (x_h + \ldots + x_H)) + \sum_{k=h}^{H} \mathcal{S}^{k,\varepsilon_s}(x_k)\Big\} \,. \tag{28}$$

Setting $h = 1$ in Eqn. (28), and using the definitions of A^{net}, D^{net}, and $\mathcal{S}^{net,\varepsilon}$, we obtain

$$D^{net}(t) \geq \inf_{x_h \leq T^h,\; h=1,\ldots H} \Big\{A^{net}(t - (x_1 + \cdots + x_H)) + \mathcal{S}^{net,\varepsilon}(x_1 + \cdots + x_H)\Big\} \,. \tag{29}$$

Thus, we have shown that Eqn. (27) implies

$$D^{net}(t) \geq A^{net} * \mathcal{S}^{net,\varepsilon}(t). \tag{30}$$

We conclude proof of the theorem by

$$Pr\Big\{D^{net}(t) \geq A^{net} * \mathcal{S}^{net,\varepsilon}(t)\Big\} \geq Pr\Big\{ \text{Eqn. (27) holds} \Big\} \tag{31}$$

$$\geq 1 - \varepsilon_s \cdot \sum_{h=1}^{H} \Big(1 + \sum_{k=h+1}^{H} T^k\Big) . \tag{32}$$

In Eqn. (31) we have used that Eqn. (27) implies Eqn. (30). In Eqn. (32), we have applied Eqn. (24) and added the violation probabilities of Eqn. (27) over all possible values of $h = 1, \ldots, H$. Exchanging the order of summation completes the proof. □

2.6 Busy Period Analysis

We now 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

$$T = \sup\{\tau \geq 0 \mid A^*(\tau) > S(\tau)\} < \infty . \tag{33}$$

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. (10) with $\varepsilon = 0$. The above 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 later 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. Note that for any workconserving scheduler, 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

$$\underline{t} = \max\{\tau \leq t : \ B_{\mathcal{C}}(\tau) = 0\} . \tag{34}$$

Our 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. (10) it suffices to prove that

$$Pr\big\{t - \underline{t} \leq T\big\} \geq 1 - \varepsilon . \tag{35}$$

The following lemma establishes such a busy period bound for a scheduler that operates at a constant rate C. We will show in Section 4 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. (10).

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

$$\sum_{\tau=1}^{\infty} \sup_{t\geq 0} Pr\left\{A_{\mathcal{C}}(t+\tau) - A_{\mathcal{C}}(t) > C\tau\right\} < \infty \,. \tag{36}$$

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

$$\sum_{\tau=T+1}^{\infty} \sup_{t\geq 0} Pr\left\{A_{\mathcal{C}}(t+\tau) - A_{\mathcal{C}}(t) > C\tau\right\} \leq \varepsilon \,. \tag{37}$$

Then T is a probabilistic bound on the busy period that satisfies Eqn. (35).

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

$$\begin{aligned} Pr\big\{t-\underline{t} > T\big\} &\leq Pr\big\{\exists\tau > T:\ A_{\mathcal{C}}(t) - A_{\mathcal{C}}(t-\tau) > C\tau\big\} & (38)\\ &\leq \sum_{\tau=T+1}^{\infty} Pr\big\{A_{\mathcal{C}}(t) - A_{\mathcal{C}}(t-\tau) > C\tau\big\} & (39)\\ &\leq \varepsilon \,, & (40)\end{aligned}$$

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* [46] with $S = K(t-L)$ for a rate K and a latency L.

The assumption in Eqn. (36) 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 [13, 45, 51]. This includes some long-range dependent processes but can fail for heavy-tailed arrival models. Some examples are discussed in Section 3.2.

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 $\overline{\mathcal{G}}^{net,\varepsilon}_{\mathcal{C}}$ for the arrivals to the network satisfying

$$Pr\big(A^{net}_{\mathcal{C}}(t+\tau) - A^{net}_{\mathcal{C}}(t) > \overline{\mathcal{G}}^{net,\varepsilon}_{\mathcal{C}}(\tau)\big) \leq \frac{2\varepsilon}{\pi(1+\tau^2)} \,. \tag{41}$$

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

At the first node on the path of the flow, we set

$$T^1 = \sup\{\tau \geq 0 \ | \ \overline{\mathcal{G}}^{net,\varepsilon}_{\mathcal{C}}(\tau) > C\tau\} \,, \tag{42}$$

in analogy with Eqn. (33). This choice of T^1 satisfies Eqn. (37), because

$$\begin{aligned}
&\sum_{\tau=T^1+1}^{\infty} \sup_{t\geq 0} Pr\left\{A_{\mathcal{C}}(t+\tau) - A_{\mathcal{C}}(t) > C\tau\right\} \\
&\quad\leq \sum_{\tau=T^1+1}^{\infty} \sup_{t\geq 0} Pr\left\{A_{\mathcal{C}}(t+\tau) - A_{\mathcal{C}}(t) > \overline{\mathcal{G}}^{net,\varepsilon}_{\mathcal{C}}(\tau)\right\} && (43)\\
&\quad\leq \frac{2\varepsilon}{\pi}\sum_{\tau=0}^{\infty}(1+\tau^2)^{-1} && (44)\\
&\quad\leq \varepsilon \,. && (45)
\end{aligned}$$

By Lemma 1, this provides the desired time scale bound at the ingress node.

The arrivals to the h-th downstream node are bounded in terms of the arrivals to the ingress node by

$$A^h_{\mathcal{C}}(t) - A^h_{\mathcal{C}}(t-\tau) \leq A^{net}_{\mathcal{C}}(t) - A^{net}_{\mathcal{C}}(t-\tau-(h-1)d^*) \,,$$

where d^* is the delay threshold for dropping a packet. It follows that

$$\sum_{\tau=0}^{\infty} Pr\big\{A^h_{\mathcal{C}}(t) - A^h_{\mathcal{C}}(t-\tau) > \overline{\mathcal{G}}^{net,\varepsilon}_{\mathcal{C}}(\tau+(h-1)d^*)\big\} \leq \varepsilon \,. \tag{46}$$

Invoking Lemma 1 as above, we obtain the time scale bound

$$T^h = \sup\{\tau \geq 0 \ | \ \overline{\mathcal{G}}^{net,\varepsilon}_{\mathcal{C}}(\tau+(h-1)d^*) > C\tau\} \,. \tag{47}$$

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. Even though these choices are clearly very conservative, leading to rather loose bounds on T^h, the numerical results on backlog and delay are satisfactory.

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 [49]. 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. (11) and Eqn. (26) 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 $\varepsilon \to 0$.

3 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 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 [15, 30, 33] 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.

In this paper we use the general definition from [30], which defines the *effective bandwidth* of an arrival process A as

$$\alpha(s,\tau) = \sup_{t\geq 0}\left\{\frac{1}{s\tau}\log E[e^{s(A(t+\tau)-A(t))}]\right\}, \qquad s,\tau \in (0,\infty) . \tag{48}$$

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. Generally, the effective bandwidth of a traffic flow varies between the mean and peak rates of the traffic and provides a link between the traffic characteristics of a flow and the resources in terms of bandwidth and buffer size necessary to support a required level of service. 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).

3.1 Overview of Effective Bandwidth

The notion of effective bandwidth emerged in the early 1990s in [26, 27, 28, 29] as a method to characterize and exploit the statistical multiplexing gain of traffic flows and, thereby, increase the utilization of network resources. The effective bandwidth of a traffic flow can be related to the minimum bandwidth needed to satisfy service guarantees for that flow. Then, one can verify that a link with capacity C is able to provide the required service to N traffic flows by testing if $\sum_{i=1}^{N}\alpha_i < C$, where α_i is the effective bandwidth of the i-th flow for suitable choices of s and τ.

Early work on effective bandwidth focused on Markov modulated fluid flow and on-off traffic models [26, 27, 29]. By relating the effective bandwidth concept to the theory of large deviations in [13, 21, 50], the effective bandwidth theory could be extended to a wide range of network traffic models including general Markovian and self-similar traffic models [21, 22]. The theory has also been generalized from FIFO scheduling algorithm to non-FIFO scheduling algorithms such as Static Priority (SP) [4, 5, 23, 34] and Generalized Processor Sharing (GPS) [47, 52], and has become an elegant and powerful framework with many applications.

A crucial result in the effective bandwidth theory concerns the large buffer asymptotics for links with FIFO scheduling. The result states that $\sum_{i=1}^{N}\alpha_i(s) < C$ if and only if $Pr\big(B > x\big) \sim e^{-sx}$ as $x \to \infty$, where $\alpha_i(s) = \lim_{\tau\to\infty}\alpha_i(s,\tau)$ and B is the steady-state backlog of the traffic. In other words, 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 may either overestimate or underestimate the actual backlog behavior by several orders of magnitude, especially if arriving traffic is bursty [17]. Furthermore, in the asymptotic regime, the bandwidth requirements given by the effective bandwidth are additive, and, hence, do not reflect the gains due to statistical multiplexing [17].

The asymptotic bounds from the effective bandwidth literature are not directly applicable in a network calculus context. Instead, when we insert effective bandwidth expressions in the network calculus we need to work explicitly with finite buffer sizes. Such non-asymptotic bounds have been presented by Chang [13, 15] for a class of linear envelope processes with parameters $\big(\sigma(s), \rho(s)\big)$, characterized by

$$\frac{1}{s}\log\big(e^{sA(t,t+\tau)}\big) \leq \sigma(s) + \rho(s)\tau \ . \tag{49}$$

If $\rho(s) < C$ for these processes, Chang [13] 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 relates these and other results on envelope processes to draw analogies to the deterministic network calculus [19]. Chang [13] 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 (exponentially) in the number of nodes. (Closely related results, without referring to effective bandwidth, are obtained by Yaron and Sidi for the class of exponentially bounded burstiness [51]).

The motivation for our work is to further develop the relationship between effective bandwidth and the network calculus. Our results, all expressed as explicit (non-asymptotic) bounds, extend the relationships established by Chang in several directions. First, we do not restrict ourselves to a specific class of arrival models, but consider all arrival models for which effective bandwidth expressions are available. For example, we consider FBM traffic which has been used to model self-similar characteristics of network traffic, but which cannot be characterized by a linear envelope process. Second, using the network calculus from Section 2, our results can be related to a (effective) network service curve which yields end-to-end backlog and delay bounds over multiple nodes. Lastly, we will (in Section 4) consider a number of commonly used scheduling algorithms, which are more complex than FIFO scheduling used predominantly in the effective bandwidth literature.

3.2 Relating Effective Bandwidth and Effective Envelopes

The choice of the term 'effective envelope' as introduced in [7] 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*

$$\mathcal{G}^{\varepsilon}(\tau) \quad = \quad \inf_{s>0}\Big\{\tau\alpha(s,\tau) - \frac{\log\varepsilon}{s}\Big\} \ . \tag{50}$$

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

$$\alpha(s,\tau) \quad \leq \quad \frac{1}{s\tau}\log\left(\int_0^1 e^{s\mathcal{G}^{\varepsilon}(\tau)}d\varepsilon\right) \ . \tag{51}$$

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. As another way to see the generality of the 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 an effective envelope $\mathcal{G}^{\varepsilon}(\tau)$ according to Eqn. (3), which is finite for all values of ε and τ.

Proof. To prove the first statement, fix $t, \tau \geq 0$. By the Chernoff bound [39],[1] we have for any x and any $s \geq 0$

$$Pr\Big\{A(t+\tau) - A(t) \geq x\Big\} \leq e^{-sx} E\left[e^{s(A(t+\tau)-A(t))}\right] \tag{52}$$

$$\leq e^{s(-x+\tau\alpha(s,\tau))} . \tag{53}$$

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

$$x^{\varepsilon,s}(\tau) = \tau\alpha(s,\tau) - \frac{\log \varepsilon}{s} \tag{54}$$

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

$$F^{t,\tau}(x) = Pr\big\{A(t+\tau) - A(t) \leq x\big\} \tag{55}$$

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

$$E\left[e^{s(A(t+\tau)-A(t))}\right] = \int_0^\infty e^{sx} dF^{t,\tau}(x) . \tag{56}$$

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

$$Pr\left\{A(t+\tau) - A(t) > G^{t,\tau}(\varepsilon)\right\} = \varepsilon , \tag{57}$$

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

$$E\left[e^{s\left(A(t+\tau)-A(t)\right)}\right] = \int_0^1 e^{sG^{t,\tau}(\varepsilon)} d\varepsilon \leq \int_0^1 e^{s\mathcal{G}^{\varepsilon}(\tau)} d\varepsilon . \tag{58}$$

It follows that

$$\alpha(s,\tau) \leq \frac{1}{s\tau} \int_0^1 e^{s\mathcal{G}^{\varepsilon}(\tau)} d\varepsilon , \tag{59}$$

as claimed. $\square$

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., [15, 30]), Lemma 2 provides a useful tool to apply the presented network calculus to a wide range of traffic models. 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 $Pr\Big\{X \geq x\Big\} < e^{-sx} E\left[e^{sX}\right]$.

3.3 Regulated Arrivals

We refer to arrivals that are bounded by an arrival envelope A^* (see Subsection 2.3) as 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

$$\forall t, \tau \geq 0: \quad A(t+\tau) - A(t) \leq A^*(\tau) \tag{60}$$

holds for every sample path. The peak rate and the average rate of regulated traffic, denoted by P and ρ, are defined as

$$P = A^*(1)\,, \quad \rho = \lim_{t\to\infty} \frac{A^*(t)}{t}\,. \tag{61}$$

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. 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

$$E\big[A_i(t+\tau) - A_i(t)\big] \leq \rho_i \tau\,, \tag{62}$$

and that the arrivals from different flows are independent. The effective bandwidth for such a collection of flows $A_\mathcal{C}$ satisfies [30]

$$\alpha_\mathcal{C}(s,t) \leq \frac{1}{st} \sum_{i\in\mathcal{C}} \log\left(1 + \frac{\rho_i t}{A_i^*(t)}(e^{sA_i^*(t)} - 1)\right)\,. \tag{63}$$

By Lemma 2, the corresponding effective envelope is given by

$$\mathcal{G}_\mathcal{C}^\varepsilon(t) \;=\; \inf_{s>0}\Big\{\sum_{i\in\mathcal{C}} \frac{1}{s}\log\Big(1 + \frac{\rho_i t}{A_i^*(t)}(e^{sA_i^*(t)} - 1)\Big) - \frac{\log\varepsilon}{s}\Big\}\,. \tag{64}$$

This effective envelope satisfies

$$\rho_\mathcal{C} t \leq \mathcal{G}_\mathcal{C}^\varepsilon(t) \leq A_\mathcal{C}^*(t) \tag{65}$$

for all $t \geq 0$.

3.4 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. 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 [15]

$$\alpha_\mathcal{C}(s,t) = \frac{1}{s}\sum_{i\in\mathcal{C}} \log\left(1 + \frac{\rho_i}{P_i}(e^{P_i s} - 1)\right)\,. \tag{66}$$

Lemma 2 gives the corresponding effective envelope as

$$\mathcal{G}_\mathcal{C}^\varepsilon(t) \;=\; \inf_{s>0}\Big\{\frac{t}{s}\sum_{i\in\mathcal{C}} \log\Big(1 + \frac{\rho_i}{P_i}(e^{P_i s} - 1)\Big) - \frac{\log\varepsilon}{s}\Big\}\,. \tag{67}$$

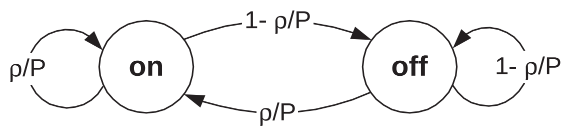

Figure 2: On-Off Transition Model.

3.5 Fractional Brownian Motion (FBM) traffic

As pointed out in [38], the self-similarity properties of measured traffic data can sometimes be modeled by processes of the form

$$A(t) = \rho t + \beta Z_t \; , \tag{68}$$

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 [38], we will refer to Eqn. (68) 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 $\beta_{\mathcal{C}}^2 = \sum_{i \in \mathcal{C}} \beta_i^2$. 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 [25].

We remark that the FBM 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 has been derived as [30]

$$\alpha_{\mathcal{C}}(s,t) = \rho_{\mathcal{C}} + \frac{1}{2} \beta_{\mathcal{C}}^2 s t^{2H-1} \; . \tag{69}$$

By Lemma 2, this results in an effective envelope of

$$\mathcal{G}_{\mathcal{C}}^{\varepsilon}(t) \quad = \quad \rho_{\mathcal{C}} t + \sqrt{-2 \log \varepsilon}\, \beta_{\mathcal{C}} t^H \; . \tag{70}$$

4 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 $q = 1, \ldots, Q$ 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 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 deadlines can become negative.

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 [41]. 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 [36] such leftover service curves where used to derive lower bounds on the service for an individual flow when the scheduling algorithms are not known ([9], 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 $q = 1, \ldots, Q$, let $\mathcal{G}_q^{\varepsilon_g}$ 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. (35) 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 $\mathcal{G}_p^{\varepsilon_p}$ are concave. Define functions $\mathcal{S}_q^{\varepsilon_s}$ as follows:*[2]

1. **SP:**
$$\mathcal{S}_q^{\varepsilon_s}(t) = \Big[Ct - \sum_{p<q} \mathcal{G}_p^{\varepsilon_g}(t)\Big]_+ , \quad \varepsilon_s = \varepsilon_b + (q-1)T\varepsilon_g \,. \tag{71}$$

2. **EDF:**
$$\mathcal{S}_q^{\varepsilon_s}(t) = \Big[Ct - \sum_{p\neq q} \mathcal{G}_p^{\varepsilon_g}(t - [d_p - d_q]_+)\Big]_+ , \quad \varepsilon_s = \varepsilon_b + (Q-1)T\varepsilon_g \,. \tag{72}$$

3. **GPS:**
$$\mathcal{S}_q^{\varepsilon_s}(t) = \lambda_q\Big(Ct + \sum_{p\neq q} \Big[\lambda_p Ct - \mathcal{G}_p^{\varepsilon_g}(t)\Big]_+\Big) , \quad \varepsilon_s = \varepsilon_b + (Q-1)T\varepsilon_g \,, \tag{73}$$

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

$$Pr\Big\{D_q(t) \geq \inf_{\tau \leq T}\{A_q(t-\tau) + \mathcal{S}_q^{\varepsilon_s}(\tau)\}\Big\} \geq 1 - \varepsilon_s \,. \tag{74}$$

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

By setting all violation probabilities ε_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 $\mathcal{S}_q^{\varepsilon_s}$ is 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 in Eqs. (71)–(73). 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.

The formulas in Eqs. (71)–(73) 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. (71) describes the performance of an SP scheduler rather closely. Eqn. (73) 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. (72) 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. (72) approaches the service guarantees of an SP scheduler for the lowest priority class, while the actual EDF scheduler approaches FIFO.

Proof. We show that Eqn. (74) holds separately for each of the scheduling algorithms.

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

$$\underline{t}_{\leq q} = \max\big\{x \leq t: \ B_{\leq q}(x) = 0\big\} \tag{75}$$

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

$$\begin{aligned} D_q(t) &= D_q(\underline{t}_{\leq q}) + \big(D_{\leq q}(t) - D_{\leq q}(\underline{t}_{\leq q})\big) - \big(D_{<q}(t) - D_{<q}(\underline{t}_{\leq q})\big) && (76)\\ &\geq A_q(\underline{t}_{\leq q}) + \Big[C(t - \underline{t}_{\leq q}) - \big(A_{<q}(t) - A_{<q}(\underline{t}_{\leq q})\big)\Big]_+ . && (77) \end{aligned}$$

In Eqn. (77), 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

$$\begin{aligned} &Pr\Big\{D_q(t) \geq \inf_{\tau \leq T}\big(A_q(t-\tau) + \mathcal{S}_q^{\varepsilon_s}(\tau)\big)\Big\} \\ &\quad\geq Pr\Big\{t - \underline{t}_{\leq q} \leq T \text{ and } D_q(t) \geq A_q(\underline{t}_{\leq q}) + \Big[C(t - \underline{t}_{\leq q}) - \sum_{p<q} \mathcal{G}_p^{\varepsilon_g}(t - \underline{t}_{\leq q})\Big]_+\Big\} && (78)\\ &\quad\geq Pr\Big\{t - \underline{t}_{\leq q} \leq T \text{ and } A_{<q}(t) - A_{<q}(\underline{t}_{\leq q}) \leq \sum_{p<q} \mathcal{G}_p^{\varepsilon_g}(t - \underline{t}_{\leq q})\Big\} && (79)\\ &\quad\geq Pr\Big\{t - \underline{t} \leq T \text{ and } \forall p<q, \ \forall \tau \leq T: \ A_p(t) - A_p(t-\tau) \leq \mathcal{G}_p^{\varepsilon_g}(\tau)\Big\} && (80)\\ &\quad\geq 1 - \Big(\varepsilon_b + (q-1)T\varepsilon_g\Big) , && (81) \end{aligned}$$

where $\underline{t}$ is the beginning of the busy period of the scheduler. In Eqn. (78), we have set $\tau = t - \underline{t}_{\leq q}$ and inserted the definition of $\mathcal{S}_q^{\varepsilon_s}$, and in Eqn. (79), we have used Eqn. (77). In Eqn. (80), 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

$$\begin{aligned} D_q(t) &= D_q(\underline{t}) + \big(D_{\mathcal{C}}(t) - D_{\mathcal{C}}(\underline{t})\big) - \sum_{p\neq q}\big(D_p(t) - D_p(\underline{t})\big) && (82)\\ &\geq A_q(\underline{t}) + \Big[C(t-\underline{t}) - \sum_{p\neq q}\big(A_p(t-(d_p-d_q)_+) - A_p(\underline{t})\big)\Big]_+ . && (83)\end{aligned}$$

We argue as in Eqs. (78)-(81) that

$$\begin{aligned} &Pr\Big\{D_q(t) \geq \inf_{\tau\leq T}\Big(A_q(t-\tau) + \mathcal{S}_q^{\varepsilon_s}(\tau)\Big)\Big\}\\ &\quad\geq Pr\Big\{t-\underline{t}\leq T \text{ and } \forall p\neq q,\ \forall\tau\leq T:\ A_p(t) - A_p(t-\tau) \leq \mathcal{G}_p^{\varepsilon_g}(\tau)\Big\} && (84)\\ &\quad\geq 1 - \Big(\varepsilon_b + (Q-1)T\varepsilon_g\Big) . && (85)\end{aligned}$$

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

$$\underline{t}_p = \max\{x \leq t:\ B_p(x) = 0\} \tag{86}$$

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

$$B_p(t) = A_p(t) - D_p(t) \leq A_p(t) - A_p(\underline{t}_p) - \lambda_p C(t-\underline{t}_p) \tag{87}$$

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

$$\underline{t}_{qp} = \max\{x \leq \underline{t}_q:\ B_p(x) = 0\} , \tag{88}$$

then Eqn. (87) with t replaced by $\underline{t}_q$ and $\underline{t}_p$ replaced by $\underline{t}_{pq}$ implies that

$$\begin{aligned} D_p(t) - D_p(\underline{t}_q) &\leq A_p(t) - A_p(\underline{t}_q) + B_p(\underline{t}_q) && (89)\\ &\leq A_p(t) - A_p(\underline{t}_{qp}) - \lambda_p C(\underline{t}_q - \underline{t}_{qp}) . && (90)\end{aligned}$$

It follows that

$$\begin{aligned} D_q(t) - D_q(\underline{t}_q) &\geq \lambda_q\Big(C(t-\underline{t}_q) + \sum_{p\neq q}\big[\lambda_p C(t-\underline{t}_q) - D_p(t) + D_p(\underline{t}_q)\big]_+\Big) && (91)\\ &\geq \lambda_q\Big(C(t-\underline{t}_q) + \sum_{p\neq q}\big[\lambda_p C(t-\underline{t}_{qp}) - A_p(t) + A_p(\underline{t}_{qp})\big]_+\Big) . && (92)\end{aligned}$$

Fix $t \geq 0$, let $\underline{t}$ be the beginning of the busy period containing time t, and assume for the moment that

$$t - \underline{t} \leq T \text{ and } \forall p\neq q, \forall\tau\leq T:\ A_p(t) - A_p(t-\tau) \leq \mathcal{G}_p^{\varepsilon_g}(\tau) . \tag{93}$$

Since $\underline{t} \leq \underline{t}_{qp} \leq \underline{t}_q$, it follows with by Eqn. (92) that

$$\begin{aligned} D_q(t) &\geq D_q(\underline{t}_q) + \lambda_q\Big(C(t-\underline{t}_q) + \sum_{p\neq q}\big[\lambda_p C(t-\underline{t}_{qp}) - A_p(t) + A_p(\underline{t}_{qp})\big]_+\Big) && (94)\\ &\geq A_q(\underline{t}_q) + \lambda_q\Big(C(t-\underline{t}_q) + \sum_{p\neq q}\big[\lambda_p C(t-\underline{t}_{qp}) - \mathcal{G}_p^{\varepsilon_g}(t-\underline{t}_{pq})\big]_+\Big) . && (95)\end{aligned}$$

(t)

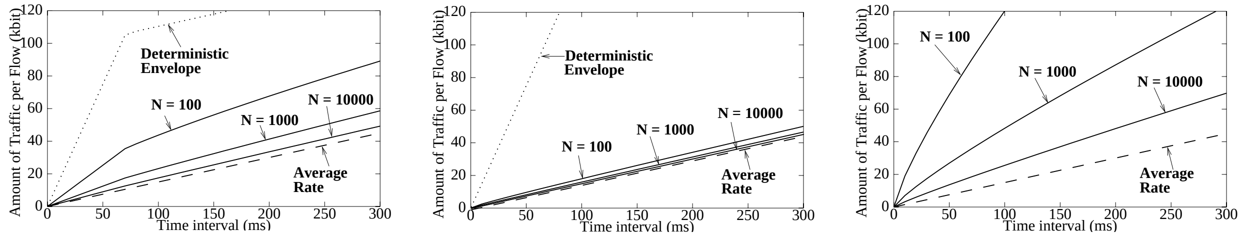

(a) (P, σ, ρ) Regulated. (b) On-Off. (c) FBM.

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

Since $\mathcal{G}_p^{\varepsilon_g}$ is concave, the function $\left[\lambda C t - \mathcal{G}_p^{\varepsilon_g}(t)\right]_+$ is nondecreasing in t. Replacing $t - \underline{t}_{pq}$ with the smaller value $t - \underline{t}_q$ in Eqn. (95) and using the definition of $\mathcal{S}_q^{\varepsilon}$ yields

$$D_q(t) \geq A_q(\underline{t}_q) + \mathcal{S}_q^{\varepsilon}(t - \underline{t}_q) \; . \tag{96}$$

Finally, we estimate

$$\begin{aligned} Pr\Big\{ D_q(t) \geq \inf_{\tau \leq T} \Big(A_q(t-\tau) + \mathcal{S}_q^{\varepsilon}(\tau) \Big) \Big\} &\geq Pr\Big\{ t - \underline{t}_q \leq T \text{ and Eqn. (96) holds} \Big\} && (97) \\ &\geq Pr\Big\{ \text{Eqn. (93) holds} \Big\} && (98) \\ &\geq 1 - \Big(\varepsilon_b + (Q-1) T \varepsilon_g \Big) \; . && (99) \end{aligned}$$

This completes the proof. □

5 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.

	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 1: Source Traffic Parameters.

For each of the three traffic models, we consider two types of flows. The parameters are given in Table 1. 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 $A^*(t) = \min\left(Pt\, , \sigma + \rho t\right)$, with parameters as in [7]. 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 [38], and select $\beta = 4.5$.

(t)

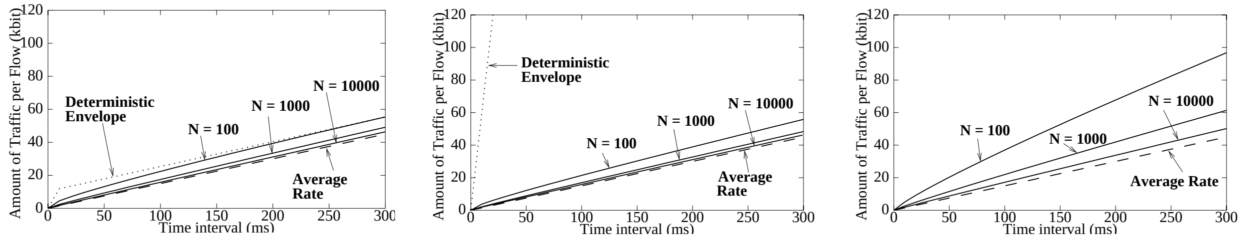

(a) (P, σ, ρ) Regulated. (b) On-Off. (c) FBM.

Figure 4: **Example 1:** Per-flow effective envelopes $\mathcal{G}_N^{\varepsilon}(t)/N$ for Type-2 flows (with $\varepsilon = 10^{-9}$).

5.1 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 $\mathcal{G}_N^{\varepsilon}(t)$ 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 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.

5.2 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.

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 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 5 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

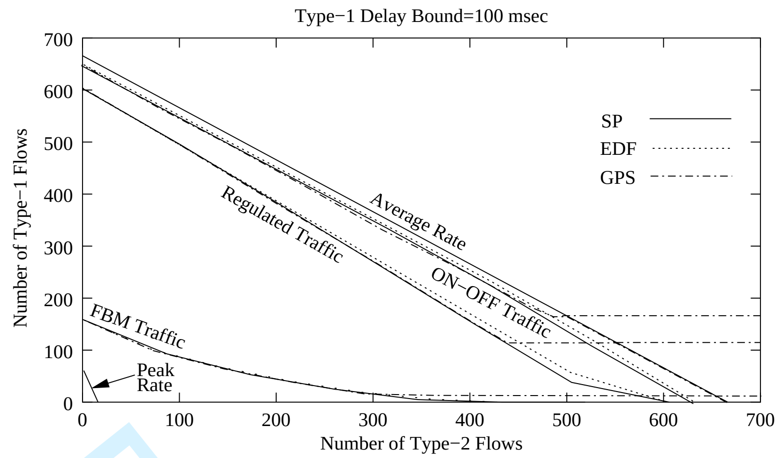

Figure 5: **Example 2:** Number of admissible Type-1 flows as a function of the number of Type-2 flows (C = 100 Mbps) for different schedulers and traffic models with $\varepsilon = 10^{-6}$, $d_1 = 100$ ms, $\phi_1 = 0.25$, $\phi_2 = 0.75$.

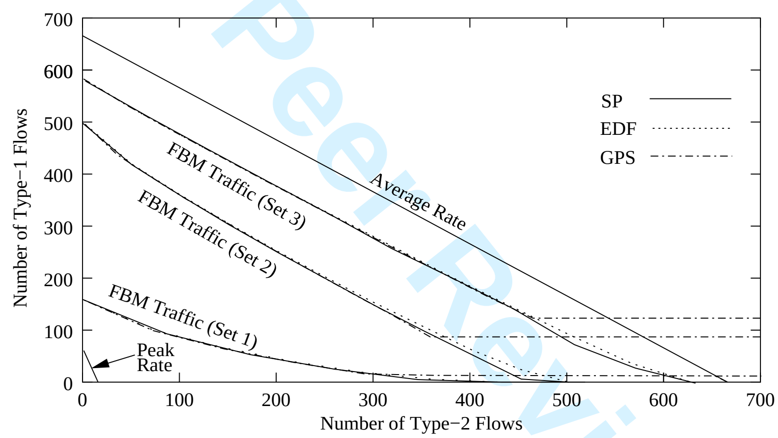

Figure 6: **Example 2:** Number of admissible Type-1 flows as a function of the number of Type-2 flows (C = 100 Mbps) for FBM traffic with different choices of β with $\varepsilon = 10^{-6}$, $d_1 = 100$ ms, $\phi_1 = 0.25$, $\phi_2 = 0.75$.

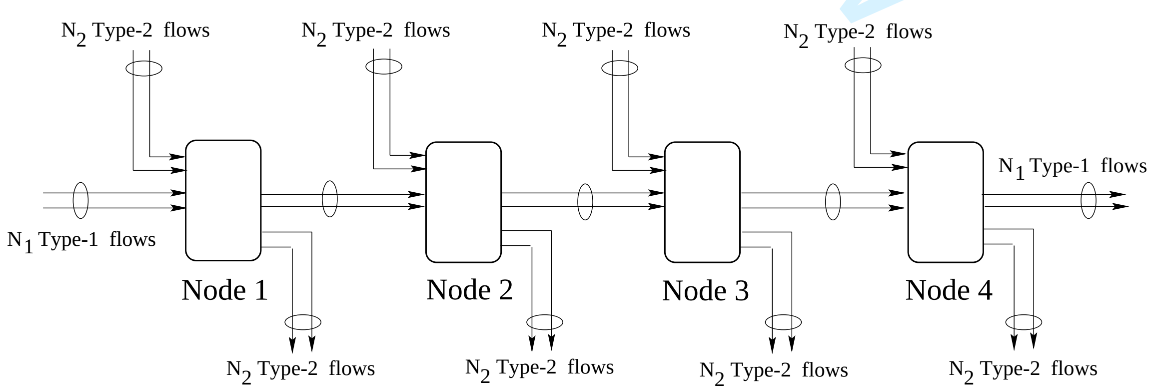

Figure 7: **Example 3:** A network with four nodes and with cross traffic.

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. Given a traffic model, the number of admitted Type-1 flows is similar for all scheduling algorithms, 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 due to the rate guarantee provided by GPS, which guarantee a minimum number of Type-1 flows: 114 flows for regulated traffic, 165 for On-Off traffic, and 12 for FBM traffic.

We emphasize again that the low multiplexing gain of FBM traffic is a result of our choice of parameters H and β. To illustrate this point, we present results for FBM traffic with different parameters, shown in Table 2. We consider three different sets of parameters. In Set 1, we use the same parameters as in Example 1. For Set 2, we select β so that the variance of FBM traffic is matched with the variance of regulated sources at a time scale corresponding to the delay bounds. This is 100 ms for Type-1 traffic and 10 ms for Type-2 traffic. For Set 3, we match the variance of FBM traffic at a time scale of 1000 ms, which is comparable to the longest busy period observed in these experiments. The results for the number of flows that can be admitted, shown in Figure 5, illustrate the dependency of the results on the parameter selection. For Set 3, FBM traffic exhibits a similar multiplexing gain as On-Off traffic.

	SET 1	SET 2	SET 3
Type	β (Mbps)	β (Mbps)	β (Mbps)
1	4.5	1.04	0.40
2	0.94	0.65	.13

Table 2: Parameters for FBM traffic.

5.3 Example 3: Multiple Nodes with Cross Traffic.

In this example, we consider a network with four nodes, as shown in Figure 7. 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 $N_1 = N_2$.

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 each node for violation probabilities $\varepsilon = 10^{-3}, 10^{-6}, 10^{-9}$ using the approach outlined in Subsection 2.6 with the number of classes $Q = 2$. We use the formula for the effective envelope given in Eqn. (50), 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. (41). At the h-th node on the route of the through flows, we set $\overline{\mathcal{G}}_1^{h,\varepsilon/2}(\tau) = \overline{\mathcal{G}}_1^{net,\varepsilon/2}(\tau + (h-1)d^*)$, as given in Eqn. (47). 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 $\overline{\mathcal{G}}_2^{h,\varepsilon/2}(\tau) = \overline{\mathcal{G}}_2^{net,\varepsilon/2}(\tau)$ for the Type-2 flows at each node, where $\overline{\mathcal{G}}^{\varepsilon/2}$ is the function computed above. We then apply Lemma 1,

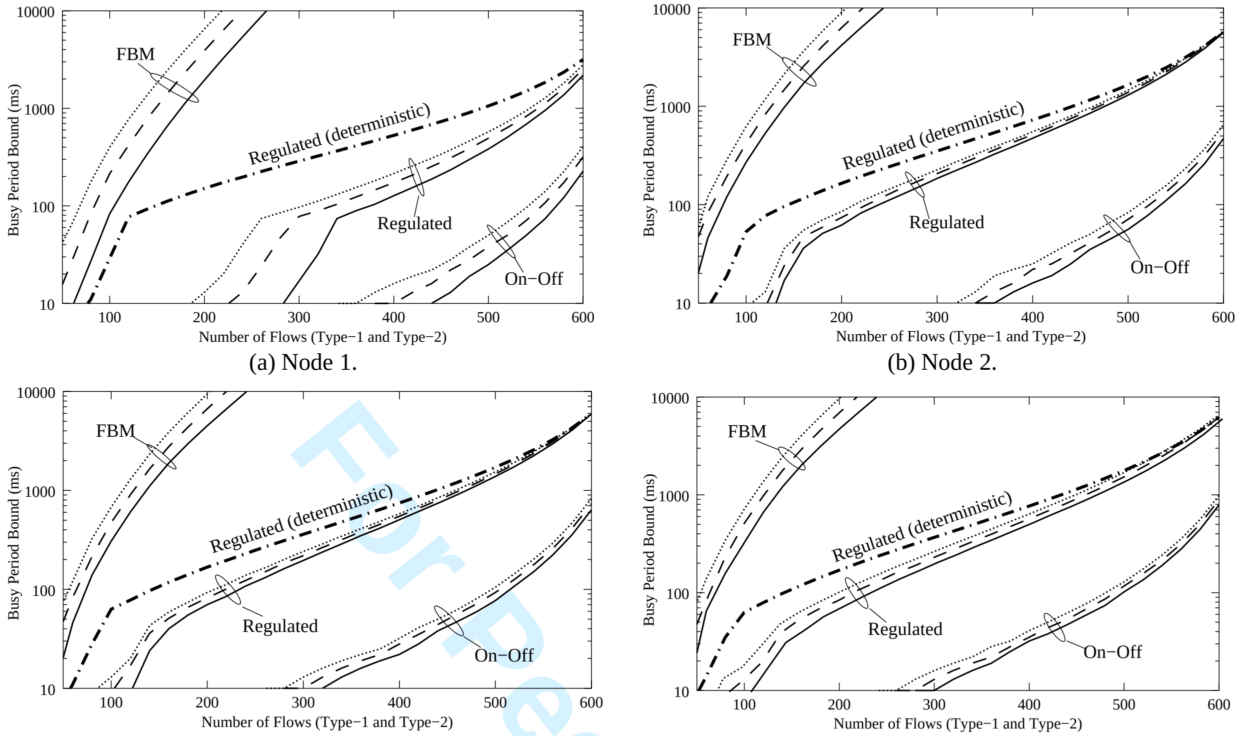

(a) Node 1. (b) Node 2.

(c) Node 3. (d) Node 4.

Figure 8: **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.

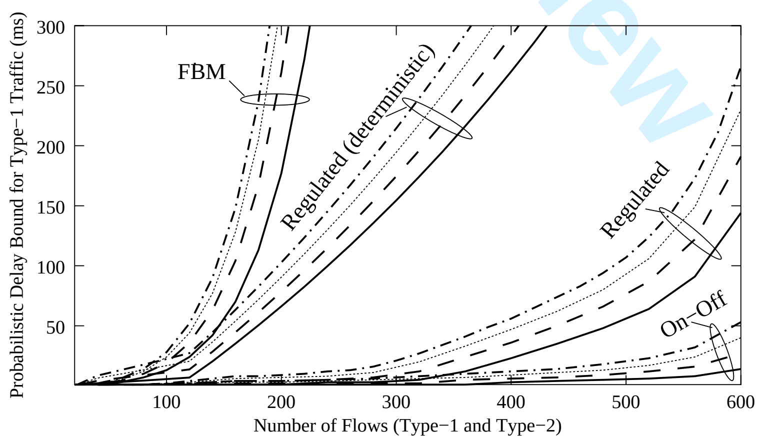

Figure 9: **Example 3:** Probabilistic bounds for the total queueing delay experienced by Type-1 traffic when leaving Node 1 (solid line), Node 2 (dashed line), Node 3 (dotted line), and Node 4 (dotted-dashed line) with violation probability $\varepsilon = 10^{-6}$. The x-axis corresponds to $N_1 + N_2$, the number of Type-1 and Type-2 flows, where we assume $N_1 = N_2$.

with $\overline{\mathcal{G}}_{\mathcal{C}}^{h,\varepsilon} = \overline{\mathcal{G}}_1^{h,\varepsilon/2} + \overline{\mathcal{G}}_2^{1,\varepsilon/2}$ to obtain bounds on the busy periods T^h at each node. 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 8 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 7. 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 9 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 9, 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 9 and omitted.

5.4 Example 4: Comparison to SBB calculus.

The next example compares network calculus results from this paper with the *stochastic bounded burstiness* analysis developed in [45], referred to as *SBB calculus*. In the SBB calculus, arrival bounds take the form

$$Pr\Big\{A(t+\tau) - A(t) \geq \rho\tau + \sigma\Big\} \leq f(\sigma)$$

where $f(\sigma)$ is a function such that the n-fold integration of f, denoted by $(\int_\sigma^\infty du)^n f(u)$, is finite. Arrival models in this class include the FBM traffic model. In the network calculus, the effective envelope for SBB arrivals of a flow are given by $\mathcal{G}_i^\varepsilon(t) = t + \sigma(\varepsilon)$, where $\sigma(\varepsilon)$ is obtained by solving $f(\sigma) = \varepsilon$.

The analysis in [45] considers a single-node work-conserving system, and derives bounds on backlog and output burstiness. The following example uses parameters from an example in ([45], Sec. IV.C). As in [45], we consider a single-node system with service rate C, where all flows have a rate $\rho = 1$ and a burstiness bound of $f_i(\sigma) = e^{-2.197\sigma} + 10^{-4} \cdot e^{-0.543\sigma}$. We modify the example from [45] in that we consider a node with capacity $C = 6$ and with five flows, indexed $i = 1, \cdots, 5$. For any work-conserving service discipline, the backlog bound is computed with Theorem 3 in [45].

We first consider the aggregate backlog. In Table 3, we compare the aggregate backlog from all flows, as obtained from Theorem 2 in [45] with those obtained with our Theorem 1 and Lemma 1 for various values of the violation probability ε. The table shows that the SBB calculus provides tighter backlog bounds for

Table 3: Comparison of bounds for the aggregate backlog.

$\varepsilon = 10^{-3}$		$\varepsilon = 10^{-6}$		$\varepsilon = 10^{-9}$	
THIS PAPER	SBB	THIS PAPER	SBB	THIS PAPER	SBB
30.2	20.5	100.4	66.8	168.5	130.4

the aggregate. The reason is that the derivation for the total backlog in the SBB calculus are done in a single estimate, whereas the network calculus makes one estimate for the busy period and another estimate for the backlog bound in Theorem 1.

The advantages of the network calculus approach become evident when we investigate the backlog of individual flows. Here, we obtain an effective service curve for a flow using the leftover service curves from Section 4. The resulting service curves are functions of the form $S_i^\varepsilon(t) = \Big[R \cdot t - X(\varepsilon)\Big]_+$, where R and $X(\varepsilon)$ are obtained from the SBB characteristics of the other flows. The backlog bound for a flow leaving the system is given by $b(\varepsilon) = \mathcal{G}^\varepsilon \oslash S_i^\varepsilon(0)$ following our Theorem 1.2. We analyze backlog bounds for all scheduling algorithms considered in this paper. For the SP service discipline, we assign flow i a priority index i. For GPS, we set the weight parameter equal at each node. For EDF, we set the flow-i delay index equal to i.

Table 4: Backlog bounds for individual flows.

	$\varepsilon = 10^{-3}$			$\varepsilon = 10^{-6}$			$\varepsilon = 10^{-9}$		
FLOW ID	SP	GPS	EDF	SP	GPS	EDF	SP	GPS	EDF
1	3.72	6.06	8.19	12.97	20.06	50.2	27.21	33.7	91.1
2	5.09	6.06	12.19	18.24	20.06	54.2	34.91	33.7	95.1
3	7.07	6.06	15.19	25.09	20.06	57.2	45.75	33.7	98.1
4	10.37	6.06	17.19	36.38	20.06	59.2	63.84	33.7	100.1
5	18.19	6.06	18.19	60.2	20.06	60.2	101.1	33.7	101.1

Table 4 shows the results of the backlog analysis. A comparison of the per-flow backlog bounds in Table 4 with the backlog bounds for the total traffic indicate that the per-flow bounds are much improved. In particular, note that with SP scheduling the backlog bounds all flows, including that for the lowest priority flow (Flow 5), are below the aggregate backlog bounds from Table 1. This demonstrates that the service description in Lemma 3 captures properties of the particular scheduling algorithm.

Remarks: The SBB calculus in [45] does not offer delay bounds or multi-node results, and has not been developed for non-FIFO scheduling algorithms. While it may be feasible to extend the SBB calculus framework to consider per-flow bounds in various scheduling algorithms, and derive delay bounds, such derivations will require a similar effort as the derivations in a min-plus algebra as done in this paper. It remains open whether the statistical calculus can be strengthened to a degree that it yields backlog bounds that are identical to those of the SBB calculus (from Table 1). For a subclass of so-called *exponentially bounded burstiness* (EBB) [51] the question has recently been answered in [18], which showed that a statistical network calculus can faithfully reproduce single node results of the EBB calculus. For a multi-node setting, a comparison of end-to-end performance bounds computed with the techniques from [51] to those obtained with the statistical network calculus showed that delay bounds from [51] scale with $\mathcal{O}\big(H^3\big)$, where H is the number of nodes, whereas the corresponding results in the statistical network calculus are bounded by $\mathcal{O}\left(H \log H\right)$.

6 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. Technical Report LA-UR-03-3939, Los Alamos National Laboratory, June 2003.

[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] A. W. Berger and W. Whitt. Effective bandwidths with priorities. *IEEE/ACM Transactions on Networking*, 6(4):447–460, August 1998.

[5] A. W. Berger and W. Whitt. Extending the effective bandwidth concept to networks with prioirty classes. *IEEE Communications Magazine*, 36(8):78–84, August 1998.

[6] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss. An architecture for differentiated services. IETF RFC 2475, December 1998.

[7] 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.

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

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

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

[11] 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.

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

[13] 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.

[14] 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.

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

[16] J. Choe and N. Shroff. Queueing analysis of high-speed multiplexers including long-range dependent arrival processes. In *Proceedings of IEEE INFOCOM '99*, New York, NY, March 1999.

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

[18] 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.

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

[20] R. Cruz. SCED+: efficient management of quality of service guarantees. In *Proceedings of IEEE INFOCOM '98*, San Francisco, CA, March 1998.

[21] N. Duffield and N. O'Connell. Large deviations and overflow probabilities for the general single server queue, with applications. *Mathematical Proceedings of the Cambridge Philosophical Society*, 118:363–374, 1995.

[22] A. Elwalid and D. Mitra. Effective bandwidth of general Markovian traffic sources and admission control of high speed networks. *IEEE/ACM Transactions on Networking*, 1(3):329–43, June 1993.

[23] A. Elwalid and D. Mitra. Analysis, appoximations and admission control of a multi-service multiplexing system with priorities. In *Proceedings of IEEE INFOCOM'95*, pages 463–472, Boston, MA, April 1995.

[24] 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.

[25] 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.

[26] 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.

[27] 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.

[28] J. Y. Hui. Resource allocation for broadband networks. *IEEE Journal on Selected Areas in Communications*, 6(9):1598–1608, 1988.

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

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

[31] E. Knightly. Second moment resource allocation in multi-service networks. In *Proceedings of ACM SIGMETRICS '97*, pages 181–191, Seattle, WA, June 1997.

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

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

[34] V. Kulkarni, L. Gun, and P. Chimento. Effective bandwidth vector for two-priority ATM traffic. In *Proceedings of IEEE INFOCOM'94*, pages 1056–1064, Toronto, Ontario, June 1994.

[35] 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.

[36] 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.

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

[38] 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.

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

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

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

[42] M. Reisslein, K. Ross, and S. Rajagopal. Guaranteeing statistical QoS to regulated traffic: The multiple node case. In *Proceedings of IEEE Conference on Decision and Control*, pages 531–538, Tampa, FL, December 1998.

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

[44] V. Sivaraman and F. Chiussi. Providing end-to-end statistical delay guarantees with earliest deadline first scheduling and per-hop traffic shaping. In *Proceedings of IEEE INFOCOM 2000*, Tel Aviv, Israel, March 2000.

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

[46] 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.

[47] G. De Veciana and G. Kesidis. Bandwidth allocation for multiple qualities of service using generalized processor sharing. *IEEE Transactions on Information Theory*, 42(1):268–272, 1996.

[48] M. Vojnovic and J.-Y. Le Boudec. Stochastic analysis of some expedited forwarding networks. In *Proceedings of IEEE Infocom 2002*, pages 1004–1013, New York, June 2002.

[49] 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.

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

[51] 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.

[52] Z. Zhang. Large deviations and the generalized processor sharing scheduling for a two queue system. *Queueing Systems, Theory and Applications*, 26(3-4):229–264, 1997.

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

To the Associate Editor:

Enclosed please find our revised paper and our responses to the reviewers. The editorial letter requested that we address Comment 2 from Reviewer 1 and Comments 1 – 4 from Reviewer 3. In addition, we have addressed *all* of the reviewers' comments, and, whenever warranted, have revised the paper to accommodate their suggestions.

In the following we discuss the comments of the reviewers in their order of appearance. Note that, as a consequence of the extraordinary delay of almost 2 years of returning the reviews, research on the network calculus that was done after the submission of this manuscript (sometimes citing the unpublished manuscript) has appeared in print.

The paper has been extensively revised to address the issues raised by the reviewers. We emphasize that the paper did not require technical corrections. Specifically, the technical concern in Comment 2 of Reviewer 1 assumes that a certain *sufficient* condition on the utilization that guarantees finite delay bounds is also *necessary*. However, the condition raised by the reviewer is not a necessary condition.

Response to Reviewer 1

```
1) A Possible Lack of New Main Contributions

A main contribution of the paper is claimed to be the established
relation between the notion of effective bandwidth and a  network
calculus.  I  am afraid that this may not be a sufficient finding
to make the paper a candidate for publication in  IEEE  ToN.  The
paper  appears  to  me incremental in the view of a large body of
the literature on probabilistic guarantees. What does the  rela-
tion with effective bandwidth bring us? In one view, it would en-
able us to obtain probabilistic envelopes for  bit  arrival  pro-
cesses  by  leveraging  on  effective bandwidth characterizations
that are known for some arrival processes. Some examples are giv-
en  in Section 3.3-3.5 in the paper. However, some of these enve-
lope characterizations are either known  or  follow  easily  from
some  related  work.  First,  for  regulated arrivals, we can use
Chernoff bound on the complementary distribution of bit  arrivals
over  a  given  time interval. The last is indeed maximized if we
replace the original process with  another  process  whereas  the
constituting  arrival  flows  are  Bernoulli with the probability
mass assigned to the end-points of the supports determined by the
arrival  curve  constraints;  this would yield (60) in the paper.
Second, for the on-off Markov example, the same applies. The  en-
velope for fractional Brownian motion may be new, but may also be
derived directly. For an  arrival  process,  for  which  we  know
bounds  on its increments and a bound on means of the increments,
we can obtain an envelope function by using Hoeffding's inequali-
ties.  This  was  found in some related work. The arguments above
```

may suggest that making a connection to effective bandwidth may not yield us much.

We argue that establishing a connection between teletraffic theories, such as network calculus and effective bandwidth, is significant. With the results of this paper much of the literature on the effective bandwidth theory can be directly applied in the network calculus context. There are many works that can now be replaced by simply applying effective bandwidth results.

With this paper, it is also feasible to analyze models that could not be analyzed before. As an example, there are no results in the literature that have analyzed non-trivial scheduling algorithm such as GPS or EDF with FBM traffic arrivals. Further, the calculus approach in the paper with its separate arrival and service descriptions makes it possible to make a sensitivity analysis that considers the impact of varying scheduling algorithms and arrival models and the multiplexing gain. For example, our paper permits us to make statements that changing the scheduling algorithm has a limited impact on the multiplexing gain.

The reviewer points out that for many *specific* systems a direct analysis is feasible. We strongly agree and emphasize that *for any given system* a direct analysis may lead to tighter bounds. Without arguing about the importance of having tight bounds, a direct analysis generally does not permit a comparative sensitivity analysis of the impact of various scheduling disciplines (as pointed out above). This is now discussed in the introduction.

We next respond in detail to the reviewer's comments on Sections 3.3 – 3.5. The purpose of these sections is to demonstrate with a few examples that envelope functions can be obtained easily with Lemma 2 from known effective bandwidth expressions.

- *Section 3.3:* It is a misconception that Hoeffding's inequalities can yield tighter bounds than those obtained via the Chernoff bound from known effective bandwidth bounds. The bound on the effective bandwidth of a collection of independent regulated flows in Section 3.3 is in fact equivalent to Lemma 1 in [H63]. The subsequent application of our Lemma 2 in this special case is equivalent to the Chernoff bound appearing in the proofs of the main theorems in [H63][1]. The famous Hoeffding's inequality (see, e.g., http://en.wikipedia.org/wiki/Hoeffding's_inequality) is the result of a further simplification and thus provides a bound that is *less* tight.

  The reviewer mentions in passing related work where Hoeffding's inequalities are used to obtain performance bounds for independent regulated arrivals. He may be referring to [V03] which applies Lemma 1 from [H63] and the Chernoff bound to obtain backlog bounds, and [B00], where a different construction of the envelope in Section 3.3 is used to analyze scheduling conditions for various schedulers. Since these are strictly single-node results and apply only to multiplexed regulated traffic models, the present paper is not the right place for an extensive numerical comparison.

- *Section 3.4:* The above discussion of Hoeffding's inequality applies also to Section 3.4. Moreover, the analysis of Section 3.4 is easily extended to the more general case of Markov-Modulated On-Off processes by using known results on the effective bandwidth of such a process [K96].

[1] See references at the end of this document.

```
Comment 2: A Technical Problem for Network Case

The  paper  may  give  impression that the proposed probabilistic
network calculus applies well for a network of  nodes.  This  may
not  be quite true. I appreciate an honest remark made by the au-
thors along this line in Section 6. However,  it  appears  to  me
that  there is still an additional difficulty to apply the calcu-
lus to the network case. This I point out now.

Consider a tandem of nodes, assumed to be FIFO,  work-conserving,
with  a  service  rate of c. Let the nodes be labeled 1 to h. Let
there be one bit arrival process that traverse the nodes 1 to  h.
Assume  the  bit  arrival  process  at the input to the node 1 is
leaky-bucket (rho,sigma) constrained. Now,  recall  Section  2.6,
page  11.  It tells us to bound the bit arrival process to a node
over a given time interval with the bit arrival  process  to  the
node  1  over an appropriately enlarged time interval. This gives
us an envelope for the arrival process to a node in the  network.
Now,  our node is FIFO and offers a service curve, so that we can
compute the worst-case delay through a node as a maximum horizon-
tal deviation between the arrival envelope and the service curve.
If the arrival envelope is with high  probability  (w.h.p),  then
the  delay bound is with a high probability. In a loose notation,
this gives us

d <= sigma/c + rho/c(h-1)d (w.h.p)

The result is that we have a bound on delay that holds  only  for
sufficiently  small  loads.  The  paper  does not state to have a
problem with this and definitely does not address  the  issue.  I
consider  this  to be a serious technical problem that would pre-
clude the proposed probabilistic calculus to  apply  for  network
case. From the results of deterministic network calculus, we know
that for a network with aggregate  scheduling,  under  some  more
general  assumptions  than  introduced  insofar, we know a finite
worst-case bound on delay for rho/c < 1/(h-1). Under the same as-
sumptions, we do not know a bound for any rho/c<1.
```

There is no technical problem. The results in our paper are correct and consistent with deterministic and statistical delay bounds published in the literature. Presumably, the reviewer refers to a result that appears in [Charny00]. Specifically, the cited delay bound and utilization threshold cited above can be directly related to Theorem 1 in [Charny00], in the special case where no bound on the peak rate is available and packetization effects are neglected. The theorem states that as long as each flow visits at most h nodes in the network and the link utilization is less than $1/(h-1)$, a certain explicit delay bound holds. The theorem requires no further assumptions on arrival models, network topology, or routing. It is known to be sharp in the sense that for higher utilizations, explicit delay bounds require additional information on the network topology and routing.

However, the theorem does *not* say that finite delay bounds are generally impossible for $\rho/c > 1/(h-1)$. In fact, such a statement would be false regardless whether the analysis is probabilistic or deterministic. Specifically, in any feedforward network topology, finite delay bounds are known to exist for any value of the utilization $\rho/c < 1$. In other words, the utilization bound cited by the reviewer is a sufficient condition for the above delay bound, but it is not necessary.

Even though the concern about correctness is not warranted, the reviewer raises an important, but almost certainly very challenging problem for the stochastic network calculus: Is it possible to devise a utilization bound in a stochastic setting that yields finite delays regardless of topology? Can the result of [Charny00] be extended beyond the regime $\rho/c < 1/(h-1)$ by taking advantage of statistical multiplexing? Can it be applied to arrival models that do not have finite worst-case bounds? The current state-of-the-art of the stochastic network calculus does not provide a handle to address these questions.

3) Miscellaneous Comments

- It would be nice to see some non-academic references to exemplify the use of the concept of the effective bandwidth in practice in order to better support the claimed importance.

As an example of non-academic references that use notions of the Effective Bandwidth, we offer the following US patent filings:

- 6,922,564, Admitting data flows to a multiple access network, Filed: May 30, 2003 (Motorola)
- 6,982,964, High performance ECL-to-ATM protocol network gateway, Filed: October 15, 2001
- 6,891,798, Estimating equivalent bandwidth capacity for a network transport device using on-line measurements, Filed: Feb 2, 2000 (Cisco)
- 6,697,369 Admission control adjustment in data networks using maximum cell count, Filed: Sep 28, 1999

In total, we found more than 70 US patents issued that relate to various forms of ”effective bandwidth” in their abstracts. We will be happy to comply with the reviewers’ request and include references to these or other patents. However, since it is uncommon to reference patent filings in IEEE or ACM journals, we refer to the Associate Editor for a decision whether to add these citations.

- Some care may be exercised with the assumption that 0 falls into a queue idle period.

There are two equivalent ways for obtaining bounds for the steady-state backlog, delay, and output burstiness:

- Assume the system is in the steady state. Then the backlog $B(t)$, the delay $W(t)$, and the output $D(t, t+\tau)$ over an interval of length τ are stationary. Hence it suffices to study the distributions of $B(0)$, $W(0)$ and $D(0, \tau)$. By [Chang00], Lemma 9.1.4, the steady-state backlog and delay stochastically dominate the backlog and delay observed in a system that started with empty queues. This approach is used, for example, in [V03].

- Start with empty queues at $t = 0$, and prove bounds on $B(t)$, $W(t)$ and $D(t, t+\tau)$ that do not depend on t. Since backlog and delay are stochastically increasing in t by ([Chang00], Lemma 9.1.4), such stationary bounds are automatically valid also in the steady-state. This is the approach used in our paper.

Choosing one or the other method is strictly a matter of technical convenience. We have added an explaining paragraph in Section 2.1, directly after the assumption that all network queues are empty at $t = 0$.

`- The statement in Section 2.3, while referring to some related work, "... have been derived for FIFO schedulers with a fixed service rate" is incorrect. Some of the cited work do indeed assume nodes to offer a service curve, with neither FIFO nor work-conserving assumption.`

The sentence has been corrected.

`- I suggest making some of the assumptions of the paper transparent in the Introduction whereas some of the results are announced. In particular, the fact that the paper redefines min-plus convolution and the fact that the paper assumes a packet discard whenever delay exceeds a fixed value.`

Presumably this refers to the appearance of the time scale T at the beginning of Section 2.5. The paper does not redefine the min-plus convolution. Assumptions on T are stated explicitly every time they are used (see Theorems 1 and 2). The paper has been revised in the introduction to state assumptions at the beginning of the paper.

Please see the response to Reviewer 3 for a detailed discussion of the role of T, and the pertinent revisions.

`- I do not quite favour proposing a network calculus for service curve nodes, and then making additional assumptions, such as FIFO, work-conserving, constant service rate. This is done at some places in the paper in an ad-hoc manner. I would prefer more to see an analysis that would carry on under the assumption of service curve nodes. Other node abstractions may be too restrictive in practice, such as for example assuming a strict service curve, as assumed for some of the results.`

The paper is rigorous and there are no hidden or ad-hoc introduced assumptions. Lemma 1 and Lemma 3 extend to nodes that provide a strict service curve $S_{\mathcal{C}}(t)$, if Ct is replaced with $S_{\mathcal{C}}(t)$ throughout. If a node offers to an aggregate of flows a service curve $S_{\mathcal{C}}(t)$ that is not strict, then the conclusion of Lemma 3 is weakened to

$$S_q^{\varepsilon_s}(t) = S_{\mathcal{C}}(t) - \sum_{p<q} \mathcal{G}_p^{\varepsilon_g}(t)$$

in the case of a SP scheduler, and correspondingly for EDF and GPS. The busy period bound in Lemma 1 remains a bound on the range of the convolution.

```
- There is no consideration of the packetization effects.
```

It is possible to account for packetization by convolving at each node with an appropriate (deterministic or stochastic) pure delay service curve, but this paper does not explicitly address packetization.

```
-  The  proposal  to bound the arrival process of bits as done in
the paper to deal with the network case was already  proposed  in
an  early  paper  by  Kurose [32], and used by others since then.
This is not appropriately referenced.
```

We insist that the discussion of related work is adequate. We provide full and generous credit to related work, including the contributions of Kurose.

```
- One may perhaps relate the priority-multiplexing formulas found
in Lemma 3 with those of  deterministic  network  calculus,  e.g.
Theorem 2.4.1 in the reference [9] of the paper.
```

We have added a reference to the corresponding deterministic left over service curve.

Response to Reviewer 3

The reviewer refers in the comments to the implications of the definition of the service curve and the relationship to the SBB framework. The main concern in the review is the role of the time scale T in Eqn. (10), and its relation to the definition of the service curve in Eqn. (4).

Comment 1:

```
First,  as  a  probabilistic  extension  of deterministic service
curve, a reader would expect to recover the deterministic service
curve  from Eqn. (10). However, this can hardly be done because of
the additional requirement on the existence of the time scale  T.
```

We have revised the paper and now explain in Section 2.6 how to recover the deterministic service curve as $\varepsilon \to 0$. If a node offers a (deterministic) service curve to a flow (i.e., $D(t) \geq A * S(t)$), we argue that the desired time scale bound is provided by

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

Indeed, $A * S(t) = \inf_{\tau \leq T} A(t-\tau) + S(\tau)$, and thus Eqn. (1) implies that Eqn. (10) holds with $\varepsilon = 0$. Furthermore, as $\varepsilon \to 0$, the value of T becomes irrelevant in Theorem 1 and 2, so that these also reduce to their classical deterministic counterparts as $\varepsilon \to 0$.

The above value of T is finite under the stability condition that the service curve eventually catches up on the arrival envelope. If the stability condition is violated, backlogs and delays need not be bounded. Note that this argument uses only the definition of the service curve in Eqn. (1) and does not require a work-conserving system or the concept of a busy period.

Comment 2:

```
Second,  while  not  explicitly stated, Eqn. (10) makes additional
requirement on traffic.  This is because, as stated in the paper,
the  time  scale is related to the maximum length of busy periods
of the considered server, and it is known that the  length  of  a
busy  period highly depends on the input traffic. Indeed, this is
also implied by Lemma 1, where, in order to  determine  the  time
scale,  the traffic needs to satisfy Eqn. (34). In other words, if
Lemma 1 is the only result that can be relied on to find the time
scale  T,  the paper does also (implicitly) have a priori assump-
tion on the traffic. The paper should clarify this, since  other-
wise  a reader might be misled that results in the paper had made
no assumption on the traffic arrival.
```

This comment has been thoroughly addressed in the revised paper, and several clarifications have been added. In particular, a discussion has been added to the Introduction, and the entire Section 2.6 has been rewritten.

The motivation for the time bounds, while mentioned, was not stressed in the original paper. Note that in Section 2.4 (What makes Calculus hard?), we list issues that make the network calculus difficult. At several places (below Eqn. (7) and Eqn. (9)) we mentioned that the difficulties can be dealt with if appropriate time bounds are available. Now, in this paper, we explore a network calculus that exploits such time scale bounds throughout. As pointed out in the review, time scale bounds can be obtained from the properties of traffic. However, there are other ways to obtain time scale bounds. For example, in the second half of Section 2.6 (after Eqn. (41) of the original submission, we discuss (and later exploit) that a priori delay thresholds or dropping policies can also result in time scale bounds. Based on the above comments, we have added a discussion in the Introduction that states that this paper attempts to explore a network calculus with time scale bounds.

Comment 3:

```
Third,  Lemma  1  plays the critical role in deriving the maximum
```

```
time scale throughout the paper. However, as  stated  above,  Eqn.
(34) in it is indeed a requirement on the traffic.  This require-
ment has also been used in the literature (e.g. SBB  [S00]),  with
which,  similar  bounds  as  in  Theorem 1 have been derived.  It
would be interesting to compare bounds in Theorem 1 and the  cor-
responding ones in [S00].  (Although in [S00], the considered serv-
er is FIFO, the extension to other types of servers can  also  be
done by using, say, the approach in Lemma 3 of the paper or other
approaches in the literature.)  The comparison  could  help  con-
vince readers the rationale of introducing the requirement of the
time scale.
```

We have reformulated Lemma 1 to make the conditions on A explicit (the previous formulation has moved to the second half of Section 2.6). As we explain after the proof of the lemma, the assumption amounts to two conditions, the network stability condition that the average arrival rate should lie below the capacity of the server ($\rho < C$), and an assumption on the tail of the arrival distribution. The condition on the tail of the arrival distribution is more technical and can possibly be relaxed by using different analytical techniques. We point out that this condition is satisfied by many commonly used arrival models. It is less restrictive than the SBB condition used in [S00] (which requires the tail of the arrival distribution to decay faster than any polynomial).

Following the reviewer's suggestion, we have added a numerical experiment comparing our single-node results (Lemma 3 + Lemma 1 + Theorem 1) with the main result in [S00], which is stated for a single-server system with any work-conserving service discipline.

We concur with the reviewer that the approach of Lemma 3 could be used to extend the results in [S00] to other types of servers. Such an extension might also be able to provide end-to-end delay bounds. Similarly, it would be interesting to compare out multinode bounds (Lemma 3 + Lemma 1 + Theorem 2 + Theorem 1) with the backlog bounds obtained by iterating the input-output relation from Theorem 3 in [S00]. These in-depth comparisons lie beyond the scope of the current paper, and we leave them for future work.

We note in passing that no delay bounds are stated in [S00]; the related work [Y93] suggests to estimate per-flow delay bounds with the length of the busy period, which is comparable to our time scale T.

Comment 4:

```
Fourth, since Lemma 1 used for obtaining the time scale is indeed
an additional requirement  on traffic,  it is hence required  for
the authors to provide other results that not only can be used to
determine the time scale but also really decorrelate arrivals and
departures.  Otherwise, using Eq. (10) as a general service model
for the probabilistic extension of network calculus is  not  con-
vincing.  (As  stated in the paper, there are special cases where
with additional requirements on the server and/or the tracc,  the
time  scale can be easily determined. The paper could be re-posi-
tioned for such cases, while not for general cases.)
```

Response: This is addressed in the revised Introduction and Section 2.6. We stress that Eq. (10) does not constitute a new service model; we identify the time scale T as a relevant quantity that needs to be controlled. Ideally, the Statistical Network Calculus should be able to provide estimates of T along with backlog and delay bounds at each node from effective envelopes at the ingress nodes; else, bounds on T can be deduced from external assumptions such as a priori limits on delay or buffer size. We now discuss the role of such external assumptions in the introduction as well as in Section 2.6.

References Cited in the Response

[B00] R.R. Boorstyn, A. Burchard, J. Liebeherr, and C. Oottamakorn. Statistical service assurances for traffic scheduling algorithms. JSAC 18(12):2651-2664, December 2000.

[Chang00] C.S. Chang, Performance guarantees in communication networks. Springer, 2000.

[Charny00] ”Delay Bounds in a Network with Aggregate Scheduling” by A. Charny and J.-Y. Le Boudec. Proceedings of the First COST 263 International Workshop on Quality of Future Internet Services table of contents, Pages 1 - 13, 2000.

[H63] W. Hoeffding Probability Inequalities for Sums of Bounded Random Variables Journal of the American Statistical Association, Vol. 58, No. 301. (Mar., 1963), pp. 13-30.

[K96] F. Kelly Notes on effective bandwidths. In: *Stochastic networks: Theory and applications.* Oxford University Press, 1996.

[S99] D. Starobinski, Quality of service in high speed networks with multiple time-scale traffic, Ph. D. dissertation, Technion-Israel Inst. Technol., Haifa, Israel, May 1999.

[S00] D. Starobinski and M. Sidi, Stochastically Bounded Burstiness for Communication Networks, IEEE Trans. On Information Theory, 2000.

[V03] Bounds for independent regulated inputs multiplexed in a service curve network element, M. Vojnovic and J.-Y. Le Boudec, IEEE Transactions on Communications, 51(5):735 - 740, May 2003.

[Y03] 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.