

A Network Service Curve Approach for the Stochastic Analysis of Networks*

Florin Ciucu
Dept. of Computer Science
University of Virginia
Charlottesville, VA 22904
florin@cs.virginia.edu

Almut Burchard
Department of Mathematics
University of Virginia
Charlottesville, VA 22904
burchard@virginia.edu

Jörg Liebeherr
Dept. of Computer Science
University of Virginia
Charlottesville, VA 22904
jorg@cs.virginia.edu

ABSTRACT

The stochastic network calculus is an evolving new methodology for backlog and delay analysis of networks that can account for statistical multiplexing gain. This paper advances the stochastic network calculus by deriving a network service curve, which expresses the service given to a flow by the network as a whole in terms of a probabilistic bound. The presented network service curve permits the calculation of statistical end-to-end delay and backlog bounds for broad classes of arrival and service distributions. The benefits of the derived service curve are illustrated for the exponentially bounded burstiness (EBB) traffic model. It is shown that end-to-end performance measures computed with a network service curve are bounded by $\mathcal{O}(H \log H)$, where H is the number of nodes traversed by a flow. Using currently available techniques that compute end-to-end bounds by adding single node results, the corresponding performance measures are bounded by $\mathcal{O}(H^3)$.

Categories and Subject Descriptors

C.4 [Computer System Organization]: Performance of Systems; G.3 [Probability and Statistics]: Stochastic Processes

General Terms

Performance

Keywords

Stochastic network calculus, Quality-of-Service, network service curve

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

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGMETRICS'05, June 6–10, 2005, Banff, Alberta, Canada.
Copyright 2005 ACM 1-59593-022-1/05/0006 ...\$5.00.

1. INTRODUCTION

The network calculus is a framework for analyzing delays and backlog in a network where the traffic, and sometimes also the service, are characterized in terms of envelope functions. Pioneered as deterministic network calculus in the early 1990s for the computation of worst-case performance bounds in packet networks [9], it has played an important role in the development of algorithms that support Quality-of-Service guarantees in packet networks.

The elegance of the network calculus becomes evident in the *min-plus algebra* formulation developed in [1, 8, 14], where service guarantees to a flow at a node (switch) are expressed in terms of *service curves* [10, 19]. In this formulation, bounds for single nodes can be easily extended to end-to-end bounds. More concretely, suppose a flow is assigned a service curve S^h at the h -th node on its route ($h = 1, \dots, H$). Then the service given to the flow by the network as a whole can be expressed in terms of a network service curve S^{net} as

$$S^{net} = S^1 * S^2 * \dots * S^H, \quad (1)$$

where $*$ is a convolution operator. With this remarkable property, bounds for the output burstiness, backlog and delay for the entire network are computed in the same fashion as single node results. The resulting end-to-end delay bounds are generally tighter than the sum of the per-node delay bounds. For example, if the service curve at the h -th node is given as a constant rate function, $S^h(\tau) = c\tau$, one obtains from Eq. (1) that $S^{net}(\tau) = c\tau$. As a result, the end-to-end backlog and delay bounds are identical to the bounds at the first node. In this way, the min-plus version of the network calculus provides simple end-to-end estimates for delay and backlog.

A drawback of the worst-case view of traffic in the deterministic network calculus is that it does not reap the benefits of statistical multiplexing, which can result in an overestimation of the actual resource requirements and a low utilization of network resources. This has motivated the search for a *stochastic network calculus* which describes arrivals and service probabilistically while preserving the elegance and expressiveness of the original framework. By allowing even a small fraction of traffic to violate its traffic description or performance guarantees, one can achieve significant resource savings.

Most work on extending the network calculus to a probabilistic setting has been concerned with deriving statistical performance bounds for a single node. In a stochastic net-

work calculus framework, traffic arrivals and sometimes also service at network nodes are random processes which are bounded by probabilistic envelope functions. The first, and probably most widely known envelope function is the exponentially bounded burstiness (EBB) characterization for traffic arrivals [23]. The EBB model, which has been generalized in [3, 6, 22, 24], has been shown to imply delay and backlog bounds at simple traffic multiplexers. In [4, 20], probabilistic arrival envelopes were used to derive schedulability conditions for a variety of scheduling algorithms. The authors of [6, 16] have established a link between envelope functions and the theory of effective bandwidth [12], which estimates bandwidth requirements to satisfy given performance guarantees. Probabilistic envelope functions that specify the amount of service made available to a flow at a network node have appeared in [11, 16, 20].

A number of studies have used probabilistic single node bounds on delay, backlog, or the burstiness of traffic departing from a node to derive multi-node performance bounds, e.g., [6, 13, 22, 23]. Indeed, by relating output descriptions of traffic at a node to corresponding input descriptions, one can obtain end-to-end bounds by adding the per-node bounds. However, such results tend to degrade rapidly as the number of nodes traversed by a flow is increased.

The promise of a min-plus algebra formulation of the stochastic network calculus lies in the development of probabilistic network service curves that yield multi-node performance bounds without the poor scaling properties associated with adding per-node results. However, the development of such a calculus has shown to be difficult. One particular challenge is the formulation of the multi-node convolution expression of a network service curve within a probabilistic context. In [5] it was shown that a straightforward probabilistic extension of deterministic concepts yields a network service curve that deteriorates with time. As a solution, the authors of [5] formulated a probabilistic service curve that takes the form of Eq. (1), however, this service curve is difficult to apply in numerical examples. In [2], a probabilistic network service curve was derived under the assumption that each node drops traffic that locally violates a given delay guarantee. This dropping policy requires that packets in each buffer are sorted according to a deadline computed from the arrivals envelope and the service curve. Another network service curve expression, developed in [16], assumes that a bound is available which limits the busy period over multiple nodes. Such a bound, however, is difficult to obtain.

Until now, the derivation of a network service curve for a stochastic network calculus that does not depend on implicit or explicit a priori upper limits on delay, backlog, or busy periods has remained an open problem. The main contribution of this paper is the construction of a probabilistic network service curve that does not rely on a priori bounds. We show that the network service curve derived here is applicable to a broad class of arrival and service distributions. The presented probabilistic service curve formulation is based on a definition introduced in [11]. A recent study [18] also attempts to construct a network service curve with a service curve as defined in [11], however, the presented network service curve is not correct.¹

¹The paper fails to account for the fact that the convolution of probabilistic arrival functions with service curves over multiple nodes requires a sample path view. This and other pitfalls in the

We illustrate the benefits of the network service curve for a statistical end-to-end analysis of multiplexed EBB traffic. By contrasting end-to-end delay bounds obtained with our service curve with bounds obtained by iterating single node results, as proposed in [23], we show the improvements attainable through our stochastic network calculus approach. We will show that the calculus approach with network service curve renders bounds for delay, backlog, and output burstiness of the order $\mathcal{O}(H \log H)$ in the number of nodes H in the network, as opposed to $\mathcal{O}(H^3)$ bounds obtained by adding per-node results. Thus, this paper, for the first time, quantifies the benefits of using network service curves in a probabilistic setting. This presents a significant step forward towards the goal of developing the stochastic network calculus into a practical methodology for the analysis of networks. As a remark, in the deterministic calculus, a network service curve leads to end-to-end bounds that scale with $\mathcal{O}(H)$, while summing up single-node results gives bounds that scale with $\mathcal{O}(H^2)$ [15]. Thus, network service curves have comparable benefits in a deterministic and a stochastic setting.

The remainder of the paper is organized as follows. In Section 2, we define our notion of a statistical service curve, and present our main result, i.e., a service curve that expresses the service received by a flow in a network. In Section 3 we use envelope functions for a probabilistic characterization of traffic and, together with our notion of probabilistic service, obtain performance bounds on backlog, delay and output burstiness. In Section 4 we discuss an application of our network service curve to EBB traffic, and compare the resulting end-to-end delay bounds with the bounds implied by existing single node results. In Section 5 we present a numerical example with Markovian On-Off traffic for illustration. We present brief conclusions in Section 6.

2. NETWORK SERVICE CURVES

The input and output traffic of a flow at a network node is described by two stochastic processes $A = (A(t))_{t \geq 0}$ and $D = (D(t))_{t \geq 0}$ that are defined on a joint probability space. $A(t)$ represents the cumulative arrivals, and $D(t)$ the cumulative departures in $[0, t)$. We require that A and D are non-decreasing, left continuous functions with $A(0) = D(0) = 0$, and that $D(t) \leq A(t)$ for any time $t \geq 0$. In this paper we use a continuous-time framework. Extensions to a discrete-time setting are discussed in remarks.

In a packet-switching network, the service available to a flow at a node is determined by a scheduling algorithm (e.g., FIFO, Fair Queueing) which sets the order of packet transmissions. A service curve, first presented in [10, 19], is an alternate method to describe the service received by a flow in terms of a function which specifies a lower bound on the service. In the min-plus algebra formulation of the deterministic calculus, a service curve is a function $S(\cdot)$, such that $D(t) \geq A * S(t)$ for all $t \geq 0$, where the *convolution* of two real-valued functions f and g is defined as $f * g(t) = \inf_{0 \leq s \leq t} \{f(s) + g(t - s)\}$ for all $t \geq 0$.

Next we define our measure of a probabilistic service guarantee for a flow. We adopt a variation of the definition of a statistical service curve from [11], where we add a positivity requirement. We use the notation $[x]_+ = \max(x, 0)$ for the positive part of a real number x .

stochastic network calculus are discussed in [16].

DEFINITION 1. (STATISTICAL SERVICE CURVE). A function $\mathcal{S}(t)$ is a statistical service curve for an arrival process A if for every choice of σ and for all $t \geq 0$

$$P\left(D(t) < A * [\mathcal{S} - \sigma]_+(t)\right) \leq \varepsilon(\sigma), \quad (2)$$

where $\varepsilon(\sigma)$ is a non-increasing function.

We refer to the bound $\varepsilon(\sigma)$ on the violation probability as the *error function*. Note that the condition is void whenever $\varepsilon(\sigma) \geq 1$, and that Eq. (2) for $t = 0$ implies that $\varepsilon(\sigma) \geq 1$ for all $\sigma < \mathcal{S}(0)$. We frequently require that the function $\varepsilon(\sigma)$ satisfies the integrability condition

$$\int_0^\infty \varepsilon(u) du < \infty. \quad (3)$$

Comparing Definition 1 to probabilistic service descriptions in the literature, we see that for each choice of σ , the function $[\mathcal{S} - \sigma]_+$ is an effective service curve in the sense of [5, 16, 17]. Here, choosing σ large amounts to increasing the latency and decreasing the violation probability of the service guarantee. Compared to the service curve in [11], the enforced positivity of the statistical service curve can lead to tighter performance bounds on backlog and output burstiness. Lastly, if $\varepsilon(\sigma) = 0$ for some value of σ , then $[\mathcal{S} - \sigma]_+$ defines a deterministic service curve almost surely.

In the continuous-time setting, we find it convenient to replace Eq. (2) by the stronger requirement that the service curve satisfies

$$P\left(D(t) < A * [\mathcal{S} - \sigma]_+(t + \tau_0)\right) \leq \varepsilon(\sigma) \quad (4)$$

for all $t \geq 0$, and that $\varepsilon(\sigma) \geq 1$ for all $\sigma < \mathcal{S}(\tau_0)$. Here, $\tau_0 > 0$ is a parameter that specifies a discretization of the time scale. If $\mathcal{S}(t)$ is a statistical service curve in the sense of Eq. (2), then Eq. (4) holds for $\mathcal{S}(t - \tau_0)$.

We now state the main result of this paper. Consider a flow with a network path through $H > 1$ nodes, as shown in Figure 1. At each node, we assume that the flow receives a probabilistic service guarantee in terms of a statistical service curve. The following theorem provides an expression for an end-to-end statistical network service curve in terms of the per-node service curves. In the theorem, we use the notation $f_\delta(t) = f(t) + \delta t$ for a real function f and a real number δ .

THEOREM 1. (STATISTICAL NETWORK SERVICE CURVE). Consider a flow with traffic through H nodes. Let $A^{net} := A^1$ denote the arrivals to the first node, $A^{h+1}(t) = D^h(t)$ for $h = 1, \dots, H-1$ the arrivals at the $(h+1)$ -th node, and $D^{net} := D^H$ the departures from the last node. Assume that each node $h = 1, \dots, H$ provides a service guarantee in form of a statistical service curve \mathcal{S}^h satisfying Eq. (4) for some $\tau_0 > 0$. Assume that for all $h < H$, the error functions ε^h satisfy the integrability condition in Eq. (3). Then, for every choice of $\delta > 0$, the function

$$\mathcal{S}^{net} = \mathcal{S}^1 * \mathcal{S}_{-\delta}^2 * \dots * \mathcal{S}_{-(H-1)\delta}^H \quad (5)$$

is a network service curve in the sense of Eq. (4), with an error function given by

$$\varepsilon^{net}(\sigma) = \inf_{\sigma^1 + \dots + \sigma^H = \sigma} \left\{ \varepsilon^H(\sigma^H) + \sum_{h=1}^{H-1} \frac{1}{\delta \tau_0} \int_{\sigma^h}^\infty \varepsilon^h(u) du \right\}.$$




Figure 1: Traffic of a flow through a set of H nodes.

The formula for the network service curve in Eq. (5) corresponds closely to the formulation of the deterministic calculus given in Eq. (1). The difference is that in Eq. (5), the service curve at the h -th node is reduced by a rate $(h-1)\delta$. The convolution expression in the deterministic calculus from Eq. (5) is recovered almost surely by setting $\varepsilon(\sigma) = 0$ for all σ and then taking $\delta, \tau_0, \sigma \rightarrow 0$.

Applications of the theorem require to choose the parameters δ and τ_0 . Large values of these parameters lead to pessimistic performance bounds that hold with small violation probability, while smaller values increase the violation probability. In addition, the parameter σ is chosen so that the value of the error $\varepsilon^{net}(\sigma)$ does not exceed a desired violation probability, e.g., 10^{-6} or 10^{-9} .

The proof of the theorem relies on sample path arguments inspired by the deterministic calculus. To estimate the probability of departure events involving entire sample paths, we need to transform Definition 1, which makes statements about departures $D(t)$ at any given time t , into a sample path expression. This crucial step is provided in the next lemma. While it is known how to obtain related sample path bounds for traffic arrivals (as discussed in the next section), this is the first time that such a sample path bound is established for service descriptions.

LEMMA 1. Let \mathcal{S} be a statistical service curve for an arrival process A . Assume that \mathcal{S} satisfies Eq. (4) for some $\tau_0 > 0$, and that $\varepsilon(\sigma)$ satisfies Eq. (3). Then, for any real number $\delta > 0$ and all $t, \sigma \geq 0$,

$$P\left(\sup_{0 \leq s \leq t} \left\{ A * [\mathcal{S} - \delta(t + \tau_0 - s) - \sigma]_+(s) - D(s) \right\} > 0\right) \leq \frac{1}{\delta \tau_0} \int_\sigma^\infty \varepsilon(u) du. \quad (6)$$

We have slightly abused notation in order to simplify the statement of the lemma. The convolution term on the left hand side of Eq. (6) should be read as

$$\inf_{0 \leq u \leq s} \left\{ A(u) + [\mathcal{S}(s-u) - \delta(t + \tau_0 - s) - \sigma]_+ \right\}. \quad (7)$$

PROOF. Given $t > 0$, we discretize the event in Eq. (6) on a time scale τ_0 . Let $0 \leq s \leq t$, and let $j = \lfloor \frac{t-s}{\tau_0} \rfloor$ be the integer part of $\frac{t-s}{\tau_0}$. Since A , \mathcal{S} , and D are nondecreasing, we have

$$A * [\mathcal{S} - \delta(t + \tau_0 - s) - \sigma]_+(s) \leq A * [\mathcal{S} - \delta(j+1)\tau_0 - \sigma]_+(t - j\tau_0)$$

and

$$D(s) \geq D([t - (j+1)\tau_0]_+).$$

It follows that

$$\begin{aligned}
& P\left(\sup_{0 \leq s \leq t} \left\{ \left(A * [\mathcal{S} - \delta(t + \tau_0 - s) - \sigma]_+ \right) (s) - D(s) \right\} > 0 \right) \\
& \leq P\left(\max_{j=0, \dots, \lfloor \frac{t}{\tau_0} \rfloor} \left\{ A * [\mathcal{S} - \delta(j+1)\tau_0 - \sigma]_+(t - j\tau_0) \right. \right. \\
& \quad \left. \left. - D([t - (j+1)\tau_0]_+) \right\} > 0 \right) \\
& \leq \sum_{j=0}^{\lfloor \frac{t}{\tau_0} \rfloor} P\left(A * [\mathcal{S} - \delta(j+1)\tau_0 - \sigma]_+(t - j\tau_0) \right. \\
& \quad \left. > D([t - (j+1)\tau_0]_+) \right),
\end{aligned}$$

where we have applied Boole's inequality in the last step. By Eq. (4), the last sum is bounded by

$$\sum_{j=0}^{\infty} \varepsilon(\sigma + (j+1)\delta\tau_0) \leq \frac{1}{\delta\tau_0} \int_{\sigma}^{\infty} \varepsilon(u) du. \quad \square$$

PROOF OF THEOREM 1. Let $t > 0$, $\delta > 0$ and $\sigma > 0$ be given, and choose $\sigma^1, \dots, \sigma^H$ with $\sum \sigma^h = \sigma$. Suppose that for a particular sample path, the inequalities

$$D^h(s) \geq A^h * \left[\mathcal{S}^h - \delta(t + \tau_0 - s) - \sigma^h \right]_+(s) \quad (8)$$

hold for all $0 \leq s \leq t$ and $h = 1, \dots, H-1$, and that

$$D^H(t) \geq A^H * \left[\mathcal{S}^H - \sigma^H \right]_+(t + \tau_0). \quad (9)$$

Inserting the bound for $A^H = D^{H-1}$ from Eq. (8) into Eq. (9) and expanding the convolutions yields with Eq. (7)

$$\begin{aligned}
D^{net}(t) & \geq \inf_{0 \leq u \leq s \leq t} \left\{ A^{H-1}(u) + \left[\mathcal{S}^{H-1}(s-u) \right. \right. \\
& \quad \left. \left. - \delta(t + \tau_0 - s) - \sigma^{H-1} \right]_+ + \left[\mathcal{S}^H(t + \tau_0 - s) - \sigma^H \right]_+ \right\}.
\end{aligned}$$

Note that our definition of \mathcal{S} justifies the restriction of the infimum to $s \leq t$. Collecting terms and using that $[x]_+ + [y]_+ \geq [x+y]_+$, we obtain

$$D^{net}(t) \geq A^{H-1} * \left[\mathcal{S}^{H-1} * \mathcal{S}_{-\delta}^H - (\sigma^{H-1} + \sigma^H) \right]_+(t + \tau_0).$$

Iterating the argument for the remaining values of h and using the definition of \mathcal{S}^{net} , we arrive at

$$D^{net}(t) \geq A^{net} * \left[\mathcal{S}^{net} - \sum_{h=1}^H \sigma^h \right]_+(t + \tau_0).$$

Using Lemma 1 for $h = 1, \dots, H-1$ and Eq. (4) for $h = H$, we bound the violation probability by

$$\begin{aligned}
& P\left(D^{net}(t) < A^{net} * \left[\mathcal{S}^{net} - \sum_{h=1}^H \sigma^h \right]_+(t + \tau_0)\right) \\
& \leq \sum_{h=1}^{H-1} P(\text{Eq. (8) fails for } \mathcal{S}^h \text{ for some } s \leq t) \\
& \quad + P(\text{Eq. (9) fails}) \\
& \leq \sum_{h=1}^{H-1} \frac{1}{\delta\tau_0} \int_{\sigma^h}^{\infty} \varepsilon^h(u) du + \varepsilon^H(\sigma^H).
\end{aligned}$$

The proof is completed by minimizing over $\sigma^1, \dots, \sigma^H$. \square

In a discrete time setting, there is no need for the parameter τ_0 appearing in Eq. (4), and we use the definition of the statistical service curve in Eq. (2). If each node provides a service curve \mathcal{S}^h in the sense of Eq. (2) with an error function ε^σ , then for any choice of $\delta > 0$

$$\mathcal{S}^{net} = \mathcal{S}^1 * \mathcal{S}_{-\delta}^2 * \dots * \mathcal{S}_{-(H-1)\delta}^H$$

is a network service curve which again satisfies Eq. (2), with error function

$$\varepsilon^{net}(\sigma) = \inf_{\sigma^1 + \dots + \sigma^H = \sigma} \left\{ \varepsilon^H(\sigma^H) + \sum_{h=1}^{H-1} \sum_{k=0}^{\infty} \varepsilon^h(\sigma^h + k\delta) \right\}.$$

3. PERFORMANCE BOUNDS

The derivation of the statistical network service curve in Theorem 1 does not make assumptions on the arrival functions A at a node, and holds for all deterministic or probabilistic descriptions of traffic. However, applying the network service curve to compute performance bounds for a traffic flow requires a characterization for the arrivals from the flow. In the deterministic network calculus, it is generally assumed that the arrivals from a traffic flow A are bounded or regulated by an *arrival envelope* A^* , such that $A(t + \tau) - A(t) \leq A^*(\tau)$ for all $t, \tau \geq 0$. A frequently used envelope is $A^*(t) = \sigma + \rho t$, which corresponds to a leaky bucket with rate ρ and burst size σ . In a stochastic network calculus, traffic arrivals are usually described in terms of probabilistic extensions [4, 22, 23] of this envelope concept. The following definition specifies such an extension.

DEFINITION 2. (STATISTICAL ENVELOPE). A *nondecreasing function* $\mathcal{G}(t)$ is said to be a *statistical envelope* for an arrival process A if, for all σ and all $0 \leq s \leq t$

$$P\left(A(t) - A(s) > \mathcal{G}(t-s) + \sigma\right) \leq \varepsilon(\sigma), \quad (10)$$

where $\varepsilon(\sigma)$ is a nonnegative, nonincreasing function.

Note that Eq. (10) is formulated for negative as well as positive values of σ . If Eq. (10) is only known to hold for non-negative values of σ , we set $\varepsilon(\sigma) = \max\{1, \varepsilon(0)\}$ for $\sigma < 0$.

Definition 2 is inspired by the Exponentially Bounded Burstiness (EBB) model [23], which is the special case where $\mathcal{G}(t) = \rho t$ is a constant-rate function and the error function decays exponentially. The Stochastically Bounded Burstiness (SBB) model [22] is a generalization of the EBB model where the condition that the error function decays exponentially is relaxed and replaced by the assumption that it decays faster than polynomially. These decay conditions are used to bound the violation probability of events involving entire arrival sample paths.

An alternate approach, introduced in [24], requires that the arrivals satisfy an a priori sample path bound, leading to the generalized Stochastic Bounded Burstiness (gSBB) model. It is shown in [24] that an arrival process which conforms to the SBB traffic model is also gSBB, though with more pessimistic bounds. More generally, gSBB can accommodate any arrival model satisfying a tail estimate analogous to Eq. (3). The corresponding result in our paper is provided by Lemma 2. Several recent studies [3, 18, 24] have established performance bounds for the gSBB arrival model with both deterministic and statistical service descriptions. At a single node these bounds take an especially appealing, concise form. However, in order to bound delay and

backlog along a path through a network, these results need to be combined with a network service curve, which is not provided in [3, 24] and incorrect in [18].

Let us comment briefly on the scope of the various traffic models. The EBB model includes important classes such as multiplexed regulated traffic, Markovian On-Off traffic, and the Poisson process, but does not include distributions with heavy tails or long-range correlations. In particular, fractional Brownian motion traffic belongs to SBB but not EBB, satisfying Eq. (10) with $\mathcal{G}(t) = (\rho + \delta)t$ and $\varepsilon(\sigma) \sim e^{-\sigma^\alpha}$, where ρ is the average traffic rate, $\delta > 0$ is arbitrarily small, and $0 < \alpha < 1$ is related to the Hurst parameter. The gSBB model was shown to contain an even richer category of traffic patterns [24], in particular, heavy-tailed distributions where the error function decays as a power law $\varepsilon(\sigma) \sim \sigma^{-\alpha}$ with $1 < \alpha \leq 2$.

Definition 2 allows statistical envelopes to be given by arbitrary nondecreasing functions, which can provide a tighter description of arrivals than the linear envelopes used in the EBB, SBB, and gSBB traffic models. Still, linear envelopes provide an important class of examples, for two reasons. First, to imply any bounds on the delay and backlog at a link which provides a constant service rate, a statistical envelope cannot grow faster than linearly. Second, linear envelopes lend themselves to efficient computations.

Definition 2 is also closely related to other probabilistic envelope characterizations of traffic used in the literature. In [11], an arrival process satisfying Eq. (10) is called \mathcal{G} -smooth with overflow profile ε . For each value of σ , the function $\mathcal{G} + \sigma$ provides an effective envelope, as defined in [4], with violation probability $\varepsilon(\sigma)$. Finally, as shown in [6, 16], statistical envelopes can be connected with the notion of *effective bandwidth* [12].

We next describe statistical performance guarantees on the backlog $B(t) = A(t) - D(t)$ and the delay $W(t) = \inf\{d \geq 0 : A(t) \leq D(t + d)\}$, where arrivals are bounded by statistical envelopes satisfying Definition 2, and service is expressed in terms of statistical service curves satisfying Definition 1. We also derive a bound on output burstiness in the form of a statistical envelope for $D(t)$ satisfying Definition 2. We point out that similar bounds have been derived in other statistical network calculus papers [2, 5, 16, 18] for a variety of arrival and service characterizations. We state the following results for completeness, since there are technical differences between our definition of a statistical service curve in Eqs. (2) and (4) and those used in the literature. The theorem uses the *deconvolution* operator \oslash , which is defined for two real functions f and g as $f \oslash g(t) = \sup_{s \geq 0} \{f(t + s) - g(s)\}$ for all $t \geq 0$.

THEOREM 2. (PERFORMANCE BOUNDS). *Let $A(t)$ and $D(t)$ denote the arrivals and departure processes at a node which provides a service curve $\mathcal{S}(t)$ satisfying Eq. (4) with some $\tau_0 > 0$ and an error function $\varepsilon^s(\sigma)$. Let the arrivals be bounded by a statistical envelope \mathcal{G} with an error function $\varepsilon^g(\sigma)$ that satisfies the integrability condition from Eq. (3). Fix $\delta, \tau_0 > 0$, and define*

$$\varepsilon(\sigma) = \inf_{\sigma^g + \sigma^s = \sigma} \left\{ \varepsilon^s(\sigma^s) + \frac{1}{\delta\tau_0} \int_{\sigma^g}^{\infty} \varepsilon^g(u) du \right\}.$$

Then we have the following bounds:

1. OUTPUT BURSTINESS: $\mathcal{G} \oslash \mathcal{S}_{-\delta}$ provides a statistical

envelope for D , i.e.,

$$P(D(t) - D(s) > \mathcal{G} \oslash \mathcal{S}_{-\delta}(t - s) + \sigma) \leq \varepsilon(\sigma) \quad (11)$$

for all $0 \leq s \leq t$.

2. BACKLOG BOUND: A statistical bound on the backlog at a node is given, for $t \geq 0$, by

$$P(B(t) > \mathcal{G}_\delta \oslash \mathcal{S}(0) + \sigma) \leq \varepsilon(\sigma) . \quad (12)$$

3. DELAY BOUND: A statistical bound on the delay is given, for $t \geq 0$, by

$$P(W(t) > d(\sigma)) \leq \varepsilon(\sigma) , \quad (13)$$

where

$$d(\sigma) = \inf \{d : \mathcal{S}(s + d) \geq \mathcal{G}_\delta(s) + \sigma \text{ for all } s \geq 0\} . \quad (14)$$

In the theorem, the rate correction parameter δ may appear either in the statistical envelope or in the statistical service curve. In the above bounds, we have placed δ wherever it gives the better performance bound.

The proof of the theorem relies on a sample path argument analogous to the proof of Theorem 1. This is provided by the following lemma.

LEMMA 2. *Assume that \mathcal{G} is a statistical envelope for A , with an error function $\varepsilon(\sigma)$ satisfying Eq. (3). Then, for every choice of $\delta, \tau_0 > 0$,*

$$\begin{aligned} P\left(A(t) > \inf_{0 \leq u \leq s} \left\{ A(u) + \mathcal{G}(t + \tau_0 - u) + \delta(s + \tau_0 - u) + \sigma \right\}\right) \\ \leq \frac{1}{\delta\tau_0} \int_{\sigma}^{\infty} \varepsilon(u) du \end{aligned} \quad (15)$$

for all $0 \leq s \leq t$ and $\sigma \geq 0$.

An implication of Lemma 2 is that an arrival process which has a linear statistical envelope $\mathcal{G}(t) = \rho t$ and an error function satisfying Eq. (3) is of class gSBB, as in [24].

PROOF. Fix $s \leq t$. As in Lemma 1, we discretize the event in Eq. (15) on a time scale τ_0 by setting $j = \lfloor \frac{s-u}{\tau_0} \rfloor$, for a given u with $0 \leq u \leq s \leq t$. By the monotonicity of A and \mathcal{G} , we have

$$A(u) \geq A(\lfloor s - (j+1)\tau_0 \rfloor_+)$$

and

$$\begin{aligned} \mathcal{G}(t + \tau_0 - u) + \delta(s + \tau_0 - u) \\ \geq \mathcal{G}(t - \lfloor s - (j+1)\tau_0 \rfloor_+) + \delta(j+1)\tau_0 . \end{aligned}$$

We estimate

$$\begin{aligned} P\left(A(t) > \inf_{0 \leq u \leq s} \left\{ A(u) + \mathcal{G}(t + \tau_0 - u) + \delta(s + \tau_0 - u) + \sigma \right\}\right) \\ \leq P\left(A(t) > \min_{j=0, \dots, \lfloor \frac{s}{\tau_0} \rfloor} \left\{ A(\lfloor s - (j+1)\tau_0 \rfloor_+) \right. \right. \\ \left. \left. + \mathcal{G}(t - \lfloor s - (j+1)\tau_0 \rfloor_+) + \delta(j+1)\tau_0 + \sigma \right\}\right) \\ \leq \sum_{j=0}^{\lfloor \frac{s}{\tau_0} \rfloor} P\left(A(t) - A(\lfloor s - (j+1)\tau_0 \rfloor_+) \right. \\ \left. > \mathcal{G}(t - \lfloor s - (j+1)\tau_0 \rfloor_+) + \delta(j+1)\tau_0 + \sigma \right) , \end{aligned}$$

where we have used Boole's inequality in the last step. Since \mathcal{G} is a statistical envelope for A by assumption and $\varepsilon(\sigma)$ is nondecreasing, we can bound the last sum by

$$\sum_{j=0}^{\infty} \varepsilon(\sigma + (j+1)\delta\tau_0) \leq \frac{1}{\delta\tau_0} \int_{\sigma}^{\infty} \varepsilon(u) du . \quad \square$$

PROOF OF THEOREM 2. The argument follows the corresponding derivations in the deterministic calculus (see, e.g., [1, 7, 14]). For given $0 \leq s \leq t$, assume that for a particular sample path,

$$D(s) \geq A * [\mathcal{S} - \sigma^s]_+(s + \tau_0) , \quad (16)$$

with $\mathcal{S}(\tau_0) \leq \sigma^s$, and

$$A(t) - A(u) \leq \mathcal{G}(t + \tau_0 - u) + \delta(s + \tau_0 - u) + \sigma^g \quad (17)$$

for all $0 \leq u \leq s$. Then

$$\begin{aligned} D(t) - D(s) &\leq A(t) - \inf_{0 \leq u \leq s + \tau_0} \left\{ A(u) + [\mathcal{S}(s + \tau_0 - u) - \sigma^s]_+ \right\} \\ &\leq \sup_{0 \leq u \leq s} \left\{ \mathcal{G}(t + \tau_0 - u) + \delta(s + \tau_0 - u) + \sigma^g \right. \\ &\quad \left. - [\mathcal{S}(s + \tau_0 - u) - \sigma^s]_+ \right\} \\ &\leq \mathcal{G} \circ \mathcal{S}_{-\delta}(t - s) + \sigma . \end{aligned}$$

In the first step, we have used that $D(t) \leq A(t)$ and the assumption in Eq. (17). Note that values $u > s$ do not contribute to the infimum since $\mathcal{S}(\tau_0) \leq \sigma^s$. In the second step, we have taken $A(t)$ into the infimum and used Eq. (16), and in the last step we have collected terms and applied the definition of the deconvolution.

Using Lemma 2 and the assumption that \mathcal{S} is a statistical service curve, we conclude that

$$\begin{aligned} P\left(D(t) - D(s) > \mathcal{G} \circ \mathcal{S}_{-\delta}(t - s) + \sigma\right) \\ \leq \varepsilon^s(\sigma^s) + \frac{1}{\delta\tau_0} \int_{\sigma^g}^{\infty} \varepsilon^g(u) du . \end{aligned} \quad (18)$$

The proofs of the delay and backlog bounds proceed along the same lines and are omitted. \square

The proof actually shows that the output burstiness satisfies the tighter bound

$$\begin{aligned} P\left(D(t) - D(s) > \mathcal{G} \circ ([\mathcal{S} - \sigma^s]_+)_-\delta(t - s) + \sigma^g\right) \\ \leq \varepsilon(\sigma^g, \sigma^s) , \end{aligned}$$

where $\varepsilon^s(\sigma^g, \sigma^s)$ is given by the right hand side of Eq. (18). This is a consequence of the positivity property of the statistical service curve in Definition 1. A similar bound can be proven for the backlog,

$$P(B(t) > \mathcal{G}_{\delta} \circ [\mathcal{S} - \sigma^s]_+(0) + \sigma^g) \leq \varepsilon(\sigma^g, \sigma^s) ,$$

but not for the delay. The improvement over Theorem 2 can be noticeable at low utilizations but disappears at high utilizations.

Given a workconserving scheduler at an output link that operates at a constant rate, we can describe the service available to a tagged flow at a scheduler serving multiple flows in terms of a service curve that expresses the capacity left unused by other flows with traffic at this link. We refer to

such a service curve as a *leftover service curve*. This concept has been applied in a deterministic setting in [8, 15], and in a probabilistic setting in [16, 17, 20]. We point out that the leftover service curve is a pessimistic estimate of the service available to the flow, which amounts to serving the tagged flow at a lower priority than all other flows. For particular scheduling algorithms, less pessimistic estimates have been obtained in [16, 20].

The following theorem shows that we can obtain a leftover service curve compliant with Definition 1, whenever traffic arrivals are characterized by statistical envelopes in the sense of Definition 2. We will use the leftover service curve in Section 4 to determine the service available to a single flow at a multiplexer with EBB traffic arrivals.

THEOREM 3. (LEFTOVER SERVICE CURVE). *Consider a workconserving scheduler with constant rate C serving multiple flows. Fix a collection of flows and denote its aggregate arrival and departure processes by $A(t)$ and $D(t)$. Let $A_c(t)$ and $D_c(t)$ denote the arrivals and departures of the aggregate of the remaining flows. Assume that \mathcal{G}_c is a statistical envelope for the arrivals A_c that satisfies*

$$\limsup_{t \rightarrow \infty} \frac{\mathcal{G}_c(t)}{t} < C ,$$

with an error function $\varepsilon_c^g(\sigma)$ that satisfies the integrability condition in Eq. (3). Then, for any choice of $\delta > 0$ and $\tau_0 > 0$,

$$\mathcal{S}(t) = C(t - \tau_0) - \mathcal{G}_{c,\delta}(t)$$

is a statistical service curve for A that satisfies Eq. (4) with error function

$$\varepsilon^s(\sigma) = \frac{1}{\delta\tau_0} \int_{\sigma}^{\infty} \varepsilon_c^g(u) du .$$

If the stronger integrability condition

$$\int_0^{\infty} \sigma \varepsilon_c^g(\sigma) d\sigma < \infty \quad (19)$$

holds, then $\varepsilon^s(\sigma)$ satisfies Eq. (3).

In the theorem, we used the notation $\mathcal{G}_{c,\delta}(t) = \mathcal{G}_c(t) + \delta t$.

PROOF. Since the link provides a service curve $R(t) = Ct$ to the aggregate of through flows (described by A) and cross flows (described by A_c), and $D_c(t) \leq A_c * R(t)$, we have

$$\begin{aligned} D(t) &\geq (A + A_c) * R(t) - A_c * R(t) \\ &= \inf_{0 \leq s \leq t} \left\{ A(s) + A_c(s) + R(t - s) - A_c * R(t) \right\} . \end{aligned}$$

Inserting the inequality

$$A_c * R(t) \leq \min\{A_c(t), A_c(s) + R(t - s)\}$$

and collecting terms, we obtain

$$D(t) \geq \inf_{0 \leq s \leq t} \left\{ A(s) + [A_c(s) - A_c(t) + R(t - s)]_+ \right\} . \quad (20)$$

Fix $t \geq 0$ and $\delta, \tau_0 > 0$. Suppose that, for a particular sample path, we have

$$A_c(t) \leq \inf_{0 \leq s \leq t} \left\{ A_c(s) + \mathcal{G}_{c,\delta}(t + \tau_0 - s) + \sigma \right\} . \quad (21)$$

Using this inequality in Eq. (20) yields

$$\begin{aligned} D(t) &\geq \inf_{0 \leq s \leq t} \left\{ A(s) + [R(t - s) - \mathcal{G}_{c,\delta}(t + \tau_0 - s) - \sigma]_+ \right\} \\ &\geq A * [\mathcal{S} - \sigma]_+(t + \tau_0) . \end{aligned}$$

By Lemma 2 with $s = t$, the violation probability is bounded by

$$\begin{aligned} P\left(D(t) < A * [S - \sigma]_+(t + \tau_0)\right) &= P(\text{Eq. (21) fails}) \\ &\leq \frac{1}{\delta\tau_0} \int_{\sigma}^{\infty} \varepsilon_c^g(u) du. \end{aligned}$$

If $\sigma < \mathcal{S}(\tau_0)$, then

$$\varepsilon^s(\sigma) \geq \frac{1}{\delta\tau_0} \int_{-\mathcal{G}_c(0) - \delta\tau_0}^{\infty} \varepsilon_c^g(u) du \geq \inf_{\mathcal{G}_c(0) + u < 0} \varepsilon_c^g(u) \geq 1,$$

as required for Eq. (4). We have used that the definition of the statistical envelope in Eq. (10) implies, by setting $s = t$, that $\varepsilon_c^g(\sigma) \geq 1$ for $\mathcal{G}_c(0) + \sigma < 0$.

If the error function ε_c^g satisfies Eq. (19), then

$$\begin{aligned} \int_0^{\infty} \varepsilon^s(\sigma) d\sigma &= \frac{1}{\delta\tau_0} \int_0^{\infty} \int_{\sigma}^{\infty} \varepsilon_c^g(u) du d\sigma \\ &= \frac{1}{\delta\tau_0} \int_0^{\infty} u \varepsilon_c^g(u) du < \infty \end{aligned}$$

by exchanging the order of integration. \square

Finally, we note that a leftover service curve and performance bounds can also be obtained in a discrete time setting. As in the case of Theorem 1, the only modification needed in Theorems 2 and 3 is that the statistical service curve needs to satisfy Eq. (2) in place of Eq. (4), and that discrete sums $\sum_{k=0}^{\infty} \varepsilon(\sigma + k\delta)$ appear in place of the integrals in the error functions.

4. AN APPLICATION WITH EBB ARRIVALS

We now present an application that relates the network service curve and the performance bounds developed in this paper to the literature on statistical service guarantees from the early 1990s. We demonstrate that statistical network service curves can faithfully reproduce the single node results in [23] which predate the service curve concept. In a multi-node setting, we show the benefits of the statistical network service curve by comparing statistical end-to-end performance bounds computed with the techniques from [23] (without a network service curve) to those obtained with the results of this paper.

The network scenario that we consider is shown in Figure 2. We will refer to the flows which traverse the network as the *through* flows, and the flows which transit the network as the *cross* flows. We are interested in statistical multi-node performance measures for the through flows, such as an output envelope at the last node in the network and bounds on the total delay experienced along the path through the network. Our performance bounds hold for all work-conserving scheduling algorithms that serve traffic from the same flow in the order of arrival.

Network arrivals are described in terms of the exponentially bounded burstiness (EBB) model defined in [23] which is given for the arrival process A by the condition that

$$P(A(t) - A(s) > \rho(t - s) + \sigma) \leq M e^{-\theta\sigma} \quad (22)$$

for any $0 \leq s \leq t$. Here, ρ represents a bound on the long-term arrival rate, θ is the decay rate of the error function, and $M \geq 1$ is a constant.

By choosing the EBB traffic model we can compare our stochastic network calculus results with existing performance




Figure 2: A network with cross traffic.

bounds, specifically [23]. Moreover, the EBB model can be used to describe arrival processes that have relevance in practice, e.g., Markov-modulated On-Off processes and multiplexed regulated arrivals. Finally, and most importantly, the EBB model lends itself to simple, closed-form performance bounds that permit us to gain insight into the scaling properties of network-wide bounds obtained with the stochastic network calculus.

4.1 Analysis at a single node

We will show that the network calculus developed in this paper can recover the results of the EBB analysis for a single node from [23]. This demonstrates that statistical service curves are an adequate representation of the service received by a flow with respect to cross traffic. Reproducing the EBB analysis with a service curve approach may help to dispel a myth that service curves necessarily lead to inaccurate descriptions of service.

Consider the first node in the network from Figure 2. We assume that the aggregate arrivals of the through flows, denoted by A , and the aggregate arrivals of the cross flows, denoted by A_c , are each EBB arrival processes with bounds ρ and ρ_c on their long-term rates. For simplicity, both arrival processes are assumed to have the same parameters $\theta > 0$ and $M \geq 1$. Following [23], this subsection uses a discrete time domain. In Subsection 4.2 we will return to a continuous time domain, as it results in simpler expressions for the desired bounds.

It was shown in ([23], Theorem 2) that the output traffic D has the EBB characterization

$$P(D(t) - D(s) > \rho(t - s) + \sigma) \leq \frac{2M}{1 - e^{-\frac{\theta}{2}(C - (\rho + \rho_c))}} e^{-\frac{\theta}{2}\sigma} \quad (23)$$

for all $0 \leq s \leq t$.

The same output bound can be obtained in the network calculus with a statistical service curve. The EBB characterization of the cross traffic says that $\mathcal{G}_c(t) = \rho_c t$ is a statistical envelope in the sense of Definition 2, with error function $\varepsilon_c^g(\sigma) = M e^{-\theta\sigma}$. To characterize the service available to the through flows, we use a leftover service curve. As explained at the end of Section 3, the discrete-time version of Theorem 3 implies that for any choice of $\delta > 0$,

$$\mathcal{S}(t) = (C - \rho_c - \delta)t$$

is a statistical service curve for A in the sense of Eq. (2), with error function

$$\varepsilon^s(\sigma) = \sum_{k=0}^{\infty} M e^{-\theta(\sigma + k\delta)} = \frac{M}{1 - e^{-\theta\delta}} e^{-\theta\sigma}.$$

Next we use that $\mathcal{G}(t) = \rho t$ is a statistical envelope for the through flows. The discrete-time version of Theorem 2 implies that, for $\delta \leq \frac{C-(\rho+\rho_c)}{2}$, the function

$$\mathcal{G} \oslash \mathcal{S}_{-\delta}(t-s) = \rho(t-s)$$

is a statistical envelope for the departure process of the through flows, with error function

$$\varepsilon(\sigma) = \frac{2M}{1-e^{-\theta\delta}} e^{-\frac{\theta}{2}\sigma}.$$

By setting $\delta = \frac{C-(\rho+\rho_c)}{2}$, we obtain Eq. (23). The backlog bound of [23] can be similarly recovered, and the discrete counterpart of Theorem 2 provides a stronger delay bound than the one suggested in [23].

This example demonstrates that the separation of the statistical analysis of the through flows (in terms of a statistical envelope) from those made for the service of the cross flows (in terms of a statistical service curve) does not lead to more pessimistic bounds than a direct analysis of the backlog process.

4.2 Analysis of end-to-end delay bounds

We turn to the problem of deriving network performance bounds for the through flows in the network from Figure 2. The purpose of this section is to compare two strategies for computing statistical end-to-end bounds. The technique suggested in [23] is to analyze bounds at each node iteratively, using the EBB characterization of the output in Eq. (23). Network performance bounds are obtained by adding the per-node delay bounds. This approach has also been used in related works that compute statistical end-to-end delay bounds for other traffic models, e.g., [6, 13, 22, 23]. Network service curves offer an alternate technique of reducing the analysis of a network to the analysis of a single node. In the deterministic network calculus under worst-case traffic scenarios it is known that network service curves render better bounds than adding per-node delay bounds. This observation is appropriately called ‘pay-bursts-only-once’ in [15]. The results in this section establish the corresponding result in the probabilistic EBB context. We note that the statistical network service curves derived in [2, 5] share this property. However, this is the first time that the benefit of the ‘pay-bursts-only-once’ property in a stochastic context is established by a direct comparison with results from the literature that rely on summing up per-node bounds. (The assumptions on the service curves in [2] and [5] make such a comparison difficult). Our results will show that adding per-node delay bounds results in a network delay bound that increases with $\mathcal{O}(H^3)$ in the number of nodes, whereas the delay bounds with network service curve grow with $\mathcal{O}(H \log H)$.

All derivations in this subsection, including those of the per-node service bounds, are done in the framework of the network calculus, that is, we use statistical envelopes to describe arrivals and service curves to describe service. This is justified by the result of the previous subsection that our network calculus can reproduce the input-output relation and performance bounds obtained by a direct analysis of the backlog process at a single node.

As shown in Figure 2, the network has H workconserving nodes arranged in series, with the capacity of each node set to $C > 0$. We use A^h and D^h to denote the arrivals and departures of the through flows at the h -th node, with

$A^{h+1} = D^h$ for $h = 1, \dots, H-1$. As before, we assume that the through and cross flows have EBB characterizations with the same values for M and θ and long-term bounds on their rates given by ρ and ρ_c .

At each node, the service available to the through flows is given by a leftover service curve. Since the cross flows at each node h are described in terms of the EBB model, it follows that

$$\mathcal{G}_c(t) = \rho_c t$$

is a statistical envelope for the cross flows with error function

$$\varepsilon_c(\sigma) = M e^{-\theta\sigma}.$$

By Theorem 3, for any choice of $\delta > 0, \tau_0 \geq 0$, the function

$$\mathcal{S}^h(t) = (C - \rho_c - \delta)t - C\tau_0$$

is a statistical service curve for the through flows at node h . The service curve can be formulated more elegantly by replacing σ with $\sigma - C\tau_0$, yielding

$$\mathcal{S}^h(t) = (C - \rho_c - \delta)t. \quad (24)$$

The resulting error function is

$$\varepsilon^{s,h}(\sigma) = \frac{e^{\theta C\tau_0}}{\theta\delta\tau_0} e^{-\theta\sigma}.$$

4.2.1 Adding per-node bounds

Here, we calculate delay bounds according to Theorem 2 for each node, and then add the per-node delay bounds to obtain network bounds. The service available to the through flows at each node is given by Eq. (24). To apply the theorem we need to also have a statistical envelope of the arrivals at each node. We will show by induction that the arrivals A^h at each node are characterized by a statistical envelope

$$\mathcal{G}^h(t) = \rho t \quad (25)$$

in the sense of Definition 2, with error function

$$\varepsilon^{g,h}(\sigma) = M \left(\frac{C\theta}{\delta} \right)^{\frac{h^2+h-2}{2h}} e^{-\frac{\theta}{h}\sigma}, \quad (26)$$

where $\delta \leq \frac{C-(\rho+\rho_c)}{2}$ is a free parameter. In other words, the arrivals at each node comply with the EBB traffic model, with identical rates at each node, but with a reduced decay rate and an increased constant at subsequent nodes ($h > 1$).

For $h = 1$, there is nothing to show, since the arrivals of the through flows at the first node are characterized by the EBB traffic model from Eq. (22). To make the inductive step, we will use the following lemma.

LEMMA 3. *For any positive numbers M_k, θ_k ($k = 1, \dots, K$) and any $\sigma \geq 0$,*

$$\inf_{\sigma_1 + \dots + \sigma_K = \sigma} \sum_{k=1}^K M_k e^{-\theta_k \sigma_k} = \prod_{k=1}^K (M_k \theta_k w)^{\frac{1}{\theta_k w}} e^{-\frac{\sigma}{w}},$$

where $w = \sum_{k=1}^K \frac{1}{\theta_k}$.

A proof of the lemma is given in the Appendix.

Assume that for some $h \geq 1$, we already know that a statistical envelope for the arrivals at node h is given by Eqs. (25) and (26). Since the service of the through flows

at the h -th node is described by the service curve \mathcal{S}^h in Eq. (24), Theorem 2.1 implies that

$$\mathcal{G}^{h+1}(t) = \mathcal{G}^h \circ \mathcal{S}_{-\delta}^h(t) = \rho(t-s)$$

is a statistical envelope for the output from the h -th node, with error function

$$\varepsilon^{h+1}(\sigma) = \inf_{\sigma_1 + \sigma_2 = \sigma} \frac{M}{\theta \delta \tau_0} \left\{ \left(\frac{C\varepsilon}{\delta} \right)^{\frac{h^2+h-2}{2h}} h e^{-\frac{\theta}{h}\sigma_1} + e^{\theta C \tau_0} e^{-\theta \sigma_2} \right\}.$$

Applying Lemma 3 with $K = 2$, $\theta_1 = \frac{\theta}{h}$ and $\theta_2 = \theta$, and then minimizing over τ_0 by setting $\tau_0 = \frac{h+1}{C\theta}$, we conclude that

$$\varepsilon^{h+1}(\sigma) = M \left(\frac{C\varepsilon}{\delta} \right)^{\frac{h^2+3h}{2(h+1)}} e^{-\frac{\theta}{h+1}\sigma}, \quad (27)$$

which completes the induction.

We next obtain a statistical bound for the delay $W^h(t)$ at node h from Theorem 2.3. Set

$$\begin{aligned} d(\sigma) &= \inf \{ d : \mathcal{S}(s+d) \geq \mathcal{G}(s) + \sigma \} \\ &= \frac{\sigma}{C - \rho_c - \delta} \end{aligned}$$

according to Eq. (14). Then we get from Theorem 2.3 that

$$P(W^h(t) > d(\sigma)) \leq \varepsilon^{h+1}(\sigma),$$

where the error function ε^{h+1} is given by Eq. (27). Replacing σ by $(C - \rho_c - \delta)d$ gives

$$P(W^h(t) > d) \leq \varepsilon^{h+1}((C - \rho_c - \delta)\sigma). \quad (28)$$

To obtain a bound for the network delay W^{net} , we add the delay bounds at each node,

$$P(W^{net}(t) > d) \leq \inf_{d_1 + \dots + d_H = d} \sum_{h=1}^H P(W^h(t) > d_h).$$

Inserting the bounds from Eq. (28) and using Lemma 3 to take the infimum, we get

$$P(W^{net}(t) > d) \leq M^{net} e^{-\frac{2\theta}{H(H+3)}(C - \rho_c - \delta)d}, \quad (29)$$

with

$$M^{net} = \frac{H(H+3)}{2} M \left(\frac{C\varepsilon}{\delta} \right)^{\frac{(H+1)(H+5)}{3(H+3)}} \prod_{h=1}^H (h+1)^{-\frac{2(h+1)}{H(H+3)}}. \quad (30)$$

The value of the free parameter δ that minimizes the violation probability in Eq. (29) under the constraint $0 < \delta \leq \frac{C - (\rho + \rho_c)}{2}$ is given by

$$\delta = \min \left\{ \frac{H(H+1)(H+5)}{6\theta d}, \frac{C - (\rho + \rho_c)}{2} \right\}. \quad (31)$$

At this point, we can compute the probability that a given delay bound is exceeded. In our numerical examples, we determine a delay bound d so that $P(W(t) > d) \leq \varepsilon$, where ε is a given violation probability. Setting the right hand side of Eq. (29) to ε and solving for d gives

$$d = (C - \rho_c - \delta) \frac{H(H+3)}{2\theta} \log \left(\frac{M^{net}}{\varepsilon} \right). \quad (32)$$

Since the optimal value for δ depends on d by Eq. (31), we first set $\delta_0 = \frac{C - (\rho + \rho_c)}{2}$ and compute the corresponding delay

d_0 from Eq. (32). We then determine δ from Eq. (31) with d_0 in place of d , and finally obtain the desired delay bound from Eq. (32).

From Eq. (30) it follows that $\log(M^{net})$ grows linearly in the number of nodes H . Thus, with the quadratic term in Eq. (32), we conclude that the delay bound in Eq. (32) scales as $O(H^3)$.

4.2.2 Using the network service curve

We now derive network delay bounds for W^{net} for the same scenario using the statistical network service curve from Section 2.

Fix $\delta \leq \frac{C - (\rho + \rho_c)}{H+1}$. Starting from the leftover service curves in Eq. (24), we apply Theorem 1 to obtain the statistical network service curve

$$\mathcal{S}^{net}(t) = (C - \rho_c - H\delta)t,$$

with error function

$$\varepsilon^{s,net}(\sigma) = M \frac{H e^{\theta C \tau_0}}{(\theta \delta \tau_0)^{\frac{2H-1}{H}}} e^{-\frac{\theta}{H}\sigma},$$

where we have again used Lemma 3 to compute the infimum appearing in the formula for ε^s from Theorem 1.

Recall the statistical envelope formulation of the through flows at the first node, which is given by $\mathcal{G}(t) = \rho t$, with error function $\varepsilon^g(\sigma) = M e^{-\theta \sigma}$. Having an arrival characterization and a network service curve, we can apply Theorem 2.3 to obtain a bound on the network delay. From Eq. (14), we find that

$$d(\sigma) = \frac{\sigma}{C - \rho_c - H\delta}.$$

It follows from Theorem 2 with σ replaced by $(C - \rho_c - H\delta)d$ that

$$P(W^{net}(t) > d) \leq M^{net} e^{-\frac{\theta}{H+1}(C - \rho_c - H\delta)d} \quad (33)$$

with

$$M^{net} = M e(H+1) \left(\frac{HC}{(H+1)\delta} \right)^{\frac{2H}{H+1}}, \quad (34)$$

where we have once more used Lemma 3 and optimized over the parameter τ_0 in the formula for the error function in Theorem 2. The minimizing value of the free parameter δ for a given value of d is given by

$$\delta = \min \left\{ \frac{2}{\theta d}, \frac{C - (\rho + \rho_c)}{H+1} \right\}. \quad (35)$$

Solving for the delay bound in Eq. (33) with the right hand side set equal to ε , we get

$$d = (C - \rho_c - H\delta) \frac{H+1}{\theta} \log \left(\frac{M^{net}}{\varepsilon} \right). \quad (36)$$

In our numerical example, we use the same two-step procedure as in Subsection 4.2.1 to determine δ . We compute d_0 from Eq. (36) with δ set equal to $\delta_0 = \frac{C - (\rho + \rho_c)}{H+1}$, then insert d_0 for d into Eq. (35) for an improved value for δ , which we use to determine the delay bound d from Eq. (36). It is apparent from Eq. (34) that M^{net} grows polynomially in the number of nodes, and hence the delay bound in Eq. (36) is of order $O(H \log H)$.

The same techniques can be used to obtain statistical bounds on the network backlog. Adding per-node bounds

for the backlog yields

$$P(B^{net}(t) > \sigma) \leq M^{net} e^{-\frac{2\theta}{H(H+3)}\sigma}, \quad (37)$$

where M^{net} is given by Eq. (30). Here, the optimal choice is $\delta = (C - (\rho + \rho_c))/2$. Using the network service curve yields

$$P(B^{net}(t) > \sigma) \leq M^{net} e^{-\frac{\theta}{H+1}\sigma},$$

with M^{net} given by Eq. (34). In this case, the optimal choice for δ is

$$\delta = \frac{C - (\rho + \rho_c)}{H + 1}.$$

With this choice, M^{net} evaluates to

$$M^{net} = Me(H + 1) \left(\frac{HC}{C - (\rho + \rho_c)} \right)^{\frac{2H}{H+1}}.$$

So, a backlog bound b can be explicitly calculated as follows:

$$\begin{aligned} b &\leq \frac{H + 1}{\theta} \log \frac{M^{net}}{\varepsilon} \\ &\leq \frac{2H}{\theta} \log \left(\frac{HC}{C - (\rho + \rho_c)} \right) + \frac{H + 1}{\theta} \log \left(\frac{Me(H + 1)}{\varepsilon} \right). \end{aligned}$$

5. NUMERICAL EXAMPLE

We next give a numerical example that illustrates the benefits of using network service curves for the computation of statistical end-to-end delay bounds. We consider the network and arrival scenario shown in Figure 2. Arrivals of cross flows and arrivals of through flows at the first node are each described as an aggregate of independent Markov Modulated On-Off processes. This type of process, which has been used for modeling voice channels, falls into the category of the EBB traffic model. We will show plots that compare the delay bounds obtained through adding per-node bounds (from Subsection 4.2.1) to those obtainable with a network service curve (from Subsection 4.2.2).




Figure 3: On-Off Traffic model.

The Markov Modulated On-Off arrival process of an arrival flow, illustrated in Figure 3, is a continuous time process with support given by a homogeneous two-state Markov chain $X(t)$ which is described in terms of the generator matrix

$$G = \begin{pmatrix} -\mu & \mu \\ \lambda & -\lambda \end{pmatrix}.$$

Here, μ denotes the transition rate from the ‘On’ state to the ‘Off’ and λ to denote the transition rate from the ‘Off’ state to the ‘On’ state. In the ‘On’ state, the arrival process transmits at the peak rate P , and no arrivals occur in the ‘Off’ state.

We assume that there are N through flows and N_c cross flows at each node. Through flows at the first node and

cross flows are stochastically independent. For the sake of simplicity, we assume that all arrival processes are homogeneous and that there is an equal number of through and cross flows ($N = N_c$). We want to emphasize that computing examples for heterogeneous flows does not pose a problem, other than increasing notation.

Next, following [6], we quickly derive an EBB characterization of N independent On-Off flows. The moment-generating function of a single On-Off flow is bounded by $E[e^{\theta A(t)}] \leq e^{\theta t \rho(\theta)}$, where

$$\rho(\theta) = \frac{1}{2\theta} \left(P\theta - \mu - \lambda + ((P\theta - \mu + \lambda)^2 + 4\lambda\mu)^{\frac{1}{2}} \right),$$

(see [12]). The quantity $\rho(\theta)$ is called ‘effective capacity’ in [21] and has the property that $\rho(0) \leq \rho(\theta) \leq P$. The rate $\rho(0) = \frac{\lambda}{\lambda + \mu} P$ represents the average rate of the flow.

For N independent flows we can bound the moment generating function by $e^{N\theta t \rho(\theta)}$. By the Chernoff bound, we can write

$$\begin{aligned} P(A(t) - A(s) > N\rho(\theta)(t - s) + \sigma) \\ \leq e^{-\theta(N\rho(\theta)(t-s) + \sigma)} E[e^{\theta A(t-s)}] \leq e^{-\theta\sigma}. \end{aligned}$$


For each choice of $\theta > 0$, this provides us with an EBB characterization for the through flows and the cross flows in the network from Figure 2. Delay bounds for the through flows are provided by Eqs. (29) and (33) from the previous section. Finally, we numerically optimize the resulting delay bounds over θ .

In the example, the capacity of each node in the network is set to $C = 100$ Mbps, and time is measured in milliseconds. The parameters of the flows are given in Table 1. We consider two types of flows, with identical peak rate ($P = 1.5$ Mbps) and average rate ($\rho(0) = 0.15$ Mbps). We introduce a parameter $T = \frac{1}{\mu} + \frac{1}{\lambda}$ to describe the burstiness of a flow. T is the expected time for the Markov chain to change states twice. For flows with given peak rate P and given mean rate $\rho(0)$, a larger value of T indicates a higher degree of burstiness. In Table 1, we use $T = 10$ for flows with low burstiness and $T = 100$ for flows with high burstiness. Lastly, the violation probability for the end-to-end delay bounds is set to $\varepsilon = 10^{-9}$.


Burstiness	T (ms)	P (Mbps)	$\rho(0)$ (Mbps)	λ (ms ⁻¹)	μ (ms ⁻¹)
low	10	1.5	0.15	0.11	1.0
high	100	1.5	0.15	0.01	0.1

Table 1: Parameters of On-Off sources.


In Figure 4 we show the probabilistic end-to-end delay bounds of flows as a function of the number of flows $N + N_c$. We consider networks where the number of nodes traversed by the through flows is set to $H = 1, 2, 5$, and 10. The maximum number of flows at each node is given by $\lfloor \frac{C}{\rho} \rfloor = 666$ flows. In Figures 4(a) and (b) we show the delay bounds obtained by adding per-node bounds from Subsection 4.2.1 and by using the network service curves from Subsection 4.2.2, respectively. For a single node ($H = 1$), both techniques yield the same delay bounds. The benefits of network service curves become pronounced when the number of nodes H traversed by the through flows is increased. Figures 4(c) and (d) show similar plots for flows that are more bursty.




(a) $T=10$ (Adding Per-Node Bounds).



(b) $T=10$ (Network Service Curve).



(c) $T=100$ (Adding Per-Node Bounds).



(d) $T=100$ (Network Service Curve).

Figure 4: End-to-end delay bounds of Markov Modulated On-Off arrivals as a function of the number of flows $N + N_c$ ($H = 1, 2, 5, 10$, $\varepsilon = 10^{-9}$, $T = 10\text{ms}$ (low burstiness), $T = 100\text{ms}$ (high burstiness), $N = N_c$).

Here, the delay bounds are higher, underlining that bursty flows have a lower statistical multiplexing gain.

In Figure 5 we present end-to-end delay bounds for the same setting as in Figure 4. The delay bounds are represented as a function of the number of nodes H . We set the number of flows so that the node utilization level is set to $U = 10\%$, 50% , and 90% of the capacity, where $N = N_c$. We only consider flows where $T = 10$ ms. In Figure 5, we depict the delay bounds obtained by adding per-node bounds as dashed lines, and the bounds obtained with network service curves as solid lines. The figure illustrates the $\mathcal{O}(H \log H)$ bounds of end-to-end delays with network service curves, and the polynomial $\mathcal{O}(H^3)$ bounds seen when adding per-node results. When the number of nodes H grows large, the scaling properties dominate even the effects of the traffic load. For $H > 50$, the bounds obtained with network service curves on a highly (90%) loaded link is smaller than the delay bounds obtained without network service curves at a lightly (10%) loaded link.

6. CONCLUSIONS

We have extended the state-of-the-art of the stochastic network calculus by deriving a network service curve formulation that is applicable to a broad class of traffic and service characterizations. The formulation of such a service curve in the presented general form has been a long-standing research problem. Using the network service curve, we calculated statistical end-to-end delay and backlog bounds which are vastly superior to bounds obtained by adding delay bounds of single nodes. For EBB traffic arrivals, we showed that for a flow that traverses H nodes and encounters cross traffic at each node, our network calculus with statistical network service curves gives statistical end-to-end delays that are bounded by $\mathcal{O}(H \log H)$, as opposed to $\mathcal{O}(H^3)$ bounds rendered by the method of adding per-node bounds. An immediate research problem suggested by this paper relates to the tightness of the $\mathcal{O}(H \log H)$ bounds. Another extension is a systematic investigation that studies the impact of traffic and service assumptions on scaling properties of network bounds.




Figure 5: End-to-end delay bounds for fixed network utilization levels of $U = 10\%$, 50% , and 90% as a function of the number of nodes H in the network ($T = 10$ ms, $\varepsilon = 10^{-9}$, $N = N_c$).

7. 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, 1999.
- [2] S. Ayyorgun and R. Cruz. A service-curve model with loss and a multiplexing problem. In *24th International Conference on Distributed Computing System (ICDCS'04)*, pages 756–765, March 2004.
- [3] S. Ayyorgun and W. Feng. A probabilistic definition of burstiness characterization. Technical Report LA-UR 03-3668, Los Alamos National Laboratory, May 2003.
- [4] R. Boorstyn, A. Burchard, J. Liebeherr, and C. Oottamakorn. Statistical service assurances for traffic scheduling algorithms. *IEEE Journal on Selected Areas in Communications. Special Issue on Internet QoS*, 18(12):2651–2664, 2000.
- [5] A. Burchard, J. Liebeherr, and S. D. Patek. A calculus for end-to-end statistical service guarantees. Technical Report CS-2001-19, University of Virginia, Computer Science Department, May 2002.
- [6] C. S. Chang. Stability, queue length, and delay of deterministic and stochastic queueing networks. *IEEE Transactions on Automatic Control*, 39(5):913–931, 1994.
- [7] 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, 1998.
- [8] C. S. Chang. *Performance Guarantees in Communication Networks*. Springer Verlag, 2000.
- [9] R. Cruz. A calculus for network delay, parts I and II. *IEEE Transactions on Information Theory*, 37(1):114–141, 1991.
- [10] R. L. Cruz. Quality of service guarantees in virtual circuit switched networks. *IEEE Journal on Selected Areas in Communications*, 13(6):1048–1056, 1995.
- [11] R. L. Cruz. Quality of service management in integrated services networks. In *Proceedings of the 1st Semi-Annual Research Review, CWC*, UCSD, June 1996.
- [12] F. Kelly. Notes on effective bandwidths. In *Stochastic Networks: Theory and Applications*. (Editors: F.P. Kelly, S. Zachary and I.B. Ziedins) Royal Statistical Society Lecture Notes Series, 4, pages 141–168. Oxford University Press, 1996.
- [13] J. Kurose. On computing per-session performance bounds in high-speed multi-hop computer networks. In *Proceedings of ACM Sigmetrics'92*, pages 128–139, June 1992.

- [14] J. Y. Le Boudec. Application of network calculus to guaranteed service networks. *IEEE/ACM Transactions on Information Theory*, 44(3):1087–1097, 1998.
- [15] J. Y. Le Boudec and P. Thiran. *Network Calculus*. Springer Verlag, Lecture Notes in Computer Science, LNCS 2050, 2001.
- [16] C. Li, A. Burchard, and J. Liebeherr. A network calculus with effective bandwidth. Technical Report CS-2003-20, University of Virginia, Computer Science Department, Nov. 2003.
- [17] J. Liebeherr, S. Patek, and A. Burchard. Statistical per-flow service bounds in a network with aggregate provisioning. In *Proceedings of IEEE Infocom 2003*, March 2003.
- [18] Y. Liu, C. Tham, and Y. Jiang. A stochastic network calculus. Technical report, ECE-CCN-0301, Dept. of Electrical and Computer Engineering, National University of Singapore, Nov. 2003.
- [19] A. K. Parekh and R. G. Gallager. A generalized processor sharing approach to flow control in integrated services networks: the single-node case. *IEEE/ACM Transactions on Networking*, 1(3):344–357, 1993.
- [20] J. Qiu and E. Knightly. Inter-class resource sharing using statistical service envelopes. In *Proceedings of IEEE Infocom '99*, pages 36–42, March 1999.
- [21] M. Schwartz. *Telecommunication Network: Protocols, Modeling, and Analysis*. Addison-Wesley, 1986.
- [22] D. Starobinski and M. Sidi. Stochastically bounded burstiness for communication networks. *IEEE Transaction of Information Theory*, 46(1):206–212, 2000.
- [23] O. Yaron and M. Sidi. Performance and stability of communication networks via robust exponential bounds. *IEEE/ACM Transactions on Networking*, 1(3):372–385, 1993.
- [24] Q. Yin, Y. Jiang, S. Jiang, and P. Y. Kong. Analysis on generalized stochastically bounded bursty traffic for communication networks. In *Proc. IEEE Local Computer Networks 2002*, pages 141–149, November 2002.

Appendix: Proof of Lemma 3

We need to minimize

$$f(\sigma_1, \dots, \sigma_K) = \sum_{k=1}^K M_k e^{-\theta_k \sigma_k} \quad (38)$$

subject to the constraint that $g(\sigma_1, \dots, \sigma_K) = \sigma_1 + \dots + \sigma_K = \sigma$, where $\sigma > 0$ is given. This problem has a unique solution by the strict convexity of f . The Lagrange multiplier condition for the minimum takes the form

$$M_k \theta_k e^{-\theta_k \sigma_k} = \lambda, \quad k = 1, \dots, K, \quad (39)$$

where the Lagrange multiplier λ is determined by the constraint. Eq. (39) allows to express the minimizing values $\sigma_1, \dots, \sigma_K$ in terms of λ . Inserting these values into Eq. (38) shows that the minimum of f is given by

$$f(\sigma_1, \dots, \sigma_K) = \sum_{k=1}^K \frac{\lambda}{\theta_k}. \quad (40)$$

Let w be as in the statement of the lemma and set $p_k = 1/(\theta_k w)$. Since $p_1 + \dots + p_K = 1$, Eq. (39) implies

$$\lambda = \prod_{k=1}^K \lambda^{p_k} = \prod_{k=1}^K \left(\frac{M_k}{w p_k} \right)^{p_k} e^{-\sigma/w}, \quad (41)$$

where we have used the constraint and the definition of w and p_k in the last line. Inserting Eq. (41) into Eq. (40) completes the proof.