A Min-Plus Calculus for End-to-end Statistical Service Guarantees

Almut Burchard, Jörg Liebeherr, *Senior Member, IEEE*, Stephen Patek, *Member, IEEE*

***Abstract*— The network calculus offers an elegant framework for determining worst-case bounds on delay and backlog in a network. This paper extends the network calculus to a probabilistic framework with statistical service guarantees. The notion of a *statistical service curve* is presented as a probabilistic bound on the service received by an individual flow or an aggregate of flows. The problem of concatenating per-node statistical service curves to form an end-to-end (network) statistical service curve is explored. Two solution approaches are presented that can each yield statistical network service curves. The first approach requires the availability of time scale bounds at which arrivals and departures at each node are correlated. The second approach considers a service curve that describes service over time intervals. Although the latter description of service is less general, it is argued that many practically relevant service curves may be compliant to this description.**

***Index Terms*— Stochastic network calculus, Quality-of-Service, network service curve.**

I. INTRODUCTION

Beginning with Cruz's seminal work [14] the deterministic network calculus has evolved into a comprehensive theory for the worst-case performance analysis of packet networks. It has provided tools for reasoning about delay and backlog in a network with service guarantees to individual or aggregate traffic flows. A strength of the network calculus is that it can be used to determine delays and backlog across multiple network nodes and that it can accurately describe the behavior of a broad class of scheduling algorithms. Using the notion of arrival envelopes and service curves [15], delay and backlog bounds can be concisely expressed in a min-plus algebra [1], [7], [11]. The most powerful characteristic of the network calculus is that bounds for single nodes can be easily extended to multi-node bounds using a convolution operation of the min-plus algebra [5].

The deterministic calculus reflects worst-case scenarios where arrivals and service in the network conspire to create maximum delays, backlogs, and bursts. Since traffic is statistically multiplexed at network nodes, such scenarios are extremely rare. Therefore, a deterministic approach to provisioning service generally overestimates the actual resource needs, resulting in a low utilization of network resources. A probabilistic view which considers that traffic in a packet network is statistically multiplexed increases the achievable utilization in the network by tolerating rare adversarial events. This is referred to as *statistical multiplexing gain*. The statistical network calculus seeks to quantify the statistical multiplexing gain while maintaining the algebraic aspects of the deterministic calculus.

Research on statistical network calculus has the potential of providing new insights into fundamental trade-offs of packet network architectures. For example, in [29] it is argued that in networks with high data rates, statistical multiplexing dominates the effects of scheduling. If it is indeed true that the selection of scheduling algorithms has little impact on the delay performance of networks, then the need for complex link scheduling algorithms in Internet routers should be re-evaluated. A statistical network calculus could also justify recent empirical evidence that buffer sizes in Internet routers are overprovisioned [2]. Another potential application area is the verification of service-level agreements (SLAs) between network customers and service providers. One can use service curves to describe statistical lower bounds for the service experienced by a single flow when resources are managed for aggregates of flows by expressing the service seen by the flow in terms of the capacity not used by other flows [30]. This concept can be applied for verifying an SLA as follows: If a network customer measures the aggregate input to the service provider and the throughput of even a single flow, the customer can determine whether the service provider has allocated the capacity specified in the agreement. Finally, probabilistic descriptions of service are suitable to express the time-varying conditions of a wireless channel, which are subject to random losses, noise, and interference between users.

The contribution of this paper is a statistical network calculus in the min-plus algebra formulation with convolution and deconvolution operators [5]. The potential of using the min-plus algebra is that, as in the deterministic context, end-to-end guarantees can be expressed as a simple concatenation (convolution) of single-node guarantees, leading to end-to-end bounds for latency and backlog that are much smaller than the sum of the corresponding single-node bounds. We introduce the concept of a *statistical service curve* as a probabilistic bound on the service received by a single flow or group of flows at a node.[1] We will show that single-node

The research in this paper is supported in part by the National Science Foundation under grants ECS-9875688 (CAREER), ANI-0085955, CNS-0435061, and DMS-0308040, by an Alfred P. Sloan research fellowship, and two NSERC Discovery grants.

Almut Burchard (almut@math.toronto.edu) is with the Department of Mathematics at the University of Toronto. Jörg Liebeherr (jorg@comm.utoronto.ca) is with the Department of Electrical and Computer Engineering at the University of Toronto. Stephen Patek (patek@virginia.edu) is with the Department of Systems and Information Engineering at the University of Virginia.

[1] We use the terms 'statistical', 'probabilistic', and 'stochastic' interchangeably throughout the paper.

performance bounds on output burstiness, backlog and delay of the deterministic network calculus carry over to a probabilistic framework based on statistical service curves.

At the same time, the derivation of end-to-end statistical service curves turns out to be significantly more difficult, and requires the availability of time scale bounds at which arrivals and departures at each node are correlated. We present two solutions to deal with this requirement. In the first solution we assume that *a priori* time scale limits are available. Such limits can be computed by obtaining bounds on the busy period or the maximum backlog at a node. In the second solution, we modify the characterization of service curves in such a way that they imply the required bounds.

Much of the literature on statistical service guarantees investigates performance bounds at a single node, assuming bounds on the distribution of arrivals, such as exponentially bounded burstiness [38], linear envelope processes [10], stochastically bounded burstiness [35], effective bandwidth characterizations [23], general burstiness characterizations [4], [13], [39], stochastic domination by a given random variable [27], or regulation by worst-case arrival envelopes [24]. Under the additional assumption that arrivals from multiple flows are stochastically independent, the statistical multiplexing gain can be captured by applying the Central Limit Theorem [25], the Chernoff Bound [19], or the Hoeffding Bound [37]. Some works have iterated probabilistic bounds to yield end-to-end bounds [35], [38]. However, end-to-end bounds obtained by adding single-node results degrade quickly with the number of nodes. We refer to [21], [26], [33] for reviews of the literature on statistical multiplexing in packet networks.

The concept of probabilistic service curves was first suggested by Cruz in [16] but no network service curve was derived there. Chang suggested a network service curve for a time-variable service description in the problem set of ([12], Chapter 7), where the service at different nodes is described by independent random processes and traffic is characterized by moment-generating functions. After the manuscript for this paper became available in 2002 [9], several studies have appeared on the network calculus, some extending the framework developed in this paper, e.g., [3], [13], [20], [22], [28], [30]. The statistical network calculus has also been related to other analytical techniques. For example, in [28] it has been shown that the effective bandwidth theory [23] can be expressed in the statistical network calculus. This was used to derive end-to-end performance bounds for networks with self-similar arrival traffic.

The remaining sections of this paper are structured as follows. In Section II, we review the notation and main results of the network calculus. In Section III we introduce statistical service curves and present the results for a statistical network calculus in terms of statistical service curves. The section concludes with a discussion of key difficulties in the statistical network calculus. In Section IV we present a solution for end-to-end service curves that exploit time scale limits. In Section V we provide a different end-to-end service curve, with a revised description of probabilistic service. In Section VI we draw brief conclusions.

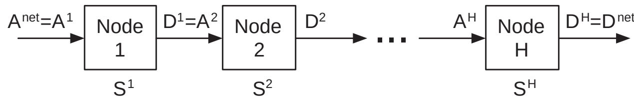

Fig. 1. Traffic of a flow through a set of H nodes. Let A^h and D^h denote the arrival and departures at the h-th node, with $A^{net} = A^1$, $D^{h-1} = A^h$ for $h = 2, \ldots, H$ and $D^H = D^{net}$.

II. Network Calculus Preliminaries

In this section we review the notation and main results of the deterministic network calculus [8], [12]. Throughout this paper we use a continuous time model with left-continuous traffic arrival functions. In our framework, a node represents a packet switch in a network. Packetization delays and other effects of discrete-sized packets, such as the non-preemption of packet transmission, are ignored. We also do not account for processing overhead and propagation delays. We refer to [12] for issues involved in relaxing these assumptions.

Let us consider the traffic of a single flow at a network node. The arrivals to and departures from the node in the time interval $[0, t)$ are denoted by non-negative, non-decreasing, left-continuous functions $A(t)$ and $D(t)$ with $D(t) \leq A(t)$ and $A(0) = D(0) = 0$. The backlog of a flow at time t, denoted by $B(t)$, is given by $B(t) = A(t) - D(t)$. The delay at time t, denoted as $W(t)$, is the delay experienced by an arrival which departs at time t, given by $W(t) = \inf\{d \geq 0 \mid A(t-d) \leq D(t)\}$. When A and D are represented as curves, $B(t)$ and $W(t)$ are the vertical and horizontal differences between the curves.

In the min-plus algebra formulation of the network calculus, developed by [1], [7], [12], the service guarantees to a flow at a node are expressed in terms of *service curves*. The algebra uses the *convolution operator* $*$ and *deconvolution operator* $\oslash$, which are defined by setting, for any two functions f and g,

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

A service curve for a flow is a function S that specifies a lower bound on the service given to the flow such that, for all $t \geq 0$,

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

A frequently encountered class of service curves is the rate-latency server $S(t) = c[t-d]_+$ [36], which appears in the Guaranteed Service specification for the Internet [34].

When arrivals are bounded by an *arrival envelope* A^*, such that $A(t+\tau) - A(t) \leq A^*(\tau)$ for all $t, \tau \geq 0$, a service curve yields bounds for the departures, the backlog and the delay at a node. A bound on the departures is given by the envelope $A^* \oslash S$, in the sense that $D(t+\tau) - D(t) \leq A^* \oslash S(\tau)$, for all t and $\tau \geq 0$. A bound for the backlog $B(t)$ can be given as $B(t) \leq A^* \oslash S(0)$. The delay $W(t)$ is bounded by the smallest number d that satisfies $\sup_{\tau \geq 0}\{A^*(\tau - d) - S(\tau)\} \leq 0$.

It is also possible to establish a relationship between service curves and link scheduling algorithms. Using deterministic envelopes, we can write the service curve of a flow corresponding to a given scheduling algorithm by expressing the link capacity that is left unused by other flows. This is frequently called a *leftover service curve*. For a priority scheduling algorithm at an output link with a fixed-rate capacity, the leftover service curve seen by the traffic from priority class p can be expressed as

$$\begin{pmatrix} \text{Service} \\ \text{of class-}p \end{pmatrix} = \begin{pmatrix} \text{Link} \\ \text{capacity} \end{pmatrix} - \begin{pmatrix} \text{Arrivals from} \\ \text{higher priority} \\ \text{classes-}q \end{pmatrix} .$$

If we use A_q to denote the aggregate arrivals from priority q, the leftover service curve for priority p has the form

$$S_p(t) = \Big[Ct - \sum_{q>p} A^*_q(t)\Big]_+ ,$$

where C is the transmission rate of the link, $\sum_{q>p} A^*_q(t)$ denotes the envelope of the arrivals from higher priority classes, and $[x]_+ = \max(0, x)$ [8].

An attractive feature of the network calculus is that 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 ($h = 1, \ldots, H$), as shown in Figure 1. 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

$$S^{net} = S^1 * S^2 * \ldots * S^H . \qquad (2)$$

With network service curves, bounds for the output burstiness, backlog and delay along a path through the network follow directly from the single-node results. End-to-end delay bounds obtained in this fashion are generally better than the sum of the per-node delay bounds. For example, when the service curve at each node is given by a constant-rate function, $S^h(t) = r_h t$, we obtain $S^{net}(t) = S^1 * S^2 * \ldots * S^H(t) = \min(r_1, r_2, \ldots, r_H)t$. At this time, the deterministic calculus has been extensively explored [8], [12] and has led to the development of new scheduling algorithms [17], [31] and network service specifications [21].

III. Statistical Network Calculus

When we approach the network calculus in a probabilistic framework, arrivals $A(t)$ and departures $D(t)$ of a flow at a node in the time interval $[0, t)$ are described by random processes. The random processes are defined over an underlying joint probability space that we suppress in our notation. We define a probabilistic version of a service curve, which defines a probabilistic measure of the available service. A *statistical service curve* is a non-decreasing non-negative function $\mathcal{S}$ that satisfies for all $t \geq 0$, [2]

$$P\Big\{D(t) \geq A * \mathcal{S}(t)\Big\} \geq 1 - \varepsilon . \qquad (3)$$

While statistical service curves $\mathcal{S}$ provide a probabilistic description of the service at a node, they are not themselves random. In the following we present statistical calculus bounds for output envelopes, backlog, and delay. Here, the role of statistical service curves in the derivations mirrors the corresponding deterministic arguments.

[2]We use the convention that 'S' denotes a deterministic service curve, and '$\mathcal{S}$' denotes a statistical service curve.

A. Statistical Performance Bounds

The following theorem is a probabilistic counterpart to the deterministic results discussed in the previous section, which assumes that arrivals of a flow are bounded by a deterministic envelope A^*.

Theorem 1: **Performance Bounds.** Given a flow with arrival process A and envelope A^* that receives a statistical service curve $\mathcal{S}$, the following hold:

1) **Output Envelope.** The function $A^* \oslash \mathcal{S}$ is a probabilistic bound for the departures on $[0, t]$, in the sense that, for all $t, \tau > 0$, $P\{D(t, t+\tau) \leq A^* \oslash \mathcal{S}(\tau)\} \geq 1 - \varepsilon$.
2) **Backlog Bound.** The function $A^* \oslash \mathcal{S}(0)$ is a probabilistic bound for the backlog, in the sense that, for all $t > 0$, $P\{B(t) \leq A^* \oslash \mathcal{S}(0)\} \geq 1 - \varepsilon$.
3) **Delay Bound.** A probabilistic upper bound for the delay is given by

$$d_{max} = \inf\{d \geq 0 \mid \forall t \geq 0 : \ A^*(t-d) \leq \mathcal{S}(t)\}$$

in the sense that for all $t \geq 0$, $P\{W(t) \leq d_{max}\} \geq 1 - \varepsilon$.

Note that the probabilistic bounds are quite similar to those of the deterministic calculus. In fact, by setting $\varepsilon = 0$ in Theorem 1, we recover the corresponding deterministic bounds.

Proof. Let $t \geq 0$ be given. Since A^* is an envelope for A, by definition, we have that

$$A(t) - A(x) \leq A^*(t - x) , \quad \forall x \leq t . \qquad (4)$$

Let $0 \leq s \leq t$ and assume that

$$D(s) \geq A * \mathcal{S}(s). \qquad (5)$$

Then we obtain the output bound

$$\begin{aligned} D(t) - D(s) &\leq A(t) - \inf_{x\in[0,s]}\{A(s-x) + \mathcal{S}(x)\} \\ &\leq \sup_{x\leq s}\{A(t) - A(s-x) - \mathcal{S}(x)\} \\ &\leq \sup_{x\leq s}\{A^*(t-s+x) - \mathcal{S}(x)) \\ &\leq A^* \oslash \mathcal{S}(t-s) . \end{aligned}$$

In the first step, we have used Eqn. (5) and have expanded the convolution operation. The next line rearranges terms. Then we have applied the envelope property as specified in Eqn. (4). Lastly, we have expanded the range of the supremum and applied the deconvolution operator.

Similarly, assuming that Eqn. (5) holds for $s = t$, the backlog satisfies

$$\begin{aligned} B(t) &= A(t) - D(t) \\ &\leq A(t) - \inf_{x\leq t}\{A(t-x) + \mathcal{S}(x)\} \\ &\leq \sup_{x\leq t}\{A^*(x) - \mathcal{S}(x)\} \\ &\leq A^* \oslash \mathcal{S}(0) . \end{aligned}$$

In the first line we have used the definition of the backlog. The second line uses the assumption for time t. The next step rearranges terms and applies the envelope property of

A^*. Finally, we have expanded the range of the supremum and applied the deconvolution operator.

To prove the delay bound we assume again that Eqn. (5) holds for $s = t$. If d is chosen such that $A^*(x) \leq \mathcal{S}(x+d)$ for all $x \geq 0$, then

$$\begin{aligned} A(t-d) - D(t) &\leq A(t-d) - \inf_{x \leq t}\{A(t-x) + \mathcal{S}(x)\} \\ &\leq \sup_{x \leq t}\{A(t-d) - A(t-x) - \mathcal{S}(x)\} \\ &\leq \sup_{x \geq 0}\{A^*(x-d) - \mathcal{S}(x)\} \\ &\leq 0 \ . \end{aligned}$$

The first two lines make the same arguments as the proof of the output bound. Then, we apply Eqn. (4) and expand the range of the supremum. The last step uses the choice of d, and shows that the delay at time t does not exceed d.

For any fixed time s, Eqn. (5) is violated with probability ε or less. Therefore, the above bounds for output, backlog, and delay are violated with probability ε or less, and the theorem follows. □

Similar performance bounds can be derived if the deterministic arrival envelope A^* is replaced by a probabilistic bound on the arrivals. An *effective envelope*, defined in [6], for an arrival process A is a function $\mathcal{G}$ such that for all t and for all $s \leq t$

$$P\Big\{A(t) - A(s) \leq \mathcal{G}(t-s)\Big\} \geq 1 - \varepsilon_a \ .$$

Effective envelopes can be constructed for multiplexed arrivals as well as for individual flows. The effective envelope for an aggregate of flows is generally smaller than the sum of the effective envelopes of the individual flows. This expresses statistical multiplexing. In [6], effective envelopes are constructed using the Chernoff bound and the Central Limit Theorem. Effective envelopes can be related to other probabilistic descriptions of traffic that can capture statistical multiplexing, such as the exponential bounded burstiness model [38], and its generalizations [4], [13], [35], [39], and the effective bandwidth theory [28]. For the min-plus formulation of single-node bounds, we need a stronger envelope which provides bounds for entire sample paths of the arrivals. Let us define a *sample-path effective envelope* as a function $\mathcal{H}$ that satisfies for all $t \geq 0$

$$P\Big\{\sup_{s \leq t}\big\{A(t) - A(s) - \mathcal{H}(t-s)\big\} \leq 0\Big\} \geq 1 - \varepsilon_a \ . \quad (6)$$

With this definition, the condition that sample paths satisfy

$$A(t) - A(s) \leq \mathcal{H}(t-s) \ , \quad \forall s \leq t \ , \quad (7)$$

is violated with probability ε_a. Arrival models in the literature that have been shown to satisfy such sample path bounds include exponentially bounded burstiness [38], linear envelope processes [10], stochastically bounded burstiness [35], and its generalizations [4], [39].

With the sample-path effective envelope, the bounds for output, backlog and delay are directly obtained by replacing A^* in the proof of Theorem 1 with $\mathcal{H}$, and by increasing the violation probability by ε_a. We illustrate this for the output bound. We obtain

$$\begin{aligned} &P\left\{D(t) - D(s) \leq \mathcal{H} \oslash \mathcal{S}(t-s)\right\} \\ &\quad \geq P\{\text{ Eqs. (5) and (7) are satisfied }\} \\ &\quad \geq 1 - \varepsilon - \varepsilon_a \ , \end{aligned}$$

where we have used Boole's inequality in the last line.

In the statistical calculus, leftover service curves have been used to perform a delay analysis of networks with non-trivial scheduling algorithms, such as Static Priority, Earliest-Deadline-First, and Generalized Processor Sharing [28], [30], [32]. In [28], a statistical leftover service curve for a flow j of the form

$$\mathcal{S}_j(t) = \Big[Ct - \sum_{i \neq j} \mathcal{H}_i(t)\Big]_+$$

was used to derive a lower bound for the service seen by a single flow at a link with fixed capacity C that serves an aggregate of flows. Since the sample-path effective envelope $\mathcal{H}$ captures the statistical multiplexing gain, a leftover service curve is generally much larger than its deterministic version.

B. Statistical Network Service Curve

A probabilistic counterpart to the statistical network service curve is provided by the next theorem. Since the theorem does not make assumptions on the arrivals to the network, the network service curve holds for all feasible sample paths of arrivals. In particular, it does not rely on a deterministic or probabilistic description of traffic. To state the theorem, define for $a \geq 0$ the *impulse function*

$$\delta_a(t) = \begin{cases} \infty \ , & \text{if } t > a \ , \\ 0 \ , & \text{if } t \leq a \ . \end{cases}$$

We have that $f * \delta_a(t) = f(t-a)$ for any $t, a \geq 0$ and for any nondecreasing function with $f(0) = 0$.

Theorem 2: **Statistical Network Service Curve.** Consider a flow that passes through H network nodes in series, as shown in Figure 1. Assume that at each node, $\mathcal{S}^h$ is a statistical service curve for the h-th node in the sense of Eqn. (3). Then, for any choice of $a > 0$, a statistical network service curve is given by

$$\mathcal{S}^{net} = \mathcal{S}^1 * \ldots * \mathcal{S}^H * \delta_{(H-1)a}$$

in the sense that

$$P\Big\{D^{net}(t) \geq A^{net} * \mathcal{S}^{net}(t)\Big\} \geq 1 - \varepsilon\left(1 + (H-1)\frac{t}{a}\right) \ . \quad (8)$$

This theorem recovers the deterministic network service curve from Eqn. (2). By setting $\varepsilon = 0$, the inequality in Eqn. (8) holds with probability one. Then by letting $a \to 0$, we obtain the deterministic network service curve from Eqn. (2) almost surely.

The parameter a that appears in Eqn. (8) is a small time scale that is used in the proof to discretize time. Every choice of a provides a valid bound. Larger values of a lead to more pessimistic network service curves, while smaller values of a lead to higher violation probabilities. In applications, it is useful to treat the choice of a in the network service curve

$\mathcal{S}^{net}$ and the choice of ε in the individual service curve $\mathcal{S}^h$ as an optimization problem where one seeks to maximize $\mathcal{S}^{net}(t)$ while keeping the resulting violation probability $\varepsilon\big(1+(H-1)t/a\big)$ fixed. In a discrete-time setting, the parameter a is not needed and the network service curve is given simply by $\mathcal{S}^{net}(t) = \mathcal{S}^1 * \mathcal{S}^2 \ldots * \mathcal{S}^H$ with violation probability $\varepsilon(1+(H-1)t)$.

Proof. We proceed in three steps. In the first step, we modify the effective service curve to give lower bounds on the entire sample path of arrivals and departures over the interval $[0,t]$. In the second step, we manipulate sample paths using techniques of the deterministic calculus. The proof concludes with an estimate for the violation probability.

Step 1: Sample path bound. Suppose that $\mathcal{S}$ is a nondecreasing statistical service curve satisfying Eqn. (3). We will show that, for any choice of $a > 0$,

$$P\big\{\forall x \le t : D(x) \ge A * \mathcal{S}(x-a)\big\} \ge 1 - \varepsilon\frac{t}{a} \ . \qquad (9)$$

To see this, fix a number $a > 0$ to serve as a discretization parameter. Let $j = \lfloor x/a \rfloor$ be the integer part of x/a, set $x_j = ja$, and consider the event

$$E_j = \big\{D(x_j) \ge A * \mathcal{S}(x_j)\big\} \ .$$

If $x \ge a$ and E_j occurs, then

$$D(x) \ge D(x_j) \ge A * \mathcal{S}(x_j) \ \ge A * \mathcal{S}(x-a) \ , \qquad (10)$$

where we have used the fact that both A and $\mathcal{S}$ are nondecreasing in the last step. On the other hand, for $x \le a$, we have $D(x) \ge A*\mathcal{S}(x-a)$, since $A(0) = \mathcal{S}(0) = 0$. It follows that

$$\begin{aligned} & P\big\{\forall x \le t : D(x) \ge A * \mathcal{S}(x-a)\big\} \\ & \quad \ge \ P\big\{\forall j = 1,\ldots,\lfloor t/a \rfloor :\ D(x_j) \ge A * \mathcal{S}(x_j)\big\} \\ & \quad = \ P\Big\{\bigcap_{1 \le j \le \lfloor t/a \rfloor} E_j\Big\} \\ & \quad \ge \ 1 - \varepsilon\frac{t}{a} \ , \end{aligned}$$

which proves Eqn. (9). The first step uses Eqn. (10), the second the definition of the events E_j and the third step uses Boole's inequality.

Step 2. Suppose that, for a particular sample path,

$$\begin{cases} \forall x \le t : & D^h(x) \ge A^h * \mathcal{S}^h * \delta_a(x) \ , \quad h < H \ , \\ & D^H(t) \ge A^H * \mathcal{S}^H(t) \ , \quad\quad h = H \ . \end{cases} \qquad (11)$$

Since $A^H = D^{H-1}$ we can insert the first line of Eqn. (11) with $h = H-1$ into the second line, which yields

$$\begin{aligned} D^H(t) &\ge \inf_{x\in[0,t]}\big\{A^{H-1} * \mathcal{S}^{H-1} * \delta_a(t-x) + \mathcal{S}^H(x)\big\} \\ &= A^{H-1} * \big(\mathcal{S}^{H-1} * \mathcal{S}^H * \delta_a\big)(t) \ . \end{aligned}$$

By induction over the number of nodes, Eqn. (11) implies that $A^{net} = A^1$ and $D^{net} = D^H$ satisfy

$$D^{net}(t) \ge A^{net} * \big(\mathcal{S}^1 * \ldots * \mathcal{S}^H * \delta_{(H-1)a}\big)(t) \ . \qquad (12)$$

Step 3. From the above, it follows that

$$\begin{aligned} & P\big\{D^{net}(t) \ge A^{net} * \big(\mathcal{S}^1 * \ldots * \mathcal{S}^H * \delta_{(H-1)a}\big)(t) \ \big\} \\ & \quad \ge P\{ \text{ Eqn. (11) is satisfied } \} \\ & \quad \ge 1 - \varepsilon\left(1 + (H-1)\frac{t}{a}\right) \ . \end{aligned}$$

The first inequality follows from the fact that Eqn. (11) implies Eqn. (12). The last inequality uses Eqn. (9) for $h = 1,\ldots,H-1$, and the definition of the statistical service curve for $h = H$. □

C. What Makes Statistical Network Calculus Hard?

Since the bound in Eqn. (8) deteriorates as t becomes large, the statistical network service curve in Theorem 2 appears to have limited practical value. It turns out that devising time-independent network service curves that retain the convolution operation is challenging. Here we discuss the key problem encountered when extending end-to-end service curves to a probabilistic setting.

To explain the difficulty in replacing Eqn. (8) with a time-independent bound, consider a network as shown in Figure 1, with $H = 2$ nodes. A statistical service curve $\mathcal{S}^2$ at the second node guarantees that, for any given time t, the departures from this node are with high probability bounded below as

$$D^2(t) \ge A^2 * \mathcal{S}^2(t) = \inf_{\tau\in[0,t]}\big\{A^2(t-\tau) + \mathcal{S}^2(\tau)\big\} \ . \qquad (13)$$

Suppose that the infimum in Eqn. (13) is assumed at some value $\hat{\tau} \le t$. Since the departures from the first node are random even if the arrivals satisfy a deterministic bound, $\hat{\tau}$ is a random variable. A statistical service curve $\mathcal{S}^1$ 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

$$D^1(x) \ge A^1 * \mathcal{S}^1(x) \ . \qquad (14)$$

Since $\hat{\tau}$ is a random variable, we cannot simply evaluate Eqn. (14) for $x = t - \hat{\tau}$ and insert Eqn. (14) into Eqn. (13). Furthermore, there is no general time-independent bound on the distribution of $\hat{\tau}$. Note that this issue does not arise in the deterministic calculus, since deterministic service curves make service guarantees that hold for all values of x.

The problem can be resolved by considering service curves that, by definition, make sample path guarantees, e.g.,

$$P\Big\{\forall t \ge 0 \ : \ D(t) \ge A * \mathcal{S}(t)\Big\} \ge 1 - \varepsilon \ . \qquad (15)$$

However, a service curve of this form, which can be found in [22], has essentially the characteristics of a deterministic service curve, i.e., it expresses worst-case lower bounds on service. In fact, under the broad assumption of stationarity and ergodicity of the arrival function A, for any given service curve $\mathcal{S}$, one can show that the violation probability ε in Eqn. (15) is either 0 or 1.

An important conclusion is that the concept of a network service curve does not carry over to a probabilistic setting by simply transcribing the corresponding expressions from the deterministic network calculus. Additional assumptions are required to establish time-independent bounds on the range of the infimum in Eqn. (13). One example of such an assumption is to require that for all $t \geq 0$:

$$P\Big\{D(t) \geq \inf_{x\in[0,T]}\big\{A(t-x)+\mathcal{S}(x)\big\}\Big\} \geq 1-\varepsilon \ . \qquad (16)$$

This assumption imposes a limit T on the range of the convolution, and hence on the maximal time scale on which the service curve relates service guarantees to arrivals.

In the next section we show that the availability of a time scale bound resolves the problems of the statistical network service curve. Then, in Section V, we approach the same problem differently, by modifying the statistical service curve so that it yields a meaningful multi-node convolution without requiring *a priori* time-scale limits.

IV. Network Service Curves with Time Scale Bounds

We now consider a service curve that has the additional property in Eqn. (16). This permits us to state a statistical network service curve of the following form.

Theorem 3: With the assumptions made in Theorem 2 and assuming that there exists a number $T \geq 0$ such that A^h and $\mathcal{S}^h$ satisfy Eqn. (16) for each $h = 1,\ldots,H$, for any choice of $a > 0$, a statistical network service curve is given by

$$\mathcal{S}^{net} = \mathcal{S}^1 * \ldots * \mathcal{S}^H * \delta_{(H-1)a}$$

in the sense that

$$P\Big\{D^{net}(t) \geq A^{net} * \mathcal{S}^{net}(t)\Big\} \geq 1-\varepsilon H\left(1+\frac{(H-1)T}{2a}\right) \ .$$

Different from Theorem 2, this service guarantee does not degrade as a function of time. Rather, the bounds depend on the time scale T as used in Eqn. (16). When a deterministic arrival envelope A^* is available, the condition on T can be satisfied at a node with a given statistical service curve $\mathcal{S}$ by choosing T such that $A^*(T) \leq \mathcal{S}(T)$. This guarantees that

$$A * \mathcal{S}(t) = \inf_{x\in[0,T]}\big\{A(t-x)+\mathcal{S}(x)\big\} \ .$$

Note, however, that a deterministic envelope, even if it is available, holds only at the first node, leaving the problem open of finding T for subsequent nodes.

A similar time scale bound can be obtained from a sample-path effective envelope $\mathcal{H}$ that satisfies Eqn. (6). While such arrival bounds are sometimes easy to obtain at the first node of a network, the analysis for downstream nodes requires additional assumptions. For a wide range of service curves, including probabilistic versions of rate-latency service curves [36], it is feasible to obtain a time scale bound when buffer constraints impose a limit on the maximum backlog or when traffic that exceeds a maximum waiting time is discarded. Since the manuscript for this paper became available [9], several papers have explored assumptions under which time-scale bounds can be obtained. For example, [3] obtains time scale bounds by assuming that each node drops traffic that locally violates a given delay guarantee. In [28], [30], time scale bounds are obtained by devising bounds on the busy period for service curves at downstream nodes.

Proof. The proof is analogous to the proof of Theorem 2, and proceeds in the same three steps.

Step 1: Sample path bound. Suppose that $\mathcal{S}$ is a statistical service curve satisfying Eqn. (16), and let $\ell > 0$. Fix $a > 0$, set $x_j = t - \ell + ja$, and consider the events

$$E_j = \Big\{D(x_j) \geq \inf_{y\in[0,T]}\big\{A(x_j-y)+\mathcal{S}(y)\big\}\Big\} \ , \ j=0,\ldots,n,$$

where $n = \lfloor \ell/a \rfloor$. If $x \in [x_j, x_{j+1})$ and the event E_j occurs, we see as in Eqn. (10) that

$$D(x) \ \geq \ \inf_{y\in[0,T]}\big\{A * \delta_a(x-y)+\mathcal{S}(y)\big\}.$$

As in the proof of Theorem 2 it follows that

$$\begin{aligned} &P\big\{\forall x \in [t-\ell,t] : D(x) \geq \inf_{y\in[0,T]}\big\{A*\delta_a(x-y)+\mathcal{S}(y)\big\}\big\} \\ &\geq \ P\Big\{\bigcap_{0\leq j\leq n} E_j\Big\} \\ &\geq \ 1-\varepsilon\left(1+\frac{\ell}{a}\right) \ . \end{aligned}$$

Step 2. Suppose that

$$\begin{cases} \forall x \in [t-(H-h)T,t] : D^h * \delta_{(H-h-1)a}(x) & \\ \quad \geq \inf_{y\in[0,T]}\big\{A^h * \delta_{(H-h)a}(x-y)+\mathcal{S}^h(y)\big\} \ , & h<H, \\ D^H(t) \geq \inf_{y\in[0,T]}\big\{A^H(t-y)+\mathcal{S}^H(y)\big\} \ , & h=H \ . \end{cases} \qquad (17)$$

Repeatedly inserting the first equation into the second equation, and using that $A^{net} = A^1$ and $D^{net} = D^H$, we see as in Eqn. (12) that

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

Step 3. Combining Steps 1 and 2, we obtain

$$\begin{aligned} &P\big\{D^{net}(t) \geq A^{net} * \mathcal{S}^{net}(t)\big\} \\ &\qquad \geq P\big\{\text{Eqn. (17) is satisfied}\big\} \\ &\qquad \geq 1-\varepsilon\left(1+\sum_{h=1}^{H-1}\left(1+\frac{(H-h)T}{a}\right)\right) \ . \end{aligned}$$

Here, the first line follows from Step 2, and the second line follows by using Step 1 with $\ell_h = (H-h-1)T$ for node $h = 1,\ldots,H-1$ and from the assumption on $\mathcal{S}^H$. Evaluating the sum yields the claim. □

V. Network Service Curves with Service Guarantees for Intervals

In this section, we present an alternative approach to the development of a statistical network service curve. Rather than trying to formulate conditions under which the statistical service curve from Eqn. (3) results in a useful network service curve, we seek a modified statistical service curve that can provide end-to-end guarantees without such external bounds.

It turns out that we can dispense with the need for a time scale limit by defining a class of service curves that provide service guarantees on the departures over intervals of a given length. In the following we will use $A(x,y)$ and $D(x,y)$ to denote the arrivals and departures in the time interval $[x,y)$, with $A(x,y)=A(y)-A(x)$ and $D(x,y)=D(y)-D(x)$. We define a modified convolution operator $A *_t S(\ell)$ by setting, for $\ell > 0$,

$$A *_t S(\ell) = \min\Big\{S(\ell), B(t) + \inf_{x \le \ell}\{A(t,t+\ell-x)+S(x)\}\Big\} .$$

For $\ell \le 0$ we set $A *_t S(\ell) = 0$. The essential property of this modified operator is that the infimum only involves arrivals in the interval $[t,t+\ell)$. Note that $A *_t S(\ell)$ depends also on the backlog at time t. When $t=0$, it reduces to the usual convolution operator, since by definition $B(0)=0$.

In the deterministic network calculus, this type of convolution has appeared in [1], [8], [18] to define service curves of the form

$$D(t,t+\ell) \ge A *_t S(\ell) \ , \quad (18)$$

which are called *adaptive service curves* in [8]. Eqn. (18) is equivalent to requiring that S satisfies Eqn. (1) for the time-shifted arrivals and departures defined by

$$\tilde{A}(\ell) = B(t) + A(t,t+\ell) \ , \quad \tilde{D}(\ell) = D(t,t+\ell) \ . \quad (19)$$

Figure 2 illustrates the time-shifted arrivals. When $t=0$, Eqn. (18) implies that S is a service curve that satisfies Eqn. (1). The difference between the two definitions is that Eqn. (18) assigns no special role to the time $t=0$.

As has been remarked in [1], [8], [18], many service curves with applications in packet networks, including shapers, schedulers with delay guarantees, and rate-controlled schedulers such as GPS, can be expressed in terms of Eqn. (18). In particular, this includes the group of so-called *strict* service curves, which guarantee departures of at least $S(\ell)$ over any interval where the backlog remains positive. The main motivation for defining a service curve as in Eqn. (18) is that a (deterministic) network service curve obtained by a convolution of strict service curves is generally adaptive but not strict.

Now we turn to a probabilistic version of the adaptive service curve. We say that a function $\mathcal{S}$ defines a statistical service curve for intervals of length ℓ, if the departures on any interval $[t,t+\ell)$ satisfy

$$P\Big\{D(t,t+\ell) \ge A *_t \mathcal{S}(\ell)\Big\} \ge 1-\varepsilon_\ell \ . \quad (20)$$

Here, the violation probability ε_ℓ depends on the length of the interval. For $\varepsilon_\ell \to 0$, Eqn. (20) turns into the deterministic condition in Eqn. (18).

We will show that statistical service curves satisfying Eqn. (20) yield a statistical network service curve with desirable properties. The modified service curves permits us to calculate performance bounds for departure envelopes, backlog and delay. Then, we will make a case that service curves of this type are in fact practical and abound by showing that the prototypical leftover service curve complies with the modified service curve definition.

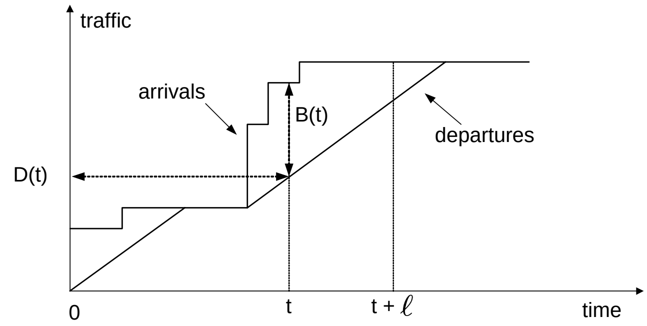

Fig. 2. Illustration for the modified convolution operator. The operator $*_t$ uses the backlog at time t and the arrivals in the interval $[t,t+\ell)$.

We now present a probabilistic version of a statistical network service curve that is obtained via a concatenation of service curves satisfying Eqn. (20). This is the content of the following theorem.

Theorem 4: Consider a flow that passes through H network nodes in series. Assume that the service at each node h is given by a statistical service curve $\mathcal{S}^h$ satisfying Eqn. (20). Then, for any choice of $a>0$, we have that

$$\mathcal{S}^{net} = \mathcal{S}^1 * \mathcal{S}^2 * \ldots * \mathcal{S}^H * \delta_{(H-1)a}$$

is a statistical service curve for intervals that satisfies

$$P\Big\{D^{net}(t,t+\ell) \ge A^{net} *_t \mathcal{S}^{net}(\ell)\Big\} \ge 1-\varepsilon_\ell\Big(1+(H-1)\frac{\ell}{a}\Big) \ . \quad (21)$$

A comparison with Theorem 3 makes clear the advantages of the interval-based service curve. Both theorems involve a time scale, T in Theorem 3 and ℓ in Theorem 4, and the violation probability deteriorates for large values of the time scale. The key difference is that the time scale T of Eqn. (16) appears in the *assumptions* of Theorem 3, while the conclusion of Theorem 4 holds for *all* choices of ℓ. In other words, in Theorem 4 we only need to make a good choice of ℓ. There is no need to derive the time scale bound from computations on the busy period or *a priori* limits on the backlog or delay. Below we will explain how to choose the parameter ℓ.

Proof. Fix $t \ge 0$. At the h-th node define the time-shifted arrivals $\tilde{A}^h$ and departures $\tilde{D}^h$ according to Eqn. (19), and rewrite the assumption in Eqn. (20) as

$$P\Big\{\tilde{D}^h(\ell) \ge \tilde{A}^h * \mathcal{S}^h(\ell)\Big\} \ge 1-\varepsilon_\ell \ .$$

Theorem 2 implies that

$$P\Big\{\tilde{D}^{net}(\ell) \ge \tilde{A}^{net} * \mathcal{S}^{net}(\ell)\Big\} \ge 1-\varepsilon_\ell\Big(1+(H-1)\frac{\ell}{a}\Big) \ .$$

Reversing the time shift yields the claim in Eqn. (21). □

The theorem can be combined with a bound on arrivals to obtain probabilistic bounds on output envelopes, backlog, and delay. The next theorem shows that Eqn. (20) provides performance bounds analogous to Theorem 1.

Theorem 5: **Performance bounds.** Given a flow with a stationary arrival process A and a deterministic envelope A^*. Assume that the service to the flow at a node over intervals of length ℓ is guaranteed by a statistical service curve that satisfies Eqn. (20). If ℓ is chosen such that $A^*(\ell) < \mathcal{S}(\ell)$ then the following bounds hold:

1) **Output Envelope.** The function $A^* \oslash \mathcal{S}$ is a probabilistic bound for the departures, in the sense that, for all $t, \tau \geq 0$, $P\{D(t, t+\tau) \leq A^* \oslash \mathcal{S}(\tau)\} \geq 1 - \varepsilon_\ell \frac{\mathcal{S}(\ell)}{\mathcal{S}(\ell) - A^*(\ell)}$.
2) **Backlog Bound.** The function $A^* \oslash \mathcal{S}(0)$ is a probabilistic bound for the backlog, in the sense that, for all $t > 0$, $P\{B(t) \leq A^* \oslash \mathcal{S}(0)\} \geq 1 - \varepsilon_\ell \frac{\mathcal{S}(\ell)}{\mathcal{S}(\ell) - A^*(\ell)}$.
3) **Delay Bound.** A probabilistic upper bound for the delay is given by
$$d_{max} = \inf \{d \geq 0 \mid \forall t \geq 0: \ A^*(t-d) \leq \mathcal{S}(t)\}$$
in the sense that for all $t \geq 0$, $P\{W(t) \leq d_{max}\} \geq 1 - \varepsilon_\ell \frac{\mathcal{S}(\ell)}{\mathcal{S}(\ell) - A^*(\ell)}$.

In the theorem, ℓ should be selected so that the probability of violating the bounds is minimized. Thus, the theorem permits to select ℓ using only information on the arrival bounds and the service curve at a node. The value of ℓ is related to the largest time scale on which arrivals and departures are correlated. In general, the parameter provides a bound on the range of the infimum analogous to Eqn. (16).

Proof. We show only the backlog bound, as the proofs of the other claims are similar. Since the arrival process is stationary and $A(0) = D(0) = 0$, the backlog process $B(t)$ is stochastically increasing in the sense that
$$P\{B(t) > b\} \leq P\{B(t+\ell) > b\} .$$
This can be found, for example, in ([12], p. 293). Thus, it suffices to prove the claim for $B(t+\ell)$.

Assume that $A^*(\ell) < \mathcal{S}(\ell)$ and set $b_{max} = A^* * \mathcal{S}(0)$. If
$$D(t, t+\ell) \geq A *_t \mathcal{S}(\ell) , \tag{22}$$
then
$$\begin{aligned} B(t+\ell) &= B(t) + A(t, t+\ell) - D(t, t+\ell) \\ &\leq \max\Big\{B(t) + A(t, t+\ell) - \mathcal{S}(\ell), \\ &\qquad \sup_{x \leq \ell}\big\{A(t, t+\ell) - A(t, t+\ell-x) - \mathcal{S}(x)\big\}\Big\} \\ &\leq \max\Big\{B(t) - \big(\mathcal{S}(\ell) - A^*(\ell)\big), b_{max}\Big\} . \end{aligned}$$
Here, the first step uses the definition of the backlog, the second step uses the assumption in Eqn. (20), and the last step uses the arrival envelope and the definition of b_{max}. For $B(t+\ell) > b_{max}$ this implies
$$B(t+\ell) \leq B(t) - \big(\mathcal{S}(\ell) - A^*(\ell)\big) . \tag{23}$$
If, on the other hand, Eqn. (22) fails, we still have the estimate
$$\begin{aligned} B(t+\ell) &= B(t) + A(t, t+\ell) - D(t, t+\ell) \\ &\leq B(t) + A^*(\ell) . \end{aligned} \tag{24}$$
Since Eqn. (22) is violated with probability at most ε_ℓ, it follows from Eqs. (23) and (24) that the expected value $E[B(t+\ell) - B(t)]$ satisfies
$$\begin{aligned} 0 &\leq \mathrm{E}\big[B(t+\ell) - B(t)\big] \\ &\leq \varepsilon_\ell A^*(\ell) \\ &\quad -\big(1 - \varepsilon_\ell - P\big\{B(t+\ell) \leq b_{max}\big\}\big)\big(S(\ell) - A^*(\ell)\big) . \end{aligned}$$
Solving for $P\big\{B(t+\ell) \leq b_{max}\big\}$ yields $P\{B(t+\ell) \leq A^* \oslash \mathcal{S}(0)\} \geq 1 - \varepsilon_\ell \frac{\mathcal{S}(\ell)}{\mathcal{S}(\ell) - A^*(\ell)}$, which proves the claim. □

As was done following Theorem 1, the arrival envelope can be replaced by a sample-path effective envelope $\mathcal{H}$ that satisfies Eqn. (6). Recall that an effective envelope can express the result of statistically multiplexed arrivals from many sources. As before, replacing A^* by $\mathcal{H}$ increases the violation probability by ε_a.

Theorems 4 and 5 can be combined to provide end-to-end service guarantees over multiple nodes. Compared with Theorem 3, the advantage is that the assumption on ℓ can be checked by considering only the network service curve and the arrivals to the network at the ingress node; no separate analysis of downstream nodes is needed.

The free parameters ε (which determines the individual service curves $\mathcal{S}^h$) and ℓ (which determines the violation probability in Theorem 4) need to be chosen so that $A^*(\ell) < \mathcal{S}^{net}(\ell)$ and the resulting violation probability takes an acceptable value (typically on the order of 10^{-6} or 10^{-9}). In applications one seeks to optimize the resulting performance bounds under these constraints. For simple arrival and service models this can sometimes be done analytically, but in general numerical methods are needed to find good choices of ε and ℓ.

Finally, we want to argue that the interval-based calculus applies to a large class of practically relevant service curves. To make this case, we show that a frequently used leftover service curve, i.e., a service curve experienced by a flow at a link scheduler with static priorities, complies to the interval-based definition of service. We emphasize that similar results can be derived for service curves that relate to other scheduling algorithms, such as GPS.

Theorem 6: **Statistical leftover service at a priority scheduler.** Consider a workconserving priority scheduler that operates at a constant rate C. Arriving flows are grouped into Q classes labeled $q = 1, \ldots, Q$, with higher values of q denoting higher priority. Let $\mathcal{H}_q(t)$ be a sample-path effective envelope for the aggregate arrivals from flows in class q, satisfying Eqn. (6). For each class define a function
$$\mathcal{S}_p(t) = \inf_{s \geq t}\Big[Cs - \sum_{q>p} \mathcal{H}_q(s)\Big]_+ .$$
Then the service available to class p over any time interval of length ℓ is bounded below as
$$P\{D_p(t, t+\ell) \geq A_p *_t \mathcal{S}_p(\ell)\} \geq 1 - (Q-p)\varepsilon_a . \tag{25}$$

Proof. For a given priority index p and a given time interval $[t, t+\ell)$, consider a sample path that satisfies
$$A_q(s, t+\ell) \leq \mathcal{H}_q(t+\ell-s) \tag{26}$$

for all $s \le t+\ell$ and all $q > p$. Let $\underline{x}$ be the beginning of the busy period of classes $q \ge p$ that contains $t+\ell$, i.e.,

$$\underline{x} = \sup\Big\{ s \le t+\ell : \sum_{q \ge p} B_q(s) = 0 \Big\} \ .$$

If $\underline{x} \ge t$, then

$$\begin{aligned} & D_p(t, t+\ell) \\ & \quad = D_p(t, \underline{x}) + D_p(\underline{x}, t+\ell) \\ & \quad = D_p(\underline{x}) - D_p(t) + \Big(C(t+\ell-\underline{x}) - \sum_{q>p} D_q(\underline{x}, t+\ell) \Big) \\ & \quad \ge A_p(\underline{x}) - D_p(t) + \Big[C(t+\ell-\underline{x}) - \sum_{q>p} A_q(\underline{x}, t+\ell) \Big]_+ \\ & \quad \ge \inf_{s \in [t, t+\ell)} \Big(B_p(t) + A_p(t,s) \\ & \qquad + \Big[C(t+\ell-s) - \sum_{q>p} A_q(s, t+\ell) \Big]_+ \Big) \ . \end{aligned}$$

In the first line, we have split the interval $[t, t+\ell)$ into two disjoint subintervals. In the second line, we have used that the scheduler is workconserving and that classes $q > p$ have higher priority than p. In the third line, we have used that $A_q(\underline{x}) = D_q(\underline{x})$ for all $q \ge p$, and that $D_q(t+\ell) \le A_q(t+\ell)$ for all $q > p$. In the last line we have rearranged terms and replaced evaluation at $\underline{x}$ by an infimum. Eqn. (26) implies that we can replace $A_q(s, t+\ell)$ with $\mathcal{H}_q(t+\ell-s)$ for any $s \le t+\ell$ and all $q > p$. It follows that in this case

$$\begin{aligned} D_p(t, t+\ell) & \ge B_p(t) + \inf_{s \in [t, t+\ell)} \big(A_p(t,s) + \mathcal{S}_p(t+\ell-s) \big) \\ & \ge A_p *_t \mathcal{S}_p(\ell) \ . \end{aligned} \tag{27}$$

If, on the other hand, $\underline{x} < t$, we let

$$\underline{t} = \sup\Big\{ s \le t : \sum_{q>p} B_q(s) = 0 \Big\}$$

be the beginning of the busy period for classes with priority higher than p that contains t. Then

$$\begin{aligned} D_p(t, t+\ell) & = D_p(\underline{t}, t+\ell) \\ & = C(t+\ell-\underline{t}) - \sum_{q>p} D_q(\underline{t}, t+\ell) \\ & \ge \Big[C(t+\ell-\underline{t}) - \sum_{q>p} A_q(\underline{t}, t+\ell) \Big]_+ \\ & \ge \inf_{s \le t} \Big[C(t+\ell-s) - \sum_{q>p} \mathcal{H}_q(t+\ell-s) \Big]_+ \end{aligned}$$

In the first line we have used that $D_p(\underline{t}, t) = 0$ by the defining property of the priority server, since the backlog of the higher priority classes remains positive in that interval. In the second line we have used that the server is workconserving. In the third line we have used that $D_q(\underline{t}) \le A_q(\underline{t})$ since $B_q(\underline{t}) = 0$ and applied Eqn. (26). In this case, we arrive at

$$D_p(t, t+\ell) \ge \mathcal{S}_p(\ell) \ge A_p *_t \mathcal{S}_p(\ell) \ . \tag{28}$$

Combining Eqs. (27) and (28) and using the assumptions on $\mathcal{H}_q$, we obtain

$$\begin{aligned} & P\big\{ D_p(t, t+\ell) < A_p *_t \mathcal{S}_p(\ell) \big\} \\ & \quad \le \sum_{q>p} P\Big\{ \sup_{s \le t+\ell} \big(A_q(s, t+\ell) - \mathcal{H}_q(t+\ell-s) > 0 \big) \Big\} \\ & \quad \le (Q-p)\varepsilon_a \ , \end{aligned}$$

as claimed. □

VI. CONCLUSIONS

We have presented a network calculus with probabilistic service guarantees. We have introduced the notion of *statistical service curves* as a probabilistic bound on the service received by an aggregate of flows or an individual flow. While performance bounds from the deterministic network calculus can be straightforwardly carried over to a probabilistic framework by inserting appropriate probabilistic arguments, the same is not true for the formulation of a network service curve. The results in this paper show that a multi-node version of the statistical network calculus requires us to control the range of the convolution operation when concatenating statistical service curves. Such limits on a maximum relevant time scale can follow from external assumptions (as discussed in Section IV) or from appropriately modified interpretations of service curves (as shown in Section V).

ACKNOWLEDGMENTS

We thank Rene Cruz for pointing out problems in an earlier version of [9]. We also gratefully acknowledge valuable comments from Chengzhi Li.

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] G. Appenzeller, I. Keslassy, and N. McKeown. Sizing router buffers. In *Proc. of ACM Sigcomm 04*, pages 281–292, August 2004.

[3] 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.

[4] S. Ayyorgun and W. Feng. A systematic approach for providing end-to-end probabilistic QoS guarantees. In *Proc. of ICCCN 2004*, pages 115–122, October 2004.

[5] F. L. Baccelli, G. Cohen, G. J. Olsder, and J.-P. Quadrat. *Synchronization and Linearity: An Algebra for Discrete Event Systems*. John Wiley and Sons, 1992.

[6] 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, December 2000.

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

[8] J. Y. L. Boudec and P. Thiran. *Network Calculus*. Springer Verlag, Lecture Notes in Computer Science, LNCS 2050, 2001.

[9] 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. *arxiv.org/abs/cs.NI/0205001*.

[10] 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.

[11] 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.

[12] C. S. Chang. *Performance Guarantees in Communication Networks*. Springer Verlag, 2000.

[13] F. Ciucu, A. Burchard, and J. Liebeherr. Scaling properties of statistical end-to-end bounds in the network calculus. *IEEE Transaction on Information Theory and ACM/IEEE Transactions on Networking. Joint Special Issue on Networking and Information Theory*, June 2006.

[14] R. L. Cruz. A calculus for network delay, Parts I,II. *IEEE Transaction of Information Theory*, 37(1):114–141, 1991.

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

[16] R. L. Cruz. Quality of service management in integrated services networks. In *Proceedings of the 1st Semi-Annual Research Review, CWC*, UCSD, June 1996.

[17] R. L. Cruz. SCED+: Efficient management of quality of service guarantees. In *IEEE Infocom*, pages 625–634, 1998.

[18] R. L. Cruz and C. Okino. Service gurantees for flow control protocols. In *Proceedings of the 34th Allerton Conference on Communications, Control and Computating*, October 1996.

[19] 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, Aug. 1995.

[20] M. Fidler. An end-to-end probabilistic network calculus with moment generating functions for efficient utilization of independence. Technical Report NTNU Trondheim, Norway, 2005, *arxiv.org/abs/cs.IT/0507004*.

[21] V. Firoiu, J.-Y. L. Boudec, D. Towsley, and Z.-L. Zhang. Theories and models for internet quality of service. *Proceedings of the IEEE*, 90(9):1565–1591, Sept. 2002.

[22] Y. Jiang and P. J. Emstad. Analysis of stochastic service guarantees in communication networks: A server model. In *Proc. of the International Workshop on Quality of Service (IWQoS 2005)*, pages 233–245, June 2005.

[23] 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.

[24] G. Kesidis and T. Konstantopoulos. Extremal shape-controlled traffic patterns in high-speed networks. *IEEE Transaction of Communications*, 48(5):813–819, 2000.

[25] E. Knightly. Enforceable quality of service guarantees for bursty traffic streams. In *Proceedings of IEEE INFOCOM'98*, pages 635–642, Mar. 1998.

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

[27] J. Kurose. On computing per-session performance bounds in high-speed multi-hop computer networks. In *ACM Sigmetrics'92*, pages 128–139, 1992.

[28] 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.

[29] J. Liebeherr. A note on statistical multiplexing and scheduling in video networks at high data rates. In *In: The Internet as a Large-Scale Complex System (Editors: K. Park, W. Willinger), Oxford University Press, SFI Book Series*, pages 179–201, 2005.

[30] 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*, pages 1680–1690, May 2003.

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

[32] J. Qiu and E. Knightly. Inter-class resource sharing using statistical service envelopes. In *Proceedings of IEEE Infocom '99*, pages 36–42, March 1999.

[33] J. Roberts, U. Mocci, and J. V. (Eds.). *Broadband Network Traffic: Performance Evaluation and Desgin of Broadband Multiservice Networks. Final Report of Action. COST 242. (Lecture Notes in Computer Science. Vol. 1152)*. Springer Verlag, 1996.

[34] S. Shenker, C. Partridge, and R. Guerin. Specification of guaranteed quality of service. IETF RFC 2212, September 1997.

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

[36] 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.

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

[38] 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.

[39] 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.

Almut Burchard received the Diplom-Mathematiker degree from the University of Heidelberg, Germany in 1989, and the Ph.D. degree in Mathematics from the Georgia Institute of Technology in 1994. From 1994-1998, she was a junior member of the Department of Mathematics at Princeton University. She was employed at the University of Virginia as an Assistant Professor in 1998, and was promoted to Associate Professor in 2004. In 2005, she joined the University of Toronto as Associate Professor in the Department of Mathematics. Her main research interest is sharp inequalities in Analysis. Dr. Burchard received the Sigma Xi Ph.D. thesis award at Georgia Tech in 1994, and was awarded a Sloan Research Fellowship in 1999. She is a member of the American Mathematical Society and the International Association of Mathematical Physics.

Jörg Liebeherr (S'88, M'92, SM'03) received the Ph.D. degree in Computer Science from the Georgia Institute of Technology in 1991. After a Postdoc at the University of California, Berkeley, he joined the Department of Computer Science at the University of Virginia in 1992. From 1997 to 1998 he was an Associate Professor in the Department of Electrical Engineering at Polytechnic University. Since Fall 2005, he is with the Department of Electrical and Computer Engineering of the University of Toronto as the Nortel Chair of Network Architecture and Services. He has served on editorial boards and program committees of several journals and conferences in computer networking. He was elected Member-at-Large of the IEEE Communications Society Board of Governors from 2003 to 2005, and chair of the IEEE Communications Society Technical Committee on Computer Communications from 2004 to 2005.

Stephen Patek (S'88, M'98) received the Ph.D. degree in Electrical Engineering and Computer Science from the Massachusetts Institute of Technology in 1997. He joined the Department of Systems and Information Engineering at the University of Virginia in 1998, where he is now an Associate Professor. He has served as associate editor and technical program committee member for a number of journals and conferences at the intersection of systems theory, automatic control, and computer networks.