

Statistical Analysis of Link Scheduling on Long Paths

Yashar Ghiassi-Farrokhfal, Jörg Liebeherr, Almut Burchard

Abstract

We study how the choice of packet scheduling algorithms influences end-to-end performance on long network paths. Taking a network calculus approach, we consider both deterministic and statistical performance metrics. A key enabling contribution for our analysis is a significantly sharpened method for computing a statistical bound for the service given to a flow by the network as a whole. For a suitably parsimonious traffic model we develop closed-form expressions for end-to-end delays, backlog, and output burstiness. The deterministic versions of our bounds yield optimal bounds on end-to-end backlog and output burstiness for some schedulers, and are highly accurate for end-to-end delay bounds.

I. INTRODUCTION

Link scheduling algorithms for packet switches have been studied extensively, e.g., [19], however, little is known about the impact of link scheduling in large networks with long end-to-end paths. As a case in point, the packet dispersion of a CBR traffic flow at an overloaded First-in-First-Out (FIFO) link with cross traffic [15] is given by

$$r_{out} = \frac{r_{in}}{r_{in} + r_c} C,$$

where C is the constant-rate link capacity, r_c is the arrival rate of CBR cross traffic, and r_{in} and r_{out} , respectively, are the arrival and output rate of the traffic flow (with $r_{in} + r_c > C$). Considering a path of a large number of overloaded FIFO links (with homogeneous cross traffic and link capacities), the output rate was shown in [9] to converge to

$$r_{out} \xrightarrow{H \rightarrow \infty} [C - r_c]_+,$$

where H is the number of traversed links and $[x]_+ = \max\{x, 0\}$. This is the same output rate observed in a network with priority scheduling where cross traffic is given higher priority. Thus, the question arises whether the role of link scheduling diminishes on long paths in more general settings.

This paper develops an analysis that can assess the impact of a broad class of link scheduling algorithms on end-to-end performance. Taking a network calculus approach [11], we characterize traffic in terms of *envelope* functions and service by *service curve* functions. We present bounds on end-to-end performance metrics for statistical, as well as worst-case assumptions on traffic. Our study considers all scheduling algorithms for which the transmission order of backlogged packets is entirely determined by the difference of their arrival times. These schedulers, referred to as Δ -schedulers, include FIFO, priority scheduling,

Y. Ghiassi-Farrokhfal (yashar@comm.utoronto.ca) and J. Liebeherr (jorg@comm.utoronto.ca) are with the Department of Electrical and Computer Engineering, University of Toronto. A. Burchard (almut@math.utoronto.ca) is with the Department of Mathematics, University of Toronto.

The research in this paper is supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC).

and deadline based scheduling. A key contribution of this paper is a significantly sharpened method for computing a statistical bound on the service given to a flow on an entire network path from statistical characterizations of the service at each node of the path. For traffic models that are characterized by a long-term rate and a short-term burst, we derive closed-form expressions for statistical and worst-case bounds on end-to-end delay and backlog. We evaluate the tightness of the bounds by computing performance metrics for specific adversarial arrival scenarios. In some cases, e.g., for backlog bounds for some of the considered schedulers, we are able to show that our bounds are optimal.

The analysis in this paper is made possible by two advances in the network calculus. The first is a tight description of schedulers by service curves at a single nodes from [14]. The second is presented here. We advance the state-of-the-art of the stochastic network calculus [10] by improving methods for concatenating statistical per-node service descriptions to achieve service descriptions for an entire network path. Compared to the best available methods [4, 14], we achieve significantly improved end-to-end performance bounds for many scheduling algorithms. For the special case of deterministic scenarios (i.e., where bounds are never violated), we recover worst-case end-to-end delay bounds known for FIFO networks [8, 12], while extending these bounds to more complex scheduling algorithms. For FIFO scheduling, researchers have found conditions when the output of a node has similar characteristics as the input, justifying a decomposition analysis where each node can be analyzed in isolation [1, 7, 18]. Our results indicate that the conditions for decomposition are often met in non-FIFO networks, as long as the scheduling in the network is sufficiently distinct from a network with strict priorities.

Our study provides new insight into the role of scheduling algorithms in networks with long paths. For a network that is not permanently overloaded, we find that the differences between delay and backlog bounds at a single node persist in a multi-node setting. In fact, for the traffic and service models considered in this paper, the differences between bounds for various schedulers grow proportionally with the path lengths.

The remaining sections of this paper are structured as follows. In Sec. II, we discuss the probabilistic traffic and service characterization in the network calculus, with traffic envelopes and service curves, and present our result of a new network service characterization. In Sec. III, we describe the class of Δ -scheduling algorithms, and derive a network service description for a path of nodes with such schedulers. In Sec. IV, we derive closed-form end-to-end bounds on backlog, delay and the burstiness of output from a network. In Sec. V, we derive lower bounds, thus enabling a discussion of the tightness of our results. We present numerical examples in Sec. VI and conclude the paper in Sec. VII.

II. NETWORK CALCULUS FRAMEWORK

We take a network calculus modeling and analysis approach where arrivals and service of a flow are expressed in terms of deterministic or probabilistic bounds, referred to as *traffic envelopes* and *service curves*, respectively [11]. We next discuss needed definitions and concepts of the deterministic and stochastic network calculus. The last subsection contains the key contribution of this paper: a new result to compute statistical service curves for a network path, that enables us to provide a sharpened analysis of end-to-end performance bounds.

A. Traffic Envelopes

Consider a buffered link with a link scheduling algorithm, referred to as a node. Using a left-continuous time model, a sample path of the arrivals to the node in the time interval $[0, t)$ is denoted by $A(t)$. We assume that traffic arrivals can be described by a stationary random process, or at least satisfy stationary bounds. Traffic departing from a node in $[0, t)$ is denoted by $D(t)$. Both $A(t)$ and $D(t)$ are nondecreasing functions with $A(0) = D(0) = 0$, and we have $D(t) \leq A(t)$ for all $t \geq 0$. For brevity, we use the notation $A(s, t) = A(t) - A(s)$ and $D(s, t) = D(t) - D(s)$ for arrivals and departures in the time interval $[s, t)$.

We characterize arrivals to a node in terms of bounds on sample paths. A *deterministic sample path envelope* E provides a worst-case description of traffic in the sense that it satisfies for all $t \geq 0$

$$\sup_{0 \leq s \leq t} \{A(s, t) - E(t - s)\} \leq 0, \quad (1)$$

with $E(t) = 0$ if $t \leq 0$. By definition, traffic arrivals never violate a deterministic envelope.

For a statistical network analysis we need a probabilistic analogue, i.e., a bound for all sample paths of the arrivals which may be violated with a small probability. A *statistical sample path envelope* \mathcal{G} is a function that satisfies for all $t \geq 0$

$$\mathbb{P}\left(\sup_{0 \leq s \leq t} \{A(s, t) - \mathcal{G}(t - s; \sigma)\} > 0\right) \leq \varepsilon(\sigma), \quad (2)$$

where σ represents one or several parameters, and $\varepsilon(\sigma) \geq 0$ is a bounding function. Further, $\mathcal{G}(t; \sigma) = 0$ if $t \leq 0$. We can view $\varepsilon(\sigma)$ as a bound on the probability of violating the envelope. The notation $\mathcal{G}(t; \sigma)$ and $\varepsilon(\sigma)$ indicate that both the envelope and the bounding function depend on the same parameter(s). Setting $\varepsilon(\sigma) = 0$ recovers a deterministic sample path envelope.

For numerical evaluations and optimizations, we will use a parsimonious envelope characterization for traffic flows consisting of a rate parameter $\rho > 0$ and a burst parameter $\sigma \geq 0$, such that for all $s, t \geq 0$,

$$A(s, t) \leq \rho(t - s) + \sigma. \quad (3)$$

A deterministic sample path bound for such an arrival characterization is given by the envelope $E(t) = \rho t + \sigma$ (for $t > 0$). This bound corresponds to the maximal output of a leaky bucket traffic regulator with rate ρ and burst size σ .

The probabilistic version of this bound takes the form that for all $s, t \geq 0$,

$$\mathbb{P}(A(s, t) - \rho(t - s) > \sigma) \leq \varepsilon(\sigma), \quad (4)$$

which can be read as the probability that the arrivals violate a long-term rate ρ by more than σ being bounded by $\varepsilon(\sigma)$. When we assume that $\varepsilon(\sigma)$ has an exponential decay, i.e., $\varepsilon(\sigma) = M e^{-\alpha \sigma}$ for given constants M and α , we obtain the Exponential Bounded Burstiness (EBB) traffic model [17], which is related to the linear envelope characterization by Chang [2]. We will write $A \sim (\rho, \alpha, M)$ to indicate that A is an EBB flow with parameters ρ , α , and M . The EBB model can express non-trivial processes, such as Markov-modulated On-Off processes [2], but it does not apply to heavy-tailed or long-range correlated traffic. Extensions of the EBB traffic model have been proposed for bounding functions with faster than polynomial decay [16] and even heavy-tailed decay [13].

While a bound as in Eq. (4) is not a statistical sample path envelope (satisfying Eq. (2)), such an envelope can be constructed with an application of the union bound [4]. For the EBB traffic model this yields a statistical sample path envelope

$$\mathcal{G}(t; \sigma) = (\rho + \gamma)t + \sigma, \quad \varepsilon(\sigma) = M e \left(1 + \frac{\rho}{\gamma}\right) e^{-\alpha\sigma} \quad (5)$$

for any choice of $\gamma > 0$, where γ can be viewed as a rate relaxation.

B. Service Curves

We use the concept of a service curve [6] to describe a lower bound on the service available to a flow at a node or sequence of nodes. A node offers a *deterministic service curve* \mathcal{S} , if the input-output relationship of traffic is such that for all $t \geq 0$

$$D(t) \geq \inf_{0 \leq s \leq t} \{A(s) + \mathcal{S}(t - s)\}. \quad (6)$$

The term on the right-hand side is referred to as min-plus convolution of A and \mathcal{S} , and denoted by ‘ $A * \mathcal{S}$ ’. Simple examples of service curves describe constant rate links, with $\mathcal{S}(t) = Ct$, and latency-rate servers, with $\mathcal{S}(t) = [r(t - d)]_+$, for suitable non-negative constants C , r , and d .

A probabilistic extension of this concept leads to the formulation of a *statistical service curve* which satisfies for all $t \geq 0$ that

$$\mathbb{P}\left(D(t) < A * \mathcal{S}(t; \sigma)\right) \leq \varepsilon(\sigma), \quad (7)$$

where the violation probability and the service curve may depend on one or more parameters indicated by σ .

C. Performance Bounds

The network calculus seeks to provide upper bounds on the delay, backlog, and output burstiness for a flow at a node or network path. For arrival and departure functions A and D , the backlog $B(t)$ and the delay $W(t)$ at time t are given by

$$\begin{aligned} B(t) &= A(t) - D(t), \\ W(t) &= \inf\{s \geq 0 \mid D(t + s) \geq A(t)\}. \end{aligned}$$

The backlog B and the delay W can be depicted, respectively, as the vertical and horizontal distance between the graphs of the arrival and departure functions. We use the term *output burstiness* to characterize a bound on the output process D . Bounds on the above metrics can be formulated using the min-plus deconvolution operator ‘ \oslash ’, which for any non-negative functions f and g is defined as $f \oslash g(t) = \sup_{s \geq 0} \{f(t + s) - g(s)\}$. For deterministic arrival and service characterizations, given by Eqs. (1) and (6), such bounds can be found in [11]. Below we provide the corresponding expressions for statistical arrival envelopes and service curves [4]. The service curve definition in [4] is not identical to that in Eq. (7). At the same time, the proof of the bounds are close enough to be omitted here.

Given a flow with a statistical sample path envelope with $\mathcal{G}(t; \sigma_g)$ and $\varepsilon_g(\sigma_g)$, and a statistical service curve with $\mathcal{S}(t; \sigma_s)$ and $\varepsilon_s(\sigma_s)$. Then with bounding function

$$\varepsilon(\sigma_g, \sigma_s) = \varepsilon(\sigma_g) + \varepsilon(\sigma_s) ,$$

the following bounds hold for all s, t with $0 \leq s \leq t$.

- **OUTPUT BURSTINESS:**

$$\mathbb{P}(D(s, t) > \mathcal{G} \circledast \mathcal{S}(t - s; \sigma_g, \sigma_s)) \leq \varepsilon(\sigma_g, \sigma_s) ;$$

- **BACKLOG:**

$$\mathbb{P}(B(t) > \mathcal{G} \circledast \mathcal{S}(0; \sigma_g, \sigma_s)) \leq \varepsilon(\sigma_g, \sigma_s) ;$$

- **DELAY:**

$$\mathbb{P}(W(t) > d(\sigma_g, \sigma_s)) \leq \varepsilon(\sigma_g, \sigma_s) ;$$

where

$$d(\sigma_g, \sigma_s) = \min \{s \mid \forall t \geq 0 : \mathcal{S}(t + s; \sigma_s) \geq \mathcal{G}(t; \sigma_g)\} .$$

Setting $\varepsilon(\sigma_g, \sigma_s) = 0$ yields the deterministic worst-case bounds of the deterministic network calculus.

D. A New Convolution Theorem

The main advantage of working with service curves is the ability to express the end-to-end service available on a network path as a min-plus convolution of the service available at each node of the path. In the deterministic network calculus, if each node on a path of H nodes offers a service curve of \mathcal{S}_h ($1 \leq h \leq H$), a service curve for the entire path is given by $\mathcal{S}_{\text{net}} = \mathcal{S}_1 * \mathcal{S}_2 * \dots * \mathcal{S}_H$. We refer to \mathcal{S}_{net} as *network service curve*.

Finding a network service curve in a statistical setting is difficult, and generally requires to introduce additional assumptions [10]. A main result of this paper is a novel convolution theorem for a statistical network service curve, which will enable us to create improved statistical end-to-end bounds for many scheduling algorithms.

Theorem 1 (Network Service Curve) Consider a flow passing through a sequence of H nodes. Assume that at each node, the flow receives a statistical service curve $\mathcal{S}_h(t; \sigma_h)$ with bounding function $\varepsilon_h(\sigma_h)$, where σ_h is a single non-negative parameter. For each $h < H$, assume that the bounding functions satisfy the tail estimate

$$\int_0^\infty \varepsilon_h(y) dy < \infty ,$$

and that the service curves \mathcal{S}_h are non-increasing in σ_h with

$$\mathcal{S}_h(t; \sigma_h + y) \geq [\mathcal{S}_h(t; \sigma_h) - y]_+ . \tag{8}$$

Let $\tau_1, \dots, \tau_{H-1}$ and $\gamma_1, \dots, \gamma_{H-1}$ be arbitrary positive parameters. Then, for the entire path, the flow is guaranteed the statistical service curve

$$\mathcal{S}_{\text{net}}(t; \sigma_1, \dots, \sigma_H) = \left[\mathcal{S}_1(\cdot; \sigma_1) * \dots * \mathcal{S}_H(\cdot; \sigma_H) * \delta_{\sum_{h<H} \tau_h}(t) - \sum_{h<H} \gamma_h t \right]_+, \quad (9)$$

with bounding function

$$\varepsilon_{\text{net}}(\sigma_1, \dots, \sigma_H) = \varepsilon_H(\sigma_H) + \sum_{h<H} \frac{1}{\gamma_h \tau_h} \int_{\sigma_h}^{\infty} \varepsilon_h(y) dy.$$

In the theorem, the function δ_a is defined as

$$\delta_a(t) = \begin{cases} 0, & t \leq a, \\ \infty, & t > a, \end{cases} \quad (10)$$

for a constant $a \geq 0$. The min-plus convolution of a function f with δ_a corresponds to a right-shift of the function by a , that is, $f * \delta_a(t) = f(t - a)$.

In case of a deterministic bound, obtained by first setting $\varepsilon_h(\sigma) = 0$ for all $\sigma > \sigma_h$, and then taking $\tau_h = \gamma_h = 0$, we recover the classical deterministic convolution result for composing a network service curve.

The theorem above generalizes a convolution theorem from [4], by allowing more freedom in the construction of service curves. In [4], service curves have the specific form $\mathcal{S}_h(t; \sigma_h) = [S_h(t) - \sigma_h]_+$, but [4, Theorem 1] performs the convolution of service curves only for the functions $\mathcal{S}_h(t)$, without the σ_h 's. In contrast, the σ_h 's enter into the convolution of Eq. (9). While carrying these parameters increases the complexity of the convolution, it yields network service curves that can more closely describe the behavior of certain schedulers.

In applications, we will make specific choices for τ_h and γ_h that simplify the computation. We will also choose $\sigma_1, \dots, \sigma_H$ as explicit functions of a single parameter σ . The impact of the new convolution on end-to-end delay computations is evaluated in Sec. VI.

Proof: We first show without using the assumption in Eq. (8), that a network service curve is given by

$$S_{1\dots H}(t; \sigma_1, \dots, \sigma_H) = \inf_{x_1, \dots, x_H} \left\{ S_H(x_H; \sigma_H) + \sum_{h<H} S_h \left(x_h; \sigma_h + \gamma_h \sum_{k>h} (\tau_{k-1} + x_k) \right) \right\}, \quad (11)$$

where the infimum ranges over the set $\left\{ x_1, \dots, x_H \geq 0 \mid \sum_{h \leq H} x_h + \sum_{h < H} \tau_h \leq t \right\}$, and the bounding function is as in the statement of the theorem. The claim follows from there by observing that $\sum_{k>h} (\tau_{k-1} + x_k) \leq t$ in the argument of S_h , and then using Eq. (8).

To establish Eq. (11), we proceed by induction. For $H = 1$, there is nothing to show. For $H = 2$ nodes, we need to show that

$$S_{12}(t; \sigma_1, \sigma_2) = \inf_{x_1 + x_2 \leq t - \tau_1} \left\{ S_2(x_2; \sigma_2) + S_1 \left(x_1; \sigma_1 + \gamma_1(\tau_1 + x_2) \right) \right\}$$

satisfies

$$P\{D_2(t) < A_1 * S_{12}(t; \sigma_1, \sigma_2)\} \leq \varepsilon_2(\sigma_2) + \frac{1}{\gamma_1 \tau_1} \int_{\sigma_1}^{\infty} \varepsilon_1(y) dy + \varepsilon_2(\sigma_2) =: \varepsilon_{12}(\sigma_1, \sigma_2). \quad (12)$$

To see this, note that for every sample path where $D_2(t) \geq A_2 * S_2(t; \sigma_2)$, there exists a point $x_2 \leq t$ such that

$$D_2(t) \geq A_2(t - x_2) + S_2(x_2; \sigma_2). \quad (13)$$

If, additionally,

$$\forall x \leq t \exists x_1 \leq x : D_1(t - x) \geq A_1(t - x_1 - x - \tau_1) + S_1(x_1; \sigma_1 + \gamma_1(x + \tau_1)), \quad (14)$$

then, since $A_2 = D_1$, it follows that

$$D_2(t) \geq A_1 * S_{12}(t; \sigma_1, \sigma_2). \quad (15)$$

We claim that the event in Eq. (15) occurs at least with probability $1 - \varepsilon_{12}(\sigma_1, \sigma_2)$. By definition, the probability that the event in Eq. (13) fails is bounded by

$$P(D_2(t) < A_2 * S_2(t; \sigma_2)) \leq \varepsilon_2(\sigma_2).$$

To estimate the event in Eq. (14), we discretize the time variable x by choosing an integer $k \geq 1$ such that $(k - 1)\tau_1 \leq x < k\tau_1$. If Eq. (14) fails, then the monotonicity of D_1 , A_1 , and S_1 implies that there exists a $k \geq 1$ such that

$$D_1(t - k\tau_1) < A_1(t - x_1 - k\tau_1) + S_1(x_1; \sigma_1 + \gamma_1\tau_1).$$

We conclude with the union bound that Eq. (14) fails with probability at most

$$\begin{aligned} & \sum_{k=0}^{\infty} P\left(D_1(t - k\tau_1) < \inf_{x_1 \leq t - k\tau_1} \{A_1(t - x_1 - k\tau_1) + S_1(x_1; \sigma_1 + k\gamma_1\tau_1)\}\right) \\ & \leq \sum_{k=1}^{\infty} \varepsilon_1(\sigma_1 + k\gamma_1\tau_1) \\ & \leq \frac{1}{\gamma_1\tau_1} \int_{\sigma_1}^{\infty} \varepsilon_1(y) dy. \end{aligned}$$

Eq. (12) follows with another application of the union bound.

For the inductive step, let $H > 2$, and assume that Eq. (11) holds for every sequence of $H - 1$ nodes. Given H nodes, we first apply the inductive hypothesis to obtain for the last $H - 1$ nodes the service curve

$$S_{2\dots H}(t; \sigma_2, \dots, \sigma_H) = \inf_{x_2, \dots, x_H} \left\{ S_H(x_H; \sigma_H) + \sum_{1 < h < H} S_h\left(x_h; \sigma_h + \gamma_h \sum_{k > h} (\tau_{k-1} + x_k)\right) \right\},$$

where the infimum ranges over $\{x_2, \dots, x_H \geq 0 \mid \sum_{h \leq H} x_h + \sum_{h < H} \tau_h \leq t\}$, and the bounding function is

$$\varepsilon_{2\dots H}(\sigma_2, \dots, \sigma_H) = \varepsilon_H(\sigma_H) + \sum_{1 < h < H} \frac{1}{\gamma_h \tau_h} \int_{\sigma_h}^{\infty} \varepsilon_h(y) dy,$$

and then use the $H = 2$ case with $S_{2\dots H}$ and $\varepsilon_{2\dots H}$ in place of S_2 and ε_2 . ■




Fig. 1

BUFFERED LINK WITH FIXED CAPACITY C AND CROSS TRAFFIC.

III. NETWORKS OF Δ -SCHEDULERS

In this section, we obtain a statistical network service curve that describes the service available on an end-to-end path of network links, where each link realizes a packet scheduling algorithm. To characterize link schedulers, we use a general description of scheduling algorithms, referred to as Δ -scheduling, recently developed in [14].

A. Δ -Schedulers

Consider a buffered link with a fixed capacity of rate C and without buffer size constraints as shown in Fig. 1. We are interested in describing the performance of a (through) flow with arrival function A_0 at this link. The flow experiences cross traffic from a flow with arrival function A_c . We do not assume independence of through and cross traffic. On the contrary, the traffic can be correlated in an adversarial fashion. A scheduling algorithm at the link determines the order of transmission of backlogged traffic.

We consider the class of work-conserving scheduling algorithms whose operation can be entirely described by a set of constants Δ_{ij} where i and j indicate flow indices. For two backlogged packets p_i and p_j from flows i and j , the scheduler selects p_i over p_j for transmission, if p_j arrives more than Δ_{ij} time units after p_i . Schedulers that can be described in this fashion are referred to as Δ -schedulers. Special cases are FIFO ($\Delta_{ij} = 0, \forall i, j$) and Static Priority ($\Delta_{ij} = \infty$, if the priority of flow i is lower than that of flow j , and $\Delta_{ij} = -\infty$, if the priority of flow i is higher than that of flow j). A Δ -scheduler with non-zero but finite values of Δ_{ij} can be implemented by an algorithm that tags each packet arrival with a timestamp consisting of the sum of the arrival time and a flow-specific value (e.g., a target delay), and transmits packets in the order of timestamps. Here, the value of Δ_{ij} is the difference of the tags assigned to through and cross traffic. For example, the Earliest-Deadline-First can be realized by setting $\Delta_{ij} = d_i^* - d_j^*$, where d_i^* and d_j^* are the a priori delay bounds of flows i and j .

For any Δ -scheduler operating at a buffered link with capacity C where the cross traffic is bounded by a statistical sample path envelope given by $\mathcal{G}_c(t; \sigma)$ and $\varepsilon_c(\sigma)$, it was shown in [14] that, for any choice of $\theta \geq 0$, the function

$$\mathcal{S}(t; \sigma) = \left[Ct - \mathcal{G}_c(t - \theta + \Delta(\theta); \sigma) \right]_+ I_{t > \theta}, \quad (16)$$

is a statistical service curve with bounding function $\varepsilon_c(\sigma)$, and I_{expr} denotes the indicator function, with $I_{\text{expr}} = 1$ if the expression 'expr' is true, and $I_{\text{expr}} = 0$ otherwise. Also, $\Delta(\theta) = \min\{\theta, \Delta\}$, where Δ




Fig. 2

NETWORK PATH WITH CROSS TRAFFIC AT EACH LINK.

is the constant for the precedence of the through traffic with arrival function A_0 relative to the cross traffic with arrival function A_c . (Note: We denote $\Delta = \Delta_{0c}$ since we only work with the constant for arrivals A_0 relative to A_c .) For a single node, the service curve characterization from Eq. (16) was shown to yield a tight delay analysis, provided that cross traffic and through traffic are characterized by concave deterministic sample path envelopes [14].

B. A Network Service Curve for Δ -Schedulers

We next apply Theorem 1 to construct a network service curve for Δ -schedulers along a network path as shown in Fig. 2, where the through flow traverses a tandem of H buffered links with Δ -scheduling, and with cross traffic. The arrivals of the cross flow and the Δ -scheduler can be different at each link. We use Δ^h to denote the constant of the Δ -scheduler at the h -th node. Arrivals to the cross flow at the h -th node, A_c^h , conform to an EBB model, as given by Eq. (4), with rate ρ_h and exponentially decaying violation probability $\varepsilon_h(\sigma) = M_h e^{-\alpha_h \sigma}$.

Inserting the statistical sample-path envelope for an EBB cross flow from Eq. (5) into Eq. (16), we obtain for the through flow at node h the service curve and the bounding function

$$\begin{aligned} \mathcal{S}_h(t; \sigma_h) &= \left[C_h t - \left((\rho_h + \gamma)(t - \theta_h + \Delta^h(\theta_h)) + \sigma_h \right) I_{t > \theta_h - \Delta^h} \right]_+ I_{t > \theta_h}, \\ \varepsilon_h(\sigma_h) &= M_h e \left(1 + \frac{\rho_h}{\gamma} \right) e^{-\alpha_h \sigma_h}. \end{aligned}$$

Here, $\gamma > 0$ is an arbitrary constant. To simplify the computation of the convolution, we replace this service curve with the lower bound

$$\mathcal{S}_h(t; \sigma_h) = \left[C_h t - [(\rho_h + \gamma)(t - \theta_h + \Delta^h(\theta_h)) + \sigma_h]_+ \right]_+ I_{t > \theta_h}. \quad (17)$$

The replacement comes without loss for $\Delta^h \geq 0$, and a reduction by at most $(C_h - \rho_h)[\Delta^h]_-$ for $\Delta^h < 0$, where we use $[x]_- = \max\{-x, 0\}$. The key point is that the right hand side of Eq. (17) is concave and strictly increasing in t as soon as it is positive. We apply Theorem 1 with the free parameters set to $\gamma_h = \gamma$ and $\tau_h = \frac{\alpha_h^{-1}}{\sum_{k < H} \alpha_k^{-1}} \tau_{\text{net}}$ for $h = 1, \dots, H - 1$, where τ_{net} and γ are chosen below. By Eq. (9),

$$\begin{aligned} \mathcal{S}_{\text{net}}(t; \sigma) &= [\mathcal{S}_1 * \dots * \mathcal{S}_H * \delta_{\tau_{\text{net}}}(t) - (H - 1)\gamma t]_+ \\ &\geq [\mathcal{S}_1 * \dots * \mathcal{S}_H - (H - 1)\gamma t]_+ * \delta_{\tau_{\text{net}}}(t) - (H - 1)\gamma \tau_{\text{net}}, \end{aligned} \quad (18)$$

where we use σ to mean $\sigma = (\sigma_1, \dots, \sigma_H)$. To evaluate the term inside the square brackets in the second line, we may equivalently convolve the functions $[\mathcal{S}_h(t; \sigma) - (H-1)\gamma t]_+$. For each of these factors, we write

$$[\mathcal{S}_h(t; \sigma_h) - (H-1)\gamma t]_+ = \left[(C_h - (H-1)\gamma)t - [(\rho_h + \gamma)(t - \theta_h + \Delta^h(\theta_h)) + \sigma_h]_+ \right]_+ I_{t > \theta_h}. \quad (19)$$

If we choose θ_h large enough so that

$$\underbrace{(C_h - (H-1)\gamma)\theta_h - [(\rho_h + \gamma)\Delta^h(\theta_h) + \sigma_h]_+}_{=: U_h} \geq 0, \quad (20)$$

then we can re-write Eq. (19) as

$$[\mathcal{S}_h(t; \sigma_h) - (H-1)\gamma t]_+ = \tilde{\mathcal{S}}_h * \delta_{\theta_h}(t),$$

with

$$\tilde{\mathcal{S}}_h(t; \sigma_h) = \min \{ (C_h - (H-1)\gamma)(t + \theta_h), (C_h - \rho_h - H\gamma)t + U_h \} I_{t > 0}. \quad (21)$$

Since each of the $\tilde{\mathcal{S}}_h$ is concave and strictly increasing in t , their convolution is given by the pointwise minimum

$$\tilde{\mathcal{S}}_{\text{net}}(t; \sigma) := \tilde{\mathcal{S}}_1 * \dots * \tilde{\mathcal{S}}_H(t; \sigma) = \min_{h=1, \dots, H} \{ \tilde{\mathcal{S}}_h(t; \sigma_h) \}.$$

Using the associativity and commutativity of the min-plus convolution, we obtain from Eq. (18) the network service curve

$$\mathcal{S}_{\text{net}}(t; \sigma) \geq (\tilde{\mathcal{S}}_{\text{net}} * \delta_{\tau_{\text{net}} + \sum_{h=1}^H \theta_h})(t) - (H-1)\gamma \tau_{\text{net}}. \quad (22)$$

The bounding function is computed with Theorem 1 as

$$\varepsilon_{\text{net}}(\sigma) = M_H e \left(1 + \frac{\rho_H}{\gamma} \right) e^{-\alpha_H \sigma_H} + \frac{\sum_{h < H} \alpha_h^{-1}}{\gamma \tau_{\text{net}}} \sum_{h=1}^{H-1} M_h e \left(1 + \frac{\rho_h}{\gamma} \right) e^{-\alpha_h \sigma_h}.$$

A probabilistic network service curve for a cascade of Δ -schedulers has been computed in [14] by applying a convolution theorem for service curves from [4]. This required to bring the service curves into the form $\mathcal{S}_h(t; \sigma_h) = [\mathcal{S}_h(t) - \sigma_h]_+$, which resulted in a pessimistic service curve of the form

$$\mathcal{S}_h(t; \sigma_h) = \left[C_h t - (\rho_h + \gamma)[t - \theta_h + \Delta^h(\theta_h)]_+ - \sigma_h \right]_+ I_{t > \theta_h}.$$

While the difference in the placement of σ_h compared to Eq. (17) may appear subtle, the evaluation in Sec. VI will show that the service curve together with the new convolution theorem (Theorem 1) results in substantially improved bounds.

IV. OUTPUT, BACKLOG, AND DELAY BOUNDS

In this section, we apply the performance bounds (from Sec. II-C) with the network service curve for Δ -schedulers (from Sec. III-B) to a through flow that traverses a network as shown in Fig. 2. We will make choices for the free parameters θ_h appearing in Eq. (16), so that we can obtain explicit bounds on the distribution of the output burstiness from the last node, the end-to-end backlog, and the end-to-end delay. The bounds apply to both deterministic (worst-case) and statistical assumptions.

Consider the same network as described in Sec. III. The through traffic is an EBB flow with arrival function $A_0 \sim (\rho_0, \alpha_0, M_0)$. We assume that the stability condition

$$\rho_0 < \min_{h=1, \dots, H} \{C_h - \rho_h\}$$

holds, and choose γ such that

$$0 < H\gamma < \min_{h=1, \dots, H} \{C_h - \rho_h\} - \rho_0.$$

To formulate the performance bounds, we need to introduce some more notation. Let

$$\begin{aligned} \alpha_{\text{net}} &= \left(\sum_{h=0}^H \alpha_h^{-1} \right)^{-1}, C_{\text{net}} = \min_{h=1, \dots, H} \{C_h\}, \tau_{\text{net}} = \frac{1}{\alpha_{\text{net}} C_{\text{net}}}, \\ M_{\text{net}} &= M_0 e \left(1 + \frac{\rho_0}{\gamma} \right) + M_H e \left(1 + \frac{\rho_H}{\gamma} \right) + \frac{C_{\text{net}}}{\gamma} \sum_{h=1}^{H-1} M_h e \left(1 + \frac{\rho_h}{\gamma} \right), \end{aligned} \quad (23)$$

and choose

$$\sigma_h = \frac{\alpha_{\text{net}}}{\alpha_h} \sigma, \quad h = 0, \dots, H.$$

For any given value of σ , the smallest value of θ_h that makes $\tilde{S}_h(t) > 0$ for all $t > 0$ (see Eq. (21)) is given by

$$\theta_h^* = \min \left\{ \frac{\sigma_h}{C_h - \rho_h - H\gamma}, \frac{[\sigma_h + (\rho_h + \gamma)\Delta^h]_+}{C_h - (H-1)\gamma} \right\}, \quad (24)$$

and the corresponding value for U_h is given by

$$U_h^* = [\sigma_h + (\rho_h + \gamma)\Delta^h]_-. \quad (25)$$

For the network in Fig. 2, let D_{net} denote the output of the through traffic from the last (H -th) node. Let $B_{\text{net}}(t) = D_{\text{net}}(t) - A_0(t)$ and $W_{\text{net}}(t)$, respectively, denote the total backlog and delay of the through traffic in the network of H nodes.

Theorem 2 (Closed-form Bounds) With the above notation and definitions, the following bounds hold.

- **OUTPUT BURSTINESS:**

$$\mathbb{P}(D_{\text{net}}(s, t) - (\rho_0 + \gamma)(t - s) > b_{\text{net}}(\sigma)) \leq M_{\text{net}} e^{-\alpha_{\text{net}} \sigma},$$

where

$$b_{\text{net}}(\sigma) = (\rho_0 + H\gamma)\tau_{\text{net}} + \sigma_0 + (\rho_0 + \gamma) \sum_{h=1}^H \theta_h^* ;$$

- END-TO-END BACKLOG:

$$\mathbb{P}(B_{\text{net}}(t) > b_{\text{net}}(\sigma)) \leq M_{\text{net}} e^{-\alpha_{\text{net}}\sigma} ;$$

- END-TO-END DELAY:

$$\mathbb{P}(W_{\text{net}}(t) > d_{\text{net}}(\sigma)) \leq M_{\text{net}} e^{-\alpha_{\text{net}}\sigma} ,$$

where

$$d_{\text{net}}(\sigma) = \tau_{\text{net}} + \max_{h=1, \dots, H} \left\{ \max \left\{ \frac{\sigma_0 + (H-1)\gamma\tau_{\text{net}}}{C_h - (H-1)\gamma}, \frac{\sigma_0 + (H-1)\gamma\tau_{\text{net}} - U_h^*}{C_h - \rho_h - H\gamma} \right\} \right\} + \sum_{h=1}^H \theta_h^* .$$

The output bound shows that the departing traffic satisfies the EBB model. In the expressions for $b_{\text{net}}(\sigma)$ and $d_{\text{net}}(\sigma)$, the second term quantifies the contribution of the through flow's burstiness to the delay at the bottleneck link, and the last term sums up the contributions of the burstiness of the cross flows. On long paths, i.e., for H large, the last term will dominate the second term.

An important observation is that all bounds increase with Δ^h , and that the difference between end-to-end delay and backlog bounds for different schedulers grows linearly with the number of nodes. This clearly indicates that the impact of the scheduling algorithm on end-to-end performance bounds does not diminish on long network paths for these schedulers.

In a discrete-time setting (where $t = 0, 1, \dots$), the bounds in the theorem hold with $\tau_{\text{net}} = 0$, and with Eq. (23) replaced by

$$M_{\text{net}} = \frac{M_0}{1 - e^{-\alpha_0\gamma}} + \frac{M_H}{1 - e^{-\alpha_H\gamma}} + \sum_{h=1}^{H-1} \frac{M_h}{(1 - e^{-\alpha_h\gamma})^2} .$$

Proof: We insert the network service curve from Theorem 1 into the performance bounds from Sec. II-C. Since $\alpha_h\tau_h \leq 1/C_{\text{net}}$, the bounding function evaluates to $\varepsilon_0(\sigma_0) + \varepsilon_{\text{net}}(\sigma) = M_{\text{net}} e^{-\alpha_{\text{net}}\sigma}$.

The backlog bound is given by $b_{\text{net}}(\sigma) = \mathcal{G}_0 \circledast \mathcal{S}_{\text{net}}(0)$. For $\theta_h \geq \theta_h^*$, we compute this deconvolution as

$$\begin{aligned} b_{\text{net}}(\sigma) &= \mathcal{G}_0 \circledast \left[\left(\tilde{\mathcal{S}}_{\text{net}} * \delta_{\tau_{\text{net}} + \sum_{h=1}^H \theta_h} \right) - (H-1)\gamma\tau_{\text{net}} \right]_+ (0) \\ &\leq \mathcal{G}_0 \left(\tau_{\text{net}} + \sum_{h=1}^H \theta_h \right) + (H-1)\gamma\tau_{\text{net}} . \end{aligned} \quad (26)$$

In the first line, we have used our representation for the network service curve from Eq. (22). The key step is the second line, where we use that $\tilde{\mathcal{S}}_{\text{net}}$ is concave and \mathcal{G}_0 is convex for $t > 0$, and that $\mathcal{G}_0(t) < \mathcal{S}_{\text{net}}(t)$ for t is sufficiently large by the stability condition and our choice of γ . Plugging in $\theta_h = \theta_h^*$, we obtain the desired backlog bound. The same computation gives the bound on the output burstiness.

We turn to the delay bound. For $\theta_h \geq \theta_h^*$, the delay bound from Theorem 2 is given by

$$d_{\text{net}}(\sigma) = \min\{d \geq 0 \mid \forall t \geq 0 : \mathcal{S}_{\text{net}}(t+d; \sigma) \geq \mathcal{G}_0(t; \sigma_0)\} .$$

By Eq. (22), and using as above the concavity of $\tilde{\mathcal{S}}_{\text{net}}$ and the stability assumption, we arrive at

$$d_{\text{net}}(\sigma) = \tau_{\text{net}} + \sum_{h=1}^H \theta_h + \min \left\{ X \geq 0 \mid \tilde{\mathcal{S}}_{\text{net}}(X) \geq \sigma_0 + (H-1)\gamma\tau_{\text{net}} \right\}. \quad (27)$$

Expressing the function $\tilde{\mathcal{S}}_{\text{net}}$ in terms of the $\tilde{\mathcal{S}}^h$, we see that the minimum is assumed for

$$X = \max_{h=1, \dots, H} \left\{ \max \left\{ \frac{\sigma_0 + (H-1)\gamma\tau_{\text{net}}}{C_h - (H-1)\gamma} - \theta_h, \frac{\sigma_0 + (H-1)\gamma\tau_{\text{net}} - U_h}{C_h - \rho_h - H\gamma} \right\} \right\}.$$

Inserting $\theta_h = \theta_h^*$ and $U_h = U_h^*$ for $h = 1, \dots, H$ completes the proof. \blacksquare

We note that $\theta_h = \theta_h^*$ is the optimal choice for the backlog and output bounds. To see this, observe that a selection in the range $0 \leq \theta_h < \theta_h^*$ decreases $\tilde{\mathcal{S}}^h$, resulting in a smaller network service curve. On the other hand, from Eq. (26) it is apparent that $b_{\text{net}}(\sigma)$ increases for $\theta_h > \theta_h^*$. We will see in the next section that the backlog bound is highly accurate, and in many cases tight.

For the delay bound in Eq. (27), the optimal choice for θ_h is not obvious, since U_h can be increased by taking $\theta_h > \theta_h^*$, which may result in a smaller value for $d_{\text{net}}(\sigma)$. The choice $\theta_h = \theta_h^*$ is clearly suboptimal when $C_h - \rho_h$ is small and the path is short, i.e., when the cross traffic consumes most of the bandwidth. In the case of a single node ($H = 1$) and deterministic arrival bounds, choosing $\theta_1 = d_{\text{net}}$ recovers known necessary and sufficient conditions for delay bounds e.g., [5]. For long paths ($H > 1$ large), it is plausible that the choice $\theta_h = \theta_h^*$, which is optimal for backlog and output, also leads to acceptable delay bounds. We will see in the next section that this is indeed the case.

Stronger delay bounds than those stated in the theorem can be obtained by solving the following optimization problem:

$$\text{minimize } d_{\text{net}}(\sigma) = \tau_{\text{net}} + X + \sum_{h=1}^H \theta_h, \quad (28)$$

subject to the constraints that $X \geq 0$ and

$$\forall h : \begin{cases} (C_h - (H-1)\gamma)(X + \theta_h) \geq \sigma_0 + (H-1)\gamma\tau_{\text{net}}, \\ (C_h - \rho_h - H\gamma)X + U_h \geq \sigma_0 + (H-1)\gamma\tau_{\text{net}}, \\ \theta_h \geq \theta_h^*, \end{cases}$$

where U_h and θ_h^* are defined in Eqs. (20) and (24). This optimization problem is convex if all Δ^h are non-positive, but convexity does not hold if $\Delta^h > 0$ for at least one node. Nevertheless, the problem is tractable, because the objective function is linear in the variables X and $\theta_1, \dots, \theta_H$. We will see in the evaluation that this optimization sometimes (i.e., for H small) improves the closed-form delay bound of Theorem 2. For FIFO networks (all $\Delta^h = 0$) and deterministic traffic ($\gamma = \tau_{\text{net}} = 0$) the optimization recovers existing deterministic end-to-end delay bound for feedforward networks in [8, 12].

V. TIGHTNESS OF END-TO-END BOUNDS

We now address the tightness of the bounds in Theorem 2. We consider the same network as in the previous section, but assume a deterministic context, where the arrivals of the through traffic and cross traffic at the h -th node satisfy the deterministic sample path envelopes $E_0(t) = \sigma_0 + \rho_0 t$ and $E_h(t) = \sigma_h + \rho_h t$ ($t > 0$), respectively. We denote a deterministic sample path envelope for the departing through traffic at the H -th node by $E_{\text{net}}(t)$. We obtain worst-case bounds by setting $\gamma = \tau_{\text{net}} = 0$ in the expressions for θ_h^* in Eq. (24) and U_h^* in Eq. (25), yielding

$$\theta_h^* = \min \left\{ \frac{\sigma_h}{C_h - \rho_h}, \frac{[\sigma_h + \rho_h \Delta^h]_+}{C_h} \right\},$$

and

$$U_h^* = [\sigma_h + \rho_h \Delta^h]_-.$$

Also, setting $\gamma = \tau_{\text{net}} = 0$ in Theorem 2, we obtain the following bounds, denoted by E_{net}^* , b_{net}^* , and d_{net}^* .

- OUTPUT BURSTINESS:

$$E_{\text{net}}^*(t) = \sigma_0 + \rho_0 \sum_{h=1}^H \theta_h^* + \rho_0 t;$$

- END-TO-END BACKLOG:

$$b_{\text{net}}^* = \sigma_0 + \rho_0 \sum_{h=1}^H \theta_h^*; \quad (29)$$

- END-TO-END DELAY:

$$d_{\text{net}}^* = \max_{h=1, \dots, H} \left\{ \max \left\{ \frac{\sigma_0}{C_h}, \frac{\sigma_0 - U_h^*}{C_h - \rho_h} \right\} \right\} + \sum_{h=1}^H \theta_h^*. \quad (30)$$

For the special cases of priority scheduling where through traffic has lower priority ($\Delta^h = +\infty$) or higher priority ($\Delta^h = -\infty$), we obtain

$$W_{\text{net}}(t) \leq \begin{cases} \frac{\sigma_0}{\min\{C_h - \rho_h\}} + \sum_{h=1}^H \frac{\sigma_h}{C_h - \rho_h}, & \Delta^h = +\infty, \\ \frac{\sigma_0}{\min\{C_h\}}, & \Delta^h = -\infty. \end{cases}$$

These bounds are known to be tight, which can be verified with elementary methods of the deterministic network calculus [11].

In the remainder of this section, we evaluate the tightness of the bounds for Δ -schedulers with finite values of Δ . For this, we consider specific adversarial sample path scenarios for through and cross traffic, and compute the maximal observed end-to-end delay and backlog. For any specific arrival scenario, let $B_{\text{max}} = \max_t \{B_{\text{net}}(t)\}$ and $W_{\text{max}} = \max_t \{W_{\text{net}}(t)\}$ denote the maximal end-to-end backlog and delay, respectively, of the through traffic. Clearly, any B_{max} and W_{max} , respectively, are lower bounds for the worst-case end-to-end backlog and delay.

Theorem 3 (Tightness of Bounds) For $-\infty < \Delta^h < \infty$, there exists a scenario of through and cross traffic that satisfies the assumptions stated above, such that

$$B_{\max} \geq \sigma_0 + \rho_0 \sum_{h=1}^H L_h, \quad (31)$$

$$W_{\max} \geq \frac{\sigma_0}{\min_h \{C_h\}} + \sum_{h=1}^H L_h, \quad (32)$$

with

$$L_h = \min \left\{ \frac{\sigma_h}{C_h - \rho_h}, \frac{[\sigma_h + \rho_h [\Delta^h]_+ - C_h [\Delta^h]_-]_+}{C_h} \right\}. \quad (33)$$

The theorem implies that the backlog bound in Eq. (29) is sharp when $\Delta^h \geq 0$ for $h = 1, \dots, H$, and, in particular, for FIFO. For $\Delta^h < 0$, observe that Eq. (33) implies that

$$0 \leq \theta_h^* - L_h \leq (1 - \frac{\rho_h}{C_h}) [\Delta^h]_-.$$

Here, the deviation $\theta_h^* - L_h$ is due to the relaxation of the service curve in Eq. (17). The delay bound, although not sharp, describes actual delays in adversarial scenarios rather well, particularly on long paths. We conjecture that an optimal choice of the θ_h would result in a delay bound that is as good as the backlog bound.

Proof: For the arrivals at the first node, we set

$$A_0(t) = E_0(t).$$

We time the cross traffic to have a maximal burst just before time $t = 0$, if $\Delta^1 \geq 0$, and just before $t = \Delta^1$, if $\Delta^1 < 0$, resulting in

$$A_c^1(t) = E_1(t + \nu - [\Delta^1]_-),$$

where $\nu > 0$ is an infinitesimally small number.

A careful analysis of the backlog from the cross flow at time $t = [\Delta^1]_+$, yields that the first bit of the through flow experiences a latency of L_1 . We can construct similar scenarios for cross traffic arrivals at subsequent nodes. For the h -th node, we time the arrival of a cross traffic burst relative to the arrival of the first bit of through traffic, by setting

$$A_c^h(t) = E_h(t + \nu - \sum_{k=1}^{h-1} L_k - [\Delta^h]_-).$$

Then, the initial latency of the through flow is increased at the h -th node by L_h . It follows that the first bit of the through flow departs from the last node at time $\sum_{h=1}^H L_h$, and so

$$B_{\max} \geq E_0 \left(\sum_{h=1}^H L_h \right),$$

proving Eq. (31). Now, consider the delay of the initial burst of the through flow. From the backlog analysis, we know that the first bit of the burst experiences a delay of $\sum_{h=1}^H L_h$. Since the rest of the burst cannot be transmitted faster than at the rate of the slowest link, we obtain Eq. (32). ■

VI. EVALUATION

We present a numerical evaluation of the end-to-end bounds derived in this paper for deterministic and statistical arrival scenarios. We consider a network as in Fig. 2, where we assume for simplicity that all nodes have the same capacity of $C = 100$ Mbps and use the same Δ -scheduler. We use a discrete-time setting with a step size of 1 ms. The traffic arrivals of a flow follow an alternation of On and Off phases, where, in the On phase, a flow transmits at rate $P = 1.5$ Mbps, and the long term average arrival rate is $\rho = 0.15$ Mbps. The total traffic load at each node is set to 600 flows with $N_0 = 10$ through flows and $N_c = 590$ cross flows, resulting in a link utilization of 90%.

For deterministic arrival scenarios, we characterize flows by a deterministic sample path envelope $E(t) = \rho t + \sigma$. We choose the value of σ such that the aggregate burstiness of all cross flows is 300 Kb, and likewise for the through flows. (Selecting equal burst sizes for aggregate through and cross traffic provides a good separation of our upper and lower bounds for worst-case end-to-end delays.) Note that the deterministic scenario does not account for statistical multiplexing, i.e., the burst size of N flows is N times the burst of one flow.

For statistical traffic, we assume that the arrivals of each flow are given by a discrete-time Markov-Modulated On-Off process, with an On state ($= 1$) and an Off state ($= 0$). The probabilities to change states are given by $p_{1 \rightarrow 0} = 0.9$ and $p_{0 \rightarrow 1} = 0.1$. The EBB characterization for the aggregate through and cross traffic can be obtained using the effective bandwidth $eb(\alpha)$ from (e.g., [3], page 292), resulting in EBB parameters for the aggregate through traffic of $\rho_0 = N_0 eb(\alpha_0)$, $M_0 = 1$, and α_0 selected to achieve high statistical multiplexing gain. The parameters for the cross traffic are chosen accordingly.

We assume that the through flows are independent of each other, and that the cross flows are independent; however, through and cross traffic is not independent. Thus, we observe statistical multiplexing for the group of through flows and the group of cross flows, but not across the groups.

For the discussion of the examples, recall that $\Delta = 0$ corresponds to FIFO, $\Delta = \infty$ to priority scheduling with low priority for through traffic, and $\Delta = -\infty$ to priority scheduling with high priority for through traffic.

A. Tightness of Delay Bounds

We evaluate the tightness of our delay bounds for the deterministic arrival scenario described above. We compute the end-to-end delay bounds of the through traffic for different Δ -schedulers as a function of the number of nodes H . We present the closed-form delay bound d_{net}^* according to Eq. (30), the bound computed with the optimization according to Eq. (28) (with $\tau_{\text{net}} = \gamma = 0$, since we are in a deterministic scenario), and the lower bound for the worst-case delay from Eq. (32).

The results are shown in Fig. 3. The closed-form delay bounds track the bounds from the optimization very closely, and are virtually identical for long paths. Both upper bounds are highly accurate compared to the lower bound on the worst-case delay. Note that the differences between the delay bounds for different scheduling algorithms increase proportionally to the number of traversed nodes. Thus, in contrast to the

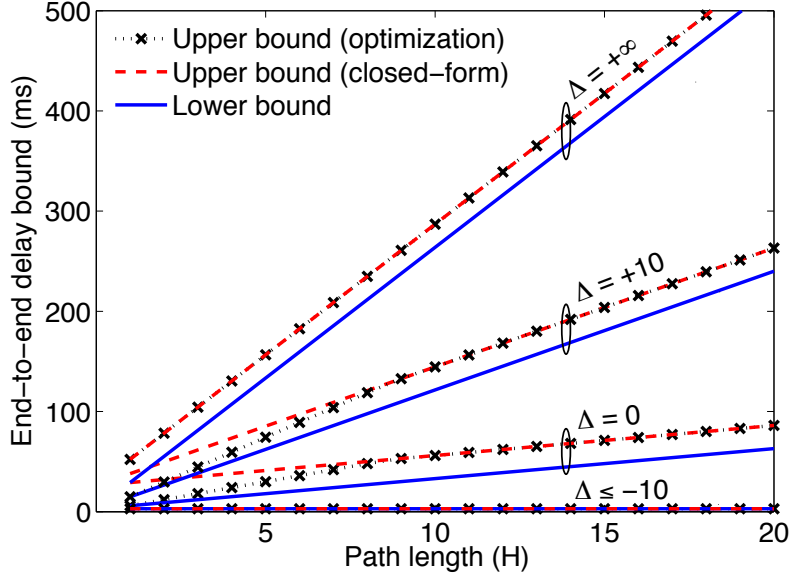


Fig. 3

DETERMINISTIC END-TO-END DELAY BOUNDS OF THROUGH TRAFFIC.

scenario discussed in the introduction of the paper, we observe that the impact of scheduling is just as strong on long paths as on a single node.

B. Improvement due to Theorem 1

For the remaining examples, we use the statistical traffic scenario. We first evaluate the benefit our convolution result, by comparing statistical delay bounds computed with the previously best convolution result from [4], to the delay bounds obtained with the new convolution result in Theorem 1. The delay bound computation for Δ -schedulers with the convolution result from [4] uses an optimization, which is described in [14]. For delay bounds computed with Theorem 1, we evaluate the closed-form expression from Theorem 2, as well as the optimization in Eq. (28). We use a fixed violation probability of $\varepsilon = 10^{-9}$.

In Fig. 4, we show the end-to-end delay bounds of the through traffic for different Δ -schedulers for increasing path length. In all cases, the closed-form delay bounds are very close to the bounds obtained by the optimization. For $\Delta = \infty$, i.e., a priority scheduler where through traffic has lowest priority, the delay bounds obtained with the new convolution theorem give the same results as [4]. Otherwise, the new convolution theorem results in dramatically improved delay bounds, especially for long paths. Our results clearly motivate the need for our new statistical network service curve expression for schedulers with $\Delta < \infty$. In fact, the delay bounds obtained with the ‘old’ convolution method for service curves show a similar growth rate for large H , thus, systematically underestimating the difference between schedulers.

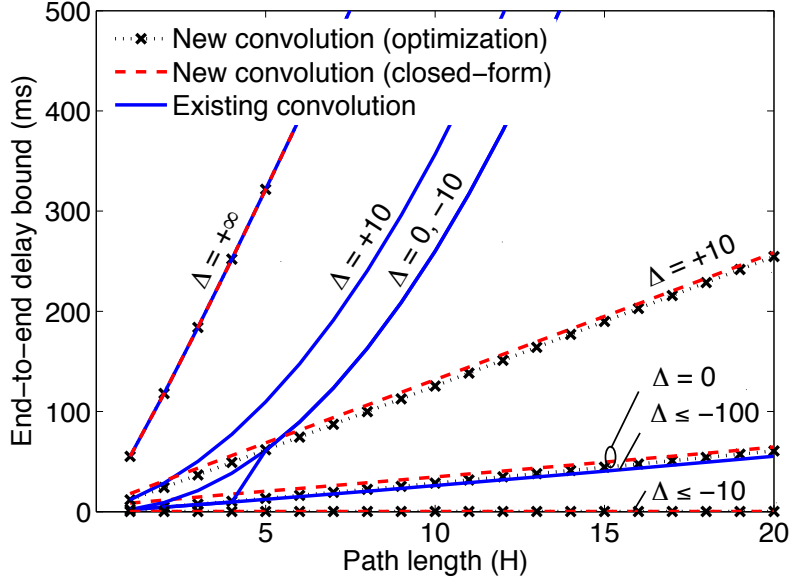


Fig. 4

STATISTICAL END-TO-END DELAY BOUNDS ($\varepsilon = 10^{-9}$).

C. Tail Bound of Backlog Distribution

Fig. 5 shows the tail distribution for the backlog of the through traffic for a path of $H = 10$ nodes. The graph shows that the backlog distribution divides the Δ -schedulers into two groups: one for large values of Δ , and one for small and negative values of Δ . In the first group, through traffic is treated effectively as low-priority traffic. The second group has a backlog distribution similar to FIFO ($\Delta = 0$), indicating that this desirable property of FIFO scheduling extends to some non-FIFO algorithms.

D. Burstiness Increase of Output Traffic

Studying the burstiness of the output traffic further corroborates the findings of the previous example. Here, we compute bounds for the output of the through traffic for different choices of Δ and H , by applying the closed-form bound for D_{net} from Theorem 2.

In Fig. 6 we show the burst size of the output bound for a violation probability $\varepsilon = 10^{-9}$ in a bar chart. (The rate values are omitted since the long-term traffic rate does not change at the output of a link.) For comparison, the leftmost bar shows the burst of the through traffic arrivals at the first node. For large values of Δ , we see that the burstiness increases sharply. For values of Δ close to zero or negative, we observe only a modest increase of the burstiness at the output, even when the number of nodes is large. As a result, the arrival process at downstream nodes is not markedly different from that at the first node. Since modest increases of output burstiness are an important condition for a network decomposition analysis, e.g., see [7, 18], our work suggests that conditions for decomposition may not be limited to FIFO networks.




Fig. 5
TAIL BOUND OF BACKLOG DISTRIBUTION ($H = 10$).

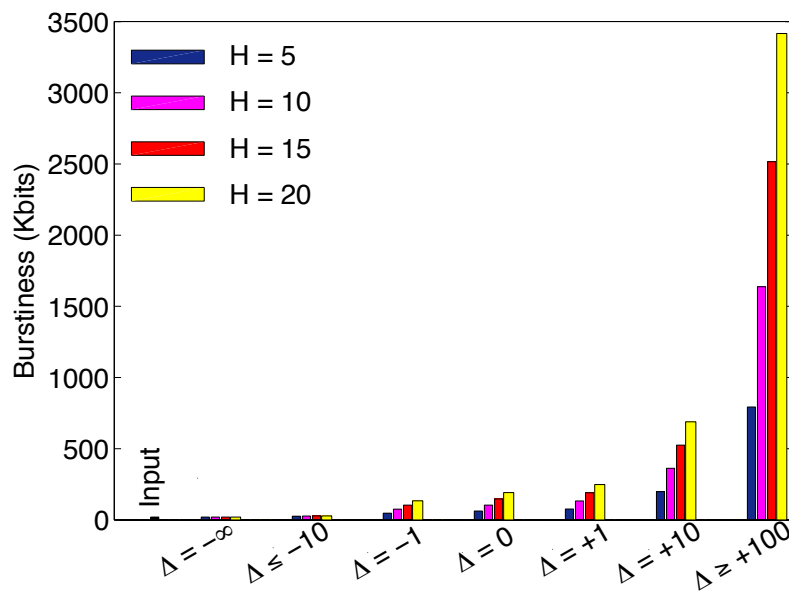


Fig. 6
BURST SIZE OF OUTPUT TRAFFIC ($\epsilon = 10^{-9}$).

VII. CONCLUSIONS

The analytical methods and closed-form expression developed in this paper vastly enhance available methods for end-to-end performance analysis in networks with deterministic and statistical traffic. With our analysis, it is now possible to make conclusive statements about the role of link scheduling in large networks. Contrary to conventional wisdom, we showed that the impact of link scheduling generally does not diminish on long network paths. We also shed light on a group of scheduling algorithms that provide service differentiation without the penalties of strict priority scheduling in terms of buffer requirements and deterioration of the burstiness of traffic.

REFERENCES

- [1] T. Bonald, A. Proutiere, and J. W. Roberts. Statistical guarantees for streaming flows using expedited forwarding. In *Proc. IEEE Infocom*, pages 1104–1112, April 2001.
- [2] C.-S. Chang. Stability, queue length, and delay of deterministic and stochastic queueing networks. *IEEE Trans. on Automatic Control*, 39(5):913–931, May 1994.
- [3] C.-S. Chang. *Performance Guarantees in Communication Networks*. Springer Verlag, 2000.
- [4] F. Ciucu, A. Burchard, and J. Liebeherr. Scaling properties of statistical end-to-end bounds in the network calculus. *IEEE Trans. on Information Theory*, 52(6):2300–2312, June 2006.
- [5] R. Cruz. A calculus for network delay, parts I and II. *IEEE Trans. on Information Theory*, 37(1):114–141, January 1991.
- [6] R. L. Cruz. Quality of service guarantees in virtual circuit switched networks. *IEEE J. on Selected Areas in Communications*, 13(6):1048–1056, August 1995.
- [7] D. Y. Eun and N. B. Shroff. Network decomposition: Theory and practice. *IEEE/ACM Trans. on Networking*, 13(3):526–539, June 2005.
- [8] M. Fidler. Extending the network calculus pay bursts only once principle to aggregate scheduling. In *Proc. QoS-IP'03*, pages 19–34, Feb. 2003.
- [9] Y. Ghiassi-Farrokhfal and J. Liebeherr. Output characterization of constant bit rate traffic in FIFO networks. *IEEE Communications Letters*, 13(8):618–620, Aug. 2009.
- [10] Y. Jiang and Y. Liu. *Stochastic Network Calculus*. Springer, 2008.
- [11] J. Y. Le Boudec and P. Thiran. *Network Calculus*. Springer, 2001.
- [12] L. Lenzini, E. Mingozzi, and G. Stea. Delay bounds for FIFO aggregates: A case study. *Computer Comm.*, 28(3):287–299, Feb. 2005.
- [13] J. Liebeherr, A. Burchard, and F. Ciucu. Non-asymptotic delay bounds for networks with heavy-tailed traffic. In *Proc. IEEE Infocom*, pages 1325–1333, March 2010.
- [14] J. Liebeherr, Y. Ghiassi-Farrokhfal, and A. Burchard. Does link scheduling matter on long paths? In *Proc. IEEE ICDCS*, pages 199–208, June 2010.
- [15] B. Melander, M. Bjorkman, and P. Gunningberg. First-come-first-served packet dispersion and implications for TCP. In *Proc. IEEE GLOBECOM*, pages 2170–2174, Nov. 2002.
- [16] D. Starobinski and M. Sidi. Stochastically bounded burstiness for communication networks. *IEEE Trans. on Information Theory*, 46(1):206–212, Jan. 2000.
- [17] O. Yaron and M. Sidi. Performance and stability of communication networks via robust exponential bounds. *IEEE/ACM Trans. on Networking*, 1(3):372–385, June 1993.
- [18] Y. Ying, R. Mazumdar, C. Rosenberg, and F. Guillemin. The burstiness behavior of regulated flows in networks. In *Proc. 4th IFIP Networking 2005*, pages 918–929, May 2005.
- [19] H. Zhang. Service disciplines for guaranteed performance service in packet-switching networks. *Proceedings of the IEEE*, 83(10):1374–1396, Oct. 1995.