

Non-asymptotic Delay Bounds for Networks with Heavy-Tailed Traffic

Jörg Liebeherr
Department of ECE
University of Toronto

Almut Burchard
Department of Mathematics
University of Toronto

Florin Ciucu
Deutsche Telekom Laboratories /
TU Berlin

Abstract—Traffic with self-similar and heavy-tailed characteristics has been widely reported in networks, yet, only few analytical results are available for predicting the delay performance of such networks. We address a particularly difficult type of heavy-tailed traffic where only the first moment can be computed, and present the first non-asymptotic end-to-end delay bounds for such traffic. The derived performance bounds are non-asymptotic in that they do not assume a steady state, large buffer, or many sources regime. Our analysis considers a multi-hop path of fixed-capacity links with heavy-tailed self-similar cross traffic at each node. A key contribution of the analysis is a probabilistic sample-path bound for heavy-tailed arrival and service processes, which is based on a scale-free sampling method. We explore how delays scale as a function of the length of the path, and compare them with lower bounds. A comparison with simulations illustrates pitfalls when simulating self-similar heavy-tailed traffic, providing further evidence for the need of analytical bounds.

I. INTRODUCTION

Traffic measurements in the 1990s provided evidence of self-similarity in aggregate network traffic [19], and heavy-tailed file sizes and bursts were found to be among the root causes [11], [28]. Since such traffic induces backlog and delay distributions whose tails decay slower than exponential, the applicability of analytical techniques based on Poisson or Markovian traffic models in network engineering has been called into question [24], thus creating a need for new approaches to teletraffic theory.

A random process X is said to have a *heavy-tailed* distribution if its tail distribution is governed by a power-law $Pr(X(t) > x) \sim Kx^{-\alpha}$, with a tail index $\alpha \in (0, 2)$ and a scaling constant K . We will consider tail indices in the range $1 < \alpha < 2$, where the distribution has a finite mean, but infinite variance. A random process X is said to be *self-similar* if a properly scaled version of the process has the same distribution as the original process. We can write this as $X(t) \sim_{dist} a^{-H} X(at)$ for any $a > 0$. The exponent $H \in (0, 1)$, referred to as the *Hurst parameter*, specifies the degree of self-similarity. (The networking literature frequently uses the weaker concept of *second-order self-similarity*. Since we will work with heavy-tailed distributions, for which higher moments are not available, we use the more general definition of self-similarity.) We refer to a process as heavy-tailed self-similar if it satisfies both criteria.

A performance analysis of networks with heavy-tailed self-similar traffic or service, where no higher moments are avail-

able, is notoriously hard, especially an analysis of a network path across multiple nodes. Single node queueing systems with heavy-tailed processes have been studied extensively [6]. However, there exist only few works that can be applied to analyze multi-node paths; these works generally consider an asymptotic regime with large buffers, many sources, or in the steady state. Tail asymptotics for multi-node networks have been derived for various topologies, such as feedforward networks [15], cyclic networks [2], tandem networks with identical service times [5], and tandem networks where packets have independent service times at nodes in the more general context of stochastic event graphs [3]. However, the accuracy of asymptotic approximations can be lacking, particularly, the quality of large buffer asymptotics for heavy-tailed service distributions was found to be lacking in [1], thus motivating a performance analysis in a non-asymptotic regime.

This paper presents a non-asymptotic delay analysis for multi-node networks with heavy-tailed self-similar traffic and heavy-tailed service. We derive the bounds for a flow or flow aggregate that traverses a network path and experiences cross traffic from heavy-tailed self-similar traffic at each node. Both fluid and packetized interpretations of service are supported; in the latter case, we assume that a packet maintains the same size at each traversed node. A key contribution of this paper is a probabilistic sample-path bound for heavy-tailed self-similar arrival processes. We obtain end-to-end delay bounds that follow (up to a logarithmic correction) the same power law tail decay as asymptotic results that exist in the literature for single nodes. Finally, we show that end-to-end delays of heavy-tailed traffic and service grow polynomially with the number of nodes. For example, for a Pareto traffic source with tail index α we find that end-to-end delays grow as $O(N^{\frac{\alpha+1}{\alpha-1}} (\log N)^{\frac{1}{\alpha-1}})$ with the number of nodes N .

Our analysis follows a network calculus approach where traffic is characterized in terms of *envelope functions*, which specify upper bounds on traffic over time intervals, and service is characterized by *service curves*, which provide lower bounds on the service available to a flow [4]. An attractive feature of the network calculus is that the service available on a path can be composed from service characterizations for each node of the path. We consider a probabilistic setting that permits performance metrics to be violated with a small probability. Probabilistic extensions of the network calculus

are available for traffic with exponential tail distributions [8], distributions that decay faster than any polynomial [26], and traffic distributions with an effective bandwidth [8]. The latter two groups include certain self-similar processes, in particular, those governed by fractional Brownian motion [23], but do not extend to heavy-tailed distributions. There are also efforts for extending the network calculus to heavy-tailed distributions [12], [13], [16], [17], which are discussed in more detail in the next section.

The remainder of this paper is organized as follows. In Section II and Section III, respectively, we discuss our characterization of heavy-tailed traffic and service by appropriate probabilistic bounds. In Section IV we present our main results: (1) a sample-path envelope for heavy-tailed self-similar traffic, (2) probabilistic bounds for delay and backlog, (3) a description of the leftover capacity at a constant-rate link with heavy-tailed self-similar cross traffic, and (4) a composition result for service descriptions at multiple nodes. In Section V we discuss the scaling properties of the derived delay bounds in terms of power laws. We give brief conclusions in Section VI.

II. THE *htss* TRAFFIC ENVELOPE

In this section we present a method for characterizing heavy-tailed self-similar network traffic that permits the derivation of rigorous backlog and delay bounds. We use a continuous time model where cumulative arrivals and departures of a traffic flow or flow aggregate in a time interval $[0, t)$ are represented by left-continuous processes $A(t)$ and $D(t)$, respectively. The arrivals in the time interval $[s, t)$ are denoted by a bivariate process $A(s, t) := A(t) - A(s)$. Backlog and delay at a node are represented by $B(t) = A(t) - D(t)$ and $W(t) = \inf \{d : A(t - d) \leq D(t)\}$, respectively. When A and D are plotted as functions of time, B and W are the vertical and horizontal distance, respectively, between these functions.

A *statistical envelope* \mathcal{G} for an arrival process A is a non-random function which bounds arrivals over a time interval such that, for all $s, t \geq 0$ and for all $\sigma > 0$:

$$\Pr\left(A(s, t) > \mathcal{G}(t - s; \sigma)\right) \leq \varepsilon(\sigma), \quad (1)$$

where the function ε is a bound on the violation probability that satisfies $\varepsilon(\sigma) \rightarrow 0$ as $\sigma \rightarrow \infty$. Statistical envelopes have been developed for many different traffic types, including regulated, Markov modulated On-Off, and Gaussian self-similar traffic. A recent survey provides an overview of existing envelope concepts [21].

The computation of performance bounds requires a statistical envelope that bounds an entire sample path $\{A(s, t)\}_{s \leq t}$. We define a *statistical sample-path envelope* $\overline{\mathcal{G}}$ as a function that satisfies for all $t \geq 0$ and for all $\sigma > 0$ [9]:

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

Clearly, a statistical sample-path envelope is also a statistical envelope, but not vice versa. In fact, only few statistical envelopes (in the sense of Eq. (1)) lend themselves easily

to the development of sample-path envelopes (as in Eq. (2)). One of the earliest such envelopes is the *Exponentially Bounded Burstiness* (EBB) model [29], which requires that $\Pr(A(s, t) > r(t - s) + \sigma) \leq Me^{-a\sigma}$, for some constants M , r and a and for all $\sigma > 0$. If r corresponds to the mean rate of traffic, an EBB envelope specifies that the deviation of the traffic flow from its mean rate has an exponential decay. A sample-path bound for EBB envelopes in the sense of Eq. (2) is obtained by applying the union bound, resulting in $\overline{\mathcal{G}}(t - s; \sigma) = (\rho + \delta)t + \sigma$ for any $\delta > 0$ and $\varepsilon(\sigma) = \frac{Me^{-a\sigma}}{1 - e^{-a\delta}}$. The EBB envelope has been generalized to distributions with moments of all orders, referred to as *Stochastically Bounded Burstiness* (SBB) [26] and corresponding sample-path bounds have been developed [30]. SBB envelopes can characterize arrival processes that are self-similar, but not heavy-tailed.

We propose an extension of the EBB and SBB concepts to capture characteristics of heavy-tailed and self-similar traffic. We define a *heavy-tail self-similar (htss) envelope* as a bound that satisfies for all $\sigma > 0$ that

$$\Pr\left(A(s, t) > r(t - s) + \sigma(t - s)^H\right) \leq K\sigma^{-\alpha}, \quad (3)$$

where K and r are constants, H indicates the Hurst parameter and α is the tail index. We generally assume $\alpha \in (1, 2)$, that is, arrivals have a finite mean but possibly infinite variance, and $H \in (0, 1)$. In Eq. (3), note that $A(s, t) - r(t - s) \leq_{dist} X(t - s)$, where X is a process satisfying the self-similar property given in the introduction. The heavy-tailed violation probability is justified by the generalized central limit theorem [14], where self-similarity arises from heavy-tailed arrival processes with independent increments.

The key characteristic of the *htss* envelope is that the probability of deviating from the average rate r follows a power law. Moreover, due to self-similarity, these deviations may increase as a function of time. Since the *htss* envelope specifies a bound, it can be used to describe any type of traffic, but the characterization may be loose unless the traffic has some heavy-tailed self-similar properties. In terms of Eq. (1), the *htss* envelope is a statistical envelope with

$$\mathcal{G}(t; \sigma) = rt + \sigma t^H, \quad \varepsilon(\sigma) = K\sigma^{-\alpha}.$$

In Section IV, we will derive a sample-path envelope for the *htss* envelope, which is necessary for the computation of probabilistic upper bounds on backlog and delay of heavy-tailed self-similar traffic.

Characterizations of self-similar and heavy-tailed traffic by envelopes have been presented before by exploiting properties of α -stable processes [12], [17]. The envelope for α -stable processes in [13] takes the same form $\mathcal{G}(t; \sigma) = rt + \sigma t^H$ as the *htss* envelope, but specifies a fixed violation probability (i.e., with a constant violation probability that does not have a heavy tail). An issue with such a characterization is that it does not easily lead to sample-path envelopes. For $H = 0$, a sample-path version of Eq. (3) has been obtained in [16] by applying an a-priori bound on the backlog process of an α -stable self-similar process from [17]. However, the backlog

bound found in Eq. (24) of [17] is a lower bound (and not an upper bound) on the tail distribution of the buffer occupancy, and hence the envelope in [16] does not satisfy Eq. (2). Since sample-path envelopes for arrivals are equivalent to backlog bounds (see Section IV) in the sense that each can be used to derive the other, our sample-path envelope also yields rigorous backlog bounds for this general class of processes.

Pareto Arrival Distribution: As an example of a heavy-tailed distribution, we consider an *htss* envelope construction for a traffic source with a Pareto arrival distribution. Traffic arrivals are assumed evenly spaced at rate λ and packet sizes are described by i.i.d. Pareto random variables X_i with tail distribution

$$\Pr(X_i > x) = \left(\frac{x}{b}\right)^{-\alpha}, \quad x \geq b, \quad (4)$$

where $\alpha \in (1, 2)$. X has finite mean $E[X] = \frac{b\alpha}{\alpha-1}$ and infinite variance. Such a traffic source could be the output of a variable-bit-rate video encoder with a fixed frame rate. We will construct an *htss* envelope for the arrival function $A(t) = \sum_{i=1}^{N(t)} X_i$, where $N(t) = \lfloor \lambda t \rfloor$ denotes the number of packets which arrive by time t . This arrival process is asymptotically self-similar with a Hurst parameter of $H = 1/\alpha$. For the *htss* envelope construction of the Pareto source, we take advantage of the *generalized central limit theorem* (GCLT) [14], which states that the α -stable distribution S_α appears as the limit of normalized sums of i.i.d. random variables.

For the sum of n independent Pareto random variables X_i , the GCLT [14] yields

$$\frac{\sum_{i=1}^n X_i - nE[X]}{c_\alpha n^{\frac{1}{\alpha}}} \xrightarrow{n \rightarrow \infty} S_\alpha \quad (5)$$

in distribution, where $c_\alpha = \left(\frac{2\Gamma(\alpha)\sin\frac{\pi\alpha}{2}}{\pi}\right)^{-\frac{1}{\alpha}}$ and $\Gamma(\cdot)$ is the Gamma function. Since the GCLT is an asymptotic limit, envelopes derived with the GCLT are approximate, with higher accuracy for larger values of n . Using that $N(t) \approx \lambda t$ for suitable large values, we can write the arrival function with Eq. (5) as

$$A(t) \approx \lambda t E[X] + c_\alpha (\lambda t)^{1/\alpha} S_\alpha.$$

We can now use the standard tail estimate for normalized α -stable distributions [25], given by

$$\Pr(S_\alpha > \sigma) \sim (c_\alpha \sigma)^{-\alpha}, \quad \sigma \rightarrow \infty, \quad (6)$$

to obtain an *htss* envelope with parameters

$$r = \lambda E[X], \quad \alpha, \quad H = \frac{1}{\alpha}, \quad K \approx \lambda. \quad (7)$$

Similar techniques can yield *htss* envelopes for other heavy-tailed processes. For example, an aggregation of independent On-Off periods, where the duration of ‘On’ and ‘Off’ periods is governed by independent Pareto random variables yields an α -stable process [22] in the limit of many flows ($N \rightarrow \infty$) and large time scales ($t \rightarrow \infty$). This aggregate process is particularly interesting since depending on the order in which

the limits of N and t are taken, one obtains processes that are self-similar, but not heavy-tailed (fractional Brownian motion), processes that are heavy-tailed, but not self-similar (α -stable Lévy motion), or a general α -stable process. An approximation by an α -stable process followed by an estimation of *htss* parameters can also be performed for the M/G/ ∞ arrival model [22].

III. SERVICE GUARANTEES WITH HEAVY TAILS

We next formulate service guarantees at a network link that have a power-law decay. In the network calculus, service guarantees are expressed in terms of functions that express for a given arrival function a lower bounds on the departures. In general, a *statistical service curve* is a function $\mathcal{S}(t; \sigma)$ such that for all $t \geq 0$ and for all $\sigma > 0$

$$\Pr(D(t) < A * \mathcal{S}(t; \sigma)) \leq \varepsilon(\sigma).$$

Here,

$$A * \mathcal{S}(t; \sigma) = \inf_{s \leq t} \{A(s) + \mathcal{S}(t - s; \sigma)\}$$

denotes the min-plus convolution of the arrivals with the service curve $\mathcal{S}(t; \sigma)$, and ε is a non-increasing function that satisfies $\varepsilon(\sigma) \rightarrow 0$ as $\sigma \rightarrow \infty$.

A *heavy-tailed (ht) service curve* is a service curve of the form

$$\mathcal{S}(t; \sigma) = [Rt - \sigma]_+, \quad \varepsilon(\sigma) = L\sigma^{-\beta} \quad (8)$$

for some β with $0 < \beta < 2$ and some constant L . The rationale for not including a Hurst parameter in the definition of service guarantees is that the form of Eq. (8) facilitates the computation of service bounds over multiple nodes.

- *Service at link with cross traffic (Leftover Service):* We are interested in deriving an *ht* service curve that describes the service available to a selected flow at a constant-rate link with capacity C , where the competing traffic at the link, referred as *cross traffic*, is given by an *htss* envelope. By considering the pessimistic case that the selected flow receives a lower priority than the cross traffic, we will obtain a lower bound for the service that holds for most workconserving multiplexers [4]. Since, with this construction, the service guarantee of the selected flow consists of the capacity that is left unused by cross traffic, we refer to the service interpretation as *leftover service*. As the derivation of an *ht* service curve for leftover service requires a sample-path bound for the cross traffic, we defer the derivation to Subsection IV-B.

- *Packet-level traffic:* Even though we use consider fluid-flow traffic, the *ht* service model is highly suitable to express a packetized view of traffic with heavy-tailed packet size distributions. We now derive a service curve for packetized traffic and use it later to relate our bounds to a queuing theoretic analysis with a packet-level interpretation of traffic. For a packet-size distribution satisfying $\Pr(X > \sigma) \leq L\sigma^{-\alpha}$, we will show that a constant-rate workconserving link of capacity C provides an *ht* service curve with rate $R = C$, and compute its violation probability $\varepsilon(\sigma)$.

Denote by $X^*(t)$ the part of the packet in transmission at time t that has already been transmitted. We can view $X^*(t)$ as the current lifetime in a renewal process. It is known from the theory of renewal processes (see [18], pp. 194) that

$$\lim_{t \rightarrow \infty} Pr(X^*(t) > \sigma) = \frac{\int_{\sigma}^{\infty} Pr(X > x) dx}{E[X]}.$$

The bound on the tail probability holds for all times t , provided that we properly randomize the arrival time of the first packet after the network is started with empty queues at $t = 0$. The departures are given by

$$D(t) = \begin{cases} A(t), & \underline{t} = t, \\ A(\underline{t}) + C(t - \underline{t}) - X^*(t), & \underline{t} < t, \end{cases}$$

where \underline{t} is the beginning of the busy period of t . Set $\mathcal{S}(t; \sigma) = [Ct - \sigma]_+$. If ρ is the utilization of A as a fraction of the link rate C , then

$$\begin{aligned} Pr(D(t) < A * \mathcal{S}(t; \sigma)) &\leq \rho Pr(X^*(t) > \sigma) \\ &\leq \frac{\rho L}{(\alpha - 1)E[X]} \sigma^{-(\alpha - 1)}. \end{aligned} \quad (9)$$

IV. NETWORK CALCULUS WITH *htss* ENVELOPES

We consider a network as in Fig. 1. A flow traverses N nodes in series. Its traffic is referred to as through traffic. At each node, the through traffic is multiplexed with arrivals from competing flows, called *cross traffic*. Both through and cross traffic are described by *htss* envelopes. We will present results that, in combination, yield bounds on the end-to-end delay and backlog in a network shown in the figure.




Fig. 1. A network with cross traffic.

In our derivations, we will frequently use properties of the function $\varepsilon(\sigma) = K\sigma^{-\alpha}$ that appears in the definition of the *htss* envelope. First, we can lower the power by using that for $K\sigma^{-\alpha} \leq 1$ and $\alpha' < \alpha$

$$K\sigma^{-\alpha} \leq K^{\frac{\alpha'}{\alpha}} \sigma^{-\alpha'}. \quad (10)$$

For $\alpha > 0$, $\sigma_0 > 0$, and $K(\sigma - \sigma_0)^{-\alpha} \leq 1$, we can remove a negative shift by

$$K(\sigma - \sigma_0)^{-\alpha} \leq 2^{[\alpha - 1]_+} (K + \sigma_0^\alpha) \sigma^{-\alpha}. \quad (11)$$

We can minimize sums of such functions by

$$\begin{aligned} \min_{\sigma_1 + \dots + \sigma_n = \sigma} \sum_{j=1}^n K_j \sigma_j^{-\alpha} &= \left(\sum_{j=1}^n K_j^{\frac{1}{1+\alpha}} \right)^{1+\alpha} \sigma^{-\alpha} \\ &\leq n^\alpha \bar{K} \sigma^{-\alpha}, \end{aligned} \quad (12)$$

where $\bar{K} = \frac{1}{n}(K_1 + K_2 + \dots + K_n)$.

A. Statistical Sample-Path Envelope

To further motivate the relevance of sample-path bounds, let us consider the backlog of a flow at a workconserving link that operates at a constant rate C . The backlog at time t is given by

$$B(t) = \sup_{s \leq t} \{A(s, t) - C(t - s)\}.$$

Notice that the backlog expression depends on the entire arrival sample path $\{A(s, t)\}_{s \leq t}$. To compute an upper bound for the tail probability $Pr(B(t) > \sigma)$, in many places in the literature, in particular, in all prior works attempting a network calculus analysis with heavy-tailed traffic [12], [13], [16], [17], the tail distribution is approximated by

$$Pr(B(t) > \sigma) \approx \sup_{s \leq t} Pr(A(s, t) - C(t - s) > \sigma).$$

Note that the right hand side is generally *smaller* than the left hand side, and thus inserting into the right hand side a statistical envelope, which only satisfies Eq. (1) but not Eq. (2), does not yield an upper bound but rather an upper bound to a lower bound. The derivation of rigorous upper bounds requires a sample-path bound for the arrivals. To derive such bounds, we discretize time by setting $x_k = \tau\gamma^k$, where $\tau > 0$ and $\gamma > 1$ are constants that will be chosen below. If $t - x_k \leq s < t - x_{k-1}$, then

$$A(s, t) - C(t - s) \leq A(t - x_k, t) - Cx_{k-1}.$$

It follows that

$$B(t) \leq \sup_k \{A(t - x_k, t) - Cx_{k-1}\}.$$

If the arrivals satisfy an *htss* envelope $\mathcal{G}(t) = rt + \sigma t^H$ with $\varepsilon(\sigma) = K\sigma^{-\alpha}$, we obtain with the union bound

$$\begin{aligned} Pr(B(t) > \sigma) &\leq \sum_{k=-\infty}^{\infty} Pr(A(t - x_k, t) > \sigma + Cx_{k-1}) \\ &\leq \tilde{K} \sigma^{-\alpha(1-H)}. \end{aligned} \quad (13)$$

In the last line, we have used Lemma 3 from the appendix to evaluate the sum. The constant is given by

$$\tilde{K} = K \cdot \inf_{1 < \gamma < 1 + \mu/r} \left\{ \left(\frac{(r + \mu)/\gamma - r}{\alpha H (1 - H) \log \gamma} \right)^{-\alpha H} \frac{\gamma^{\alpha H (1 - H)}}{\alpha H (1 - H) \log \gamma} \right\}, \quad (14)$$

where $\mu = C - r$. We remark that in typical applications, $1 < \alpha < 2$ and $1/\alpha \leq H < 1$, so that $\alpha(1 - H) < 1$. This means that the backlog is almost surely finite, but cannot be expected to have finite mean.

The main technical ingredient of the above proof of the backlog bound is the discretization of time by the geometric sequence $x_k = \gamma^k \tau$. This scale-free discretization is well-suited to capture self-similar properties of traffic. Scale-free sampling is an instance of *under-sampling*, where not every time step is used in probabilistic estimates. Commonly in the literature, time is discretized by dividing it into equal units with $x_k = k\tau$. In [27], the choice is described as a general optimization problem over arbitrary sequences $x_0 \leq x_1 \leq$

$\dots \leq t$, but not applied, since all examples in [27] only optimize over τ in uniformly spaced sequences. Note that using a uniform discretization in the derivation of Eq. (13) would cause the infinite sums to become unbounded.

An immediate consequence of the backlog bound is a sample-path bound for *htss* envelopes.

Lemma 1: ht SAMPLE-PATH ENVELOPE If arrivals to a flow are bounded by an *htss* envelope

$$\mathcal{G}(t; \sigma) = rt + \sigma t^H, \quad \varepsilon(\sigma) = K\sigma^{-\alpha},$$

then, for every choice of $\mu > 0$,

$$\bar{\mathcal{G}}(t; \sigma) = (r + \mu)t + \sigma, \quad \varepsilon(\sigma) = \tilde{K}\sigma^{-\alpha(1-H)},$$

is a statistical sample-path envelope according to Eq. (2). The constant \tilde{K} is given by Eq. (14).

The proof follows immediately from Eq. (13) by replacing C with the relaxed arrival rate $r + \mu$. The *ht* sample-path envelope is reminiscent of a leaky-bucket constraint with a single burst and rate, and does not reflect the self-similar scaling of the *htss* envelope.

We note that a small modification of the proof would yield a sample-path envelope of the form

$$\bar{\mathcal{G}}(t, \sigma) = (r + \mu)t + \sigma t^H + M, \quad \varepsilon(\sigma) = L\sigma^{-\alpha},$$

which retains the self-similar scaling properties of the *htss* envelope. The constant L depends on the parameters α, H, r, μ and on the choice of $M > 0$. The reason we prefer the simpler envelope given by Lemma 1 is that it facilitates the estimation of the service provided to a flow across multiple nodes.

B. Heavy-Tailed Leftover Service Curve

With a sample-path envelope for heavy-tailed traffic at hand, we can now derive a service curve for the heavy-tailed leftover service from Section III. Assuming that the cross traffic A_c is characterized by an *htss* envelope of the form $\mathcal{G}_c(t) = r_c(t - s) + \sigma(t - s)^{H_c}$ with $\varepsilon(\sigma) = K_c\sigma^{-\alpha_c}$ where the bound on the arrival rate satisfies $r_c < C$, we will show that the through flow is guaranteed a *ht* service curve $\mathcal{S}(t; \sigma) = [Rt - \sigma]_+$ with rate $R = C - r_c - \mu$, and a violation probability $\varepsilon(\sigma)$ that can be estimated explicitly. Here, μ is a free parameter.

We denote arrivals of the through flow by A . Departures are denoted by D and D_c , respectively. Let \underline{t} be the beginning of the busy period of t at the link. Then, the aggregate departures in $(\underline{t}, t]$ satisfy $(D + D_c)(\underline{t}, t) = C(t - \underline{t})$, and departures for the cross traffic satisfy $D_c(\underline{t}, t) \leq \min\{C(t - \underline{t}), A_c(t) - A_c(\underline{t})\}$. With this we can derive

$$\begin{aligned} D(t) &\geq A(\underline{t}) + [C(t - \underline{t}) - A_c(\underline{t}, t)]_+ \\ &\geq \inf_{s \leq \underline{t}} \{A(s) + (C - r_c - \mu)(t - s)\} \\ &\quad - \sup_{s \leq \underline{t}} \{A_c(s, t) - (r_c + \mu)(t - s)\}, \end{aligned}$$

for every choice of $\mu > 0$. We obtain

$$\begin{aligned} Pr(D(t) < A * \mathcal{S}(t; \sigma)) &\leq Pr(\sup_{s \leq t} \{A_c(s, t) - (r_c + \mu)(t - s)\} > \sigma) \\ &\leq \tilde{K}_c \sigma^{-\alpha_c(1-H_c)}, \end{aligned} \quad (15)$$

where \tilde{K}_c is given by Eq. (14). This proves that $\mathcal{S}(t; \sigma) = [Rt - \sigma]_+$ is an *ht* service curve.

The description of the leftover service in Eq. (15) can be combined with Eq. (9) to characterize the leftover service available to a packetized through flow at a node. The result (which we state without proof) is that at a link that operates at rate $C > r_c$, the through flow receives a service guarantee given by the *ht* service curve

$$\mathcal{S}(t; \sigma) = [(C - r_c - \mu)t - \sigma]_+, \quad \varepsilon(\sigma) = L\sigma^{-\beta}, \quad (16)$$

where $\beta = \min\{\alpha_p - 1, \alpha_c(1 - H_c)\}$, α_p is the tail decay in Eq. (9), and $\mu > 0$ is a free parameter. The violation probability is given by

$$\varepsilon(\sigma) = \inf_{\sigma_1 + \sigma_2 = \sigma} \left\{ \tilde{K}_c \sigma_1^{-\alpha_c(1-H)} + \frac{\rho_p L_p}{(\alpha_p - 1)E[X]} \sigma_2^{-(\alpha_p - 1)} \right\},$$

where $\rho_p \leq 1$ is the utilization of the packetized part of the through traffic as a fraction of C , $E[X]$ is the average packet size, and the constant \tilde{K}_c is defined by Eq. (14) with \tilde{K}_c in place of K . The constant in Eq. (16) can be computed explicitly by first using Eq. (10) to lower the larger exponent to β , and then applying Eq. (12).

C. Single Node Delay Analysis

We next present a delay bound at a single node where arrivals are described by *htss* envelopes and service is described by an *ht* service curve.

Theorem 1: SINGLE NODE DELAY BOUND. Consider a flow that is characterized by an *htss* envelope with $\mathcal{G}(t, \sigma) = rt + \sigma(t - s)^H$ and $\varepsilon(\sigma) = K\sigma^{-\alpha}$, and that receives an *ht* service curve at a node given by $\mathcal{S}(t; \sigma) = [Rt - \sigma]_+$ and $\varepsilon(\sigma) = L\sigma^{-\beta}$. If $r < R$, then the delay W satisfies

$$Pr(W(t) > w) \leq M(Rw)^{-\min\{\alpha(1-H), \beta\}},$$

where M is a constant that depends on $\alpha, H, r, \mu = R - r$, and β .

PROOF. Let $A(t)$ and $D(t)$ denote the arrival and departures of the (tagged) flow at the node. The delay is given by

$$W(t) = \inf\{t - s \mid A(s) \leq D(t)\}.$$

Fix $\sigma_1, \sigma_2 > 0$ with $R(\sigma_1 + \sigma_2) = w$. Suppose that on a particular sample path,

$$\sup_{s \leq t - w} \{A(s, t - w) - R(t - s - w)\} \leq \sigma_1,$$

and that

$$D(t) \geq \inf_{s \leq t} \{A(s) + [R(t - s) - \sigma_2]_+\}.$$

If the infimum is assumed for some $s \leq t - w$, then

$$\begin{aligned} D(t) &\geq A(s) + R(t - s) - \sigma_2 \\ &\geq A(t - w). \end{aligned}$$

If, on the other hand, the infimum is assumed for some $s \geq t - w$, then

$$D(t) \geq A(s) \geq A(t - w)$$

by monotonicity. In both cases, we see that $W(t) \leq w$. It follows with union bound that

$$\begin{aligned} Pr(W(t) > w) &\leq Pr\left(\sup_{s \leq t-w} \{A(s, t-w) - R(t-s-w)\} > \sigma_1\right) \\ &\quad + Pr\left(D(t) < \inf_{s \leq t} [A(s) + R(t-s) - \sigma_2]_+\right) \\ &\leq \tilde{K} \sigma_1^{-\alpha(1-H)} + L \sigma_2^{-\beta}. \end{aligned} \quad (17)$$

For the first term, we have used the sample-path bound in Lemma 1 with $\mu = R - r$, and for the second term we have used the definition of ht service curves. The proof is completed by first lowering the larger of the two exponents to $\beta' = \min\{\alpha(1-H), \beta\}$, using Eq. (10), and then minimizing explicitly over the choice of σ_1 and σ_2 using Eq. (12). For the constant, this yields the estimate

$$M \leq \left\{ \tilde{K} \frac{\beta'}{(1+\beta')\alpha(1-H)} + L_s \frac{\beta'}{(1+\beta')\beta} \right\}^{1+\beta'}, \quad (18)$$

where \tilde{K} is defined by Eq. (14). \square

Example: We compute the delay experienced by a Pareto traffic source at a 100 Mbps link. The parameters are

$$\alpha = 1.6, \quad b = 150 \text{ Byte}, \quad \lambda = 75 \text{ Mbps}. \quad (19)$$

With these values, the average data unit has a size of 400 Byte, and the link utilization is 75%. The service curve is computed from Eq. (9). The reason for selecting this example (which does not have cross traffic) is that it permits a comparison with a queueing theoretic result in [7], which presents a lower bound on the quantiles of a Pareto source in a tandem network with N nodes and no cross traffic, $w_N(z)$, as

$$w_N(z) \geq \frac{(Nb)^{\alpha/(\alpha-1)}}{((\alpha-1)|\log(1-\varepsilon)|/\lambda)^{1/(\alpha-1)}}. \quad (20)$$

In Fig. 2 we show a log-log plot of the delay distribution. The graph illustrates the power-law decay for the upper bound and the lower bound from [7]. We also show the results of four simulation runs of an initially empty system with $10^6, 10^7, 10^8$ and 10^9 packets. The simulation traces indicate that the actual delays may be closer to the lower bounds. Note that the fidelity of the simulations deteriorates at smaller violation probabilities. Since even long simulation runs do not capture sufficiently many events with large delays, they violate analytical lower bounds. Note that even the simulation run of 1 billion arrivals does not maintain the power-law decay for violation probabilities below $\varepsilon = 10^{-3}$.

D. Multi-Node Delay Analysis

We turn to the computation of end-to-end delays for a complete network path. As in the deterministic version of the network calculus [4] we express the service given by all nodes on the path in terms of a single service curve, and then apply single-node delay bounds. We start with a network of two nodes. We denote by A_1 the arrivals of the analyzed flow at the first node, and by D_1 or A_2 the departures of the first node that arrive to the second node.

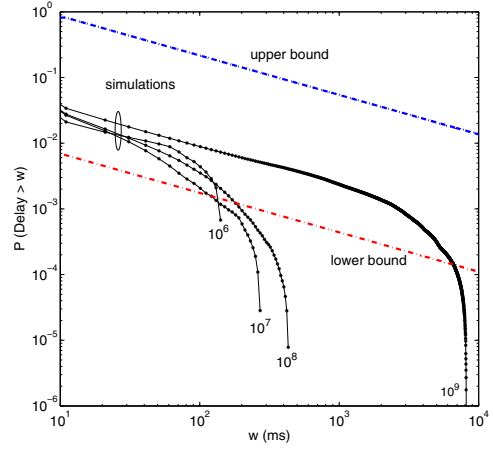


Fig. 2. Log-log plot of single-node delays for a Pareto traffic source. Upper bounds and lower bounds are compared to simulation traces with $10^6, 10^7, 10^8$ and 10^9 arrivals.

Lemma 2: CONCATENATION OF TWO ht SERVICE CURVES. Consider an arrival flow traversing two nodes in series. The first node offers an ht service curve with $S_1(t; \sigma) = [R_1 t - \sigma]_+$ and $\varepsilon_1(\sigma) = L_1 \sigma^{-\beta_1}$, and the second node offers a service curve $S_2(t; \sigma) = [R_2 t - \sigma]_+$ and an arbitrary function $\varepsilon_2(\sigma)$. Then for any $\gamma > 1$, the two nodes in sequence offer the combined service curve given by

$$\begin{aligned} S(t; \sigma) &= [\min\{R_1, R_2/\gamma\}t - \sigma]_+, \\ \varepsilon(\sigma) &= \inf_{\sigma_1 + \sigma_2 = \sigma} \left\{ \tilde{L}_1 \sigma_1^{-\beta_1} \left(\log \sigma_1 + \frac{2 - \log \tilde{L}_1}{\beta_1} \right) + \varepsilon_2(\sigma_2) \right\}, \end{aligned}$$

where $\tilde{L}_1 = \frac{L_1 2^{\max\{1, \beta_1\}}}{\log \gamma}$.

The service rate $R = \min\{R_1, R_2/\gamma\}$ in the expression for the service curve is the result of a min-plus convolution of the service curves at the individual nodes. The logarithmic term can be removed by using that for $\beta' < \beta$

$$\sigma^{-\beta} \log \sigma \leq \frac{1}{e^{(\beta - \beta')}} \sigma^{-\beta'}.$$

If the second node also offers an ht service curve, with $\varepsilon_2(\sigma) = L_2 \sigma^{-\beta_2}$, then for every choice of β with $\beta < \beta_1$ and $\beta \leq \beta_2$ there exists a constant $L = L(\beta, R_2, \gamma)$ such that $\varepsilon(\sigma) \leq L \sigma^{-\beta}$. The value of the constant L can be computed from Eqs. (10) and (12).

PROOF. We proceed by inserting the service guarantee for $D_1 = A_2$ at the first node into the service guarantee at the second node. Similar to the backlog and delay bounds, this requires an estimate for an entire sample path of the service at the first node.

Fix $t \geq 0$. We consider discretized time points $t - y_k$, where $y_0 = 0$ and $y_k = \tau + \gamma' y_{k-1}$ for some $\tau > 0$ and $\gamma' > 1$ to be chosen below. For $t - y_k \leq s < t - y_{k-1}$, we have

$$A_2(s) + [R_2(t-s) - \sigma]_+ \geq A_2(t - y_k) + [R_2 y_{k-1} - \sigma]_+,$$

and thus

$$A_2 * \mathcal{S}_2(t; \sigma) \geq \inf_{k \geq 1} \left\{ A_2(t - y_k) + [(R_2/\gamma')y_k - (\sigma + R_2\tau/\gamma')]_+ \right\}. \quad (21)$$

Set $R = \min\{R_1, R_2/\gamma'\}$ and let $\gamma' > 1$ and $\delta > 0$ be chosen so that $R_2/\gamma' - \delta = R$. Also fix $\sigma_1, \sigma_2 > 0$ and set $\sigma = \sigma_1 + \sigma_2$. If for a given sample path

$$D_2(t) \geq A_2 * \mathcal{S}_2(t; \sigma_2) \quad (22)$$

and, for all $k \geq 1$ with $y_k \leq t$,

$$D_1(t - y_k) \geq A_1 * \mathcal{S}_1(t - y_k; \sigma_1 + \delta y_k - R_2\tau/\gamma'), \quad (23)$$

then we can insert the lower bound for $D_1 = A_2$ from Eq. (23) into Eq. (21). After collecting terms, the result is $D_2(t) \geq A_1 * \mathcal{S}(t; \sigma)$.

The violation probability of Eq. (22) is given by $\varepsilon_2(\sigma)$. We estimate the violation probability of Eq. (23) as follows:

$$\begin{aligned} & Pr(\text{Eq. (23) fails for some } k \text{ with } y_k \leq t) \\ & \leq L_1 \sum_{k=1}^{\infty} Pr(D_1(t - y_k) < A_1 * \mathcal{S}_1(t; \sigma_2 - R_2\tau/\gamma')) \\ & \leq \frac{L_1}{\log \gamma'} \left(\log \frac{(\gamma' - 1)(\sigma_1 - R\tau)}{\delta\tau} + \frac{1}{\beta_1} \right) (\sigma_1 - R\tau)^{-\beta_1}. \end{aligned}$$

In the first step, we have used the union bound and the *htss* envelope. In the second step, we have used Lemma 4 to evaluate the sum (with γ' in place of γ , and $a = \frac{\delta}{\gamma' - 1}$), and recalled that $R_2/\gamma' - \delta = R$. We eliminate the shift with Eq. (11), and insert the optimal choice $\tau = R^{-1} \left(\frac{L_1}{\beta_1 \log \gamma'} \right)^{\frac{1}{\beta_1}}$. Taking $\gamma' = \sqrt{\gamma}$ and $\delta = R(\gamma' - 1)$, we arrive at

$$Pr \left(\begin{array}{l} \text{Eq. (23) fails for} \\ \text{some } k \text{ with } y_k \leq t \end{array} \right) \leq \tilde{L}_1 \left(\log \sigma_1 + \frac{2 - \log \tilde{L}_1}{\beta_1} \right) \sigma_1^{-\beta_1}.$$

Applying the union bound to the violation probabilities in Eqs. (22) and (23) gives the claim of the lemma. \square

Iterating the lemma results in the following end-to-end service guarantee, referred to as *network service curve*. To keep the statement of the theorem simple, we have assumed that each node offers an *ht* service guarantee with the same rate R , the same constant L , and the same power law β . The general case can be reduced to this with the help of Eqs. (10) and (12).

Theorem 2: *ht* NETWORK SERVICE CURVE. Consider an arrival flow traversing N nodes in series, and assume that the service at each node $n = 1, \dots, N$ satisfies an *ht* service curve

$$\mathcal{S}_n(t; \sigma) = Rt - \sigma, \quad \varepsilon(\sigma) = L\sigma^{-\beta}.$$

Then, for any choice of $\gamma > 1$, the network provides the *ht* service guarantee

$$\begin{aligned} \mathcal{S}_{net}(t; \sigma) &= [(R/\gamma)t - \sigma]_+, \\ \varepsilon_{net}(\sigma) &\leq \tilde{L}N^{2+\beta} \left(\log \frac{\sigma}{N} + \frac{2 - \log \tilde{L}}{\beta} \right) \sigma^{-\beta}, \end{aligned}$$




Fig. 3. Log-log plot of delay bounds for N nodes.

where $\tilde{L} = \frac{L2^{\max\{1, \beta\}}}{\log \gamma}$.

PROOF. We use Lemma 2 to recursively estimate the service offered by the last n nodes with $n = 2, \dots, N$. In each step, we reduce the service rate by a factor $\gamma^{1/(N-1)}$ in place of γ . Fix σ and let $\sigma_1, \dots, \sigma_N$ be such that $\sigma = \sum_{n=1}^N \sigma_n$. This yields

$$\begin{aligned} & Pr(D_N(t) < A_1 * \mathcal{S}_{net}(t; \sigma)) \\ & \leq \sum_{n=1}^N N\tilde{L} \left(\log \sigma_n + \frac{2 - \log L}{\beta} \right) \sigma_n^{-\beta}. \end{aligned}$$

We have replaced $N-1$ with N in several places and increased the $n = N$ term to simplify the result. The claim follows by setting $\sigma_n = \sigma/N$. \square

Example: We perform a multi-node delay analysis for a sequence of homogeneous nodes with the same parameters used for Fig. 2. For this scenario it is possible to draw a comparison with the lower bound for multi-node networks from [7] given in Eq. (20). In Fig. 3, we show lower and upper bounds for networks with $N = 1, 2, 4, 8$ nodes. For reference, we also include the results of individual simulation runs with 10^8 packets. The difference between lower and upper bounds is more pronounced than in the single-node analysis, and increases with the number of nodes N . For both lower and upper bounds, the straight lines make the power-law decay in w apparent. The growth of the bounds in N suggests a power-law growth in N , where the larger spacing for the upper bounds indicates a higher power. As before, we see that simulations violate analytical lower bounds. Since simulations of heavy-tailed traffic have little predictive values for larger delays, our analytical bounds provide more reliable estimates, even with the significant gap between upper and lower bounds.

V. SCALING OF DELAY BOUNDS

We now explore the scaling properties of the delay bounds from the previous section. Throughout this section, we consider a network as in Fig. 1. We assume that the network is *homogeneous*, in the sense that all nodes have the same

capacity C , and all traffic is bounded by *htss* envelopes as in Eq. (3) with the same power α and Hurst parameter H . The cross traffic at each node has rate r_c and constant K_c , and the through flow has rate r_0 and constant K_0 . Traffic can be either fluid-flow or packetized. In the latter case, the packet-size distribution of the through flow satisfies

$$Pr\{X > \sigma\} \leq L_p \sigma^{-\alpha_p}.$$

We assume the stability condition $r_0 + r_c < C$ holds at each node.

Our first result concerns the power-law decay of the delay distribution at a single node. We choose a relaxation of $\mu = \frac{1}{2}(C - r_c - r_0)$, and use the leftover service curve from Eq. (16), given by

$$S(t; \sigma) = [(C - r_c - \mu)t - \sigma]_+, \quad \varepsilon_s(\sigma) \leq L\sigma^{-\beta}, \quad (24)$$

where

$$\beta = \min\{\alpha_p - 1, \alpha(1 - H)\}, \quad (25)$$

and L is an explicitly computable constant. We then apply the delay bound of Theorem 1 with $R = r_0 + \mu$ to obtain

$$Pr(W(t) > w) \leq MR^\beta w^{-\beta}. \quad (26)$$

The constant M is determined by Eq. (18) of Theorem 1 with $\beta' = \beta = \alpha(1 - H)$. This shows that the delay decays with the same power law as the backlog bound in Eq. (13).

Now we consider scaling in networks with $N > 1$ nodes. For $\mu > 0$ (to be chosen below), we obtain at each node the service curve in Eq. (24), with β given by Eq. (25). For $\gamma > 1$ (to be chosen below), we obtain from Theorem 2 the network service curve

$$S_{net}(t; \sigma) = \left[\frac{C - r_c - \mu}{\gamma} t - \sigma \right]_+$$

with violation probability bounded by

$$\varepsilon_{net}(\sigma) \leq N^2 \left(\log z + \frac{2}{\beta} \right) z^{-\beta} \Big|_{z = \frac{\sigma}{L^{1/\beta} N}}.$$

Combining the network service curve with the arrival envelope, we obtain from Theorem 1 for the end-to-end delay W_{net} that

$$Pr(W_{net}(t) > w) \leq \tilde{K}\sigma_1^{-\beta} + \varepsilon_{net}(\sigma_2).$$

If we take $\mu = \frac{1}{3}(C - r_c - r_0)$ and $\gamma = \frac{C - r_c - \mu}{r_0 + \mu}$, then the constants \tilde{K} , L and \tilde{L} do not depend on N . We further choose $\sigma_1 = N^{-1 - \frac{2}{\beta}}(C - r_c - 2\mu)w$ and $\sigma_2 = (C - r_c - 2\mu)w - \sigma_1$, and see that there exist constants M_1, M_2 such that

$$Pr(W_{net}(t) > w) \leq M_1 N^2 (M_2 + \log z) z^{-\beta} \Big|_{z = \frac{w}{N}}.$$

It follows that for $w \rightarrow \infty$ we have the asymptotic upper bound

$$Pr(W_{net}(t) > w) = O(w^{-\beta} \log w).$$

The quantiles of the delay, defined by

$$w_{net}(\varepsilon) = \inf\{w > 0 \mid Pr(W_{net} > w) \leq \varepsilon\}$$

satisfy

$$w_{net}(\varepsilon) = O\left(N^{\frac{2+\beta}{\beta}} (\log N)^{\frac{1}{\beta}}\right). \quad (27)$$

We next compare these upper bounds with scaling results from the literature for a Pareto service time distribution and no cross traffic, where traffic arrives in the form of evenly spaced packets X_i , with an i.i.d. Pareto packet-size distribution, as characterized in Section II. We assume that service times of packets are identical at each node in the sense of [5]. By scaling the units of time and traffic, we may assume an average packet size of $E[X] = 1$ and a link rate $C = 1$, resulting in a rate $\lambda = \rho$, where ρ is the utilization.

For this model, it is known from queueing theory that the delay at a single node decays with a power law with exponent $\alpha - 1$ [10]. Theorem 1 from [10] yields for the queueing time of the k -th packet in the steady state $Q = \lim_{k \rightarrow \infty} Q_k$ that

$$Pr(Q > \sigma) \sim \frac{\rho}{1 - \rho} \frac{(\alpha - 1)^{\alpha - 1}}{\alpha^\alpha} \sigma^{-(\alpha - 1)}$$

as $\sigma \rightarrow \infty$. The delay of the packet is the sum of Q_k and its processing time X_k . This per-packet delay is related with the delay $W(t)$ at a given time by

$$W(t) = \left(Q_{k(t)} + X^*(t) \right) I_{B(t) > 0},$$

where $k(t)$ is the number of the packet being processed at time t , and $X^*(t)$ is the lifetime of the current packet, as defined in Section III. Since packets are i.i.d., $W_{k(t)}$ is independent of $X^*(t)$ and its distribution agrees with W_k , and we can compute

$$\lim_{t \rightarrow \infty} Pr(W(t) > w) \sim \frac{\rho}{1 - \rho} c(\alpha) w^{-(\alpha - 1)} \quad (28)$$

as $w \rightarrow \infty$, where $c(\alpha)$ is a constant that depends on the tail index.

If we compare this exact result with the bound from Eq. (26), we see that $\beta = \alpha - 1$, and so Eq. (26) provides – up to a logarithmic correction – the same power-law decay as Eq. (28). The constant M in Eq. (26) is of order $O((1 - \rho)^{-2})$, while the right hand side of Eq. (28) is of order $(1 - \rho)^{-1}$, which indicates that our delay bound becomes pessimistic as $\rho \rightarrow 1$.

Exploring the scaling in a tandem network, we first note that Eq. (13) states that for a single node, the tail probability of the delays decays with $O(w^{-(\alpha - 1)} \log w)$. Since end-to-end delays exceed the delay at a single node, Eq. (28) guarantees that $W(t) = \Omega(w^{-(\alpha - 1)})$. Thus, the upper and lower bound differ by at most a logarithmic correction. Eq. (27) implies furthermore that delay quantiles are bounded by $O\left(N^{\frac{\alpha + 1}{\alpha - 1}} (\log N)^{\frac{1}{\alpha - 1}}\right)$ as $N \rightarrow \infty$. From the lower bound from [7] given in Eq. (20) we can obtain that quantiles of the end-to-end delay grow at least as fast as $w_{net}(\varepsilon) = \Omega\left(N^{\frac{\alpha}{\alpha - 1}}\right)$.

Lastly, we note that end-to-end delays are expected to grow more slowly if service times are independently regenerated at each node. A large buffer asymptotic from [3] for multi-node networks could be used to obtain the scaling properties of such a network.

VI. CONCLUSIONS

We have presented an end-to-end analysis of networks with heavy-tailed and self-similar traffic. Working within the framework of the network calculus, we developed envelopes for heavy-tailed self-similar traffic and service curves for heavy-tailed service models. By presenting new sample-path bounds for arrivals and service, we were able to derive non-asymptotic performance bounds on backlog and delay, as well as network-wide service characterizations. We explored the scaling behavior of the derived bounds and showed that, for single nodes, the tail probabilities of our delay bounds observe the same power-law decay as known results for G/G/1 systems. We also described the scaling behavior of end-to-end delays. Our paper may motivate further study of the conditions under which performance bounds in a heavy-tailed regime can be tightened. A useful, possibly difficult extension is the derivation of a multi-node service curve that accounts for self-similarity, in addition to heavy-tails.

VII. APPENDIX

The following two lemmas present auxiliary estimates for two sums that involve geometric sequences. In the following, $\tau > 0$ and $\gamma > 1$ are given constants. We refer to [20] for complete proofs in a more general setting.

Lemma 3: Let $x_k = \gamma^k \tau$. For every $\alpha > 0$, $0 < H < 1$, and every $\sigma > 0$, we have

$$\sum_{k=-\infty}^{\infty} \left(\frac{\sigma + x_k}{x_k^H} \right)^{-\alpha} \leq \frac{\gamma^{\alpha H(1-H)}}{\alpha H(1-H) \log \gamma} \sigma^{-\alpha(1-H)}.$$

Lemma 4: Let y_k be defined recursively by $y_0 = 0$ and $y_{k+1} = \tau + \gamma y_k$. For every $\beta > 0$, $a > 0$, and every $\sigma > 0$,

$$\begin{aligned} \sum_{k=1}^{\infty} (\sigma + a y_k)^{-\beta} &\leq \\ &\leq \frac{1}{\log \gamma} \left(\log \left(\frac{(\gamma - 1)(\sigma + a\tau)}{a\tau} \right) + \frac{1}{\beta} \right) (\sigma + a\tau)^{-\beta}. \end{aligned}$$

ACKNOWLEDGMENTS

The research in this paper is supported in part by the National Science Foundation under grants CNS-0435061, and by the Natural Sciences and Engineering Research Council of Canada under two Discovery grants and a Strategic project grant.

REFERENCES

- [1] J. Abate, G. L. Choudhury, and W. Whitt. Waiting-time tail probabilities in queues with long-tail service-time distributions. *Queueing Systems*, 16(3-4):311–338, Sept. 1994.
- [2] H. Ayhan, Z. Palmowski, and S. Schlegel. Cyclic queueing networks with subexponential service times. *Journal of Applied Probability*, 41(3):791–801, Sept. 2004.
- [3] F. Baccelli, M. Lelarge, and S. Foss. Asymptotics of subexponential (max,+): the stochastic event graph case. *Queueing Systems*, 46(1-2):75–96, Jan. 2004.
- [4] J. Y. L. Boudec and P. Thiran. *Network Calculus*. Springer Verlag, Lecture Notes in Computer Science, LNCS 2050, 2001.
- [5] O. Boxma. On a tandem queueing model with identical service times at both counters. part 1,2. *Advances in Applied Probability*, 11(3):616–659, 1979.

- [6] O. J. Boxma and V. Dumas. Fluid queues with long-tailed activity period distributions. *Computer Communications*, 21(17):1509–1529, Nov. 1998.
- [7] A. Burchard, J. Liebeherr, and F. Ciucu. On $\Theta(H \log H)$ scaling of network delays. In *Proc. of IEEE Infocom*, May 2007.
- [8] 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.
- [9] F. Ciucu, A. Burchard, and J. Liebeherr. Scaling properties of statistical end-to-end bounds in the network calculus. *IEEE Transactions on Information Theory*, 52(6):2300–2312, June 2006.
- [10] J. Cohen. Some results on regular variation in queueing and fluctuation theory. *Journal of Applied Probability*, 10(2):343–353, June 1973.
- [11] M. Crovella and A. Bestavros. Self-similarity in World Wide Web traffic: evidence and possible causes. *IEEE/ACM Transactions on Networking*, 5(6):835–846, Dec. 1997.
- [12] J. R. Gallardo, D. Makrakis, and L. Orozco-Barbosa. Use of alpha-stable self-similar stochastic processes for modeling traffic in broadband networks. *Performance Evaluation*, 40:71–98, March 2000.
- [13] J. R. Gallardo, D. Makrakis, and L. Orozco-Barbosa. Probabilistic envelope processes for α -stable self-similar traffic models and their application to resource provisioning. *Performance Evaluation*, 61(2-3):257–279, July 2005.
- [14] B. V. Gnedenko and A. N. Kolmogorov. *Limit distributions for sums of independent random variables*. Addison-Wesley, 1968.
- [15] T. Huang and K. Sigman. Steady-state asymptotics for tandem, split-match and other feedforward queues with heavy tailed service. *Queueing Systems*, 33(1-3):233–259, Dec. 1999.
- [16] Y. Jiang and P. J. Emstad. Analysis of stochastic service guarantees in communication networks: A traffic model. In *19th International Teletraffic Congress (ITC)*, Aug. 2005.
- [17] A. Karasaridis and D. Hatzinakos. Network heavy traffic modeling using α -stable self-similar processes. *IEEE Transactions on Communications*, 49(7):1203–1214, July 2001.
- [18] S. Karlin. *A First Course in Stochastic Processes*. Academic Press, 1975.
- [19] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the self-similar nature of Ethernet traffic. *IEEE/ACM Transactions on Networking*, 2(1):1–15, Feb. 1994.
- [20] J. Liebeherr, A. Burchard, and F. Ciucu. Delay bounds for networks with heavy-tailed and self-similar traffic. Technical report, arXiv:0911.3856, 2009.
- [21] S. Mao and S. S. Panwar. A survey of envelope processes and their applications in quality of service provisioning. *IEEE Communications Surveys & Tutorials*, 8(3):2–20, 3rd Quarter 2006.
- [22] T. Mikosch, S. Resnick, H. Rootzen, and A. Stegeman. Is network traffic approximated by stable Lévy motion or fractional Brownian motion? *Annals of Applied Probability*, 12(1):23–68, Feb. 2002.
- [23] I. Norros. On the use of fractional Brownian motion in the theory of connectionless networks. *IEEE Journal on Selected Areas in Communications*, 13(6):953–962, Aug. 1995.
- [24] V. Paxson and S. Floyd. Wide-area traffic: The failure of Poisson modelling. *IEEE/ACM Transactions on Networking*, 3(3):226–244, June 1995.
- [25] G. Samorodnitsky and M. S. Taqqu. *Stable non-Gaussian random processes: stochastic models with infinite variance*. Chapman and Hall, CRC Press, 1994.
- [26] D. Starobinski and M. Sidi. Stochastically bounded burstiness for communication networks. *IEEE Transactions on Information Theory*, 46(1):206–212, Jan. 2000.
- [27] 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.
- [28] W. Willinger, M. S. Taqqu, R. Sherman, and D. V. Wilson. Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level. In *ACM Sigcomm*, pages 100–113, 1995.
- [29] 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.
- [30] Q. Yin, Y. Jiang, S. Jiang, and P. Y. Kong. Analysis on generalized stochastically bounded bursty traffic for communication networks. In *Proceedings of IEEE Local Computer Networks (LCN)*, pages 141–149, November 2002.