

Delay Bounds for Networks with Heavy-Tailed and Self-Similar Traffic

Jörg Liebeherr, Almut Burchard, Florin Ciucu

Abstract

We provide upper bounds on the end-to-end backlog and delay in a network with heavy-tailed and self-similar traffic. The analysis follows a network calculus approach where traffic is characterized by envelope functions and service is described by service curves. A key contribution of this paper is the derivation of a probabilistic sample path bound for heavy-tailed self-similar arrival processes, which is enabled by a suitable envelope characterization, referred to as *htss* envelope. We derive a heavy-tailed service curve for an entire network path when the service at each node on the path is characterized by heavy-tailed service curves. We obtain backlog and delay bounds for traffic that is characterized by an *htss* envelope and receives service given by a heavy-tailed service curve. The derived performance bounds are non-asymptotic in that they do not assume a steady-state, large buffer, or many sources regime. We also explore the scale of growth of delays as a function of the length of the path. The appendix contains an analysis for self-similar traffic with a Gaussian tail distribution.

I. INTRODUCTION

Traffic measurements in the 1990s provided evidence of self-similarity in aggregate network traffic [22], and heavy-tailed files sizes and bursts were found to be among the root causes [12], [34]. 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 [30], 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 .¹ 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 rescaled 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 every $a > 0$. The exponent $H \in (0, 1)$, referred to as the *Hurst parameter*, specifies the degree of self-similarity.² 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 available, 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 [7], [23], [29]. 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 [17], cyclic networks [2], tandem networks with identical service times [6], and tandem networks where packets have

J. Liebeherr (jorg@comm.utoronto.ca) is 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. F. Ciucu (florin@net.t-labs.tu-berlin.de) is with Deutsche Telekom Laboratories at TU Berlin.

The research in this paper is supported in part by the National Science Foundation under grant CNS-0435061, by two NSERC Discovery grants and an NSERC Strategic grant.

¹We write $f(x) \sim g(x)$, if $\lim_{x \rightarrow \infty} f(x)/g(x) = 1$.

²The 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.

independent service times at nodes in the more general context of stochastic event graphs [3]. The accuracy of some asymptotic approximations has been called into question, 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. The derivation of the sample path bound is made possible by a suitable envelope characterization for heavy-tailed self-similar traffic, referred to as *htss* envelope. We present a characterization for heavy-tailed service and show that it can express end-to-end service available on a path as a composition of the heavy-tailed service at each node. Our end-to-end service characterization enables the computation of end-to-end delay bounds using our single-node result. In an asymptotic regime, our bounds 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 are bounded by $O(N^{\frac{\alpha+1}{\alpha-1}} (\log N)^{\frac{1}{\alpha-1}})$ in 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 [5]. 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 [9], distributions that decay faster than any polynomial [32], and traffic distributions with an effective bandwidth [9]. The latter two groups include certain self-similar processes, in particular, those governed by fractional Brownian motion [28], but do not extend to heavy-tailed distributions. There are also efforts for extending the network calculus to heavy-tailed distributions [14], [15], [18], [19], 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 present brief conclusions in Section VI.

II. THE *htss* TRAFFIC ENVELOPE

In this section we present and evaluate a probabilistic envelope function, for characterizing heavy-tailed self-similar network traffic that permits the derivation of rigorous backlog and delay bounds. The proposed *htss* envelope further develops concepts that were previously studied in [15], [18], [19].

We consider arrivals and departure of traffic at a system, which represents a single node or a sequence of multiple nodes. We use a continuous time model where arrivals and departures of a traffic flow at the system for 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 $\sigma > 0$ [10]:

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

where ε is a non-increasing function of σ that satisfies $\varepsilon(\sigma) \rightarrow 0$ as $\sigma \rightarrow \infty$. The function $\varepsilon(\sigma)$ is used as a bound on the violation probability. 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 envelope concepts [24].

The computation of performance bounds, e.g., bounds on backlog delay, and output burstiness, requires a statistical envelope that bounds an entire sample path $\{A(s, t)\}_{s \leq t}$. A *statistical sample path envelope* $\bar{\mathcal{G}}$ is a statistical envelope that satisfies for all $t \geq 0$ and for all $\sigma > 0$ [10]:

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

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 appears in the *Exponentially Bounded Burstiness* (EBB) model from [35], 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 via the union bound³ by evaluating the right-hand side of Eq. (2) as $\sum_k Pr(A(s_k, t) > \bar{\mathcal{G}}(t-s_{k-1}; \sigma))$ for a suitable discretization $\{s_k\}_{k=1,2,\dots}$, yielding $\bar{\mathcal{G}}(t-s; \sigma) = Rt + \sigma$ for $R > r$ and $\bar{\varepsilon}(\sigma) = \frac{Me^{-a\sigma}}{1-e^{-a(R-r)}}$ [10]. The EBB envelope has been generalized to distributions with moments of all orders, referred to as *Stochastically Bounded Burstiness* (SBB) [32] and corresponding sample path bounds have been developed in [36]. SBB envelopes can characterize arrival processes that are self-similar, but not heavy-tailed. For instance, fractional Brownian motion processes can be fitted with an envelope function $\mathcal{G}(t) = rt + \sigma$ with a Weibullian bound on the violation probability of the form $\varepsilon(\sigma) = Ke^{-(\sigma/a)^\alpha}$ for some $a > 0$ and $0 < b < 1$.

A statistical envelope for general self-similar arrival processes can be expressed as

$$Pr\left(A(s, t) > r(t-s) + \sigma(t-s)^H\right) \leq \varepsilon(\sigma). \quad (3)$$

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. For self-similar traffic, it is natural to allow a heavy-tailed violation probability, since self-similarity can arise from heavy-tailed arrival processes with independent increments.⁴ This consideration leads to our proposed extension of the EBB and SBB concepts that 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

$$Pr\left(A(s, t) > r(t-s) + \sigma(t-s)^H\right) \leq K\sigma^{-\alpha}. \quad (4)$$

³For two events X and Y , $Pr(X \cup Y) \leq Pr(X) + Pr(Y)$.

⁴This is evident in Eqs. (13)–(14) below in an application of the Generalized Central Limit Theorem for a Pareto traffic source.

where K and r are constants, and H and α , respectively, indicate the Hurst parameter and the tail index. We generally assume that $\alpha \in (1, 2)$, that is, arrivals have a finite mean but infinite variance, and $H \in (0, 1)$. In the *htss* envelope 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 characterizations will 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

$$\mathcal{G}(t; \sigma) = rt + \sigma t^H, \quad \varepsilon(\sigma) = K\sigma^{-\alpha} \quad (5)$$

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 at a network node.

Characterizations of self-similar and heavy-tailed traffic by envelopes have been presented before, generally, by exploiting specific properties of α -stable processes [14], [19]. An envelope for α -stable processes in [15] takes the same form $\mathcal{G}(t; \sigma) = rt + \sigma t^H$ as the *htss* envelope, but specifies a fixed violation probability rather than a bound on the distribution. 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. (4) has been obtained in [18] by applying an a-priori bound on the backlog process of an α -stable self-similar process from [19]. Since the backlog bound given in Eq. (24) of [19] is a lower bound (and not an upper bound) on the tail distribution of the buffer occupancy, the envelope in [18] does not satisfy Eq. (2). In Section IV it will become evident that sample path envelopes for arrivals and backlog bounds are interchangeable, in that the availability of one can be used to derive the other. Thus, the sample path bound derived in this paper for heavy-tailed self-similar processes satisfying the *htss* envelope from Eq. (4) also provides the first rigorous backlog bounds for this general class of processes.

In the remainder of this section, we show how to construct *htss* envelopes for relevant distributions, as well as for measurements of packet traces. Ever since traffic measurements at Bellcore from the late 1980s discovered long-range dependence and self-similarity in aggregate network traffic [22], many studies have supported, refined, sometimes also repudiated (e.g., [16]) these findings. This report does not participate in the debate whether aggregate network traffic is best characterized as short-range or long-range dependent, self-similar or multi-fractal, short-tailed or heavy-tailed, and so on. Rather we wish to provide tools for evaluating the performance of networks that may see heavy-tailed self-similar traffic, and shed light on the opportunities and pitfalls of envelope descriptions for heavy-tailed traffic.

A. α -stable Distribution

Stable distributions provide well-established models for non-Gaussian processes with infinite variance. The potential of applying stable processes to data networking was demonstrated in [19] by fitting traces of aggregate traffic (i.e., the Bellcore traces studied in [22]) to an α -stable self-similar process.

A defining property of an α -stable distribution ($0 < \alpha \leq 2$) is that the linear superposition of i.i.d. α -stable random variables preserves the original distribution. That is, if X_1, X_2, \dots, X_m are independent random variables with the same (centered) α -stable distribution, then $m^{-1/\alpha} \sum_{i=1}^m X_i$ has the same distribution. A challenge of working with α -stable distributions is that closed-form expressions for the distribution are only available for a few special cases. However, there exists an explicit expression for the characteristic function of stable distributions, in terms of four parameters (see [31]): a *tail index* $\alpha \in (0, 2]$, a *skewness* parameter $\beta \in [-1, 1]$, a *scale* parameter $a > 0$, and a *location* parameter $\mu \in \mathbb{R}$. For our purposes it is sufficient to work with a normalized stable random variable S_α where $\beta = 1$, $a = 1$, and $\mu = 0$.

The point of departure for our characterization of α -stable processes with *htss* envelopes is the α -stable process proposed in [19] which takes the form

$$A(t) \stackrel{dist.}{=} rt + bt^H S_\alpha . \quad (6)$$

Here, r is the mean arrival rate and b is a parameter that describes the dispersion around the mean.

Remark: We can use Eq. (6) to observe the statistical multiplexing gain of α -stable processes. By the defining property of S_α , the superposition of N i.i.d. processes as in Eq. (6), denoted by A_{mux} , yields

$$A_{mux}(t) = Nrt + N^{1/\alpha} bt^H S_\alpha .$$

Since $1/\alpha < 1$ in the considered range $\alpha \in (1, 2)$, the aggregate of a set of flows increases slower than linearly in the number of flows, thus, giving clear evidence of multiplexing gain. The multiplexing gain diminishes as $\alpha \rightarrow 1$.

We can obtain an *htss* envelope for Eq. (6) from the tail approximation for α -stable distributions [27]

$$Pr(S_\alpha > \sigma) \sim (c_\alpha \sigma)^{-\alpha} , \quad \sigma \rightarrow \infty , \quad (7)$$

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. With Eq. (6) we can write

$$Pr\left(\frac{A(t) - rt}{bt^H} > \sigma\right) \sim (c_\alpha \sigma)^{-\alpha} , \quad \sigma \rightarrow \infty ,$$

Matching this expression with Eq. (4) we obtain the remaining parameter K of the *htss* envelope by setting

$$K = \left(\frac{b}{c_\alpha}\right)^\alpha . \quad (8)$$

By Eq. (7), this envelope only holds for large σ , or, equivalently, low violation probabilities. An alternative method to obtain an *htss* envelope for *all* values of σ is to take advantage of the quantiles of S_α . Since the density of S_α is not available in a closed form, the quantiles must be obtained numerically or by a table lookup. Let the quantile $z(\varepsilon)$ be the value satisfying

$$P(S_\alpha > z(\varepsilon)) = \varepsilon . \quad (9)$$

We obtain a statistical envelope by setting $\mathcal{G}(t; \sigma) = rt + z(\varepsilon)\sigma t^H$ with a fixed ε . In fact, this is the envelope for α -stable processes from [15]. However, since $z(\varepsilon)$ does not follow a power law, it is not an *htss* envelope. To obtain an *htss* envelope from the quantiles, we express Eq. (4) in terms of Eq. (9), which can be done by setting

$$K = \sup_{0 < \varepsilon < 1} \{\varepsilon \cdot (bz(\varepsilon))^\alpha\} . \quad (10)$$

In Fig. 1, we present envelopes for a process satisfying Eq. (6) with

$$r = 75 \text{ Mbps}, \quad \alpha = 1.6, \quad H = 0.8, \quad b = 60 \text{ Mbps} .$$

We show statistical envelopes with fixed violation probabilities $\varepsilon = 10^{-1}, 10^{-2}, 10^{-3}$. The graph compares the *htss* envelopes constructed with the tail approximation using Eq. (7) to those obtained from the quantiles via Eq. (10). As can be expected, the envelopes computed from the asymptotic tail approximation are smaller than the quantile envelopes. We add that envelopes computed from the quantiles for a fixed ε , as described in Eq. (9), are very close to the tail approximation envelopes. If the corresponding envelopes were included in the figure, they would appear almost indistinguishable, suggesting that Eq. (7) provides reasonable bounds for all values of σ .

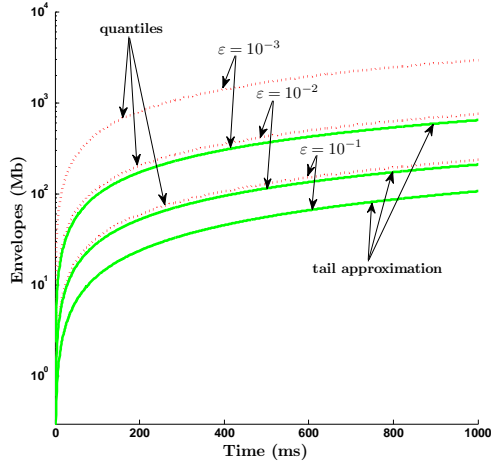


Fig. 1

COMPARISON OF *htss* ENVELOPES FOR AN α -STABLE DISTRIBUTION WITH
 $r = 75 \text{ Mbps}$, $\alpha = 1.6$, $H = 0.8$, $b = 60 \text{ Mbps}$.

B. Pareto Packet Distribution

As second case study, we present an *htss* envelope construction for a packet source with a Pareto arrival distribution. Packets arrive evenly spaced at rate λ and packet sizes are i.i.d. described by a Pareto random variable X_i for the i -th packet with tail distribution

$$\Pr(X_i > x) = \left(\frac{x}{b}\right)^{-\alpha}, \quad x \geq b, \quad (11)$$

where $\alpha \in (1, 2)$. X has finite mean $E[X] = \frac{b\alpha}{\alpha-1}$ and infinite variance. We will construct an *htss* envelope for the compound arrival process

$$A(t) = \sum_{i=1}^{N(t)} X_i, \quad (12)$$

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) [27], which states that the α -stable distribution S_α appears as the limit of normalized sums of i.i.d. random variables. For the independent Pareto random variables X_i , the GCLT yields

$$\frac{\sum_{i=1}^n X_i - nE[X]}{c_\alpha n^{\frac{1}{\alpha}}} \xrightarrow{n \rightarrow \infty} S_\alpha \quad (13)$$

in distribution. 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 in Eq. (12) with Eq. (13) as

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

Since this expression takes the same form as Eq. (6), we can now use the tail estimate of Eq. (7) to obtain an *htss* envelope with parameters

$$r = \lambda E[X], \quad \alpha, \quad H = \frac{1}{\alpha}, \quad K \approx \lambda. \quad (14)$$

The same parameters are valid when $N(t)$ is a Poisson process, according to Theorem 3.1 in [21].

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 [26] in the limit of many flows ($N \rightarrow \infty$) and large time scales ($t \rightarrow \infty$). This aggregate process is particularly interesting since dependent 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 reproduced for the M/G/ ∞ arrival model [26].

Example. We next compare envelope constructions for a Pareto source with evenly spaced packet arrivals with a size distribution given by Eq. (11). The parameters are

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

With these values, the average packet size is 400 Byte. We evaluate the following types of envelopes:

1. *htss GCLT envelope.* This refers to the envelope constructed with the GCLT according to Eq. (14). The value of σ of the *htss* envelope is set so that the right hand side of Eq. (4) satisfies a violation probability of $\varepsilon = 10^{-3}$.
2. *Deterministic trace envelope.* This envelope is computed from a simulation of a packet trace with 1 million packets drawn from the given Pareto distribution. We compute the smallest envelope for the trace that satisfies Eq. (1) with $\varepsilon(\sigma) = 0$ for all $\sigma > 0$. The deterministic trace envelope, which is computed by $\mathcal{G}(t) = \sup_{\tau} \{A(t + \tau) - A(\tau)\}$ [5], provides the smallest envelope of a trace that is never violated.
3. *htss trace envelope.* This is an *htss* envelope created from the same Pareto packet trace. We assume that the values of α and $H = 1/\alpha$ are given, but that the distribution is not known. The envelope is created directly from Eq. (4) by inspecting the relative frequency at which subintervals of the trace violate the *htss* envelope. First, K is selected as the smallest number that satisfies the right hand side of Eq. (4) for all values of σ . Then σ is found by fixing the violation probability $\varepsilon = 10^{-3}$.
4. *Average rate.* For reference, we also include the average rate of the data in the figures, which is obtained from the same packet trace as in the trace envelopes.

The resulting envelopes are plotted in Fig. 2. The discrete steps of the deterministic trace envelope around $t = 0$ ms, $t = 460$ ms, and $t = 680$ ms are due to arrivals of very large packets at certain times in the simulated trace. The *htss* GCLT envelope is quite close to the data of the average rate. However, since the GCLT is an asymptotic result, this envelope is possibly too optimistic. On the other hand, the safely conservative *htss* trace envelope is much larger than the corresponding deterministic trace envelope. The reason is that the construction of this envelope performs a heavy-tailed extrapolation of the data trace, and thus amplifies the variability of the underlying trace.

To investigate the variability of *htss* trace envelopes as a function of the length of the trace, we present envelopes obtained from subintervals of the trace used for Fig. 2. We use non-overlapping subintervals of

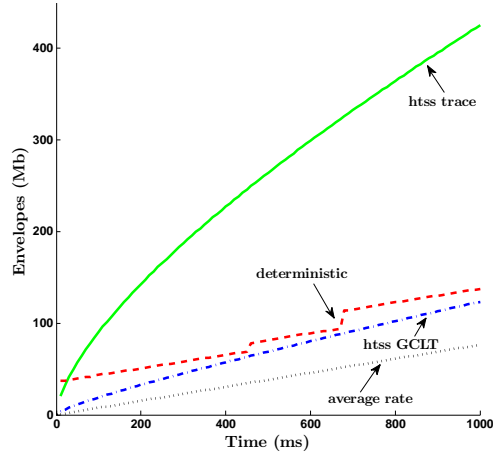


Fig. 2

ENVELOPES FOR A PARETO PACKET SOURCE ($\varepsilon = 10^{-3}$).

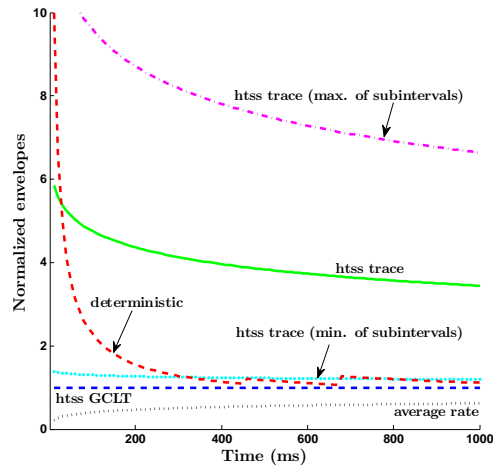


Fig. 3

NORMALIZED ENVELOPES FOR A PARETO PACKET SOURCE ($\varepsilon = 10^{-3}$).

100,000 packets from the trace and compute *htss* trace envelopes for the subintervals. In Fig. 3, we show the envelopes normalized by the values of the *htss* GCLT envelope. We plot the maximum and the minimum values of the computed *htss* envelopes for the subintervals. The figure makes drastically clear that the range of values of the *htss* envelopes for the shorter intervals cover a wide range. This illustrates an inherent problem with generating a traffic characterization for heavy-tailed traffic from limited data sets.

C. Measured Packet Traces

We next show how to obtain an *htss* envelope from measured traffic traces. The trace data was collected in October 2005 at the 1 Gbps uplink of the Munich Scientific Network, a network with more than 50,000 hosts, to the German research backbone network. The complete trace contains more than 6 billion packets, collected over a 24-hour time period. Further details on the data trace and the collection methodology can be

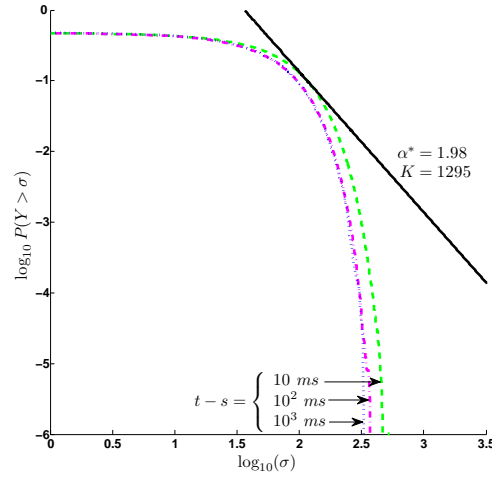


Fig. 4

NORMALIZED LOG-LOG PLOT OF *Munich* TRACE.

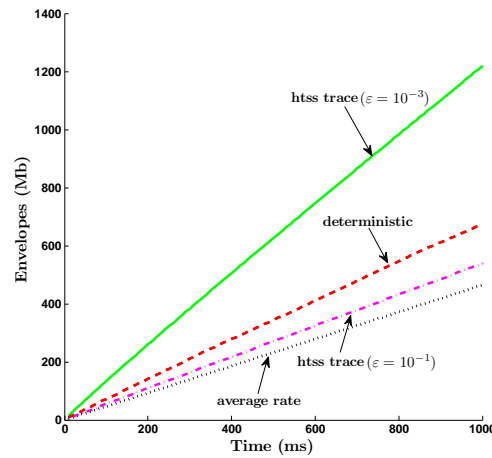


Fig. 5

htss ENVELOPES FOR *Munich* TRACE.

found in [13]. From this data, we select the first 10^9 packets corresponding to 2.75 hours worth of data, with an overall average rate of $r^* = 465$ Mbps. We refer to this data as the *Munich* trace.

To extract the tail index and the Hurst parameter from the trace, we take advantage of parameter estimation methods for stable processes from [25],⁵ which yields the following parameters for the *Munich* trace:

$$\alpha^* = 1.98, \quad H^* = 0.93.$$

The remaining parameter K needed for the *htss* envelope can now be obtained in the same way as described for the *htss* trace envelope.

To provide a sense of the trace data, we present in Fig. 4 a normalized log-log plot of the *Munich* trace

⁵We use source code provided to us by the authors of [4].

data in terms of the normalized random variable

$$Y := \frac{A(s, t) - r^*(t - s)}{(t - s)^{H^*}} .$$

Since $Pr(Y > \sigma) = Pr(A(s, t) > \mathcal{G}(t-s; \sigma))$ where \mathcal{G} is given in Eq. (5), the distribution of Y corresponds to that of violations of the *htss* envelope. In the figure, we show the log-log plot of Y for different values of $(t - s)$, namely, $t - s = 10, 100, 1000$ ms. If the trace data was self-similar with the exact Hurst parameter H^* , the log-log data curves should match perfectly for all values of $(t - s)$. (We note that by reducing the value of the Hurst parameter slightly, the curves for different values of $(t - s)$ can be made to match up almost perfectly). Since the decay of the log-log plots is obviously not linear, the distribution of the *Munich* trace does not appear to be heavy-tailed. We will see that a characterization of such a non-heavy-tailed process by an *htss* envelope leads to a pessimistic estimation.

We can also use Fig. 4 to graphically construct an *htss* envelope for the *Munich* trace. Since we already have determined the tail index α^* and the Hurst parameter H^* as given above, we only need to find K . The value of this parameter can be obtained by taking the logarithm of Eq. (4). Using the definition of Y , this yields

$$\log Pr(Y > \sigma) \leq \log K - \alpha^* \log \sigma .$$

Applying this relationship to Fig. 4, we should select K as the smallest value such that the linear function $\log K - \alpha^* \log \sigma$ lies above the log-log plots of $Pr(Y > \sigma)$ in the figure. In Fig. 4, we include the linear segment with $K = 1225$ as a thick line. Clearly, any other selection of K and α^* providing an upper bound of the log-log plots of the *Munich* trace also yields a valid *htss* envelope for all values of σ . An *htss* envelope for a fixed violation probability ε can be obtained from Fig. 4 by finding the value of σ that corresponds to the desired violation probability of the linear segment. Finally, we can use Fig. 4 to assess the accuracy of the *htss* envelope. The linear segment (the thick black line) is close to the trace data when $Pr(Y > \sigma) \approx 10^{-1}$. Otherwise, the linear segment is quite far apart from the plots of the trace. This indicates that the *htss* envelopes developed with the parameter settings used for the linear segment are accurate only when the violation probability is around 10^{-1} . If the data trace was truly heavy-tailed, the data curves would maintain a linear rate of decline at a rate around α^* , and would remain close to the linear segment for any σ sufficiently large.

In Fig. 5, we show *htss* envelopes for the *Munich* trace obtained with the linear segment from Fig. 4 for $\varepsilon = 10^{-1}$ and $\varepsilon = 10^{-3}$. For comparison, we include in Fig. 5 the average rate of the traffic trace, as well as a deterministic envelopes of the 1 sec subinterval of the trace that generates the most traffic. (Since the computation of a deterministic trace envelope as defined in Subsection II-B grows quadratically in the size of the trace, the computation time to construct a deterministic envelope for the complete *Munich* trace is prohibitive. The included deterministic envelope for a subinterval of the trace is a lower bound for the deterministic envelope of the complete trace. However, for the depicted time intervals, the deterministic envelope for the subinterval is a good representation of the deterministic envelope of the entire trace, for several reasons. First, by selection of the subinterval, at $t = 1000$ ms the envelope of the subinterval and the envelope of the complete trace are identical. Second, since any deterministic envelope is a subadditive function, the slope of the envelope decreases for larger values of time. Now, any function that satisfies these properties cannot vary significantly from the depicted envelope of the selected subinterval.) Comparing the *htss* envelopes with the reference curves confirms our earlier discussion on the accuracy of the *htss* envelopes: For $\varepsilon = 10^{-1}$, the *htss* envelope is close to the plot of the average rate. On the other hand, the envelope for $\varepsilon = 10^{-3}$ is quite pessimistic, and lies well above the deterministic envelope.

III. SERVICE GUARANTEES WITH HEAVY TAILS

We next formulate service guarantees with 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 bound 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 $\sigma \rightarrow \infty$.

We define a *heavy-tailed (ht) service curve* as a service curve of the form

$$\mathcal{S}(t, \sigma) = [Rt - \sigma]_+, \quad \varepsilon(\sigma) = L\sigma^{-\beta} \quad (15)$$

for some β with $0 < \beta < 2$ and some constant L . In analogy to the formulation of traffic envelopes in Section II, the *ht* service curve specifies that the deviation from the service rate guarantee R has a heavy-tailed decay. The rationale for not including a Hurst parameter in the definition of the *ht* service guarantees is that the form of Eq. (15) facilitates the computation of service bounds over multiple nodes. In this paper, we consider two types of *ht* service curves, one characterizing the available capacity at a link with cross-traffic, the other modeling a packetizer.

- *Service at link with cross traffic (Leftover Service)*: This service curve seeks to describe 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 guarantees for most workconserving multiplexers [5]. Since 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*. Since the derivation of an *ht* service curve for such a leftover service requires a sample path bound for the *htss* cross traffic, we defer the derivation to Subsection IV-B.

- *Packetizer*: We will also use the *ht* service model to express a packetized view of traffic with a heavy-tailed packet size distribution. We model discrete packet sizes by a service element that delays traffic until all bits belonging to the same packet have arrived, and then releases all bits of the packet at once. Such an element is referred to as a *packetizer*. By investigating packetized traffic we can relate our bounds to a queuing theoretic analysis with a packet-level interpretation of traffic (see Section V). We now derive a service curve for a packetizer. For a packet-size distribution satisfying $Pr(X > \sigma) \leq L\sigma^{-\alpha}$, we show that a constant-rate workconserving link of capacity C provides an *ht* service curve with rate $R = C$ and a suitable function $\varepsilon(\sigma)$.

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 of a renewal process. It is known from the theory of renewal processes (see [20], pp. 194) that

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

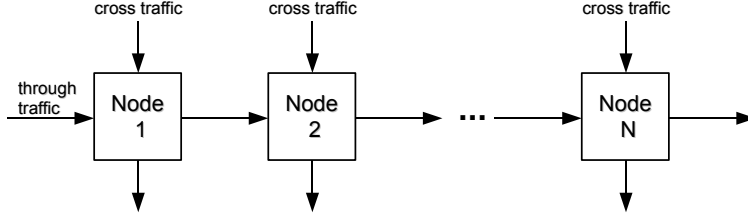


Fig. 6

A NETWORK WITH CROSS TRAFFIC.

The bound on the tail probability holds for all times t , provided that the arrival time of the first packet after the network is started with empty queues at $t = 0$ is properly randomized. Using $X^*(t)$, the departures of a packetizer 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 , that is,

$$\rho = \sup_{s \leq t} \frac{E[A(s, t)]}{C(t - s)},$$

then

$$\begin{aligned} Pr(D(t) < A * \mathcal{S}(t; \sigma)) &\leq \rho Pr(X^*(t) > \sigma) \\ &\leq \frac{\rho L}{(\alpha - 1)E[X]} \sigma^{-(\alpha-1)}. \end{aligned} \quad (16)$$

IV. NETWORK CALCULUS WITH *htss* ENVELOPES

We consider a network as in Fig. 6. 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 yield bounds on the end-to-end delay and backlog of the through traffic.

A. Statistical sample path Envelope

The network calculus for heavy-tailed traffic is enabled by a statistical sample path envelope for traffic with *htss* envelopes. To motivate the relevance of the sample path bound, 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 [14], [15], [18], [19], the tail distribution is approximated by

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

However, the right hand side is generally *smaller* than the left hand side. Applying to the right hand side a statistical envelope that 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

$$\begin{aligned} Pr(B(t) > \sigma) &\leq \sum_{k=-\infty}^{\infty} Pr(A(t - x_k, t) > \sigma + Cx_{k-1}) \\ &\leq \frac{1}{H(1-H)\log\gamma} \int_z^{\infty} Kx^{-\alpha-1} dx \Big|_{z=\frac{(C/\gamma-r)H\sigma^{1-H}}{\gamma^{H(1-H)}}} \\ &\leq \tilde{K}\sigma^{-\alpha(1-H)}. \end{aligned} \quad (17)$$

In the second line we have used Lemma 6 from the appendix to evaluate the sum. Writing $C = r + \mu$ and minimizing over γ gives the constant

$$\tilde{K} = K \cdot \inf_{1 < \gamma < 1 + \frac{\mu}{r}} \left\{ \left(\frac{r + \mu}{\gamma} - r \right)^{-\alpha H} \frac{\gamma^{\alpha H(1-H)}}{\alpha H(1-H)\log\gamma} \right\}. \quad (18)$$

We remark that, typically, we have $1 < \alpha < 2$ and $\alpha^{-1} \leq H < 1$, 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 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 [33], 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 [33] only optimize over τ in uniformly spaced sequences. Note that using a uniform discretization in the derivation of Eq. (17) would cause the infinite sum 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 $\mu > 0$,

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

is a statistical sample path envelope according to Eq. (2). The constant \tilde{K} is given by Eq. (18).

The proof follows immediately from Eq. (17) 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

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

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

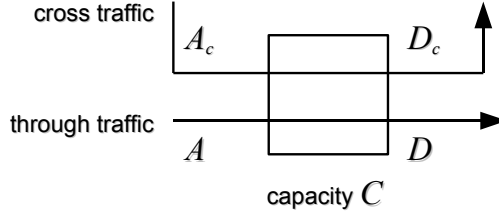


Fig. 7

CONSTANT-RATE LINK WITH CAPACITY C .

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 at a constant-rate link with capacity C , as illustrated in Fig. 7. We denote arrivals of the through flow by A and cross traffic arrivals by A_c . Departures are denoted by D and D_c , respectively. Assuming that 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.

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

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

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

$$\begin{aligned} Pr(D(t) < A * \mathcal{S}(t; \sigma)) &\leq Pr(\sup_{s \leq \underline{t}} \{A_c(s, t) - (r_c + \mu)(t - s)\} > \sigma) \\ &\leq \tilde{K}_c \sigma^{-\alpha_c(1-H_c)}, \end{aligned} \quad (19)$$

where \tilde{K}_c is given by Eq. (18). This proves that $\mathcal{S}(t; \sigma) = [Rt - \sigma]_+$ is an *ht* service curve.

The description of the leftover service in Eq. (19) can be combined with Eq. (16) 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

$$\mathcal{S}(t; \sigma) = [(C - r_c - \mu)t - \sigma]_+, \quad \varepsilon(\sigma) = L\sigma^{-\beta}, \quad (20)$$

where $\beta = \min\{\alpha_p - 1, \alpha_c(1 - H_c)\}$, α_p is the tail decay in Eq. (16), and $\mu > 0$ is a free parameter. The violation probability is given by

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

where $\rho \leq 1$ is the utilization 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. (18) with \tilde{K}_c in place of K . When traffic is not packetized, the second term in the sum above is equal to zero. The constant in Eq. (20) can be computed explicitly by first using Lemma 5 (Eq. (42)) to lower the larger exponent to β , and then applying Lemma 5 (Eq. (45)).

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

$$\begin{aligned} Pr(W(t) > w) &\leq Pr\left(\sup_{s \leq t - w} \{A(s, t - w) - R(t - s - w)\} > \sigma_1\right) \\ &\quad + Pr\left(D(t) < \inf_{s \leq t} [A(s) + R(t - s) - \sigma_2]_+\right) \\ &\leq \tilde{K} \sigma_1^{-\alpha(1-H)} + L \sigma_2^{-\beta}, \end{aligned} \tag{21}$$

where \tilde{K} is defined by Eq. (18). 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 Lemma 5 (Eq. (42)), and then minimizing explicitly over the choice of σ_1 and σ_2 using Lemma 5 (Eq. (45)). For the constant, this yields the estimate

$$M \leq \left\{ \tilde{K}^{\frac{\beta'}{(1+\beta')\alpha(1-H)}} + L_s^{\frac{\beta'}{(1+\beta')\beta}} \right\}^{1+\beta'}. \quad (22)$$

□

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

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

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. (16). The reason for selecting this example (which does not have cross traffic) is that it permits a comparison with a queueing theoretic result in [8], 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

$$w_N(z) \geq \frac{(Nb)^{\frac{\alpha}{\alpha-1}}}{((\alpha-1)\lambda^{-1}|\log(1-\varepsilon)|)^{\frac{1}{\alpha-1}}}. \quad (23)$$

In Fig. 8 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 [8]. 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 contain sufficiently many events with large delays, they violate analytical lower bounds. 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 [5] 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.

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 $\mathcal{S}_1(t, \sigma) = [R_1 t - \sigma]_+$ and $\varepsilon_1(\sigma) = L_1 \sigma^{-\beta_1}$, and the second node offers a service curve $\mathcal{S}_2(t; \sigma) = [R_2 t - \sigma]_+$ and an arbitrary function $\varepsilon_2(\sigma)$. Then for any $\gamma > 1$, the two nodes offer the combined service curve given by

$$\begin{aligned} \mathcal{S}(t, \sigma) &= \left[\min \left\{ R_1, \frac{R_2}{\gamma} \right\} t - \sigma \right]_+, \\ \varepsilon(\sigma) &= \inf_{\sigma_1 + \sigma_2 = \sigma} \left\{ \tilde{\varepsilon}_1(\sigma_1) (|\log \tilde{\varepsilon}_1(\sigma_1)| + 2) 2^{[\beta_1 - 1]_+} + \varepsilon_2(\sigma_2) \right\}, \end{aligned}$$

where $\tilde{\varepsilon}_1(\sigma) = \min \left\{ 1, \frac{2}{\beta_1 \log \gamma} L_1 \sigma_1^{-\beta_1} \right\}$.

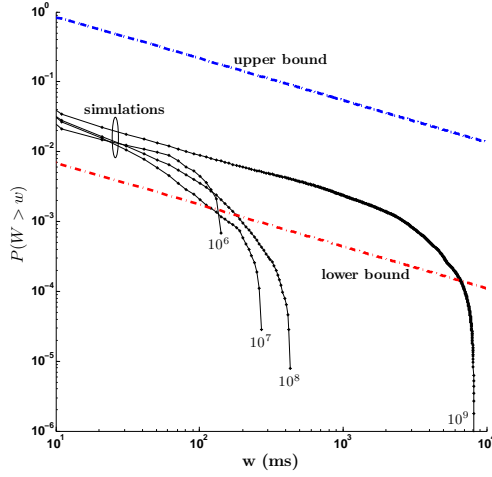


Fig. 8

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.

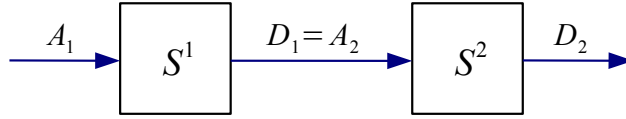


Fig. 9

TWO NODES IN SERIES.

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 at the expense of lowering the exponent using Eq. (43) from Lemma 5. 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 Lemma 5 (Eqs. (42) and (45)).

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

$$A_2 * \mathcal{S}_2(t; \sigma) \geq \inf_{k \geq 1} \left\{ A_2(t - y_k) + \left[\frac{R_2}{\gamma'} y_k - \left(\sigma + \frac{R_2}{\gamma'} \tau \right) \right]_+ \right\}. \quad (24)$$

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

$$D_2(t) \geq A_2 * \mathcal{S}_2(t; \sigma_2) \quad (25)$$

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

$$D_1(t - y_k) \geq A_1 * \mathcal{S}_1(t - y_k; \sigma_1 + \delta y_k - \frac{R_2}{\gamma'} \tau), \quad (26)$$

then we can insert the lower bound for $D_1 = A_2$ from Eq. (26) into Eq. (24). After collecting terms, the result is $D_2(t) \geq A_1 * \mathcal{S}(t; \sigma)$.

The violation probability of Eq. (25) is given by $\varepsilon_2(\sigma_2)$. Assume for the moment that $\sigma \geq \frac{\delta \tau}{\gamma' - 1}$. We estimate the violation probability of Eq. (26) by

$$\begin{aligned} & Pr(\text{Eq. (26) fails for some } k \text{ with } y_k \leq t) \\ & \leq L_1 \sum_{k=1}^{\infty} Pr(D_1(t - y_k) < A_1 * \mathcal{S}_1(t - y_k; \sigma_1 + \delta y_k - R_2 \tau / \gamma')) \\ & \leq \frac{L_1}{\log \gamma'} \left(\frac{1}{\beta_1} + \left[\log \frac{(\gamma' - 1)(\sigma_1 - R\tau)}{\delta \tau} \right]_+ \right) (\sigma_1 - R\tau)^{-\beta_1}. \end{aligned}$$

In the first step, we have used the union bound and the ht service curve \mathcal{S}_1 . In the second step, we have used Lemma 7 to evaluate the sum (with γ' in place of γ , and $a = \delta$), and recalled that $R_2/\gamma' - \delta = R$. (Here, we have used the assumption on σ given before the equation). We eliminate the shift with Lemma 5 (Eq. (44)), 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

$$Pr \left(\begin{array}{l} \text{Eq. (26) fails for} \\ \text{some } k \text{ with } y_k \leq t \end{array} \right) \leq \tilde{L}_1 \sigma_1^{-\beta_1} \left(\log(\tilde{L}_1 \sigma_1^{-\beta_1}) + 2 \right) \leq \tilde{\varepsilon}_1(\sigma_1) (|\log \tilde{\varepsilon}_1(\sigma_1)| + 2),$$

where $\tilde{L}_1 = \frac{2^{\max\{1, \beta_1\}}}{\beta_1 \log \gamma} L_1$. This bound remains valid for $\sigma < \frac{\delta \tau}{\gamma' - 1} = R\tau$, since then we have $\tilde{\varepsilon}_1(\sigma) = 1$. Applying the union bound to the violation probabilities in Eqs. (25) and (26) 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 Lemma 5 (Eqs. (42) and (45)).

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 every choice of $\gamma > 1$, the network provides the service guarantee

$$\begin{aligned} \mathcal{S}_{net}(t, \sigma) &= [(R/\gamma)t - \sigma]_+, \\ \varepsilon_{net}(\sigma) &\leq N^{2+\beta} \cdot 2^{[\beta-1]_+} \cdot \tilde{\varepsilon}(\sigma) (|\log \tilde{\varepsilon}(\sigma)| + (1 + \beta) \log N + 2), \end{aligned}$$

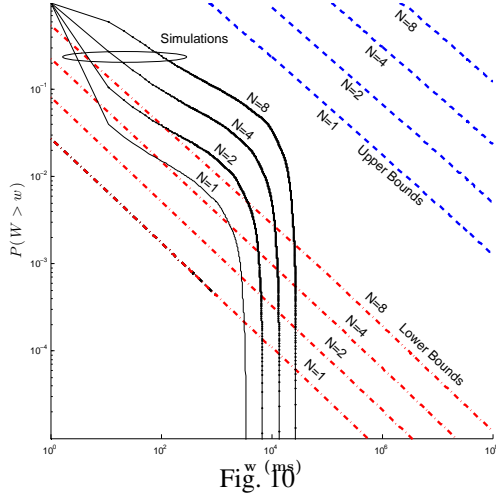


Fig. 10 LOG-LOG PLOT OF DELAY BOUNDS FOR N NODES.

where $\tilde{\varepsilon}(\sigma) = \min \left\{ 1, \frac{2^{\max\{1, \beta\}}}{\beta \log \gamma} L \sigma^{-\beta} \right\}$.

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 $\gamma^{\frac{1}{N-1}}$ in place of γ . Fix σ , and set $\sigma_n = \sigma/N$ for $n = 1, \dots, N$. If $\tilde{\varepsilon}(\sigma/N) \geq 1$, there is nothing to show. Otherwise, we obtain

$$Pr(D_N(t) < A_1 * S_{net}(t; \sigma)) \leq \sum_{n=1}^N N \tilde{\varepsilon}\left(\frac{\sigma}{N}\right) (|\log(N \tilde{\varepsilon}\left(\frac{\sigma}{N}\right))| + 2),$$

and the claim follows by collecting the factors of N . \square

Example: We perform a multi-node delay analysis for a sequence of homogeneous nodes with the same parameters used for Fig. 8. The reason for using this example, is that it permits us to draw a comparison with the lower bound for multi-node networks from [8] given in Eq. (23). In Fig. 10 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. 6. 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. (4) 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.

Single node, large delays ($w \rightarrow \infty$). 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. (20), given by

$$\mathcal{S}(t; \sigma) = [(C - r_c - \mu)t - \sigma]_+, \quad \varepsilon_s(\sigma) \leq L\sigma^{-\beta}, \quad (27)$$

where

$$\beta = \min\{\alpha_p - 1, \alpha(1 - H)\}, \quad (28)$$

and L is an explicitly computable constant. (For fluid-flow traffic, that is, without a packetizer, the first term does not appear, and we have $\beta = \alpha(1 - H)$.) We then apply the delay bound of Theorem 1 with $R = r_0 + \mu$ to obtain

$$Pr(W(t) > w) \leq MR^\beta w^{-\beta}. \quad (29)$$

The constant M is determined by Eq. (22) of Theorem 1 with $\beta' = \beta$. This shows that the delay decays with the same power law as the backlog bound in Eq. (17).

Multiple nodes, large delays ($w \rightarrow \infty$). Now we consider scaling in networks with $N > 1$ nodes. We choose $\mu = \frac{1}{3}(C - r_c - r_0)$. We obtain at each node the service curve in Eq. (27), with β given by Eq. (28). We next choose $\gamma = \frac{C - r_c \mu}{r_0 + \mu}$ and obtain from Theorem 2 the network service curve

$$\mathcal{S}_{net}(t; \sigma) = [R_{net}t - \sigma]_+,$$

where $R_{net} = r_0 + \mu$, and with violation probability bounded by

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

with an explicitly computable constant \tilde{L} that does not depend on N . Combining the network service curve with the arrival envelope, we obtain from Eq. (21) of Theorem 1 for the end-to-end delay W_{net} that

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

Here, the constant \tilde{K} is given by Eq. (18) with r_0 in place of r . We further choose $\sigma_1 = N^{-1-\frac{2}{\beta}} R_{net}w$ and $\sigma_2 = R_{net}w - \sigma_1$, and see that

$$Pr(W_{net}(t) > w) \leq N^{2+\beta} (M_1 \log w + M_2 \log N + M_3) w^{-\beta}.$$

The constants M_1 , M_2 , and M_3 are again explicitly computable, and do not depend on N . The tail of the delay distribution, i.e., when $w \rightarrow \infty$, is dominated by the first summand in the brackets, thus, we have the asymptotic upper bound

$$Pr(W_{net}(t) > w) = O(w^{-\beta} \log w), \quad (w \rightarrow \infty). \quad (30)$$

Multiple nodes, long paths ($N \rightarrow \infty$). For long paths, i.e., $N \rightarrow \infty$, the second summand dominates. The quantiles of the delay, defined by

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

satisfy

$$w_{net}(\varepsilon) = O(N^{\frac{2+\beta}{\beta}} (\log N)^{\frac{1}{\beta}}), \quad (N \rightarrow \infty). \quad (31)$$

Comparison of scaling bounds. 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 [6]. 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 $\alpha - 1$ [11]. Theorem 1 from [11] yields for the queueing time of the k -th packet in the steady state $Q = \lim_{k \rightarrow \infty} Q_k$ that

$$Pr(Q > \sigma) \sim \frac{\rho}{1 - \rho} \frac{(\alpha - 1)^{\alpha - 1}}{\alpha^\alpha} \sigma^{-(\alpha - 1)}, \quad (\sigma \rightarrow \infty).$$

The delay of the k th packet is the sum of its queueing time Q_k and its processing time X_k . This per-packet delay is related with the delay $W(t)$ at a given time by

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

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

$$\lim_{t \rightarrow \infty} Pr(W(t) > w) \sim \frac{\rho}{1 - \rho} c(\alpha) w^{-(\alpha - 1)}, \quad (w \rightarrow \infty), \quad (32)$$

where $c(\alpha)$ is a constant that depends on the tail index.

If we compare this asymptotic exact result with our bound from Eq. (29) and (30), we see that $\beta = \alpha - 1$, and so Eq. (30) provides (up to a logarithmic correction) the same power-law decay as Eq. (32). The constant M in Eq. (29) is of order $O((1 - \rho)^{-2})$, while the right hand side of Eq. (32) 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. (17) 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. (32) guarantees that $W(t) = \Omega(w^{-(\alpha - 1)})$. Thus, our upper and lower bound differ by at most a logarithmic factor. Eq. (31) implies furthermore that delay quantiles are bounded by $O(N^{\frac{\alpha + 1}{\alpha - 1}} (\log N)^{\frac{1}{\alpha - 1}})$ as $N \rightarrow \infty$. From the lower bound from [8] given in Eq. (23) we can obtain that quantiles of the end-to-end delay grow at least as fast as $w_{net}(\varepsilon) = \Omega(N^{\frac{\alpha}{\alpha - 1}})$.

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.

ACKNOWLEDGMENTS

We thank Prof. Anja Feldmann at TU Berlin for giving us access to a collection of data trace. J. Liebeherr thanks Deutsche Telekom Laboratories for hospitality during a sabbatical. We thank Prof. Stephen McLaughlin for sharing a software library for alpha-stable distributions. We thank Yashar Ghiassi for comments on the manuscript.

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] S. Bates and S. McLaughlin. The estimation of stable distribution parameters. In *Proc. of IEEE SP Workshop on Higher-Order Statistics*, pages 390–94, 1997.
- [5] J. Y. L. Boudec and P. Thiran. *Network Calculus*. Springer Verlag, Lecture Notes in Computer Science, LNCS 2050, 2001.
- [6] 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.
- [7] O. J. Boxma and V. Dumas. Fluid queues with long-tailed activity period distributions. *Computer Communications*, 21(17):1509–1529, Nov. 1998.
- [8] A. Burchard, J. Liebeherr, and F. Ciucu. On $\Theta(H \log H)$ scaling of network delays. In *Proc. IEEE Infocom*, pages 1866–1874, May 2007.
- [9] 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.
- [10] 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.
- [11] J. Cohen. Some results on regular variation in queueing and fluctuation theory. *Journal of Applied Probability*, 10(2):343–353, June 1973.
- [12] 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.
- [13] H. Dreger, A. Feldmann, M. Mai, V. Paxson, and R. Sommer. Dynamic application-layer protocol analysis for network intrusion detection. In *Proc. 15th USENIX Security Symposium*, pages 257–272, July/August 2006.
- [14] 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(1–3):71–98, March 2000.
- [15] 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.
- [16] W.-B. Gong, Y. Liu, V. Misra, and D. Towsley. Self-similarity and long range dependence on the internet: a second look at the evidence, origins and implications. *Computer Networks*, 48(21):377–399, June 2005.

- [17] 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.
- [18] Y. Jiang and P. J. Emstad. Analysis of stochastic service guarantees in communication networks: A traffic model. In *Proc. 19th International Teletraffic Congress (ITC)*, Aug. 2005.
- [19] 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.
- [20] S. Karlin. *A First Course in Stochastic Processes*. Academic Press, 1975.
- [21] C. Klüppelberg and T. Mikosch. Large deviations of heavy-tailed random sums with applications in insurance and finance. *Journal of Applied Probability*, 34(2):293–308, June 1997.
- [22] 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.
- [23] Z. Liu. Special issue on long range dependence and heavy tail distributions. *Performance Evaluation*, 61(2-3), July 2005.
- [24] 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.
- [25] H. J. McCulloch. Simple consistent estimators of stable distribution parameters. *Communications Statistics and Simulation*, 15(4):1109–1136, 1986.
- [26] 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.
- [27] J. P. Nolan. *Stable Distributions: Models for Heavy Tailed Data*. Birkhäuser, Noston, 2009.
- [28] 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.
- [29] K. Park and W. Willinger, editors. *Self-Similar Network Traffic and Performance Evaluation*. John Wiley & Sons, 2000.
- [30] V. Paxson and S. Floyd. Wide-area traffic: The failure of Poisson modelling. *IEEE/ACM Transactions on Networking*, 3(3):226–244, June 1995.
- [31] G. Samorodnitsky and M. S. Taqqu. *Stable Non-Gaussian Random Processes: Stochastic Models with Infinite Variance*. Chapman and Hall, New York, 1994.
- [32] D. Starobinski and M. Sidi. Stochastically bounded burstiness for communication networks. *IEEE Transactions on Information Theory*, 46(1):206–212, Jan. 2000.
- [33] 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.
- [34] 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, August/September 1995.
- [35] 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.
- [36] 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.

APPENDIX

I. SELF-SIMILAR TRAFFIC WITH GAUSSIAN TAILS

While self-similar traffic is expected to be heavy tailed, some self-similar traffic types appearing in network analysis exhibit a faster decay. In the following we present a derivation of a network service curve for traffic with a Gaussian bound on the violation probability. More precisely, will consider processes that satisfy for all $\sigma > 0$

$$Pr\left(A(s, t) > r(t-s) + \sigma(t-s)^H\right) \leq Ke^{-\frac{1}{2}(\sigma/b)^2}, \quad (33)$$

with some constants $K > 0$, $a > 0$, and $\alpha \geq 0$. In terms of Eq. (1), this corresponds to a statistical envelope

$$\mathcal{G}(t; \sigma) = rt + \sigma t^H, \quad \varepsilon(\sigma) = Ke^{-\frac{1}{2}(\sigma/b)^2}. \quad (34)$$

To motivate that envelopes given by Eq. (34) can provide useful traffic models, consider first the fractional Brownian motion (fBm) traffic model. In this model, arrivals are given by

$$A(t) \stackrel{dist.}{=} rt + bt^H N_{0,1}, \quad (35)$$

where $N_{0,1}$ is the standard normal distribution, and b is a parameter that relates to the standard deviation. Eq. (35) is the special case of the α -stable process from Eq. (6) with $\alpha = 2$. Using that the density of $N_{0,1}$ is given by $\phi(x) = \frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}x^2}$, we see from Eq. (35) that

$$Pr\left(A(s, t) > r(t-s) + \sigma(t-s)^H\right) = \int_{\sigma}^{\infty} \frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}x^2} dx \leq \frac{1}{2}e^{-\frac{1}{2}\sigma^2}, \quad (36)$$

i.e., Eq. (34) with $K = 1/2$ is a statistical envelope.

An alternative description of arrival processes with Gaussian tails is through their *effective bandwidth*, defined by

$$eb(\theta, t) = \frac{1}{\theta t} \log E\left(e^{\theta A(t)}\right).$$

The effective bandwidth of the self-similar process in Eq. (35) is given by $eb(\theta, t) = \rho + \frac{1}{2}b^2\theta t^{2H-1}$. Let us consider more generally processes whose effective bandwidth satisfies the self-similar bound

$$eb(\theta, t) \leq \rho + \frac{1}{2}b^2\theta t^{2H-1}. \quad (37)$$

By the Chernoff bound, Eq. (37) implies that

$$Pr\left(A(s, t) > r(t-s) + \sigma(t-s)^H\right) \leq e^{-\frac{1}{2}(\sigma/b)^2},$$

i.e., Eq. (34) with $K = 1$ is a statistical envelope. Note that for $H > \frac{1}{2}$ and any given choice of θ , Eq. (37) does not define a linear bounded envelope process as defined in [9].

Lemma 3: SAMPLE PATH ENVELOPE FOR SELF-SIMILAR PROCESSES WITH GAUSSIAN TAIL. Assume that arrivals to a flow are bounded by a statistical envelope, given by

$$\mathcal{G}(t; \sigma) = rt + \sigma t^H, \quad \varepsilon(\sigma) = K e^{-\frac{1}{2}(\sigma/b)^2}. \quad (38)$$

Then, for any choice of $\mu > 0$, a statistical sample path envelope for the arrival process is given by

$$\bar{\mathcal{G}}(t; \sigma) = (r + \mu)t + \sigma, \quad \bar{\varepsilon}(\sigma) = L e^{-(\sigma/c)^\beta},$$

with parameters $\beta = 2(1 - H)$, $c = \left(\frac{2b}{\mu^H}\right)^{\frac{1}{1-H}}$, and $L = e \cdot \max\left\{1, 4^H K \frac{r/\mu + 2 - H}{H(1-H)}\right\}$.

PROOF. Fix $\mu > 0$, let $1 < \gamma < 1 + \frac{\mu}{r}$, and set $b = \frac{r+\mu}{\gamma} - r > 0$. As done for the backlog bound from Eq. (17), we argue that

$$\begin{aligned} Pr\left(\sup_{s \leq t} \left\{A(s, t) - \bar{\mathcal{G}}(t - s; \sigma)\right\} > 0\right) &\leq \sum_{k=-\infty}^{\infty} \varepsilon\left(\frac{\sigma + bx_k}{x_k^H}\right) \\ &\leq \frac{1}{H(1-H) \log \gamma} \int_z^{\infty} \frac{\varepsilon(x)}{x} dx \\ &\leq \frac{K}{H(1-H) \log \gamma} z^{-2} e^{-\frac{1}{2}z^2}, \end{aligned} \quad (39)$$

where $z = \frac{b^H \sigma^{1-H}}{\gamma^H (1-H)}$. Here, the first inequality follows by discretization. In the second step, we have used Lemma 6 and replaced the integration variable x with x/b . In the third step, we have evaluated the integral.

We want to replace the variable z by $y = \mu^H \sigma^{1-H} / b$, so that $\frac{1}{2}y^2 = (\sigma/c)^\beta$. Suppose for the moment that $y^2 \geq 2$, and choose γ such that $\log \gamma = (r/\mu + 2 - H)^{-1} y^{-2}$. Then $\gamma < 1 + r/\mu$, as required. Moreover,

$$z^2 = y^2 \left(\frac{b}{\mu \gamma^{1-H}}\right)^{2H} \geq y^2 (1 - (r/\mu + 2 - H) \log \gamma)^{2H} \geq \frac{1}{2} y^2,$$

and

$$z^2 \geq y^2 (1 - 2(r/\mu + 2 - H) \log \gamma) \geq y^2 - 2,$$

and we obtain the bound

$$Pr\left(\sup_{s \leq t} \left\{A(s, t) - \bar{\mathcal{G}}(t - s; \sigma)\right\} > 0\right) \leq L e^{-(\sigma/c)^\beta}.$$

This proves the claim for $y^2 \geq 2$. On the other hand, for $y^2 < 2$, the claim holds trivially since the right hand side exceeds 1. \square

As in Section IV-B, the sample path envelope immediately yields a leftover service curve. If fluid-flow traffic arrives to a link of constant rate C , where it is subject to cross traffic that has a statistical envelope with Gaussian tails, as in Eq. (34), then for every choice of $\mu > 0$,

$$\mathcal{S}(t; \sigma) = [(C - r_c - \mu)t - \sigma]_+, \quad \varepsilon(\sigma) = L e^{-(y/c)^\beta} \quad (40)$$

provides a statistical service curve for the through flow with a Weibullian bound on the violation probability. Here, the parameters β , c , and L are given by Lemma 3. Next, we show how to concatenate such service curves:

Lemma 4: CONCATENATION OF WEIBULL-TAILED SERVICE CURVES. Consider two consecutive nodes on the path of a flow through a network. Assume that the first node offers a statistical service curve $\mathcal{S}_1(t; \sigma) = [R_1 t - \sigma]_+$ with $\varepsilon_1(\sigma) = L_1 e^{-(\sigma/c_1)^{\beta_1}}$, and the second node offers a statistical service curve $\mathcal{S}_2(t; \sigma) = [R_2 t - \sigma]_+$ with an arbitrary function $\varepsilon_2(\sigma)$. Then, for every choice of $\gamma > 1$, the two nodes offer the combined service curve given by

$$\begin{aligned} \mathcal{S}(t; \sigma) &= [\min\{R_1, R_2/\gamma\} - \sigma]_+, \\ \varepsilon(\sigma) &= \inf_{\sigma_1 + \sigma_2 = \sigma} \left\{ \tilde{L}_1 e^{-(\sigma_1/c_1)^{\beta_1}} + \varepsilon_2(\sigma_2) \right\}, \end{aligned}$$

where $\tilde{L}_1 = \max \left\{ e^2, \frac{\gamma}{\gamma-1} \frac{(2e)^{[\beta_1-1]_+}}{c_1} L_1 \right\}$.

PROOF. We follow the proof of Lemma 2, with the Weibullian function $\varepsilon_1(\sigma) = L_1 e^{-(\sigma/c_1)^{\beta_1}}$ in place of the power law. Set $R = \min\{R_1, R_2/\gamma\}$, $\delta = R_2 - R$, and consider the event

$$D_1(t - y_k) \geq A_1 * \mathcal{S}_1(t - y_k; \sigma_1 + \delta y_k - R_2 \tau), \quad (41)$$

where y_k is a sequence of discretization time steps. We argue as in the proof of Lemma 2 that \mathcal{S} is a service curve, with violation probability given by

$$\varepsilon(\sigma) \leq \inf_{\sigma_1 + \sigma_2 = \sigma} \left\{ Pr(\text{Eq. (41) fails for some } k \text{ with } y_k \leq t) + \varepsilon_2(\sigma_2) \right\}.$$

We could now use Lemma 7 to bound the violation probability of Eq. (41). However, since $\varepsilon_1(\sigma_1)$ is an integrable function, we can obtain a slightly stronger bound by using the arithmetic sequence $y_k = k\tau$, as in [10]. The resulting estimate is

$$\begin{aligned} Pr(\text{Eq. (41) fails for some } k \text{ with } y_k \leq t) &\leq \sum_{k=1}^{\infty} \varepsilon_1(\sigma_1 + \delta y_k - R_2 \tau) \\ &\leq \frac{1}{\delta \tau} \int_{\sigma_1 - R_2 \tau}^{\infty} \varepsilon_1(x) dx \\ &\leq L_1 \frac{c_1}{\beta_1 \delta \tau} z^{-(\beta_1-1)} e^{-z^{\beta_1}} \Bigg|_{z=\frac{\sigma_1 - R_2 \tau}{c_1}} \end{aligned}$$

We have first used the union bound and the ht service curve, then replaced the sum by an integral, and finally evaluated the integral. Suppose for the moment that $(\sigma_1/c_1)^{\beta_1} \geq 2$. We choose $\tau = \frac{c_1^{\beta_1}}{\beta_1 R_2 \sigma_1^{\beta_1-1}}$ and use that

$$(\sigma_1 - R_2 \tau)^{\beta_1} \geq \sigma_1^{\beta_1} - \max\{1, \beta\} R_2 \tau \sigma_1^{\beta_1-1}, \quad \frac{\sigma_1}{\sigma_1 - \tau} \leq 2$$

to obtain

$$Pr(\text{Eq. (41) fails for some } k \text{ with } y_k \leq t) \leq \tilde{L}_1 e^{-(\sigma_1/c_1)^{\beta_1}}.$$

But this inequality clearly also holds for $(\sigma_1/c_1)^{\beta_1} < 2$. The lemma follows by combining this with the service guarantee at the second node. \square

With this result, we finally present a network service curve that concatenates an arbitrary number of Weibullian service curves.

Theorem 3: WEIBULLIAN 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 a service curve

$$\mathcal{S}_n(t, \sigma) = [Rt - \sigma]_+, \quad \varepsilon(\sigma) = Le^{-(\sigma/c)^\beta}.$$

Then, for every choice of $\gamma > 1$, the network provides the service curve

$$\begin{aligned} \mathcal{S}_{net}(t, \sigma) &= [(R/\gamma)t - \sigma]_+, \\ \varepsilon_{net}(\sigma) &\leq N \left(1 + \frac{N}{\log \gamma}\right) \tilde{L} e^{-(\sigma/Nc)^\beta}, \end{aligned}$$

where $\tilde{L} = \max \left\{ \frac{e^2}{N} \log \gamma, \frac{(2e)^{\lfloor \beta-1 \rfloor_+}}{c} L \right\}$.

PROOF. See the proof of Theorem 2. We use Lemma 4 to iteratively derive the service guarantee for the last n nodes. In each step, we reduce the rate by a factor of $\gamma^{\frac{1}{N-1}}$. To simplify the formula, we have replaced $N-1$ with N whenever it appeared convenient, and used that $\frac{N}{\log \gamma} \leq \frac{\gamma^{\frac{1}{N}}}{\gamma^{\frac{1}{N}-1}} \leq 1 + \frac{N}{\log \gamma}$. \square

Consider now a network as in Figure 6 where all flows have envelopes with a Gaussian tail, given by Eq. (33). Assume that the network is homogeneous, i.e., all nodes have the same link rate, and the cross traffic has the same parameters at each node. Combining the results from this section, we obtain, similarly to the analysis in Section V, the end-to-end delay bound

$$Pr(W_{net}(t) > w) = O(N^2 e^{-(\frac{w}{Nw_0})^\beta}), \quad (w \rightarrow \infty),$$

where M and w_0 are constants that do not depend on N , and $\beta = 2(1-H)$. It follows that, over long paths, the quantiles of the delays scale as

$$w_{net}(\varepsilon) = O(N(\log N)^{1/\beta}), \quad (N \rightarrow \infty).$$

For $\beta = 1$, this bounds recovers the $O(N \log N)$ bound for exponential traffic from [10].

II. TECHNICAL LEMMAS

In our derivations, we frequently use properties of the function $\varepsilon(\sigma) = K\sigma^{-\alpha}$ that appears in the definition of the *htss* envelope. The properties are summarized in the following lemma, and presented without proof.

Lemma 5:

1. (*Lower power.*) We can lower the power by using that for $K\sigma^{-\alpha} \leq 1$ and $\alpha' < \alpha$

$$K\sigma^{-\alpha} \leq K\frac{\alpha'}{\alpha}\sigma^{-\alpha'}. \quad (42)$$

2. (*Eliminate logarithmic factor.*) For $\beta' < \beta$,

$$\sigma^{-\beta} \log \sigma \leq \frac{1}{e(\beta - \beta')} \sigma^{-\beta'}. \quad (43)$$

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

$$K(\sigma - \sigma_0)^{-\alpha} \leq 2^{[\alpha-1]_+} (K + \sigma_0^\alpha) \sigma^{-\alpha}. \quad (44)$$

4. (*Minimize sum.*) We can minimize sums of such functions by

$$\min_{\sigma_1 + \dots + \sigma_n = \sigma} \sum_{j=1}^n K_j \sigma_j^{-\alpha} = \left(\sum_{j=1}^n K_j^{\frac{1}{1+\alpha}} \right)^{1+\alpha} \sigma^{-\alpha} \leq n^\alpha \bar{K} \sigma^{-\alpha}, \quad (45)$$

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

The following derives auxiliary estimates for two sums that involve geometric sequences.

Lemma 6: Assume that $\varepsilon(x)$ is a nonincreasing nonnegative function. Fix $\gamma > 1$ and $\tau > 0$, and set $x_k = \tau\gamma^k$. Then, for every $\sigma \geq 0$ and every $c > 0$,

$$\sum_{k=-\infty}^{\infty} \varepsilon\left(\frac{\sigma + cx_k}{x_k^H}\right) \leq \frac{1}{H(1-H)\log\gamma} \int_z^{\infty} \frac{\varepsilon(x)}{x} dx \Bigg|_{z=\frac{cH\sigma^{1-H}}{\gamma^{H(1-H)}}}. \quad \square$$

PROOF. Consider first the case where $c = \tau = 1$, i.e., $x_k = \gamma^k$. Each summand in the series satisfies

$$\varepsilon\left(\frac{\sigma + \gamma^k}{\gamma^{Hk}}\right) \leq \min\{\varepsilon(\sigma\gamma^{-Hk}), \varepsilon(\gamma^{(1-H)k})\}.$$

Since the first term on the right hand side increases with k while the second term decreases, we can bound the series by the sum of two integrals

$$\sum_{k=-\infty}^{\infty} \varepsilon\left(\frac{\sigma + \gamma^k}{\gamma^{Hk}}\right) \leq \int_{-\infty}^{T+1} \varepsilon(\sigma\gamma^{-Ht}) dt + \int_T^{\infty} \varepsilon(\gamma^{(1-H)t}) dt,$$

where the overlap between the intervals of integration compensates for the change of monotonicity. The optimal choice for the limit of integration is $T = -H + \frac{\log \sigma}{\log \gamma}$, so that $\sigma \gamma^{-H(T+1)} = \gamma^{(1-H)t}$. In the first integral, the change of variables $x = \sigma \gamma^{-Ht}$ yields

$$\int_{-\infty}^{T+1} \varepsilon(\sigma \gamma^{-Ht}) dt = \frac{1}{H \log \gamma} \int_z^{\infty} \frac{\varepsilon(x)}{x} dx,$$

where $z = \sigma^{1-H} \gamma^{-H(1-H)}$. In the second integral, the change of variables $x = \gamma^{(1-H)t}$ yields

$$\int_T^{\infty} \varepsilon(\gamma^{(1-H)t}) dt = \frac{1}{(1-H) \log \gamma} \int_z^{\infty} \frac{\varepsilon(x)}{x} dx,$$

Adding the two integrals proves the claim for $a = \tau = 1$. For other values of c and τ , we rescale $\sigma = c\tau^{1-H} \tilde{\sigma}$, and apply the first case to the function $\tilde{\varepsilon}(x) = \varepsilon(c\tau^{1-H}x)$. \square

Lemma 7: Assume that $\varepsilon(x)$ is a nonincreasing nonnegative function. Fix $\tau > 0$ and $\gamma > 1$, and define recursively $y_0 = 0$, $y_k = \tau + \gamma y_{k-1}$. Then, for every $c > 0$ and $\sigma \geq \frac{c\tau}{\gamma-1}$,

$$\sum_{k=1}^{\infty} \varepsilon(\sigma + cy_k) \leq \frac{1}{\log \gamma} \left(\varepsilon(z) \log \left(\frac{\gamma-1}{a\tau} z \right) + \int_z^{\infty} \frac{\varepsilon(x)}{x} dx \right) \Bigg|_{z=\sigma+c\tau}.$$

PROOF. Consider first the case where $c = 1$ and $\tau = \gamma - 1$, i.e., $y_k = \gamma^k - 1$, and set $z = \sigma + \gamma - 1$. For $\sigma \geq 1$, each summand is bounded by

$$\varepsilon(\sigma + \gamma^k - 1) \leq \min \left\{ \varepsilon(\sigma + \gamma - 1), \varepsilon(\gamma^k) \right\}.$$

Since both terms are nonincreasing, we can bound the series by

$$\sum_{k=1}^{\infty} \varepsilon(\sigma + \gamma^k - 1) \leq \int_0^T \varepsilon(\sigma + \gamma - 1) dt + \int_T^{\infty} \varepsilon(\gamma^t) dt.$$

We choose $T = \frac{\log(\sigma + \gamma - 1)}{\log \gamma} \geq 1$, so that $\gamma^T = \sigma + \gamma - 1$, and change variables $x = \gamma^t$ in the second integral to obtain

$$\sum_{k=1}^{\infty} \varepsilon(\sigma + \gamma^k - 1) \leq \frac{1}{\log \gamma} \left(\varepsilon(z) \log z + \int_z^{\infty} \frac{\varepsilon(x)}{x} dx \right) \Bigg|_{z=\sigma+\gamma-1}.$$

This proves the claim in the special case $a = 1$, $\tau = \gamma - 1$. For other values of c and τ , we rescale $\sigma = \frac{c\tau}{\gamma-1} \tilde{\sigma}$, $z = \frac{c\tau}{\gamma-1} \tilde{z}$, and apply the first case to $\tilde{\varepsilon}(x) = \varepsilon\left(\frac{c\tau}{\gamma-1}x\right)$. \square