# Integrating Fractional Brownian Motion Arrivals into the Statistical Network Calculus 

Paul Nikolaus*, Sebastian Henningsen ${ }^{\dagger}$, Michael Beck $^{\ddagger}$ and Jens Schmitt*<br>*Distributed Computer Systems (DISCO) Lab<br>TU Kaiserslautern, Germany<br>\{nikolaus,jschmitt\}@cs.uni-kl.de<br>${ }^{\dagger}$ Humboldt-Universität zu Berlin, Germany<br>sebastian.henningsen@hu-berlin.de<br>$\ddagger$ University of Manitoba, Canada<br>michaelalexbeck@gmail.com


#### Abstract

Stochastic network calculus (SNC) is a versatile framework to derive probabilistic performance bounds. Recently, it was proposed in [1] to replace the typical a priori assumptions on arrival processes with measurement observations and to incorporate the corresponding statistical uncertainty into calculation of the bounds. This so-called statistical network calculus (StatNC) opens the door for many applications with limited traffic information. However, the important traffic class of self-similar processes such as fractional Brownian Motion (fBm) was left open in [1], thus, e.g., depriving the usage of the StatNC for Internet traffic. In this work, we close this gap by integrating fBm arrivals into the StatNC. To this end, we analyze the impact imposed by the uncertainty on the backlog bound and show in numerical evaluations that the additional inaccuracy is only of moderate size.


## I. Introduction

The framework of stochastic network calculus [2]-[7] is a versatile methodology to compute probabilistic performance bounds. It originated in the work on deterministic bounds by Cruz [8], [9]. Quickly after this, violation probabilities in order to benefit from statistical multiplexing were introduced [10][12]. Many different approaches for this stochastic extension have been proposed since, yet we focus on the seminal work of Chang's calculus with moment generating function (MGF) [2].

Network calculus' flexibility stems from its abstraction into arrival and service curves making it simple to incorporate different traffic models and scheduling disciplines. For instance, it can deal with arrivals as diverse as exponentially bounded burstiness [4], Markov-modulated processes [4], self-similar processes such as fractional Brownian motion [13], [14], or even heavy-tailed traffic [15], just to mention a few.

All of these approaches assume a priori knowledge about the system, in particular the arrival processes. In contrast, the framework of StatNC [1] releases these assumptions. Instead, StatNC provides a methodology to estimate the incoming arrival process based on measurement observations and to include the arising measurement uncertainty into the performance bounds. In [1], the StatNC framework was demonstrated to work with several traffic classes, yet, the case of self-similar and long-range dependent traffic was explicitly
left for future work (see Section IV.D in [1]). Clearly, this open question in StatNC is very relevant to the application of the framework in, e.g., an Internet scenario where selfsimilarity and long-range dependence have been observed to be ubiquitous [16]-[20].

In this paper, we aim to fill this important gap in the StatNC. To that end, we propose an estimator of the arrivals' MGF with fractional Brownian motion [21] in the framework of StatNC. In particular, we estimate the self-similarity parameter, the so-called Hurst parameter, which characterizes the behavior of a fractional Brownian motion process. By providing a confidence interval for the Hurst parameter, we are able to derive an upper bound on the MGF of fBm arrivals that we subsequently use to derive a backlog bound. We complement the theoretical analysis with a simulation-based evaluation. The empirical results indicate that our StatNC backlog bound is reasonably close to the traditional SNC bound. This not only validates our proposed estimator but also highlights the value of StatNC.
The rest of the paper is structured as follows. First, we discuss related work in II, followed by a brief collection of necessary prerequisites in Section III. In Section IV, we present the estimator for fBM traffic and derive a statistical backlog bound. We also show how the necessary preconditions for the application of the StatNC can be met. The estimation quality is then evaluated in Section V. Section VI concludes the paper.

## II. Related Work

The StatNC in [1] lays the foundation for our work, in which we aim to integrate the important traffic class of fractional Brownian motion ( fBm ) arrivals. The fBm arrival process has seen some treatment in the general field of network calculus (though not in the StatNC framework). [13] derives an effective envelope by making use of effective bandwidth [21]. The work of [14] derives a sample path envelope for the same traffic class together with end-to-end performance bounds under fBm cross-traffic. [15] stands out in this context as it introduces heavy-tailed self-similar arrival in conjunction with
heavy-tailed service to derive sample path bounds for multinode networks. Moreover, it estimates the Hurst parameter of a 24 -hour backbone network trace in order to provide a statistical envelope. Yet, as it was pointed out in [1], the uncertainty that stems from the parameter estimation is not incorporated into the stochastic bound, which is exactly what we target at in this paper.

Another related work can be found in [22], where, in the same spirit as in [1], it is proposed to accommodate statistical dependencies between traffic flows using copula analysis based on measurement observations. Yet, the treatment of the fBm traffic class is not targeted in [22].

## III. BACKGROUND

## A. Network Calculus with Moment Generating Functions

We use the moment-generating function (MGF)-based SNC in order to calculate backlog bounds. To be precise, we bound the probability that the backlog exceeds a given value, typically denoted by $b$. The MGF bound on a probability is established by applying Chernoff's bound

$$
\mathrm{P}(X>a) \leq e^{-\theta a} \mathrm{E}\left[e^{\theta X}\right]
$$

with $\mathrm{E}\left[e^{\theta X}\right]$ as the MGF of a random variable $X$ and the free parameter $\theta>0$. Further, we denote $\phi_{X}(\theta, s, t):=$ $\mathrm{E}\left[e^{\theta X(s, t)}\right]$.

Definition 1 (Arrival Process). We define an arrival flow by the stochastic process $A$ with discrete time space $\mathbb{N}$ and continuous state space $\mathbb{R}_{0}^{+}$as

$$
\begin{equation*}
A(s, t):=\sum_{i=s+1}^{t} a(i) \tag{1}
\end{equation*}
$$

with $a(i)$ as the traffic increment process in time slot $i$.
Network calculus provides an elegant system-theoretic analysis by employing min-plus algebra.

The characteristics of the service process are captured by the notion of a dynamic $S$-server.

Definition 2 (Dynamic $S$-Server). Assume a service element has an arrival flow $A$ as its input and the respective output is denoted by $A^{\prime}$. Let $S(s, t), 0 \leq s \leq t$, be a stochastic process that is non-negative and increasing in $t$. The service element is a dynamic $S$-server iff for all $t \geq 0$ it holds that:

$$
A^{\prime}(0, t) \geq \inf _{0 \leq s \leq t}\{A(0, s)+S(s, t)\}
$$

Throughout this paper, we assume the service element to be a dynamic $S$-server that is work-conserving.

Theorem 3 (Backlog Bound). [5], [23] Let $q(t):=A(t)-$ $A^{\prime}(t)$ be the backlog at time $t, S(s, t)$ be the service provided by a dynamic $S$-server. Then it holds that

$$
\begin{equation*}
q(t) \leq \sup _{0 \leq s \leq t}\{A(s, t)-S(s, t)\} \tag{2}
\end{equation*}
$$

## B. StatNC - A Framework For Arrival-Estimation

Stochastic network calculus is based on a priori knowledge of the involved arrival and service processes. In practical applications, this knowledge may often not be available and therefore has to be estimated through measurements. Statistical estimation, however, introduces a new source of uncertainty that, conventionally, has not been taken into account when performance bounds are obtained with SNC. In order to obtain more precise bounds, this systematic uncertainty should be included in the calculation.

As a first step in this direction, the framework of StatNC [1] was proposed, which enables the estimation of arrival processes while incorporating the arising uncertainty into the backlog bound. In simple terms, StatNC provides a set of different estimators for various classes of arrival processes (exponential, On/Off MMAPs, ...). Each estimator has to provide an upper bound on the arrivals' MGF with probability at least $1-\alpha$. Taking the upper bound is necessary to guarantee valid performance bounds, since the backlog bound is increasing in the arrival processes' MGF.
Technically, the estimator $\phi_{A}(\theta, \cdot, \cdot)$ of the MGF has to be from the set of all functions with $\left\{f \mid f: \mathbb{R}^{+} \times \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{R}^{+}\right\} \quad=: \quad \mathcal{F}$. We use this definition to introduce a statistic on a sample of size $\left|t_{0}\right|, a=\left(a_{-t_{0}}, \ldots, a_{-1}\right)$. This sample represents the measured observations we are given for estimation.

Theorem 4. [1] Let $\theta^{*}=\sup \left\{\theta: \phi_{A}(\theta, s, t)<\infty\right\}$ and $\Phi: \mathbb{R}^{\left|t_{0}\right|} \rightarrow \mathcal{F}$ be a statistic on $a=\left(a_{-t_{0}}, \ldots, a_{-1}\right)$ such that

$$
\sup _{\theta \in\left(0, \theta^{*}\right)} \mathrm{P}\left(\bigcup_{s \leq t} \Phi_{a}(\theta, s, t)<\phi_{A}(\theta, s, t)\right) \leq \alpha
$$

Then for all $t \in \mathbb{N}_{0}, \theta<\theta^{*}$

$$
\mathrm{P}(q(t)>b) \leq \alpha+e^{-\theta b} \sum_{s=0}^{t} \Phi_{a}(\theta, s, t) \mathrm{E}\left[e^{-\theta S(s, t)}\right]
$$

This means that for the backlog bound, we can replace the arrivals' MGF with its estimate if we can ensure that the estimate is a probabilistic upper bound on the MGF. That is, the error of underestimation is bounded by probability $\alpha$. Note that the backlog bound given in this theorem is only valid for a finite time horizon. This is often of disadvantage in general, yet in this case, it emphasizes the dynamic nature of StatNC. We only have a finite set of observation data and therefore cannot predict the entire future but rather make an update on our information by a new estimation, e.g., based on a sliding window procedure (see [1] for more details).

## IV. Integrating Fractional Brownian Motion into StatNC

## A. FBM Definition and Confidence Interval for Hurst Param-

 eterIn this section, we derive an MGF backlog bound for fBm arrivals. A statistic on the MGF is provided and is shown to satisfy the requirements of Theorem 4.

Definition 5 (Fractional Brownian Motion). [24], [25] A stochastic process $Z(t)$ is called normalized fractional Brownian motion ( $f B m$ ) with (self-similarity) Hurst parameter $H \in$ $(1 / 2,1)$, if it can be characterized by the following properties:

- $Z(t)$ has stationary increments,
- $Z(0)=0$ and $\mathrm{E}[Z(t)]=0$ for all $t$,
- $\mathrm{E}\left[Z(t)^{2}\right]=|t|^{2 H}$ for all $t$,
- $Z(t)$ has continuous paths,
- $Z(t)$ is Gaussian, i.e., all its finite-dimensional marginal distributions are Gaussian.
The increments of this process $Z(t+1)-Z(t)$ are called fractional Gaussian noise (fGn).

The application of StatNC requires a statistic on the MGF in order to obtain the StatNC backlog bound. As we see in Subsection IV-B, a confidence interval for the Hurst parameter $H$ is key to obtain this statistic. A variety of methods to compute estimators $\hat{H}$ are known (see for example [26]). Yet, not all of them enables the calculation of a confidence interval. In this paper, we opted for a periodogram-based maximum likelihood estimate that yields confidence intervals [27]. The book of [28] provides well-tested code in $\mathbf{S +}$.

Theorem 6 (Confidence Interval for Hurst Parameter). [29] Let $f(l, \Theta)$ be the spectral density of fractional Gaussian noise for a sample of size $n, \mathbf{x}=\left(x_{1}, \ldots, x_{n}\right)$, where the parameter vector of the process $\Theta=\left(\Theta_{1}, \Theta_{2}\right)$ is structured as follows: $\Theta_{1}$ is a scale parameter (see [29] for more details) and $\Theta_{2}$ is the Hurst parameter. $A(1-\alpha)$-confidence interval for the estimate of $H$ can be obtained by

$$
\hat{H} \pm q_{1-\frac{\alpha}{2}} \cdot \sqrt{\frac{V_{11}}{n}}
$$

where $V=2 D^{-1}$, matrix $D$

$$
D_{i j}:=\frac{1}{2 \pi} \int_{-\pi}^{\pi} \frac{\partial}{\partial \Theta_{i}} \log f(l) \frac{\partial}{\partial \Theta_{j}} \log f(l) \mathrm{d} l
$$

and $q_{1-\frac{\alpha}{2}}$ is the $\left(1-\frac{\alpha}{2}\right)$-quantile of the normal distribution.
The matrix $1 / 2 D$ is known as the asymptotic Fisher information matrix [29] guaranteeing asymptotic efficiency for Gaussian data.

We assume the arrivals to be of the form (as in [13], [25])

$$
\begin{equation*}
A(t):=\lambda t+\sigma Z(t), \tag{3}
\end{equation*}
$$

where $Z(t)$ is a normalized fractional Brownian motion, $\lambda$ is the mean arrival rate and $\sigma^{2}$ is the variance of $A(1)$. The fBm $Z(t)$ is governed by the Hurst parameter $H$ that, in turn, is independent of the other parameters [25]. This arrival process has been shown to have the following MGF [21]:

$$
\begin{equation*}
\mathrm{E}\left[e^{\theta A(t)}\right]=e^{\lambda \theta t+\frac{\theta^{2} \sigma^{2}}{2} t^{2 H}} \tag{4}
\end{equation*}
$$

Given the continuous nature of fBm , the arrivals (3) are a continuous-time process

$$
A(s, t)=\int_{s}^{t} a(x) \mathrm{d} x, \text { where } 0 \leq s \leq t
$$

that has to be discretized in order to be applicable to our discrete-time arrival model (cf. Equation (1)). In the following, we derive a backlog bound for this arrival model. The difference to existing network calculus results is that we have to derive it inside the StatNC framework set up in Theorem 4. The denotation as SNC backlog bound is just for an easier comparison to the StatNC bound below and is not meant to make any pretenses on its tightness in comparison to other performance bounds for this traffic class. Again, its nonstationarity is justified by the dynamic nature of StatNC over a finite-time observation period.
Theorem 7 (SNC Backlog Bound for FBM Traffic). Let $q(t)$ be the backlog at time $t$. We assume that the fBm arrivals $A$ with rate $\lambda$ and variance $\sigma^{2}$ are independent of the service $S$. We have a work-conserving server with constant rate $C$ and the stability condition $C>\lambda$ holds. Moreover, we assume that $b>\lambda \tau$ with arbitrary discretization parameter $\tau>0$. Then we have that

$$
\begin{equation*}
\mathrm{P}(q(t)>b) \leq \sum_{k=1}^{\left\lfloor\frac{t}{\tau}\right\rfloor+1} e^{-\frac{(b-C \tau+(C-\lambda) k \tau)^{2}}{2 \sigma^{2}(k \tau)^{2 H}}} \tag{5}
\end{equation*}
$$

Proof: See Appendix A.

## B. A Statistic for FBM Arrivals

For the application of StatNC, our statistic has to be an upper bound on the MGF of the arrival process due to the monotonic behavior of the backlog bound as mentioned in [1]. To that end, it suffices to provide a probabilistic upper bound on the estimated Hurst parameter (as it is given in Theorem 6).

In the following, we give a statistic for the fBm arrival model (cf. Equation (3)) which provides an upper bound on the MGF with probability $1-\alpha$. This constitutes the basis for the application of the StatNC framework.
Theorem 8. Let the confidence interval's upper endpoint of a given sample set $\mathrm{x}=\left(x_{1}, \ldots, x_{n}\right)$ be defined as

$$
\begin{equation*}
\hat{H}_{1-\alpha}^{\text {up }}:=\hat{H}+q_{1-\alpha} \cdot \sqrt{\frac{V_{11}}{n}} . \tag{6}
\end{equation*}
$$

Then the statistic

$$
\begin{equation*}
\Phi_{a}(\theta, s, t):=e^{\lambda \theta(t-s)+\frac{\theta^{2} \sigma^{2}}{2}(t-s)^{2 \hat{H}_{1-\alpha}^{\mathrm{up}}}} \tag{7}
\end{equation*}
$$

## fulfills the condition of Theorem 4.

Note, that we only need the one-sided confidence interval since we are interested in quantifying the probability of underestimating the MGF. That is, we can replace the $1-\frac{\alpha}{2}$ by $1-\alpha$, which we also call the quantile of $\hat{H}$.

Proof: It can easily be shown that the MGF (4) is increasing in $H$. Subsequently, in case of $\hat{H}_{1-\alpha}^{\text {up }} \geq H$, this yields

$$
\Phi_{a}(\theta, s, t) \geq \phi_{A}(\theta, s, t)=e^{\lambda \theta(t-s)+\frac{\theta^{2} \sigma^{2}}{2}(t-s)^{2 H}}
$$

Thus, we conclude that


Fig. 1. Backlog distribution with quantile (blue) and SNC (red) / StatNC bound (black) for a quantile $=0.998,1-\alpha=0.999, \lambda=10^{-2}, H=$ $0.7, \sigma=1$, utilization $=\lambda / C=2 / 3$, sample paths of length $2^{16}$, and a time horizon of $t=200$. The dashed green lines represent the $95 \%$ quantile of the StatNC bounds in these 500 iterations. The estimated Hurst parameter is $\hat{H}=0.703$ with confidence interval $(0.695,711)$.

$$
\begin{aligned}
1-\alpha & =\mathrm{P}\left(\hat{H}_{1-\alpha}^{\mathrm{up}} \geq H\right) \\
& \leq \inf _{\theta>0}\left\{\mathrm{P}\left(\bigcap_{s \leq t} \Phi_{a}(\theta, s, t) \geq \phi_{A}(\theta, s, t)\right)\right\}
\end{aligned}
$$

which, in turn, is equivalent to the necessary condition in Theorem 4:

$$
\sup _{\theta>0}\left\{\mathrm{P}\left(\bigcup_{s \leq t} \Phi_{a}(\theta, s, t)<\phi_{A}(\theta, s, t)\right)\right\} \leq \alpha
$$

The result of Theorem 8 is crucial for our approach of embedding the fBm arrival estimation into the framework of StatNC. It enables us to provide a StatNC backlog bound with the quantile of $\hat{H}$ by simply combining Theorem 4 and Theorem 8.

Corollary 9 (StatNC Backlog Bound for FBM Traffic). Assume the same setting as in Theorem 7 except that the Hurst parameter $H$ is now unknown. Then we have for $b>\lambda \tau$

$$
\begin{equation*}
\mathrm{P}(q(t)>b) \leq \alpha+\sum_{k=1}^{\left\lfloor\frac{t}{\tau}\right\rfloor+1} e^{-\frac{(b-C \tau+(C-\lambda) k \tau)^{2}}{2 \sigma^{2}(k \tau)^{2 \hat{H}_{1-\alpha}^{u p}}}} \tag{8}
\end{equation*}
$$

## V. Evaluation

In the following, we evaluate the quality of the StatNC backlog bound in various scenarios. To that end, we simulate a single constant rate server with sample paths of the fBm arrival process (cf. Equation (3)) and repeatedly observe the backlog of the system. The backlog from the simulations is compared to the SNC and StatNC backlog bounds. The SNC bound is simply the computation of Equation (5), whereas the StatNC bound (provided in Equation (8)) requires the computation of the estimator $\hat{H}$ and its confidence interval. As discussed above, we opted for an approach based on periodograms. We
chose this approach because of its ability to provide confidence intervals and its tendency to overestimate the Hurst parameter in our simulation giving us a "safer" backlog bound. The periodogram is defined as

$$
I\left(\omega_{n}\right):=\frac{1}{2 \pi n}\left|\sum_{k=1}^{n} x_{k} e^{i k \omega_{k}}\right|^{2}
$$

for a given data vector $\mathbf{x}=\left(x_{1}, \ldots, x_{n}\right)$ with frequency $\omega_{k}=2 \pi k / n$. The basic idea is to plot $I\left(\omega_{n}\right)$ in a log-log grid and estimate its slope. This should give an estimate of $1-2 H$ [27]. The confidence interval is then obtained via Theorem 6. As both backlog bounds (Equation (5) and Equation (8)) only yield the backlog's violation probability, the bound on the system's backlog is obtained by numerical inversion of the probability bound. We conducted the simulations in the statistical programming language $\mathbf{R}$ [30], version 3.5.

First, we consider the tightness of the backlog bound as well as the additional inaccuracy that is implied by eliminating the a priori assumptions using StatNC. For each simulation, the observed cumulative backlog is captured, effectively yielding an empirical backlog distribution at this point in time. Additionally, we repeatedly compute the StatNC backlog bound, since it varies depending on the arrival samples used for estimation.

The results are shown in Figure 1. It depicts the empirical backlog distribution with the backlogs on the $x$-axis and their respective relative frequencies on the $y$-axis. The vertical blue line indicates the distribution's 0.998 -quantile, the red line the SNC backlog bound and the black one the mean of the StatNC bound. Both bounds use the same 0.998-quantile; the StatNC estimates the Hurst parameter with confidence level $1-\alpha=$ 0.999. The dashed greed lines indicate the empirical $95 \%$ quantile of the StatNC bounds.

We observe the gap between the empirical quantile and the SNC to be of non-negligible size. This confirms previous observations of additional inaccuracy for traffic with high correlations when using standard SNC methods such as the Union bound [31], [32]. On the other hand, the StatNC gap to the standard SNC is moderate, even for all bounds in the $95 \%$ quantile. Our evaluation also indicates that the Hurst parameter estimation, and thus the StatNC bound, is very sensitive to path lengths and the chosen estimation approach. The StatNC's quality might be further enhanced by using more advanced techniques than periodograms.

Beck et al. [1] also analyzed the robustness of StatNC. That is, for a capped i.i.d. Pareto distributed traffic, both the SNC and the StatNC backlog bound were computed under the false assumption of capped i.i.d. exponentially distributed arrivals (with the same mean, though). It was shown that the standard SNC bound clearly underestimated the bound due to the assumption of a very light tail. However, StatNC managed to provide a valid upper bound using a non-parametric estimator to derive the MGF based on the observed traffic. We conduct a similar experiment by applying the two StatNC estimators to the fBm simulations from above. For one bound, we use the


Fig. 2. Backlog distribution and StatNC bounds for quantile $=0.998,1-\alpha=$ $0.999, \lambda=10^{-2}, H=0.7, \sigma=1$, utilization $2 / 3$, sample paths of length $2^{16}$, and a time horizon of $t=200$ for 500 iterations.
correct assumption as in Figure 1, and for the other, we use the non-parametric estimator from [1], hence, wrongly assuming i.i.d. arrivals.

The results are shown in Figure 2. It can be observed that the non-parametric estimator clearly underestimates the distribution, since the arrivals are not i.i.d. We conclude that StatNC's adaptivity as proposed in [1] is limited to processes with short-range dependencies as the backlog bound is significantly violated. Therefore, a dedicated integration of the fBm arrivals into the StatNC is an important step towards extending the application scope of the StatNC framework.

The previous discussions focused on a fixed choice of server and arrival rates, next we investigate the performance at different utilizations $\frac{\lambda}{C}$. Figure 3 shows the empirical backlog as well as the backlog bounds at different utilizations. Again, the dashed green lines represent the $95 \%$ quantile of the StatNC bounds. The simulated backlog at each utilization corresponds to the 0.998 -quantile of the empirical backlog distribution.

As observed in the previous experiments, the standard techniques of SNC do not deliver tight performance bounds for long-range dependent traffic classes. At the same time, since StatNC makes less assumptions on the traffic than the SNC, the additional inaccuracy of StatNC is clearly visible, yet arguably not contributing too much further conservatism.

## VI. Conclusion

In this paper we have shown how the StatNC framework can be extended to arrival processes following a fractional Brownian motion. This is an important step towards the applicability of StatNC as fBm arrivals allow to capture the frequent traffic characteristics of self-similarity and long-range dependence. Technically, the integration of fBm arrivals into the StatNC framework was achieved by using a confidence interval estimator for the Hurst parameter of fBm to derive a bounding statistic for the arrivals' moment-generating function. The estimated MGF bound was then used to derive a backlog bound in the StatNC framework. Numerical evaluations highlighted the use of this fBm estimator. Additionally,


Fig. 3. Empirical backlog bound together with SNC and StatNC bounds for quantile $=0.998,1-\alpha=0.999, \lambda=10^{-2}, H=0.7, \sigma=1$, sample paths of length $2^{16}$, and a time horizon of $t=200$ for 5000 iterations.
the calculated bounds exhibit a relatively moderate additional inaccuracy imposed by the StatNC.

## ACKNOWLEDGEMENTS

We gratefully acknowledge the help of Jürgen Franke who gave valuable input on estimating the Hurst parameter and its confidence interval.

## References

[1] M. A. Beck, S. A. Henningsen, S. B. Birnbach, and J. Schmitt, "Towards a statistical network calculus - dealing with uncertainty in arrivals," in Proc. IEEE International Conference on Computer Communications (INFOCOM'14), Toronto, Canada, 2014.
[2] C.-S. Chang, Performance guarantees in communication networks. London: Springer, 2000.
[3] A. Burchard, J. Liebeherr, and S. D. Patek, "A min-plus calculus for end-to-end statistical service guarantees," IEEE Transactions on Information Theory, vol. 52, no. 9, pp. 4105-4114, 2006.
[4] F. Ciucu, A. Burchard, and J. Liebeherr, "Scaling properties of statistical end-to-end bounds in the network calculus," IEEE/ACM Transactions on Networking (ToN), vol. 14, no. SI, pp. 2300-2312, 2006.
[5] M. Fidler, "An end-to-end probabilistic network calculus with moment generating functions," in Proc. IEEE IWQoS'06, Jun. 2006, pp. 261-270.
[6] Y. Jiang and Y. Liu, Stochastic network calculus. Springer, 2008, vol. 1.
[7] F. Ciucu and J. Schmitt, "Perspectives on network calculus - no free lunch, but still good value," in Proc. ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM'12), New York, NY, USA, Aug. 2012, pp. 311-322.
[8] R. L. Cruz, "A calculus for network delay, part I: Network elements in isolation," IEEE Transactions on Information Theory, vol. 37, no. 1, pp. 114-131, 1991.
[9] ——, "A calculus for network delay, part II: Network analysis," IEEE Transactions on Information Theory, vol. 37, no. 1, pp. 132-141, 1991.
[10] J. Kurose, "On computing per-session performance bounds in high-speed multi-hop computer networks," in Proc. ACM SIGMETRICS '92. ACM, 1992, pp. 128-139.
[11] O. Yaron and M. Sidi, "Performance and stability of communication networks via robust exponential bounds," IEEE/ACM Transactions on Networking (ToN), vol. 1, no. 3, pp. 372-385, 1993.
[12] C.-S. Chang, "Stability, queue length, and delay of deterministic and stochastic queueing networks," IEEE Transactions on Automatic Control, vol. 39, no. 5, pp. 913-931, 1994.
[13] C. Li, A. Burchard, and J. Liebeherr, "A network calculus with effective bandwidth," IEEE/ACM Transactions on Networking, vol. 15, no. 6, pp. 1442-1453, 2007.
[14] A. Rizk and M. Fidler, "Non-asymptotic end-to-end performance bounds for networks with long range dependent fbm cross traffic," Elsevier Computer Networks, vol. 56, no. 1, pp. 127-141, 2012.
[15] J. Liebeherr, A. Burchard, and F. Ciucu, "Delay bounds in communication networks with heavy-tailed and self-similar traffic," IEEE Transactions on Information Theory, vol. 58, no. 2, pp. 1010-1024, 2012.
[16] 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 (ToN), vol. 2, no. 1, pp. 1-15, 1994.
[17] V. Paxson and S. Floyd, "Wide area traffic: the failure of Poisson modeling," IEEE/ACM Transactions on Networking (ToN), vol. 3, no. 3, pp. 226-244, 1995.
[18] W. Willinger, M. S. Taqqu, R. Sherman, and D. V. Wilson, "Selfsimilarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level," IEEE/ACM Transactions on Networking (ToN), vol. 5, no. 1, pp. 71-86, 1997.
[19] M. E. Crovella and A. Bestavros, "Self-similarity in world wide web traffic: evidence and possible causes," IEEE/ACM Transactions on Networking (ToN), vol. 5, no. 6, pp. 835-846, 1997.
[20] A. Feldmann, A. C. Gilbert, P. Huang, and W. Willinger, "Dynamics of IP traffic: A study of the role of variability and the impact of control," in Proc. ACM SIGCOMM'99, vol. 29, no. 4, 1999, pp. 301-313.
[21] F. P. Kelly, "Notes on Effective Bandwidths," in Stochastic Networks: Theory and Applications, ser. Royal Statistical Society Lecture Notes Series, F. P. Kelly, S. Zachary, and I. Ziedins, Eds. Oxford University Press: Oxford, 1996, vol. 4, pp. 141-168.
[22] F. Dong, K. Wu, and V. Srinivasan, "Copula analysis for statistical network calculus," in Proc. IEEE International Conference on Computer Communications (INFOCOM'15), 2015, pp. 1535-1543.
[23] M. A. Beck, "Advances in theory and applicability of stochastic network calculus," Ph.D. dissertation, TU Kaiserslautern, 2016.
[24] I. Norros, "A storage model with self-similar input," Springer Queueing Systems, vol. 16, no. 3, pp. 387-396, 1994.
[25] - "On the use of fractional Brownian motion in the theory of connectionless networks," IEEE Journal on selected Areas in Communications, vol. 13, no. 6, pp. 953-962, 1995.
[26] M. S. Taqqu, V. Teverovsky, and W. Willinger, "Estimators for longrange dependence: an empirical study," Fractals, vol. 3, no. 04, pp. 785-798, 1995.
[27] O. Rose, "Estimation of the Hurst parameter of long-range dependent time series," University of Würzburg, Institute of Computer Science Research Report Series, Tech. Rep., 1996.
[28] J. Beran, Statistics for long-memory processes. New York: Chapman and Hall, 1994.
[29] _-, "Statistical methods for data with long-range dependence," Statistical Science, vol. 4, no. 7, pp. 404-427, 1992.
[30] R Core Team, R: A Language and Environment for Statistical Computing, R Foundation for Statistical Computing, Vienna, Austria, 2018. [Online]. Available: https://www.R-project.org/
[31] F. Ciucu, F. Poloczek, and J. Schmitt, "Sharp per-flow delay bounds for bursty arrivals: The case of FIFO, SP, and EDF scheduling," in Proc. IEEE International Conference on Computer Communications (INFOCOM'14), Toronto, Canada, 2014.
[32] P. Nikolaus and J. Schmitt, "On Per-Flow Delay Bounds in Tandem Queues under (In)Dependent Arrivals," in Proc. IFIP Networking 2017 Conference (NETWORKING'17), Stockholm, Sweden, 2017.

## Appendix A

## Proof of Theorem 7

Proof: We introduce the time scale parameter $\tau>0$ (compare to [4]) to discretize the event from (5) at the points

$$
\begin{equation*}
j:=\left\lfloor\frac{t-s}{\tau}\right\rfloor \tag{9}
\end{equation*}
$$

for all $0 \leq s \leq t$. Then we obtain for the arrivals $A(s, t) \leq$ $A(t-(j+1) \tau, t)$ and for the service $S(s, t) \geq S(t-j \tau, t)$. This gives us

$$
\begin{aligned}
& \mathrm{P}(q(t)>b) \\
& \stackrel{(2)}{\leq} \mathrm{P}\left(\sup _{0 \leq s \leq t}\{A(s, t)-S(s, t)\}>b\right)
\end{aligned}
$$

$$
\begin{align*}
& \stackrel{(9)}{\leq} \mathrm{P}\left(\max _{0 \leq j \leq\left\lfloor\frac{t}{\tau}\right\rfloor}\{A(t-(j+1) \tau, t)-S(t-j \tau, t)\}>b\right) \\
& \leq \sum_{j=0}^{\left\lfloor\frac{t}{\tau}\right\rfloor} \mathrm{P}(A(t-(j+1) \tau, t)-S(t-j \tau, t)>b)  \tag{10}\\
& \leq \sum_{j=0}^{\left\lfloor\frac{t}{\tau}\right\rfloor} \inf _{\theta_{j}>0}\left\{e^{-\theta_{j} b} \mathrm{E}\left[e^{\theta_{j}(A(t-(j+1) \tau, t)-S(t-j \tau, t))}\right]\right\} \\
& \stackrel{(4)}{=} \sum_{j=0}^{\left\lfloor\frac{t}{\tau}\right\rfloor} \inf _{\theta_{j}>0}\left\{e^{-\theta_{j} b} e^{-\theta_{j} C j \tau} e^{\lambda \theta_{j}(j+1) \tau+\frac{\theta_{j}^{2} \sigma^{2}}{2}((j+1) \tau)^{2 H}}\right\} \\
& =\sum_{k=1}^{\left\lfloor\frac{t}{\tau}\right\rfloor+1} \inf _{\theta_{k}>0}\left\{e^{\theta_{k}((\lambda-C) k \tau-(b-C \tau))} e^{\frac{\theta_{k}^{2} \sigma^{2}}{2}(k \tau)^{2 H}}\right\}  \tag{11}\\
& =\sum_{k=1}^{\left\lfloor\frac{t}{\tau}\right\rfloor+1}\left(e^{\frac{b-C \tau+(C-\lambda) k \tau}{\sigma^{2}(k \tau)^{2 H}}((\lambda-C) k \tau-(b-C \tau))}\right.  \tag{12}\\
& \left.\cdot e^{\left(\frac{b-C \tau+(C-\lambda) k \tau}{\sigma^{2}(k \tau)^{2 H}}\right)^{2} \frac{\sigma^{2}}{2}(k \tau)^{2 H}}\right) \\
& =\sum_{k=1}^{\left\lfloor\frac{t}{\tau}\right\rfloor+1} e^{-\frac{(b-C \tau+(C-\lambda) k \tau)^{2}}{\sigma^{2}(k \tau)^{2 H}}} e^{\frac{1}{2} \frac{(b-C \tau+(C-\lambda) k \tau)^{2}}{\sigma^{2}(k \tau)^{2 H}}} \\
& =\sum_{k=1}^{\left\lfloor\frac{t}{\tau}\right\rfloor+1} e^{-\frac{(b-C \tau+(C-\lambda) k \tau)^{2}}{2 \sigma^{2}(k \tau)^{2 H}}},
\end{align*}
$$

where we used the Union bound in (10) and in the subsequent line Chernoff bound to each summand. In (11), we shift the index, whereas in (12) we computed the minimum $\theta_{k}=\frac{b-C \tau+(C-\lambda) k \tau}{\sigma^{2}(k \tau)^{2 H}}$ (since $\left.b>\lambda \tau\right)$. Term manipulations finish the proof.

