

Towards a Statistical Network Calculus— Dealing with Uncertainty in Arrivals

Michael A. Beck, Sebastian A. Henningsen, Simon B. Birnbach, Jens B. Schmitt
Distributed Computer Systems (DISCO) Lab, University of Kaiserslautern, Germany

Abstract—The stochastic network calculus (SNC) has become an attractive methodology to derive probabilistic performance bounds. So far the SNC is based on (tacitly assumed) exact probabilistic assumptions about the arrival processes. Yet, in practice, these are only true approximately—at best. In many situations it is hard, if possible at all, to make such assumptions a priori. A more practical approach would be to base the SNC operations on measurements of the arrival processes (preferably even on-line). In this paper, we develop this idea and incorporate measurements into the framework of SNC taking the further uncertainty resulting from estimation errors into account. This is a crucial step towards a statistical network calculus (StatNC) eventually lending itself to a self-modelling operation of networks with a minimum of a priori assumptions. In numerical experiments, we are able to substantiate the novel opportunities by StatNC.

I. INTRODUCTION

A. Motivation

Over the last two decades the stochastic network calculus (SNC) has evolved as a valuable methodology to compute probabilistic performance bounds [10]. It has found numerous and diverse usage in important network design and control problems: smart grid control [26], delay control in cognitive radio networks [15], and as foundation for bandwidth estimation on Internet end-to-end paths [22], to name a few recent examples.

SNC originated from its deterministic counterpart as conceived by Cruz [12], [13] to provide stochastically relaxed performance bounds, mainly in order to capture the statistical multiplexing gain as is characteristic for packet-switched networks. Some of the earliest work on SNC can be traced back to [27], [7], [20]. In particular Chang’s sigma-rho calculus based on moment-generating functions (MGF) received much attraction in the field and was refined in [14] to match with the latest advances in the min-plus algebraic formulation of network calculus (alternative SNC formulations can be found in [14], [6], [19], see [10] for some perspectives about these). The core modelling abstractions of SNC are arrival envelopes and service curves. Arrival envelopes provide probabilistic bounds on how much traffic arrives within a time interval of a given length; service curves essentially do the same for the amount of work done by a system serving those arrivals.

One of the strengths of SNC is its versatility with respect to traffic models that can be treated, ranging from short-range dependent traffic with exponentially bounded burstiness (see e.g. [9]) to long-range dependent traffic such as fractional Brownian motion [25], or even heavy-tailed self-similar traffic [21].

Yet, all of these works start from “clean”, a priori and exact probabilistic assumptions. In practice, however, the question arises: where do these assumptions come from? In most cases the answer must be: observation of the past traffic behaviour, in the form of measurements and subsequent statistical inference. However, statistical inference involves errors and, thus, another source of uncertainty besides traffic variations themselves. To the best of our knowledge, none of the existing work on SNC has taken this uncertainty into account and integrated it into the SNC operations. We take this missing first step, i.e., measuring the arrivals and making statistical inferences, and integrate it into the SNC, thereby moving towards a statistical network calculus (StatNC)¹. Moving from SNC to StatNC can be viewed as going from stochastic processes to time series. Of course, we still have to make assumptions for the time series corresponding to past traffic arrivals with respect to the underlying stochastic process, but we can adapt them dynamically (possibly on-line) and some deviations from the assumptions may be tolerable (depending on the robustness of our statistical estimators). Clearly, the goal of our StatNC framework is to cope with as few assumptions as possible while still providing accurate performance bounds.

To illustrate where a statistical network calculus can be very beneficial, let us briefly sketch two application scenarios:

- 1) **Traffic engineering in an MPLS domain** [2]: traffic is measured at ingress nodes to an MPLS domain and label-switched paths are dynamically dimensioned according to service level agreements based on StatNC; an immediate benefit is that time-of-day effects or any other seasonal effects are automatically taken into account.
- 2) **Self-modelling in wireless sensor networks**: traffic is measured at sensor nodes and the resulting estimates are delivered towards a sink (in the simplest case) which can then base decisions such as, e.g., topology control on the respective StatNC models; an immediate benefit is that no a priori traffic description is necessary any more, which is very helpful in many WSN applications as the behaviour of the physical phenomena to be observed is often not well-understood before deployment and thus the traffic induced by them is hard to predict.

¹The term statistical network calculus has been used before to indicate that the SNC takes into account statistical multiplexing gains [6], whereas here we use it to indicate the usage of statistical methods instead of purely probabilistic reasoning.

Overall, we make the following contributions

- development of a uniform framework for a statistical network calculus which allows to plug in a large class of traffic estimators (\rightarrow Section III);
- design of several traffic estimators with differing amount of presumed knowledge and probabilistic assumptions (\rightarrow Section IV);
- in numerical examples, the practicality, precision, and robustness of the StatNC is investigated and contrasted against the performance of SNC alongside with simulative results (\rightarrow Section V).

B. Related Work

In the SNC literature, there are only few papers that discuss the fact that arrival envelopes could be derived from measurements: for example, [10] provides a brief sketch how a measured packet trace could be fitted to a weighted hyperexponential traffic envelope, while [21] even does it for a heavy-tailed self-similar traffic envelope. Yet, none of these integrates the measurements with the SNC operations such that the uncertainty resulting from estimation errors is factored into the stochastic bounds. In the paper at hand, we perform this integration in a rigorous and uniform manner (see Theorem 5 in Section III). Furthermore, to the best of our knowledge there is no previous work in the SNC literature about an on-line estimation of the arrival envelope as it is enabled by our StatNC framework.

In a larger context, somewhat related work can be found in the domain of measurement-based admission control (see e.g. [18], [17], [24], [16]). However, because at that time the SNC was not yet fully developed, these works are restricted to admission control rather than basing on a general performance evaluation framework like SNC. Furthermore, they typically assume a known (deterministic) traffic envelope and then measure to what extent this envelope is used and how statistical multiplexing helps to reduce resource demands, whereas in our work we basically start one step earlier by estimating the probabilistic arrival envelopes themselves.

Also slightly related is the work by Lübben et al. [22] on the identification of stochastic service curves to represent Internet end-to-end paths. Clearly, measurements (though active ones) play a central role here as well, yet the target is different in our case as we deal with the uncertainty about arrival rather than service processes.

On a very high level, the vision of autonomic networking (see e.g. [5] for a prominent large-scale project in that domain) could be related especially to the self-modelling aspect of the StatNC when used in an on-line fashion, yet no use of SNC within this domain is known to us, although it appears to be a very promising idea.

II. PRELIMINARIES ON STOCHASTIC NETWORK CALCULUS

In this paper, we focus on the SNC formulation as originally presented in [8] and later on generalized by [14], which is also known as $(\sigma(\theta), \rho(\theta))$ -calculus. In this setup, time is discrete while data is allowed to be continuous (i.e., we deal with

infinitesimally small data units). For convenience, we make a few small modifications to definitions and notations from [8], and therefore repeat the most important of them together with the main results needed in this paper. For brevity, we focus on the backlog as performance measure in this paper and only present the corresponding results. Results concerning other performance measures (i.e., virtual delay and output bounds) or reducing the complexity of networks with multiple flows and service elements, can be derived in a similar fashion.

In SNC, data flows arrive at service elements and after being processed leave them again. We represent such a flow, by a real non-negative stochastic process $(a_k)_{k \in \mathbb{Z}}$ and the bivariate cumulatives

$$A(m, n) := \sum_{k=m+1}^n a_k.$$

We henceforth call the random variables a_k *increments* of the flow A . The basic idea of StatNC is to apply statistical methods on past observations and hence we think of increments with time index $k < 0$ as lying in the past (the so-far observed time series of arrivals). The increments with indices $k \geq 0$ are upcoming arrivals. Further, we assume a value $n_0 \leq 0$ such that $a_k = 0$ for all $k < n_0$, this is the time when we started our observations.

The service element is also abstracted by a doubly indexed stochastic process S with the properties:

$$\begin{aligned} 0 &\leq S(m, n) && \forall m, n \in \mathbb{N}_0 \\ S(m, n) &\leq S(m, n') && \forall m, n, n' \in \mathbb{N}_0 \text{ and } n \leq n' \end{aligned}$$

Note that we define S only on $\mathbb{N}_0 \times \mathbb{N}_0$, which is—as we will see—sufficient. The service process S , arrival flow A and the departure flow D of a service element are linked with each other in the following way:

Definition 1. If for all $n \in \mathbb{N}_0$ it holds that

$$D(0, n) \geq \min_{0 \leq k \leq n} \{A(0, k) + S(k, n)\},$$

we call the service element a dynamic S -server. Here D is defined as a flow with $n_0 = 0$.

Before we can provide stochastic bounds on the backlog of a system, we need some bounds on the arrivals and the dynamic S -server. More precisely, we need bounds on the moment generating functions (MGF) of the corresponding stochastic processes.

Definition 2. Let $\theta > 0$. An arrival is $(\sigma_A(\theta), \rho_A(\theta))$ -bounded iff

$$\sup_{m \in \mathbb{Z}} \{\mathbb{E}(e^{\theta A(m, m+k)})\} \leq e^{k\theta\rho_A(\theta) + \theta\sigma_A(\theta)} \quad \forall k \in \mathbb{N}$$

A dynamic S -server is $(\sigma_S(\theta), \rho_S(\theta))$ -bounded iff

$$\sup_{m \geq 0} \{\mathbb{E}(e^{-\theta S(m, m+k)})\} \leq e^{k\theta\rho_S(\theta) + \theta\sigma_S(\theta)} \quad \forall k \in \mathbb{N}$$

We are now able to provide stochastic bounds on a service element's *backlog* process defined by $q(n) := A(0, n) - D(0, n)$.

Theorem 3. Let A be an arrival flow served by a dynamic S -server and $\theta > 0$. Assume A is $(\sigma_A(\theta), \rho_A(\theta))$ -bounded and S

is $(\sigma_S(\theta), \rho_S(\theta))$ -bounded. If A is stochastically independent of S , the following probabilistic bound holds²:

$$\mathbb{P}(q(n) > x) \leq e^{-\theta x} e^{\theta(\sigma_A(\theta) + \sigma_S(\theta))} \sum_{k=0}^n e^{k\theta(\rho_A(\theta) + \rho_S(\theta))}.$$

Proof: By definition of the dynamic S -server we have:

$$\begin{aligned} q(n) &\leq A(n) - \min_{0 \leq k \leq n} \{A(0, k) + S(k, n)\} \\ &= \max_{0 \leq k \leq n} \{A(k, n) - S(k, n)\} \end{aligned}$$

from which we can derive, using Chernoff's inequality³:

$$\begin{aligned} \mathbb{P}(q(n) > x) &\leq e^{-\theta x} \mathbb{E}(e^{\theta \max_{0 \leq k \leq n} \{A(k, n) - S(k, n)\}}) \\ &\leq e^{-\theta x} \sum_{k=0}^n \mathbb{E}(e^{\theta A(k, n)}) \mathbb{E}(e^{-\theta S(k, n)}) \\ &\leq e^{-\theta x} e^{\theta(\sigma_A(\theta) + \sigma_S(\theta))} \sum_{k=0}^n e^{k\theta(\rho_A(\theta) + \rho_S(\theta))} \end{aligned}$$

Where the independence of A and S has been used in the second line. ■

Here, problems arise when we face uncertainty in the description of the arrival flow A . If the exact distribution of the increments is unknown, we cannot calculate the expression $\mathbb{E}(e^{\theta A(k, n)})$, which in turn prohibits calculation of the backlog bound. Hence, we use statistics to deal with $\mathbb{E}(e^{\theta A(k, n)})$, effectively replacing the $(\sigma_A(\theta), \rho_A(\theta))$ -bound in the above proof.

III. A FRAMEWORK FOR A STATISTICAL NETWORK CALCULUS

In this section, we present the framework of StatNC (\rightarrow Theorem 5). Technically, it can be seen as a sufficient condition on the employed statistics, allowing calculations of performance bounds. For brevity, denote by $\phi_{m, n}(\theta) := \mathbb{E}(e^{\theta A(m, n)})$ the MGF of $A(m, n)$ at point θ . First, we need a small lemma proving the monotonic behaviour of the backlog bound in the MGF of A .

Lemma 4. *Let $\hat{\phi}_{m, n}(\theta) \geq \phi_{m, n}(\theta)$ for some $\theta > 0$ and all $m, n \in \mathbb{N}_0$ with $m \leq n$. Assume A being stochastically independent⁴ from S . Then*

$$\mathbb{P}(q(n) > x) \leq e^{-\theta x} \sum_{k=0}^n \hat{\phi}_{k, n}(\theta) \mathbb{E}(e^{-\theta S(k, n)})$$

for all $n \in \mathbb{N}_0$.

²For the case of A not being stochastically independent of S bounds also exist, yet, for the sake of clarity and brevity, we leave these out here. See [8], [3] for more information.

³Chernoff's inequality states that for some real random variable X and every $\theta > 0$: $\mathbb{P}(X > x) \leq e^{-\theta x} \mathbb{E}(e^{\theta X})$.

⁴A corresponding result for the dependent case can be found in [4].

Proof: From the proof of Theorem 3 we know

$$\begin{aligned} \mathbb{P}(q(n) > x) &\leq e^{-\theta x} \sum_{k=0}^n \phi_{k, n}(\theta) \mathbb{E}(e^{-\theta S(k, n)}) \\ &\leq e^{-\theta x} \sum_{k=0}^n \hat{\phi}_{k, n}(\theta) \mathbb{E}(e^{-\theta S(k, n)}). \end{aligned}$$

Define \mathcal{F} to be the space of functions mapping from $\mathbb{N}_0 \times \mathbb{N}_0 \times \mathbb{R}^+$ to \mathbb{R}_0^+ . In expression, if $f \in \mathcal{F}$, then:

$$f : \mathbb{N}_0 \times \mathbb{N}_0 \times \mathbb{R}^+ \rightarrow \mathbb{R}_0^+$$

An already familiar example for a member of \mathcal{F} is the MGF $\phi_{m, n}(\theta)$ of the arrival flow A .

We now provide a theorem on how the uncertainties of using statistics can be combined with the probabilistic bounds derived from SNC.

Theorem 5. *Let $\theta^* = \sup\{\theta : \phi_{m, n}(\theta) < \infty\}$ and $\Phi : \mathbb{R}^{|n_0|} \rightarrow \mathcal{F}$ be a statistic on $a = (a_{n_0}, \dots, a_{-1})$ such that*

$$\sup_{\theta \in (0, \theta^*)} \mathbb{P}\left(\bigcup_{m \leq n} \Phi(a)(m, n, \theta) < \phi_{m, n}(\theta)\right) \leq \alpha.$$

Then for all $n \in \mathbb{N}_0$, $\theta < \theta^*$

$$\mathbb{P}(q(n) > x) \leq \alpha + e^{-\theta x} \sum_{k=0}^n \Phi(a)(k, n, \theta) \mathbb{E}(e^{-\theta S(k, n)}).$$

Proof: Fix some $\theta > 0$:

$$\begin{aligned} &\mathbb{P}(q(n) > x) \\ &= \mathbb{P}\left(q(n) > x \cap \bigcup_{m \leq n} \Phi(a)(m, n, \theta) < \phi_{m, n}(\theta)\right) \\ &\quad + \mathbb{P}\left(q(n) > x \cap \bigcap_{m \leq n} \Phi(a)(m, n, \theta) \geq \phi_{m, n}(\theta)\right) \\ &\leq \alpha + \mathbb{P}\left(q(n) > x \cap \bigcap_{m \leq n} \Phi(a)(m, n, \theta) \geq \phi_{m, n}(\theta)\right) \\ &\leq \alpha + \mathbb{P}\left(q(n) > x \mid \bigcap_{m \leq n} \Phi(a)(m, n, \theta) \geq \phi_{m, n}(\theta)\right) \\ &\leq \alpha + e^{-\theta x} \sum_{k=0}^n \Phi(a)(m, n, \theta) \mathbb{E}(e^{-\theta S(k, n)}) \end{aligned}$$

From the proof, the nature of the condition in the theorem becomes clearer. We need the intersection of the events $\Phi(a)(m, n, \theta) \geq \phi_{m, n}(\theta)$ to leverage from the monotonic behaviour of the backlog bound (\rightarrow Lemma 4). We achieve this intersection by partitioning the event $q(n) > x$ and hence have to deal with the corresponding complement, which is the union appearing in the second line of the proof. This union describes the event, that our statistic delivers a value lying below the real MGF of A at least once. We bound this kind of (estimation) error by a confidence level of α . The confidence level α can be seen as a parameter of optimization. ■

Please note, that, for our theorem, we defined the statistic Φ in the most general way, i.e., being a function on the complete history $a = (a_{n_0}, \dots, a_{-1})$. This does not mean, however, that one has to use all this information to construct a Φ satisfying the above condition. Assume a subsample $a' = (a'_{n'_0}, \dots, a'_{-1})$ of a , such that: $a'_i = a_{j_i}$ for some index $j_i \in \{n_0, \dots, -1\}$ and a statistic $\Phi' : \mathbb{R}^{|n'_0|} \rightarrow \mathcal{F}$ on a' . If Φ' meets the assumption of Theorem 5 for a' , we can extend it canonically to a statistic Φ by setting: $\Phi(a) = \Phi'(a')$ for all a , such that $a_{j_i} = a'_i$ holds. This shows that using only a part of the history a is just a special case of the above theorem.

Often one may want to give more recent observations a larger impact on the sample and diminish the influence of observations as they get older, e.g., $w_i(a_i) = \beta^{|i+1|} a_i$, $0 < \beta < 1$ (\rightarrow exponential smoothing). Such transformations of the sample are also covered by the above theorem: for some weighting function $w : \mathbb{R}^{|n_0|} \rightarrow \mathbb{R}^{|n_0|}$ on the sample a , the concatenation $\Phi \circ w(a)$ needs to meet the above assumption in order to apply the statistic Φ on the weighted sample.

Another typical and very practical way of subsampling would be to use a sliding window of length l on the observations, i.e., $a'_i = a_i$, with $i \in \{-l, \dots, -1\}$. The sliding window is particularly interesting for on-line estimation, since it allows the statistical network calculus to dynamically adapt to changes in the arrivals' characteristics. In Section V, we investigate how the versatility of this dynamic view can be leveraged to achieve better bounds than the static ones of SNC.

For the rest of the paper, we adhere to the general notation as in Theorem 5, unless otherwise mentioned.

IV. EXAMPLES OF STATISTICAL ESTIMATORS

As stated above, estimating the quantity $\phi_{m,n}(\theta)$ for an arbitrary $\theta \in (0, \theta^*)$ and $m \leq n \in \mathbb{N}$ is key for StatNC. In the following subsections, we present different scenarios and their corresponding $\phi_{m,n}(\theta)$ -estimators. The crucial point is to meet the condition from Theorem 5. Starting with a fairly simple, but illustrative example (exponential i.i.d. increments), we move on to more complex scenarios, involving non-i.i.d. behaviour of the increments a_k .

A. Exponential Traffic

Assume the $(a_k)_{k \geq n_0}$ to be i.i.d. exponentially distributed with some unknown parameter λ . The idea to construct Φ in this scenario, is to estimate λ first. For this, note that a lower bound on the real distribution parameter λ with confidence level α can be computed by

$$\bar{\lambda} := \frac{\chi_\alpha^2(2|n_0|)}{2 \cdot A(n_0 - 1, -1)}.$$

Here $\chi_\alpha^2(2|n_0|)$ is the one-sided α -quantile of a Chi-Squared distribution with $2|n_0|$ degrees of freedom (scaling $A(n_0 - 1, -1)$ by 2λ results in a random variable, which is Chi-squared-distributed with $2|n_0|$ degrees of freedom). Moreover, the MGF of the exponential distribution is given by $\phi(\lambda, \theta) = \left(\frac{\lambda}{\lambda - \theta}\right) = 1 + \left(\frac{\theta}{\lambda - \theta}\right)$, for all $\theta < \lambda$. With the simple implication

$$\bar{\lambda} \leq \lambda \Rightarrow \frac{\bar{\lambda}}{\bar{\lambda} - \theta} \geq \frac{\lambda}{\lambda - \theta},$$

we obtain for $\bar{\lambda} \leq \lambda$ that

$$\Phi(a)(m, n, \theta) := \left(\frac{\bar{\lambda}}{\bar{\lambda} - \theta}\right)^{n-m} \geq \left(\frac{\lambda}{\lambda - \theta}\right)^{n-m} = \phi_{m,n}(\theta)$$

applies for all $\theta < \lambda$ and $m \leq n \in \mathbb{N}_0$.

Hence we obtain

$$\begin{aligned} 1 - \alpha &= \mathbb{P}(\bar{\lambda} \leq \lambda) \\ &\leq \inf_{\theta \in (0, \lambda)} \mathbb{P}\left(\bigcap_{m \leq n} \Phi(a)(m, n, \theta) \geq \phi_{m,n}(\theta)\right). \end{aligned}$$

Or, equivalently

$$\sup_{\theta \in (0, \lambda)} \mathbb{P}\left(\bigcup_{m \leq n} \Phi(a)(m, n, \theta) < \phi_{m,n}(\theta)\right) \leq \alpha,$$

which is exactly the condition of Theorem 5.

B. Bandwidth-Limited i.i.d. Traffic

Assume again the increments of A to be i.i.d., now also adhering to a bandwidth limitation M , i.e., no more than M data units per time slot can arrive. The latter is a valid assumption as any real access link has such a restriction. In contrast to the previous subsection, we now lack any further knowledge about the distribution of the increments a_k . Yet, it turns out that we can still construct a statistic Φ fitting Theorem 5. The Dvoretzky-Kiefer-Wolfowitz inequality [23] is very useful in this setup:

Lemma 6. *Let $F_{n_0}(x) := 1/|n_0| \sum_{k=n_0}^{-1} \mathbf{1}_{\{a_k \leq x\}}$ be the empirical distribution function of the sample $a = (a_{n_0}, \dots, a_{-1})$ and F be the distribution of one increment of A . Then we have for all $\varepsilon > 0$ the following:*

$$\mathbb{P}\left(\sup_{x \in [0, M]} |F_{n_0}(x) - F(x)| \leq \varepsilon\right) \geq 1 - 2e^{-2|n_0|\varepsilon^2}.$$

Since the arrivals are bounded we can fix some arbitrary $\theta > 0$ for the rest of this subsection. The next theorem constructs Φ :

Theorem 7. *Let $\alpha \in (0, 1)$ be given. The statistic Φ defined by*

$$\Phi(a)(m, n, \theta) := (\bar{A} + \varepsilon(e^{\theta M} - 1))^{n-m}$$

satisfies the condition in Theorem 5. Here

$$\bar{A} := \frac{1}{|n_0|} \sum_{k=n_0}^{-1} e^{\theta a_k}.$$

and $\varepsilon := \sqrt{\frac{\log(\alpha/2)}{2|n_0|}}$.

Proof: Let a_k be an arbitrary increment of A . Assuming the event in the left-hand side of the Dvoretzky-Kiefer-Wolfowitz inequality (\rightarrow Lemma 6), we can derive successively:

$$\begin{aligned} F(x) &\geq F_{n_0}(x) - \varepsilon \quad \forall x \in [0, M], \\ 1 - F(x) &\leq 1 - F_{n_0}(x) + \varepsilon \quad \forall x \in [0, M], \end{aligned}$$

$$\mathbb{P}(e^{\theta a_k} > x) \leq 1 - F_{n_0}(1/\theta \log(x)) + \varepsilon \quad \forall x \in [1, e^{\theta M}],$$

$$\begin{aligned} \mathbb{E}(e^{\theta a_k}) &= 1 + \int_1^{e^{\theta M}} \mathbb{P}(e^{\theta a_k} > x) dx \\ &\leq 1 + \int_1^{e^{\theta M}} 1 - F_{n_0}(1/\theta \log(x)) + \varepsilon dx. \end{aligned}$$

Hence, we have for all $\theta > 0$:

$$\begin{aligned} \mathbb{P}\left(\mathbb{E}(e^{\theta a_k}) \leq 1 + \int_1^{e^{\theta M}} 1 - F_{n_0}\left(\frac{1}{n} \log(x)\right) + \varepsilon dx\right) \\ \geq \mathbb{P}\left(\sup_{x \in [0, M]} |F_{n_0}(x) - F(x)| \leq \varepsilon\right) \\ \geq 1 - 2e^{-2|n_0|\varepsilon^2}. \end{aligned}$$

This means we have constructed a one-sided confidence interval for $\mathbb{E}(e^{\theta a_k})$ with a significance level smaller than $\alpha = 2e^{-2|n_0|\varepsilon^2}$, which is given by:

$$\left[0, 1 + \int_1^{e^{\theta M}} 1 - F_{n_0}(1/\theta \log(x)) + \varepsilon dx\right].$$

Simplifying the integral and inserting the corresponding ε for a significance level smaller than α (for details see [4]), the confidence interval becomes:

$$\left[0, \bar{A} + \sqrt{\frac{-\log(\frac{\alpha}{2})}{2|n_0|}}(e^{\theta M} - 1)\right].$$

For the statistic Φ we indeed have:

$$\begin{aligned} &\inf_{\theta} \mathbb{P}\left(\bigcap_{m,n} \Phi(a)(m, n, \theta) \geq \phi_{m,n}(\theta)\right) \\ &= \inf_{\theta} \mathbb{P}\left(\bigcap_{m,n} \prod_{k=m+1}^n \bar{A} + \varepsilon(e^{\theta M} - 1) \geq \prod_{k=m+1}^n \mathbb{E}(e^{\theta a_k})\right) \\ &\geq \inf_{\theta} \mathbb{P}(\bar{A} + \varepsilon(e^{\theta M} - 1) \geq \mathbb{E}(e^{\theta k})) \\ &\geq 1 - \alpha. \end{aligned}$$

Or, equivalently:

$$\sup_{\theta} \mathbb{P}\left(\bigcup_{m,n} \Phi(a)(m, n, \theta) < \phi_{m,n}(\theta)\right) \leq \alpha.$$

Again, the statistic $\Phi(a)(m, n, \theta)$ satisfies the condition of Theorem 5 and thus can be used to calculate the desired performance bounds. \blacksquare

C. Markov-Modulated Arrivals

Next, we discuss a traffic class, in which the i.i.d. assumption is dropped and multiple statistics are combined to construct Φ . For this consider a Markov-modulated arrival, with a Markov chain $(Y_k)_{k \in \{n_0, \dots\}}$ corresponding to the transition matrix

$$T = \begin{pmatrix} \mu & 1 - \mu \\ 1 - \nu & \nu \end{pmatrix}$$

and denote the first (second) state as *On*-state (*Off*-state). The increments of the arrival A are now defined by:

$$a_k = \begin{cases} 0 & \text{if } Y_k \text{ is in } \textit{Off}\text{-state} \\ x_k & \text{if } Y_k \text{ is in } \textit{On}\text{-state} \end{cases}$$

where $(x_k)_{k \in \{n_0, \dots\}}$ is a sequence of i.i.d. random variables bounded by the bandwidth limitation M and $x_k \neq 0$ almost surely. Note that in general the increments a_k are neither identically distributed nor stochastically independent. This model generalizes the well-known and popular Markov-modulated On-Off traffic model [1], with the distinction, that in the *On*-state the arrivals are defined by a random process, instead of a constant rate.

Before we can construct Φ , we need a lemma, showing the monotonicity of $\phi_{m,n}(\theta)$ with respect to the parameters μ and ν . For this define $\theta^* = \sup\{\theta : \mathbb{E}(e^{\theta x_k}) < \infty\}$.

Lemma 8. *For the above model, all $m \leq n \in \mathbb{N}$ and all $\theta \in (0, \theta^*)$ it holds that*

$$\mu \leq \tilde{\mu} \quad \Rightarrow \quad \mathbb{E}(e^{\theta A_{\mu, \nu}(m, n)}) \leq \mathbb{E}(e^{\theta A_{\tilde{\mu}, \nu}(m, n)}),$$

and

$$\nu \leq \tilde{\nu} \quad \Rightarrow \quad \mathbb{E}(e^{\theta A_{\mu, \nu}(m, n)}) \leq \mathbb{E}(e^{\theta A_{\mu, \tilde{\nu}}(m, n)}).$$

Proof: Although this statement is very intuitive, a rigorous proof is surprisingly involved. Please see [4] for details. \blacksquare

Without knowing the transition probabilities μ and ν , nor the distribution of the (x_i) , we can still construct a Φ satisfying Theorem 5. In order to do so, denote for an arrival sample $a = (a_{n_0}, \dots, a_{-1})$ the observed number of transitions from *Off*-state to *Off*-state by $X_{0,0}$, from *Off*-state to *On*-state by $X_{0,1}$, from *On*-state to *Off*-state by $X_{1,0}$ and from *On*-state to *On*-state by $X_{1,1}$. Further, denote the observed number of visits in the *On*-state in a time interval $(m, n]$ by $O(m, n)$ and the number of visits in the *Off*-state by $P(m, n)$. Further, define $O := O(n_0, -1)$ and $P := P(n_0, -1)$.

Theorem 9. *Choose some confidence level $\alpha = \alpha_\mu + \alpha_\nu + \alpha_d$ and consider some sample $a = (a_{n_0}, \dots, a_{-1})$ with $O \neq 0$. Define the statistics*

$$\begin{aligned} \mu_l &:= \beta^{-1}(\alpha_\mu; X_{0,0}, P - X_{0,0} + 1) \\ \nu_u &:= \beta^{-1}(1 - \alpha_\nu; X_{1,1} + 1, O - X_{1,1}) \end{aligned}$$

where β^{-1} is the inverse of the beta-distribution. Further define the transition matrix

$$T^* = \begin{pmatrix} \mu_l & 1 - \mu_l \\ 1 - \nu_u & \nu_u \end{pmatrix}.$$

Define

$$A^* = \bar{A} + \left(\frac{-\log(\frac{\alpha_d/2}{2|O|})}{2|O|}\right)^{1/2} (e^{\theta M} - 1).$$

Then the statistic $\Phi : \mathbb{R}^{|n_0|} \rightarrow \mathcal{F}$ defined by:

$$\Phi(a)(m, n, \theta) := A^* \frac{\bar{x}_{On} \vee \bar{x}_{Off}}{\bar{x}_{On} \wedge \bar{x}_{Off}} \rho(\bar{E}T^*)^{n-m-1}$$

satisfies the condition of the statistical framework (\rightarrow Theorem 5). Here $\bar{A} := \frac{1}{O} \sum_{i: Y_i = On} e^{\theta a_i}$, \bar{E} , \bar{x}_{On} and \bar{x}_{Off} are from the $\sigma(\theta)$, $\rho(\theta)$ -bound for Markov-modulated arrivals (see [4]) with fixed arrivals in the *On*-states equal to $1/\theta \log(A^*)$.

Proof: First we show that $[\mu_l, 1]$ and $[0, \nu_u]$ are confidence intervals for μ and ν and the confidence levels α_μ and α_ν ,

respectively. Basically, these are the well-known Clopper-Pearson intervals [11], which are constructed as follows: Interpret $X_{0,0}$ as the number of successes in a $Bin(\mu, P)$ -distributed random variable. It is known that for some $X \sim Bin(p, n)$ and β being the beta-distribution it holds that:

$$\mathbb{P}(X \leq k) = \sum_{i=0}^k \binom{n}{i} p^i (1-p)^{n-i} = \beta(1-p; n-k, k+1).$$

From that we can continue with:

$$\beta(1-p, n-k, k+1) = 1 - \beta(p; k+1, n-k).$$

We now ask for the smallest μ' such that a random variable $X \sim Bin(\mu', P)$ meets:

$$\mathbb{P}_{\mu'}(X \geq X_{0,0}) \geq \alpha_\mu,$$

or, equivalently:

$$\mathbb{P}_{\mu'}(X \leq X_{0,0} - 1) \leq 1 - \alpha_\mu.$$

Hence μ' must fulfill:

$$1 - \beta(\mu'; X_{0,0}, P - X_{0,0} + 1) = \mathbb{P}_{\mu'}(X \leq X_{0,0} - 1) = 1 - \alpha_\mu,$$

which is solved by $\mu_l = \beta^{-1}(\alpha_\mu; X_{0,0}, P - X_{0,0} + 1)$. Using a simple coupling argument and the definition of μ_l one obtains the implication:

$$\mathbb{P}(\mu_l > \mu) \leq \mathbb{P}(X \geq X_{0,0}) = \alpha_\mu.$$

Very similarly, we obtain:

$$\mathbb{P}(\nu_u < \nu) \leq \mathbb{P}(X \leq X_{1,1}) = \alpha_\nu,$$

where X is now a $Bin(\nu_u, O)$ -distributed variable.

Now fix some arbitrary $\theta \in (0, \theta^*)$ and assume for the moment

$$\begin{aligned} \mu_l &\leq \mu, \\ \nu_u &\geq \nu, \\ A^* &\geq \mathbb{E}(e^{\theta a_{O_n}}). \end{aligned}$$

Then, for all $m, n \in \mathbb{N}$ it would hold that:

$$\begin{aligned} \mathbb{E}(e^{\theta A(m,n)}) &\leq \mathbb{E}(e^{\theta A_{\mu_l, \nu}(m,n)}) \\ &\leq \mathbb{E}(e^{\theta A_{\mu_l, \nu_u}(m,n)}) \leq \Phi(a)(m, n, \theta), \end{aligned}$$

where we still have to show the last inequality. Putting the proof of the last inequality on hold, we henceforth have for all $\theta \in (0, \theta^*)$:

$$\begin{aligned} \alpha &= \alpha_\mu + \alpha_\nu + \alpha_d \geq \mathbb{P}(\mu_l > \mu \cup \nu_u < \nu \cup A^* < \mathbb{E}(e^{\theta a_{O_n}})) \\ &\geq \mathbb{P}\left(\bigcup_{m \leq n} \mathbb{E}(e^{\theta A(m,n)}) > \Phi(a)(m, n, \theta)\right) \end{aligned}$$

which is what we wanted to show.

To show the missing inequality define a new Markov-modulated arrival with a constant rate $a_{O_n}^* = 1/\theta \log(A^*)$ and T^* as transition matrix. Then, we then have $\mathbb{E}(e^{\theta a_{O_n}^*}) = A^*$

and the $\sigma(\theta), \rho(\theta)$ -bound is given by $\Phi(a)(m, n, \theta)$. Further we have:

$$\begin{aligned} \mathbb{E}(e^{\theta A_{\mu_l, \nu_u}(m,n)}) &= \sum_{k=0}^{n-m} \mathbb{P}(O_{\mu_l, \nu_u}(m, n) = k) \mathbb{E}(e^{\theta a_{O_n}})^k \\ &\leq \sum_{k=0}^{n-m} \mathbb{P}(O_{\mu_l, \nu_u}(m, n) = k) \mathbb{E}(e^{\theta a_{O_n}^*})^k \\ &= \Phi(a)(m, n, \theta). \end{aligned}$$

This completes the proof. \blacksquare

D. Summary

We have provided three examples with differing degrees of assumed knowledge and complexity, when constructing the statistic Φ . From the formulation of the framework, it is clear that the technically hard part in applying the StatNC lies in constructing such estimators. Taking care of other, potentially more complex arrival processes is thus a question of finding the corresponding Φ , i.e., meeting the condition of the framework theorem. Admittedly, this can be hard in some circumstances and, for example, we leave the construction of Φ for long-range dependent traffic types for future work.

V. NUMERICAL EVALUATION – STATNC AT WORK

In this section, we compare the statistical network calculus with its stochastic counterpart. To that end, we investigate how high the costs of involving statistics are (in terms of looser bounds). Furthermore, we study special properties of StatNC which the SNC lacks; these are its dynamic view on the measurements, as well as its robustness against false assumptions.

Scenario 1: The Price of StatNC

In the first scenario, we study if the additional uncertainties resulting from the statistical part of the performance bounds are acceptable. In expression, we calculate the smallest b such that we can still guarantee

$$\mathbb{P}(q(n) > b) \leq \varepsilon,$$

with our methods of StatNC (or with the methods of standard SNC). For a perfect bound, we expect that for a large number of simulations N to see roughly $N \cdot \varepsilon$ of the simulations producing a backlog larger than b at time n . Hence, we simulate the backlog process for time n in N repetitions and compare the empirical distribution of the observed backlogs with the b obtained from the StatNC/SNC formula above. The bounds are the better, the closer they lie to the $(1-\varepsilon)$ -quantile of the empirical backlog distribution.

To that end, we simulate a Markov-modulated arrival process, as described in Section IV-C, with x_i being exponentially distributed, but capped by a bandwidth limitation M . The parameter λ of the exponential distribution is chosen to be 0.2, while the bandwidth limitation is set to $M = 20$ (which means a hypothetical access link is maximally utilized at 25%). The transition probabilities of the Markov chain are set to $\mu = 0.9$ and $\nu = 0.9$. We use a constant rate server with rate $c = 5$,

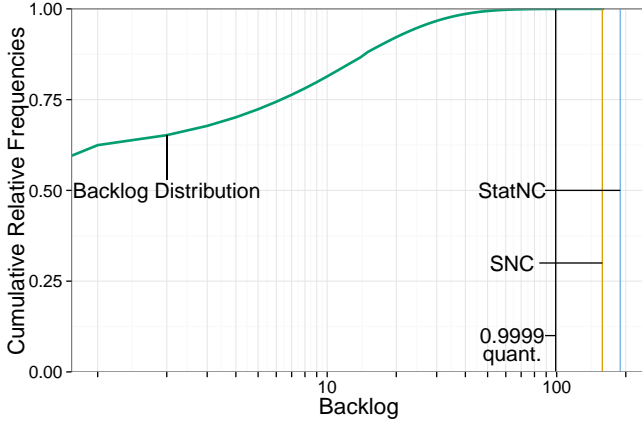


Figure 1. StatNC and SNC backlog bounds as well as the empirical distribution of the backlog measured at time $n = 100$ resulting from $N = 10^6$ simulation runs.

which means during the *On*-state, considering the bandwidth limitation, we see a peak utilization of $\frac{1}{\lambda c}(1 - e^{-\lambda M}) \approx 98\%$. Considering the bandwidth limitation and the fact, that we are not always in the *On*-state, we compute an average utilization of roughly $\approx 49\%$. The computation of the backlog bound based on SNC can be found in [4], while the StatNC bounds are computed according to Section IV-C. For illustration, we have simulated $N = 10^6$ runs of this system and evaluated the backlog at time slot 100 (at which time in all simulation runs the initial distribution of the Markov chain had faded out and steady-state was reached). In Figure 1, the empirical distribution function of the backlog is plotted; for the bounds a violation probability of $\varepsilon = 10^{-4}$ was assumed. As can be observed, both bounds are reasonably close to the $(1 - \varepsilon)$ -quantile; but, even more importantly, the bounds are pretty close to each other. This demonstrates that the price we pay for using StatNC is not too high.

As a side remark, we point out that the confidence level α , being a fraction of the violation probability ε , does not affect the quality of the StatNC bounds significantly, if not chosen at the extremes (α extremely close to 0 or ε). For a more detailed discussion, why α does only marginally affect the bounds, please refer to [4].

Scenario 2: Exploiting the Dynamic Behaviour of StatNC

In the second scenario, we investigate StatNC’s dynamic point of view. In particular, we use a sliding window approach over the last l observations (as discussed in Section III). Using this kind of sub-sampling, we eventually forget old measurements and “learn” from new arrivals instead. As such, the observation window allows to track changes in the arrival process (stemming, e.g., from non-stationarities such as time-of-day or other seasonal effects), which take place on longer timescales. For example, if we imagine a flow starting with large increments and diminishing over time, standard SNC faces problems; it lacks the adaptability to track this behaviour and its bounds get looser over time. On the other hand, StatNC can adapt by forgetting about the first large increments as time passes.

To investigate the adaptability, we use a Markov-modulated arrival process, similar to the previous one, but instead of having *On*- and *Off*-states, we use states *High* and *Low*. For both of these, arrivals are drawn from an exponential distribution with a parameter λ_{x_i} (and then capped by M); here, the parameter λ_{x_i} depends on the state of the Markov chain.

In this scenario, we use the estimators from Subsections IV-A and IV-B and not the estimator presented in Subsection IV-C. The goal is to use the observation window for tracking changes of states, instead of “learning” the Markov chain itself. Therefore, we use transition probabilities of $\mu = 0.999$ and $\nu = 0.999$; this means in expectation we stay 1000 time slots in one of the states until we change into the other one; this emulates a non-stationary behaviour of the arrival process. Further, we set $\lambda_{Low} = 5$ and $\lambda_{High} = 0.2$ and a bandwidth limitation of $M = 10$. With a service rate of $c = 5$, we have—taking the bandwidth limitation into account—a utilization of 4%, while residing in the *Low*-state and 86% in the *High*-state. In the simulations, we started the arrival process at time $n_0 = -1000$ to provide StatNC with an initial observation window. A typical run of this scenario is plotted in Figures 2 (for the exponential traffic estimator) and 3 (for the i.i.d. bandwidth-limited estimator). In addition to the bounds, we have plotted the simulated backlog process over time, to see how close the bounds are. Due to their dynamic nature, the StatNC bounds evolve over time. Similar to the SNC bounds, they are computed for a violation probability of $\varepsilon = 10^{-4}$ and for a time which lies n time slots after the point they have been computed; we provide the results for $n = 10$ and $n = 1000$ time slots, representing a short and long prediction horizon, respectively. As can be seen clearly, the bounds react and ultimately adapt to the observed arrivals: If the arrival intensity is high (indicated by larger backlogs), the statistical bounds also increase, while they decrease, when the Markov chain changes to the *Low*-state. One can also observe that the StatNC bounds track the changes of states with some delay, since old measurements need to be discarded from the observation window first. The effect of the prediction horizon n is such that a larger n results in higher bounds for both methods. For comparison, we also provide the SNC bounds calculated by exactly modelling the Markov chain. As can be observed, although the SNC uses complete information its bounds lie far above the StatNC bounds, which perform very well in terms of staying close, but not being violated too often (in accordance with the violation probability $\varepsilon = 10^{-4}$).

Another effect that can be observed, when comparing the two plots with each other, is how helpful the additional information about the exponential distribution is. We see, for the same run, that the StatNC bounds using knowledge about the type of distribution in the *High*-state, performs moderately better, compared to the non-parametric i.i.d. estimator of Subsection IV-B. We will demonstrate in the next scenario, however, that taking more assumptions about the arrivals into account, bears the risk of making false assumptions, which in turn can be fatal for the bounds.

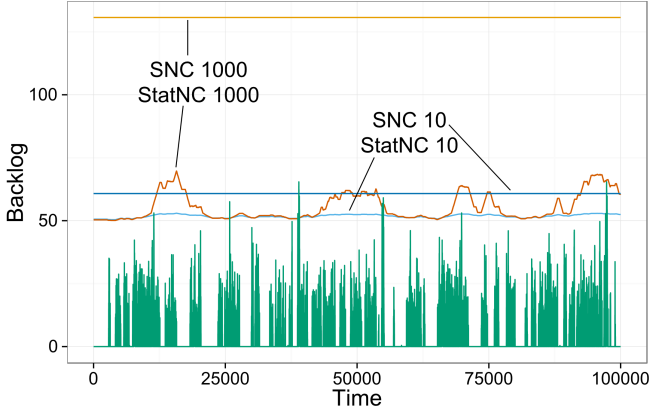


Figure 2. The backlog process for a typical simulation run; further the StatNC and SNC bounds for $n = 10, 1000$. For the StatNC bound the parametric estimator of Subsection IV-A was used.

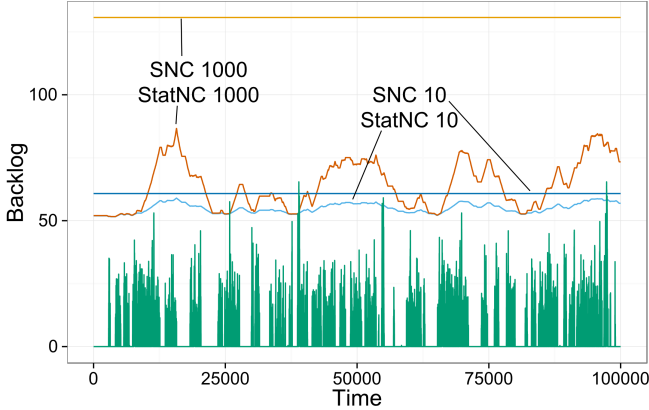


Figure 3. The backlog process for a typical simulation run; further the StatNC and SNC bounds for $n = 10, 1000$. For the StatNC bound the non-parametric estimator of Subsection IV-B was used.

Scenario 3: Robustness of StatNC

In the third scenario, we investigate the robustness of StatNC and SNC bounds against false assumptions on the arrival process. This reveals another feature of StatNC when using the estimator of Subsection IV-B: StatNC can cope with rather few assumptions about the arrivals and is therefore more robust than SNC.

To illustrate this, we let SNC make a false assumption about the distribution of the i.i.d. increments of the arrival process: For this we have chosen the increments to be i.i.d. Pareto distributed with parameters x_{min} and s , again capped by the bandwidth limitation M ; yet, for the calculation of the SNC bound we assume the increments to be i.i.d. exponentially distributed with parameter λ (again capped by M), with λ set such that the expectations of the capped Pareto-distributed arrivals and the assumed capped exponentially distributed arrivals coincide (for details see [4]). On the other hand, the StatNC using the non-parametric estimator from Subsection IV-B cannot—by definition—make such a false assumption about the arrivals.

In Figure 4, it can be observed that such a false assumption

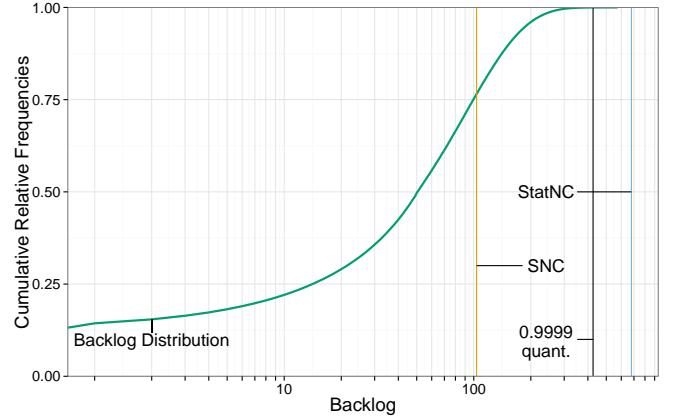


Figure 4. Under a false distribution assumption (exponential instead of Pareto distributed increments), SNC delivers a grossly invalid bound while StatNC remains correct.

on the arrival process can lead to disastrous results. As in the first scenario, the empirical distribution of the backlog is displayed and compared against the bounds calculated by StatNC and SNC. For the plot, parameters have been set as $x_{min} = 1, s = 1, M = 55$; further, we used a violation probability of $\varepsilon = 10^{-4}$ at time $n = 100$. The plot shows the empirical backlog distribution for 10^6 simulation runs, which in turn means, we would expect a tight 10^{-4} bound to be violated 100 times in expectation. The SNC bound however is broken by 234,526 runs, i.e., in approximately 23% of the simulations! This lies far below the empirical $(1 - \varepsilon)$ -quantile, the location of a sharp bound. This means the SNC is far too optimistic and hence is rendered useless. In contrast, the StatNC remains valid and stays reasonably close to the empirical quantile—a very satisfying result.

VI. CONCLUSION

By integrating statistical methods into the network calculus framework in order to deal with the frequent uncertainty about arrivals, we believe to have made an important step towards a better applicability of network calculus. In particular, the dynamic mode of operation of the newly developed statistical network calculus is attractive for many application fields where uncertainty and permanent change rules and modelling assumptions are already outdated when the actual system is under operation. After providing the basic technical results for StatNC—a framework theorem providing a sufficient condition for statistical estimators of the arrival process to connect them with the SNC framework and several matching estimators—we were able to make a case for the promising opportunities of the novel StatNC framework in a set of numerical experiments.

Given the positive results from this paper, there are many opportunities for future work within the StatNC framework: besides the already mentioned estimator for long-range dependent traffic, there are many more useful estimators that can be conceived; also, more sophisticated sub-sampling techniques than sliding windows, e.g., optimally weighted estimators, could provide even better reaction times; and last, but not least, a validation of the framework in a practical setting like

the ones mentioned in the introduction (MPLS domain, WSN) should provide new insights.

REFERENCES

- [1] D. Anick, D. Mitra, and M. M. Sondhi. Stochastic theory of a data-handling system with multiple sources. *The Bell System Technical Journal*, 61(8):1871–1894, 1982.
- [2] D. Awduche, J. Malcom, J. Agogbua, M. O’Dell, and J. McManus. Requirements for traffic engineering over mpls. RFC 2702, September 1999.
- [3] M. A. Beck. A first course in stochastic network calculus. Online lecture notes, April 2013. <http://disco.cs.uni-kl.de/index.php/projects/snc-course>.
- [4] M. A. Beck, S. A. Henningsen, S. B. Birnbach, and J. B. Schmitt. Towards a statistical network calculus—dealing with uncertainty in arrivals. Technical report, TU Kaiserslautern, 2013. <http://disco.informatik.uni-kl.de/discofiles/publicationfiles/BHBS13.pdf>.
- [5] G. Bouabene, C. Jelger, C. Tschudin, S. Schmid, A. Keller, and M. May. The autonomic network architecture (ana). *IEEE Journal on Selected Areas in Communications*, 28(1):4–14, 2010.
- [6] A. Burchard, J. Liebeherr, and S.D. Patek. A min-plus calculus for end-to-end statistical service guarantees. *IEEE Transactions on Information Theory*, 52(9):4105–4114, 2006.
- [7] 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.
- [8] C.-S. Chang. *Performance Guarantees in Communication Networks*. Telecommunication Networks and Computer Systems. Springer-Verlag, 2000.
- [9] F. Ciucu, A. Burchard, and J. Liebeherr. A network service curve approach for the stochastic analysis of networks. In *Proc. ACM SIGMETRICS*, pages 279–290, June 2005.
- [10] F. Ciucu and J. B. Schmitt. Perspectives on network calculus: no free lunch, but still good value. *Proc. of ACM SIGCOMM*, August 2012.
- [11] C. J. Clopper and E. S. Pearson. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*, 26(4):404–413, 1934.
- [12] R. L. Cruz. A calculus for network delay, Part I: Network elements in isolation. *IEEE Transactions on Information Theory*, 37(1):114–131, January 1991.
- [13] R. L. Cruz. A calculus for network delay, Part II: Network analysis. *IEEE Transactions on Information Theory*, 37(1):132–141, January 1991.
- [14] M. Fidler. An end-to-end probabilistic network calculus with moment generating functions. In *Proc. of IEEE IWQoS*, pages 261–270, June 2006.
- [15] Y. Gao and Y. Jiang. Analysis on the capacity of a cognitive radio network under delay constraints. *IEICE Transactions on Communications*, E95-B(4):1180 – 1189, April 2012.
- [16] R. Gibbens and F. Kelly. Measurement-based connection admission control. In *Proc. 15th International Teletraffic Congress*, pages 781–790, 1997.
- [17] M. Grossglauser and D. N.C. Tse. A framework for robust measurement-based admission control. *IEEE/ACM Transactions on Networking*, 7:293–309, June 1999.
- [18] S. Jamin, P.B. Danzig, S. J. Shenker, and L. Zhang. A measurement-based admission control algorithm for integrated services packet networks. *IEEE/ACM Transactions on Networking*, 5:56–70, February 1997.
- [19] Y. Jiang and Y. Liu. *Stochastic Network Calculus*. Springer, 2008.
- [20] J. Kurose. On computing per-session performance bounds in high-speed multi-hop computer networks. In *Proc. of ACM SIGMETRICS*, pages 128–139, 1992.
- [21] J. Liebeherr, A. Burchard, and F. Ciucu. Delay bounds in communication networks with heavy-tailed and self-similar traffic. *IEEE Transactions on Information Theory*, 58:1010–1024, 2012.
- [22] R. Lübben, M. Fidler, and J. Liebeherr. A foundation for stochastic bandwidth estimation of networks with random service. In *Proc. of IEEE INFOCOM*, pages 1817–1825, 2011.
- [23] P. Massart. The tight constant in the dvoretzky-kiefer-wolfowitz inequality. *Ann. Probab.*, 18(3):1269–1283, 1990.
- [24] J. Qiu and E. W. Knightly. Measurement-based admission control with aggregate traffic envelopes. *IEEE/ACM Transactions on Networking*, 9(2):199–210, April 2001.
- [25] A. Rizk and M. Fidler. Non-asymptotic end-to-end performance bounds for networks with long range dependent fbm cross traffic. *Computer Networks*, 56(1):127–141, January 2012.
- [26] K. Wang, F. Ciucu, C. Lin, and S. Low. A stochastic power network calculus for integrating renewable energy sources into the power grid. *IEEE Journal on Selected Areas in Communications: Smart Grid Communications Series*, 30:1037–1048, 2012.
- [27] 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.