Unleashing the Power of Paying Multiplexing Only Once in Stochastic Network Calculus

Anne Bouillard
anne.bouillard@huawei.com
Huawei Technologies France
Boulogne-Billancourt, France

Paul Nikolaus
nikolaus@cs.uni-kl.de
Distributed Computer Systems Lab (DISCO), TU Kaiserslautern
Kaiserslautern, Germany

Jens Schmitt
jschmitt@cs.uni-kl.de
Distributed Computer Systems Lab (DISCO), TU Kaiserslautern
Kaiserslautern, Germany

ABSTRACT

The stochastic network calculus (SNC) holds promise as a versatile and uniform framework to calculate probabilistic performance bounds in networks of queues. A great challenge to accurate bounds and efficient calculations are stochastic dependencies between flows due to resource sharing inside the network. However, by carefully utilizing the basic SNC concepts in the network analysis the necessity of taking these dependencies into account can be minimized. To that end, we unleash the power of the pay multiplexing only once principle (PMOO, known from the deterministic network calculus) in the SNC analysis. We choose an analytic combinatorics presentation of the results in order to ease complex calculations. In tree-reducible networks, a subclass of general feedforward networks, we obtain an effective analysis in terms of avoiding the need to take internal flow dependencies into account. In a comprehensive numerical evaluation, we demonstrate how this unleashed PMOO analysis can reduce the known gap between simulations and SNC calculations significantly, and how it favourably compares to state-of-the art SNC calculations in terms of accuracy and computational effort. Motivated by these promising results, we also consider general feedforward networks, when some flow dependencies have to be taken into account. To that end, the unleashed PMOO analysis is extended to the partially dependent case and a case study of a canonical topology, known as the diamond network, is provided, again displaying favourable results over the state of the art.

CCS CONCEPTS

• **Networks → Network performance modeling**.

KEYWORDS

stochastic network calculus; pay multiplexing only once

ACM Reference Format:
Anne Bouillard, Paul Nikolaus, and Jens Schmitt. 2022. Unleashing the Power of Paying Multiplexing Only Once in Stochastic Network Calculus. In *Abstract Proceedings of the 2022 ACM SIGMETRICS/IFIP PERFORMANCE Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS/PERFORMANCE '22 Abstracts), June 6–10, 2022, Mumbai, India.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3489048.3530964

SIGMETRICS/PERFORMANCE '22 Abstracts, June 6–10, 2022, Mumbai, India

ACM ISBN 978-1-4503-9141-2/22/06.
https://doi.org/10.1145/3489048.3530964

1 INTRODUCTION

Stochastic network calculus (SNC) is a mathematical framework with the goal to control tail probabilities for the end-to-end (e2e) delay, i.e., probabilities for rare events shall be bounded, e.g., $\mathbf{P}(\text{e2e delay} \geq 10\text{ms}) \leq 10^{-6}$. SNC originates in the deterministic analysis by Rene Cruz [5] and was later transferred to a stochastic setting, see [4] for a perspective.

Analysing more general networks of queues, in particular feedforward networks, usually requires the consideration of stochastically dependent flows. Even if all external arrival and service processes are independent, the sharing of resources by individual flows at queues generally makes them stochastically dependent at subsequent queues. How much this kind of dependencies has to be taken into account is affected by the network analysis method because different methods require different levels of knowledge about the internal characterization of flows. Further on, we call these dependencies *method-pertinent*. SNC analysis methods with less method-pertinent dependencies are strongly favourable as they are more accurate and more efficient. In fact, in [8], it has been observed that the PMOO analysis known from deterministic network calculus [1, 6, 9, 10] leads to less method-pertinent dependencies compared to the state of the art and is thus also promising in an SNC analysis. However, the application of PMOO in the SNC has, so far, been limited to so-called nested interference structures – this is very restrictive. The overall goal of our paper [2] is therefore to unleash the power of the PMOO principle in the SNC framework in order to not widen the known simulation-calculation gap [3] further, especially in more complex and larger networks of queues.

2 MAIN RESULT

We present a PMOO-based SNC end-to-end analysis for a subclass of feedforward networks, so-called tree-reducible networks; the main result is given in Theorem 1 below. It achieves *zero* method-pertinent stochastic dependencies when external arrivals and service processes are independent. I.e., if all input flows are assumed to be independent, we can derive bounds without resorting to Hölder's inequality. Also, Theorem 1 allows us to calculate the residual service in one big step avoiding the sequencing penalty in previous network analysis methods.

We introduce the necessary definition in order to present Theorem 1. A bivariate (arrival) process $A(s,t) = \sum_{i=s+1}^{t} a_i$ denotes the amount of data of a flow traversing at some server between discrete times s and t, $0 \leq s \leq t$. We define the departure process D accordingly. Let S be a non-negative bivariate function. A server is called a dynamic S-server if the relation between its arrival and departure processes satisfies $\forall t \geq 0, D(0,t) \geq \inf_{0 \leq s \leq t} \{A(0,s) + S(s,t)\}$ and

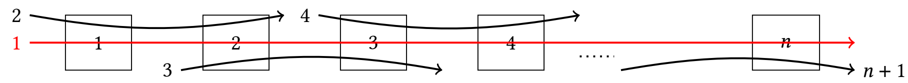

Figure 1: Extended interleaved tandem.

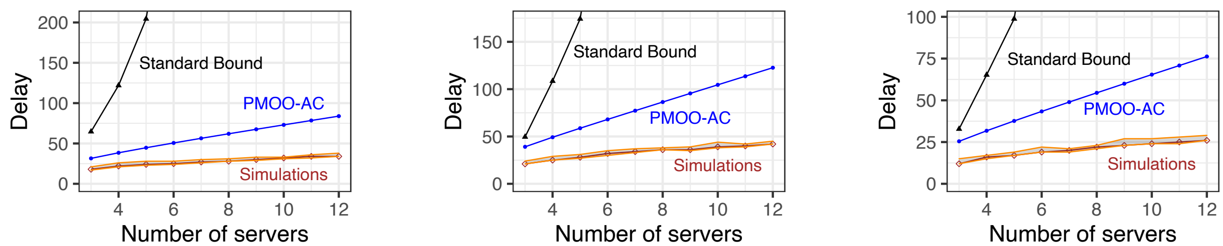

(a) Exponential distribution with $\lambda_i = 2.0$

(b) Weibull distribution with $k = 2$ **and** $\lambda_i = 0.7$

(c) MMOO with $p_{\text{on},i} = 0.5$, $p_{\text{off},i} = 0.5$, $P_i = 1.0$

Figure 2: Delay bounds for the extended interleaved tandem with server rates $C_i = 2.0$ **for** $i = 1, \ldots, 12$.

is called work-conserving if it offers at least service $S(s,t)$ during any backlogged period $[s,t]$.

For all flows f_i, $i \in \{1,\ldots,m\}$, we denote its path by $\pi_i = (\pi_i(1),\ldots,\pi_i(\ell_i))$, where ℓ_i is the length of the path of flow f_i.

THEOREM 1. *Assume a tree network with flows* $f_1,\ldots,f_m$ *and work-conserving* S*-servers* $S_1,\ldots,S_n$. *Assume a flow of interest* f_1 *traversing the servers* $(\pi_1(1),\ldots,\pi_1(\ell_1)=n)$. *We denote by* $S_{j^\bullet}$ *the successor of* S_j *(successors are unique for tree networks), with convention* $n^\bullet = n+1$, *and denote the indices of the time variables accordingly. By abuse of notation, each server offers the service* $S_j(s,t)$, $j \in \{1,\ldots,n\}$ *for* $0 \le s \le t$. *Then flow* f_1 *is at least offered a dynamic* S_{e2e}*-server* $\forall 0 \le t_{\pi_1(1)} \le t_{n+1}$ *with*

$$S_{\text{e2e}}(t_{\pi_1(1)}, t_{n+1}) = \left[\inf_{\forall j,\ t_j \le t_{j^\bullet}} \sum_{j=1}^{n} S_j(t_j, t_{j^\bullet}) - \sum_{i=2}^{m} A_i\Big(t_{\pi_i(1)}, t_{\pi_i(\ell_i)^\bullet}\Big) \right]_+ .$$

Based on bounds for the moment-generating function of arrivals and the Laplace transform of service processes, Theorem 1 is then used to calculate bounds of the form $\mathbf{P}(d \ge T) \le \varepsilon(T)$ for a given tree-reducible network. E.g. for a tandem with a frequently used interference pattern as in Figure 1, we calculate a stochastic delay bound T such that $\varepsilon(T) \le 10^{-6}$ and vary the tandem length. The results are displayed in Figure 2. While the state-of-the-art bound explodes in the number of servers (we have only included the results from 3 to 5 servers), the new PMOO-technique (PMOO-AC) scales significantly better.

3 FURTHER CONTRIBUTIONS

The following further contributions are elaborated in [2].

- We apply analytic combinatorics (AC) [7] to recover bounds from state-of-the-art analysis methods in simple networks and enable a generalization to more complex settings.
- We conduct an extensive numerical evaluation with respect to the accuracy of the new bounds for several traffic classes and different network topologies.
- We discuss first results to extend our new method from tree-reducible to general feedforward networks, still striving for the goal to minimize method-pertinent dependencies.

ACKNOWLEDGMENTS

This work was partially supported by Huawei Technologies Co., Ltd.

REFERENCES

[1] Anne Bouillard, Bruno Gaujal, Sébastien Lagrange, and Eric Thierry. 2008. Optimal routing for end-to-end guarantees using network calculus. *Performance Evaluation* 65, 11-12 (2008), 883–906.
[2] Anne Bouillard, Paul Nikolaus, and Jens Schmitt. 2022. Fully Unleashing the Power of Paying Multiplexing Only Once in Stochastic Network Calculus. *Proc. of the ACM on Measurement and Analysis of Computing Systems* 3, 3 (2022), 1–27.
[3] Florin Ciucu, Felix Poloczek, and Jens Schmitt. 2014. 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)*.
[4] Florin Ciucu and Jens Schmitt. 2012. Perspectives on Network Calculus – No Free Lunch, But Still Good Value. In *Proc. ACM Conf. on Applications, Technologies, Architectures, and Protocols for Computer Commun. (SIGCOMM'12)*. 311–322.
[5] Rene L Cruz. 1991. A calculus for network delay, part I: Network elements in isolation. *IEEE Transactions on information theory* 37, 1 (1991), 114–131.
[6] Markus Fidler. 2003. Extending the network calculus pay bursts only once principle to aggregate scheduling. In *International Workshop on Quality of Service in Multiservice IP Networks*. Springer, 19–34.
[7] Philippe Flajolet and Robert Sedgewick. 2009. *Analytic combinatorics*. Cambridge University Press.
[8] Paul Nikolaus and Jens Schmitt. 2017. On Per-Flow Delay Bounds in Tandem Queues under (In)Dependent Arrivals. In *Proc. 16th IFIP Networking Conference (NETWORKING'17)*. 1–9.
[9] Jens Schmitt and Frank A. Zdarsky. 2006. The DISCO Network Calculator - A Toolbox for Worst Case Analysis. In *Proceedings of the First International Conference on Performance Evaluation Methodologies and Tools (VALUETOOLS'06)*.
[10] Jens Schmitt, Frank A Zdarsky, and Ivan Martinovic. 2008. Improving Performance Bounds in Feed-Forward Networks by Paying Multiplexing Only Once. In *Proc. GI/ITG Conference on Measurement, Modeling, and Evaluation of Computer and Communication Systems (MMB'08)*. 1–15.