Sharp Bounds in Stochastic Network Calculus

Florin Ciucu, Felix Poloczek
T-Labs / TU Berlin

Jens Schmitt
University of Kaiserslautern

ABSTRACT

The practicality of the stochastic network calculus (SNC) is often questioned on grounds of potential looseness of its performance bounds. In this paper it is uncovered that for bursty arrival processes (specifically Markov-Modulated On-Off (MMOO)), whose amenability to *per-flow* analysis is typically proclaimed as a highlight of SNC, the bounds can unfortunately indeed be very loose (e.g., by several orders of magnitude off). In response to this uncovered weakness of SNC, the (Standard) per-flow bounds are herein improved by deriving a general sample-path bound, using martingale based techniques, which accommodates FIFO, SP, and EDF scheduling disciplines. The obtained (Martingale) bounds capture an additional exponential decay factor of $\mathcal{O}\left(e^{-\alpha n}\right)$ in the number of flows n, and are remarkably accurate even in multiplexing scenarios with few flows.

Categories and Subject Descriptors

H.1.1 [**Systems and Information Theory**]: Information Theory; C.4 [**Performance of Systems**]: Modeling techniques

Keywords

Stochastic network calculus, scheduling

1. INTRODUCTION

The stochastic network calculus (SNC) is a relatively recent methodology to solve queueing problems (see Chang [2] and Jiang and Liu [7]). From a technical point of view, SNC is a combination between the deterministic network calculus conceived by Cruz [5] and the effective bandwidth theory. Because SNC solves queueing problems in terms of *bounds*, it is often regarded as an unconventional approach, especially by the queueing theory community.

Based on its ability to partially solve hard queueing problems (i.e., in terms of bounds), SNC is justifiably proclaimed as a valuable alternative to the classical queueing theory (see Ciucu and Schmitt [4]). At the same time, SNC is also justifiably questioned on the tightness of its bounds. While the asymptotic tightness generally holds (see Chang [2], p. 291, and Ciucu *et al.* [1]), doubts on the bounds' numerical tightness shed skepticism on the practical relevance of SNC. This skepticism is supported by the fact that SNC largely employs

the same probability methods as the effective bandwidth theory, which was argued to produce largely inaccurate results for non-Poisson arrival processes (see Choudhury *et al.* [3]). Moreover, although the importance of accompanying bounds by simulations has already been recognized in some early works (see Zhang *et al.* [11] for the analysis of GPS), the SNC literature is scarce in that respect.

In this paper we reveal what is perhaps 'feared' by SNC proponents and expected by others: the bounds are very loose for the class of MMOO processes. In addition to providing numerical evidence for this fact (the bounds can be off by arbitrary orders of magnitude), we also prove that the bounds are asymptotically loose in most multiplexing regimes. Concretely, we (analytically) prove that the Standard bounds are 'missing' an exponential decay factor of $\mathcal{O}\left(e^{-\alpha n}\right)$ in the number of flows n, where $\alpha > 0$; this missing factor has been indicated through numerical experiments in Choudhury *et al.* [3] in the context of effective bandwidth bounds (which scale identically as the SNC bounds).

While this paper convincingly uncovers a major weakness in the SNC literature, it also shows that the looseness of the bounds is generally not inherent in SNC but it is due to the 'temptatious' but 'poisonous' elementary tools from probability theory leveraged in its application. In this sense, we prove that by leveraging more advanced tools (i.e., martingale based techniques), the SNC bounds improve dramatically to the point that they almost match simulation results. Concretely, we show these improvements to hold for per-flow delay bounds in FIFO, SP, and EDF scheduling scenarios with MMOO flows. Based on these improvements we argue that the core analysis in SNC, being reminiscent from the deterministic network calculus, is not only asymptotically but also numerically tight.

The sharp bounds obtained in this paper are the first in the conventional stochastic network calculus literature (i.e., involving service processes) concerning bursty arrivals. Their significance, relative to existing sharp bounds in the effective bandwidth literature (e.g., Duffield [6] and Chang [2], pp. 339-343, using martingale inequalities, or Liu *et al.* [9] by extending an approach of Kingman involving integral inequalities [8]), is that they apply at the *per-flow* level for various scheduling disciplines; in turn, existing sharp bounds only apply at the *aggregate* level. Our sharp bounds thus generalize existing ones by accounting for scheduling.

A weakness of our results, from a purely *network calculus perspective*, is that they are restricted to a *specific* class of processes, i.e., MMOO; we point out that one of the conceptual promises of the SNC is to provide general bounds for

much broader classes. While we thus deliberately sacrifice this conceptual generality, we also advocate a conceptual shift in running the SNC. Concretely, based on the results obtained in this paper, we believe that 1) SNC must be coupled with the mainstream queueing literature, in particular by "getting a firm grip on arrivals", and 2) the main two features of SNC (i.e., dealing with scheduling and multi-node) must be carefully leveraged in order to obtain sharp bounds.

2. MARTINGALE BOUNDS FOR MMOO

We consider a single queue whereby two cumulative arrival processes $A_1(t)$ and $A_2(t)$, each containing n_1 and n_2 MMOO processes, are served by a server with constant-rate $C = nc$, where $n = n_1 + n_2$. The time model is continuous. The following general result enables the per-flow analysis, in particular of the aggregate $A_1(t)$, for several scheduling algorithms (FIFO, SP, and EDF).

THEOREM 1. (MARTINGALE SAMPLE-PATH BOUND) *Consider the previous queueing system in which all $n_1 + n_2$ subflows are independent MMOO processes with transition rates μ and λ, and peak rate P, and starting in the steady-state. The aggregate arrival processes are $A_1(t)$ and $A_2(t)$, each being modulated by the (stationary) Markov processes $Z_1(t)$ and $Z_2(t)$ with n_1 and n_2 states, respectively. Assume that the utilization factor $\rho := \frac{pP}{c}$ satisfies $\rho < 1$ for stability, where $p := \frac{\mu}{\mu+\lambda}$ is the steady-state 'On' probability; assume also that $P > c$ to avoid a trivial scenario with zero delay. Then the following sample-path bound holds for all $0 \leq u \leq t$ and σ*

$$\mathbb{P}\left(\sup_{0 \leq s < t-u}\{A_1(s,t-u) + A_2(s,t) - C(t-s)\} > \sigma\right)$$
$$\leq K^n e^{-\gamma(C_1 u + \sigma)}, \quad (1)$$

where $C_1 = n_1 c$, $K = \rho\left(\frac{\rho-p}{1-p}\right)^{\frac{p}{\rho}-1}$, and $\gamma = \frac{(\lambda+\mu)(1-\rho)}{P-c}$.

The theorem generalizes a result by Palmowski and Rolski [10], which is restricted to an *aggregate* analysis under FIFO. The key to our proof is the construction of a single martingale M_t, from two existing martingales, such that the per-flow analysis for the different scheduling algorithms becomes possible. The sample-path bound from Eq. (1) then follows from a standard technique based on the Optional Sampling theorem, applied to the martingale M_t.

Using existing service processes for $A_1(t)$ for each of the three scheduling algorithms, the sample path bound from Eq. (1) lends itself to bounds on the virtual delay process $W_1(t) := \inf\{d \geq 0 : A_1(t-d) \leq D_1(t)\}$ of $A_1(t)$:

FIFO : $\mathbb{P}(W_1(t) > d) \leq K^n e^{-\gamma C d}$

SP : $\mathbb{P}(W_1(t) > d) \leq K^n e^{-\gamma C_1 d}$

EDF : $\mathbb{P}\left(W_1(t) > d\right) \leq K^n e^{\gamma C_2 \min\{d_1^*-d_2^*, d\}} e^{-\gamma C d}$.

The EDF bound holds for $d_1^* \geq d_2^*$, where d_1^* and d_2^* are the relative deadlines associated to $A_1(t)$ and $A_2(t)$, respectively. The bound for $d_1^* < d_2^*$ is similar and is omitted here. The parameters K and γ are as in Theorem 1, and $C_2 = n_2 c$.

Figure 1 illustrates the Martingale delay bounds obtained using Theorem 1, in contrast to the Standard bounds obtained using existing SNC methods, and simulations. The
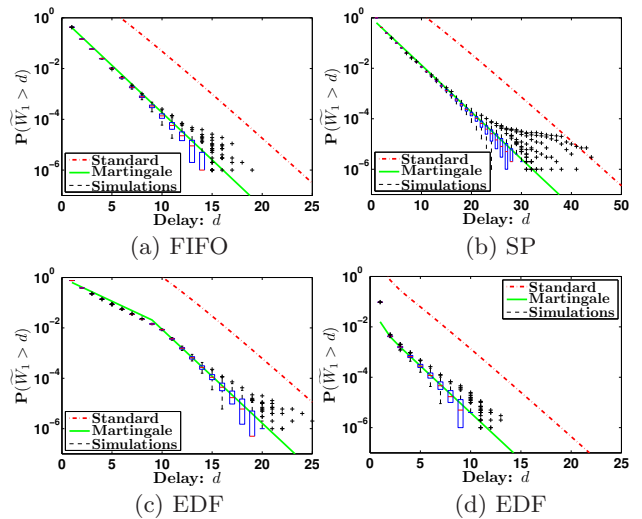

(a) FIFO (b) SP

(c) EDF (d) EDF

Figure 1: Delay bounds ($\lambda = 0.5$, $\mu = 0.1$, $P = 1$, $n_1 = n_2 = 10$, $\rho = 75\%$, ($d_1^* = 10$, $d_2^* = 1$) in (c) and $d_1^* = 1$, $d_2^* = 10$ in (d))

figure convincingly shows that the Standard bounds are very loose whereas the new Martingale bounds are quite sharp[1].

3. REFERENCES

[1] A. Burchard, J. Liebeherr, and F. Ciucu. On superlinear scaling of network delays. *IEEE/ACM Transactions on Networking*, 19(4):1043–1056, Aug. 2011.

[2] C.-S. Chang. *Performance Guarantees in Communication Networks*. Springer Verlag, 2000.

[3] G. Choudhury, D. Lucantoni, and W. Whitt. Squeezing the most out of ATM. *IEEE Transactions on Communications*, 44(2):203–217, Feb. 1996.

[4] F. Ciucu and J. Schmitt. Perspectives on network calculus - No free lunch but still good value. In *ACM Sigcomm*, 2012.

[5] R. Cruz. A calculus for network delay, parts I and II. *IEEE Transactions on Information Theory*, 37(1):114–141, Jan. 1991.

[6] N. G. Duffield. Exponential bounds for queues with markovian arrivals. *Queueing Systems*, 17(3-4):413–430, Sept. 1994.

[7] Y. Jiang and Y. Liu. *Stochastic Network Calculus*. Springer, 2008.

[8] J. F. C. Kingman. Inequalities in the theory of queues. *Jour. Royal Slat. Soc. Series B*, 32(1):102–110, 1970.

[9] Z. Liu, P. Nain, and D. Towsley. Exponential bounds with applications to call admission. *Journal of the ACM*, 44(3):366–394, May 1997.

[10] Z. Palmowski and T. Rolski. A note on martingale inequalities for fluid models. *Statistics & Probability Letters*, 31(1):13–21, Dec. 1996.

[11] Z.-L. Zhang, D. Towsley, and J. Kurose. Statistical analysis of generalized processor sharing scheduling discipline. In *ACM Sigcomm*, pages 68–77, 1994.

---

[1]Outliers are depicted in the box-plots with the '+' symbol; on each box, the central mark is the median, and the edges of the box are the 25th and 75th percentiles.