Exponential Supermartingales for Evaluating End-to-End Backlog Bounds

Florin Ciucu
Department of Computer Science, University of Virginia

Abstract

A common problem arising in network performance analysis with the stochastic network calculus is the evaluation of $(min, +)$ convolutions. This paper presents a method to solve this problem by applying a maximal inequality to a suitable constructed supermartingale. For a network with D/M input, end-to-end backlog bounds obtained with this method improve existing results at low utilizations. For the same network, it is shown that at utilizations smaller than a certain threshold, fluid-flow models may lead to inaccurate approximations of packetized models.

1 Introduction

Stochastic network calculus is a theory for the performance analysis of networks in terms of probabilistic bounds. A feature of the calculus is that end-to-end backlog and delay bounds can be derived using a $(min, +)$ convolution operation. The bounds obtained in this manner are generally much tighter than corresponding results obtained by adding per-node bounds (see Ciucu *et. al.* [6], Fidler [7]).

A common technique for the problem of evaluating backlog bounds in network calculus is to invoke Boole's inequality, i.e., $Pr(\sup X_s > x) \leq \sum Pr(X_s > x)$. For the same problem restricted to the class of Lévy processes with bounded moment generating functions, an alternative technique is to apply Doob's maximal inequality, i.e., $Pr(\sup e^{\theta X_s} > e^{\theta x}) \leq E[e^{\theta X_0}]e^{-\theta x}$, where $e^{\theta X_s}$ is a supermartingale. This technique is reminiscent of the derivation of backlog bounds in GI/GI/1 queues, in a classic note by Kingman [10]. In previous work [5], we have shown in a network calculus setting that Doob's inequality leads to improved bounds than those obtained with Boole's inequality, in the single-node case.

In the multi-node case, multiple $(min, +)$ convolutions make difficult the construction of supermartingales to extend the benefits shown in the single-node case. In this paper we resolve this problem in scenarios with low utilizations. As a consequence, for a non-trivial fluid-flow network scenario, we find an utilization threshold below which the end-to-end backlog is uniformly bounded in the number of nodes. In contrast, packetization causes the end-to-end backlog to increase (see Burchard *et. al.* [1]).

We thus identified a scenario where fluid-flow models may be unsuitable to approximate packet networks. This finding presents interest since fluid-flow models are frequently used to simplify network analysis with theories such as effective bandwidth (Kelly [8]) or network calculus (Chang [3]). Another attractive feature of fluid-flow models is that they can speed up packet-based simulations in small networks, but this advantage may disappear in networks with large number of nodes (see Kesidis *et. al.* [9]).

A notable characteristic of the utilization threshold found in this paper is that it decreases with the number of nodes H as $1/\log H$. This raises an interesting parallel with a result obtained by Charny and Le Boudec [4]. There, it is shown that the end-to-end delay of a flow traversing H nodes with packetization grows with H, as long as the utilization decreases as $1/H$; otherwise, the delays may grow unbounded. The result holds for high priority and deterministically regulated flows with aggregate scheduling. In contrast, our result holds for statistical arrivals and any workconserving scheduling. The importance of the result in [4] is that it holds for networks with general topologies; in turn, we only consider a single-path network.

In Section 2 we derive an end-to-end backlog bound with Doob's inequality, and for comparison we present a second bound obtained with Boole's inequality. The scaling and numerical comparisons of these bounds are discussed in Section 3 for networks with D/M input.

2 End-to-End Backlog Bounds

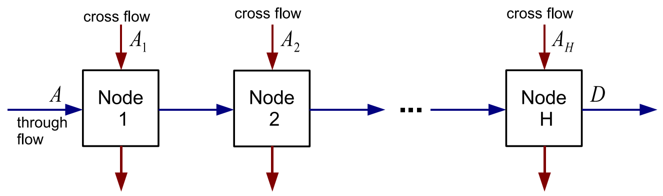

Figure 1: A network with cross traffic

We consider the network scenario from Figure 1. A through flow traverses H nodes in series, and additional cross flows arrive at the nodes. All flows are discrete time processes, *i.i.d.*, have stationary and independent increments, and also have bounded moment generating functions. Each node is workconserving, and has capacity $C > 0$.

The next theorem gives our main result, i.e., a bound on the end-to-end backlog process $B(t) = A(t) - D(t)$ of the through flow traversing the network.

Theorem 1 *Given the network scenario above, assume that there exits* $\theta > 0$ *such that*

$$E\left[e^{\theta A(1)}\right] E\left[e^{\theta \max_h \{A_h(1)\}}\right] \leq e^{\theta C} \tag{1}$$

holds. Then we have for all $t, \sigma \geq 0$

$$Pr\Big(B(t) > \sigma\Big) \leq e^{-\theta\sigma} \ . \tag{2}$$

Proof. Fix $t \geq 0$. Following [7] we have that $\mathcal{S}^h(s,t) = C(t-s) - A_h(s,t)$ are statistical service curves for the through flow at each node h, and furthermore $D(t) \geq A * \mathcal{S}^1 * \ldots * \mathcal{S}^H(t)$, where $f * g(s,t) := \inf_{s \leq u \leq t}(f(s,u) + g(u,t))$. Then we can write for the backlog: $B(t) \leq \sup_s \left[A(s,t) - \mathcal{S}^1 * \ldots * \mathcal{S}^H(s,t)\right]$. Denoting $T(s)$ the process in the braces, with s replaced by $t-s$, we make the critical observation that

$$T(s+1) \leq T(s) + A(t-s-1, t-s) - C + \max_h \{A_h(t-s-1, t-s)\} \ .$$

Let us now choose $\theta > 0$ satisfying Eq. (1). Using the stationary and independent increments properties, we can write for an appropriate filtration $\mathcal{F}_s$

$$E\left[e^{\theta T(s+1)} \parallel \mathcal{F}_s\right] = e^{\theta T(s)} e^{-\theta C} E\left[e^{\theta A(1)}\right] E\left[e^{\theta \max_h \{A_h(1)\}}\right] \leq e^{\theta T(s)} \ ,$$

thus showing that $e^{\theta T(s)}$ is a supermartingale. Invoking Doob's maximal inequality we obtain the end-to-end backlog bound for all $\sigma \geq 0$

$$Pr\left(B(t) > \sigma\right) \leq Pr\left(\sup_s e^{\theta T(s)} > e^{\theta\sigma}\right) \leq e^{-\theta\sigma} \ ,$$

that completes the proof. □

For comparison, we next follow [3, 7] to obtain a second end-to-end backlog bound with Boole's inequality. We first apply the Chernoff bound in $Pr(\sup_s T_s > \sigma)$ and then apply Boole's inequality in the form $E[\sup X_s] \leq \sum E[X_s]$ for nonnegative X_s. Evaluating the sums and letting $\theta > 0$ such that

$$\max_h E\left[e^{\theta A(1)}\right] E\left[e^{\theta A_h(1)}\right] < e^{\theta C} \ , \tag{3}$$

we obtain the end-to-end backlog bound for all $t, \sigma \geq 0$

$$Pr\Big(B(t) > \sigma\Big) \leq \left(\frac{1}{1 - e^{-\theta C} E\left[e^{\theta A(1)}\right] E\left[e^{\theta A_1(1)}\right]}\right)^H e^{-\theta\sigma} \ . \tag{4}$$

For fixed values of θ, Eq. (2) yields tighter bounds than Eq. (4). However, the condition from Eq. (1) poses more constraints on θ than the condition from Eq. (3). To see the additional constraints in Eq. (1), Jensen's inequality (i.e. $exp(\theta E[X]) \leq E\left[exp(\theta X)\right]$) gives a necessary condition for Eq. (1) to hold:

$$E\left[A(1)\right] + E\left[\max_h \{A_h(1)\}\right] \leq C \ . \tag{5}$$

This is stronger than a stability condition and it implies that the utilization should decrease as the number of nodes grows. In contrast, Eq. (3) implies only the stability condition

$$E\left[A(1)\right] + \max_h E\left[A_h(1)\right] \leq C \ . \tag{6}$$

3 Application to a Network with D/M Input

We now consider a special case of the network from Figure 1, where each flow consists of packets with constant interarrival times equal to $1/\lambda$, and exponentially distributed sizes with mean $1/\mu$. We denote the utilization factor $\rho = 2\lambda/(\mu C)$ and assume $\rho < 1$. The time unit is $1/\lambda$. We assume a fluid-flow service model, i.e., each fraction of a packet becomes available for service as soon as processed upstream.

Applying Theorem 1, the condition from Eq. (1) reduces for $0 < \theta < \mu$ to

$$\left(1 - \frac{\theta}{\mu}\right)^{-1} \prod_{h=1}^{H} \left(1 - \frac{\theta}{h\mu}\right)^{-1} \leq e^{\frac{2\theta}{\rho\mu}} \ . \tag{7}$$

Here we used a result by Rényi on order statistics stating that $\max\{X_h\} = \sum \frac{X_h}{h}$ in distribution, when X_h are independent and exponentially distributed [11]. Hence, for values of θ satisfying Eq. (7), the end-to-end backlog bound is given by Eq. (2). To illustrate the scaling of this bound we choose $\theta = \frac{\mu}{2}$ which permits an explicit computation of the left-hand side of Eq. (7):

$$\left(1 - \frac{1}{2}\right)^{-1} \prod_{h=1}^{H} \left(1 - \frac{1}{2h}\right)^{-1} = \frac{2\left(H! \cdot 2^H\right)^2}{(2H)!} = 2\sqrt{\pi H}\left(1 + \mathcal{O}\left(\frac{1}{H}\right)\right) \ ,$$

after using Stirling's approximation for the factorial. Solving for ρ in Eq. (7) yields

$$\rho \leq \frac{1}{1.3 + \log\sqrt{H} + \mathcal{O}\left(\frac{1}{H}\right)} \ . \tag{8}$$

Some numerical values for the above threshold are given in Table 1. The threshold agrees with the observation we made about Eq. (5). Moreover, since we chose $\theta = \frac{\mu}{2}$, it follows from Eq. (2) that the end-to-end backlog is uniformly bounded in the number of nodes H at utilizations smaller than the threshold. A similar conclusion can be drawn by plugging-in the threshold into Eq. (4).

Number of nodes	10	10^2	10^3	10^6
Utilization threshold	0.41	0.27	0.21	0.12

Table 1: Threshold in Eq. (8), ignoring the asymptotic error, as a function of the number of nodes.

Let us next derive the scaling of the end-to-end backlog when accounting for packetization, i.e., each packet becomes available for service as soon as fully processed upstream. We also assume that each node is locally FIFO, and each packet maintains its size constant at each traversed node. Using results from [1] with ρ chosen as in Eq. (8), we obtain that end-to-end delays (hence end-to-end backlogs as well) grow as $\Theta\left(H \log \frac{H}{\log H}\right)$. Besides packetization effects, this order of growth is also determined by statistical correlations arising from maintaining the packets' sizes constant. In contrast, we note that such correlations are not present in the fluid-flow model.

At last, Figures 2.(a,b,c) numerically compare the end-to-end backlog bounds obtained with Doob's inequality (Eq. (2)) and Boole's inequality (Eq. (4)), adapted to the D/M case. The parameter θ is numerically optimized in the conditions from Eq. (7) and Eq. (3), respectively. In (a), both bounds show that the end-to-end backlog is uniformly bounded in H, at very small utilizations; the bounds are preserved even at extreme values of H (e.g. 10^6, which is not displayed here), agreeing with the last pair of values from Table 1.

Figure 2.(b) shows that the bounds from Eq. (2) are tighter than the bounds from Eq. (4) at utilizations smaller than approximatively 25%. The reason is that at utilizations smaller than the threshold from Eq. (8), similar θ's hold in Eqs. (2), (4), but the exponential in Eq. (2) has a smaller prefactor than in Eq. (4). We observed that at utilizations around 30%, the bounds from Eq. (2) start to degrade significatively for $H \geq 100$; this agrees with the second pair of values from Table 1. The blow-up is illustrated in (c) at $\rho = 40\%$; we remark that bigger ρ implies an earlier blow-up in H.

To conclude, this paper shows that supermartingales based techniques can improve upon techniques based on Boole's inequality at small utilizations. Other sufficient conditions under which this advantage holds are provided by Chang [2] for stochastic linear systems under the $(max, +)$ algebra. Whether this advantage holds in general for the multi-node network analysis remains open.

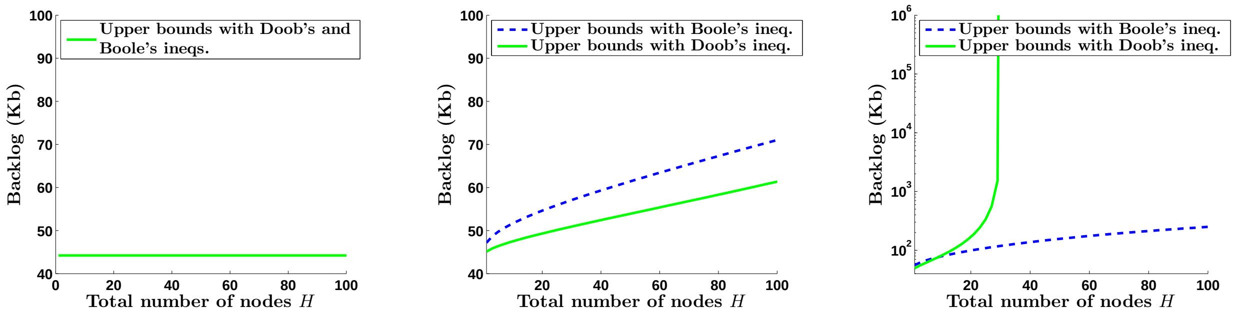

(a) $\rho = 5\%$ (b) $\rho = 25\%$ (c) $\rho = 40\%$

Figure 2: End-to-end backlog bounds in a fluid-flow network with cross traffic and D/M input, as a function of the number of nodes H (capacity $C = 100\ Mbps$, utilization factors ($\rho = 5\%$, $\rho = 25\%$, and $\rho = 40\%$), average packet size $1/\mu = 400\ Bytes$, violation probability $\varepsilon = 10^{-6}$)

Acknowledgements

The research in this paper is supported in part by the National Science Foundation under grant CNS-0435061. The author gratefully acknowledges valuable comments from Almut Burchard and Jörg Liebeherr.

References

[1] A. Burchard, J. Liebeherr, and F. Ciucu. On $\Theta\left(H \log H\right)$ scaling of network delays. In *Proceedings of IEEE Infocom*, May 2007.

[2] C.-S. Chang. On the exponentiality of stochastic linear systems under the max-plus algebra. *IEEE Transactions on Automatic Control*, 41(8):1182–1188, Aug. 1996.

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

[4] A. Charny and J.-Y. Le Boudec. Delay bounds in a network with aggregate scheduling. In *Proceedings of the First International Workshop on Quality of Future Internet Services*, pages 1–13, 2000.

[5] F. Ciucu. Network calculus delay bounds in queueing networks with exact solutions. In *20th International Teletraffic Congress (ITC)*, June 2007.

[6] F. Ciucu, A. Burchard, and J. Liebeherr. A network service curve approach for the stochastic analysis of networks. In *Proceedings of ACM SIGMETRICS*, volume 33, pages 279–290, 2005.

[7] M. Fidler. An end-to-end probabilistic network calculus with moment generating functions. In *IEEE 14th International Workshop on Quality of Service (IWQoS)*, pages 261–270, June 2006.

[8] F. P. Kelly. Notes on effective bandwidths. In *Stochastic Networks: Theory and Applications. (Editors: F.P. Kelly, S. Zachary and I.B. Ziedins) Royal Statistical Society Lecture Notes Series, 4*, pages 141–168. Oxford University Press, 1996.

[9] G. Kesidis, A. Singh, D. Cheung, and W. Kwok. Feasibility of fluid event-driven simulation for ATM networks. In *IEEE Globecom*, pages 2013 – 2017, Nov. 1996.

[10] J. F. C. Kingman. A martingale inequality in the theory of queues. *Cambridge Philosophical Society*, 59:359–361, 1964.

[11] A. Rényi. On the theory of order statistics. *Acta Math. Acad. Sci. Hungarica*, 4:191–232, 1953.