Price of Asynchrony: Queuing under Ideally Smooth Congestion Control

Maxim Podlesny and Sergey Gorinsky
{*podlesny,gorinsky*}*@arl.wustl.edu*
Applied Research Laboratory, Department of Computer Science and Engineering, Washington University
One Brookings Drive, Campus Box 1045, Saint Louis, Missouri 63130-4899, United States of America

Introduction Smooth fair lossless transmission at high bitrates is an aspiration for many explicit and delay-driven congestion control protocols [1]–[4]. Potential benefits of such transmission include short queues and no data loss at bottleneck links. These properties are particularly important for interactive multimedia and other applications that would suffer from excessive queuing delays at shared network links.

In this paper, we present a model for investigating lower bounds on queuing under smooth congestion control with overprovisioned buffers. We consider an idealized protocol where all flows always transmit at equal rates. The ideally smooth transmission does not eliminate queuing altogether because packets of different flows might arrive to a node simultaneously due to asynchronous arrivals of the flows, which is intrinsic to computer networks. A prominent feature of our model is its simplicity, making analysis tractable and experiments scalable. Our results reveal steady-state queues of at least $O(\sqrt{N})$ packets, where N is the number of flows. Hence, no congestion control protocol is able to avoid losses at a fully utilized link with a constant buffer shared by arbitrarily many flows.

Model We model a steady-state scenario where N flows share a bottleneck link with bitrate C and a FIFO (First-In First-Out) buffer. We denote arrival time of flow i as t_i, where $i = 1, \ldots, N$. Without loss of generality, we assume $t_1 = 0$. Average utilization of the link is U, where $0 < U \leq 1$. Each flow transmits packets of size S periodically at the same constant bitrate R equal to:

$$R = \frac{U \cdot C}{N}. \tag{1}$$

Hence, subsequent packets within any flow are separated by the same time interval T:

$$T = \frac{N \cdot S}{U \cdot C} = \frac{N \cdot D}{U} \tag{2}$$

where D is per-packet transmission time.

This pattern of packet transmissions is the smoothest possible under asynchronous congestion control where distributed senders of different flows do not deliberately schedule packets to arrive at a shared link at non-overlapping times. After the last flow arrives, the imperfect alignment of the flows creates a queue oscillation pattern that repeats with period T.

We consider three smooth distributions of flow interarrival times: Exponential, Uniform, and Pareto. All three distributions have the same average value:

$$\mu = \frac{D}{U} = \frac{T}{N}, \tag{3}$$

i.e., the N flows are expected to arrive over a time interval that has the same duration T as the period of the steady-state queue oscillations. The variances of the distributions are $\left(\frac{T}{N}\right)^2$, $\frac{1}{3}\left(\frac{T}{N}\right)^2$, and $\frac{1}{k(k-2)}\left(\frac{T}{N}\right)^2$ respectively, where $k = 2.1$ is Pareto index.

Analysis We conduct a stochastic analysis of steady-state queuing in the overprovisioned buffer of a fully utilized link with N flows. We show that the number of flows arriving outside time interval $[0; T)$ is negligible in comparison to the total number of flows. Therefore, we assume that all flows arrive within time interval $[0; T)$. Then, we express queue size q_i encountered by the i-th flow during steady-state time interval $[T; 2T)$ as:

$$q_i = q_0 + i - \frac{t_i}{D} \tag{4}$$

where q_0 represents the queue size at time T, i is the number of packets that have arrived during time interval $[T; T + t_i)$, and $\frac{t_i}{D}$ denotes the number of packets transmitted into the link during this interval $[T; T + t_i)$.

Let $\theta(i)$ denote the probability that the i-th flow encounters a steady-state queue longer than Q. Then, applying the Central Limit Theorem, we express $\theta(i)$ as:

$$\theta(i) = \frac{1}{2}\left(1 - \operatorname{erf}\left(\frac{Q}{A_\psi \sqrt{i}}\right)\right) \tag{5}$$

1-4244-1588-8/07/$25.00 ©2007 IEEE

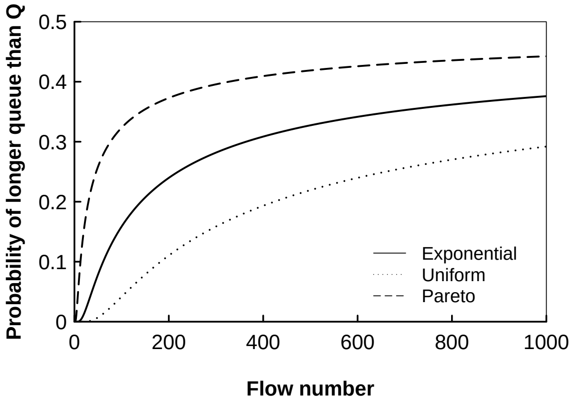

Fig. 1. Probability that the steady-state queue encountered by a flow is longer than Q = 10 packets (fully utilized link with 1,000 flows, overprovisioned buffer, and various interarrival time distributions).

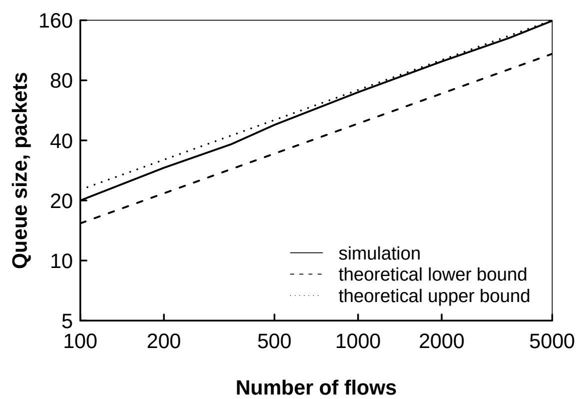

Fig. 2. Steady-state queue size encountered by the top 1% flows in the overprovisioned buffer of a fully utilized link with Exponential flow interarrival times.

where A_ψ is a coefficient associated with the interarrival time distribution; A_ψ equals $\sqrt{2}$, $\sqrt{\frac{2}{3}}$, and $\sqrt{\frac{2}{k(k-2)}}$ respectively for the considered Exponential, Uniform, and Pareto distributions. Figure 1 shows profiles of $\theta(i)$ for $N = 1,000$ and $Q = 10$ packets. We represent the probability of a steady-state queue longer than Q as the ratio of the average number of flows experiencing a steady-state queue longer than Q to the total number of flows:

$$\theta = \frac{1}{N}\sum_{i=1}^{N}\theta(i). \quad (6)$$

We derive lower $Q_{\min}$ and upper $Q_{\max}$ bounds on the steady-state queue size for the top θ flows as:

$$Q_{\min} = L_\theta A_\psi \sqrt{N} \text{ and } Q_{\max} = E_\theta A_\psi \sqrt{N}. \quad (7)$$

where L_θ and E_θ depend only on fraction θ of flows; $L_\theta \approx 1.1$ and $E_\theta \approx 1.6$ for $\theta = 1\%$. Lower bound $Q_{\min}$ reveals impossibility to avoid packet losses at a fully utilized link with a constant buffer and arbitrarily many flows:

Theorem 1: The steady-state queue in the overprovisioned buffer of a fully utilized link is at least $O(\sqrt{N})$ packets, where N is the number of flows sharing the link.

Simulations To validate the above bounds, we conduct simulations within our model. We vary N from 100 to 5,000 flows. For each value of N, we perform 1,000 experiments with the following parameters: $U = 1$, $C = 100$ Mbps, and $S = 1{,}000$ bytes (neither S nor C affects queuing in our model). Our simulation methodology allows us to capture the steady state exactly by looking at only $2N$ packets, or 2 packets per flow.

Figure 2 reports analytical and experimental results for the queue size encountered by the top 1% flows for Exponential flow interarrival times. This and our other experiments are generally consistent with the above theoretical conclusion that steady-state queuing in the overprovisioned buffer of a fully utilized link is at least $O(\sqrt{N})$.

Discussion We exposed lower bounds on steady-state queuing at highly multiplexed links under any congestion control protocol. Our results imply impossibility to avoid packet loss at a fully utilized link with a constant buffer size and arbitrarily many flows. In subsequent work [5], we explore avoiding the losses by underutilizing the bottleneck link. We also investigate fully utilized links with small buffers and surprisingly show that loss rates under our ideally smooth congestion control have lower bounds that are independent of N. This suggests possibility of practical congestion control where small buffers provide the double benefit of short queues and bounded loss rates at fully utilized links.

References

[1] K.K. Ramakrishnan, R. Jain, “A Binary Feedback Scheme for Congestion Avoidance in Computer Networks with a Connectionless Network Layer,” in *Proceedings ACM SIGCOMM 1988*, August 1988.

[2] R. Jain, “A Delay-Based Approach for Congestion Avoidance in Interconnected Heterogeneous Computer Networks,” in *ACM Computer Communications Review*, 19(5), pp. 56-71, 1989.

[3] L. Brakmo, S. O’Malley, and L. Peterson, “TCP Vegas: New Techniques for Congestion Detection and Avoidance,” in *Proceedings ACM SIGCOMM 1994*, August 1994.

[4] M. Podlesny and S. Gorinsky, ”Multimodal Congestion Control for Low Stable-State Queuing” in *Proceedings IEEE INFOCOM 2007*, May 2007

[5] M. Podlesny and S. Gorinsky, “Lower Bounds on Queuing and Loss at Highly Multiplexed Links,” *Technical Report WUCSE-2007-42*, Department of Computer Science and Engineering, Washington University in St. Louis, July 2007.