

Burstiness as Traffic Metric in Next-Generation Optical Core Networks

Tilman Wolf, Weibo Gong, Yan Cai

Department of Electrical and Computer Engineering
University of Massachusetts, Amherst, MA, USA
{wolf,gong,ycai}@ecs.umass.edu

Abstract—Next-generation all-optical packet-switched core networks use very small packet buffers. We argue that it is essential to focus on burstiness rather than bandwidth when conditioning traffic for effective network operation.

I. INTRODUCTION

The all-optical core of the next-generation Internet exhibits characteristics that are the exact opposite of traditional network. Since the first days of the Internet, networks have been limited in communication bandwidth but have had large packet buffers. In contrast, emerging optical networks provide vast amounts of bandwidth, but are very limited in their ability to implement packet buffers efficiently. This change in network design causes problems when applying obsolete operational principles that are based on outdated technology assumptions. In particular, the commonly used practice of using bandwidth as the sole metric to describe network traffic is problematic. We argue that *traffic burstiness is equally as important as bandwidth when characterizing traffic for small-buffer optical networks*. To operate a small-buffer optical network, it is essential to actively condition network traffic to reduce burstiness.

Packet buffers in routers have traditionally been large, sized in the order of hundreds of megabytes to accommodate the product of bandwidth and end-to-end delay. This size was believed to be necessary to compensate for contention due to statistical multiplexing and to avoid underflow due to TPC congestion control. More recently, it has been argued that smaller buffers, scaling with the logarithm of the congestion window size of flows, are sufficient [4]. However, all-optical packet-switched networks have even lower limits on the number of packets that can be buffered due to the cost and complexity of optical buffers [7].

A key problem in the operation of small-buffer networks is to ensure that network traffic can be transmitted effectively despite the lack of buffer space. One potential approach is to reduce the load on the network by limiting the aggregate bandwidth of traffic. However, in the following section we show that it is not bandwidth, but burstiness that causes queuing on routers and the resulting packet losses. Therefore, we propose a network-wide pacing architecture for next-generation optical core networks.

This material is based upon work supported by the National Science Foundation under Grant No. DRL-0632586.

II. IMPACT OF BURSTINESS ON SMALL-BUFFER NETWORKS

We illustrate the impact of traffic burstiness on small-buffer networks by analyzing the impact of bursts on queue lengths. The more queuing is caused by a burst, the more packets are dropped in network with small buffers. We have analyzed this impact using a tandem queue model. This model represents typical Internet connections, which are multi-hop over decreasing link capacity (when approaching the bottleneck link and due to competing cross-traffic).

There are two basic approaches to analyzing such queuing models: classical queuing theory and stochastic fluid model. Classical queuing theory focuses on the packet-level behavior of a queuing system with certain assumptions on the distributions of inter-arrival time and service time. It is widely used in performance evaluation of network protocols (e.g., in the context of buffer sizing [6]). Stochastic fluid models treat packet arrival bursts as a continuous fluid. As demonstrated for tandem queuing models [5], Markov On-Off fluid models are able to capture the impact of source correlation on the average queue size.

We first focus on a fluid model representation of the tandem queue network. In this model, input traffic is modeled as a Markov On-Off flow. Each of M nodes is equipped with a queue to buffer packets to prevent packet drop when the input rate exceeds the outgoing capacity. The instantaneous queue length of the buffer at node i is denoted as $v_i(t)$.

The dynamics of such a model can be described by a group of Poisson Counter Driven Stochastic Differential Equations (PCSD). In particular, we have

$$\begin{cases} dx &= (1-x)dN_1 - xdN_2; \\ dv_1 &= -c_1 I_{v_1} dt + hxdx; \\ dv_2 &= -c_2 I_{v_2} dt + c_1 I_{v_1} dt; \\ \dots & \\ dv_M &= -c_M I_{v_M} dt + c_{M-1} I_{v_{M-1}} dt, \end{cases} \quad (1)$$

where I_{v_i} is an indicator function with value 1 whenever $v_i(t) > 0$ and 0 otherwise. Two constraints on parameters h and c_i 's are $h > c_1 > \dots > c_M$ and $c_M > hE[x]$, where c_i are the link capacities and h is the peak input rate. The first inequality constraint is due to the fact that for node i with $c_i > h$, $v_i(t)$ is 0 all the time. The second constraint is to make the whole system stable.

The expected queue length $E[v_1]$ is derived in Equation (23) in [1] as

$$E[v_1] = \frac{h - c_1}{(c_1 - hE[x])(\lambda_1 + \lambda_2)} hE[x], \quad (2)$$

where λ_1 and λ_2 are average arrival rates of On periods and Off periods respectively. An explicit expression of $E[v_i]$ for $i = 2, \dots, M$ is given by Equation (39) in [1]. It can be further simplified and we obtain a simpler form as follows.

$$E[v_i] = \frac{(c_{i-1} - c_i)\lambda_2 h^2 E^2[x]}{\lambda_1(\lambda_1 + \lambda_2)(c_{i-1} - hE[x])(c_i - hE[x])}. \quad (3)$$

When $E[x] = \frac{\lambda_1}{\lambda_1 + \lambda_2}$ and the c_i s and h are kept constant, the mean queue length at node i is proportionally impacted by the inverse of $\lambda_1 + \lambda_2$, which is the time constant of the autocorrelation function of $x(t)$ [5]. Let $\lambda'_1 = \alpha\lambda_1$ and $\lambda'_2 = \alpha\lambda_2$. Substituting such λ'_1 and λ'_2 into Eq. (2) and Eq. (3), we obtain

$$E[v'_i] = \frac{E[v_i]}{\alpha}, \quad (4)$$

which indicates a linear impact while the the average input rate $hE[x]$ is kept constant.

While the above formula demonstrates the linear impact of the arrival burstiness to the mean queue lengths in a tandem queue for Markov on-off sources, such impact to the instantaneous queue length in the tandem queue is not clear. To further demonstrate this impact and also include more general arrival processes, we have carried out an infinitesimal perturbation analysis to a tandem queue with a marked point process arrival with i.i.d. inter-arrival distribution and i.i.d. mark distribution in [3].

Our analysis shows that the instantaneous queue length at every node in the tandem queue is proportionally impacted by the perturbations of the scale parameters of either the inter arrival distribution or the mark distribution while the product of the average arrival rate and average mark is kept constant.

These results indicate that burstiness in network traffic leads directly to queuing and thus to packet losses in small-buffer networks.

III. PACING SYSTEMS IN NEXT-GENERATION NETWORKS

Based on the observation that burstiness leads to packet losses, we propose a packet pacing architecture for the next-generation Internet architecture (see Figure 1). In particular, we propose to use a novel pacing algorithm that paces traffic indiscriminately and opportunistically at any node in the network [2]. This approach differs from prior approaches where pacing is limited to the edge and uses more complex algorithms [8].

Figure 2 shows the effectiveness of our pacing approach as presented in prior work [2]. Traffic bursts traversing the multiple pacing nodes are smoothed out to nearly match constant bit-rate traffic. Using this pacing technology throughout access networks helps operate optical core networks more effectively. Our ongoing research focuses on making such pacing an integral part of the next-generation Internet architecture.




Fig. 1. Network Architecture with Pacing.




Fig. 2. Effectiveness of Pacing on Multiple Nodes. From our prior work [2].

REFERENCES

- [1] R. W. Brockett, W. Gong, and Y. Guo. Stochastic analysis for fluid queueing systems. In *Proc. of the 38th IEEE Conference on Decision and Control (CDC)*, pages 3077–3082, Phoenix, AZ, Dec. 1999.
- [2] Y. Cai, Y. S. Hanay, and T. Wolf. Practical packet pacing in small-buffer networks. In *Proc. of IEEE International Conference on Communications (ICC)*, Dresden, Germany, June 2009.
- [3] Y. Cai, Y. Liu, W. Gong, and T. Wolf. Impact of arrival burstiness to queue length: An infinitesimal perturbation analysis. Technical report, Department of Electrical and Computer Engineering, University of Massachusetts, Amherst, MA, 2009.
- [4] M. Enachescu, Y. Ganjali, A. Goel, N. McKeown, and T. Roughgarden. Routers with very small buffers. In *Proc. of the Twentyfifth Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2006)*, Barcelona, Spain, Apr. 2006.
- [5] Y. Huang, Y. Liu, W. Gong, and D. Towsley. Two-level stochastic fluid tandem queueing model for burst impact analysis. In *Proc. of IEEE Conferences on Decision and Control (CDC)*, pages 3042–3047, New Orleans, LA, Dec. 2007.
- [6] A. Lakshminantha, R. Srikant, and C. Beck. Impact of file arrivals and departures on buffer sizing in core routers. In *Proc. of the IEEE Conference on Computer Communications (INFOCOM)*, pages 86–90, Phoenix, AZ, Apr. 2008.
- [7] R. Langenhorst, M. Eiselt, W. Pieper, G. Grosskopf, R. Ludwig, L. Kuller, E. Dietrich, and H. G. Weber. Fiber loop optical buffer. *Journal of Lightwave Technology*, 14(3):324–335, Mar. 1996.
- [8] V. Sivaraman, H. Elgindy, D. Moreland, and D. Ostry. Packet pacing in short buffer optical packet switched networks. In *Proc. of the Twentyfifth Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2006)*, Barcelona, Spain, Apr. 2006.