Tuning Self-Similar Traffic to Improve Loss Performance in Small Buffer Routers

Yongfei Zang[1], Jinyao Yan[2]
[1]Information Engineering School, Communication University of China, 100024 Beijing, China
[2]Computer Engineering and Networks Lab, ETH Zurich, CH-8092 Zurich, Switzerland
Emails: {zangyongfei123@163.com, jyan@ee.ethz.ch}

Abstract—**The issue of router buffer sizing is an important research problem and is still open though researchers have debated this for many years. The research method can be classified into two kinds: one is based on queuing theory, the other uses TCP as model. From the point of TCP model, many researchers concluded that buffer size can be significantly reduced. It's desirable that the buffers are so small that fast memory technology and all-optical buffering can be used. But queuing model with self-similar incoming traffic suggested that extremely large buffers are needed to achieve acceptable packet loss rate. In this paper, we will first exam the performance of non-TCP and self-similar traffic with small router buffers, and then address the question how to improve the packet loss rate performance for self-similar traffic. Through a combination of simulation and analysis, we found that packet arrivals' burstiness has a significant influence on loss rate performance. We further point out a simple and effective approach, which smoothes the packet injections to the network, to improve the performance of small buffers at Internet core router for self-similar traffic.**

Keywords-buffer size; TCP; self-similarity; traffic smoothing.

I. INTRODUCTION AND MOTIVATION

All Internet routers need buffers to hold packets when TCP connections back off due to the congestion of the network and buffer the transient bursts that naturally occurred due to the characteristics of strong bursts and self-similarity of the Internet traffic, so the router buffers can keep high utilization of output link and reduce the packet loss rate. Meanwhile, buffers introduce queuing delay and jitter, and increase the router cost and power dissipation inevitably. The issue of sizing router buffer properly has generated much debate in the past few years. Different assumptions and objects have led to different conclusion. However, some of the recent reasearch all claimed that the router buffer can be significantly reduced in some cases. It is a good news for manufacturers and the development of all optical routers considering that recent advanced technology can at best hold a few dozen packets in an integrated opto-electronic chip [5].

The rule-of-thumb commonly used by router manufacturers today was proposed by Villamizar and Song in 1994 [1]. It claims that in order to make full utilization of the bottle link, a router needs a bandwidth-delay production buffering because of the sawtooth-like of TCP's congestion control algorithm, i.e., $B = RTT \times C$, where C is the capacity of the bottleneck link, B is the buffer of the bottleneck router and RTT is the average round trip time of a single and persistent TCP flow that attempts to saturate the link. The amount of buffers is in direct proportion to C and it will be a very large value considering that nowadays that backbone links commonly operate at Gbps magnitude.

In 2004, Appenzeller et al. from Stanford University challenged role-of-thumb. They concluded that when a large number of long TCP flows go through a bottleneck link in the core of the network, the buffer requirement decreases with the square root of the number of long TCP flows [2], i.e., $B = RTT \times \mathrm{C}/\sqrt{\mathrm{N}}$. According to their conclusion, a core router carrying 10000 long-lived flows needs only 1% of buffers proposed by rule-of-thumb.

In 2005, Enachescu et al. showed that if the TCP sources are paced or the access network is much slower than the backbone, and the maximum widow size has upper bound, O(logW) buffers (a few dozen packets) are sufficient if we are willing to sacrifice a small amount of the link capacity (say 10-20%), where W is the window size of each flow [3]. This result has made a useful exploration for the building of all optical routers with small integrated optical buffers.

In 2007, authors of [4] used a different metric and parameter to revisit the issue of router buffer sizing. Instead of only focusing on aggregate metrics such as link utilization and packet loss rate, they used average per-flow throughput to assess TCP performance. They claimed that the ratio of output/input capacity at a network link largely determines the required buffers. If the ratio is larger than 1, the loss rate drops exponentially with the buffer size and the optimal buffer size is extremely small (a few packets in practice). Otherwise, if the ratio is lower than one, the loss rate follows a power-law reduction with the buffer size and significantly large buffering is needed, especially with long-lived TCP flows which spend most of their time in congestion-avoidance.

The sizing router buffer formulas above is concluded based on closed-loop TCP congestion control model. Statistics shows that about 90% Internet traffic are TCP based while the rest traffic is transmitted over UDP and considered as open-loop traffic. Authors in [15] examined the dynamics of UDP and TCP interaction at a core router with few tens of packets of buffering and discovered the anomaly of UDP traffic's loss performance. From the view of queuing theory with specific incoming traffic model, the buffer sizing for open-loop traffic is quite different from

ISBN:978-1-61208-113-7

TCP. In the open-loop model, the router buffer is often modeled as a single queue with constant service rate (i.e., the capacity of the output link) and buffer size. The overflowing rate of the buffer depends on not only the buffer size and the capacity of the output link, but also the packet arrivals' patterns and the traffic's statistical features [6]. Various studies [7, 8, 9] have shown that network traffic exhibit ubiquitous properties of self-similarity. Analysis on video traffic, which is often transmitted over UDP, shows that self-similarity is also an inherent feature of VBR video traffic [10, 11]. The self-similar nature of network traffic has a significant influence on the queuing performance of router buffer. Authors of [12] pointed out that the packet loss rate in a network with self-similar traffic might be several orders of magnitude higher than that predicted by the traditionally used Markovian traffic models.

In this paper, we will exam how the burstiness of self-similar traffic affects the queuing performance in the condition of small router buffers, and propose methods to improve the performance of small buffers in Internet core router for self-similar traffic.

The rest of the paper is organized as follows. In Section II, we compare the loss performances of self-similar traffic with varied traffic burstiness with Poisson traffic. We study how the burstiness of data sources influences the queuing performance of router buffer. In Section III, real video traces from the Internet and CBR traffic are used to validate our finding. We summarize our work and point out directions for future work in Section IV.

II. THE BURSTINESS OF SELF-SIMILAR TRAFFIC AND PERFORMANCE

For self-similar traffic, bursts will exist across a range of scales and the positive correlations in traffic will adversely affect the QoS provided to network users [13]. Simply increasing the routers' buffer sizes will have marginal impact on the packet loss rate. The heavy tailed nature of the burst size distribution [11] implies that only extremely large buffers are effective in reducing packet loss rate [12]. The queuing delay introduced by large buffers will impact the transfer delay performance of delay-sensitive traffic such as streaming media.

To present how the traffic's self-similar influences router's queuing performance, we use NS2 simulator on the commonly used dumbbell topology to simulate self-similar traffic and Poisson traffic.

The aggregation of many On/Off sources with heavy-tailed ON periods exhibits Long-range Dependence (LRD) [14]. In our simulation, we aggregate many Pareto On/Off Traffic Generators in NS2 to generate self-similar traffic. We use Poisson traffic generator in NS2 to generate Poisson traffic. Because of Poisson process' additive property, we can use a single flow to represent the aggregation of many individual ones passing through the bottleneck link. UDP is used for both the self-similar and Poisson traffic. The capacity of the access links is 10Mbps, and the propagation delays on the access links uniformly distributed between [1, 25] ms. The capacity and propagation delay of the core link are 10Mbps and 50ms respectively. We employ FIFO queue with drop-tail queue management, which is commonly used in most router today. There are 100 On/Off source nodes each with the same configuration (burst_time_ 500ms, idle_time_ 500ms, rate_ 200Kbps, packetSize_ 200, shape_ 1.5). The mean rate during an ON-Off pair is 100Kbps. We set the Poisson rate to 10Mbps. So, in all simulations, the output link is lightly saturated. We examine the packet loss rates of self-similar traffic and Poisson traffic while increasing the buffer size at bottleneck link.

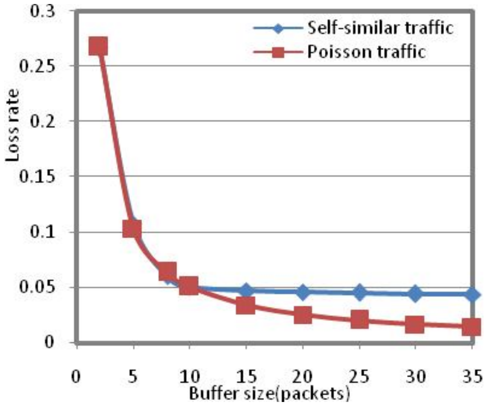

Figure 1. Loss rate for different buffer size

Figure 1 shows that in both cases, packets loss rate falls rapidly to a low value (5%, when the buffer size is 10 packets) as the buffer increases. After that, the self-similar traffic's loss rate curve drops very gently with the increase of the buffer size, while the loss rate of Poisson traffic falls faster than self-similar traffic. For self-similar traffic, increasing the buffer size simply will not get a good gain at loss rate.

One important reason for the self-similarity of network traffic is the statistical property of the size of the data blocks to be transferred, such as Web size, the size of Internet video's frames or GoPs. It's infeasible to change the statistical properties of the data to be transferred. From the observation on the loss rate performances' differences between the self-similar and the Poisson traffic, we consider the burstiness of packet arrivals leads to these differences. If we can make data sources to send data more smoothly, then what will happen?

In the next simulation, we will check the loss performance with different burstiness of self-similarity. We keep both the duration of On-Off pair (1000ms in our simulation) and the mean data rate at a constant value for self-similar traffic. By changing the length of On period ranging from 1ms to 1000ms, we adjust the burstiness for self-similar traffic. Let the mean data rate unchanged. Figure 2 shows the loss rate as a function of mean On time with buffer size 10 packets. We observe that the loss rate nearly falls exponentially with the increase of mean On time, which means that the data sources' burstiness has a significant influence on the loss performance. When the mean burst time is 1000ms, the Off time becomes to 0ms and each data source sends data with a low constant rate, namely with the lowest burstiness. Correspondingly, the loss rate achieves to the least value.

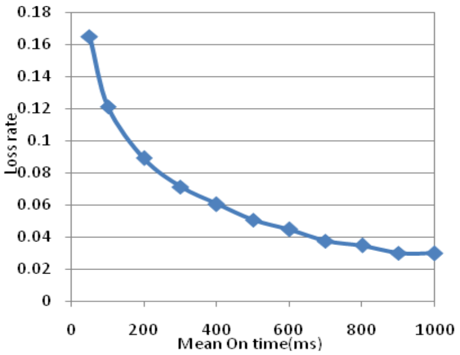

Figure 2. Loss rate for different mean On time

From the above, we can conclude that making the senders to space out packets evenly to weaken the burstiness of the data sources can improve the loss performance remarkably, even though the traffic to be transferred is self-similar. Therefore, we propose a simple and effective approach, which smoothes the packet injections to the network, to improve the performance of small buffers at Internet core router for self-similar traffic. For self-similar video streams, we send the frames with a constant rate continuously instead of sending the whole frame instantaneously as soon as we receive the frame from the application program. In the next section, we will validate our conclusion for co-existing TCP and UDP traffic with small buffer size routers.

III. SIMULATION VALIDATION

Let us consider a more realistic case of non-persistent TCP flows co-existed with UDP flows. We keep the fraction of UDP traffic fixed at about 8% in our simulations as that in the Internet. We use TCP traffic as background flows and focus on the loss performance of UDP traffic. Modified Harpoon system is used to generate closed-loop TCP flows. The size of TCP transfers follows Pareto distribution. After each download, an idle time which follows an exponential distribution with mean duration of 1 second follows until next TCP transfer starts [4]. Each TCP source can be seen as an On/Off model. The aggregation of many these TCP sources exhibit LRD property.

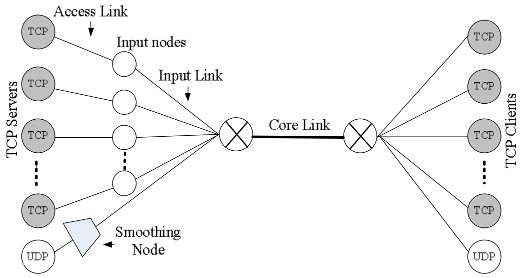

Figure 3. Simulation topology

Figure 3 shows our NS2 simulation topology. Our simulation setting has referred to [15], which focused on the anomalous loss performance of mixed real-time and TCP traffic. However, we simulate and compare the loss performance of the Internet video traffic and smoothed traffic over UDP. In [15], the TCP traffic was generated from persistent TCP flows while we use more realistic non-persistent TCP flows to generate it.

TABLE I. SIMULATION CONFIGURATION

Param Value Type		Buffer size (packets)	Link Capacity (Mbps)	Propogation Delay (ms)
Core Link		variable	10	50
TCP SACK	Access links	0	1	5,7,9…43
	Input links	100000	100	5
UDP	Input link	0	10, 0.256	10

Input link means the link that directly connects to the input core node.

TABLE I shows parts of simulation configuration. We employ FIFO queue with drop-tail queue management. There are 20 servers that are connected to 20 input nodes. Each of the 100 TCP users at client-side selects a server randomly to ecreate connections through the core link. The TCP transfer follows a Pareto distribution with mean 100KBytes and shape parameter 1.5. There are 3 UDP source nodes connecting to the input core node directly. The buffers of all the access links are too large to induce loss rate, so the output link is the single bottleneck. The simulation duration is 300s.The reported results ignore the first 20s of each simulation. TCP and UDP packet sizes were fixed at 1000Bytes and 200 Bytes respectively.

In what follows, we compare the loss performance of self-similar video traffic with that of smoothed video traffic. We insert a smoothing node between the UDP source nodes and the input core node with very large buffers (1 million packets) to ensure no dropped packet and 256 Kbps capacity of output link. The smoothing node buffers video trace packes and sends them smoothly to the input core node.

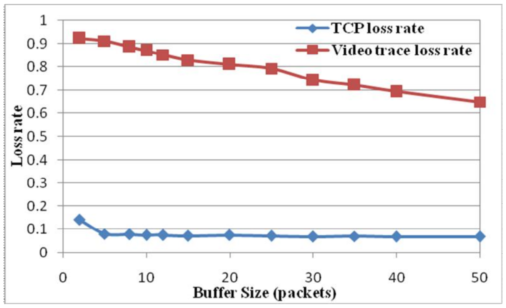

Figure 4. Video trace over UDP: Loss rate for different buffer size

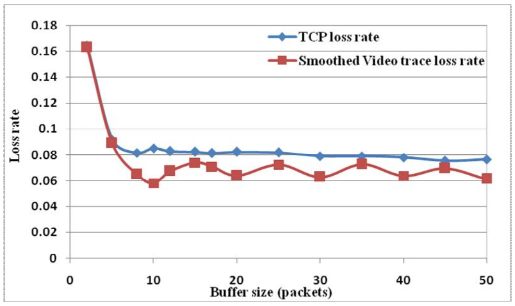

Figure 5. Smoothed Video trace over UDP: Loss rate for different buffer size

We use video traces of Jurassic Park I, Silence of The Lamb and Star Wars IV encoded by H.263 to generate UDP traffic [16]. Each video trace's mean rate is about 256 Kbps and the capacity of each UDP's input link is set to 10 Mbps. Previous study has shown that VBR video traffic is self-similar. So the UDP traffic generated from video traces is much bursty. We plot the bottleneck link utilization and loss rate of TCP and UDP as functions of the buffer size in Figure 4 and Figure 5.

We observe that loss rate of TCP are nearly identical. While there are vast differences between the UDP loss rate of Figure 4 and Figure 5. In Figure 4, the UDP loss rate drops linearly with the buffer size, while in Figure 5, the loss rate of UDP falls rapidly to a low value, then falls gently(but with high variance). For instance, the loss at 20 packets of buffering in Figure 4 is approximately 10 times higher than that of Figure 5.

We can explain with the findings in section II for the significant differences between the UDP loss rate curves in figures. In Figure 4, UDP traffic is generated from video traces. If the video frame to be transferred is larger than UDP packet size (200 Bytes in our simulation), the frame will be cut into a few of packets, then transferred simultaneously. But after smoothing, packets are sent in an approximately constant rate. So, UDP traffic generated from video traces is burstier than smoothed traffic though it originated from self-similar traffic. Therefore, our approach smoothing the packet injections to the network can improve the loss performance of small buffers at Internet core router for self-similar traffic.

IV. CONCLUSION AND FUTURE WORK

The study of sizing router buffer has generated much debate over the past few years. Researchers have questioned the commonly used rule-of-thumb which leads to a huge packet buffers in core routers today and have argued that small buffers at core routers are sufficient to meet acceptable performance. Various studies have shown that network traffic exhibit ubiquitous properties of self-similarity, from the point of queuing theory, extremely large buffer is needed. In this paper, we exploit how to improve the queuing performance of self-similar traffic at a bottleneck link router equipped with small buffers.

Through a combination of simulation and analysis, we found that there exists huge difference of loss performance between Poisson traffic and self-similar traffic due to the different bursty strength of packet arrivals. We can smooth packets injections at the edge of the network for self-similar traffic. Our realistic simulation mixed with TCP and UDP traffic shows that smoothed video traffic has a much better loss performance than VBR video streaming. We suggest that to adapt self-similar traffic to the small buffers at Internet core router, a simple and effective way to improve the queuing performance is smoothing the packet injections to the network.

As an important part of our future work, we will design the algorithm for our approach and implement the algorithm to test the performance in real network. After all, the "smoothing node" in our simulation was primary and used to make qualitative analysis. We will also revisit the issue of sizing router buffer with comprehensive consideration of TCP model and queuing model for self-similar traffic.

ACKNOWLEDGMENT

We would like to thank Yue Zhou and Botao Bai for their help in discussion and edit. Yongfei Zang is supported by NSFC under grant No. 60970127 and Key Project of Chinese Ministry of Education (No. 109029). Jinyao Yan is supported by Swiss National Science Foundation (No.200020_121753) and by Program for New Century Excellent Talents in Chinese University (NCET-09-0709).

REFERENCES

[1] C. Villamizar and C. Song, "High performance TCP in ANSNET," ACM Computer Communications Review, 24(5):45–60, 1994.

[2] G. Appenzeller, I. Keslassy, and N. McKeown, "Sizing router buffers," In Proc. of the SIGCOMM 2004. New York: ACM Press, 2004. 281−292.

[3] M. Enachescu, Y. Ganjali, A. Goel, N. McKeown, and T. Roughgarden, "Routers With Very Small Buffers," Proc. IEEE INFOCOM, Barcelona, Spain, Apr 2006.

[4] R. S. Prasad, C. Dovrolis, and M. Thottan, "Router Buffer Sizing Revisited: The Role of the Output/Input Capacity Ratio," In ACM CoNEXT, USA, 2007.

[5] H. Park, E. F. Burmeister, S. Bjorlin, and J. E. Bowers, "40-Gb/s optical buffer design and simulation," Proc. Numerical Simulation of Optoelectronic Devices (NUSOD), California, USA, Aug 2004.

[6] I. Norros, "A Storage Model with Self-similar Input," Queueing System, vol. 16, pp. 387-396, 1994

[7] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, "On the self-similar nature of Ethernet traffic," IEEE/ACM Trans. Networking, vol. 2, no. 1, pp. 1 - 15, 1994.

[8] V. Paxson, "Empirically-derived analytic models of wide-area TCP connections," IEEE/ACM Trans. Networking, vol. 2, pp. 316 - 336, 1994.

[9] M. Crovella and A. Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes," Proceedings of the 1996 ACM SIGMETRICS Conference, Philadelphia, PA, pp. 160-169, May 1996.

[10] J. Beran, R. Sherman, M. Taqqu, and W. Willinger, "Long-range Dependence in Variable Bit-Rate Video Traffic," IEEE Transactions on Communications, Volume 43, pp. 1566-1579, 1995.

[11] M. Garrett and W. Willinger, "Analysis, Modeling and Generation of Self-Similar VBR Video Traffic," Proceedings of ACM SIGCOMM '94, London, UK, pp. 269-280, August 1994.

[12] Y. Chen, Z. Deng, and C. Williamson, "A Model for Self-Similar Ethernet LAN Traffic: Design, Implementation, and Performance Implications," Proceedings of the 1995 Summer Computer Simulation Conference (SCSC'95), Ottawa, Ontario, pp. 831-837, July 1995.

[13] N. Duffield, J. Lewis, N. O'Connell, R. Russell, and F. Toomey, "Predicting Quality of Service for Traffic with Long Range Dependence," Proceedings of ICC'95, Seattle, WA, pp. 473-477, September 1995.

[14] W. Willinger, M. Taqqu, R. Sherman, and D. V. Wilson, "Self-Similarity Through High-Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level," In ACM Sigcomm, 1995.

[15] A. Vishwanath and V. Sivaraman, "Routers With Very Small Buffers: Anomalous Loss Performance for Mixed Real-Time and TCP Traffic," Proc. IEEE IWQoS, The Netherlands, Jun. 2008.

[16] http://trace.eas.asu.edu/TRACE/ltvt.html, November 8, 2010