

How Poisson is TCP Traffic at Short Time-Scales in a Small Buffer Core Network?

Arun Vishwanath[†], Vijay Sivaraman[†] and Diethelm Ostry[‡]

[†]School of EE&T, University of New South Wales, Sydney, Australia

Emails: {arunv@ee.unsw.edu.au, vijay@unsw.edu.au}

[‡]CSIRO ICT Centre, Sydney, Australia, Email: {Diet.Ostry@csiro.au}

Abstract—There is widespread debate regarding the nature of TCP traffic in today’s Internet - while some researchers have shown that it exhibits long-range dependent (LRD) properties, others argue that it can be modelled as a Poisson process due to the high degree of traffic aggregation that exists in the core. In this paper, we investigate the nature of TCP traffic as the Internet core moves towards an all-optical packet switched network with *very limited buffering* (few tens of KiloBytes) capability. In particular, we show that the bottleneck link buffers have a large influence on the aggregate TCP arrival process: large buffers can induce synchronisation amongst TCP flows, thus creating significant burstiness (equivalently LRD), but as buffers become smaller, the TCP aggregate can be well approximated as a Poisson process. Our work has major impact on the design and analysis of future high-speed optical packet switched networks with very small buffers.

I. INTRODUCTION

The publication of the seminal paper in [1] showed empirically that network traffic exhibits long-range dependent characteristics (i.e., bursty traffic over a wide range of time-scales). The heavy-tailed distribution of the file transfer sizes and the duration of active/inactive times of users are the two main reasons why network traffic is thought to be LRD. Since the publication of this result, there has been a keen interest in trying to understand how Internet traffic has evolved. Many have argued that due to the phenomenal growth in Internet traffic [2] (currently in the exabyte (10^{18}) range) and link speeds, there is high degree of statistical multiplexing and traffic aggregation in the core, resulting in the traffic being Poisson in sub-second time-scales [3]. There is further support from stochastic processes theory, which suggests that the superposition of a large number of *independent* flows also converges to a Poisson process [4], [5], [6].

Contrary to the above claims, there is still a large body of work showing that despite the growth in the Internet, the traffic is still bursty (LRD) and not smooth (Poisson). Through a combination of wavelet-based analysis and trace data obtained from a OC-48 Internet2 core link, the authors of [7] showed that Internet traffic is bursty in short time-scales because of the self-clocking nature of TCP and queuing at the bottleneck link. Also, as previously argued in [3], the aggregation of many TCP flows does not make the traffic smooth. Further light is shed in the recent paper [8] that uses trace data (spanning seven years) from a trans-Pacific backbone link to show that Internet traffic is indeed LRD and not Poisson.

It must be mentioned that the above measurement studies are pertinent to today’s core links that have of the order of delay-bandwidth amount of buffering, since ISPs typically run their networks with such large buffers [9]. Over the past few years, researchers have shown that core router buffers can be reduced significantly, to as small as a few tens of packets, and TCP would still achieve acceptable end-to-end performance [10]. Although the recent work advocating the use of small buffers has assumed that TCP traffic at a bottleneck link can be modelled as being Poisson [11], a claim predominantly supported by the theoretical superposition results, to the best of our knowledge, little has been done to justify this assumption using experiments and/or simulations.

The superposition results from stochastic processes is accurate only when the flows are *independent*. This is a critical assumption that easily holds for open-loop traffic (UDP for example). However, TCP is a closed-loop (i.e., feedback controlled) protocol, and the behaviour of a TCP flow is influenced by the presence of other TCP and/or UDP flows, as well as the buffer occupancy at the bottleneck link. In addition, TCP flows can synchronise their window dynamics behaviour (causing them to be in lock-step), which is a detrimental effect that breaks the independence assumption. As a result it is not very clear if the superposition results is directly applicable to TCP traffic, and as shown in [7], the aggregation of a large number TCP sources does not lend itself to being Poisson.

In this paper, we show that the nature of aggregate TCP traffic is strongly dependent on the size of the bottleneck link buffer. When buffers are small, the aggregate TCP arrival process is well approximated as being Poisson, but when buffers are large (of the order of delay-bandwidth product), TCP flows tend to synchronise, thus creating bursty LRD traffic. This seems to indicate that as we move to a core optical network in the future with very limited buffering, approximating the traffic as being Poisson is well justified.

II. SIMULATION SETUP

In the context of studying the burstiness of TCP traffic in a small buffer core network, we require a bottleneck link carrying several TCP flows. We use *ns-2* (version 2.33) simulator on the well-known dumbbell topology shown in Fig. 1, which directly captures the bottleneck link, and is commonly used to analyse the performance of various congestion control algorithms, including TCP. We consider




Fig. 1. *ns-2* simulation dumbbell topology

TCP-Reno flows corresponding to each source-destination pair $(s-tcp_i, d-tcp_i)$, $1 \leq i \leq 2000$, and employ FIFO queue with drop-tail queue management (at the bottleneck router r_0), which is commonly used in most routers today. Each TCP packet is of size 1 KB. The bottleneck link operates at 200 Mbps, while the access link speeds are uniformly distributed between [8, 12] Mbps (mean of 10 Mbps), to reflect a typical home user. The propagation delay on the bottleneck link is 25 ms with the access link delays being uniformly distributed in the interval [4, 10] ms, thus RTT varies between [58, 70] ms. All TCP flows begin their transmission in the interval [0, 10] sec and the simulation is run for 150 sec. Data in the interval [20, 150] sec is used in all our computations so as to capture the steady-state behaviour of the network.

III. IMPACT OF BUFFER SIZE ON TCP BURSTINESS

Our objective is to test if TCP packet arrivals to the bottleneck link buffer are near-Poisson or not. To do so we measure the burstiness of the arrival traffic at various time-scales. Burstiness at time-scale s is quantified by $\beta(s)$, the coefficient of variation (i.e., ratio of standard deviation to mean) of traffic volume measured over time intervals of size s . Log-log plots of $\beta(s)$ versus s are routinely used in the literature to depict traffic burstiness over various time-scales as an indicator of long-range dependency of traffic traces, and to show the influence of the Hurst parameter H . We will use estimates of H obtained from the burstiness plots as a measure of how close the traffic is to being short-range dependent (i.e., Poisson).

Consistent with prior work [7], we are also interested in short time-scales (between $2^{10} \mu s \approx 1$ ms to $2^{20} \mu s \approx 1$ s). Time-scales of sub-1 ms is not very relevant since the average transmission time of a TCP packet on the access link is 0.8 ms (because the TCP packet size is 1 KB and the mean access link rate is 10 Mbps). Further, for a given number of TCP flows, the variation at larger time-scales is due to TCP traffic arriving and departing the network, and not because of the congestion control algorithm of TCP. Therefore, we upper-bound the time-scale of interest to 1 s, which is a few multiples of the average round-trip time.

A. Long-lived TCP flows

We begin the study by considering the simple case of long-lived TCP flows, i.e., all TCP flows are persistent and have infinite data to send. Fig. 2 shows the traffic burstiness $\beta(s)$ as a function of the time-scale (s in μs) for six different values of bottleneck link buffer size and 2000 long-lived TCP flows.




Fig. 2. Burstiness for 2000 long-lived TCP flows with different buffer sizes

The bottom curve indicates burstiness when the bottleneck link has very limited buffering - 50 KB, equivalent to 50 packets, since each TCP packet is 1 KB in our simulation. The remaining five curves reflect increasing buffer sizes (ranging from 3 MB to 12.5 MB), with 6.25 MB of buffering corresponding to the delay-bandwidth product (average RTT of 250 ms times 200 Mbps link rate). The curves fairly overlap each other till time-scale of $2^{12} \mu s \approx 4$ ms. However, at time-scales beyond 4 ms, the burstiness curves for large buffers (approaching the delay-bandwidth product) flatten significantly, indicating onset of long-range dependence, with the Hurst parameter estimated at approximately 0.66 (for 5 MB buffers), 0.81 (for 6.25 MB buffers) and 0.88 (for 12.5 MB buffers) respectively. These results show that when buffers are large, TCP traffic is significantly bursty and cannot be modeled as being Poisson. On the contrary, with small buffers, the Poisson approximation seems reasonable.

B. Mix of short-lived and long-lived TCP flows

Though results from the previous section indicate that TCP packet arrivals are near-Poisson if we have a large number of long-lived TCP flows and small buffers, the reader may wonder if similar behaviour can be seen when many of the TCP flows are short-lived (or equivalently, the number of TCP flows is time varying). This is an important consideration since measurement based studies at the core of the Internet suggest that a large number ($> 80\%$) of TCP flows (e.g. HTTP requests) are short-lived and carry only a small volume of traffic, while a small number of TCP flows (e.g. FTP) are long-lived and carry a large volume of traffic.

We simulate such a scenario in *ns-2* using the same dumbbell topology shown in Fig. 1. There are a total of 2000 TCP flows with 1600 of them being short-lived and the rest long-lived. These 1600 users perform successive file transfers to their respective destination nodes, with the size of the file to be transferred following a Pareto distribution with mean 100 KB and shape parameter 1.2. These chosen values are representative of Internet traffic, and comparable with measurement data. After each file transfer, the user transitions into an idle or off state. The duration of the off period is exponentially distributed with mean 1 second. It can




Fig. 3. Burstiness for 1600 short-lived and 400 long-lived TCP flows with different buffer sizes

be noted that the above traffic generation mechanism, which is a combination of several ON-OFF sources with Pareto-distributed ON periods, is long-range dependent [1].

Fig. 3 shows the resulting burstiness curves for the same set of bottleneck link buffer sizes as before. It is interesting to note from the figure that, although a significant number of TCP flows have LRD properties, when fed into a small buffer (such as 50 KB, the bottom curve), the resulting TCP arrival process becomes near-Poisson (similar to the 50 KB curve in Fig. 2). However, we do note that as buffer size increases, the traffic tends towards LRD, as evidenced by the Hurst parameter values of 0.67 (for 4 MB), 0.79 (for 5 MB) and 0.84 (for 6.25 MB) respectively.

C. Burstiness explained via queue occupancy

We are now interested in understanding why the aggregate TCP arrivals in a small buffer network is near-Poisson while it is LRD with large buffers? To do so, we plot the queue (buffer) occupancy (for an arbitrarily chosen interval of 2 sec, consistent with our time-scale of interest) in Fig. 4 with 1600 short- and 400 long-lived TCP flows and two choices of buffer size - 50 KB (small buffers) and 6.25 MB (delay-bandwidth product large buffers). We note from the figure that for large buffers, the queue spends most of its time oscillating continuously between full and $\approx 75\%$ occupancy. The reason for this oscillation is that the TCP flows are not independent, in fact they are synchronised. When the queue is almost full, many flows experience packet drops and simultaneously back-off by reducing their windows in synchrony. This causes the packets already stored in the queue to be drained out, thus providing room for the flows to ramp-up (i.e., expand their windows) again. This periodicity, and hence synchronisation, is primarily because of large buffers as TCP sources are unable to perceive loss for relatively long periods of time (due to the large queueing delays experienced by the flows). On the contrary, the queue occupancy with small buffers looks like white noise, clearly showing that small buffers are instrumental in mitigating synchronisation effects. The TCP flows can thus be treated as being independent, and the superposition results can be invoked thereby rendering




Fig. 4. Normalised queue occupancy for small and large buffers with 1600 short- and 400 long-lived flows

credence to the fact that the aggregation of a large number of TCP flows at a bottleneck link with small buffers is near-Poisson. The synchronisation of TCP flows with large buffers is what causes them to be LRD.

IV. CONCLUSIONS AND FUTURE WORK

In this paper, we undertook a simulation study to determine the nature of TCP traffic at short time-scales when multiplexed at a bottleneck link equipped with very limited buffering. We showed that if we have a large number of TCP flows, the aggregate arrival process can be well approximated as being Poisson. However, with larger buffers, the dynamics can be significantly different, and the aggregate arrival process need not converge to the Poisson model.

We are currently exploring via analysis and simulation the effect of various system parameters - such as the number of TCP flows, core link speed to access link speed ratio, RTT, etc., on the burstiness of TCP, and hope to be able to report these results soon.

REFERENCES

- [1] W. Willinger, M. S. Taqqu, R. Sherman and D. V. Wilson, "Self-similarity through high-variability: Statistical analysis of ethernet LAN traffic at the source level," *Proc. ACM SIGCOMM*, USA, Aug/Sep 1995.
- [2] Cisco white paper, *Approaching the zettabyte era*, http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481374.pdf, Jun 2008.
- [3] T. Karagiannis et al., "A nonstationary Poisson view of Internet traffic," *Proc. IEEE INFOCOM*, Hong Kong, Mar 2004.
- [4] D. R. Cox and V. Isham, "Point processes," *Chapman and Hall*, London.
- [5] J. Cao and K. Ramanan, "A Poisson limit for buffer overflow probabilities," *Proc. IEEE INFOCOM*, USA, Jun 2002.
- [6] J. Cruise, "Poisson convergence, in large deviations, for the superposition of independent point processes," Submitted to *Annals of Operations Research* (2009), <http://www.statslab.cam.ac.uk/~rjrc3/PAPERS/poisson.html>
- [7] H. Jiang and C. Dovrolis, "Why is the Internet traffic bursty in short time scales?" *Proc. ACM SIGMETRICS*, Canada, Jun 2005.
- [8] P. Borgnat et al., "Seven years and one day: Sketching the evolution of Internet traffic," *Proc. IEEE INFOCOM*, Brazil, Apr 2009.
- [9] C. Villamizar and C. Song, "High performance TCP in ANSNet," *ACM SIGCOMM CCR*, vol. 24, no. 5, pp. 45-60, 1994.
- [10] M. Enachescu et al., "Routers with very small buffers," *Proc. IEEE INFOCOM*, Spain, Apr 2006.
- [11] A. Lakshminantha, R. Srikant and C. Beck, "Impact of file arrivals and departures on buffer sizing in core routers," *Proc. IEEE INFOCOM*, USA, Apr 2008.