

Accuracy of Probing Techniques in Estimating TCP Loss Rates*

Sushant Rewaskar, Jasleen Kaur, and F. Donelson Smith
Department of Computer Science
University of North Carolina at Chapel Hill
{rewaskar, jasleen, smithfd}@cs.unc.edu

Introduction

TCP is the dominant transport protocol in the Internet and arguably its performance fundamentally determines the performance of Internet applications. The issue of predicting the throughput of a TCP connection on a given network path has, consequently, received considerable attention in the research community. Most of existing efforts probe the network path for the round-trip time (RTT) and the loss rate, using Periodic or Poisson probe streams, and then feed these into analytical models that predict TCP throughput.

A crucial, but non-validated, assumption underlying these efforts is that the RTT and loss rates observed by the probes are the same as what TCP data segments would have experienced on the same path. While it has been recently shown that both the Poisson probing and Periodic probing observe similar loss rates and delays on network paths [1], not much is known about how these compare to the path properties experienced by TCP. Given the importance of TCP-performance prediction, it is of great interest to ask: *are the loss rate and delays obtained by Poisson/Periodic probing of a path good indicators of those that would be experienced by TCP?*

We conduct controlled experiments in an emulated lab setting to study the above question. Our preliminary results indicate that while TCP delays are estimated accurately by probing techniques, the TCP loss rate is underestimated. We also find that bursty probe streams are likely to give better estimates of TCP loss rates.

Methodology

We create a dumbbell topology with a 100 Mbps link connecting two routers and 20 machines connected to each of the routers (using an Ethernet switch). This setup ensures that the inter-router link is the bottleneck. We use Tmix [2]—which reproduces source-level transmission behavior derived from real Internet traffic—on the 40 end-machines to generate cross-traffic on the bottleneck link.

We study TCP and probing loss rates¹ by using a pair of machines (one connected to each of the routers) to generate Poisson probes, Periodic probes, and a TCP stream. Our setup allows us to control: (i) the probing rate, (ii) the RTT of the probes, (iii) the probe packet sizes, (iv) the offered load of cross traffic, and (v) the router queue length. For our preliminary investigation, we vary only the probing rate and the router queue length. The probe RTT is set to 500 ms, the probe packet sizes to 1460 B, and the offered load to 85%.

For each experiment, we start a bulk-transfer on TCP and let it stabilize for 10 minutes. We then start Poisson and Periodic prob-

*This research was supported in part by NSF CAREER grant CNS-0347814, a UNC Junior Faculty Development Award, and NSF RI grant EIA-0303590.

¹We also study end-to-end delays and find these to be similar for TCP and the probes in all of our experiments.

Queue size	Loss Rates (%)		
	Poisson	Periodic	TCP
14	4.43	4.15	3.54
48	1.16	1.23	1.08
140	0.09	0.09	0.07

Table 1: Two-way loss rate (measured at the sender)

Queue size	Direction	Loss Rates (%)			
		Poisson	Periodic	TCP	
				Actual	Effective
14	Forward	2.43	2.33	3.42	–
	Reverse	1.9	1.81	2.07	0.11
48	Forward	0.80	0.80	1.1	–
	Reverse	0.34	0.43	0.43	0.003
140	Forward	0.06	0.06	0.07	–
	Reverse	0.02	0.02	0.03	0.001

Table 2: One way loss rate

ing at the same average rate as the TCP send rate. All three streams then operate for 30 minutes. We monitor the actual packets dropped in either direction for any connection and use this to calculate the observed loss rate in each direction. We repeat each experiment several times. We also use a stateful analyzer [3] to identify the cause of TCP retransmissions.

Results

Two-way Loss Rates. We first compare the two-way loss rates observed by TCP, Periodic, and Poisson probes with three settings of TCP loss rate: 0.07%, 1%, and 3.5%. We vary the TCP loss rate by changing the queue size at the routers. Table 1 shows that the two-way loss rates observed at the sender is higher for Periodic/Poisson probing, as compared to TCP. The difference is more significant with smaller queue sizes (higher loss rates).

On analysis, we trace this behavior to the “cumulative” nature of acknowledgments (acks) used in TCP, which significantly reduces the impact of ack losses on TCP transmission. Specifically, a typical Periodic/Poisson probe stream infers a loss if it doesn’t receive an ack within an expected amount of time—thus, even if a probe is not lost and only the ack for it is lost, the sender concludes a loss. TCP, on the other hand, is much less affected by ack losses, since subsequent cumulative ack repeat the information in the lost ack.

This effect is quantified in Table 2, which lists the one-way loss rates observed in either direction by each of the three connections. For the reverse direction of TCP, we report the actual ack loss rate as well as the “effective” TCP loss rate (the rate at which TCP concludes a segment loss in the forward direction, despite the cumulative nature of acks). We see that while the actual loss rates for

Queue size	Probe Rate	Loss Rates (%)		
		Poisson	Periodic	TCP
48	0.5	0.82	0.79	1.03
48	1	0.80	0.80	1.1
48	2	0.86	0.82	1.04

Table 3: Loss rates with different probing rates

Queue size	Burst size	Loss Rates (%)		
		Poisson	Periodic	TCP
14	1	2.43	2.33	3.42
14	2	2.9	2.6	3.1
14	4	3.4	3.9	2.8
48	1	0.80	0.80	1.1
48	2	0.93	0.99	1.1
48	4	1.0	1.19	1.0

Table 4: Loss rates with different burst sizes

acks are similar in all three cases, the effective loss rate for TCP in the reverse direction is much lower than any of the probe streams.

Probing mechanisms often have control over only the source end-point of a path and rely on the availability of echo-response mechanisms at the receiver. The observations above suggest that when such probes are used to measure two-way loss rates, they are likely to overestimate the TCP loss rate.

One-way Loss Rates. We next consider the scenario when a probing mechanism has control over both the end-points of a path and is capable of sampling the one-way loss rates. Table 2 shows that the one-way forward loss rate for the Poisson/Periodic probes streams is lower than that for TCP. Recall that the probing rate is the same in all of the three connections. In order to see if a higher (or lower) probing rate helps the Probe streams achieve the same loss rate as TCP, we re-run the experiments with probes being sent at an average of 0.5 and 2 times the TCP sending rate. Table 3 lists our observations. We see that the probing rate has negligible effect on the loss rates for the probes.

A key difference between TCP and the probing connections is in the “structure” of their probe streams. While TCP transmission is inherently bursty, the Poisson probes are independently spaced (and the Periodic probes are uniform). To illustrate the impact of the probe stream structure on the observed loss process, Figure 1 plots the distribution of the inter-loss time intervals for the Poisson probes (“Poisson:1”) and TCP. We see that while for TCP, around 20% of the segment losses occur “back-to-back” (within 0.1 ms), this is hardly true for any losses in the Poisson stream.

In order to see if bursty probing will help improve the situation, we redesign the Poisson/Periodic probes to send 2 or 4 packets back-to-back, instead of a single packet—the rate of sending such groups of packets is reduced in order to maintain the same average probing rate as TCP. Table 4 shows the observed loss rates for different queue sizes and burst sizes. We see that as we increase the probing burst size, the observed loss rates for the probes increases. Figure 1 plots the inter-loss times for the bursty probes as well—we see that with a burst size of 2, the inter-loss times as well as the average loss rate for the probes closely matches those for TCP.

Conclusions. Unlike UDP-based Poisson or Periodic probes, TCP reduces its packet sending rate on experiencing packet losses. Intuition may suggest that since TCP samples the network less often when the latter is congested, it should experience fewer losses. Our preliminary results, however, defy this intuition and instead reveal two important insights for designing probing mechanisms for estimating TCP loss rates:

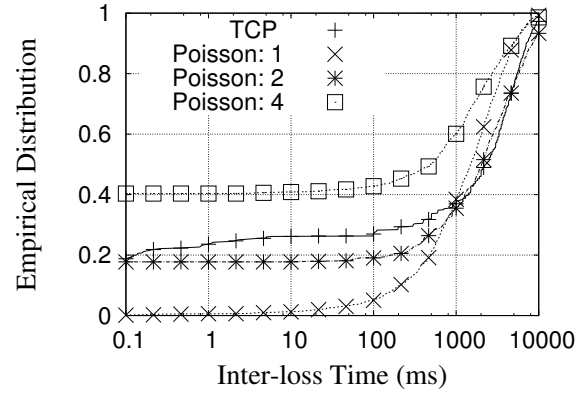


Figure 1: Distribution of inter-loss times

- Two-way loss rates measured at a node using popular probing techniques overestimate the loss rate seen by TCP, especially under heavy loss conditions. This is because of the insensitivity of TCP to the loss of cumulative acks.
- One-way TCP loss rates are underestimated by probing techniques, even when the probing rate is much higher than TCP. This is due to the bursty transmission behavior of TCP. Introducing burstiness in the structure of the probe streams helps it achieve loss rates similar to TCP.

Open Issues

Two key open issues guide our ongoing research on this topic:

- The loss process observed at a router depends on the queuing discipline used at the router—specifically, it is well-known that drop-tail queuing leads to bursty and correlated losses, while AQM schemes help in spreading out the losses over time and across connections. We are currently studying the impact of queuing discipline on the relative loss rates observed by TCP and Poisson/Periodic probing.
- According to the PASTA principle for stationary processes, the router loss probability observed using Poisson probing is the same as the time-average of the loss probability at the router queue. Our preliminary observations show, however, that this is not the case for our experimental setup. This may be because Internet traffic (which gets reproduced by Tmix) may not qualify as a stationary process due to its often-reported Long Range Dependence properties. We plan to study the relation between TCP, Poisson, and Router loss rate for such traffic.

1. REFERENCES

- [1] M.M. Bin, A. Dhamdhere, C. Dovrolis, and M. Ammar. Poisson Versus Periodic Path Probing (or, Does PASTA Matter?). In *USENIX/ACM Internet Measurements Conference*, Berkeley CA, October, 2005.
- [2] F. Hernandez-Campos, F. D. Smith, and K. Jeffay. Generating Realistic TCP Workloads. In *CMG2004 Conference*, pages 273–284, December 2004.
- [3] S. Rewaskar, J. Kaur, and F.D. Smith. A Passive State-Machine Approach for Accurate Analysis of TCP Out-of-Sequence Segments. *ACM Computer Communication Review*, 36(3):51–64, July 2006.