A HighSpeed TCP Study: Observations and Re-evaluation

Tuan Anh Trinh, Balázs Sonkoly and Sándor Molnár

Abstract—**This paper reports some results on the performance of HighSpeed TCP. Firstly, we find that the only publicly available simulation results in [6] are only *qualitatively* valid but not *quantitatively* and we show how to correct them. Secondly, our results show that the version of TCP used has *minimal* impact on the performance of HighSpeed TCP. Thirdly, we observe that as the number of flows *increases*, the advantage of HighSpeed TCP over Regular TCP *decreases*. Our results are supported by simulation using ns2.**

I. Introduction

The TCP has been widely used as a transport protocol in the Internet. Many applications such as HTTP and FTP are based on TCP. However, recent experience indicates that the congestion control of current TCP prevents it from fully utilize high-speed wide-area connection. Thus, network applications demanding high bandwidth are rarely able to take full advantage of high-speed networks and they are often not utilizing the available bandwidth. The major reason for under utilization is that the additive increase of traditional TCP is too slow and the multiplicative decrease is too harsh. In a steady-state environment, with a packet loss rate p, the current TCP's average congestion window is roughly $\frac{1.2}{\sqrt{p}}$ packets. This places a serious constraint on the congestion window that can be achieved by TCP in realistic environments.

HighSpeed TCP [1] (among others [2], [3], [4])is a recently proposed revision to the TCP congestion control mechanism. It is specifically designed for use in networks with high bandwidth-delay product. There exist very few studies on the performance implication of HighSpeed TCP so far. In this paper we report on our study into the performance issues of HighSpeed TCP.

This paper is organized as follows. Section 2 provides the background and discusses related work on HighSpeed TCP. Simulation environment and parameter setting are described in Section 3. Results are provided and discussed in Section 4. Finally, Section 5 concludes the paper.

II. Related Work

HighSpeed TCP was formally defined in RFC 3649 [1]. In that paper the basic elements of HighSpeed TCP, its motivation, and performance are described and discussed in detail. However, it should be mentioned that the RFC 3649 is based heavily on the simulation results provided in [6]. We consider this paper as the starting point of our work. We find that the parameters simulations used in the paper and in RFC 3649 are different. So the results in the paper are *qualitatively* true, but, we argue, not *quantitatively.* Specifically, in RFC 3649, Sally Floyd proposed a default set of parameter for a link with speed of *10Gbps.* E. Souza *et al* in [6] use the same set of parameters but for a link with speed *1Gbps.* To be more specific, we would mention that the High Window variable should have been set to 8300 for a 1Gbps link, not to 83000 as in [6]. In this paper we carried out simulations with *corrected* parameter settings, and *reevaluate* both qualitatively and quantitatively the performance of HS TCP in different environments. We also examine the performance of HSTCP not only with SACK TCP but also with different versions of TCP such as Reno and NewReno.

The authors are with the High Speed Networks Laboratory, Department of Telecommunications and Media Informatics, Budapest University of Technology and Economics, Hungary. E-mail: {trinh,sonkoly,molnar}@tmit.bme.hu

TABLE I
Link bandwidth vs. High Window

Bandwidth (Mbps)	High Window
1	8.3
10	83.3
100	833.3
1000	**8333.3**
10000	**83333.3**

III. Methodology

A. An Important Observation

As mentioned in the previous section, we suggest that the parameters used to investigate the performance of HighSpeed TCP in [6] are not fully in accordance with the parameters proposed in RFC 3649. Moreover, the claims in RFC 3649 are based on the simulation results in [6]. This *duality* is fine if they are in agreement with each other. However, it is not the case here. We address here the inconsistence in parameter setting. Specifically, let's consider Table I to see the dependence of link's bandwidth with proposed High Window, according to RFC 3649. As we can see in the table, for 1 Gbps link, the High Window *should have been chosen* to be 8333.3, instead of 83333.3 (or 83000, to round up) as presented in [6].

B. Simulation

We use ns2 [5] to conduct our experiments. We consider a general *dumbbell* network topology with a single bottleneck link and with multiple sender as well as multiple receivers, as shown in Figure 1. The queue

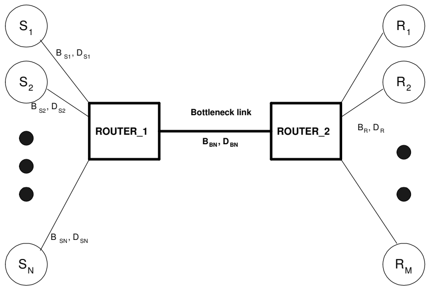

Fig. 1. Network Topology

management used at the routers was RED. We also have the ECN bit tuned, so instead of drop rate, we will mainly discuss the congestion event rate as observed by the number of packet marked. Not only SACK TCP but Reno TCP and NewReno TCP are also evaluated. We call them REGTCP. The HighSpeed version of TCP is called HSTCP.

We investigate the performance of HSTCP in different scenarios. First, we consider an isolated case to compare the dynamics of the congestion window process of both HSTCP and REGTCP. A number of performance metrics like link utilization, fairness, congestion event rate is reevaluated and discussed with different scenarios (ideal case and lossy links).

IV. Results

A. Isolated Flows

In this experiment, we investigate the dynamics of the congestion window processes of both REGTCP and HighSpeed TCP in *isolated* environment. The experiment run only one time, without external interference. As shown in Figure 2, a HighSpeed TCP flow

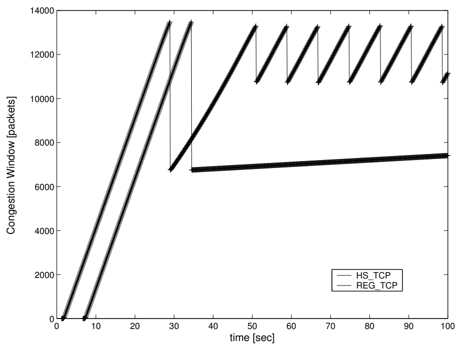

Fig. 2. Dynamics of the congestion windows

increases more quickly than a REGTCP. As a result, it takes a longer time for the REGTCP to reach the bandwidth limit (fully utilizes) in congestion avoidance. We also observe that, HSTCP *oscillates* with a very short period. This may not be desirable for a congestion control scheme.

Next, we examine the link utilization of isolated flows with different versions of TCP.

In Figure 3, the link utilization of different versions of TCP is illustrated. We can observe that in isolated condition, HSTCP has *clear* advantage over

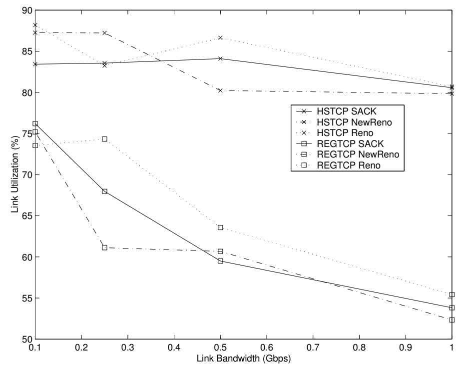

Fig. 3. Isolated flows: link utilization

REGTCP, at least in terms of link utilization, regardless of the version of TCP being considered.

B. Ideal Condition

In this experiment, we consider the behavior of a variable number of REGTCP and HSTCP flows, when there was no external interference, except the background traffic. We not only investigate the dynamics of the congestion window process but also other important performance metrics like link utilization and fairness. The results in different link speeds are compared. Figure 4 shows the link utilization of both

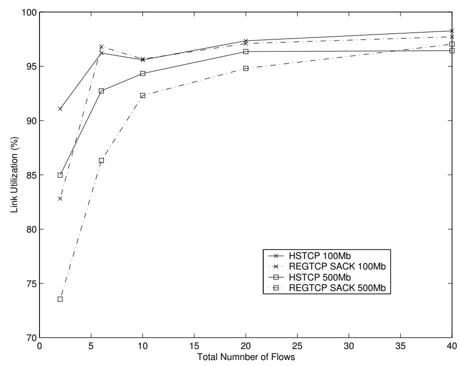

Fig. 4. Link utilization: Ideal condition

REGTCP and HSTCP under ideal condition. The number of flows investigated are varied between 2 and 40. We have a number of important observation here. First and foremost, we observe that HSTCP indeed has better link utilization than REGTCP, *as long as* the number of flows is small. As the number of flows *increases*, the advantage of HSTCP over REGTCP *decreases*. The second observation is that with current parameter setting as proposed in RFC 3649, surprisingly, both REGTCP and HSTCP seem to better utilize *lower* links, at least in the region of 100 Mbps to 500 Mbps. Now, let's consider the link utilization of other versions of TCP under ideal condition.

As shown in Figure 5, the advantage of HSTCP over REGTCP remains true not only with SACK TCP but also with Reno and NewReno.

Next, we consider a mixture of both REGTCP and HSTCP competing for bandwidth in ideal condition,

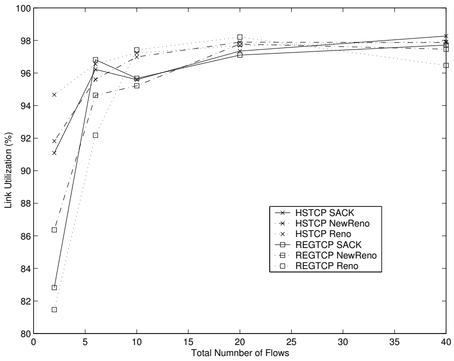

Fig. 5. Link utilization: Ideal condition with versions

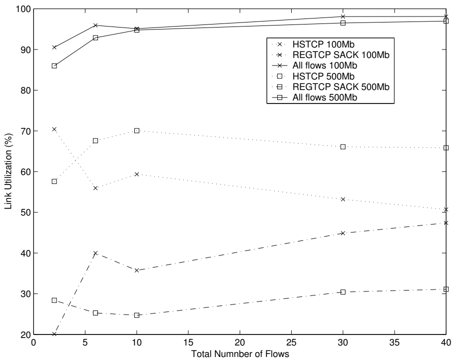

Fig. 6. Link utilization: Ideal condition (mixture)

with different link speeds. As shown in Figure 6, HSTCP has better utilization than REGTCP and in the condition of mixture of flows, this advantage *increases* as the link speed *increases*.

What we have been discussed so far is only about link utilization. Maintaining fairness among the connections in the network is an essential feature. So, new solutions *must* coexist nicely with existing solutions, or only interfere when the existing protocols are unable to use link capacity well.

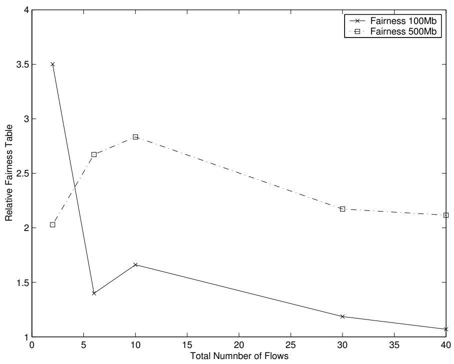

Fig. 7. Relative fairness: Ideal condition

Figure 7 shows the relative fairness between HSTCP and REGTCP (SACK TCP). We observe that is the number of flows is small, then HSTCP *unfairly steals* the bandwidth from REGTCP. This condition is changed as the number of flows increases (the relative fairness tends to 1 (fair)).

Finally, let us consider the congestion event rate with HSTCP and REGTCP in ideal condition. We say "congestion event rate" instead of "loss rate" because we use RED with ECN. This mechanism *marks* the packet instead of dropping it. A marked packet indicates a congestion event. Figure 8 shows the relative fairness of HSTCP and REGTCP. As HSTCP is more aggressive than REGTCP, we observe that the congestion event rate is higher with HSTCP than REGTCP. It is understandable since HSTCP achieves more bandwidth, and under RED, it will receive, more or less, proportional marked packets.

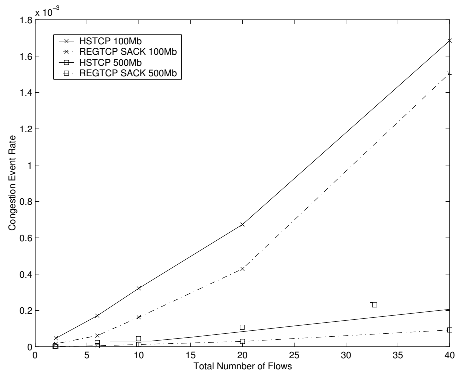

Fig. 8. Congestion event rate: Ideal condition

V. Conclusion

We have presented a HighSpeed TCP study based on our observation of the parameter setting in previous research. Results and claims in previous research are reevaluated both qualitatively and quantitatively. We find that the results published in previous research are in agreement with ours qualitatively, but not quantitatively. The results show that HighSpeed TCP indeed improve the performance of traditional TCP in high speed links environment. Analytically evaluating HighSpeed TCP is a subtle and difficult task, which is left as our future work.

References

[1] Sally Floyd, *HighSpeed TCP for Large Congestion Window,* RFC 3649, December 2003
[2] Cheng Jin, David X. Wei and Steven H. Low *FAST TCP: motivation, architecture, algorithms, performance,* IEEE Infocom, March 2004
[3] D. Katabi, M. Handley, and C. Rohrs, *Internet congestion control for high bandwidth-delay product networks* ACM Sigcomm 2002, August 2002.
[4] Tom Kelly, *Scalable TCP: Improving Performance in Highspeed Wide Area Networks,* ACM SIGCOMM Computer Communication Review, Vol. 33, Issue 2, pp. 83-91, April 2003.
[5] Network Simulator (ns2), *http://www.isi.edu/nsnam/ns*
[6] E. Souza and D. Agarwal *A HighSpeed TCP Study: Characteristics and Deployment Issues* Lawrence Berkeley National Lab's Technical Report, Berkeley, USA, 2003.