Available online at www.sciencedirect.com

ScienceDirect

Computer Networks 51 (2007) 1763–1776

Computer Networks

www.elsevier.com/locate/comnet

Can high-speed networks survive with DropTail queues management? ☆

Shan Chen *, Brahim Bensaou

Computer Science and Engineering Department, The Hong Kong University of Science and Technology, Hong Kong, China

Available online 28 November 2006

Abstract

It has been observed that TCP connections that go through multiple congested links (MCL) have a smaller transmission rate than the other connections. Such TCP behavior is a result of two components (i) the cumulative packet losses that a flow experiences at each router along its path; (ii) the longer round trip times (RTTs) suffered by such flows due to non-negligible queueing delays at congested routers. This double "bias" against connections with MCLs has been shown to approximate the so-called minimum potential delay fairness principle in the current Internet. Despite the recent proliferation of new congestion control proposals for TCP in high-speed networks, it is still unclear what kind of fairness principle could be achieved with such newly proposed congestion control protocols in high-speed networks with large-delays. Studies already show that some high-speed TCP variants may cause surprisingly severe RTT unfairness in high-speed networks with DropTail routers.

This paper studies the problem of unfairness in high-speed networks with some well-known high-speed TCP variants in presence of multiple congested links and highlights the severity of such unfairness when DropTail queue management is adopted.

Through a simple synchronized loss model analysis, we show how synchronized losses with DropTail in high-speed networks could lead to severe RTT unfairness and drop probability (DP) unfairness; while random marking AQM schemes, which break the packet loss synchrony mitigate such unfairness dramatically by ensuring that the packet loss probability of a flow is the sum of the loss probabilities on the congested routers it crosses.

Extensive simulations are carried out and the results support our findings.

Keywords: High-speed network; TCP congestion control; Multiple congested links; Fairness; DropTail; Active queue management

1. Introduction

The TCP congestion control algorithm and DropTail queue management are prevalently used in the current Internet and have been remarkably successful for many years. As the Internet evolves, wide-area high-speed networks are emerging in which the available bandwidth to a flow is significantly

☆ This work was supported in part under grant CERG HKUST6166/04E.

* Corresponding author.
E-mail addresses: chenshan@cse.ust.hk (S. Chen), brahim@cse.ust.hk (B. Bensaou).

1389-1286/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2006.11.004

higher than it is currently. Unfortunately, it has been clearly shown that the current TCP performs extremely poorly in networks with high bandwidth-delay product. The classical Additive-Increase Multiplicative-Decrease (AIMD) algorithm behind TCP congestion control, which increases the congestion window (cwnd) by one segment in each round trip time (RTT) and reduces it by half for each congestion loss, prevents TCP from effectively utilizing the network bandwidth over high-speed connections with long RTTs.

A great deal of research effort has been put into developing new transport protocols to address this limitation in TCP. Several newly proposed protocols have attracted a lot of attention, including XCP [1], FAST TCP [2], HighSpeed-TCP (HS-TCP) [3], Scalable-TCP (S-TCP) [4], BIC-TCP [5], H-TCP [6], and so on. Among these, HS-TCP, S-TCP, BIC-TCP, and H-TCP are loss-based, self-clocked, end-to-end high-speed TCP variations, which are potential candidates for a safer incremental deployment [7]. HS-TCP and S-TCP increase the cwnd much faster when the cwnd grows larger and maintain the same increase rate as the standard TCP when the cwnd size is below a certain threshold. Such bandwidth probing behavior ensures efficiency in high-speed networks and TCP friendliness in normal-speed networks where the packet loss probability is high or moderate. However, as pointed out by Xu et al. [5], HS-TCP and S-TCP suffer great RTT unfairness in high-speed networks, which means competing flows with different RTTs may finally receive dramatically unfair bandwidth shares. Further, BIC-TCP is proposed to mitigate the RTT unfairness. Another newly proposed congestion control algorithm, namely H-TCP, shows very little or no RTT unfairness and even a fast convergence speed in high-speed networks.

There has recently been some work done on evaluating the performance of these newly proposed protocols (e.g., [8,9]); unfortunately, most of the experiments and simulations carried out, focus on the "standard" dumbbell topology where only one single link is congested. Under this single bottleneck setting, we believe many characteristic behaviors of these protocols in real networks are not captured; in particular, the dramatic unfairness that appears when flows cross-multiple congested links.

Floyd and Jacobson [10] discussed the MCL unfairness of TCP/IP networks, a bias against flows passing through multiple congested routers, and also advocated that a fairness principle other than the max-min should be applied in TCP/IP networks. Later, Kelly et al. [11] argued in favor of the so-called proportional fairness principle, in which the utilization of network resources is maximized. In fact, the pros and cons of max-min fairness and proportional fairness are quite obvious. The debate on which fairness principle is preferable for TCP/IP networks is beyond the scope of this paper. Indeed which fairness principle is better has to be determined by the objectives determined by the users and the network providers. From a network stand point, maximizing network resource utilization is one of the major goals, therefore resource-centric fairness principles such as the proportional fairness are preferred. However, high resource utilization is only the penultimate goal as such resources are deployed to serve end-users. And from the user's perspective, max-min fairness is considered to be the best. Furthermore, as shown in [12] more end-user fairness is not always synonymous with throughput degradation. TCP on the current Internet is thought to achieve the so-called minimum potential delay fairness [13] principle, which is somewhere between the max-min and proportional fairness in terms of end-user fairness and network resource utilization. In our definition of fairness, we consider a protocol for high-speed networks to be fair if it achieves a similar fairness as AIMD in normal-speed networks.

In this paper we study the MCL unfairness in future high-speed networks with some of the newly proposed congestion control protocols. We show that:

- With DropTail, the MCL unfairness in high-speed networks could be far worse than it is in normal-speed networks.
- Furthermore, while the MCL unfairness in normal-speed networks is mainly due to the difference in RTTs between flows, in high-speed networks, the MCL unfairness is far worse than expected, even when the flows have similar RTTs.
- Finally, this result is mainly due to the prevailing synchronized losses in high-speed networks.
- With random marking AQM schemes, the RTT unfairness and the Drop Probability (DP) unfairness could be mitigated to a level similar to that in the current Internet.

The rest of the paper is organized as follows: Section 2 demonstrates the MCL unfairness issue by simple simulations. A synchronized loss model is