TCP Congestion Control and Asymmetric Networks

Georgi V. Hristov[1] and Mihail P. Iliev[2]

Abstract – **As the number of Internet users grows, the demand for faster networks becomes larger. To meet these demands, some new network technologies are emerging. Some of these technologies provide asymmetric bandwidth for uplink (from the client to the server) and downlink (from the server to the client). These technologies are considered to be suitable for the Internet access. It is because the user's access to the Internet is essentially asymmetric. The user usually retrieves the information from the Internet through WWW (World Wide Web) service or file transfer service. Both HTTP (Hyper Text Transfer Protocol) for WWW service and FTP (File Transfer Protocol) for file transfer service use TCP - the most popular transport layer protocol in the Internet. The problem is that TCP has not been designed for asymmetric networks.**

Keywords – TCP, asymmetric network, simulation, emulation

I. INTRODUCTION

There are various types of asymmetry, which can be related to [1, 2, 6]:

Bandwidth: The bandwidth is different in the two directions (downstream and upstream).

Latency: The latency is different in the two directions. Normally caused by data being sent through two different transmission media. The latency in this case is viewed as the delay of physical medium, i.e., the combination of propagation and processing delay.

Quality: The error rate is different in the two directions, also normally caused by two different transmission media.

We describe some known asymmetric technologies to exemplify different types of asymmetry. A given asymmetric technology is often affected by a combination of different types of asymmetry.

Cable TV networks have wide downstream bandwidth, but often there is no return channel. Some solutions use ISDN modems as return channel, which introduces a significant bandwidth asymmetry. Others use a limited return channel via cable net, which is shared with everyone else.

The technology clearly exhibits bandwidth asymmetry. The technology also exhibits a degree of latency asymmetry, which is dependent on the return channel. If ISDN is used, the two medias differ and thus the latency will probably also differ. The latency of the return channel via cable net is often also affected, due to guard times between transmissions and contention intervals on the shared media.

As the name implies, the **ADSL technology** is based on asymmetric lines, where bandwidth asymmetry is the primary factor. In ADSL the downstream capacity is larger than the upstream capacity.

There is a vast number of **satellite links** with different kinds of feedback channels. Due to the many different types of feedback channels, satellite links can exhibit all types of asymmetry. The two major asymmetries for satellite links are bandwidth and latency asymmetry. The physical distance to the satellite often introduces large propagation latency (roughly 300 ms one-way).

Some **wireless/radio** technologies suffer from media access asymmetry. Especially technologies using a central base station and several mobile clients as this requires a RTC/CTS (Ready To Send/Clear To Send) protocol, with a large turn-around time caused by the radio. The drop rate is also significantly higher in wireless networks, thus some kind of quality asymmetry is also likely to occur.

II. THEORY OF ASYMMETRY AND EFFECT ON TCP

The network model which we use in the analysis and simulation is depicted in Figure 1. The model consists of a server, a client, and two links; downlink from the server to the client and uplink from the client to the server. The uplink and the downlink have asymmetric bandwidth, denoted as μ_f [data segments/s] and μ_r [ACK segments/s], respectively. Note that μ_f and μ_r are represented by units of data/ACK segment, respectively. The buffer sizes are denoted as B_f [data segments] and B_r [ACK segments], and the propagation delays between the server and client are τ_f [s] and τ_r [s]. Whether the upstream capacity is sufficient, depends on the ratio between the downstream and upstream bandwidth divided by the ratio between the size of data packets and ACK packets. For one-way transfers, this is defined as the normalized bandwidth ratio, k:

$$k = \frac{Downstream / Upstream}{Data_{size} / ACK_{size}} = \frac{\mu_f}{\mu_r} \qquad (1)$$

[1]Georgi V. Hristov is with the Department of Communication Technique and Technologies, 8 Studentska Str., 7017 Rousse, Bulgaria, E-mail: ghristov@mbox.contact.bg

[2] Mihail P. Iliev is with the Department of Communication Technique and Technologies, 8 Studentska Str., 7017 Rousse, Bulgaria, E-mail: miliev@ecs.ru.acad.bg

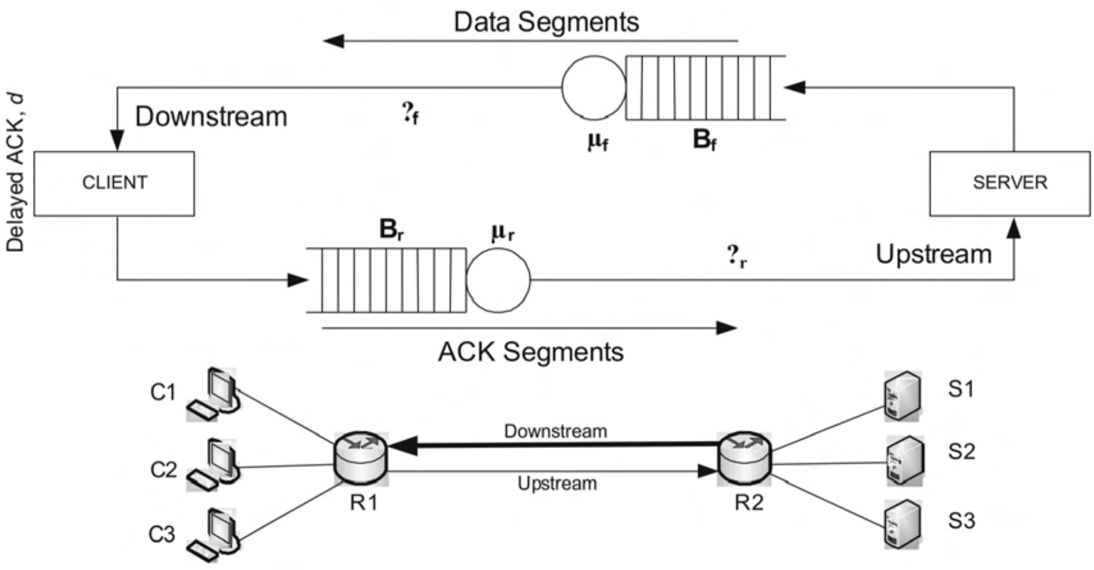

Fig. 1. Network Model

The value ***k*** indicates, how many data packets need to be acknowledged per ACK packet, to avoid saturation of the upstream by ACK packets. In other words, if there is more than one ACK packet for every k data packets, the upstream will get saturated (with ACK packets) before the downstream.

The TCP connection's delayed ACK factor, ***d*** defines the number of TCP data segments acknowledged by an ACK packet. Therefore, the ***k*** value must be evaluated in relation to the ***d*** factor. We can conclude if ***k*** > ***d*** the upstream will be saturated by ACK packets. The delayed ACK algorithm specifies that TCP should send an ACK packet (at least) every second data packet. This means that 2 is an upper bound on ***d***. Accordingly , it is a minimum requirement that ***k*** < ***d*** < 2 [3].

The above scenario only holds for one-way transfers, but most real-life traffic is bidirectional. The upstream link might be used for data packets, while at the same time delivering ACK packets. The competing reverse traffic consumes a part of the upstream capacity, effectively increasing the degree of bandwidth asymmetry.

Variable latency affects the smoothness of a TCP data flow. TCP measures the Round- Trip Time (RTT) on the path to estimate the path Retransmission Time Out (RTO) (calculated from a smoothed RTT estimate and a linear deviation). In the event of (multiple) packet loss, the RTO needs to be accurate in order to respond in time to avoid unnecessary idle periods.

The line quality and packet losses generally affect TCP and is not only an asymmetric problem. Loss of ACK packets is less significant than loss of data packets, as ACK packets are cumulative (ACK'ing all outstanding data), which generally results in stretched ACKs.

Media access asymmetry can be manifested in several ways (which we will not describe in detail). It is generally defined in literature an uneven access to a shared medium by a distributed set of nodes.

III. IMPROVING TCP PERFORMANCE OVER ASYMMETRIC NETWORKS

Several researchers have studied TCP performance on asymmetric networks [1, 2, 4, 6]. To solve the problems caused by bandwidth asymmetry, authors proposed such strategies as ACK Congestion Control (ACC), ACK Filtering (AF), Sender Adaptation (SA), and ACK Reconstruction (AR) to reduce the congestion in the upstream direction and improve TCP performance.

ACK filtering (AF) is a TCP-aware link-layer technique that reduces the number of TCP ACKs sent on the reverse channel. The challenge is to ensure that the sender does not stall waiting for ACKs, which can happen if ACKs are removed indiscriminately on the reverse path. AF removes only certain ACKs without starving the sender by taking advantage of the fact that TCP ACKs are cumulative. As far as the sender's error control mechanism is concerned, the information contained in an ACK with a later sequence number subsumes the information contained in any earlier ACK.

ACK congestion control (ACC) is an alternative to ACK filtering that operates end-to-end rather than at the upstream bottleneck router. The key idea in ACC is to extend congestion control to TCP ACKs, since they do make non-negligible demands on resources at the bandwidth-constrained upstream link. ACKs occupy slots in the reverse channel buffer, which capacity is often limited to a certain number of packets (rather than bytes).

ACC and AF alleviate the problem of congestion on the reverse bottleneck link by decreasing the frequency of ACKs, with each ACK potentially acknowledging several data packets. This can cause problems such as sender burstiness and a slowdown in congestion window growth. Sender adaptation is an end-to-end technique for alleviating this problem. A bound is placed on the maximum number of packets the sender can transmit back-to-back, even if the window allows the transmission of more data. If necessary, more bursts of data are scheduled for later points in time computed based on the connection's data rate. The data rate is estimated as the ratio cwnd/srtt, where cwnd is the TCP congestion window size and srtt is the smoothed RTT estimate. Thus, large bursts of data get broken up into smaller bursts spread out over time.

ACK reconstruction is a technique to reconstruct the ACK stream after it has traversed the reverse direction bottleneck link. AR is a local technique designed to prevent the reduced ACK frequency from adversely affecting the performance of standard TCP sender implementations. This enables us to use schemes such as ACK filtering or ACK congestion control without requiring TCP senders to be modified to perform sender adaptation. This solution can be easily deployed by Internet Service Providers (ISPs) of asymmetric access technologies in conjunction with AF to achieve good performance.

IV. ANALYSIS AND SIMULATION RESULTS

In this section, we use simulation and emulation to demonstrate the performance of TCP in asymmetric networks.

We expect bidirectional traffic to have a negative impact on the throughput achieved by the downstream transfer as the reverse ACK traffic is competing with the upstream transfer, which effectively increases the $\boldsymbol{k}$ value.

To illustrate and verify that the downstream TCP traffic is affected by upstream traffic a simple test is performed with a single continuous downstream transfer and short upstream transfers. The downstream transfer is allowed to run alone after which an upstream transfer is started. The results is shown in Figure 2.

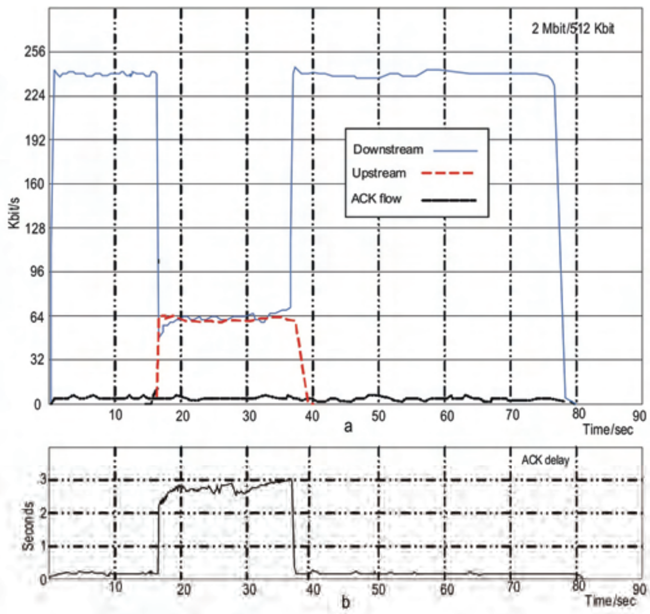

Fig. 2. Illustrating the downstream throughput problem on asymmetric link (a) and ACK delay (b)

In this item, we consider TCP level performance. Figure 3 shows the throughput of TCP as a function of the asymmetry factor of the network, $\boldsymbol{k = \mu_f / \mu_r}$. We change $\boldsymbol{\mu_f}$ while fixing $\boldsymbol{\mu_r}$.to determine $\boldsymbol{k}$. In the figure, we also plot emulation results to confirm accuracy of our simulation.

We can see from this figure that when the asymmetry factor $\boldsymbol{k}$ is small, the throughput increases in proportion to $\boldsymbol{k}$. Then, it is suddenly decreased and kept constant when the nominal value of $\boldsymbol{k}$ is too big. That is, if $\boldsymbol{k}$ is smaller than the downlink buffer size $\boldsymbol{B_f}$, segment loss occurs only when the window size exceeds the sum of bandwidth-delay product of the connection and the uplink and downlink buffer sizes. In other words, the throughput increases in proportion to the downlink bandwidth, $\boldsymbol{\mu_f}$. If $\boldsymbol{k}$ is larger than $\boldsymbol{B_f}$, on the other hand, segment loss takes place when the size of the burst generated by the server exceeds the downlink buffer size. It is independent of $\boldsymbol{\mu_f}$, and therefore the throughput is kept constant.

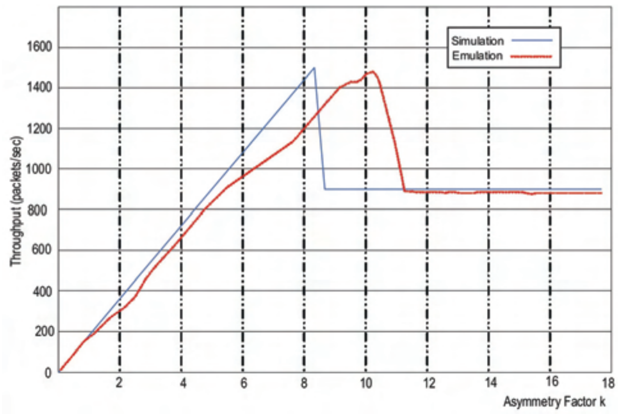

Fig 3. Throughput vs. asymmetry factor $\boldsymbol{k}$

We now discuss the results of the second set of experiments, designed to study the effects of performance, and to investigate how Reno, AF and ACC alter the inter-ACK spacing and how SA and AR prevent burst transmissions.

AF/AR and AF/SA perform the best, achieving throughputs between 15% and 21% better than Reno TCP. The degree of burstiness reduces significantly and the reverse router queue is no longer perpetually full. This is shown in Figure 4, which depicts the time-evolution of congestion windows for the different schemes. It is clear from these results that reducing the frequency of ACKs alone is not sufficient—AF alone performs even worse than Reno and techniques like SA or AR are essential to achieve good performance.

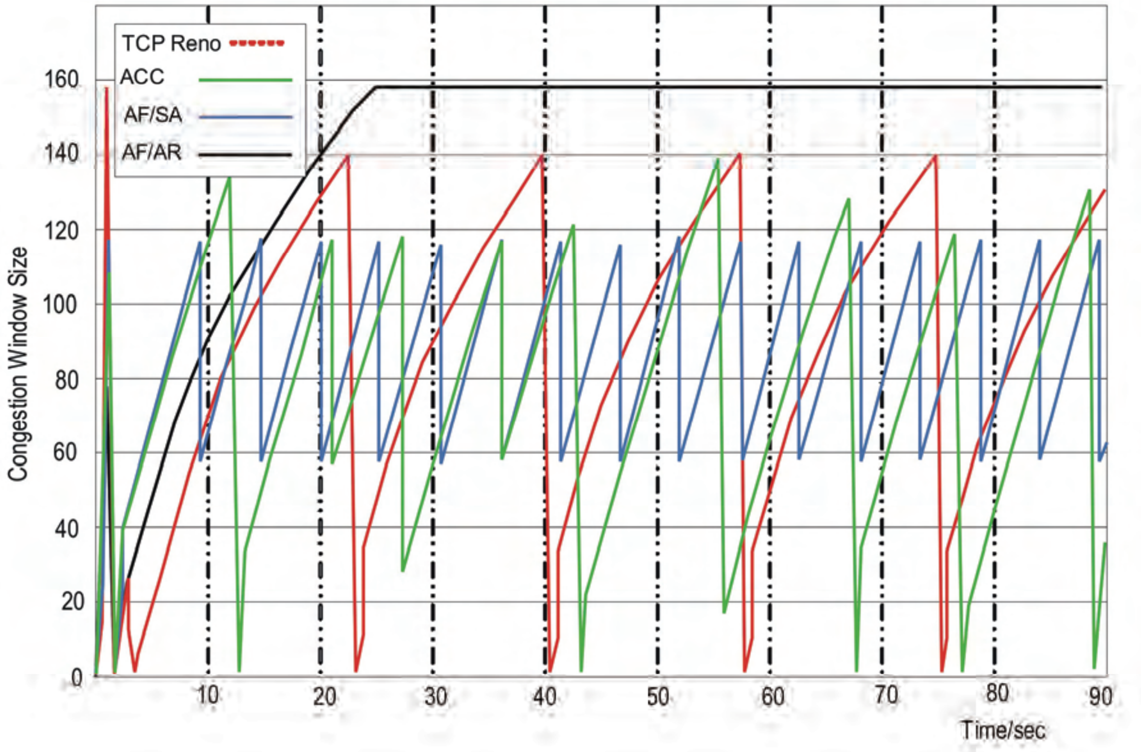

Fig. 4. Congestion Window evolution for the different scheme

V. Conclusion

In asymmetric networks, the slower upstream channel is the primary bottleneck for the TCP throughput over the faster downstream channel. The amount of throughput degradation increases with the normalized asymmetry factor ***k***. For unidirectional data transfer, the *delayed TCP ACK* mechanism helps in reducing the congestion over the slower upstream channel resulting in better TCP throughput over the forward channel. Bi-directional or two-way data transfer further increase the asymmetry problem. AF/AR and AF/SA perform the best, achieving throughputs between 15% and 21% better than standard TCP version. It is clear from these results that reducing the frequency of ACKs alone is not sufficient—AF alone performs even worse than standard TCP version and techniques like SA or AR are essential to achieve good performance.

Acknowledgement

This work is a part of the research project ВУ-ТН-105/2005 of Bulgarian Science Fund at Ministry of Education and Science.

References

[1] Hossam Afifi, Omar Elloumi, Gerardo Rubino, “A Dynamic Delayed Acknowledgement to Improve TCP Performance for Asymmetric Links” Computers and Communications, *ISCC '98*. Proceedings. Third IEEE Symposium on , 1998

[2] H. Balakrishnan, V.N. Padmanabhan and R.H. Katz, “ The effects of TCP/IP performance,” Proceedings of *INFOCOM'99*

[3] T.V. Lakshman “Window-based error recovery and flow control with a slow acknowledgement channel: a study of TCP/IP performance.” *INFOCOM '97*.

[4] Henderson, T.R.; Sahouria, E.; McCarme, S.; Katz, “On Improving the Fairness of TCP Congestion Avoidance” *GLOBECOM'98*. Volume: 1 , 1998 , Page(s): 539 -544

[5] Go Hasegawa, Masayuki Murata, and Hideo Miyahar, “ Fairness and Stability of Congestion Control Mechanism of TCP” *INFOCOM '99*. Volume: 3 , 1999 , Page(s): 1329 -1336 vol.3

[6] T.C. Kwok, “Residential Broadband Architecture over ADSL andG.Lite (G.992.2): PPP over ATM,” *IEEE Communications Magazine*, vol. 37, no. 5, pp.84-89, May 1999.