Abstract

The Transmission Control Protocol (TCP) has been the dominant transport protocol in the Internet for many years. One of the reasons to this is that TCP employs congestion control mechanisms which prevent the Internet from being overloaded. Although TCP's congestion control has evolved during almost twenty years, the area is still an active research area since the environments where TCP are employed keep on changing. One of the congestion control mechanisms that TCP uses is fast retransmit, which allows for fast retransmission of data that has been lost in the network. Although this mechanism provides the most effective way of retransmitting lost data, it can not always be employed by TCP due to restrictions in the TCP specification.

The primary goal of this work was to investigate when fast retransmit inhibitions occur, and how much they affect the performance of a TCP flow. In order to achieve this goal a large series of practical experiments were conducted on a real TCP implementation.

The result showed that fast retransmit inhibitions existed, in the end of TCP flows, and that the increase in total transmission time could be as much as 301% when a loss were introduced at a fast retransmit inhibited position in the flow. Even though this increase was large for all of the experiments, ranging from $16 - 301\%$, the average performance loss, due to an arbitrary placed loss, was not that severe. Because fast retransmit was inhibited in fewer positions of a TCP flow than it was employed, the average increase of the transmission time due to these inhibitions was relatively small, ranging from $0,3 - 20,4\%$.