

TCP Congestion Control

КОКШЕНЁВ ВЛАДИМИР ВЛАДИМИРОВИЧ

Томский государственный университет (Томск), Россия

e-mail: vladimir_finf@mail.ru

Аннотация

TCP is the most widely used transport layer protocol in the Internet. The performance of many network services and applications are influenced by TCP performance. The protocol was introduced in 70s, and since then a lot of new standards, proposals and modifications were developed and discussed. Congestion control (CC) is a feature that each TCP implements in attempt to avoid network overloading by excessive traffic. As specified in RFC793 TCP uses windowing mechanism to implement it. The main idea of CC is to adjust offered sending rate depending on network conditions. CC procedures base their decisions on different types of feedback received from a network: loss rate, round-trip or one-way delay, etc. A lot of CC algorithms were developed to address different network types and conditions: Tahoe, Reno, NewReno, Vegas, Hybla, CTCP, Illinois, BIC, CUBIC, Westwood(+), H-TCP, High Speed TCP, Scalable TCP, VenO, YeAH, CHD, etc. TCP NewReno and its predecessors Tahoe and Reno are only standardized by IETF methods.

Based on feedback type used in CC decisions all TCP implementation can be divided into three major categories: loss-based, delay-based and hybrid that uses both loss and delay. Loss-based algorithms are more common. Tahoe, Reno, NewReno, BIC, CUBIC, HSTCP are members of loss-based CC class. Vegas and FAST TCP are examples of delay-based CC approach. CTCP, Illinois and Westwood+ can be counted as hybrid

algorithms.

High-speed long distance networks with high bandwidth-delay product received special attention from engineering society. As a result of countless efforts such CC protocols as HSTCP, BIC, CUBIC, Hybla, Scalable, Illinois, YeAH, HTCP, CHD, FAST, FIT and others were developed. For wireless networks with relatively high packet loss level VenO, Westwood and FIT methods were proposed. NewReno and Vegas TCP implementations are targeted to standard networks.

Being an important networking area, TCP congestion control is an integral part of any operating system. Linux OS supports as much as thirteen CC algorithms: NewReno, Vegas, VenO, Westwood+, BIC, CUBIC, HSTCP, Hybla, Scalable, Illinois, YeAH, HTCP, LP – with CUBIC as a default method (since kernel version 2.6.18). At the same time FreeBSD implements six: NewReno (default in FreeBSD 9.0-release), CUBIC, HTCP, CHD, HSTCP, Vegas. And Microsoft Windows OS uses just two: NewReno and CTCP.

Congestion control stays one of the most important areas in TCP for more than 30 years. During this period a lot of algorithms were proposed to overcome the limitations of classic and standard TCP implementations.