

TCP Friendliness

Adrian Kosmaczewski

2007-03-20

TCP congestion control tries to bring solutions to the problem of packet loss, which happens in networks for several reasons:

- Increasing number of hosts connected to the network;
- Increased traffic to and from the hosts;
- Limited capacity of routers in between hosts and networks (Kurose & Ross, page 254).

Congestion is an important problem raised by multimedia streaming applications, where a high quantity of UDP packets (typically with audio and video information) is sent through the network with important rates. The problem with this approach is that UDP has no congestion traffic, and as such packets might be dropped by routers if a buffer overflow occurs. TCP, on the other hand, provides network congestion management, but

TCP is not well-suited for many applications such as streaming multimedia, because TCP congestion control algorithms introduce large variations in the congestion window size (and corresponding large variations in the sending rate)

(Kyasanur)

Network congestion, in turn, has complex consequences in terms of lower perceived quality of service, which is easily seen by the low quality of IP telephony communications in peak hours, bad quality video connections, and the commercial implications of these problems (difficulty to newcomers to provide innovative services, low acceptance by end users, etc).

The ideal solution to the problem of network congestion would be to be able to benefit from the congestion management features of TCP, while being able to deliver complex information at high rates, coupled with the low overhead of UDP. “TCP-friendly” is the generic name given to a set of algorithms that provide the advantages of congestion control while using UDP as the transport protocol, and is currently a hot research topic (this is easily seen by the high number of papers that are returned as search results for the terms “TCP-friendly”: <http://www.google.com/search?q=tcp%20friendly>).

The IETF has published a standard “TCP-friendly” protocol, called TFRC or “TCP Friendly Rate Control”. The official specification states that

TFRC is a congestion control mechanism for unicast flows operating in a best-effort Internet environment. It is reasonably fair when competing for bandwidth with TCP flows, but has a much lower variation of throughput over time compared with TCP, making it more suitable for applications such as telephony or streaming media where a relatively smooth sending rate is of importance.

(Network Working Group, 2003)

The issue of all current research in the subject will most likely have a huge impact on the Internet, with important consequences in terms of quality of service, reliability and speed, and also in terms of commercial and business implications.

References

Kurose, J. F.; Ross, K. W.; “Computer Networking, Third Edition”, Addison Wesley, ISBN 0-321-22735-2

Kyasanur, Pradeep; “Analysis of TCP-friendly Congestion Control Algorithms”, http://www.crhc.uiuc.edu/~kyasanur/papers/cs497rhk_tcp.pdf (Accessed December 10th, 2006)

Network Working Group, “RFC 3448 - TCP Friendly Rate Control (TFRC): Protocol Specification”, January 2003, [Internet] <http://www.faqs.org/rfcs/rfc3448.html> (Accessed December 10th, 2006)