Telecommunication Systems 30:4, 295–296, 2005

Guest Editorial

Up to recently a central objective in developing new TCP protocols has been to produce a global transport protocol that would perform well on any support and would be "friendly" to existing versions. However, as the heterogeneity of networks increases, it becomes more and more difficult to obtain satisfactory performance on all supports. This explains the need to either (i) design and develop new protocols that are more adaptive to network conditions or (ii) to design TCP protocols suitable for a specific type of network conditions, and/or abandon the TCP friendliness to one single protocol.

This special issue is devoted to the analysis of new TCP protocols which have been designed to operate well in fast long-distance networks—ones operating at 622 Mbit/s, 2.5 Gbit/s, or 10 Gbit/s and spanning several countries or states. The protocols need to allow transmission of between 10 GB and multi-TB datasets over gigabit networks. Application domains for such massive transfers include data-intensive Grids, database mirroring for Web sites, and push-based Web cache updates.

The five papers collected in this special issue analyze the following well established protocols for fast long-distance networks.

1. HighSpeed TCP proposed in [5] is studied in [1,2]. This protocol is a special attempt to achieve higher adaptiveness than standard TCP protocols by allowing the additive window increase parameter as well as the multiplicative window decrease parameter to change with the window size. This allows for solving the problem of standard TCP protocols which are too slow in recovering from losses. The adaptive mechanism preserves TCP friendliness when used in an environment which does not have a large bandwidth delay product.
2. The scalable TCP proposed in [7] is studied in [1–4]. This protocol manages to recover quickly from losses and thus solve the main problem of standard TCP operating over fast long-distance networks. On the other hand it abandons the "TCP friendliness" paradigm.
3. The FAST TCP as proposed in [6] is studied in [2]. Unlike the previous novel protocols which use loss (or packet marking) as a feedback, FAST TCP uses round trip delay information to detect congestion.

In addition to these protocols, this issue contains analytical approaches to and a study of TCP with general window increase and decrease factor [2, 8].

Several papers in this issue compare various versions of TCP. [4] studies the fairness when various TCP connections using MIMD (Multiplicative Increase and Multiplicative decrease) share a common bottleneck link, and also when MIMD connections compete with AIMD (Additive Increase Multiplicative Decrease) connections. [1] studies through simulations the performance of HSTCP and Scalable TCP over satellite links and compares them to the performance of NewReno and SACK. In [3] the authors investigate the coexistence of Fast TCP and TCP Reno, and identify various equilibria. This phenomenon is related to the fact that the bandwidth sharing between various TCP connections can be computed by assigning to each TCP an objective function which depends on the TCP version used. In the case when all connections use the same TCP version, the bandwidth sharing turns out to be the solution of a convex optimization problem which has a unique solution. But in the case of various TCP versions, the bandwidth sharing is obtained using a game formulation (where each version of TCP behaves as if it tried to maximize its own objective function). This game formulation can result in multiple solutions (equilibria). Finally, in [8], simple throughput formulae are obtained and compared for a whole family of protocols where both the increase and the decrease factors vary as a function of the window size.

Acknowledgment

We wish to thank the excellent job done by the referees.

References

[1] A Simulation Study of Scalable TCP and HighSpeed TCP in Geostationary Satellite Networks, by G. Giambene, D. Miorandi, this issue.
[2] Comparative Study of different adaptive window protocols, by E. Altman, K. E. Avrachenkov, A.A. Kherani, B. J. Prabhu, this issue.
[3] Equilibrium and fairness of networks shared by TCP Reno and FAST, by J. Wang, K. Tang, S. Hegde and S. H. Low, this issue.
[4] Fairness in MIMD Congestion Control Algorithms, by E. Altman, K. E. Avrachenkov, B. J. Prabhu, this issue.
[5] S. Floyd, HighSpeed TCP for large congestion Windows, RFC 3649, Experimental, December 2003. Available at www.icir.org/floyd/hstcp.html
[6] C. Jin, D.X. Wei, S.H. Low, G. Buhrmaster, J. Bunn, D.H. Choe, R.L.A. Cottrell, J. C. Doyle, W. Feng, O. Martin, H. Newman, F. Paganini, S. Ravot, S. Singh, FAST TCP: from theory to experiments, IEEE Network 19(1) (2005) 4–11.
[7] T. Kelly, Scalable TCP: Improving performance in highspeed wide-area networks, Submitted for publication, December 2002. Available at http://www-lce.eng.cam.ac.uk/ ctk2/scalable/
[8] Transport protocols in the TCP and their performance, by T. J. Ott, this issue.

Sándor Molnár and Eitan Altman