A Systematic Analysis of High Throughput TCP in Real Network Environments*

Yee-Ting Li†, Stephen Dallison‡, Richard Hughes-Jones‡, Peter Clarke†

†Centre of Excellence in Networked Systems
High Energy Particle Physics
University College London
UK, WC1B 6BT

‡Particle Physics Group
Department of Physics and Astronomy
The University of Manchester
UK, M13 9PL

Extended Abstract

Within the last few years, computer scientists and intensive network users have discovered that standard TCP does not perform well in high bandwidth delay environments [1]. It is this poor scalability of standard TCP that results in network throughput that is well below the network capacity even on un-congested links. In this paper, we explore the challenges of achieving high throughput TCP based transport over real environments and comment on the deployment issues of new TCP algorithms for use in high capacity networks.

As high speed networks become increasingly widespread, the problems of the Additive Increase and Multiplicative Decrease factors of standard TCP become more apparent: standard TCP is slow to take up spare capacity in high bandwidth delay environments, and congestion or even the occasional lost packet, causes massive decreases in throughput. This behaviour causes very bad overall performance.

Several proposals have been made that offer solutions to obtaining high throughput transport in high bandwidth delay environments. Amongst those, we investigate protocols that only require sender side modifications to the TCP stack. We have tested Floyd's HighSpeed TCP [2], Kelly's ScalableTCP [3] and Caltech's FAST protocol [4] on real networks with a set of systematic tests using different network conditions on both the long distance, high throughput testbed of DataTAG [5] and a low delay, high throughput network of MB-NG [6].

For each protocol mentioned above, we have studied the effects of running parallel streams of test traffic and comment on the implications of intra-protocol fairness and whether there are any practical benefits over striped standard TCP flows, which are prolific today in achieving high throughput [7]. We also look at how well these protocols perform under stressed networks with both constant bit rate and self-similar background traffic patterns and determine that whilst preliminary studies

*To be presented at the Second Internatioal Workshop on Protocols for Fast Long-Distance Networks, February 16-17, 2004, Argonne National Laboratory, Argonne, Illinois USA.

show that HighSpeed TCP is safely deployable, the effects of ScalableTCP have a greater impact on existing traffic. To complete the systematic analysis we also induce network loss to compare against the theoretical TCP response function [8] where applicable.

In order to gain wide acceptance and deployment, new TCP transport protocols have also to be able to maintain network stability and fairness to prevent Internet collapse. We therefore also include an experimental study of the convergence, stability and fairness of each protocol; both individually and when competing against each other. To quantify these, we use the popular metric Coefficient of Variance as an indication of the stability of these new TCP stacks in both intra and inter-protocol experiments; and we introduce a measure of fairness that is based in terms of the impact of new TCP flows upon existing traffic and apply it to these protocols.

In order to analyse end-to-end performance in real life scenarios, we also include a study of the behaviour and problems of transporting real data across the Internet using popular application level transport programs such as GridFTP [9], bbcp [10] and http [11]. The TCP behaviour of each application is examined and performance differences are discussed.

We have shown how various high throughput modifications of the TCP protocol behave in real network environments and have highlighted practical limitations. We have also measured the impact of these protocols with competing network traffic in a quantifiable and systematic way and show possible deployment issues in obtaining high throughput transport.

References

[1] "First International Workshop on Protocols for Fast Long-Distance Networks", http://datatag.web.cern.ch/datatag/pfldnet2003

[2] Sally Floyd, "HighSpeed TCP for Large Congestion Windows", IETF Internet Drafts, http://www.ietf.org/internet-drafts/draft-ietf-tsvwg-highspeed-01.txt

[3] Tom Kelly, "ScalableTCP: Improving Performance in HighSpeed Wide Area Networks", First International Workshop on Protocols for Fast Long-Distance Networks, February 2003

[4] Cheng Jin, David X. Wei and Steven Low, "FAST TCP for high-speed long-distance networks", http://netlab.caltech.edu/pub/papers/draft-jwl-tcp-fast-01.txt

[5] "Research and Technological Development for Transatlantic Grid", http://www.datatag.org

[6] "Managed Bandwidth - Next Generation", http://www.mb-ng.net

[7] "IEPM: Internet End-to-end Performance Monitoring", http://www-iepm.slac.stanford.edu/monitoring/bulk/window-vs-streams.html

[8] J. Padhye, V. Firoiu, D. Towsley and J. Kurose, "Modeling TCP Throughput: A Simple Model and its Empirical Validation", Proceedings of SIGCOMM'98, ftp://gaia.cs.umass.edu/pub/Padhye-firoiu98-TCP-throughput.ps

[9] W. Allcock, J. Bester, J. Bresnahan, A. Chervenak, L. Liming, S. Meder, S. Tuecke, "GridFTP Protocol Specification", GGF GridFTP Working Group Document, http://www.globus.org/research/papers.html#GridftpSpec02, September 2002

[10] "BaBar Copy", http://www.slac.stanford.edu/~abh/bbcp

[11] R. Fielding, J.Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, T. Berners-Lee, "Hypertext Transfer Protocol - HTTP/1.1", RFC2616, ftp://ftp.isi.edu/in-notes/rfc2616.txt