ECSE-489 : TELECOMMUNICATION NETWORK LAB
WINTER 2003
(2 CREDITS)

Experiment 4

Introduction

This experiment is designed to let you explore aspects of the TCP/IP. Two-thirds of the experiment involves OPNET, and the remaining third involves Internet measurement. There is a degree of flexibility in the experiment: you have a choice in what you choose to undertake, demonstrate and report. Your goal should be to design and perform experiments that enable you to address the questions below. In the demonstration session, you will be asked to justify your experiments, describe how they were performed, and demonstrate key steps in your procedure.

PART 1: OPNET SIMULATIONS

Your starting point for this section should be the *TCP* project made available in OPNET. Open this project, and copy a renamed and editable version to your workspace.

Section 1.1: TCP slow-start and congestion avoidance

Design and perform experiments to explore how the following parameters affects on the length of time spent in the slow-start phase:

(i) propagation delay
(ii) data rate

Section 1.2: TCP SACK

Design and perform experiments that compare the recovery of TCP Reno (SACK disabled) with that of SACK-enabled TCP Reno in the presence of *multiple* drops. How is this affected by propagation delay and data rate?

Section 1.3: High-speed TCP

Design experiments to explore the effectiveness of Sally Floyd's high-speed TCP proposal. In commenting on effectiveness, examine how bandwidth is shared by two competing connections commencing at different times.

We have provided HSTCP, an OPNET project that performs the high speed algorithm with the default parameters proposed in the Internet draft document. Refer to the tips sheet (posted in discussions on Monday afternoon) for our list of OPNET issues/techniques.

PART 2: INTERNET EXPERIMENTS

Section 2.1: TCP connections

The following equation has been proposed as a theoretical approximation of the *steady-state throughput* of long-duration TCP connections [1]:

$$B(p) = \min\left(\frac{W_{\max}}{RTT}, \frac{1}{RTT\sqrt{\frac{2bp}{3}} + T_0 \min\left(1, 3\sqrt{\frac{3bp}{8}}\right) p\left(1+32p^2\right)}\right)$$

where $B(p)$ is the throughput, W_{max} is the maximum congestion window size (as advertised by the receiver buffer), b is the number of packets acknowledged by a received ACK, T_0 is the timeout period, and p is the probability that a packet is lost given that it is either the first packet in its round or the preceding packet in its round is not lost.

Design experiments that explore the validity of this expression for a range of round-trip-times. In your report, state carefully how you estimated all the values in the equation. Also discuss whether there are any additional factors that may be biasing your results (rate-limiters, pacing mechanisms, proxies etc.).

Large files are available to download at: www.tucows.com. You can also set your mirror site so that you are downloading from a range of locations.

[1] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, "Modeling TCP throughput: a simple model and its empirical validation", Proc. ACM SIGCOMM, Vancouver, Canada, pp. 303-314, 1998