Identifying TCP Congestion Control Algorithms Using Active Probing

Seyedshams Feyzabadi and Jürgen Schönwälder
{s.feyzabadi, j.schoenwaelder}@jacobs-university.de
Computer Science, Jacobs University Bremen, Germany

JACOBS UNIVERSITY

Introduction:

- The Transmission Control Protocol (TCP) carries most of the Internet traffic these days
- The main difference among different TCP implementations are the algorithms for controlling congestion
- While RFCs and other documents specify how TCP should behave, it is not clear what implementations really do

Motivation:

- Understanding TCP behaviour is important
- Detecting misbehaving TCP implementations
- We do not always have access to a remote host in order to identify the TCP version running in the operating system
- Not all remote hosts respond probing requests due to firewalls and network address translators
- But web servers supporting HTTP usually respond

Previous Work:

The TCP Behavior Inference Tool, TBIT, uses active probing to characterize TCP behavior of a remote web server.

General Algorithm:

- Send ”fabricated” TCP packets over raw IP sockets
- Host firewall prevents kernel from seeing response packets
- BPF delivers blocked packets to a user-space process
- The user–space process creates different streams of TCP segments to detect the TCP version
- The software sends an HTTP GET request to ask for the index page of the web site
- The web server starts sending packets to the client
- An intentional packet drop occures
- Based on retransmissions, the user-space process detects the congestion control algorithm

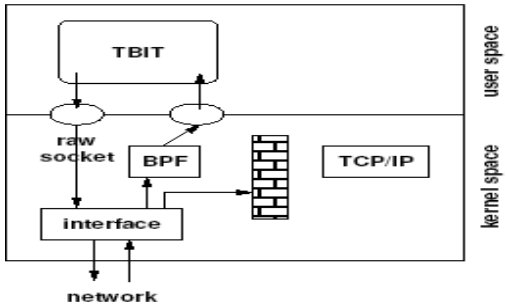

Our Approach:

Based on TBIT architecture, we use fabricated packets

- After packet drop we receive the packets
- We waits for a longer time than one RTT
- We send back the ACK messages to the server and we count the Received Congestion Window size *rcwin*

Note that *rcwin* is the same as the Congestion Window size *cwnd* with one RTT delay.

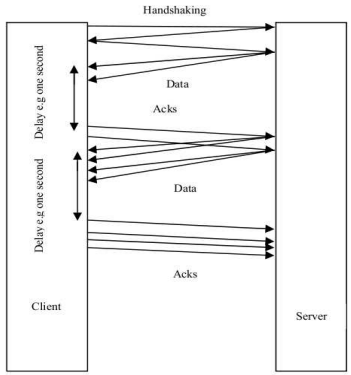

Using the sequence of recorded *rcwin* sizes, we estimate the equation of the *cwnd* growth function. Matlab functions such as *fit* do the estimations for us.

Results obtained in our lab:

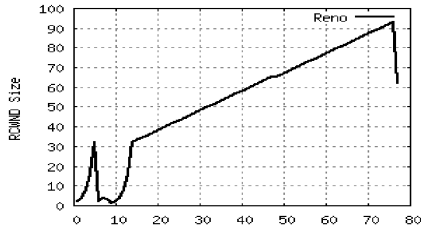

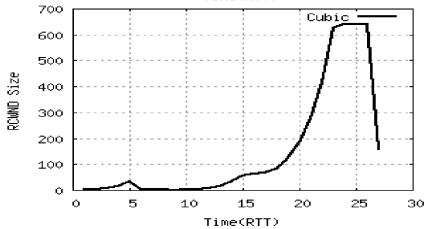

Result obtained on the Internet:

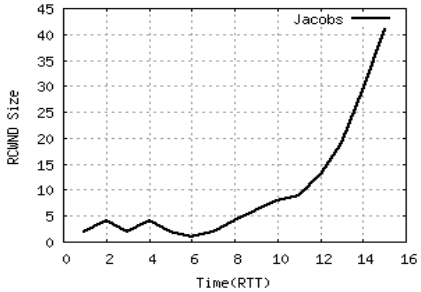

Future Work:

- Large scale tests on the Internet
- Integrate the software with peer-to-peer applications or web servers to detect the congestion control mechanisms used by TCP clients

The work reported in this poster is supported by the EC IST-EMANICS Network of Excellence (#26854).