Receiver-based Management Scheme of Access Link Resources for QoS-Controllable TCP Connections

Go Hasegawa
Osaka University
hasegawa@cmc.osaka-u.ac.jp
http://www.anarg.jp/

Background

- Evolution of the Internet
- Applications with long-lived session
  - P2P, GRID, Streaming, network storage, …
- Congestion in access network becomes apparent
  - Broadband access network is still narrow due to killer applications

Problems in access network

- Multiple network applications on the receiver host
- Congestion at the access link
- Unfairness between short-lived and long-lived TCP connections
  - Mice suffers from low performance due to elephant's occupation

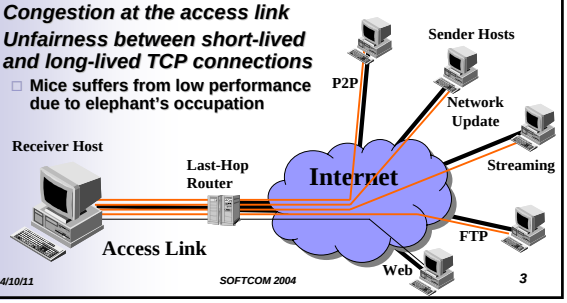

Objectives

- Effective management of access link resources
- Resolve the congestion at the access link
- Relieve short-lived TCP connections from long-lived connections' possession

Design principles

- Receiver-side control
- No cooperation of the last-hop router or sender hosts
- Monitor TCP connections' RTTs to detect the congestion
- Control the receive socket buffer of TCP connections

Detection of access link congestion

- Monitor Round Trip Values (RTT) values of all TCP connections at the receiver host
- When the RTTs of all connections increases
  - Access link congestion occurs
- When RTTs remains constant
  - No congestion at the access link

Resolving access link congestion

- *Control the total amount of the socket buffer for all TCP connections*
  - Decrease when congestion occurs
  - Increase when no congestion occurs
  - Otherwise, remains unchanged
- *Simple, but quite effective*

Buffer assignment for connections

- *Fixed assignment for short/long-lived connections is apparently ineffective*
  - Required size is quite different
- *Short-lived connections should be prioritized over long-lived connections in buffer assignment*
  - Packet loss severely degrades performance of short-lived connections

Assignment for short-lived connections

- *Definition: short-lived connections are TCP connections in slow-start phase*
- *Assign receive buffer according to the change of window size in slow-start phase*
  - $R = 2 \cdot 2^t$ [packets]
  - t : number of RTTs from beginning of the connection

Assignment for long-lived connections

- *Definition: short-lived connections are TCP connections in congestion avoidance phase*
- *Assign receive buffer according to the pre-defined priority and RTT (Round Trip Time) values*
  - $R = (B - T_{is}) p \cdot rtt / \Sigma (p \cdot rtt)$ [packets]
  - p : priority, B : total buffer size at the receiver, T_{is} : total buffer size for short-lived connections

Simulation Settings

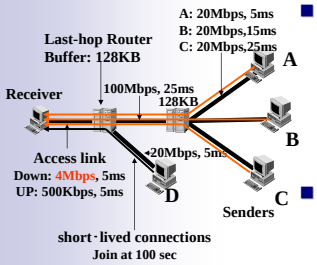

- *Bulk transfer from SenderA, B and C to Receiver*
  - 3 long-lived connections occupy the access link bandwidth
- *After 100 sec, short-lived transfers begin*
  - 100 times 30KB transfers from D to Receiver

Simulation Setting (2)

- *Performance metric*
  - For short-lived connections
    - Connection initiation time
    - Transmission time of 30KB data
  - Utilization of access link bandwidth
  - Queue length at the last-hop router buffer

Simulation Setting (3)

- *Comparison candidates*
  - Spring
    - N. T. Spring, M. Chesire, M. Berryman, V. Sahasranaman, T. Anderson, and B. N. Bershad, “Receiver based management of low bandwidth access links,” in Proccedings of IEEE INFOCOM 2000, pp. 245–254, 2000.
  - Mehra
    - P. Mehra, A. Zakhor, and C. D. Vleeschouwer, “Receiver-driven bandwidth sharing for TCP,” in Proceedings of IEEE INFOCOM 2003, Mar. 2003.
  - Original
    - Fixed assignment

Simulation results

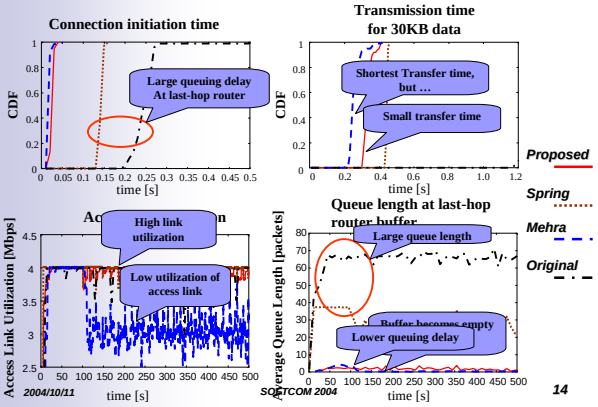

Summary and future work

- *Management scheme for Access link resources*
  - Detect/resolve access link congestion
  - Different buffer assignment for short/long-lived connections
  - Simple, but quite effective
- *Future work*
  - Implementation experiments
  - Integration with the sender-side mechanisms