A Survey on TCP Congestion Control Mechanisms

By
Sridevi Polavaram
CS 756

How does TCP react to congestion

- Each endhost controls its packet transmission rate by changing the widow size in response to network congestion.
- Basic congestion control schemes
  - Slow start
  - Congestion avoidance
  - Fast retransmission
  - Fast recovery
- TCP Reno vs FACK vs RH vs LRD vs Vegas vs…

What could go wrong?

- Misbehaving connection
  - How about TCP Fairness Issue
  - TCP-LP
- TCP starvation leading to congestion collapse
  - How about TCP-Friendly protocols

TCP Friendliness Issue

- Non-TCP flows are defined as TCP-friendly when “their long-term throughput does not exceed the throughput of a conformant TCP connection under the same conditions”.
- The throughput of TCP depends mainly on the parameters round-trip time, retransmission timeout value, segment size, and packet loss rate

Classification of TCP-Friendly Congestion control schemes

- Window-Based vs Rate-Based
- Unicast vs Multicast
- Single-rate vs Multi-rate
- End-to-end vs Router-supported

TCP-Friendly protocol family

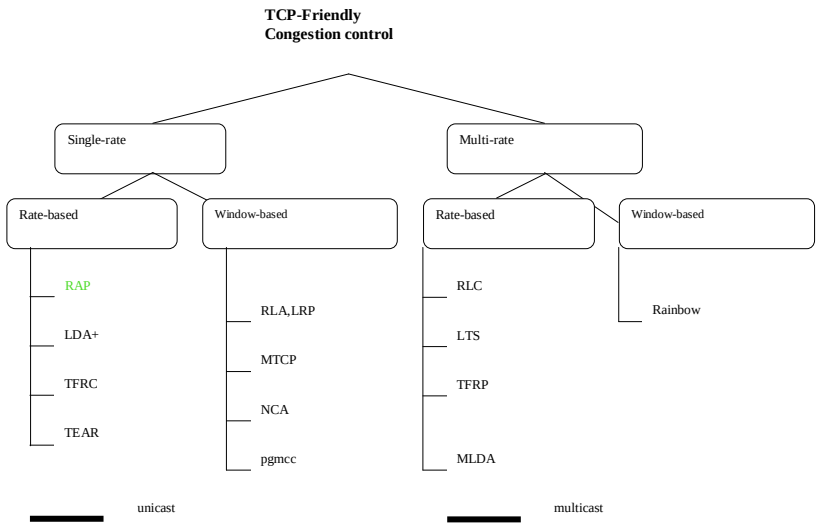

TCP-LP: A distributed algorithm

- Design objectives
- Proposed model vs Real Model
- Useful applications
- Parameter settings

TCP-LP: Design Objectives

- Goal of TCP-LP
  - Use only excess bandwidth as compared to the “fair share” of bandwidth as targeted by TCP
- Non-intrusive to TCP flows
- Fairness among multiple TCP-LP & TCP flows

TCP-LP: Proposed Model

Proposed Model

- Two class hierarchical scheduling model with a high-priority class and low-priority class.
- Strict switching between the two classes according to priority
- In each class, service is fair among competing flows.

TCP-LP: Approximated Real Model

Real Model

- End-point congestion control algorithm
  - All flows are multiplexed into single FCFS queue
  - Service class approximation using two different protocols TCP and TCP-LP
- Early congestion inferences
  - ECN messages
  - One way packet delays
  - Delay threshold
  - Delay measurement

TCP-LP: Congestion avoidance policy

- Congestion avoidance policy
  - Quickly back-off in the presence of congestion from TCP flows
  - Quickly reacting to early congestion indications
  - Achieve fairness among TCP-LP flows
  - Inference phase

TCP-LP: congestion avoidance policy

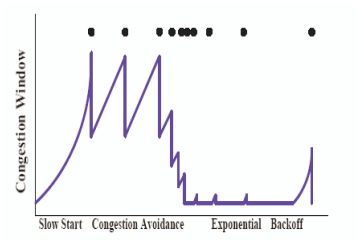

TCP congestion control

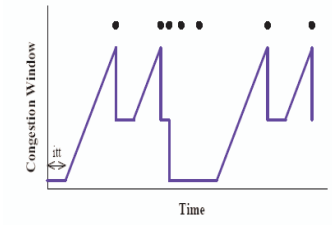

TCP-LP congestion avoidance policy

TCP-LP: Useful Applications

- Low-priority file transfer over internet
- Inference of available bandwidth by making for network monitoring
- Response times of web-connections in the best-effort class decrease by 90%
- Dynamic adaptive algorithm

TCP-LP: Remarks

- How far TCP-LP flows with large round-trip times can still infer congestion prior to TCP flows with small round-trip times
  - Queuing model approach

References

- TCP-LP: A Distributed Algorithm for Low Priority Data Transfer A.Kuzmanovic and E.W.Knightly IEEE INFOCOM 2003
- Go Hasegawa, Masayuki Murata. “Survey on fairness issues in TCP congestion control mechanisms”
- The Effectiveness of End-to-End Congestion Control Mechanisms J.Bolliger, U.Hengartner and Th. Gross
- J. Widmer, R. Denda, and M. Mauve, “A survey on TCP-friendly congestion control,“ IEEE Networks, pp. 28{37, May/June 2001.