H-TCP: TCP Congestion Control for High Bandwidth-Delay Product Paths

draft-leith-tcp-htcp-00

Doug Leith, Robert Shorten

High Bandwidth-Delay Product Paths

The pipe size of a link is roughly $BT+q_{max}$

where B is the the link rate (packets/s), T is the propagation delay and q_{max} is the queue size.

On a long distance gigabit link, B=100,000 packets/s, T=200ms, q_{max}=1000 and

$$BT+q_{max}=21,000$$

Note that the pipe size determines the peak window size of a TCP source.

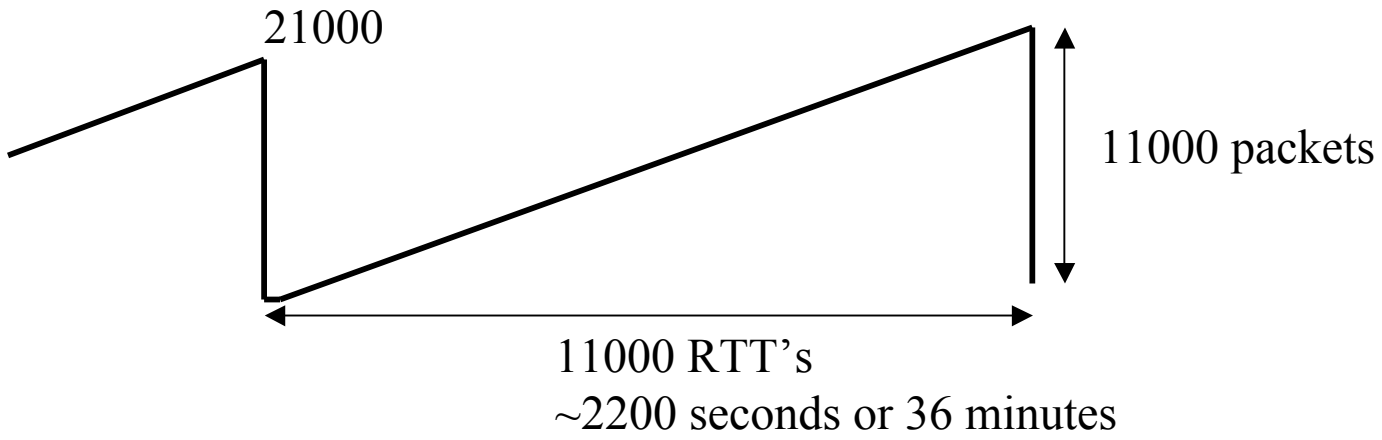

- TCP becomes sluggish, and requires v.low drop rate to achieve reasonable throughput.

Background

- Scalable TCP
- High-Speed TCP
- FAST TCP

etc

Moving forward

Seek to re-open discussion on congestion control algorithms for high BDP paths

Guiding principle - seek *smallest changes* to TCP that yield scalability with respect to BDP.

H-TCP - Adjust increase rate as function of time since last backoff

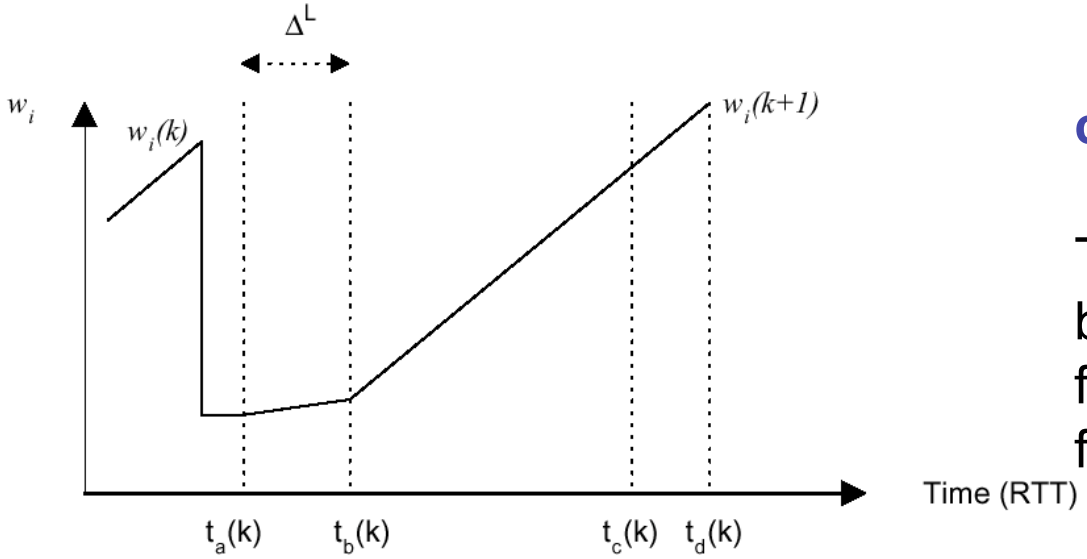

cwnd←cwnd+f(T)/cwnd

T elapsed time since last backoff
f(.) determines response function.

Preserves symmetry in network (newly started flows with small cwnd compete on level playing field).

- Responsiveness (measured in congestion epochs) similar to standard TCP
- Fairness properties (including RTT unfairness) similar.
- Backward compatibility guaranteed on low BDP paths.
- Aggressiveness/response function is a design parameter.

Current Status

Extensive experimental testing over last two years or so
-SLAC 2004
-Hamilton Institute 2005
plus misc smaller tests by various groups.

Stable algorithm/Linux implementation

Ongoing testing in more diverse environments

Initial I-D to solicit comments.

Example Scalable TCP

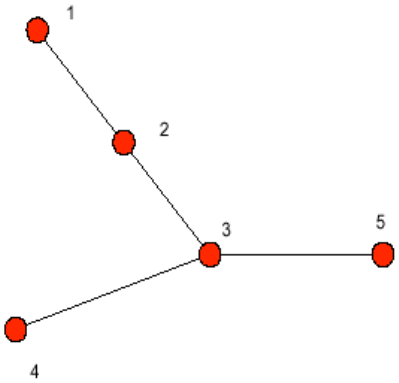

Scaleable TCP has convergence issues …