Hamilton Institute

Network Congestion Control

HEAnet Conference 2005

(David Malone for Doug Leith)

What has TCP ever done for us?

- Demuxes applications (using port numbers).
- Makes sure lost data is retransmitted.
- Delivers data to application in order.
- Engages in congestion control.
- Allows a little out-of-band data.
- Some weird stuff in TCP options.

Hamilton Institute

Congestion Control

- TCP controls the number of packets in the network.
- Packets are acknowledged, so flow of ACKs.
- Receiver advertises window to avoid overflow.
- Congestion window tries to adapt to network.
- Slow start mechanism to find rough link capacity.
- Congestion avoidance to gradually adapt.

The Congestion Window

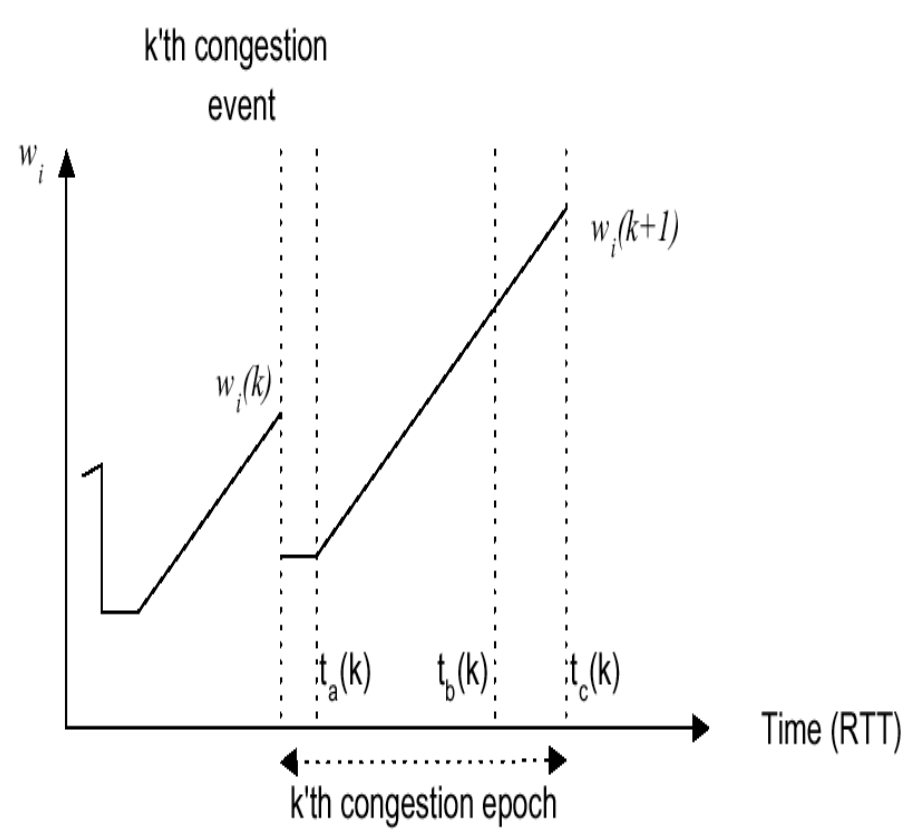

- Additive increase, multiplicative decrease (AIMD).
- To fill link need to reach BW x Delay.
- Backoff by 1/2 => buffer at bottleneck link should be BW x Delay.
- Fairness (responsiveness, stability, …)

Problems for Congestion Control

- Packet loss caused by other factors.
- Filling a big link at one-packet-per-round-trip.
- The combination is bad for high speed long distance links.
- Problem was flagged up: various solutions being studied (Scalable TCP, High-Speed TCP, FAST TCP, H-TCP, …)

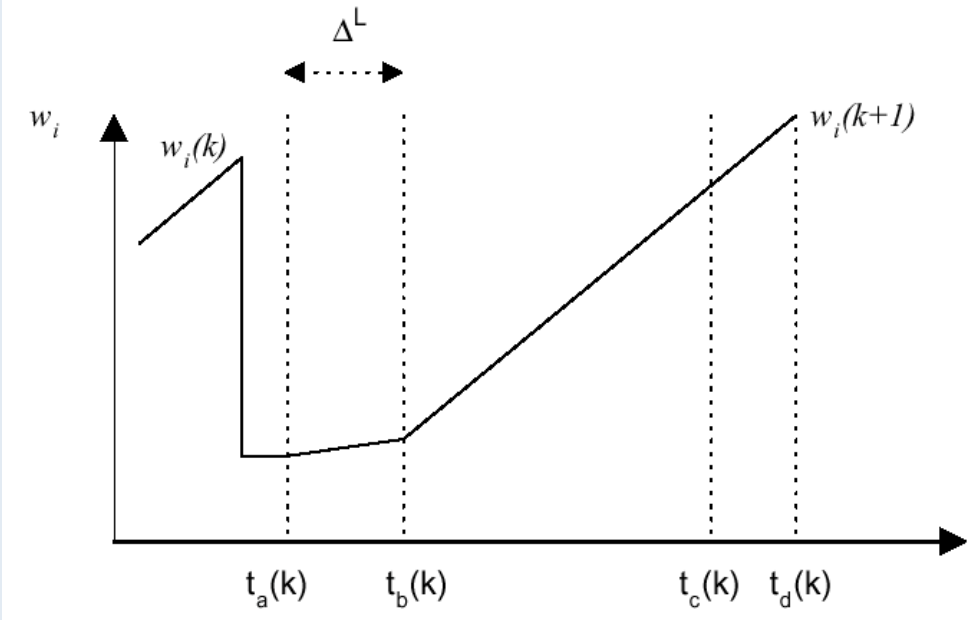

H-TCP

- Aim to make small changes we can analyse.
- Rate of increase depends on how long since last backoff.
- New flows compete on level playing field.
- Similar fairness & responsiveness.
- Competes fairly with normal TCP where it can compete.

Hamilton Institute

Practical Issues

- Congestion control isn't the only issue.
- Implementation is important.
- Land speed records.
- Web 100 project to instrument Linux.
- Important stack tuning to be done.

Before Tuning

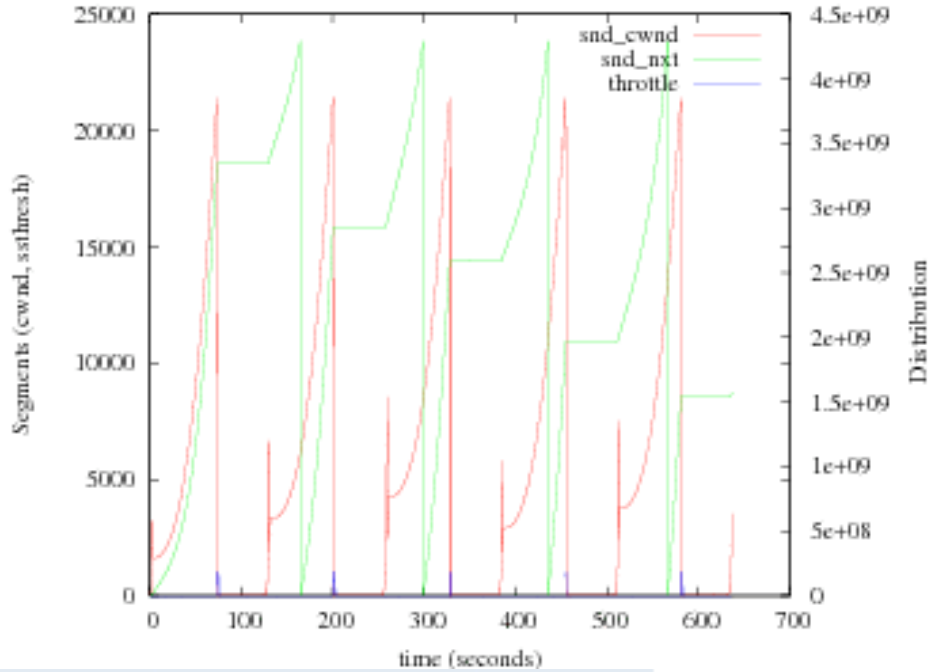

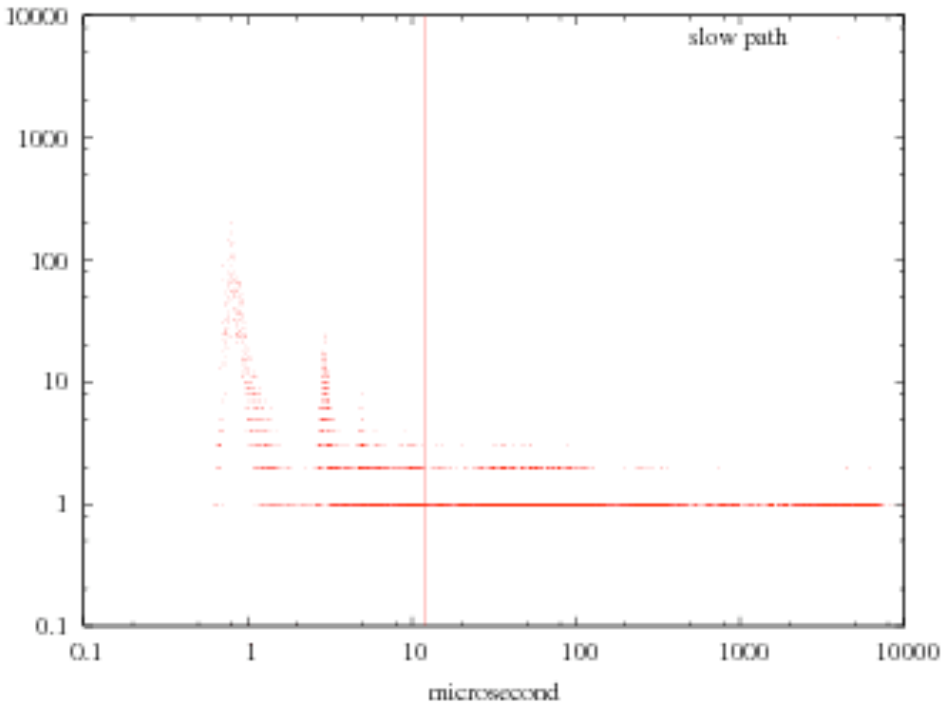

After Tuning

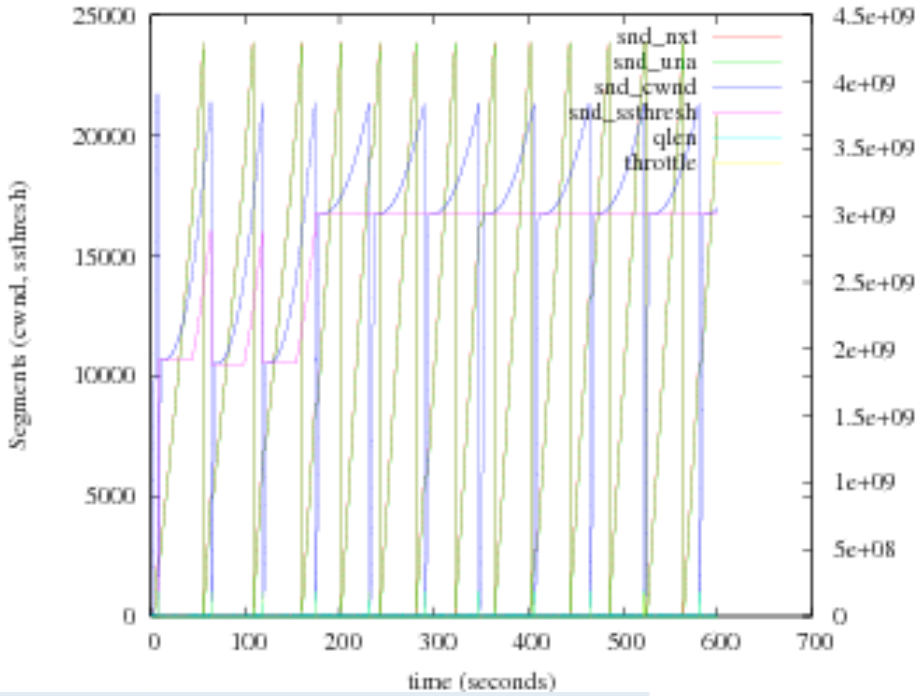

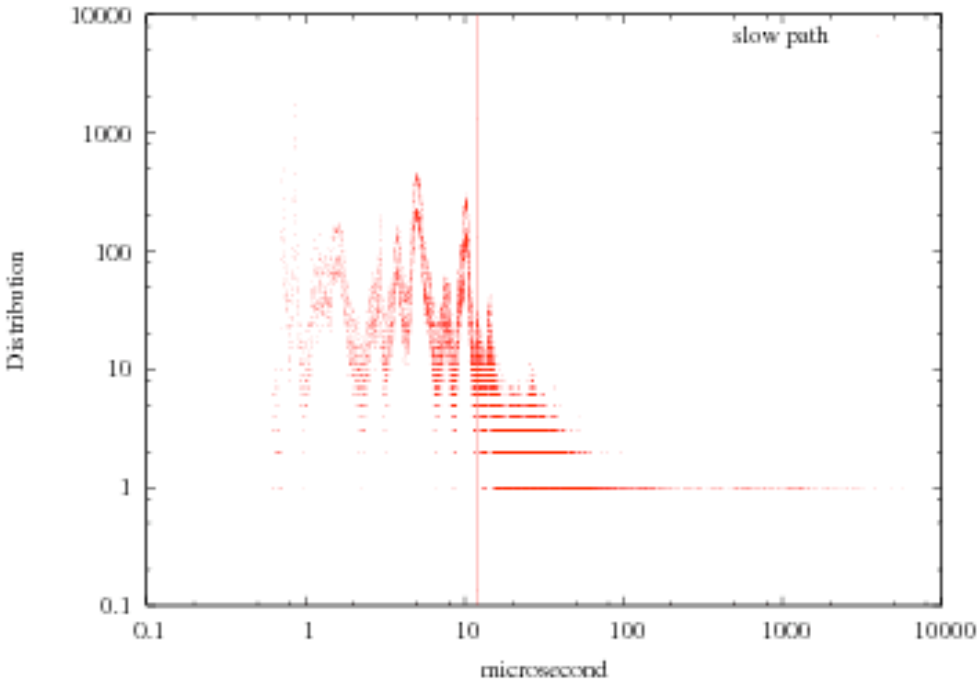

Hamilton Institute

Other issues

- High speed is important (packet switched vs. circuit switched), but not for everyone.
- Sizing router buffers important to everyone (cost, QoS, optics).
- Wireless interesting - random losses.
- Other interesting wireless issues too.

802.11 AP Before

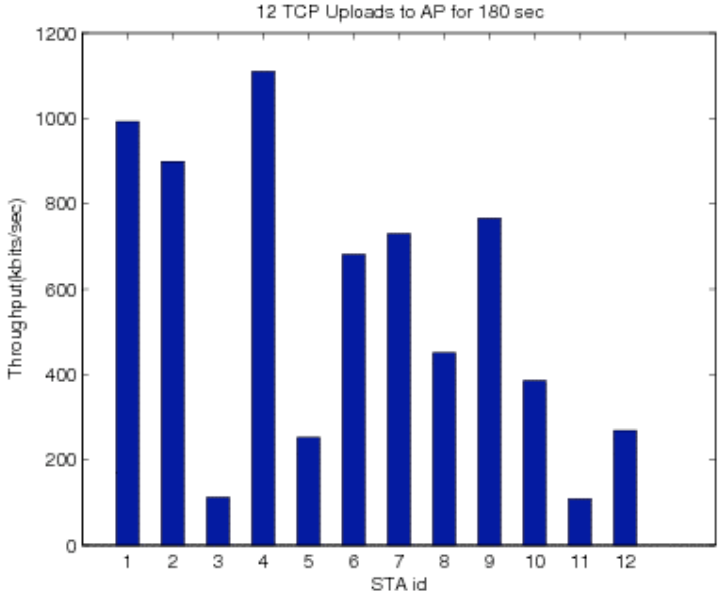

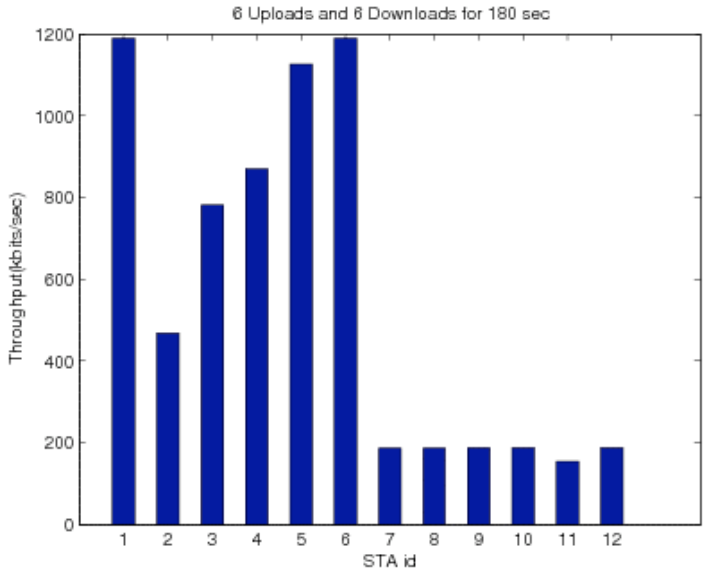

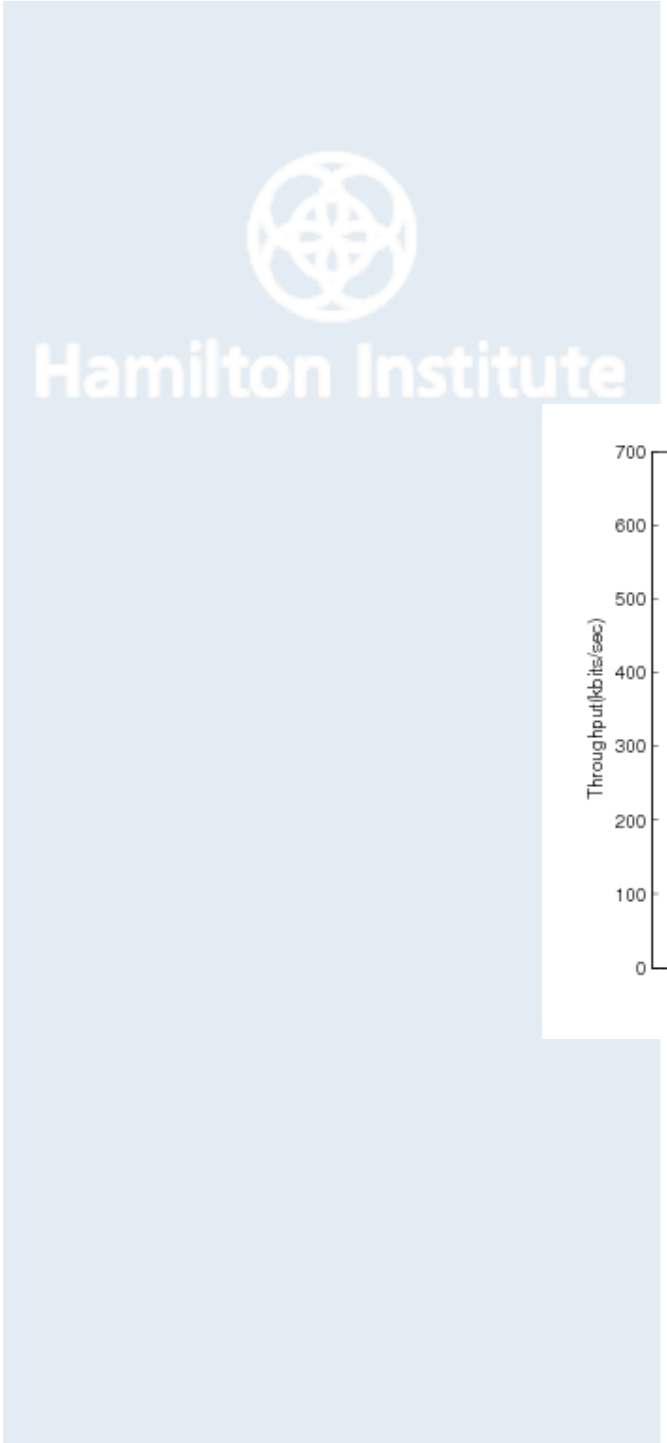
Hamilton Institute

Baseline

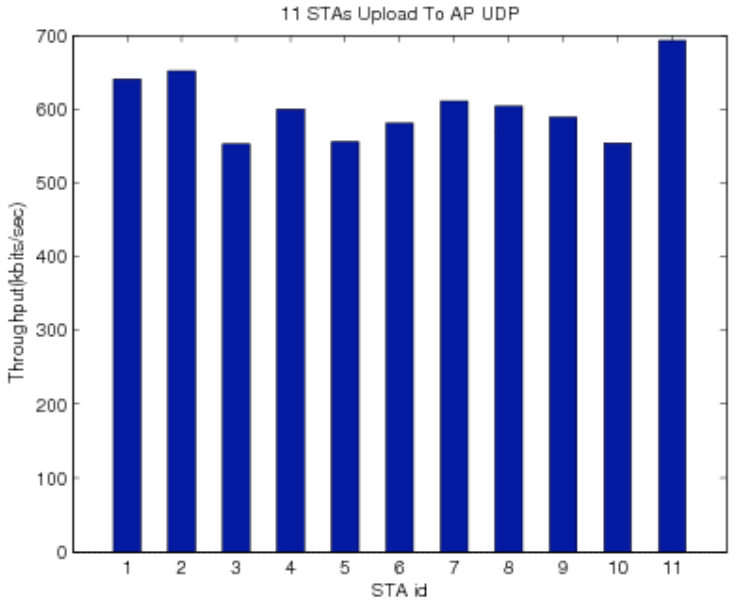
11 STAs Upload To AP UDP
Throughput(kbits/sec)
700
600
500
400
300
200
100
0
1
2
3
4
5
6
7
8
9
10
11
STA id

Adjusting 802.11e Parameters

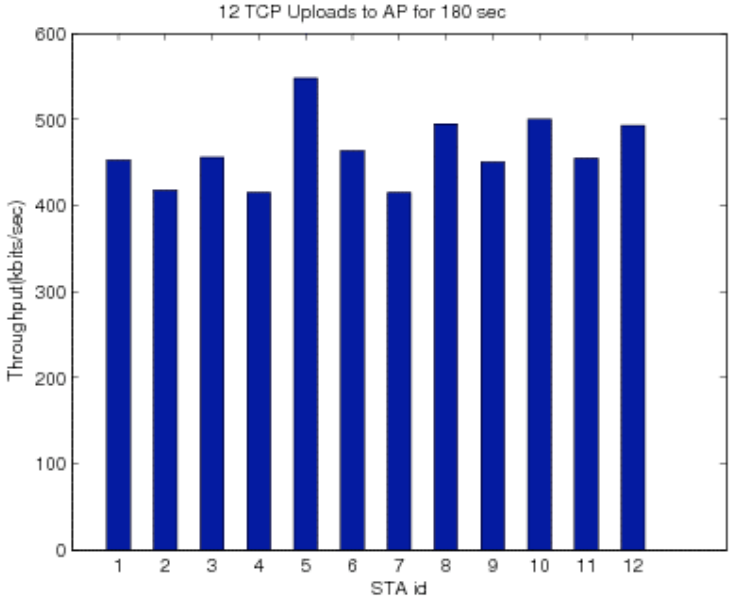

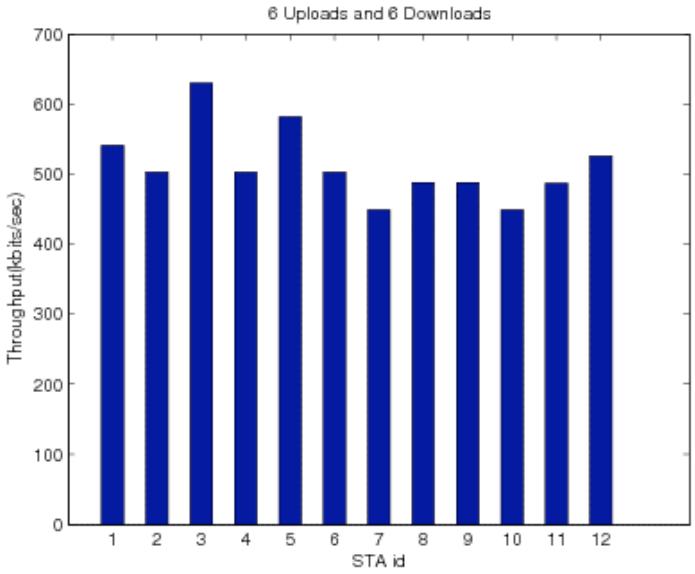

Thanks! Questions?