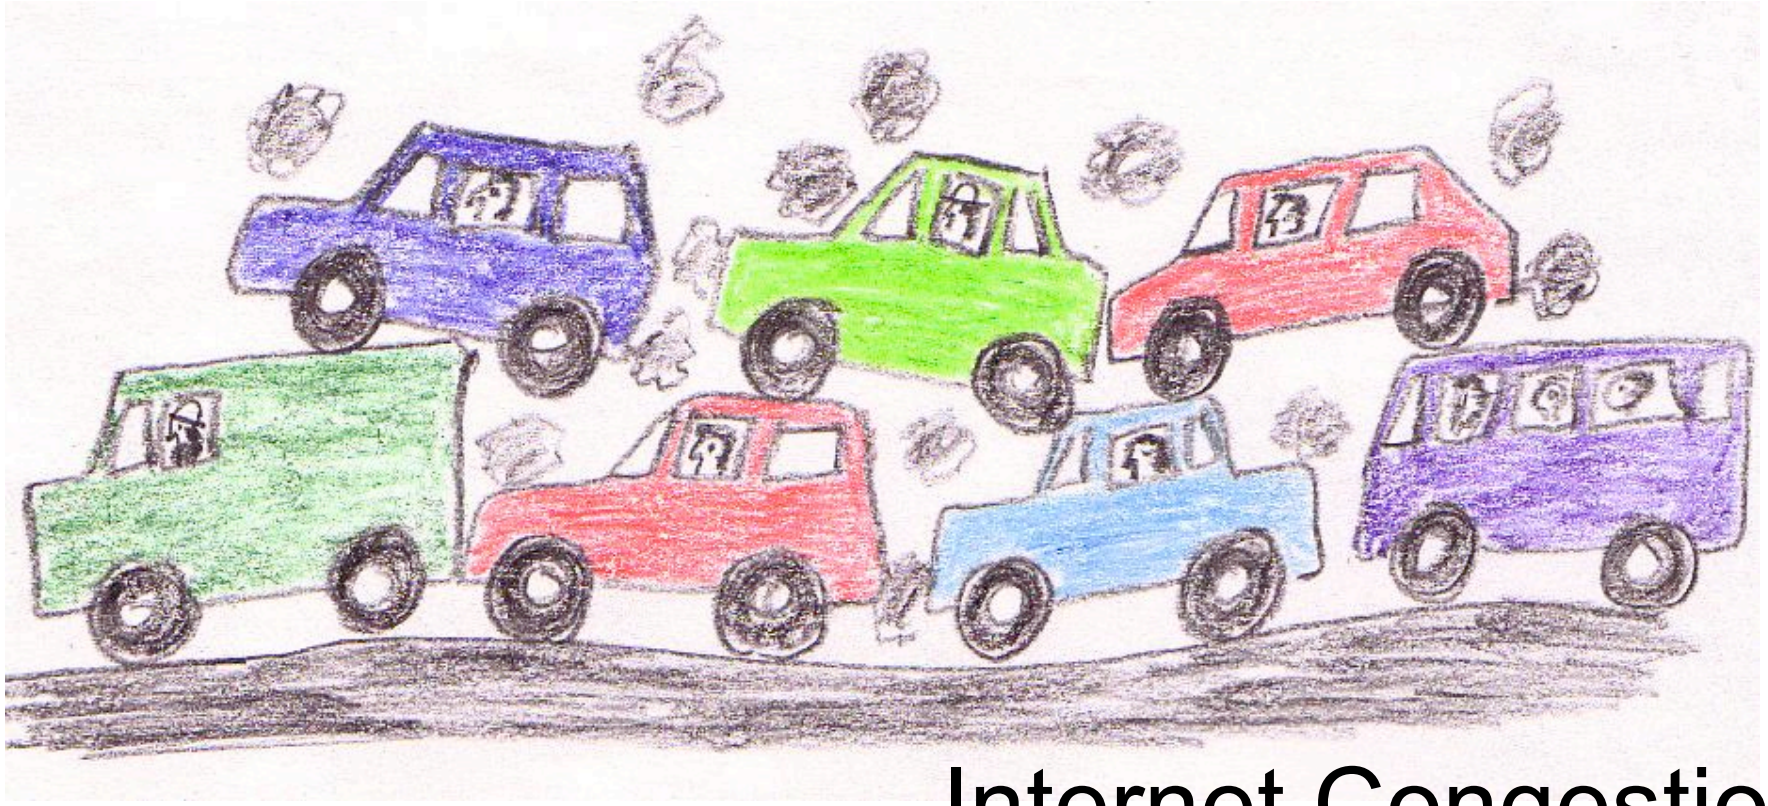

Internet Congestion Control Research Group

Mark Handley

UCL

Congestion Control

- The Internet only functions because TCP's congestion control does an effective job of matching traffic demand to available capacity.

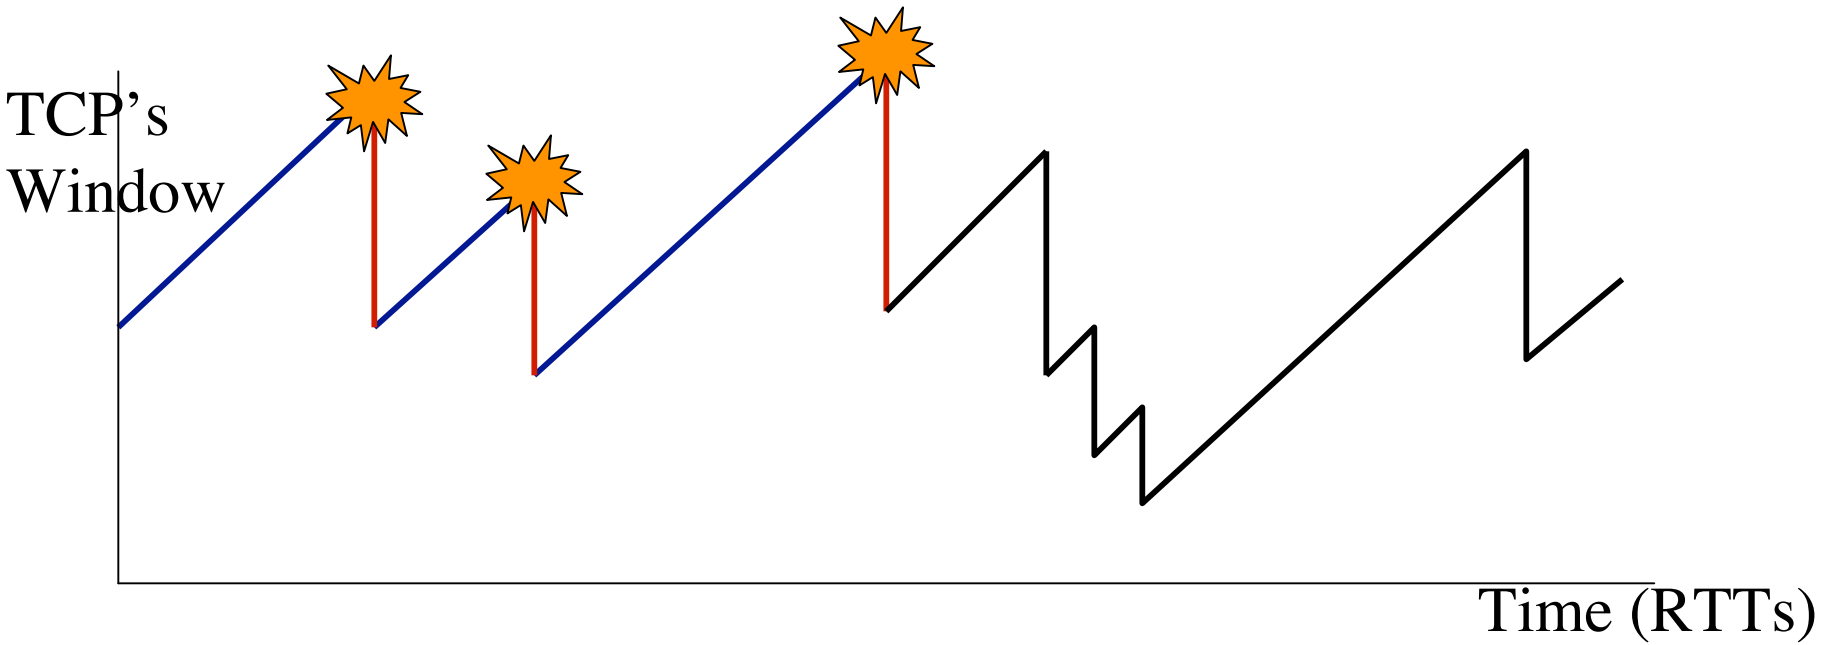

But my network doesn't have congestion!

- Maybe.
- But the end-to-end path should if we've done our job right.
- File transfer:
  - Move *x* bytes from *a* to *b* in time *t*.
  - Applications work better as $t \rightarrow 0$
- Realistically, *t* will never be zero, but our long term goal should be to make it as close to one RTT as possible.

Limitations of AIMD Congestion Control

(Additive Increase, Multiplicative Decrease)

- Very variable transmit rate is fine for bulk-transfer, but hard for real-time traffic.

  RFC3448: TCP-Friendly Rate Control (TFRC)

  RFC????: Datagram Congestion Control Protocol (DCCP)

Limitations of AIMD Congestion Control

- Failure to distinguish congestion loss from corruption loss.
  - Wireless

- Limited dynamic range.

$$\text{transmit rate} \sim= \frac{\text{packet size}}{\text{RTT}\sqrt{\text{loss rate}}}$$

AIMD: Limited Dynamic Range

One loss every half hour, 200ms RTT, 1500bytes/pkt.

⇒ 9000 RTTs increase between losses.

⇒ peak window size = 18000 pkts.

⇒ mean window size = 12000 pkts.

⇒ 18MByte/RTT

⇒ 720Mbit/s.

⇒ Needs a bit-error rate of better than 1 in 10^12.

⇒ Takes a very long time to converge or recover from a burst of loss.

Opportunity

- We *will* need to change the congestion control dynamics of the Internet.
- This presents an opportunity to do it right and solve many additional problems at the same time.
  - Wireless?
  - Smooth throughput for multimedia?
  - Low delay service?
  - DoS resistant?
- Always easier to solve only the immediate problem.

XCP: eXplicit Control Protocol

Katabi, Handley, Rohrs, Sigcomm 2002

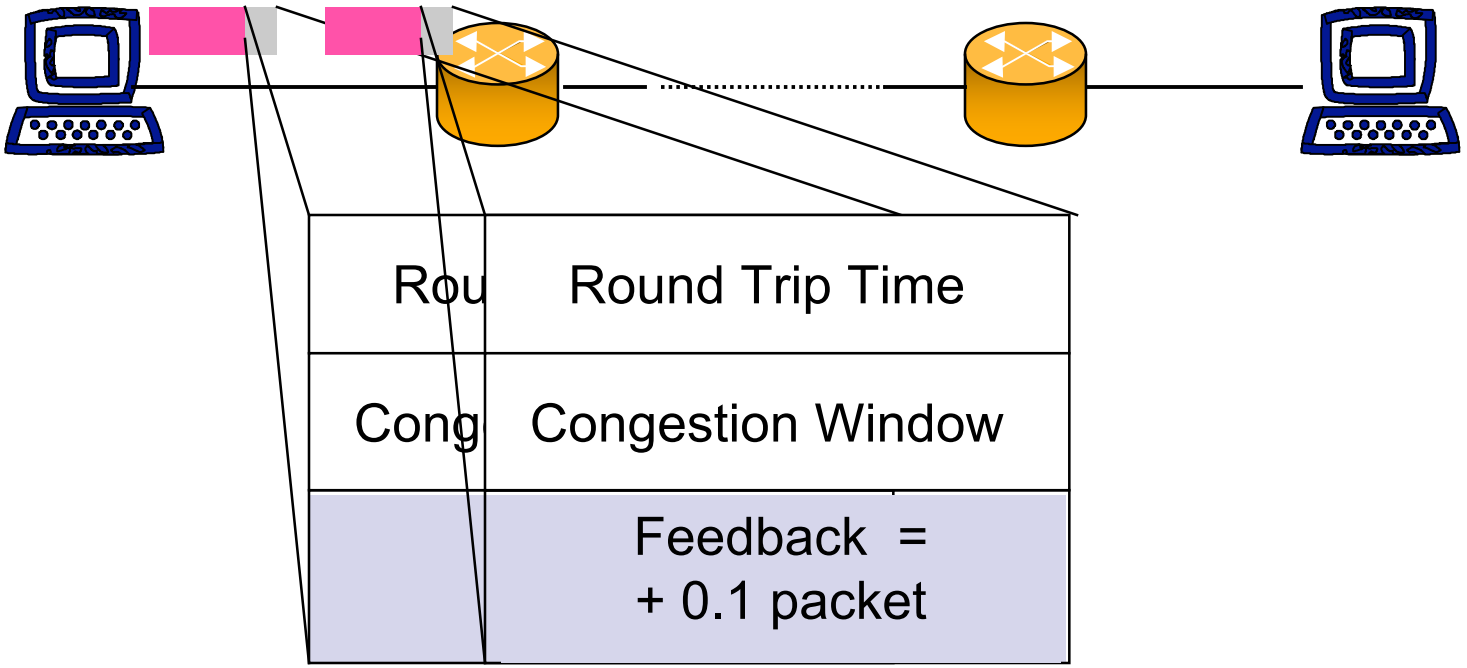

Congestion Header

XCP: eXplicit Control Protocol

Katabi, Handley, Rohrs, Sigcomm 2002

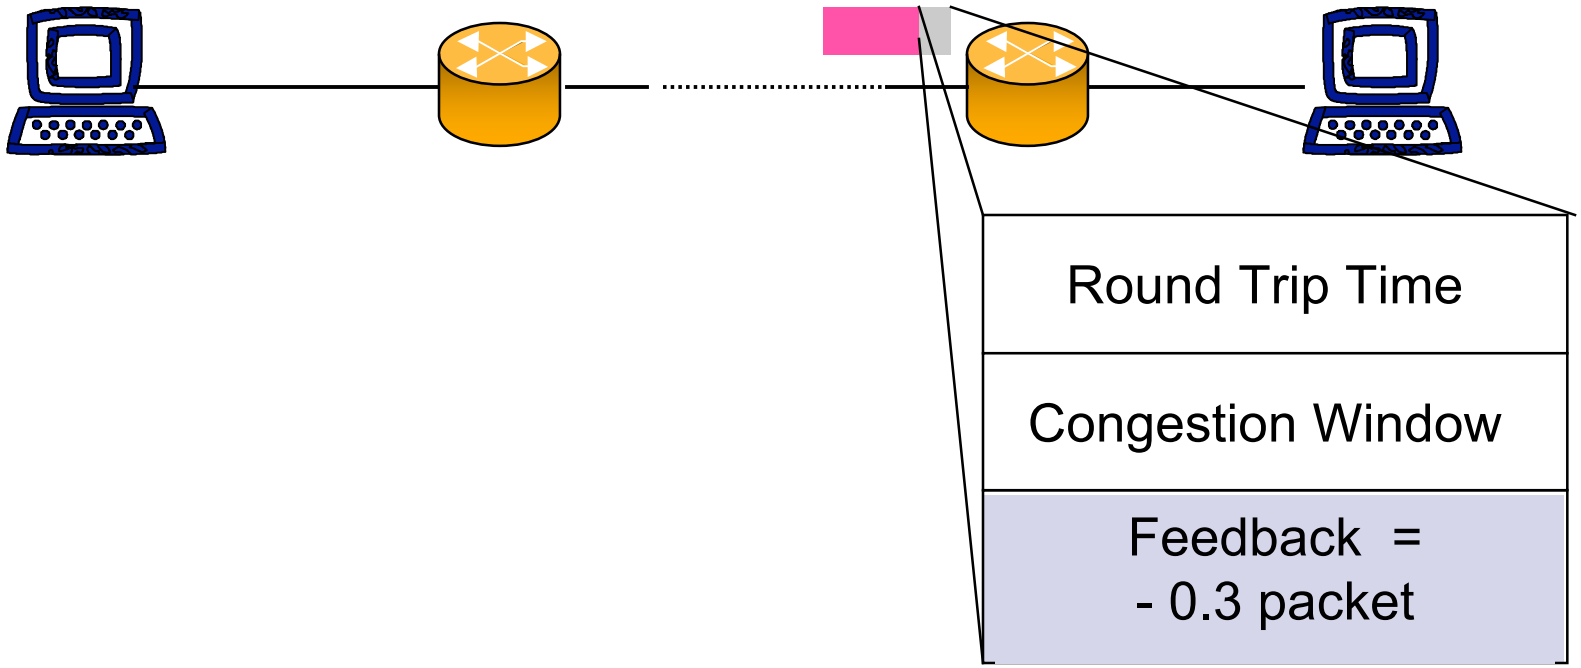

XCP: eXplicit Control Protocol

Katabi, Handley, Rohrs, Sigcomm 2002

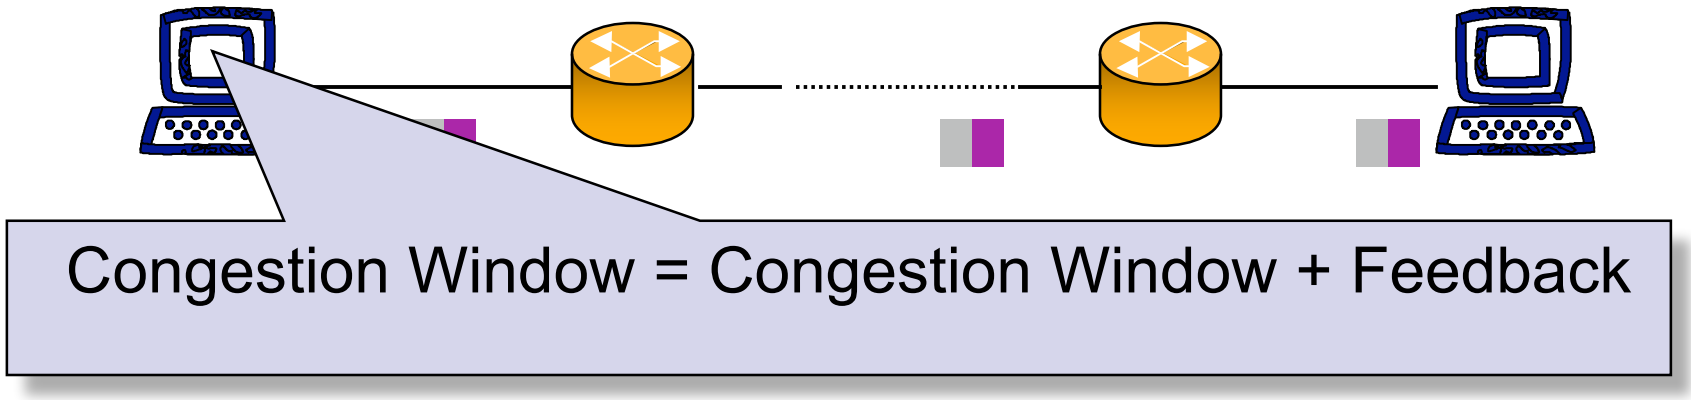

Routers compute feedback without any per-flow state

XCP vs. TCP

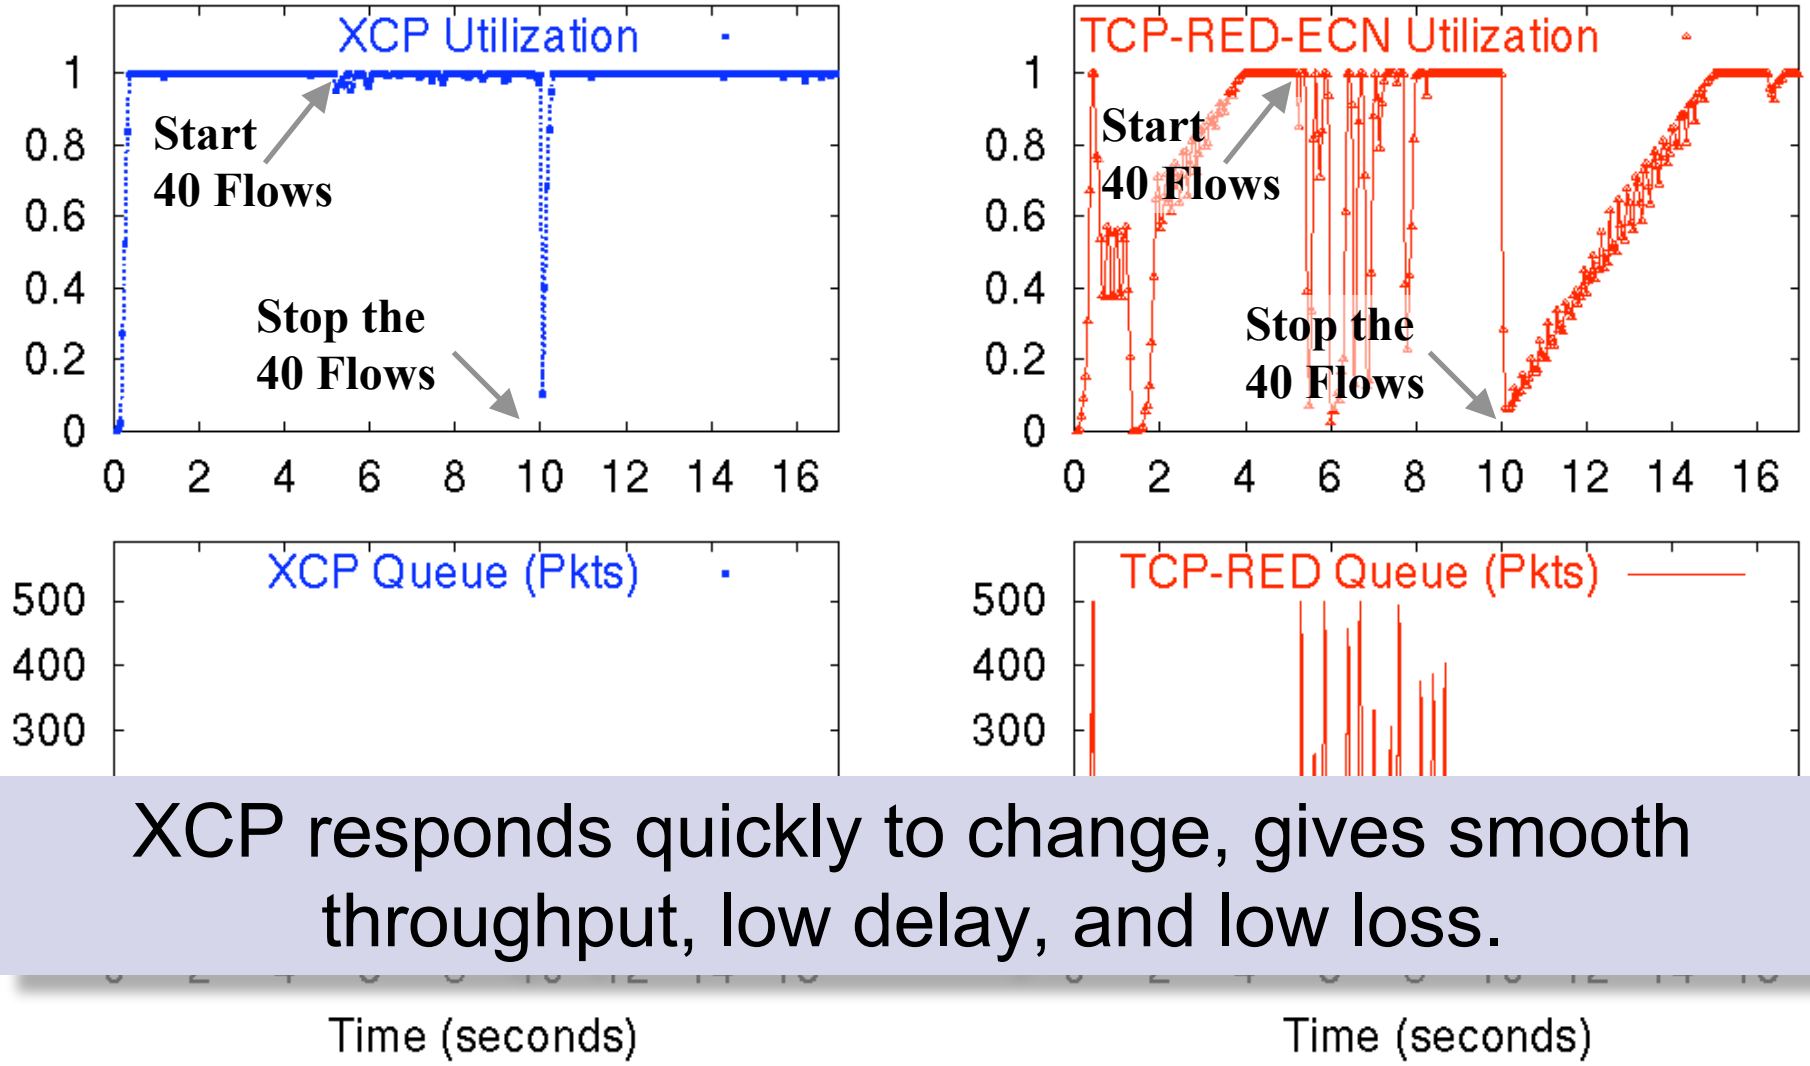

XCP responds quickly to change, gives smooth throughput, low delay, and low loss.

So why isn't everyone doing it?

- XCP was intended as a *blue-sky* idea to see what was possible.
  - Needs all the routers on the path to play.
  - Lots of bits in packet headers.
  - A couple of multiplies and a few adds per packet.

- Need phase 2: *Can we make it economically viable*?
  - Reduce costs without destroying benefits.
  - Enable incremental benefit with incremental deployment.

Plenty of Ideas

- High-speed TCP (S. Floyd)
- Scalable TCP (T. Kelly)
- FAST (S. Low)
- H-TCP (D. Leith)
- Bic-TCP (I. Rhee)
- XCP (Katabi)
- Re-feedback (Briscoe)
- VCP (Xia, Subramanian)
- Work on router buffer sizing (Appenzeller, McKeown, Wischik)

- Need a forum for evaluation and consensus that includes both researchers and equipment vendors.
  - IETF is not terribly good at this.

Internet Congestion Control Research Group

- Forum for discussion and evaluation of existing congestion control ideas, with the goal of reaching a consensus on how to move forward.
  - Researchers, vendors, operators needed to be successful.
- Influence the long-term plans of the IETF.

- Proposed charter:
  - http://nrg.cs.ucl.ac.uk/mjh/iccrg
- Mailing list:
  - http://oakham.cs.ucl.ac.uk/mailman/listinfo/iccrg