Congestion Control for High Bandwidth-Delay Product Networks

Presented by: Emad Shihab

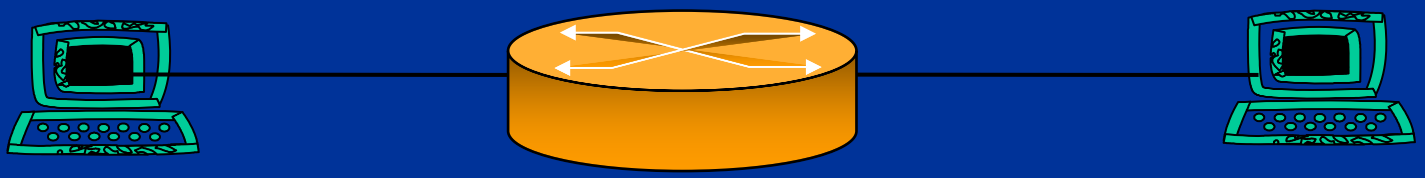

University of Victoria
British Columbia • Canada

Overview

- **Introduce the problem**
- **The zoom out view of XCP (what the protocol achieves)**
- **The zoom in view (how the protocol achieves this)**
- **Performance evaluation**
- **Conclusion**

University of Victoria
British Columbia • Canada

The problem!

- TCP's AIMD policy limits its ability to acquire spare bandwidth to one packet per RTT (in congestion avoidance stage)
- Even in slow start, short TCP flows (web flows) cannot make use of high bandwidth links

University of Victoria
British Columbia • Canada

The problem cont'd…

- Also, TCP cannot maintain stability over high-capacity or large-delay links
  - TCP becomes oscillatory and prone to instability
- Since TCP's throughput is inversely proportional to RTT, fairness becomes a problem for flows with longer RTTs

Solutions?.......

University of Victoria
British Columbia • Canada

The problem cont'd

Shown analytically in [Low01] and via simulations

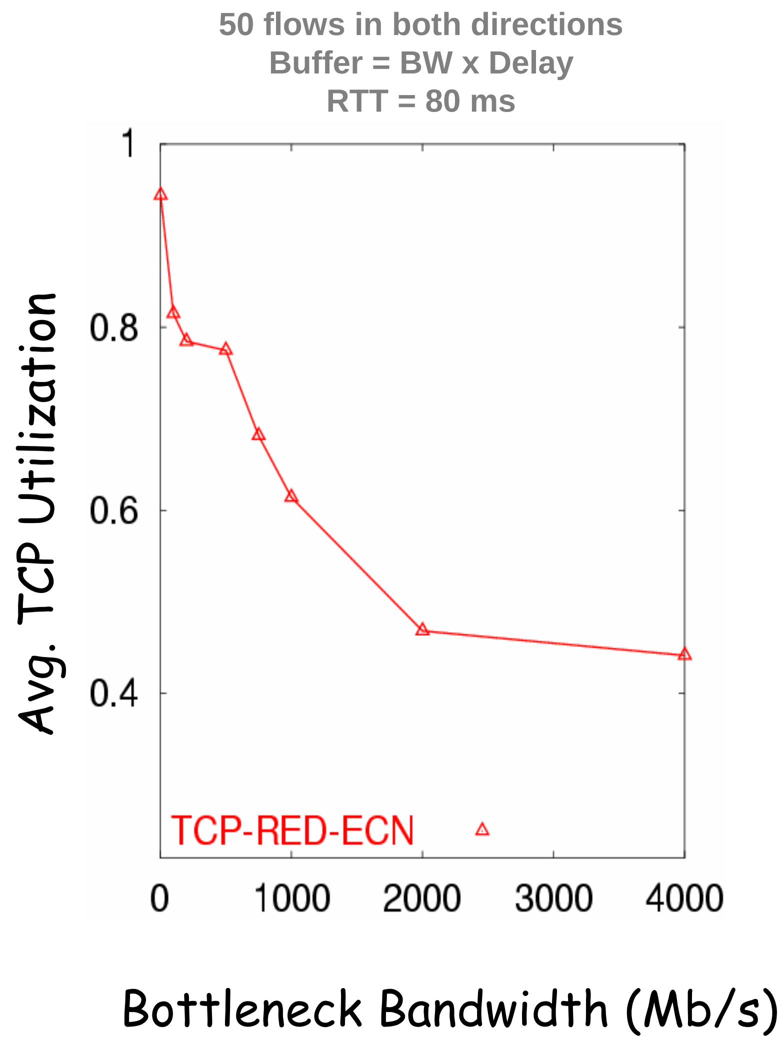

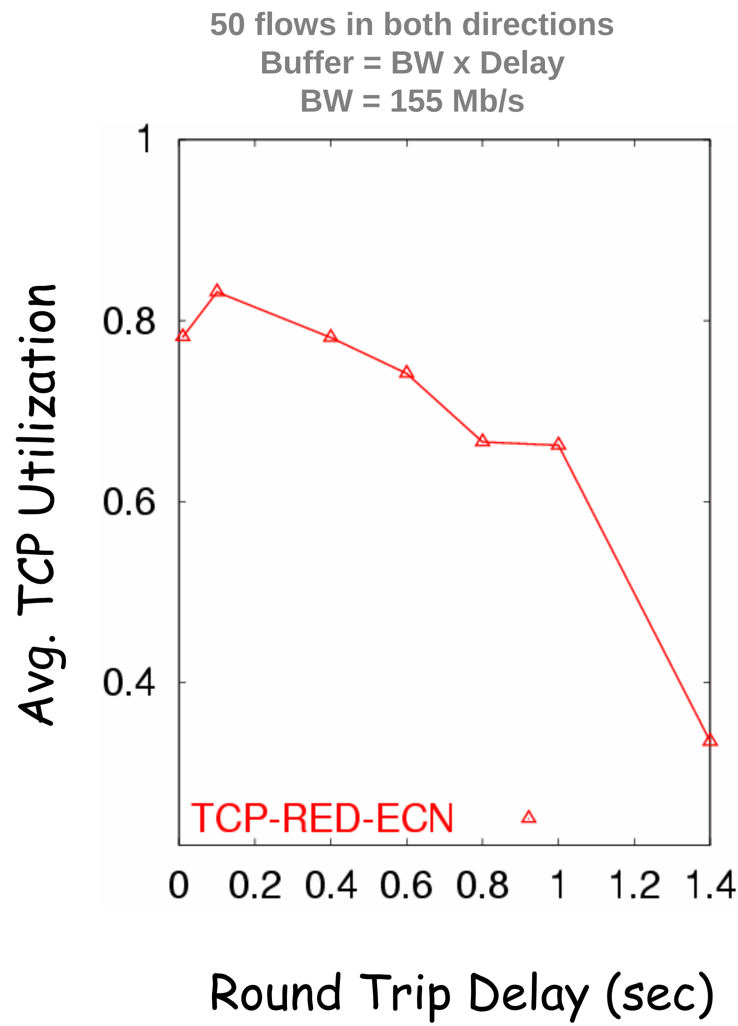

University of Victoria
British Columbia • Canada

XCP: eXplicit Control Protocol

- Main idea: Routers inform senders about *degree* of congestion at the bottleneck, not just if there is congestion or not

- Also, *decouple* congestion (utilization) control from fairness control

- In TCP, *both* utilization *and* fairness are controlled through the AIMD mechanism

XCP: eXplicit Control Protocol

- To control congestion, use MIMD for fast response
- To control fairness, use AIMD which converges to fairness

University of Victoria
British Columbia • Canada

XCP: eXplicit Control Protocol cont'd…

- Advantages of doing this:
  - Better utilization, use MIMD to achieve good utilization
  - Better fairness, use AIMD to achieve fairness amongst flows
  - Scalable, no per-flow info needed at router
  - Better support for differentiated service
  - Can distinguish between error losses and congestion losses
  - Can detect misbehaving sources

XCP congestion header

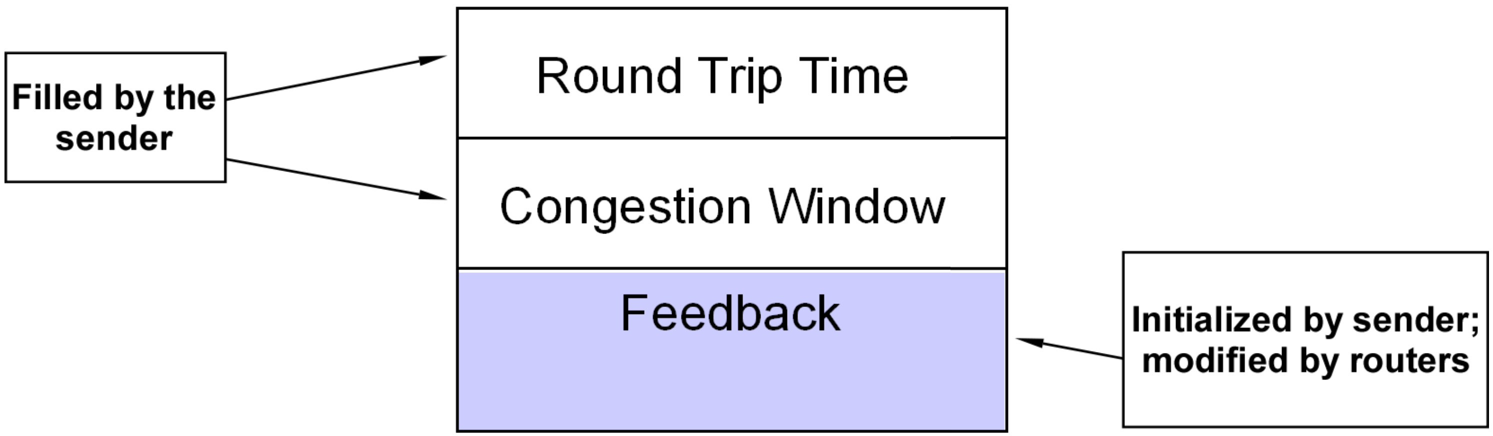

University of Victoria
British Columbia • Canada

XCP: Framework

- Senders maintain their cwnd and rtt and communicate this to routers

- Routers monitor input traffic and depending on the link, tell flows to increase or decrease cwnd
  - More congested routers later in the path can further reduce the cwnd/feedback

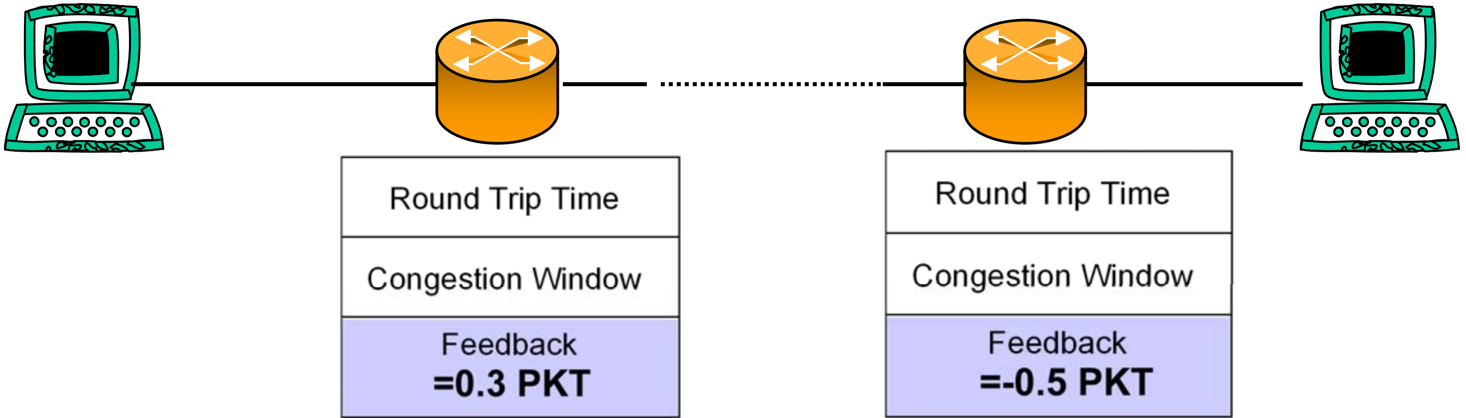

University of Victoria
British Columbia • Canada

XCP: Framework

Advance computer networks pres.

Introduction

The zoom out view

The zoom in view

Performance evaluation

Conclusion

- Feedback is divided between flows based on their cwnd and rtt so that the system converges to fairness

- Receivers return the feedback to the sender in an acknowledgment packet and senders update their cwnd accordingly

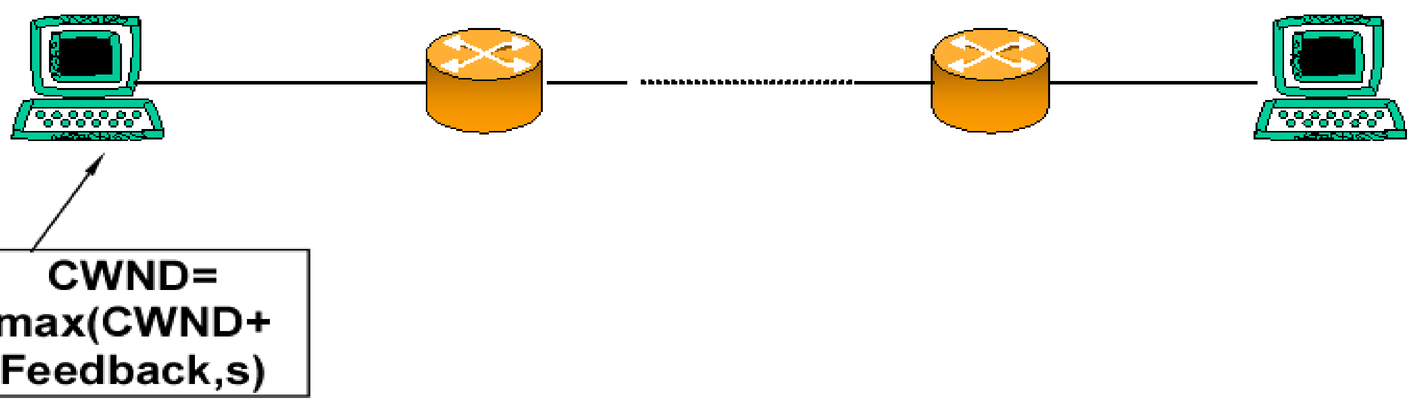

University of Victoria
British Columbia • Canada

Advance computer networks pres.

Introduction

The zoom out view

The zoom in view

Performance evaluation

Conclusion

XCP: The efficiency controller (EC)

- The purpose of EC is to maximize utilization while minimizing drop rate and persistent queues

- The EC computes desired increase/decrease based on

$$\Phi = \alpha * \text{avg. RTT} * \text{spare BW} - \beta * \text{pers. queue}$$

- For system to be stable,

$$0 < \alpha\ \pi/5.7 \text{ and } \beta = \alpha\text{\^{}}2 * 1.41$$

Q: Why do we need to consider queue here?

University of Victoria
British Columbia • Canada

XCP: The fairness controller (FC)

- You can notice, the EC does not deal with the issue of how we divide the feedback among packets. This is a fairness issue
- The FC relies on AIMD to achieve/converge fairness
- Follows a per packet policy that states:
  - If $\Phi > 0$, allocate it so that that increase in thr. of all flows is the same
  - If $\Phi < 0$, allocate it so that decrease in thr. of a flow is proportional to its current thr.

University of Victoria
British Columbia • Canada

XCP: The fairness controller (FC) cont'd…

- What if $\Phi \sim 0$, then use bandwidth shuffling

- The idea is to simultaneously allocate and deallocate BW, such that total traffic does not change, yet thr. of each individual flow changes gradually to approach its fair share

$$h = \max(0, 10\% * input - |\Phi|)$$

- h is called the shuffled traffic

- This equation ensures that every RTT 10% of traffic is redistributed according to AIMD

University of Victoria
British Columbia • Canada

XCP: The fairness controller (FC) cont'd…

- To be able to apportion the feedback correctly, we need to estimate number of flows

$$N=\Sigma_{pkt}\ 1/(d^*cwnd_i/rtt_i)$$

- Now we need to calc. the per-packet feedback

$$H\_feeback = p_i - n_i$$

University of Victoria
British Columbia • Canada

XCP: The fairness controller (FC) cont'd…

- When Φ > 0, then positive feedback is given by:

$$p_i = \text{change in cwnd/ \# of pkts router sees per d}$$
$$= \text{thr. RTT / ((cwnd/s)/RTT)}$$

- Where thr. is a const. given as

$$\text{thr.} = \text{agg. traffic rate / N}$$
$$= (h + \max(\Phi,0)) / N$$

- Therefore,

$$p_i = \text{thr.*RTT\^{}2*}s_i / cwnd_i$$

University of Victoria
British Columbia • Canada

XCP: The fairness controller (FC) cont'd…

- When $\Phi < 0$, then negative feedback is given by:
- n_i = change in cwnd / # of pkts router sees per d

  $= \Delta cwnd / ((cwnd/s)/RTT)$

- Here however, $\Delta cwnd \propto cwnd$ (multiplicative dec.)
- In fact $\Delta cwnd$ = thr. cwnd
- Where thr. is a const. given as

$$thr. = agg.\ traffic\ decrease / N = (h + \max(-\Phi, 0)) / N$$

- Therefore,

$$n_i = thr.{*}RTT{*}s_i$$

Efficiency when high BW and high Delay

Utilization as a function of Bandwidth

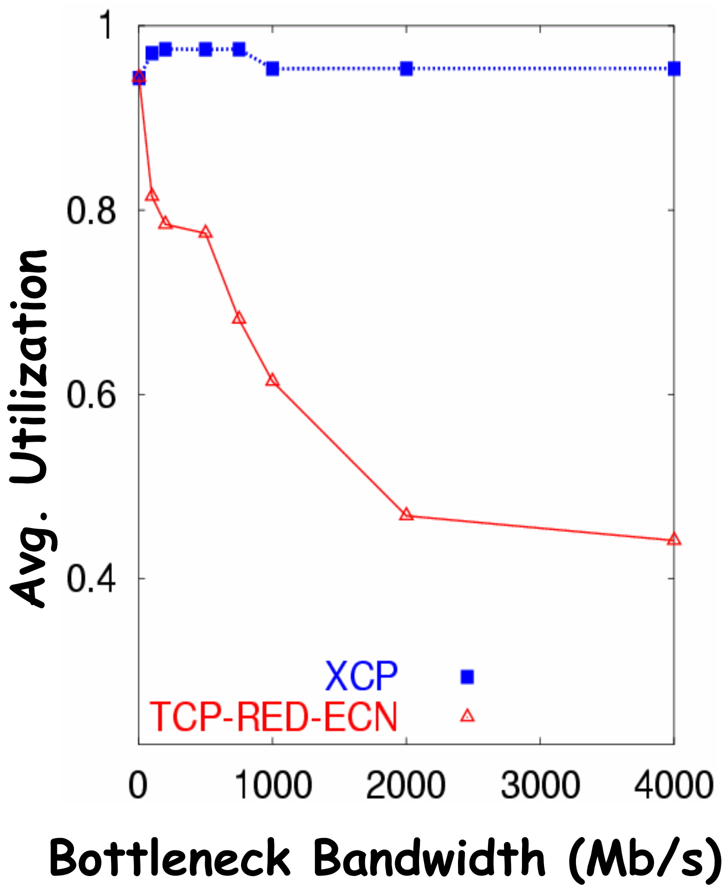

Utilization as a function of Delay

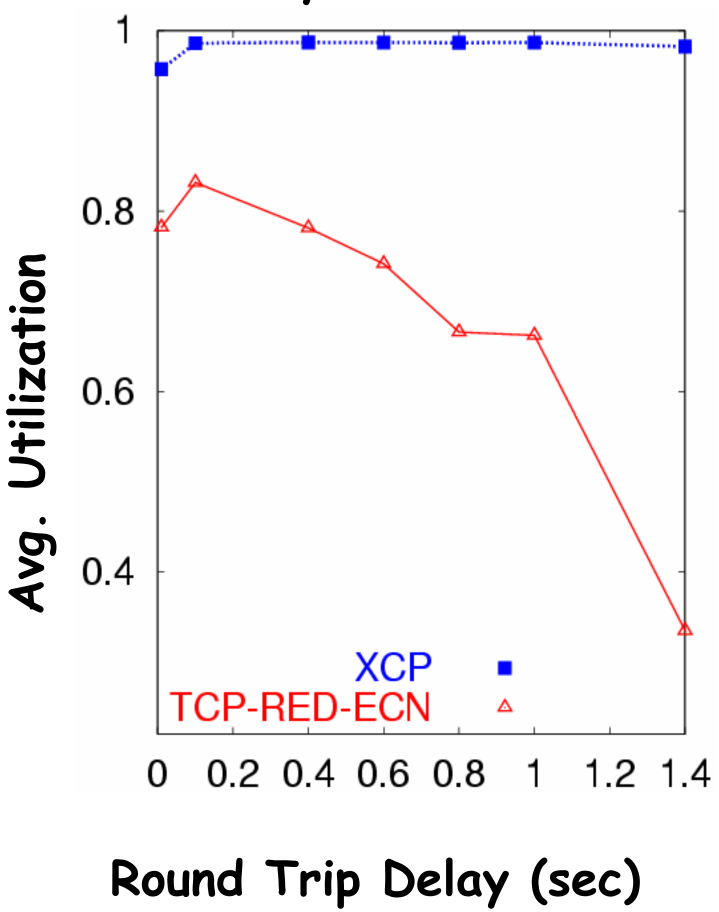

University of Victoria
British Columbia • Canada

XCP's response time

Advance computer networks pres.

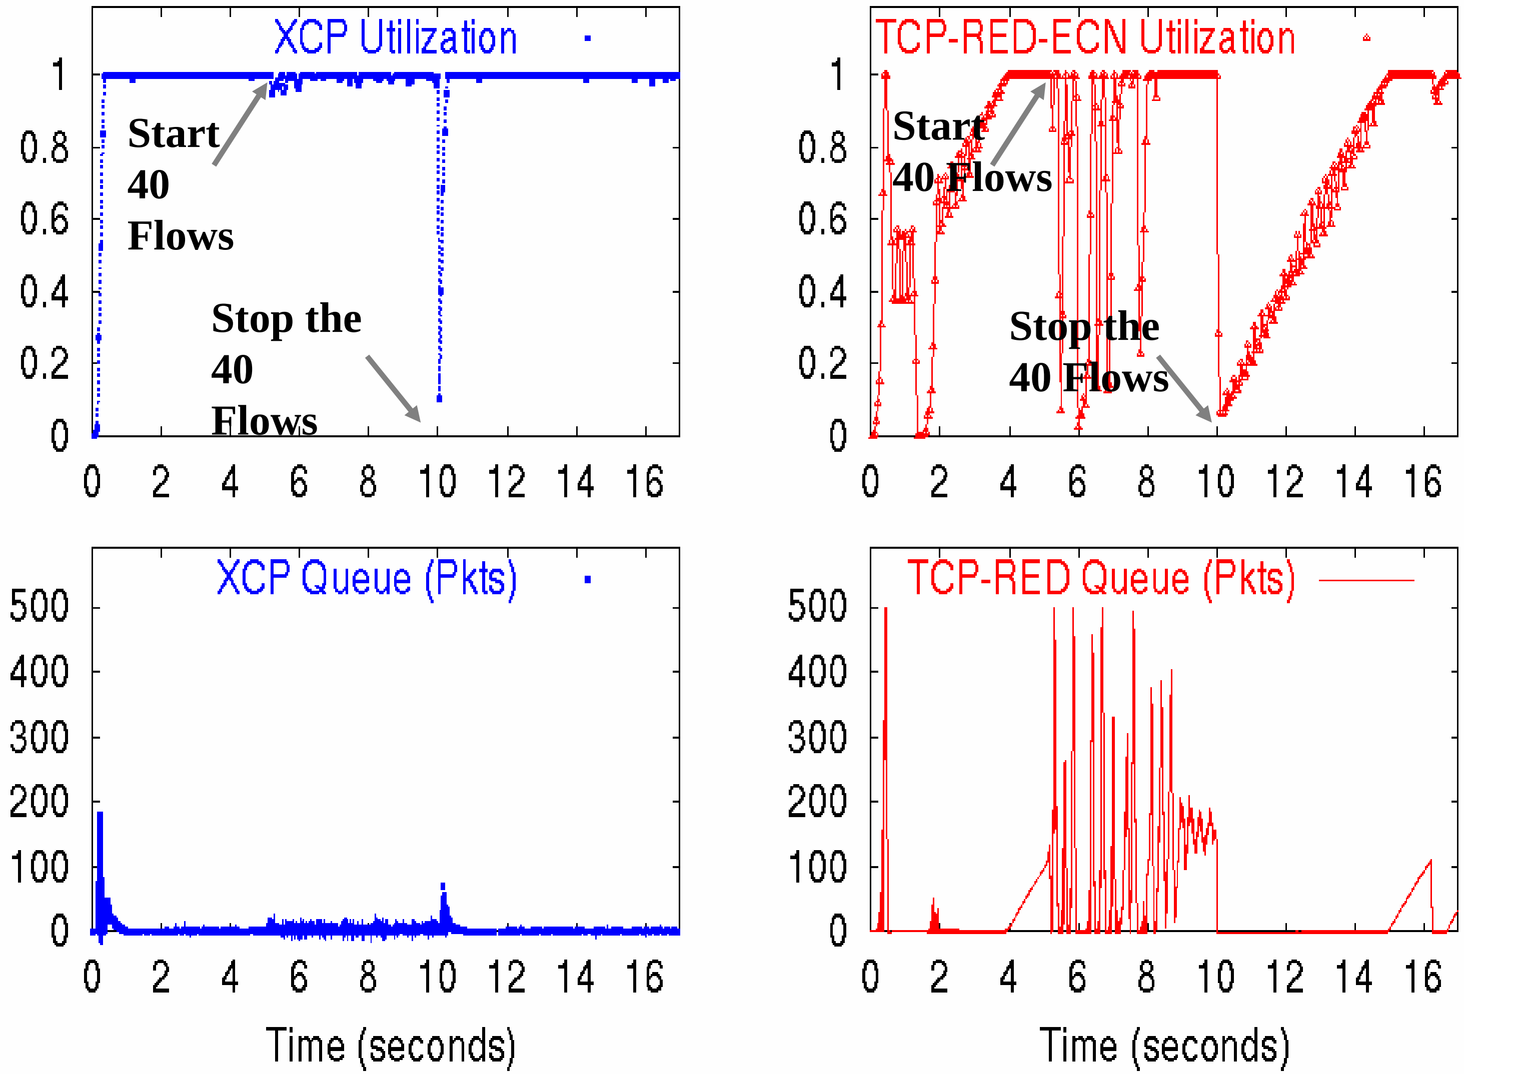

University of Victoria
British Columbia • Canada

XCP's performance in presence of short web like flows

Advance computer networks pres.

Average Utilization

Average Queue

Drops

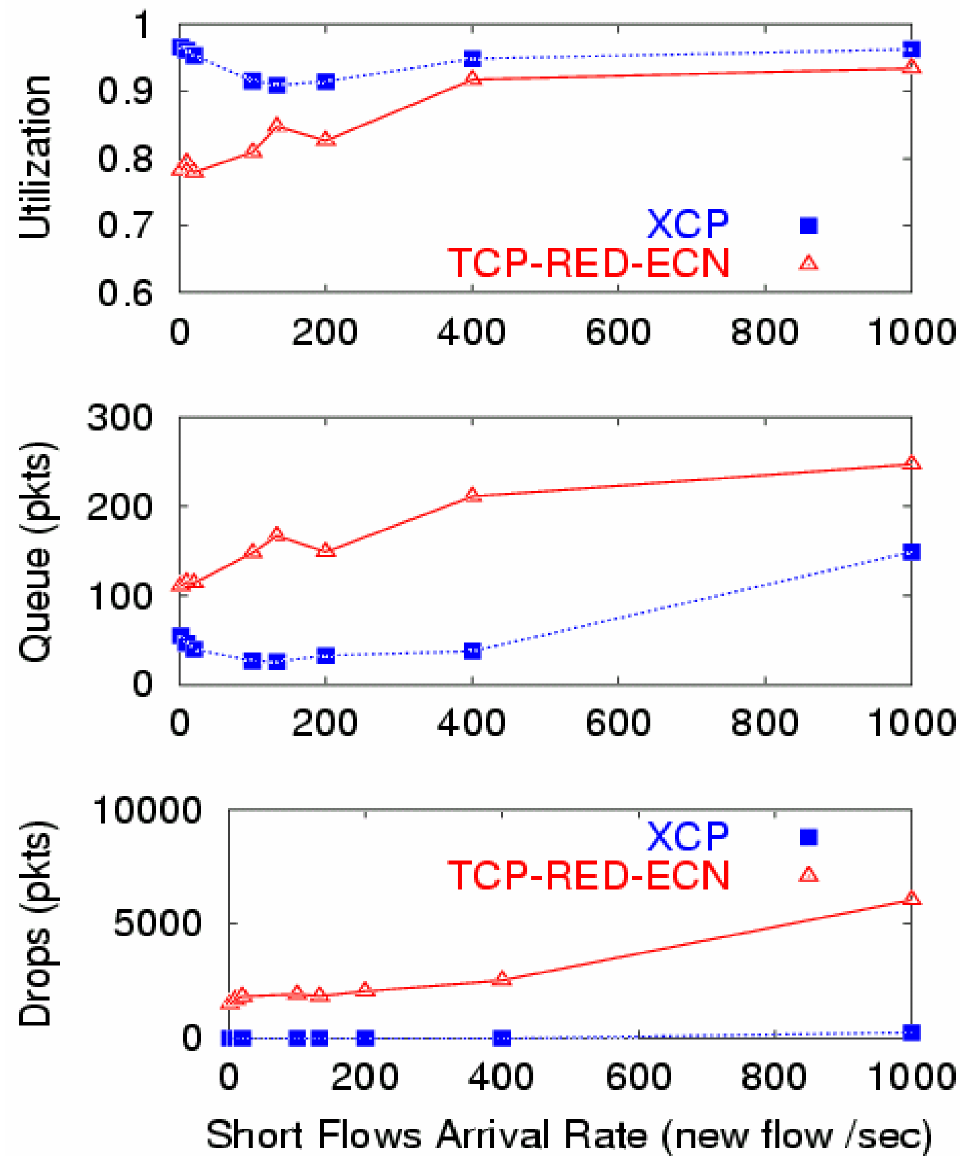

University of Victoria
British Columbia • Canada

XCP's fairness compared with TCP

Advance computer networks pres.

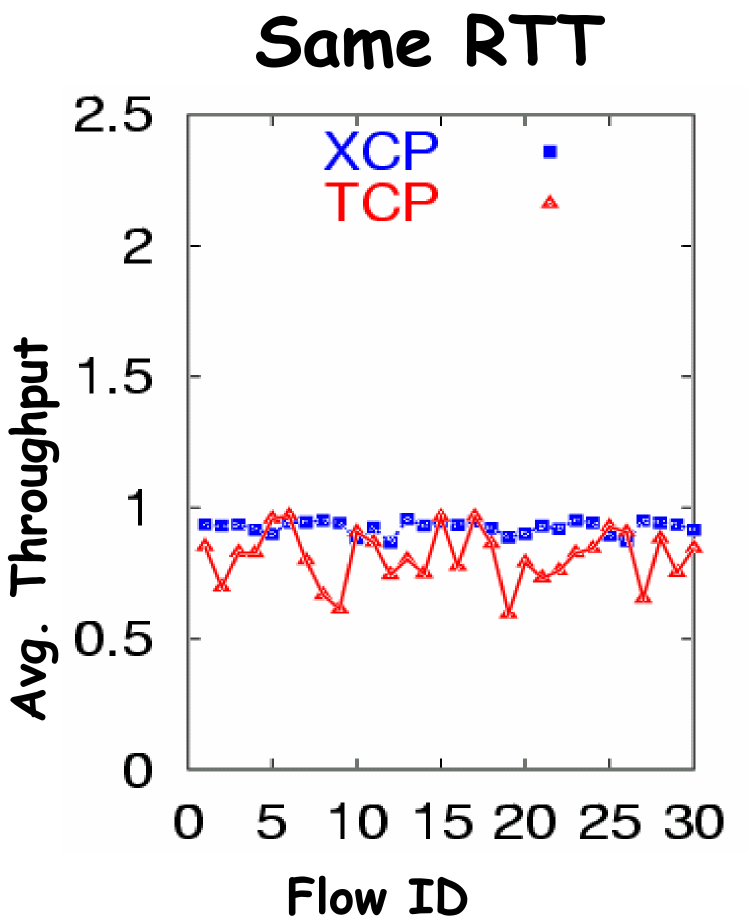

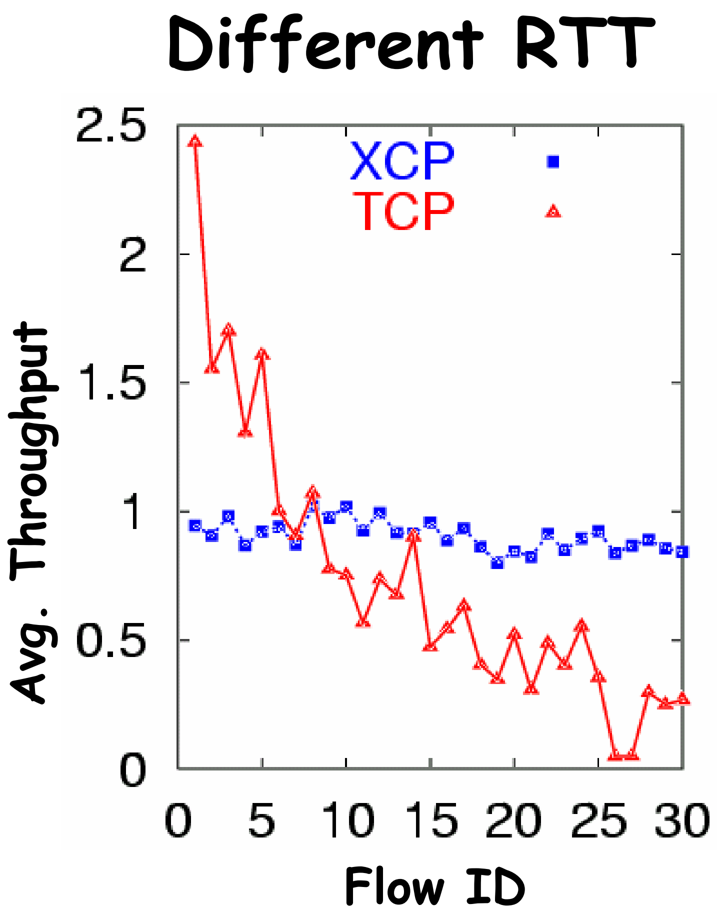

(RTT is 40 ms → 330 ms)

University of Victoria
British Columbia • Canada

Conclusion

- XCP is efficient for large bandwidth and delay
- This is because it decouples utilization and fairness
  - Uses MIMD for congestion control which can grab/release large bandwidth quickly
  - Uses AIMD for fairness which converges to fair bandwidth allocation
- XCP also proved to be effective for short, web like flows

University of Victoria
British Columbia • Canada

Questions or ideas?

Ideas...suggestions. ..comments?