Two bits are enough

Ihsan Ayyub Qazi*, Lachlan Andrew# and Taieb Znati*
*University of Pittsburgh, Pittsburgh, PA 15260, USA
#Swinburne University of Technology, Hawthorn, Vic 3122, Australia

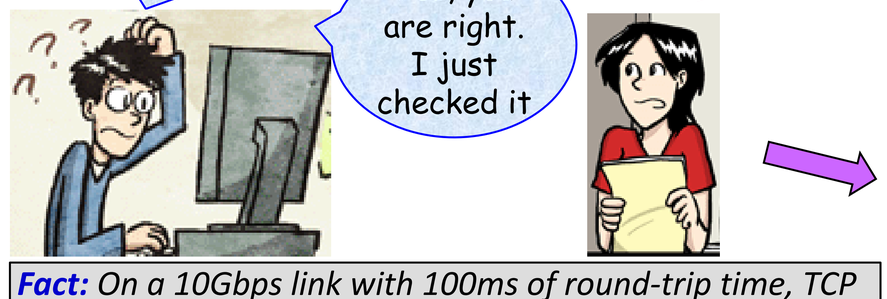

Fact: *On a 10Gbps link with 100ms of round-trip time, TCP takes ~1hr 45mins to recover from a single packet loss*

2. Introduction

- TCP performs poorly on high speed, long distance links
- End-host solutions (e.g., SACK, FAST) have fundamental performance limitations
- Network based solutions (e.g., XCP, RCP) have serious deployment and complexity issues
- This work proposes ***Binary Marking Congestion Control*** (BMCC) that achieves all the desirable goals

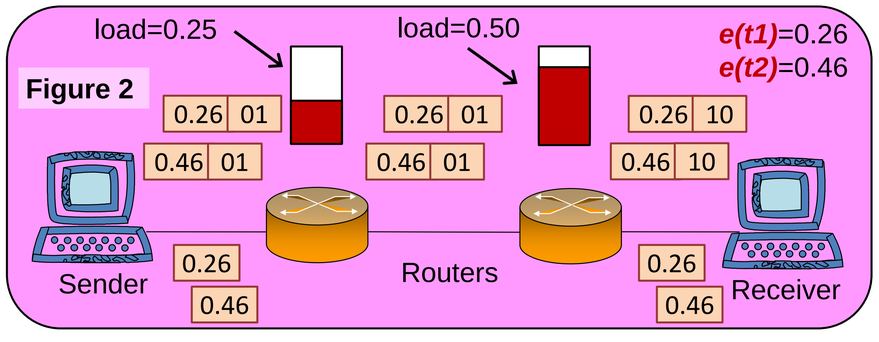

Figure 2

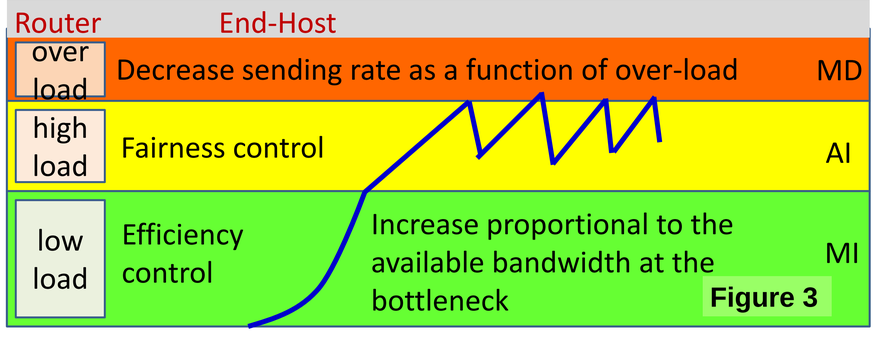

Figure 3

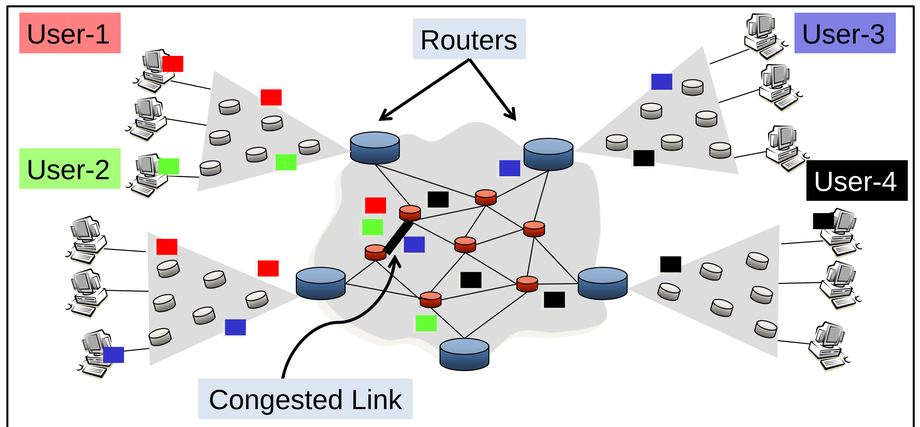

Figure 1: A network with 4 users sharing a congested link.
Problem: *How to share the link capacity amongst users?*

4. ADPM (basic idea, Figure 2)

- Interpret the 16-bit IPid field in the IP header as a number, ***i***, in [0,1]
- Router marks a packet if ***link load > i***
- Receiver maintains a load estimate, ***e***, sets it to ***i*** if the packet is marked and ***e < i*** or unmarked and ***e > i***
- Sources adjust their rates based on ***e*** (see Figure 3)

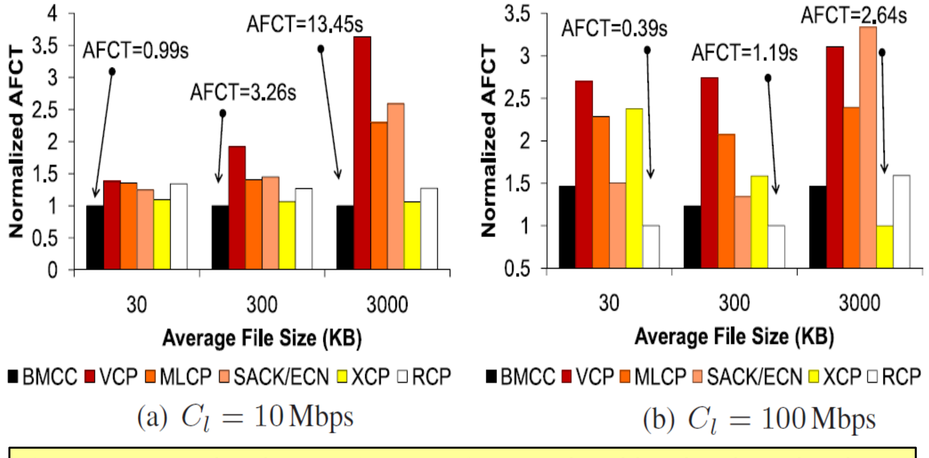

Figure 4: Download times (or AFCT) normalized by the download time of the best performing scheme (see [1] for details)

1. Wish List for a Solution

- Efficient and fair use of network links
- Low queues and negligible loss rate
- Short download times
- Deployable

3. BMCC

- Routers mark packets (using the **A**daptive **D**eterministic **P**acket **M**arking scheme) to explicitly tell sources about the link load, using the existing two IP ECN bits
- Sources increase/decrease sending rate based on the link load

5. Contributions & Results

- Developed analytical models and conducted extensive ns2 simulations
- BMCC improves download times by up to 4x over VCP, 3x over TCP+RED/ECN, 2x over XCP, 1.5x over RCP (see Figure 4)
- BMCC achieves efficiency and fairness with low queues and negligible loss rate

6. Future Work

- Improving fairness convergence
- Enabling BMCC over wireless networks

[1] Ihsan A. Qazi, Lachlan Andrew, Taieb Znati, "**Congestion Control using Efficient Explicit Feedback**" in *Proceedings of IEEE INFOCOM 2009*
[2] Ihsan A. Qazi, Lachlan Andrew, Taieb Znati, "**Two bits are enough**" in *Proceedings of ACM SIGCOMM 2008* (poster)

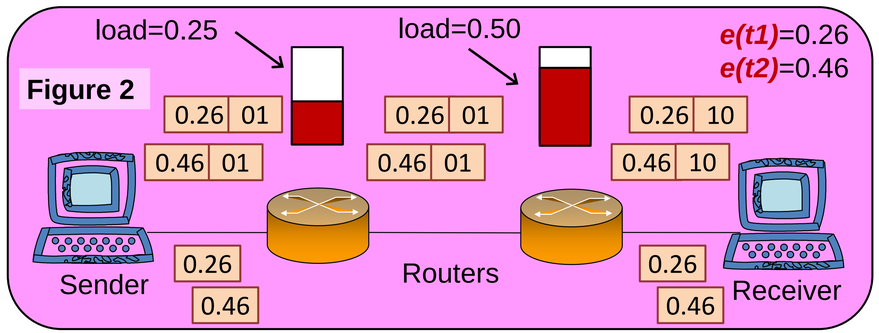
load=0.25
load=0.50
e(t1)=0.26
e(t2)=0.46
Figure 2
0.26 01
0.46 01
0.26 01
0.46 01
0.26 10
0.46 10
0.26
0.46
Sender
Routers
0.26
0.46
Receiver