Understanding the Impacts of Short-Term Dynamics on Long-Term Fairness of Competing High Speed TCP Flows: A Root-Cause Analysis

Balázs Sonkoly,
Tuan Anh Trinh and Sándor Molnár

High Speed Networks Laboratory
Department of Telecommunications and Media Informatics
Budapest University of Technology and Economics

email: sonkoly@tmit.bme.hu

Overview

- Motivation
- Goals
- Methodology
- Transient and equilibrium analysis of TCP versions
- Intra-protocol behavior
- Interaction of TCP versions
- Conclusion

Motivation

- Standard TCP (Reno): poor performance in high speed wide-area network environment
- High speed transport protocol proposals
  - modified congestion control
  - main principles
    - loss-based (e.g., HighSpeed TCP, Scalable TCP, BIC TCP)
    - delay-based (e.g., TCP Vegas, FAST TCP)
    - most recent proposals: combined algorithm (e.g., Compound TCP)
- Most important issues: operability and deployability

↓

? Fairness ?

Motivation (cont'd)

- Number of fairness metrics proposed (e.g., Jain's index, proportional fairness, etc.)
- Common aspect: all concern with the
  - long term average of the flows
  - stable & equilibrium performance
- Can not capture ***dynamic*** properties!

Goals

- Fairness analysis of loss-based high speed TCPs
- focusing on dynamic behavior
- Reveal the impacts of short term dynamics on fairness
- Explanation of the phenomena
- Investigating
  - ***MIMD*** mechanism (Scalable TCP)
  - ***AIMD-like*** mechanisms (HighSpeed TCP, BIC TCP)

Methodology

- Ns-2 simulation environment
- Different topologies
  - dumb-bell
  - simple parking-lot

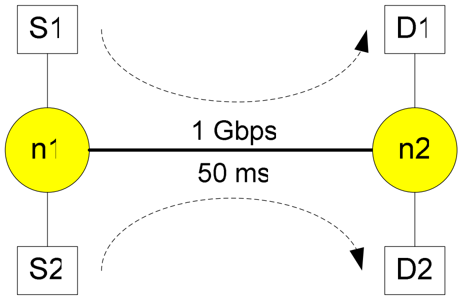

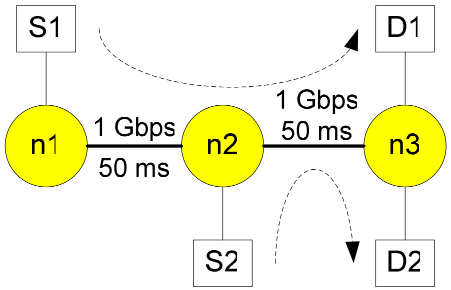

- Competition of high speed TCP flows
  - individual flows
  - single flow – traffic aggregate
- Dynamic behavior → different starting times

Initial behavior – saturation time of loss-based protocols

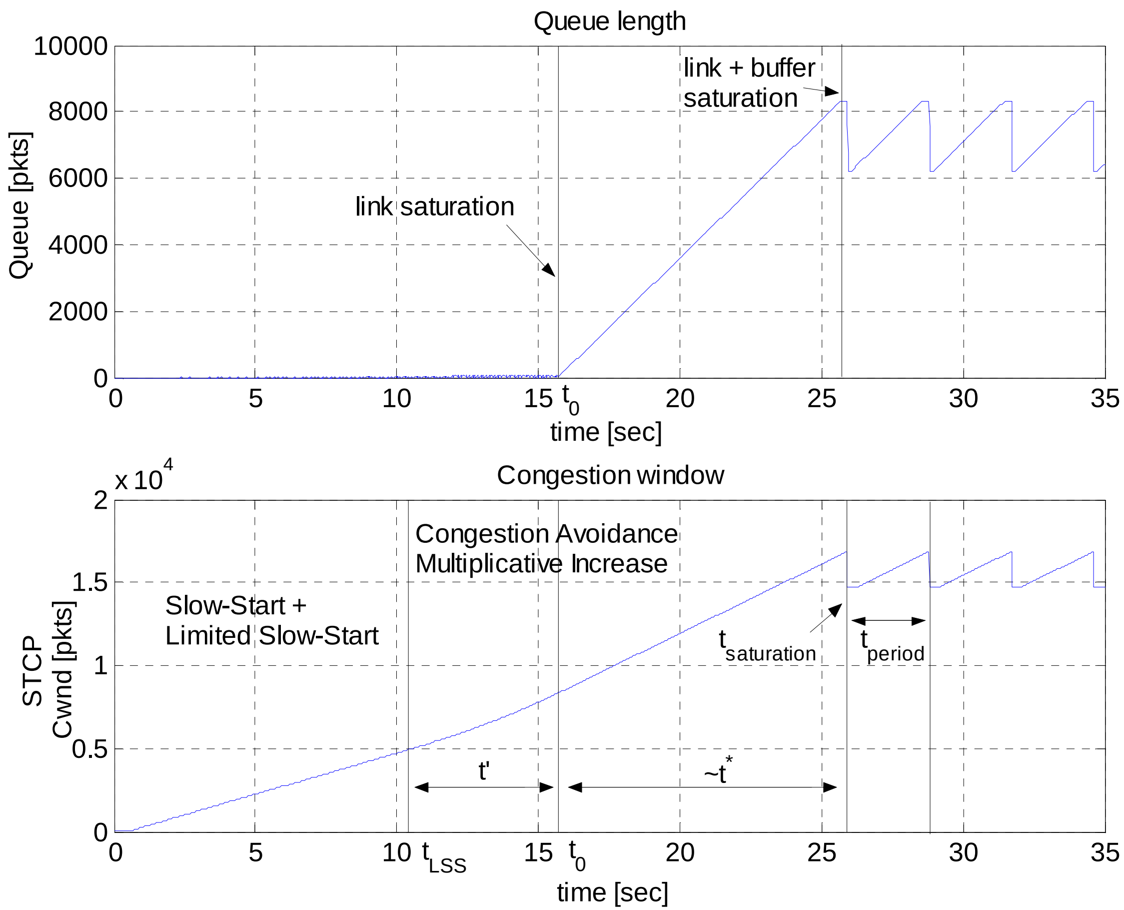

- Initial dynamics: significant role
- Saturation time
  - time till the first packet drop
  - can be derived analytically
- Example: STCP
  - Slow-Start
  - Limited Slow-Start
  - Multiplicative Increase
- After saturation time
  - periodic behavior
  - oscillation

Equilibrium behavior

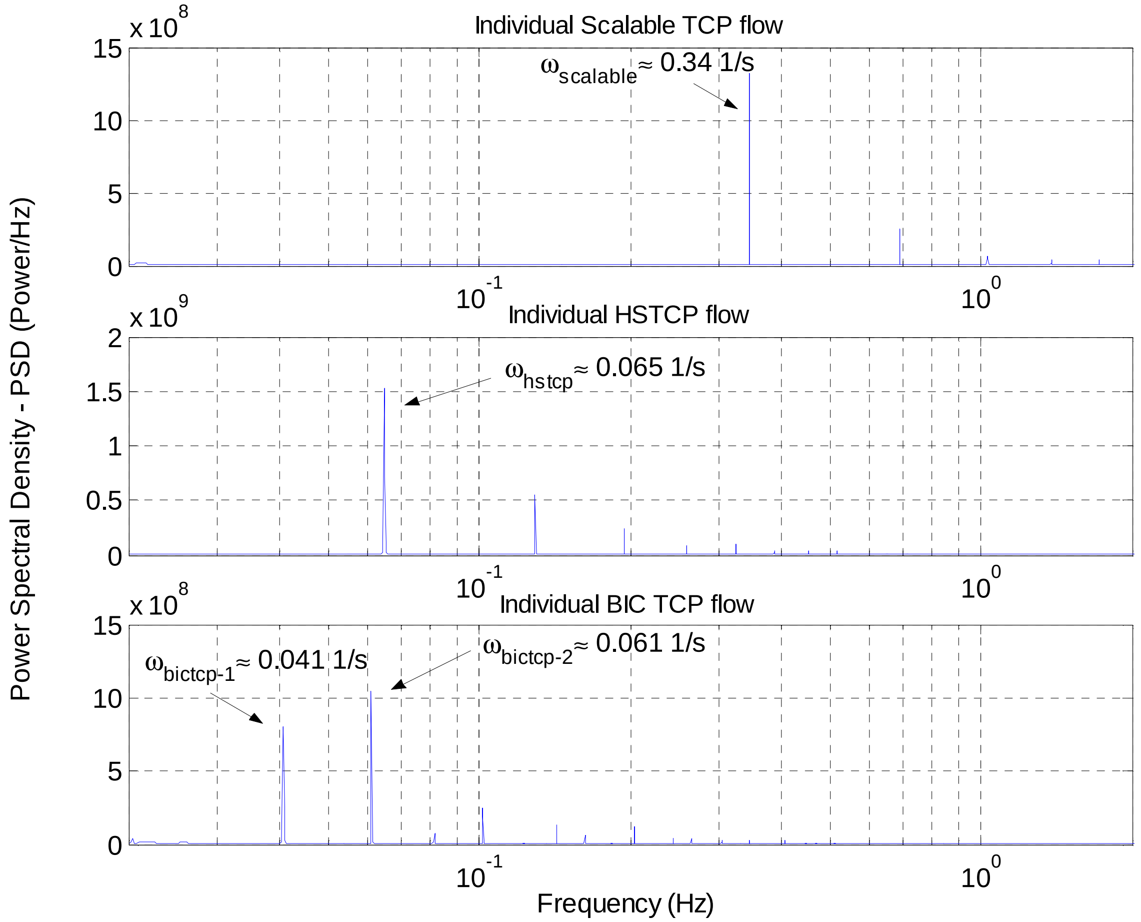

- Periodic behavior can be analyzed in the frequency domain
- main frequency spikes characterizing the equilibrium behavior
- significant differences

Intra-protocol behavior: MIMD

- Dumb-bell topology: two Scalable TCP flows
- Equilibrium behavior & fairness determined by starting time!

- delay: 50 sec
- delay: 15 sec

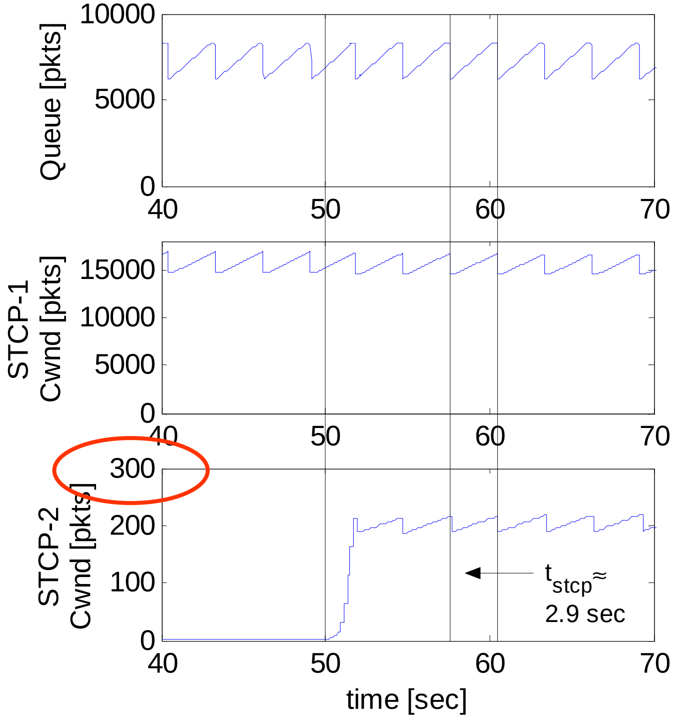

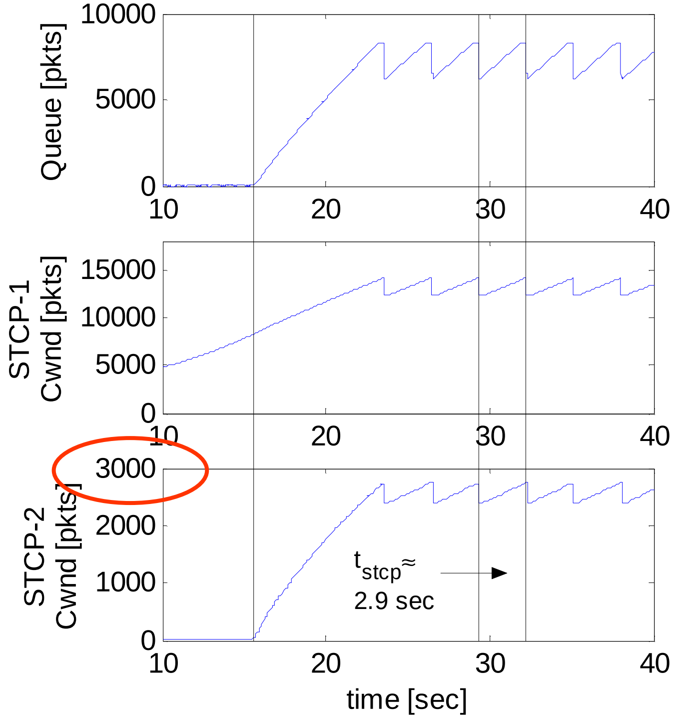

- MIMD mechanism → constant operating frequency
- Synchronized losses
- Equilibrium state depends on the length of the transient phase

Intra-protocol behavior: AIMD

- Dumb-bell topology: two HSTCP flows
- Fair behavior, but transient phase can be too long!

- delay: 50 sec

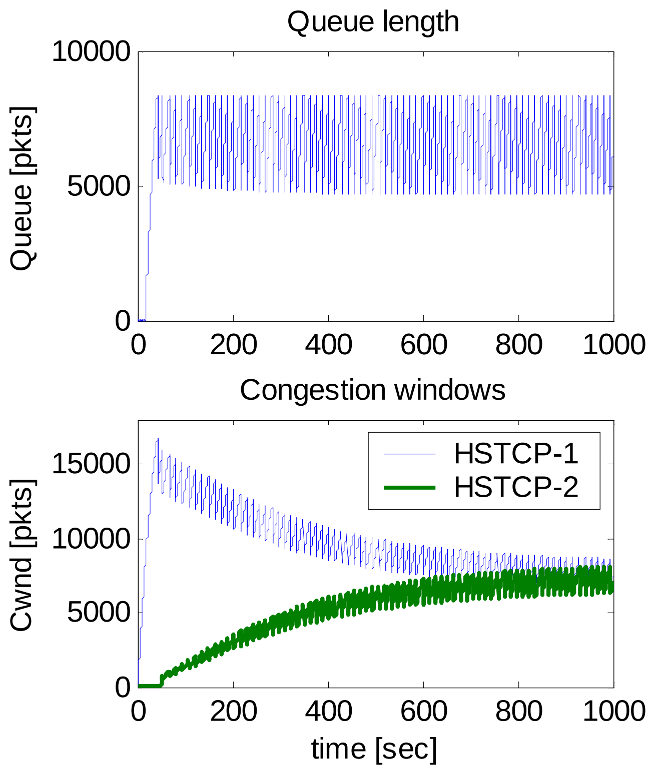

- delay: 15 sec

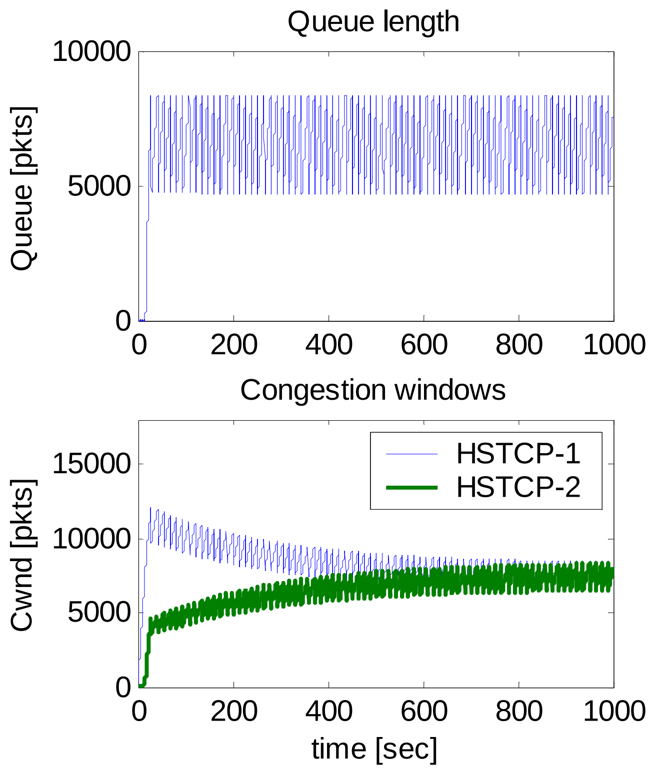

Interaction of TCP versions (1)

- Dumb-bell topology: Scalable TCP – HSTCP
- HSTCP enters later, **after** the **saturation time** of Scalable TCP

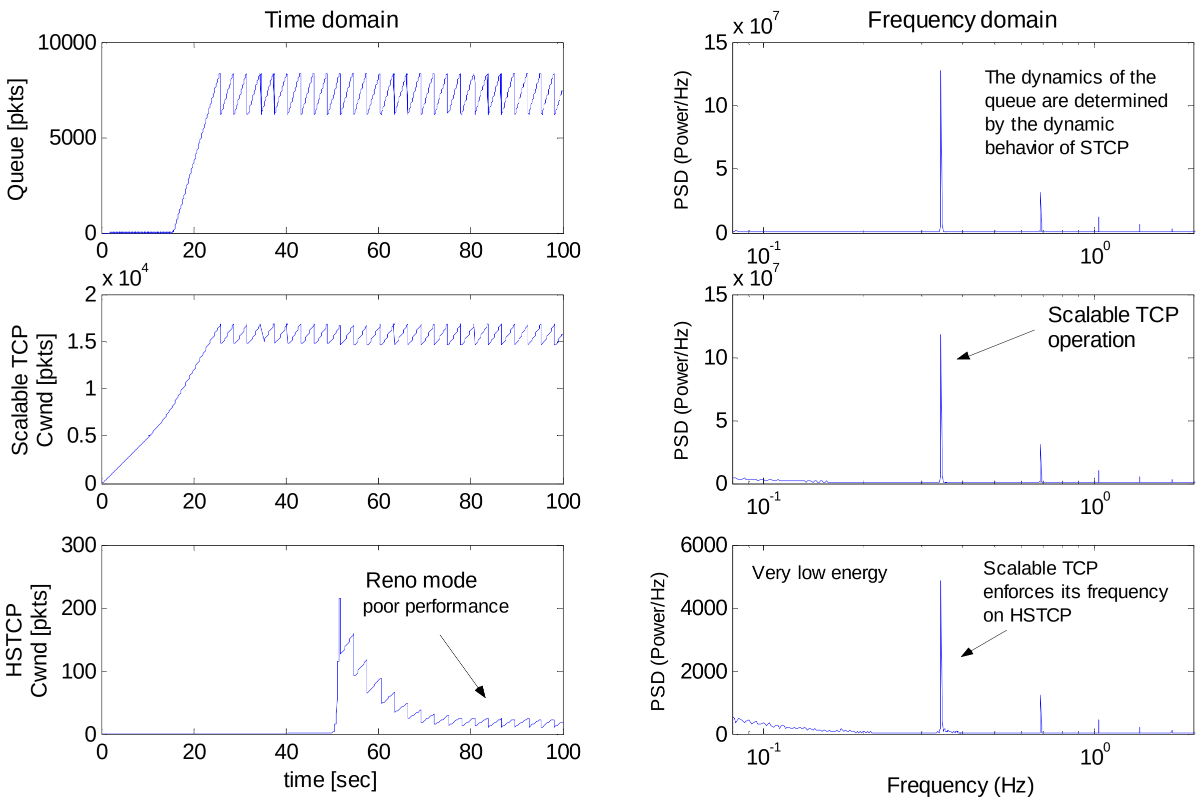

- ***Starvation!***
- Synchronized losses
- Frequency is determined by Scalable TCP
- HSTCP
  - incr./period < decr.
  - decreasing trend
- Equilibrium:
  - cwnd < Low_W
  - TCP Reno operating mode
  - poor performance

Interaction of TCP versions (2)

- Dumb-bell topology: Scalable TCP – HSTCP
- HSTCP enters first: better performance is expected ↔ but ***starvation***!

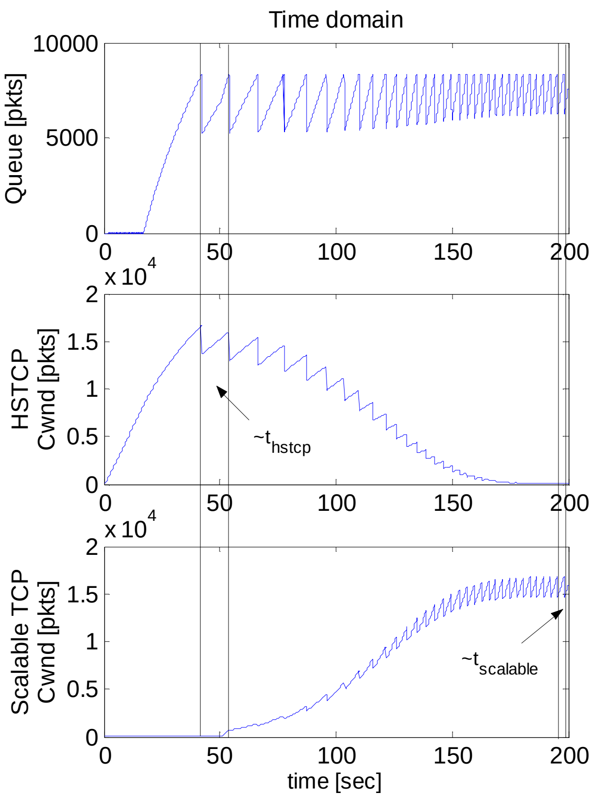

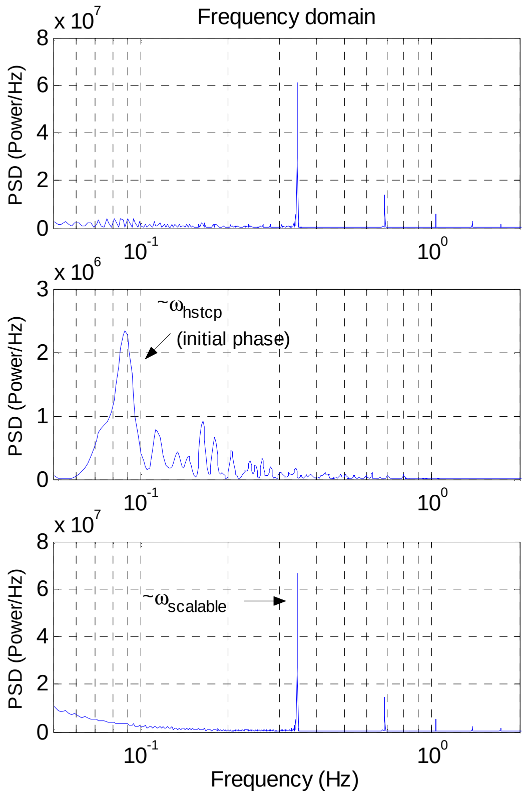

- Interesting phenomenon
- Increasing frequency
- $\omega_{hstcp} \rightarrow \omega_{scalable}$
- Network behavior will be determined by Scalable TCP
- HSTCP converges to the same equilibrium state (TCP Reno operating mode)

Interaction of TCP versions (3)

- Dumb-bell topology: HSTCP aggregate vs. single Scalable TCP flow
- 10 HSTCP flows and a later entering Scalable TCP

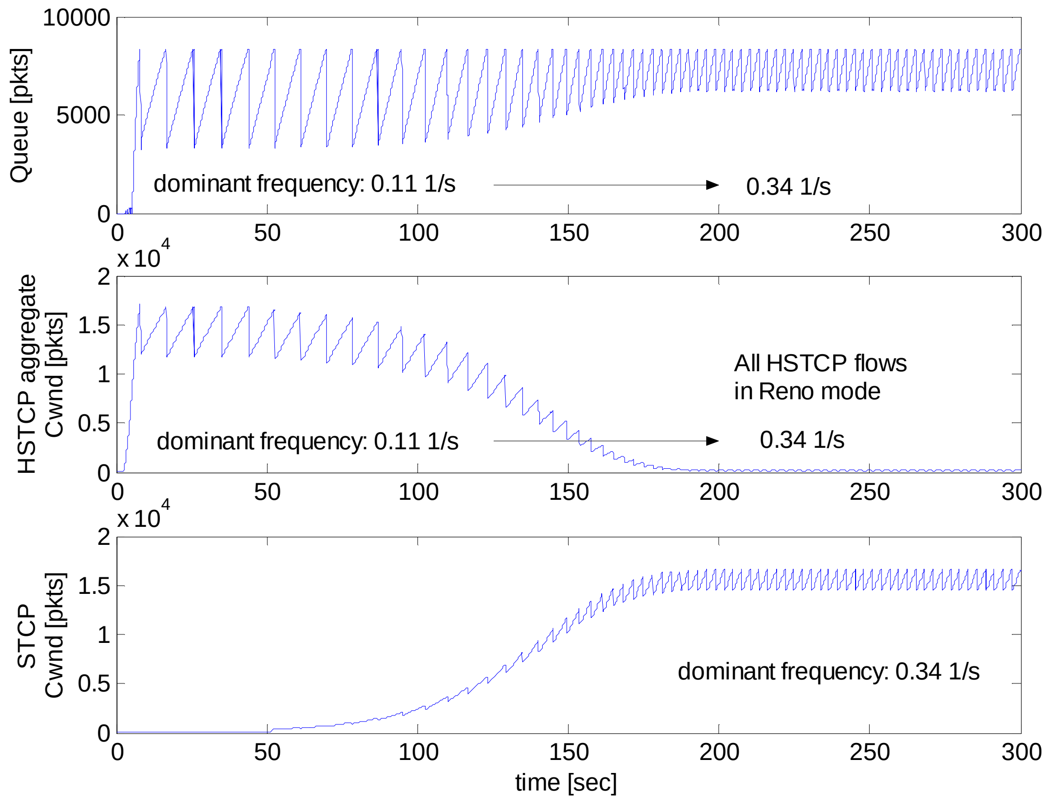

- Similar phenomenon
- Scalable TCP enforces its frequency on HSTCP aggregate
- All HSTCP flows fall back to Reno mode!

Interaction of TCP versions (4)

- Dumb-bell topology: BIC TCP aggregate vs. single Scalable TCP flow
- 10 BIC TCP flows and a later entering Scalable TCP

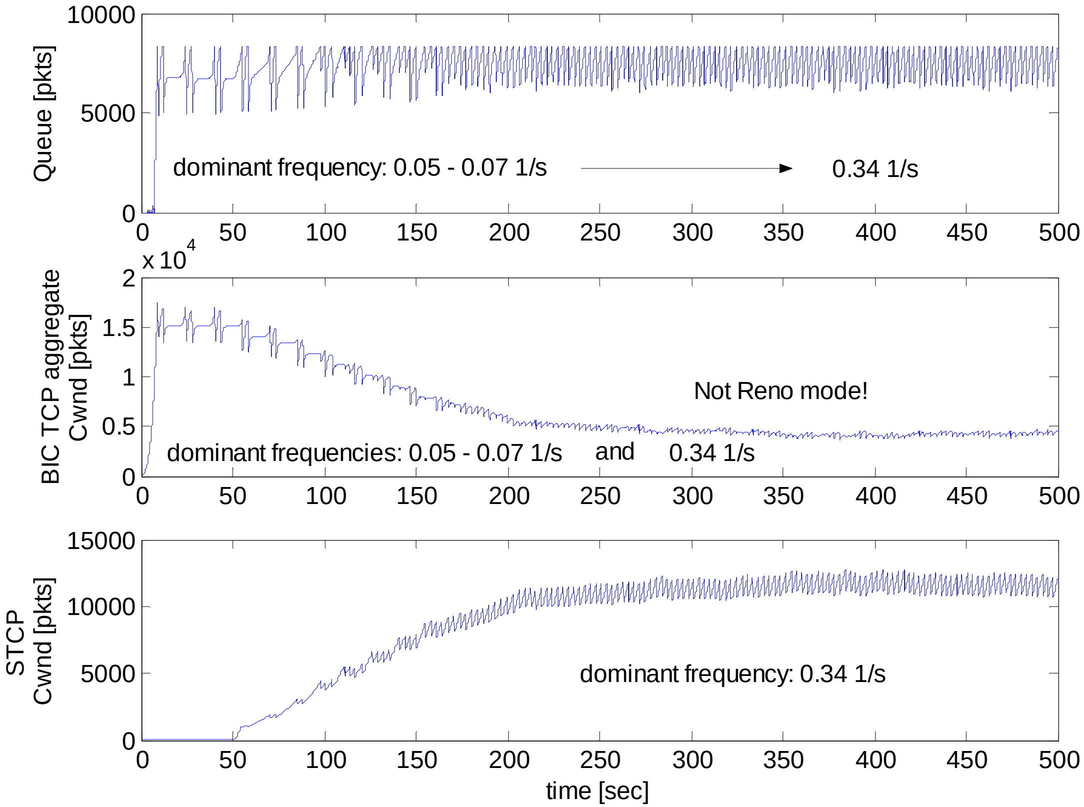

- Slightly better performance
- Not Reno mode!
- BIC TCP aggregate exhibits the dominant frequency of Scalable TCP as well
- Individual Scalable TCP flow dominates the network

Interaction of TCP versions (5)

- Simple parking-lot topology: single flows
- AIMD-like flow on the longer path
  - poor performance in all scenarios
  - Reno mode!
- MIMD flow on the longer path
  - Scalable TCP can push down the other flow
  - AIMD flow is always starved!

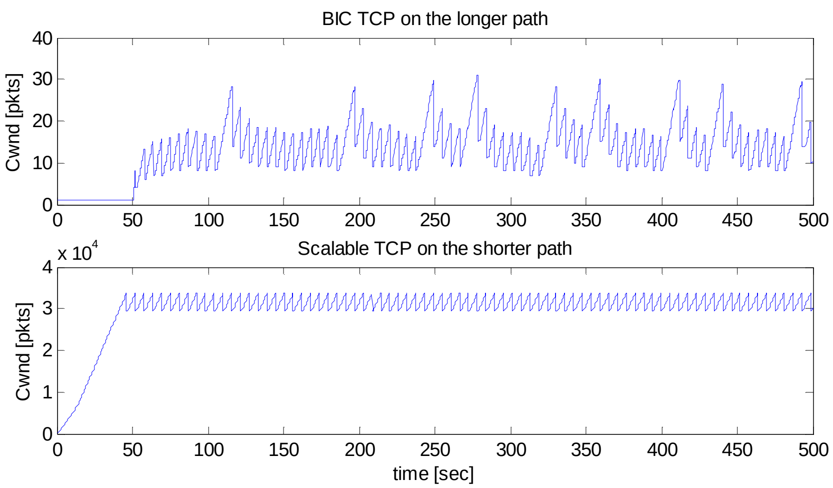

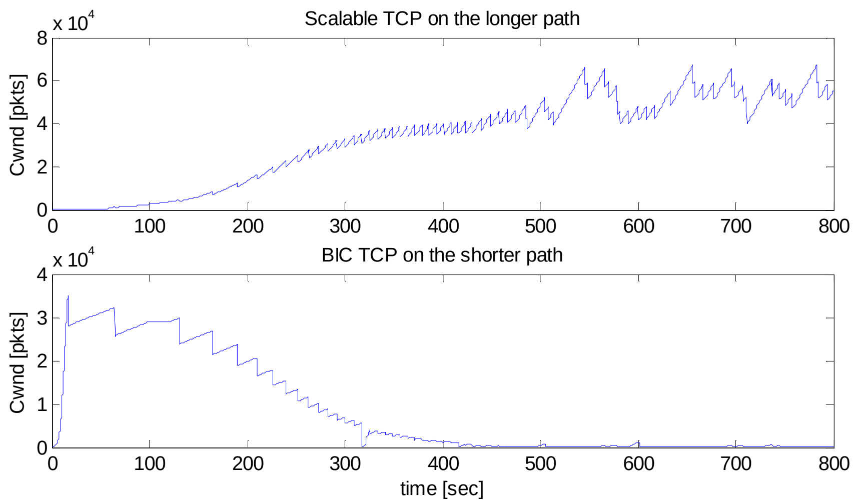

Summary of the results

- MIMD mechanism can not guarantee the fair behavior with AIMD-like protocols or other MIMD flows
- AIMD-like flows are starved in a wide range of network environments
- The intra-protocol behavior of MIMD mechanism is determined by the starting delay
- AIMD-like protocols are able to exhibit intra-protocol fairness
- but the transient phase can be very long

Conclusions

- Fairness analysis of loss-based high speed TCPs
- focusing on dynamic behavior
- Impacts of short term dynamics on fairness
- Explanation of the experienced phenomena

Thank You For Your Attention!

Simulation parameters

Network parameters

capacity	1 Gbps
RTT	100 ms
packet size	1,500 bytes

Bottleneck buffer size

dumb-bell	8,333 pkts
simple parking-lot	25,000 pkts

HSTCP parameters

Low_W	38
High_W	83,000
High_P	10^{-7}
High_Dec	0.1

Sampling periods

cwnd	0.01 sec
queue	0.01 sec
throughput	1 sec

BIC TCP parameters

beta	0.8
S_max	32
S_min	0.01
B	4

STCP parameters

a	0.01
b	0.125

Motivation (cont'd)