Probe or Wait : Handling tail losses using Multipath TCP

Kiran Yedugundla, Per Hurtig, Anna Brunstrom

Outline

- Introduction
- Handling tail losses
    - Tail loss probe (TLP)
- Emulation Setup
- Analysis
- Proposed TLP modification
- Evaluation
- Conclusions

Introduction

- Losses are common cause of performance degradation
- Response time with losses is 10 times more than lossless transfer -> Hurts latency sensitive applications
- End to End losses handled by transport layer (TCP, MPTCP, .)
- Loss recovery affected by location of packet loss in a packet train or burst,
- Short flows are more affected than long ones.
- How TCP and MPTCP handles packet losses?

Introduction - MPTCP

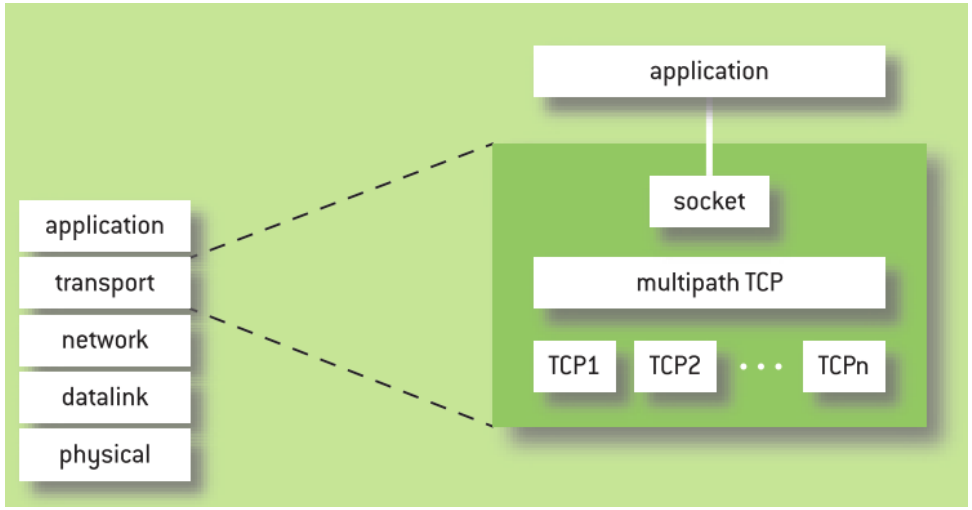

- Applications interact through the regular socket API
- MPTCP manages the underlying TCP connections (subflows)
- MPTCP acts as a ”middleware” between the socket API and the subflows
- MPTCP requires additional signaling between end hosts to set up a connection, adding subflows, and transmitting data

Handling packet losses - TCP

- Retransmission timeout

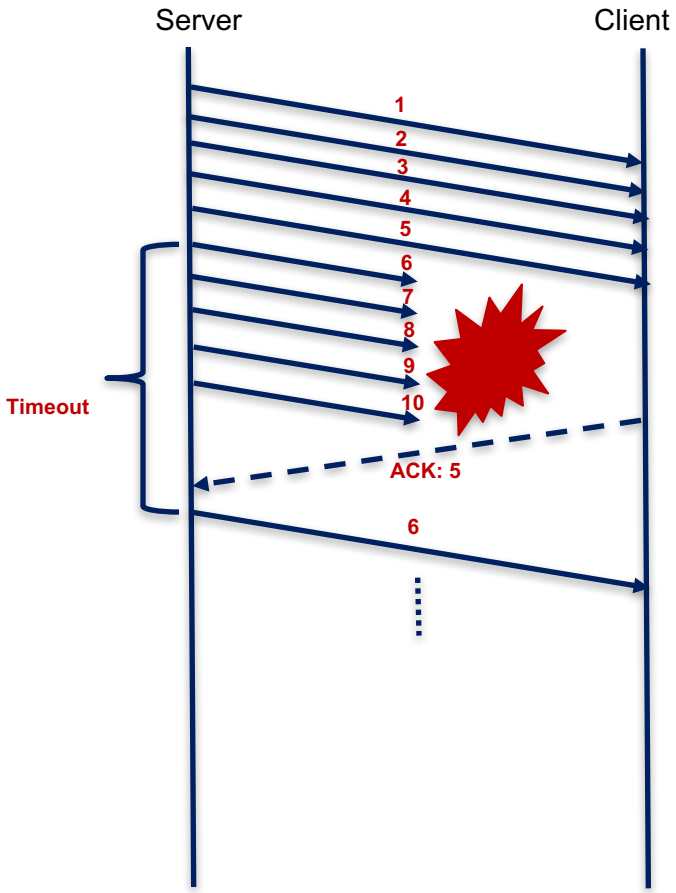

Handling packet losses - TCP

- Fast retransmit -Duplicate ACK
- Further improvements
  - Limited Transmit
- Based on congestion state of the connection
  - SACK, FACK
  - Not relevant for short flows

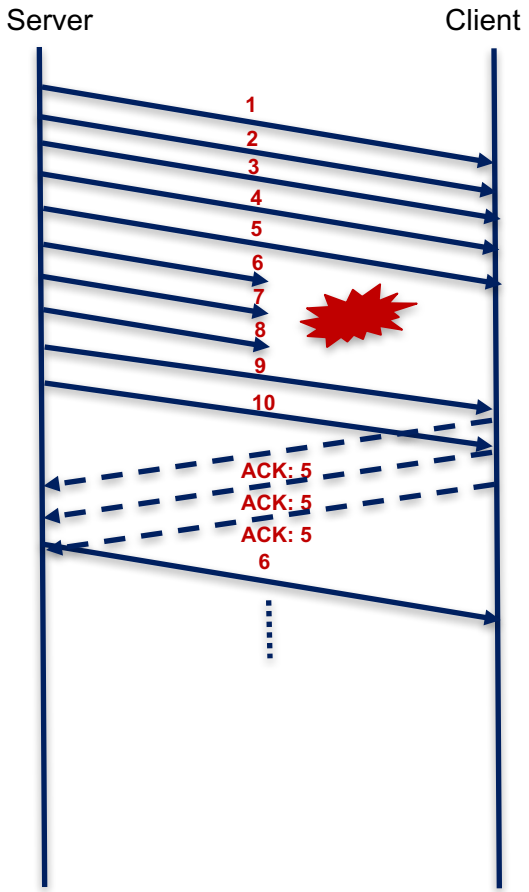

Handling tail losses - TCP

- Tail loss probe
  - Convert RTOs into fast recovery
  - Transmit loss probe after approx. 2. RTT in absence of ACKs.
  - Retransmit last packet (or new if available) to trigger fast recovery.
- Early Retransmit

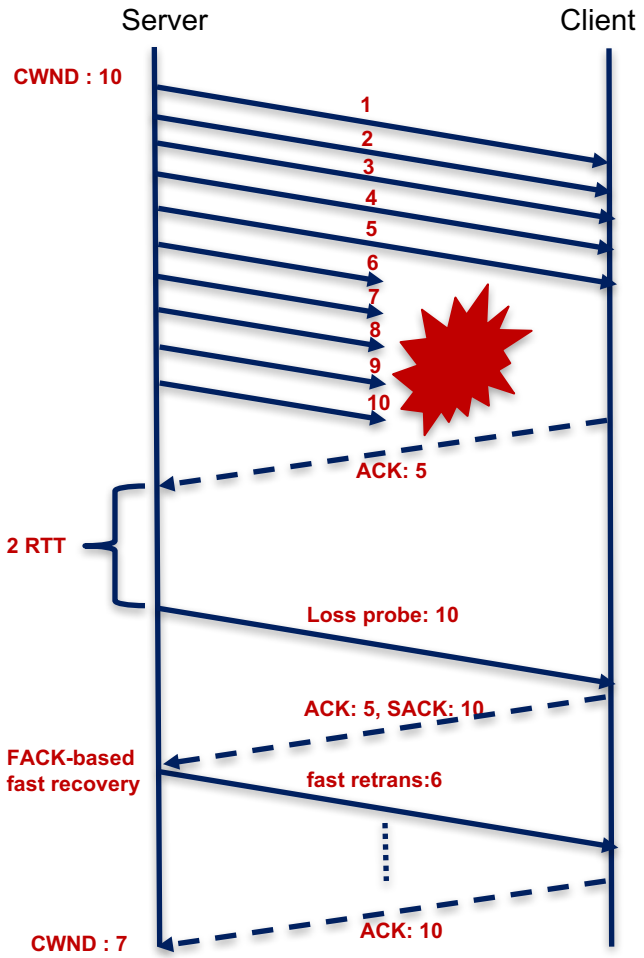

Handling tail losses - MPTCP

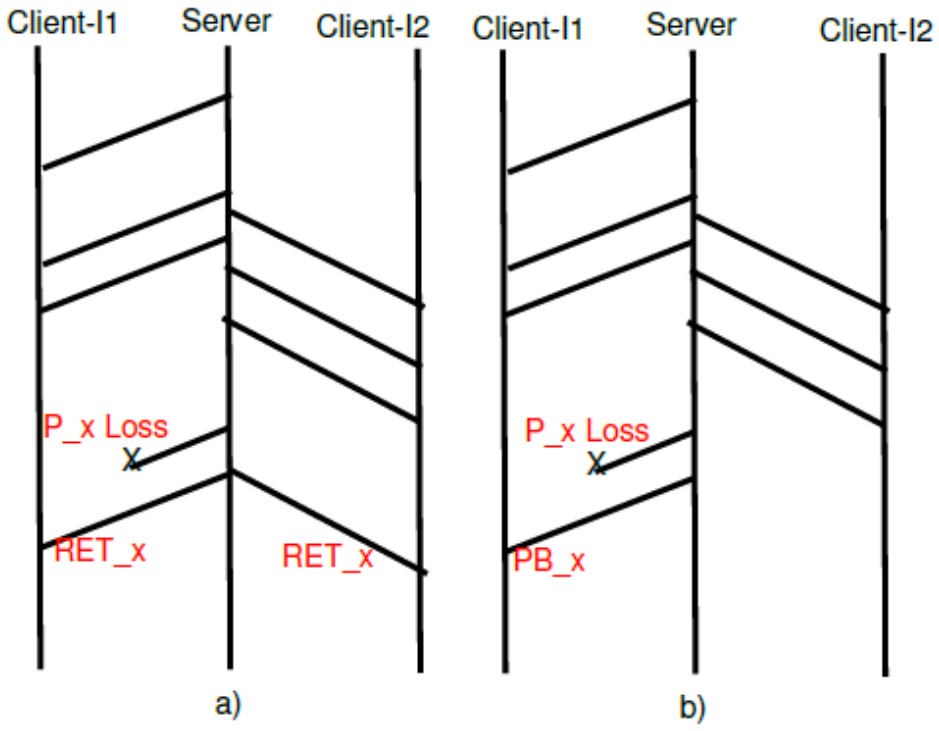

(i) Single packet

Handling tail losses - MPTCP

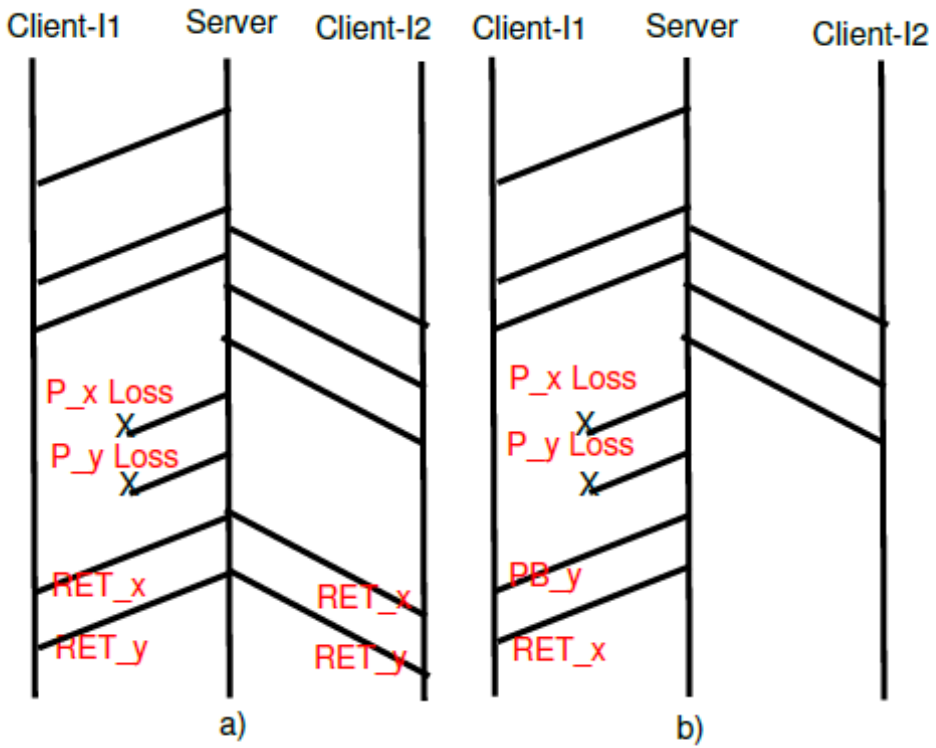

(ii) Two packets

Handling tail losses - MPTCP

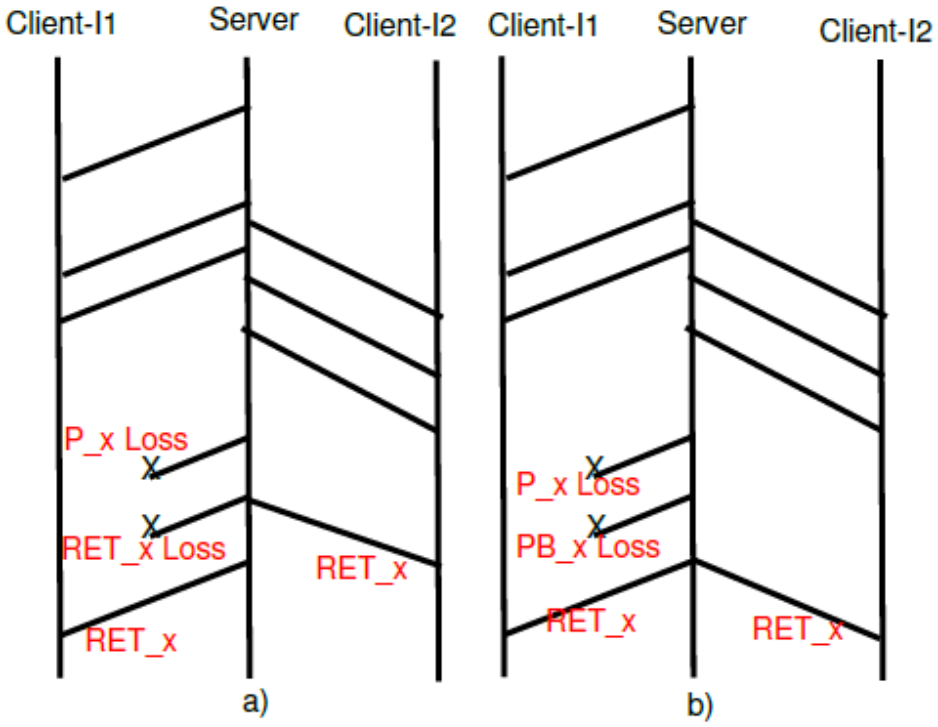

(iii) Single packet and probe loss

Emulation Setup

Topology

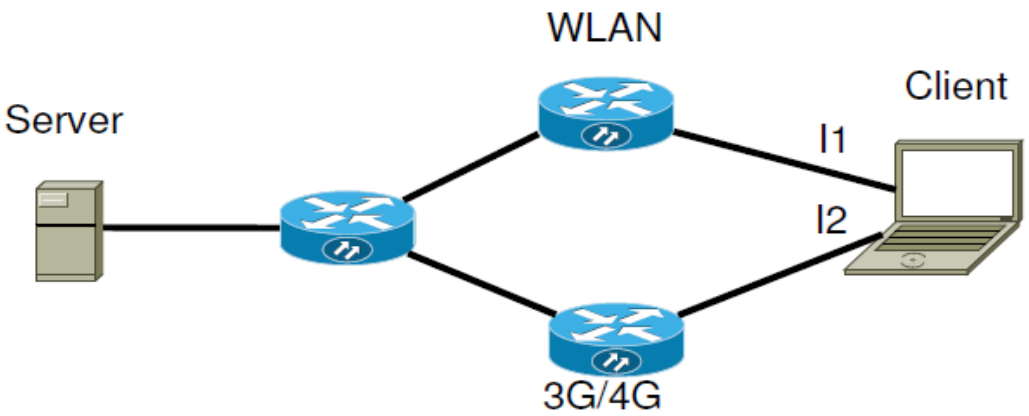

Parameters	WLAN/4G
Link Capacity (Mbps)	54
Propagation Delay (ms)	20 ms – 120 ms
Loss	Deterministic
Burst size	80 packets
Seperation time	2s

Scenarios	
Symmetric	20ms-20ms
Asymmetric	20ms-30ms 30ms-20ms 20ms-120ms 120ms-20ms

Performance analysis - TCP

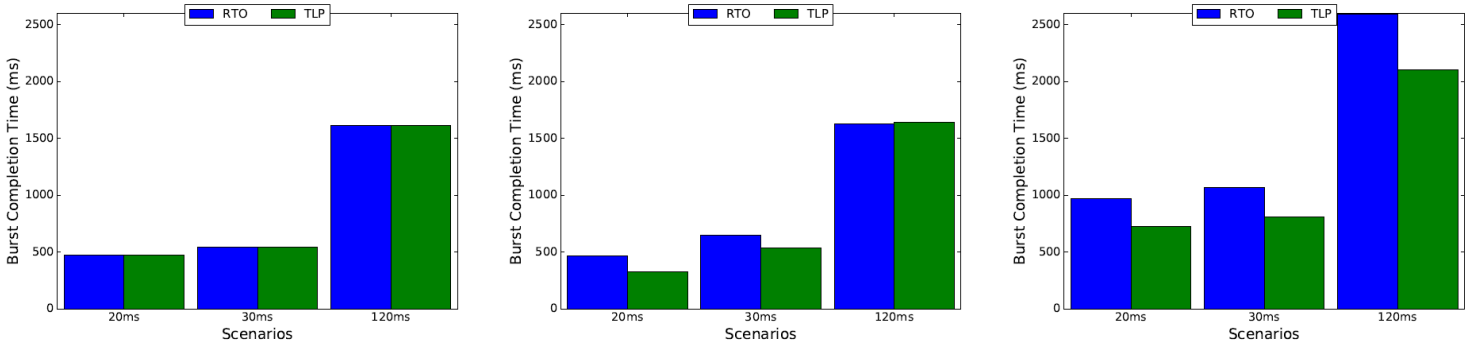

(a) Single packet

(b) Two packets

(c) Single packet and probe loss

Fig. 2: Tail loss scenarios using TCP

Performance analysis - MPTCP

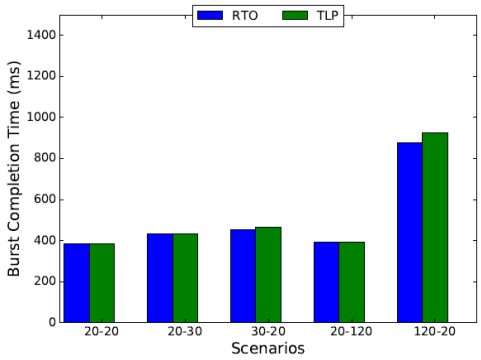

(a) Single packet

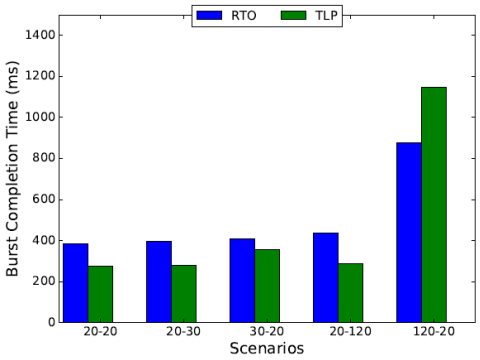

(b) Two packets

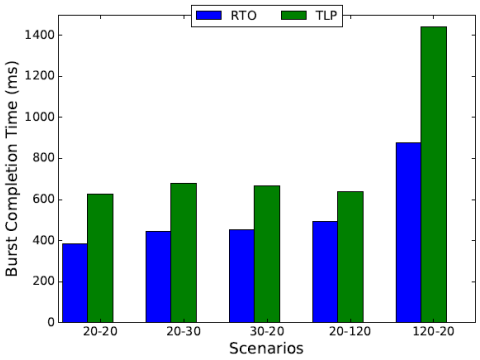

(c) Single packet and probe loss

Fig. 4: Tail loss scenarios using MPTCP

Proposed TLP in MPTCP

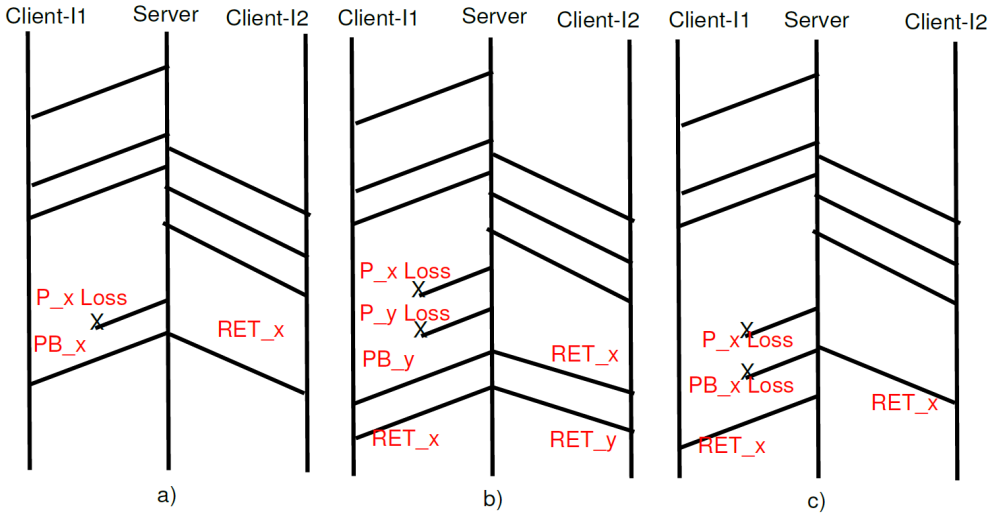

Fig. 5: Timing diagram of MPTCP behavior with proposed improvements with loss of a) Single packet b) Two packets c) Single packet and probe loss

Performance Analysis

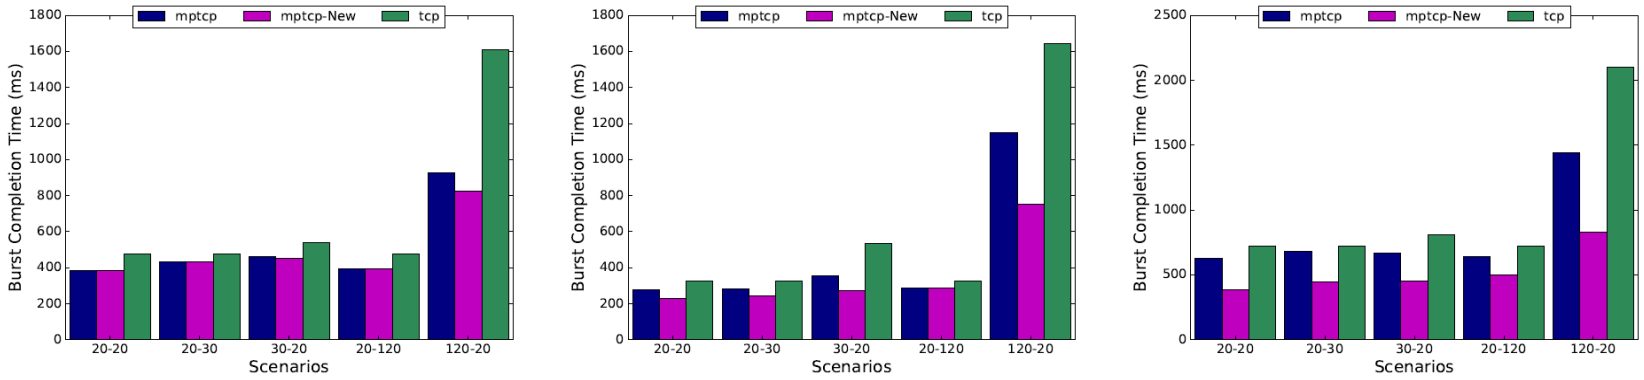

(a) Single packet (b) Two packets (c) Single packet and probe loss

Fig. 6: Performance comparison of MPTCP-New with MPTCP and TCP

Conclusions

- Study tail loss recovery
- Comparative analysis of handling tail losses with TCP and MPTCP
- Proposed a less conservative approach to handling tail losses in MPTCP
- Efficient in cases of probe loss or path loss
- Efficient in cases with large asymmetry and loss on higher delay path
- Improvements of up to 50% in some scenarios

MPTCP in Future Internet Transport

- MPTCP – Ongoing research
  - Improving protocol
  - Improving throughput
  - Improving End-to-End Latency (This paper)
- Usecases
  - Datacenters
  - 4G/WLAN Offload
  - MPTCP Proxies
- Deployments
  - Apple Siri
  - Korean Telecom
  - Netvision Korea LTE/WiFi Traffic Aggregation Solution (MPAS)

Research work supported by: SIDUS READY Project
(RESEARCH ENVIRONMENT FOR ADVANCING LOW LATENCY INTERNET)

Thanks

Q&A

Handling packet losses - TCP

- Fast retransmit - Selective ACK
  - Do not resend already SACKed packets
  - Further improvements with FACK congestion control