Redesigning MPTCP for Edge Clouds

Nitinder Mohan[T] Tanya Shreedhar[I] Aleskandr Zavodovski[T] Otto Waltari[T]
Jussi Kangsharju[T] Sanjit K. Kaul[I]
[T]University of Helsinki, Finland [I]IIIT Delhi, India

Rise of Edge Clouds

- Applications like IoT, AR/VR require low network delay that datacenters fail to provide
- Edge cloud is collection of compute-capable servers deployed close to users
- Edge server has **multiple network interfaces** of **different access technologies**[1]
  e.g. Ethernet, WiFi, LTE, .. etc.

Advantage: Low delay due to localized connectivity over public network

Problem: Cloud technologies (e.g. VM migration) are designed for managed networks

Multipath TCP over Edge

MPTCP[2] is TCP extension which forms multiple parallel TCP flows over available NICs

Benefits: Robustness, reliability, bandwidth aggregation, seamless handovers etc.

Usage: Apple Siri[3], Citrix NetScaler[4], Korea Telecom[5], ...

Naturally utilizes all NICs of Edge servers

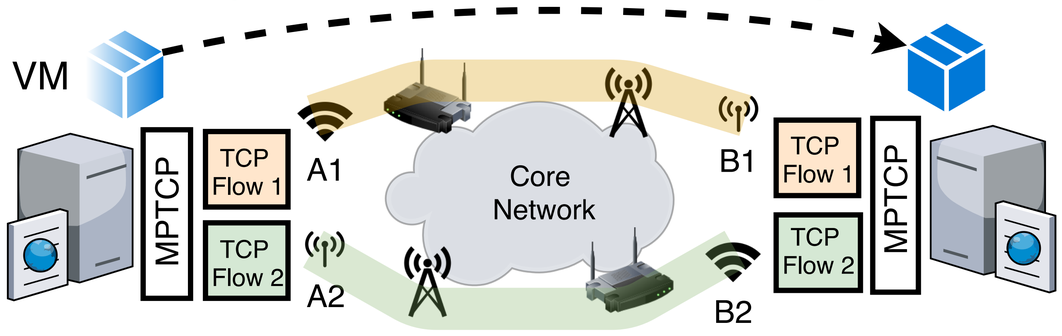

MPTCP in Edge ↔ Edge

Setup: Sender ↔ Receiver running two parallel TCP paths over 802.11g WiFi NICs

Timeline:

t_1: Packet errors [20%] start on Flow 2

t_2: MPTCP stops injecting packets on affected flow

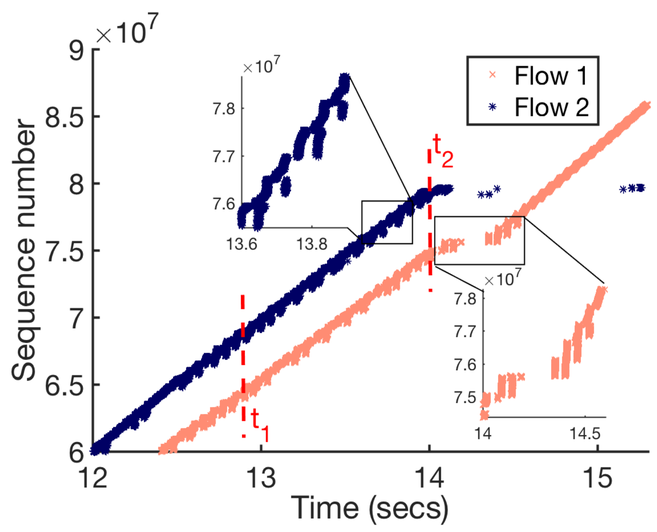

Impact: ≈74% packets on Flow 2 retransmitted

Reason? MPTCP control decisions are based on **sender-perceived Round-Trip Time**

Receiver-Assisted MPTCP (RAMPTCP)

Enables MPTCP sender to consider receiver-side metrics (along with SRTT) in control decisions

e.g. channel utilization, signal strength, path loss% ...

Goal: Reduce delays due to re-ordering & re-transmissions

System Design

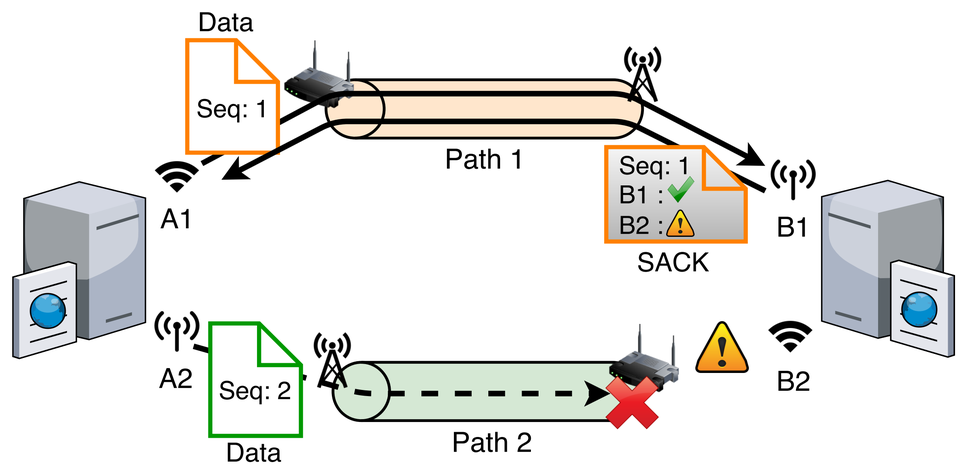

Sender sends packet to Receiver → Receiver sends back local metrics of paths in ACK → Sender breaks per-path RTT into segment delay $T_{sender} + T_{core} + T_{receiver}$ → Sender takes flow decisions considering detailed delays

Preliminary Evaluation

ns3 topology where both sender & receiver are equipped with **two 802.11g NIC**

Packet loss on **B1** NIC (flow 1) from 2-5 secs

Result:

↓Re-transmissions 58%

↑Throughput 19%

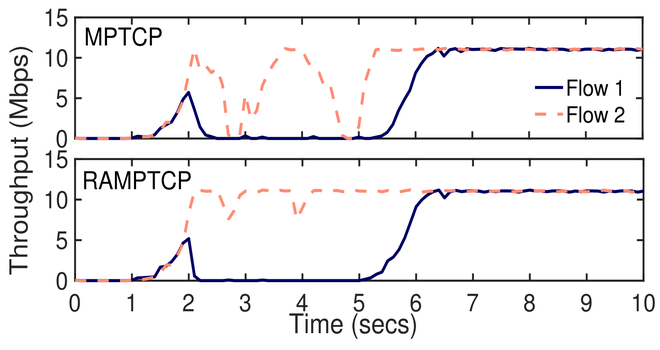

Discussion & Future Work

- Embedding receiver-metrics in packets?
  →ACK, TCP data, MPTCP Data Sequence Signal
- RAMPTCP control decisions?
  →Limit packet injections, packet duplication, TCP rate change etc.
- Network applicability?
  →Zero dependence on underlying NIC

References

[1] CISCO edge server data spec sheet
[2] https://www.multipath-tcp.org/
[3] https://support.apple.com/en-us/HT20137
[4] https://www.citrix.com/blogs/2013/05/28/maximize-mobile-user-experience-with-netscaler-multipath-tcp/
[5] http://blog.multipath-tcp.org/blog/html/2015/07/24/korea.html