LISA: A Linked Slow-Start Algorithm for MPTCP

draft-barik-mptcp-lisa-00

Runa Barik (UiO), Simone Ferlin (SRL), Michael Welzl (UiO)

Multipath TCP @94th IETF Meeting
Yokohama, Japan
2nd November 2015

R/TE

REDUCING INTERNET TRANSPORT LATENCY

[simula . research laboratory]

Problem Statement

Performance Issue to MPTCP:

- MPTCP is aggressive during slow start (SS): subflows are uncoupled.
- What will a Datacenter face when 8 subflows per MPTCP is recommended [3]?

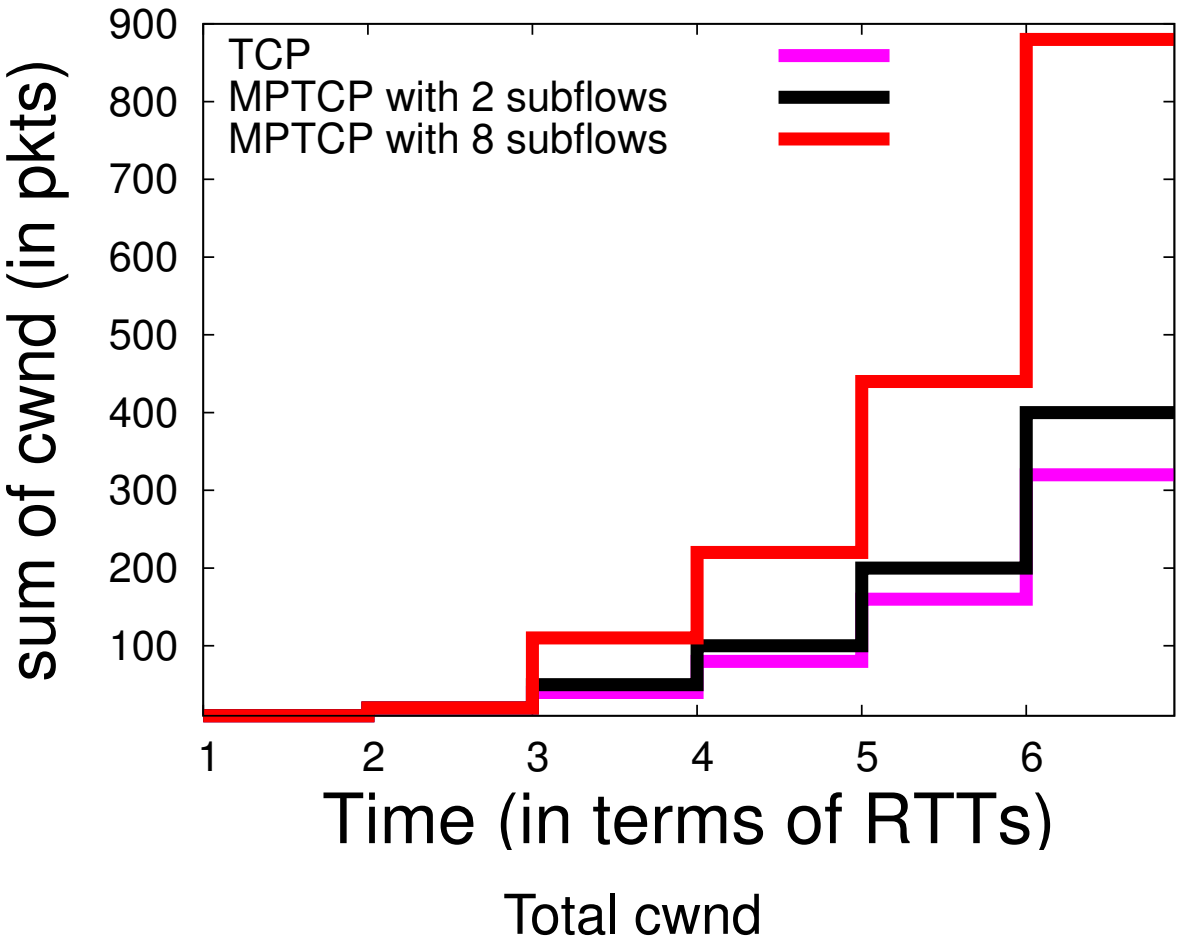

Proposed Solution for LISA

- LISA tries to behave like one TCP during SS.

Basic approach in LISA:

- From subflows in SS, select subflow with maximum sending rate (*max_subflow*).
- From *max_subflow*'s cwnd, take between 3 and IW packets as "packet credit" to give *new_subflow* as IW.
- *Max_subflow* ignores cwnd-increase for (*packets_inflight* - *cwnd*) ACKs.
- If no *max_subflow* , set the IW of new_subflow as per RFC 6928.

Contd.

Alternative approach:

- Assign smaller IW to each subflow.
  - Problem: we do not know how many subflows will be established.

Implementation:

- LISA is implemented as a patch to the Linux kernel 3.14.33+, within MPTCP's v0.89.5.
- The algorithm is explained in more detail in the draft.

LISA results in CORE emulator

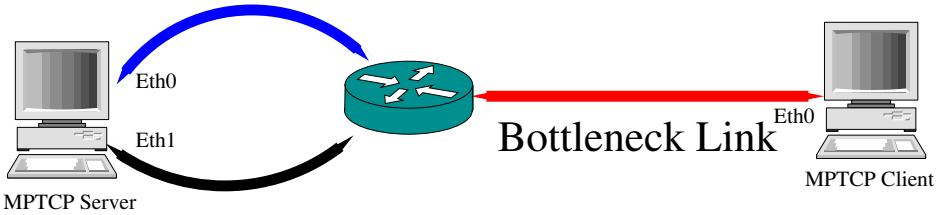

(a) Shared bottleneck: 2.5Mbps, 70ms, transfer-size: 300KB

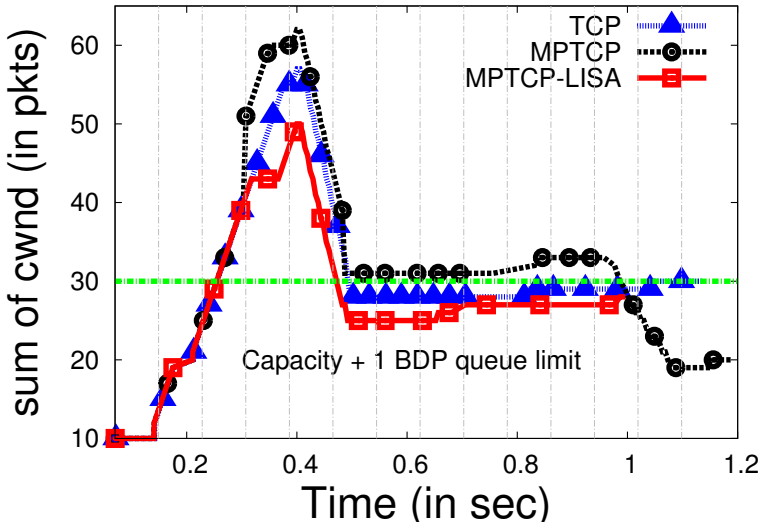

(b) Total cwnd

Performance Evaluation metrics:

- Completion time T,
- Total retransmissions R, and retransmissions during SS, R_{ss}.

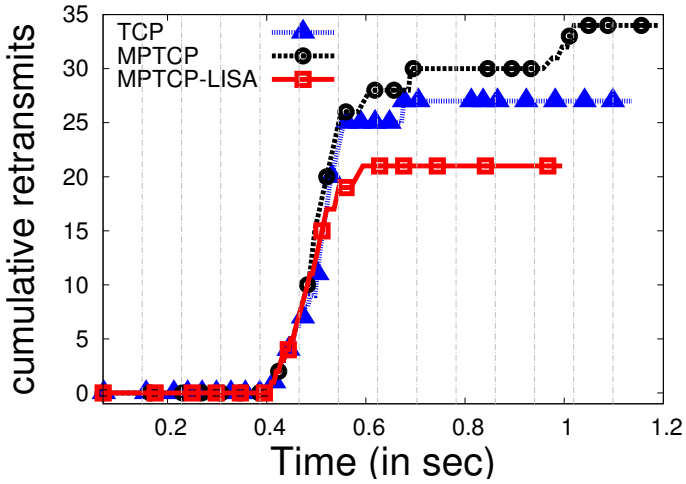

Total retransmissions

Shared Bottleneck

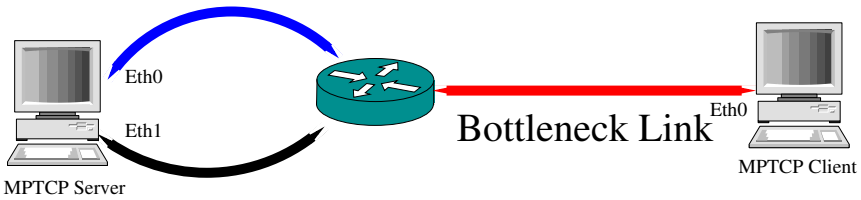

(a) Shared bottleneck:5Mbps, 40ms

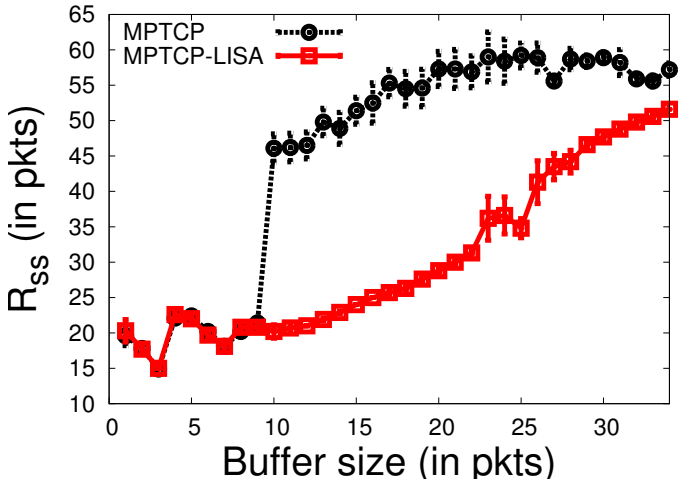

(b) Total retransmits (SS) R_{ss} vs buffer size

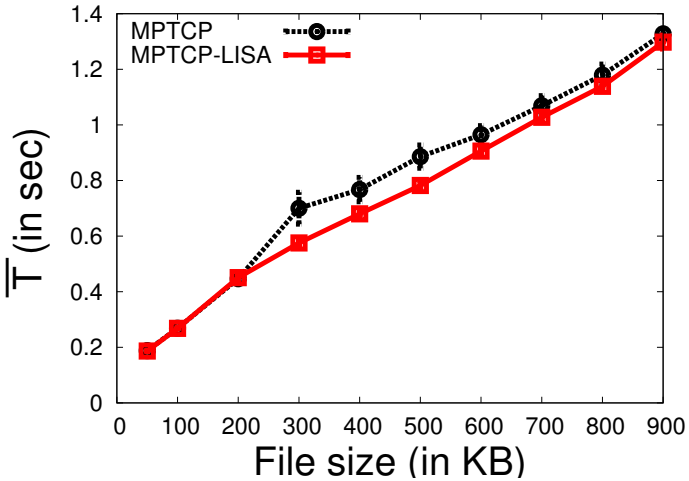

(c) Mean completion time vs file size

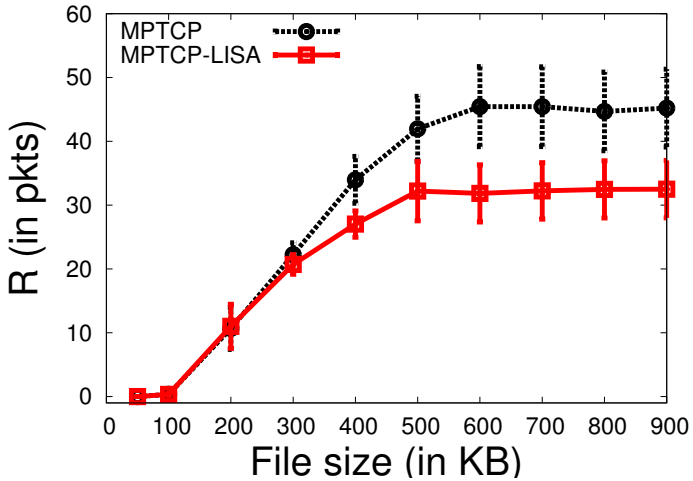

(d) Retransmissions vs file size

Datacenter

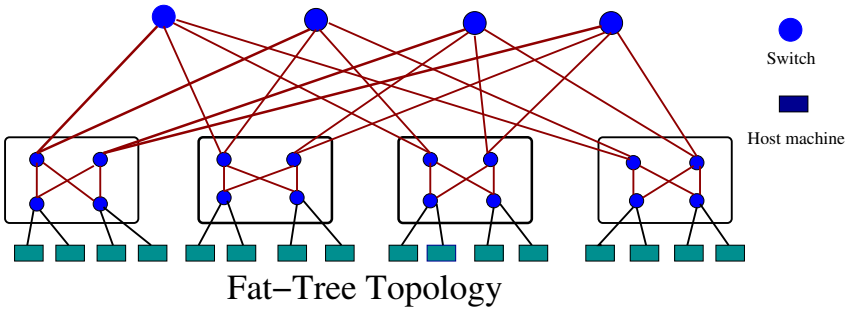

(a) Datacenter:50Mbps, 5ms

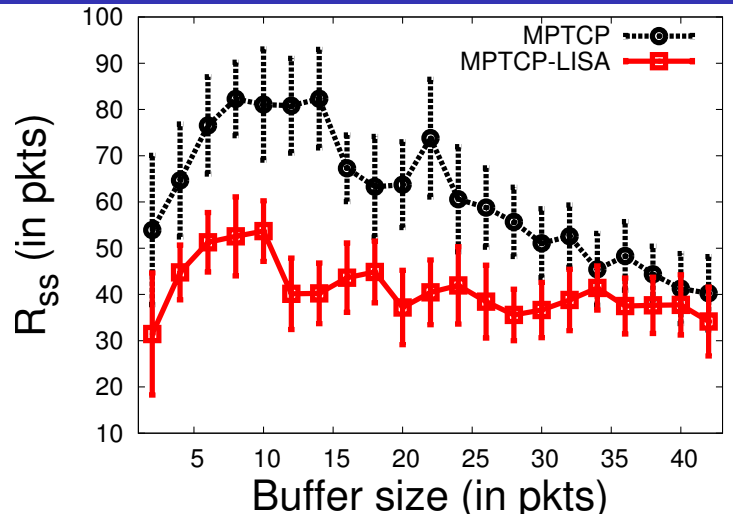

(b) R_{ss} vs buffer size

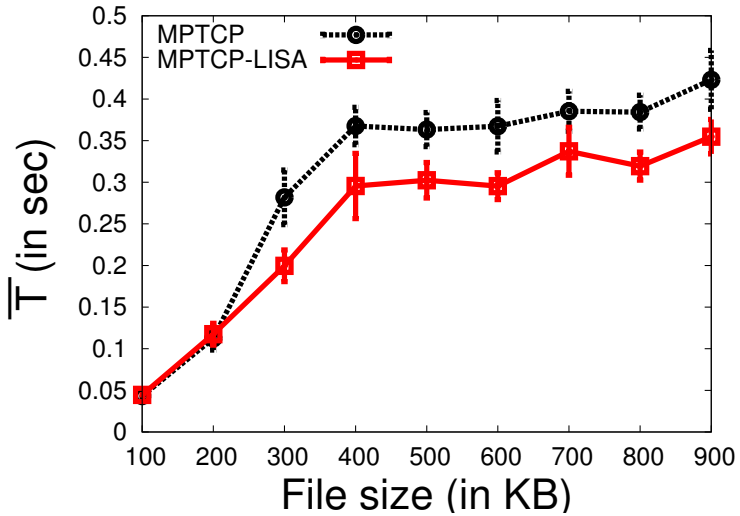

(c) Mean cmpl. time vs file size

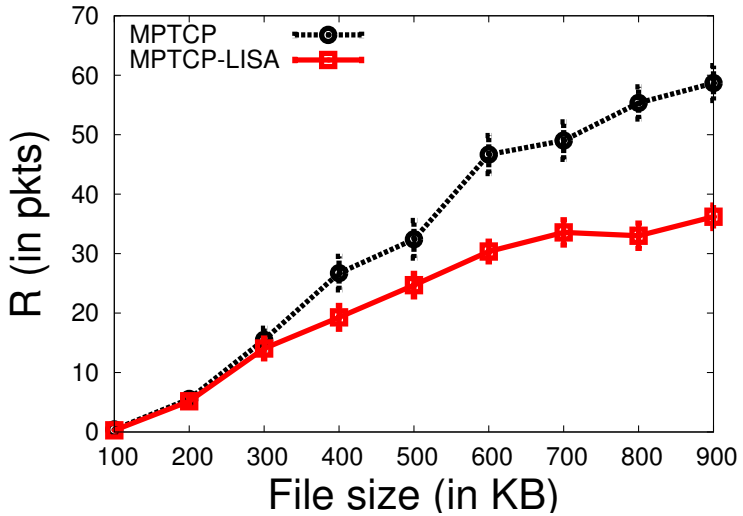

(d) R vs file size

Datacenter

LISA performance in Nornet testbed

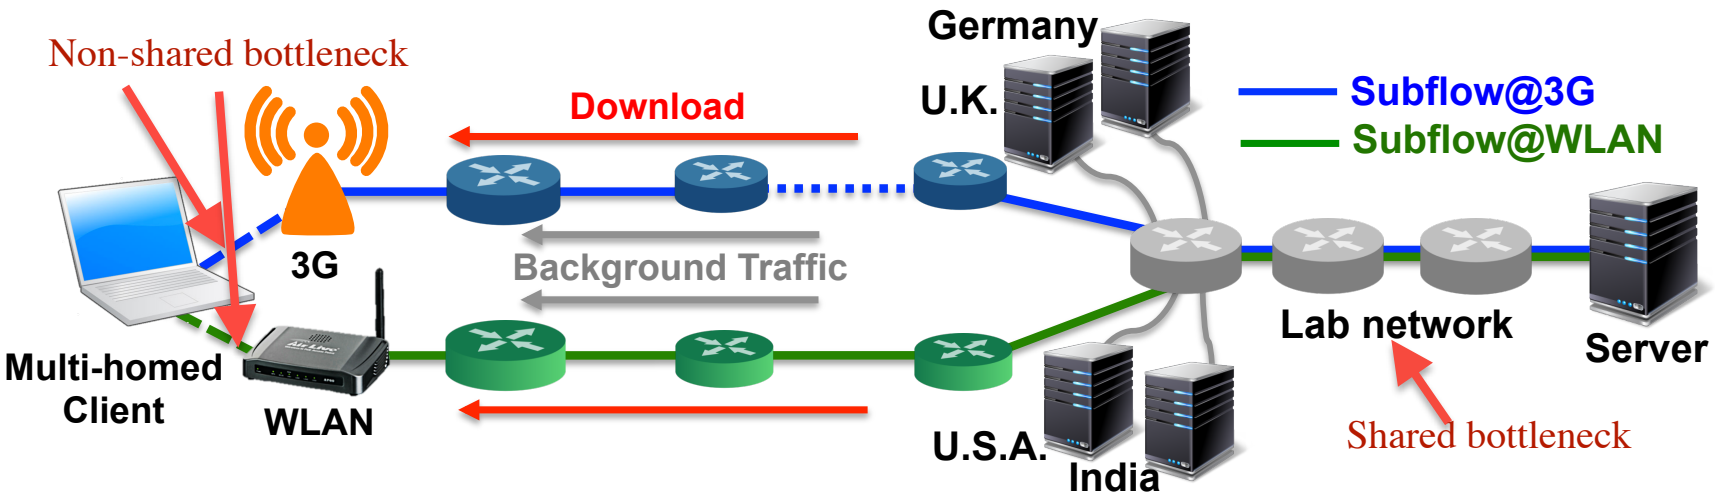

Results

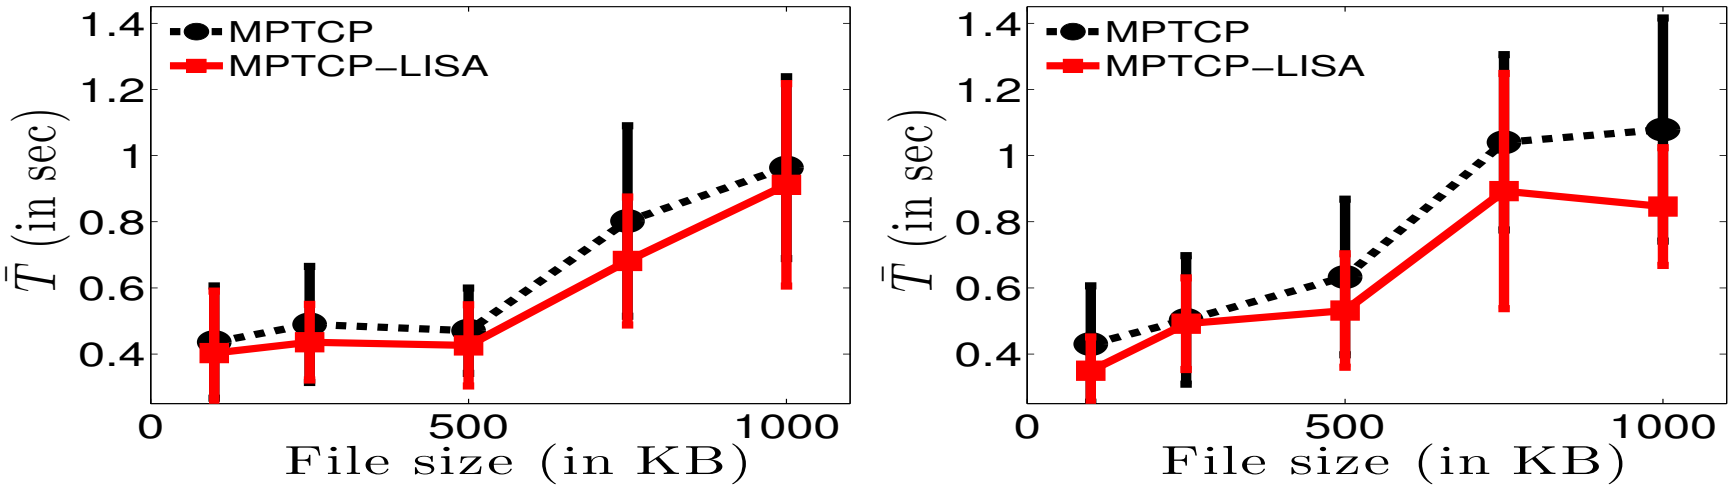

(a) Shared Bottleneck (b) Non-Shared Bottleneck

MPTCP and MPTCP-LISA: Completion Time

Conclusion

- We identified the adverse effect of MPTCP's uncoupled slow-start on the performance of MPTCP itself and on concurrent TCP traffic.
- LISA was implemented as a patch to the Linux kernel 3.14.33+, within MPTCP's v0.89.5.
- We analysed the performance of LISA in shared, non-shared bottleneck experiments in both emulation and real testbed, and in a datacenter topology in emulation.

Acknowledgements:

- This work was part-funded by the European Community under its Seventh Framework Programme through the Reducing Internet Transport Latency (RITE) project (ICT-317700).

References


M. Allman, S. Floyd, and C. Partrdge.
Increasing TCP's Initial Window.
RFC 3390, Oct 2002.


J. Chu, N. Dukkipati, Y. Cheng, M. Mathis.
Increasing TCP's Initial Window.
RFC 6928, April 2013.


C. Pearce, P. Thomas.
Multipath TCP — Breaking Today's Networks with Tomorrow's Protocol.
In: *BlackHat USA 2014*.

Thank you!