Reducing Transport Latency using Multi-path Protocols

Per Hurtig†, Özgü Alay*, Simone Ferlin*, Kiran Yedugundla†
†Karlstad University, *Simula Research Laboratory
Email: {perhurt, kirayedu}@kau.se, {ozgu,ferlin}@simula.no

***Abstract*—Multipath TCP (MPTCP) is a proposed extension of TCP that provides network resilience to failures and improved throughput by utilizing the available redundant resources. While several studies confirm the improved throughput with the use of MPTCP under various scenarios, its end to end latency performance has not been considered. In this paper, we evaluate the impact of MPTCP on web traffic transport latency.**

I. Introduction

In recent days, there has been significant growth in the use of mobile devices for Internet access. Many devices, such as smartphones, are equipped with two interfaces: a WLAN and a mobile broadband (e.g., 3G/4G). However, applications currently exploit only one of these interfaces though there may be throughput/latency improvements in simultaneous use of the interfaces. To enable such efficient resource usage, multi-path transport protocols are currently being deployed. One example is Multi-path TCP (MPTCP) [1].

Our research effort leverages the two trends discussed above. To the best of our knowledge, it is the first research effort that investigates the potential benefits of using multiple paths to reduce latency in order to improve web user experience. This paper is part of a larger study that assesses the potential latency reduction by leveraging multipath transport protocols. While the study discusses and evaluates the performance of different multipath protocols considering various delay sensitive applications, we limit the discussion in this paper to the latency performance of MPTCP for web traffic.

II. Web Transfer and Multi-path TCP

In today's Internet, we observe an increased number of applications that are extremely sensitive to latency, such as web transfers, video streaming and online gaming. In such applications, the user's experience is affected significantly when data delivery is delayed. The Internet is dominated by web traffic running on top of short-lived TCP connections [2] with around 95% of client TCP and 70% of server TCP traffic consisting of smaller than ten packets [3]. The quality of user experience when accessing a web page is linked to the download times of the corresponding short-lived flows. As an example, Google measured a 25% reduction in the number of searches done by users, as a result of adding 500 ms to web search times [4].

MPTCP is developed by the IETF MPTCP working group and provides a set of extensions to TCP [5] to enable the simultaneous use of multiple paths between endpoints. The motivation behind MPTCP is efficient resource usage and improved user experience through improved resilience and higher throughput. Once an MPTCP connection has been fully established, both end hosts can send data over any of the available subflows, transparently to the application. Thus, MPTCP is able to schedule web objects on one path even if the other path experiences loss or excessive buffering. We therefore believe that MPTCP could be a good candidate to transport web traffic.

III. Experimental Setup and Results

We evaluate the performance through emulations and real-world experiments. For the emulations, we use the Common Open Research Emulator (CORE) [6]. The setup consists of two end-hosts, one acting as a mobile device equipped with two interfaces and the other as a web server. To reach the server, each interface of the mobile client is connected to an (emulated) access network. For the real-world experiments, we use the NorNet testbed [7] that consists of nodes that are simultaneously connected to multiple 3G operators in Norway as well as multiple public WLANs at Simula Research Laboratory. In our experiments, NorNet node acts as the end-host and we use a web-server located in Oslo.

To evaluate the effect of having both homogeneous and heterogeneous network access, we performed experiments where the interfaces used a combination of emulated WLAN and 3G access (i.e., WLAN-WLAN, WLAN-3G, 3G-3G). The capacity of the WLAN was randomly chosen in the ranges 20-30Mbps, whereas for 3G we use values in the range of 3-5Mbps. The propagation delays are chosen in the range of 20-25ms and 65-75ms for WLAN and 3G, respectively. In the WLAN access we consider a random loss of $1-2\%$.

For each experiment, the server stores a small set of files from three different classes of websites: Wikipedia, Amazon and Huffpost. Each website contains a different number of objects of different sizes. Of the three, Wikipedia was the smallest (15 objects, 72Kb total size) followed by Amazon (54 objects, 1Mb total size) and then Huffpost (138 objects, 3.9Mb total size). The sites were then requested and downloaded from the mobile client using six concurrent connections each using TCP or MPTCP, in respective experiments.

Figure 1 shows the average download times for the web pages, with standard deviation, over 30 repetitions for each configuration. Each figure depicts the average download time when using MPTCP or TCP on one of the interfaces (TCP1,

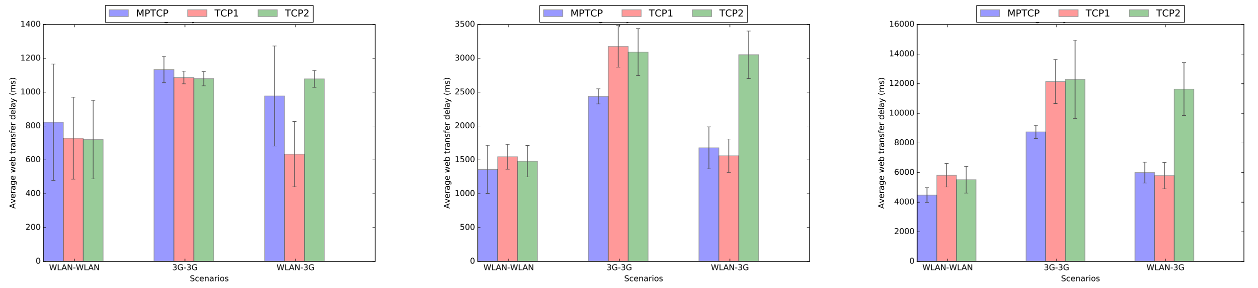

(a) Wikipedia (15 objects, 72Kb total size). (b) Amazon (54 objects, 1Mb total size). (c) HuffPost (134 objects, 3.9Mb total size).

Fig. 1: Emulation: Average web transfer delay, with standard deviation.

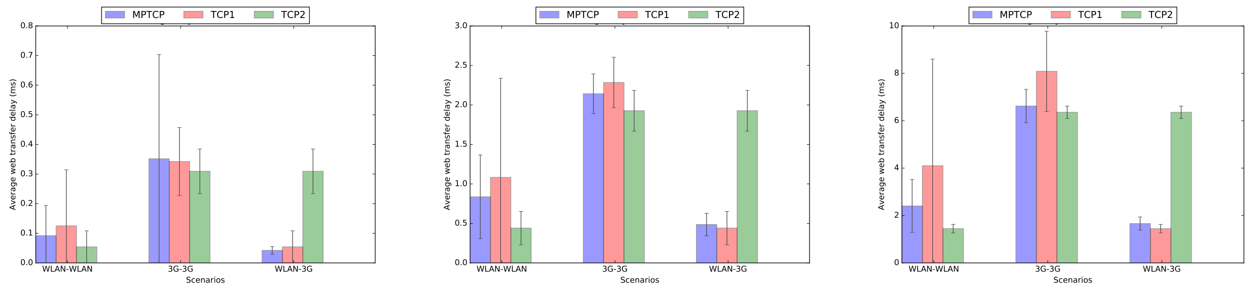

(a) Wikipedia (15 objects, 72Kb total size). (b) Amazon (54 objects, 1Mb total size). (c) HuffPost (134 objects, 3.9Mb total size).

Fig. 2: NORNET: Average web transfer delay, with standard deviation.

TCP2); the results are grouped according to the emulated access networks used (WLAN-WLAN/3G-3G/WLAN-3G). We observe that MPTCP cannot reduce download times for the Wikipedia site, but is able to do so for both Amazon and Huffpost (especially in the 3G-3G scenarios). The reason for the poor performance of retrieving the Wikipedia site is simply that the amount of data is so small that it can be transmitted within TCP's initial window (given the six concurrent connections used). Employing more paths in such scenarios is not useful as long as the path itself can sustain the traffic load. However, for the other sites the amount of data is much larger and the download time can be reduced by MPTCP's implicit load-balancing over its available subflows; especially when the paths are homogeneous, as possible head-of-line blocking effects are less prevalent.

Figure 2 illustrates the average transfer times for web traffic for real-word experiments. In our experiments, we observe that the two selected 3G networks (also the two WLAN networks) have very different path characteristics, hence, we mostly have heterogeneous paths where the performance of MPTCP is similar to or slightly worse than that of TCP, consistent with emulation results.

IV. Conclusion

In this paper, we evaluate the latency performance of MPTCP for web traffic and compare it with that of TCP using both emulations and real-world experiments. We observed that MPTCP utilizes the available paths efficiently when there is enough data to send and performs better when there is path symmetry. No improvement with MPTCP was observed when there is path asymmetry and it performs worse when there is not enough data to send. Moreover, our results indicate that MPTCP provides significant gains for websites with large object sizes both for homogeneous and heterogeneous paths.

Acknowledgments

The authors are funded by the European Community under its Seventh Framework Programme through the Reducing Internet Transport Latency (RITE) project (ICT-317700). The views expressed are solely those of the authors.

References

[1] A. Ford *et al.*, "TCP Extensions for Multipath Operation with Multiple Addresses," IETF, RFC 6824, Jan. 2013.

[2] C. Labovitz *et al.*, "Atlas Internet Observatory 2009 Annual Report," in *47th NANOG*, 2009.

[3] D. Ciullo *et al.*, "Two schemes to reduce latency in short lived TCP flows." *IEEE Communications Letters*, vol. 13, no. 10, 2009.

[4] O3b Networks and Sofrecom, "Why Latency Matters to Mobile Backhaul," Apr. 2013.

[5] J. Postel, "Transmission Control Protocol," IETF, RFC 793, Sep. 1981.

[6] J. Ahrenholz, "Comparison of core network emulation platforms," in *MILCOM 2010*, Oct 2010.

[7] T. Dreibholz and E. G. Gran, "Design and Implementation of the NorNet Core Research Testbed for Multi-Homed Systems," in *(PAMS) 2013*, Barcelona, Spain, Mar. 2013.