



HAL
open science

Measuring and Modeling Multipath TCP

Viet-Hoang Tran, Ramin Sadre, Olivier Bonaventure

► **To cite this version:**

Viet-Hoang Tran, Ramin Sadre, Olivier Bonaventure. Measuring and Modeling Multipath TCP. 9th Autonomous Infrastructure, Management, and Security (AIMS), Jun 2015, Ghent, Belgium. pp.66-70, 10.1007/978-3-319-20034-7_8 . hal-01410150

HAL Id: hal-01410150

<https://hal.science/hal-01410150>

Submitted on 6 Dec 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Measuring and Modeling Multipath TCP

Viet-Hoang Tran, Ramin Sadre, Olivier Bonaventure

ICTEAM, Université catholique de Louvain, Louvain-la-Neuve, Belgium
{tranviet,ramin.sadre,olivier.bonaventure}@uclouvain.be

Abstract. Multipath TCP, a major extension to regular TCP, allows TCP clients to utilize multiple paths to improve the transfer rate and connection robustness. Providing these benefits without requiring to upgrade network infrastructure nor applications, Multipath TCP is becoming more popular. Notably Apple iOS 7 now supports it for SIRI. However, there is still lack of a complete understanding of Multipath TCP in practice. How much can a user benefit from Multipath TCP in different scenarios? Which factors affect the performance of Multipath TCP? How well can we predict the behavior of Multipath TCP in a specific environment? Our research aims to answer these questions by large-scale measurements and model-based analysis. The answers will be an important input for designers and developers to further improve Multipath TCP.

Keywords: multipath, MPTCP, traffic measurement, modeling

1 Background

To be able to use several network paths for a single connection is a long-desired feature, since it would bring several benefits including higher transfer rate and better robustness. However, this is not possible with TCP – the dominating reliable protocol in today’s Internet. As an extension for TCP, Multipath TCP (MPTCP) is rapidly adopted by both academia and industry. At the time of writing, MPTCP implementations exist on several operating systems, including Linux [10], Apple iOS and MacOS [1], FreeBSD [14], Solaris [5] and Citrix. A notable use case of MPTCP is to enable WiFi/3G offload on mobile devices [11, 3, 4]. On recent iPhones and iPads, Apple has deployed and enabled MPTCP by default for its voice recognition application (SIRI) [1] in order to reduce end-to-end delays. It could also be used in datacenters to exploit multiple paths between hosts [13]. Before MPTCP, Stream Control Transmission Protocol (SCTP) [6] was designed with multi-homing in mind and supports concurrent multipath extensions [9]. However, SCTP is still not widely used since it requires developers to change their applications and many networking devices like NATs/Firewalls do not understand SCTP and block its traffic. Instead, Multipath TCP was designed with backward-compatible goals in mind: it provides network applications the same API as regular TCP, for that TCP applications can use MPTCP

without any modification. Behind the scene, it uses several subflows which appear to network infrastructure as separated regular TCP connections. Detailed explanations of MPTCP can be found in [7].

2 Motivation and Research Problems

Given the quick adoption and the potential applications of MPTCP, it is important to know how reliable MPTCP is and how much performance we can gain from using it in practice. A thorough understanding of MPTCP behavior and performance now becomes critical. We have identified three major questions that need to be answered:

1. *How is MPTCP currently used?*

In MPTCP, so-called *schedulers* and *path managers* control how subflows are established and how packets are distributed over them. Different implementations for scheduling and path managing are available. In order to further develop MPTCP, the designers of MPTCP need to know what users are currently using. Which configurations are used in practice? Can we find that from passive measurements? For what applications is MPTCP used?

2. *How does MPTCP behave in practice?*

This can be split into more specific questions: whether MPTCP always behaves correctly, and how well the performance of MPTCP is. SCTP has been implemented in many operating systems, but its usage is very limited since it is incompatible with conventional TCP. Learned from the painfully slow adoption/deployment of SCTP, MPTCP was designed to work with conventional applications and network infrastructure. For example, MPTCP should fall back to regular TCP when proper operation over multiple paths is not possible. Until now, there is little knowledge about the correct behavior of MPTCP in such situations over the global Internet. In terms of performance, the fact that MPTCP uses several paths does not mean it will automatically gain the sum of goodput over all paths. Moreover, the perceived delay to the user is also important to investigate.

3. *How to predict the behavior of MPTCP before deploying it in a specific environment?*

While MPTCP has potential to bring several benefits to a wide range of devices, network operators or service providers need to anticipate the benefits and the risks from deploying MPTCP on a large scale.

Generally, in order to answer the above questions, we can evaluate MPTCP through analysis, simulation, emulation [12], or real-world measurement. There are several existing efforts. Passive measurement is done by collecting the available information from the network, for example in [8], while active measurement is done by sending particular packets through networks and getting the necessary metrics [3, 4]. There are also works toward modeling and analysis. For example, Arzani et al. [2] present a simple model for MPTCP behavior for a simple 2-path topology. However, to predict MPTCP behavior in a real, heterogeneous environment is still a big challenge.



Fig. 1. Our methodology

3 Research Directions

We plan to answer research questions 1 and 2 by performing measurements. A first step was done in [8] where we collected and analyzed the MPTCP network traffic on a public Internet server for one week. This server hosts *multipath-tcp.org* site to which everyone interested in MPTCP can access, download kernel and do testing with MPTCP enabled. In order to gain a more general insight into the usage and behavior of MPTCP in practice, we need to conduct measurements with the following considerations in mind:

1. Large-scale measurement: Previous works, which mostly focused on active measurement, only tested MPTCP with a few hosts. We envisage to perform active and passive measurements on a large number and wide variety of hosts. Our current measurement [8] on *multipath-tcp.org* have users accessed from around the world, using a single or multiple network interfaces. Since the popularity of MPTCP is increasing, we have now for the first time the opportunity to observe MPTCP's behavior at large scale.
2. Dealing with a wide variety of MPTCP implementations: We want to develop techniques to identify the specific MPTCP implementation used by a host by passive measurements. This identification is an essential step in determining whether a particular observed behavior is due to the MPTCP protocol or due to a specific implementation.

3. Avoiding bias: Until recently, MPTCP was mainly used by people interested in this new technology. In order to avoid such kind of bias (and others, such as on the used operating system) in our measurements, we need to analyze passive measurement results to know more about the users of MPTCP.

To answer research question 3, we plan to build models of MPTCP behavior that can be applied to different network and host configurations. A first step could be to extend the work in [2] to more complex topologies and other schedulers. We expect that a general model will be hard to achieve. Instead, we will use the outcomes of our measurement activities to define and parametrize the most typical or interesting usage scenarios.

4 Conclusion

This paper presents the need for a thorough research of MPTCP behavior and usage in practice, and gives an overview on the challenges and directions of our research. We believe our work will bring a more comprehensive understanding of MPTCP protocol, implementations and traffic. This knowledge will be a valuable input for MPTCP designer to further improve the protocol. For the users, the work will give a more clear view on the benefit of using MPTCP and foster the widespread of MPTCP usage in practice.

Acknowledgments. This work has been carried out in the context of the METRICS project - an Initial Training Network (ITN) which is supported by the European Commission.

References

1. Apple. iOS: Multipath TCP Support in iOS 7. <http://support.apple.com/en-us/HT201373>.
2. B. Arzani, A. Gurney, Sitian Cheng, R. Guerin, and Boon Thau Loo. Deconstructing MPTCP Performance. In *Network Protocols (ICNP), 2014 IEEE 22nd International Conference on*, pages 269–274, Oct 2014.
3. Y-C Chen, Y-S Lim, R. Gibbens, E. Nahum, R. Khalili, and D. Towsley. A Measurement-based Study of Multipath TCP Performance over Wireless Networks. In *ACM SIGCOMM IMC*, 2013.
4. S. Deng, R. Netravali, A. Sivaraman, and H. Balakrishnan. WiFi, LTE, or Both?: Measuring Multi-Homed Wireless Internet Performance. In *Proceedings of the 2014 Conference on Internet Measurement Conference, IMC '14*, pages 181–194, New York, NY, USA, 2014. ACM.
5. P. Eardley. Survey of MPTCP Implementations. Internet-Draft draft-eardley-mptcp-implementations-survey-02, IETF Secretariat, July 2013.
6. R. Stewart (Ed.). Stream Control Transmission Protocol. IETF RFC 4960, September 2007.
7. A. Ford, C. Raiciu, M. Handley, and O. Bonaventure. TCP Extensions for Multipath Operation with Multiple Addresses. IETF RFC 6824, January 2013.

8. B. Hesmans, H. Tran-Viet, R. Sadre, and B. Bonaventure. A First Look at Real Multipath TCP Traffic. In *7th International Workshop on Traffic Monitoring and Analysis (TMA)*, 2015.
9. J. Iyengar, P. Amer, and R. Stewart. Concurrent Multipath Transfer Using SCTP Multihoming over Independent End-to-end Paths. *Networking, IEEE/ACM Transactions on*, 14(5):951–964, 2006.
10. C. Paasch, S. Barre, et al. Multipath TCP implementation in the Linux kernel. available from <http://www.multipath-tcp.org>, 2014.
11. C. Paasch, G. Detal, F. Duchene, C. Raiciu, and O. Bonaventure. Exploring Mobile/WiFi Handover with Multipath TCP. In *ACM SIGCOMM workshop CellNet*, pages 31–36, 2012.
12. C. Paasch, R. Khalili, and O. Bonaventure. On the Benefits of Applying Experimental Design to Improve Multipath TCP. In *Proceedings of the Ninth ACM Conference on Emerging Networking Experiments and Technologies, CoNEXT '13*, pages 393–398, New York, NY, USA, 2013. ACM.
13. C. Raiciu, S. Barre, C. Pluntke, A. Greenhalgh, D. Wischik, and M. Handley. Improving Datacenter Performance and Robustness with Multipath TCP. In *ACM SIGCOMM 2011*, 2011.
14. N. Williams, L. Stewart, and G. Armitage. FreeBSD kernel patch for Multipath TCP. available from <http://caia.swin.edu.au/urp/newtcp/mptcp/tools.html>, July 2014.