Throughput Performance Analysis of the Multipath Communication Technologies for the Cloud

SZILÁGYI Szabolcs, BORDÁN Imre, HARANGI Lajos, KISS Benjámin

University of Debrecen, Hungary,
Department of IT Systems and Networks, Faculty of Informatics,
26 Kassai Way, 4028 Debrecen, Hungary, E-Mail: szilagyi.szabolcs@inf.unideb.hu

***Abstract* – Cloud based computing is among the top present-day research areas. Datacenters serving as the backend for cloud solutions have to satisfy the demands set by the time-critical applications emerging in our rushing world. In order to provide physical redundancy, the datacenters are equipped with redundant inter-server connections that are left idle during normal operation via the traditional TCP protocol, e.g. with regards to link-capacity aggregation. Our paper provides a comparison between two technologies (MPTCP and MPT-GRE) that support multipath operation and could prove useful in datacenter environments by increasing efficiency and enhancing user experience in the cloud.**

***Keywords:* cloud network; data center; MPT-GRE; MPTCP; multipath communication; performance analysis; system throughput; tunneling.**

I. INTRODUCTION

Cloud based computing is one of the most popular areas in IT today, enabling business processes, software, platforms and infrastructures through its flexible capabilities (**BPaaS**: Business Process as a Service, **SaaS**: Software as a Service, **PaaS**: Platform as a Service, **IaaS**: Infrastructure as a Service). Since numerous applications are communicating through the cloud, cloud networking has a direct impact on user experience.

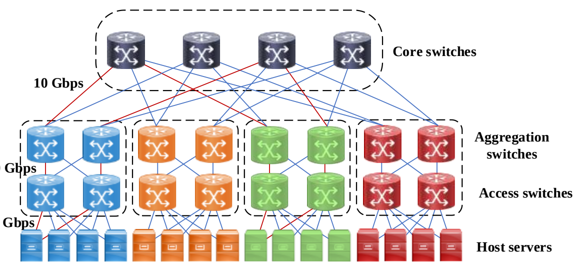

Fig.1. Datacenter topology example

In this paper, we present measurements comparing two technologies enabling multipath communication that can greatly help in providing a more efficient operation in datacenter environments, while also enhancing the user experience.

The MPT-GRE software library [1], capable of effective link-aggregation even in quad-path Gigabit Ethernet environments, was developed at the University of Debrecen. In the current scope, we examine it in a dual-path 10 Gigabit Ethernet measurement environment, providing a comparison with Multipath TCP (MPTCP) [2] from a maximum throughput and resource usage point of view.

In the second chapter, we give a short description about the operating principles of these multipath solutions, followed in the third chapter by a presentation of our measurement environment. Chapter four showcases the results of the experiments, before we get to the final chapter to draw a conclusion and touch on further development opportunities.

II. MPTCP AND MPT-GRE IN A NUTSHELL

Numerous multipath solutions operating in different layers of the computer networking model exist [3]; however, in our view, MPTCP is definitely to be regarded as the flagship of this area. MPTCP was standardized in 2013 [4]. Since then, renowned network device manufacturers and corporations developing operating systems have integrated it into their own products (e.g. Cisco, Apple) [5].

Practically speaking, MPTCP is a multipath extension of the traditional TCP, realized via the use of so-called TCP-subflows. This approach enables each physical interface to be assigned a subflow that is responsible for data transfer through that given interface. Earlier publications have shown MPTCP to be quite effective with respect to link-capacity aggregation (see e.g. [6]-[8]). The following figure shows the layered architecture of MPTCP:

```
+-----------------------------------------------+
|                  Application                  |
+-----------------------------------------------+
|                     MPTCP                     |
+-----------------------+-----------------------+
|      TCP subflow      |      TCP subflow      |
+-----------------------+-----------------------+
| IPv4/IPv6 (Physical)  | IPv4/IPv6 (Physical)  |
+-----------------------+-----------------------+
```

Fig.2. MPTCP layered architecture

The MPT-GRE software started development in 2012 at the Faculty of Informatics, University of Debrecen, and has since gone through numerous versions [9]-[10]. This solution also showed effective performance in earlier publications on multipath environments [11]-[15]. Its operation differs from that of both the traditional TCP and the MPTCP. It enables multipath interface mapping via the introduction of a logical tunnel interface. The operation of the interface as discernible to the applications above the tunnel layer is completely unchanged. However, beneath this logical layer, the mapping to physical interfaces happens via utilizing the multipath GRE tunneling technology. The following figure shows the MPT-GRE architecture:

```
+-----------------------------------------------+
|              Application (Tunnel)             |
+-----------------------------------------------+
|                TCP/UDP (Tunnel)               |
+-----------------------------------------------+
|               IPv4/IPv6 (Tunnel)              |
+-----------------------------------------------+
|                  GRE-in-UDP                   |    +-----+
+-----------------------+-----------------------+<--| MPT |
|    UDP (Physical)     |    UDP (Physical)     |    +-----+
+-----------------------+-----------------------+
| IPv4/IPv6 (Physical)  | IPv4/IPv6 (Physical)  |
+-----------------------+-----------------------+
|    Network Access     |    Network Access     |
+-----------------------+-----------------------+
```

Fig. 3. The layered architecture of MPT-GRE

III. MEASUREMENT ENVIRONMENT

We have performed our measurements using the dual-path 10 Gigabit Ethernet environment detailed below. The two servers had the following specifications:

- Gigabyte Z77-D3H motherboard with Intel Z77 chipset.
- Intel Core i7-3770K 3.50 GHz processor with 4 cores and 8 threads.
- 4 X 4 GB 1600 MHz DDR3 SDRAM.
- Intel dual 10 Gigabit Ethernet server adapter.
- Ubuntu 16.04 LTS (XenialXerus) 64-bit operating system with 4.4.0-62-generic Linux kernel module.

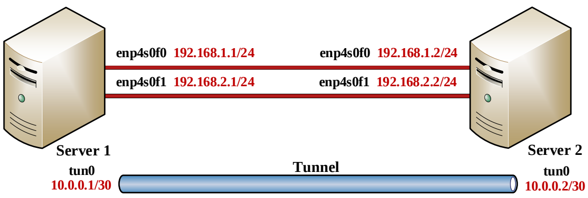

Fig. 4. Measurement environment

To establish a direct connection between the servers, we used two 10 Gigabit Ethernet cables (HP X240 10G SFP+ SFP+ 3m DA cable) and two 10 Gigabit Ethernet Intel Server dedicated network interface cards with two ports each. The motherboard's integrated NIC was used for remote management purposes, and it was always disabled for the entire duration of the measurements. For the MPTCP tests, we installed the latest, 0.95 version (see [2]). The MPT-GRE measurements were carried out using the version available from our development website. All experiments were run on the Linux Ubuntu 16.04 LTS distribution. The effectiveness of both multipath technologies was examined through *iperf3*, CPU utilization and FTP-based measurements. We wrote Python-based bash scripts to automate the process, and repeated each measurement run ten times.

IV. MEASUREMENT RESULTS

A. *iperf3-based Measurements*

We began with the *iperf3* measurements. No physical file download took place during these runs, as the direct data reads and writes were performed between the memories of the two servers. This got rid of the bottleneck of hard drive read/write speed constraints. The results were logged using a program named *tee*.

We performed the different measurements utilizing one interface, and then utilizing two interfaces as well. As it can be seen on Figure 5, increasing the number of interfaces in the 10G environment did not always go hand in hand with an obvious increase in throughput. Using a single interface with MPT-GRE, we managed to reach a throughput of 4.14 Gbps, while using two interfaces resulted in 3.83 Gbps. This is progress compared to earlier measurements (see e.g. [11]-[15]), but does not take advantage of the potential offered by the applied technologies.

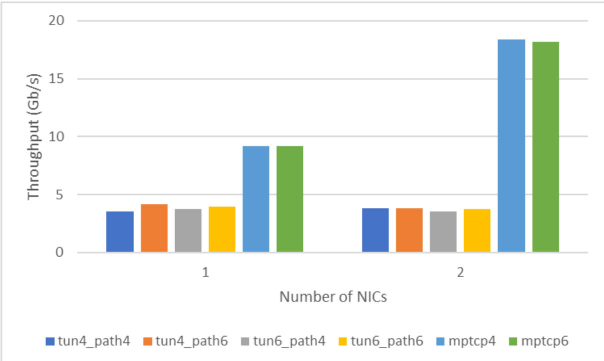

Fig. 5. Testing the MPT and MPTCP *iperf3* throughput performance

The MPTCP achieved noticeably better results with maximum speeds of 9.22 Gbps and 18.4 Gbps respectively. However, in contrast to earlier results, we can notice that MPT did not produce its fastest results during measurements ran in a purely IPv4 configuration, but instead when operating with an IPv4 tunnel over IPv6 physical paths.

B. *FTP Measurements*

The next round of scenarios involved FTP-based measurements. The server on the left of Figure 4. was set up as an FTP server that was used to download a 10GB file to the server on the right.

We have performed the ramdisk configuration using the following script:

```
sudo moun -t tmpfs -o size=11G tmpfs /var/ftp/pub/

cp /var/ftp/10GB.zip /var/ftp/pub/
```

The FTP download process itself was automated as follows:

```
!/bin/bash
#HOST="[fec0::2]"
#HOST="10.0.0.2"
#HOST="192.168.1.2"
HOST="[fec1:300::2]"
wget ftp://$HOST/pub/10GB.zip -O /dev/null
- -report-speed=bits 2>&1
```

Figure 6. shows the performance achieved by utilizing two aggregated paths. We can similarly note in this case as well that the IPv4 tunnel operating over IPv6 paths is the most effective configuration when it comes to MPT results. These settings allowed us to reach 2.97 Gbps during the download of the 10GB file. The other configurations achieved maximum speeds of 2.94 Gbps, 2.83 Gbps, and 2.84 Gbps. The MPTCP performed better in these runs as well, but the difference was not as big compared to MPT results as in the case of the *iperf3* measurements: the maximum speeds here were 5.93 Gbps over IPv4, and 6.13 Gbps over IPv6.

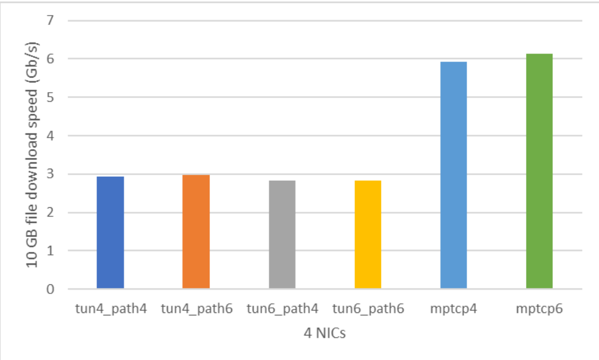

Fig. 6. *Wget* download speed comparison of the MPT and MPTCP using 2 x 10 GE interfaces

We continued with further comparisons of MPT and the MPTCP taking a look at FTP download speeds. First, we performed a baseline measurement of network performance with both MPT and MPTCP disabled. The result we got was 8.78 Gbps, which equated to a download time of 9.11s for the 10GB test file (see Figure 7.).

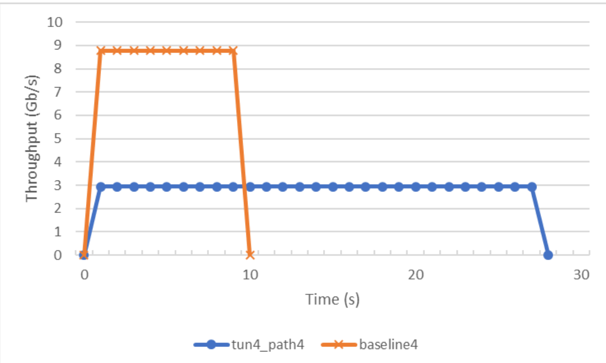

Fig. 7. The MPT-GRE IPv4-IPv4 FTP throughput performance using two 10 Gigabit Ethernet interfaces

However, increasing the number of paths did not bring increased throughput figures. This is discernible during the usage of MPT and MPTCP as well. Using MPT we achieved 2.94 Gbps, resulting in a 27.2s download time. This is close to a threefold increase compared to the baseline measurement.

The MPTCP results took a similar shape (see Figure 8.). In this case, the throughput was 5.93 Gbps, taking 13.49s to transfer the test file using two paths. This means a one and a half times increase in transfer time compared to baseline network performance.

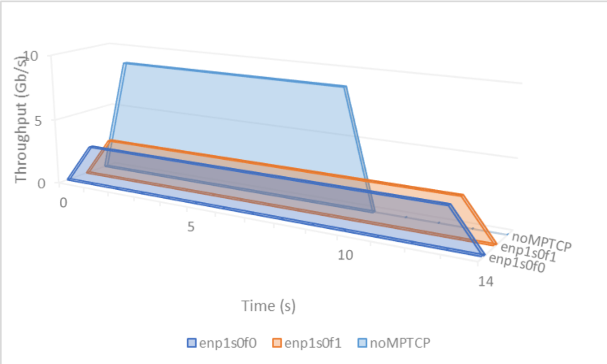

Fig. 8. MPTCP IPv4 FTP throughput performance using two 10 Gigabit Ethernet interfaces

C. *CPU Utilization*

Finally, we examined the CPU resource demand of the MPT and the MPTCP solutions. Figure 9. shows the most critical corner case, namely the CPU performance figures we experienced while operating over IPv6. Naturally, in the other less taxing cases the results were a bit better; however, to keep the paper tidy, those are not presented here.

We can see that CPU utilization per the applied technology roughly reflects the achievable throughput speeds. Processing the data transfer using MPT resulted in a higher average of 23.3% and 23.4% CPU loads while utilizing one and then two paths respectively. The MPTCP was less taxing on the CPU, consuming 8.7% and 13.7% of CPU resources depending on the number of paths.

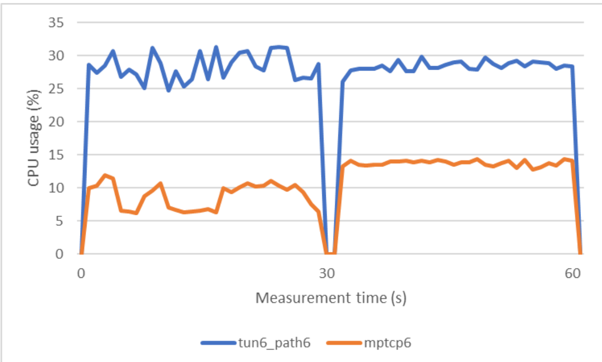

Fig. 9. MPT-MPTCP CPU utilization comparison in a two path 10 Gigabit Ethernet multipath network environment

V. CONCLUSIONS

In this paper, we presented the 10 Gigabit Ethernet dual-path environment efficiency analysis of two technologies supporting multipath communication, which have proved effective in our earlier measurements dealing with Gigabit Ethernet systems (see e.g. [16]-[22]). Both the MPT-GRE and the reference MPTCP solution performed below our expectations. Even though MPTCP managed to beat MPT-GRE in itself with respect to throughput capacity, it still could not provide a sufficient enough performance. With regards to CPU utilization, the MPTCP came out a clear winner against the MPT-GRE technology. In our opinion, these multipath solutions are capable of improving datacenter performance when it comes to Gigabit Ethernet environments; however, 10 GE still seems too tough a nut to effectively crack for both technologies. We believe that further developments of these solutions and the use of more modern server configurations could greatly improve the results that we were currently able to achieve.

ACKNOWLEDGMENTS

This work was supported by the construction EFOP-3.6.3-VEKOP-16-2017-00002. The project was supported by the European Union, co-financed by the European Social Fund.

REFERENCES

[1] "*The MPT-GRE Project official website*", [August 2019] - https://irh.inf.unideb.hu/~szilagyi/index.php/en/mpt/

[2] "The MPTCP Project official website", - http://multipath-tcp.org/ [August 2019]

[3] M. Li, A. Lukyanenko, Z. Ou, A. Yla-Jaaski, S. Tarkoma, M. Coudron, S. Secci, "*Multipath transmission for the internet: A survey*", IEEE Commun. Surveys Tutor., vol. 18, no. 4, pp. 2887–2925, 2016. DOI: 10.1109/COMST.2016.2586112

[4] A. Ford, C. Raiciu, M. Handley, O. Bonaventure, "*TCP extensions for multipath operation with multiple ad-dresses*", RFC6824, RFC Editor, 2013. http://tools.ietf.org/html/rfc/6824

[5] S. Szilágyi, I. Bordán, L. Harangi, B. Kiss, "*Throughput Performance Comparison of MPT-GRE and MPTCP in the Gigabit Ethernet IPv4/IPv6 Environment*", Journal of Electrical and Electronics Engineering, vol. 12, no. 1, pp. 57–60, 2019.

[6] C. Paasch, G. Detal, S. Barré, F. Duchêne, O. Bonaventure "*The fastest TCP connection with Multipath TCP*". http://multipath-tcp.org/pmwiki.php?n=Main.50Gbps

[7] S. Szilágyi, F. Fejes, R. Katona, "*Throughput Performance Comparison of MPT-GRE and MPTCP in the Fast Ethernet IPv4/IPv6 Environment*", Journal of Telecommunications and Information Technology, Vol. 3, No. 2, pp. 53-59, 2018. DOI: 10.26636/jtit.2018.122817

[8] Á. Kovács, "*Comparing the aggregation capability of the MPT communications library and Multipath TCP*", in Proc. 7th IEEE Int. Conference on Cognitive Infocommunications (CogInfoCom), 2016, pp. 157–161. DOI: 10.1109/CogInfoCom.2016.7804542

[9] G. Lencse, S. Szilágyi, F. Fejes, M. Georgescu, "*Internet Draft: MPT Network Layer Multipath Library*", https://tools.ietf.org/html/draft-lencse-tsvwg-mpt-04 [September 2019]
F. Fejes, "MPT source code on GitHub" – [March 2019] https://github.com/spyff/mpt/

[10] B. Almási, G. Lencse, S. Szilágyi, "*Investigating the Multipath Extension of the GRE in UDP Technology*", Computer Communications, vol. 103, issue. C, pp. 29–38, 2017. DOI: 10.1016/j.comcom.2017.02.002

[11] B. Almási, S. Szilágyi, "*Throughput Performance Analysis of the Multipath Communication Library MPT*", TSP 2013 – The 36th International Conference on Telecommunications and Signal Processing, pp. 86-90, 2013. DOI: 10.1109/TSP.2013.6613897

[12] B. Almási, S. Szilágyi, "*Multipath FTP and Stream Transmission Analysis using the MPT Software Environment*", International Journal of Advanced Research in Computer and Communication Engineering, Vol. 2, Issue 11, pp. 4267-4272, 2013.

[13] B. Almási, S. Szilágyi, "*Investigating the Throughput Performance of the MPT Multipath Communication Library in IPv4 and IPv6*", International Journal of Advances in Telecommunications, Electrotechnics, Signals and Systems, Vol. 5, No. 1, pp. 53-60, 2016. DOI: 10.11601/ijates.v5i1.148

[14] S. Szilágyi, I. Bordán, L. Harangi, B. Kiss, "*MPT-GRE: A Novel Multipath Communication Technology for the Cloud*", 9th IEEE International Conference on Cognitive Infocommunications: CogInfoCom 2018 Proceedings, pp. 81–86, 2018. DOI: 10.1109/CogInfoCom.2018.8639941

[15] G. Lencse, Á. Kovács, "*Testing the Channel Aggregation Capability of the MPT Multipath Communication Library*", World Symposium on Computer Networks and Information Security 2014 (WSCNIS 2014), Paper ID: 1569946547, 2014.

[16] G. Lencse, Á. Kovács, "*Advanced Measurements of the Aggregation Capability of the MPT Multipath Communication Library*", International Journal of Advances in Telecommunications, Electrotechnics, Signals and Systems, Vol. 4. No. 2., pp 41-48, 2015. DOI: 10.11601/ijates.v4i2.112

[17] B. Almási, M. Kósa, F. Fejes, R. Katona, L. Püsök, "*MPT: A solution for eliminating the effect of network breakdowns in case of HD video stream transmission*", In Proc. CogInfoCom 2015 Conf., pp. 121-126, 2015. DOI: 10.1109/CogInfoCom.2015.7390576

[18] F. Fejes, R. Katona, L. Püsök, "*Multipath strategies and solutions in multihomed mobile environments*", In Proc. CogInfoCom 2016 Conf., pp. 79-84, 2016. DOI: 10.1109/CogInfoCom.2016.7804529

[19] F. Fejes, S. Rácz, G. Szabó, "*Application agnostic QoE triggered multipath switching for Android devices*", Proceedings of the 2017 IEEE International Conference on Communications: Bridging People, Communities, and Cultures, IEEE Computer Society, pp. 1585-1591, Washington, 2017.

[20] Á. Kovács, "*Evaluation of the Aggregation Capability of the MPT Network Layer Multipath Communication Library and Multipath TCP*", Acta Polytechnica Hungarica, Vol. 16, No. 6, pp. 129-147, 2019. DOI: 10.12700/APH.16.6.2019.6.9

[21] G. Lencse, Y. Kadobayashi, "*Comprehensive Survey of IPv6 Transition Technologies: A Subjective Classification for Security Analysis*", IEICE TRANSACTIONS on Communications, Vol. E102-B, No. 10, pp. 2021-2035, 2019.