A Large-Scale Multipath Playground for Experimenters and Early Adopters

Felicián Németh[†], Balázs Sonkoly[*], Levente Csikor and András Gulyás[†]
HSN Lab, Department of Telecommunications and Media Informatics
Budapest University of Technology and Economics
{nemethf,sonkoly,csikor,gulyas}@tmit.bme.hu

ABSTRACT

Multipath TCP is an experimental transport protocol with remarkable recent past and non-negligible future potential. However, the lack of available large-scale testbeds and publicly accessible multiple paths grossly prohibits the adoption of the technology. Here, we demonstrate a large-scale multipath playground deployed on PlanetLab Europe, which can be used either by experimenters and researchers to test and verify their multipath-related ideas (e.g., enhancing congestion control, fairness or even the arrangement of multiple paths) and also by early adopters to enhance their Internet connection even if they are single-homed[1].

Categories and Subject Descriptors

C.2.6 [**Computer-Communication Networks**]: Internetworking—*routers*

Keywords

Multipath TCP, PlanetLab, OpenFlow, SDN

1. INTRODUCTION

Multipath TCP is a recently standardized experimental transport protocol carrying a future promise for improving the connectivity of its users in terms of resiliency, throughput and cost. The operation of multipath TCP [1] relies on a feature of the underlying network provisioning multiple and hopefully independent (edge-disjoint) paths. The currently used networking paradigm is, however, settled to provide single path service between the endpoints of a TCP connection. In such an environment, it is hard to get multiple paths[2] in a large-scale inter-domain setting, which is clearly an obstacle to experimenting with and adopting the technology. Moreover, installing, configuring and integrating the currently available software components (MPTCP [4], network settings, routing configuration, etc.) consumes a significant amount of time for experimenters. The case of end-users is even worse, since the entry cost of multipath TCP is high (install Linux, install new kernel, additional configuration, etc.) with practically no improvement in their connectivity. All in all, the current situation is by no means promoting the development and adoption of multipath TCP.

In this demonstration, we present our multipath playground deployed on PlanetLab Europe[3] (PLE), which can be used by both experimenters and educated end-users. For experimenters, we give a bunch of tools to easily set up a multipath capable overlay network of PLE nodes with conveniently configurable routing in a few simple steps. For end-users, we deploy a PLE service to supply their multipath TCP with multiple paths. We hope that this framework may catalyze the development and the adoption of multipath TCP.

2. ARCHITECTURE

The architecture and the components of our playground are shown in Figure 1. Besides the regular internet path, we can provide extra paths through a dynamically configurable overlay network on top of PLE. The architecture relies on the brand new OpenFlow [3] capability available in PLE nodes. More specifically, our OpenFlow-based overlay network consists of `sliver-ovs` instances[4] connected by UDP tunnels according to the configured topology. Our special purpose OpenFlow controller is responsible for establishing the network topology and paths given by the experimenter in simple configuration files. Hosts can connect to the overlay network in different ways. First, `qemu` virtual machines can be run on PLE nodes which can connect to `sliver-ovs` via UDP tunnels provided by qemu's socket networking backend[5]. Secondly, hosts operating outside of PLE may configure a special `tap` device and then similar UDP tunnels can be used for connecting to a dedicated `sliver-ovs`. Finally, in case of a host without MPTCP capability, a spe-

[*]MTA-BME Future Internet Research Group
[†]MTA-BME Information systems research group
[1]End-user's host machine does only have one physical upstream link with single IP address
[2]The only usable scenario is when the end-user is dual-homed and tends to use its connections simultaneously, however knowing the habits of the users such situation is extremely rare.
[3]http://www.planet-lab.eu
[4]OpenFlow support in PLE (sliver-ovs) is built around a modified version of the Open vSwitch software switch.
[5]Hosts can connect to both local and remote sliver-ovs instances.

SIGCOMM'13, August 12–16, 2013, Hong Kong, China.
ACM 978-1-4503-2056-6/13/08.

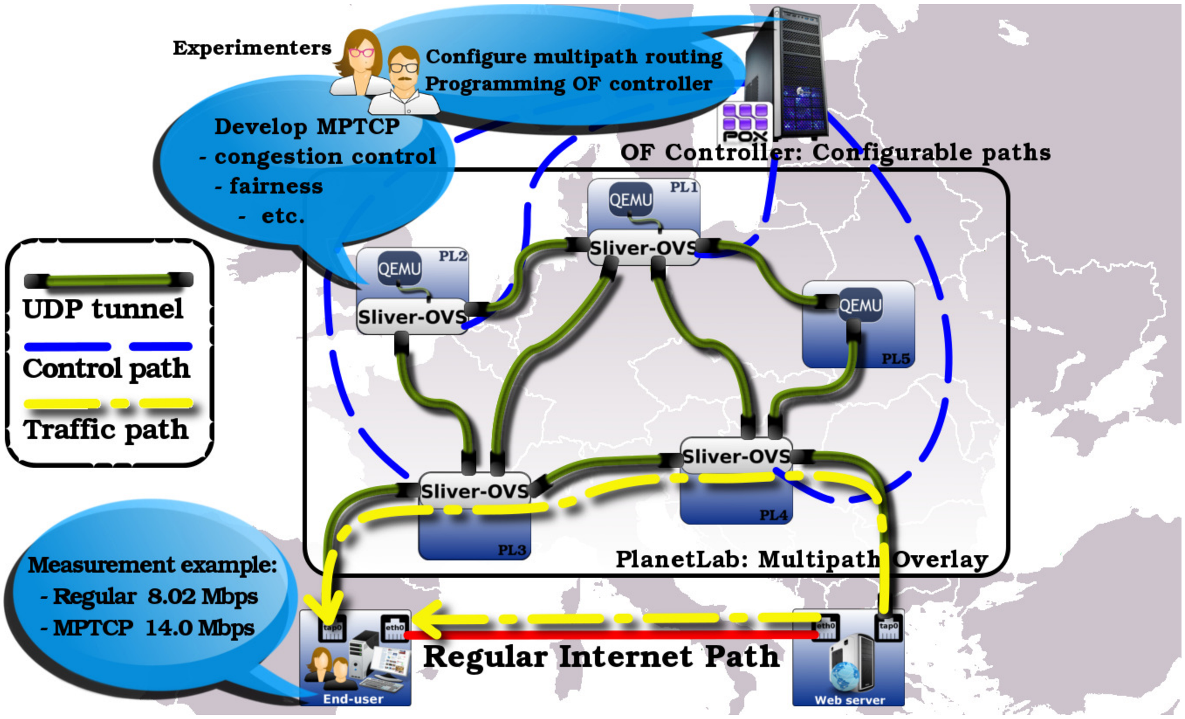

Figure 1: Architecture and Demo

cial proxy can be managed in PLE, which can terminate MPTCP connections from the remote peer and establish a single TCP connection with such host. This proxy has to be placed "close" to the corresponding hosts dynamically posing additional tasks on the control framework. The proxy mechanism clearly has impact on the performance, however, it lets multipath TCP to be used in a fully transparent manner.

For experimenters, we provide a set of tools enabling them to conduct multipath experiments in a few easy steps over this architecture. First, we supply prebuilt, PLE compliant and easily installable `qemu` images equipped with the MPTCP kernel implementation [4]. For promoting research with the multipath TCP implementation itself (e.g., solving the known issues of MPTCP's congestion control in case of large BDP [5,6]), we provide additional tools for deploying alternative virtual machines with experimental multipath TCP variants. Secondly, we give scripts in order to set up a ready-to-run multipath experiment by fully configuring the MPTCP hosts and creating and testing the desired overlay topology. We also provide a POX controller application enabling the experimenter to set up the routes in an intuitive manner and conduct measurements even without knowing much about OpenFlow. It additionally hides the technical details from researchers working on, e.g., theoretical aspects of edge-disjoint path selection. Certainly, more technical savvy researchers can use all the features the OpenFlow overlay network can provide by writing their own controller, and customize our measurement tools, as well.

For early adopter end-users, we offer a PLE service providing extra paths besides the regular internet path they use by default. We supply scripts to these inquiring users, which automatically install MPTCP to their machine and configure additional `tap` devices (besides the existing ethX device(s)) and attach them to our closest service access point in PLE. Then the users can exploit additional paths through our PLE overlay network and thus get the opportunity to improve their connectivity in terms of throughput (Figure 1) and reliability. Albeit, if the bottleneck or the failure is on the last mile, then they won't experience improvements.

During the course of the demo we present the previously introduced use-cases of our multipath playground for both the experimenters and the end-users through on-the-fly establishment of overlay network, configurable experiments and real-time measurements. Our code is available at *http://github.com/nemethf/sigcomm2013*.

Acknowledgements - The research leading to these results has received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement n°287581 - OpenLab, and from the Hungarian National Development Agency through the Economic Development Operational Programme GOP-1.1.1-11-2012-0340.

3. REFERENCES

[1] A. Ford et al. Architectural guidelines for multipath TCP development. IETF RFC 6182.

[2] PlanetLab Europe website. www.planet-lab.eu/.

[3] N. McKeown et al. OpenFlow: enabling innovation in campus networks. *SIGCOMM Comput. Commun. Rev.*, 38(2):69–74, 2008.

[4] Costin Raiciu et al. How hard can it be? designing and implementing a deployable multipath tcp. In *USENIX (NSDI'12), San Jose (CA)*, 2012.

[5] Costin Raiciu et al. Improving datacenter performance and robustness with multipath tcp. In *ACM SIGCOMM 2011*.

[6] Geant-MTA-BME flyer at ONS 2013. http://sb.tmit.bme.hu/mediawiki/index.php/ons2013.