

OFLOPS-Turbo: Testing the Next-Generation OpenFlow switch

Charalampos Rotsos*[¶], Gianni Antichi*, Marc Bruyere^{†‡§}, Philippe Owezarski^{†‡}, Andrew W. Moore*

*Computer Laboratory, University of Cambridge

[†]Université de Toulouse [‡]CNRS, LAAS, France [§]DELL Inc.

[¶] Correspondence author: *firstname.lastname@cl.cam.ac.uk*

Keywords-SDN; OpenFlow; Open-Source; High Performance; Testing; NetFPGA.

I. INTRODUCTION

Research on SDN technologies and primarily its predominant realisation, the OpenFlow protocol, has highlighted a wide range of applications to improve modern network functionality. OpenFlow control applications can improve network management, monitoring and performance, while being backwards-compatible with protocol and network stacks. As a result, many vendors, within only a few years, have integrated production-level support in their products, in a effort to transfer innovative research output to the market. Nonetheless, such continuous network innovation introduces a dilemma for network testing, as relevant platforms remain closed and proprietary and provide limited flexibility. To achieve compliant and functional equipment, effort must be put into all parts of the network-equipment life-cycle, from design to production. Furthermore, the philosophy of the OpenFlow protocol introduces new performance challenges in network device design, dissimilar to the challenges of traditional network equipment, and sets testing flexibility as a primary requirement. Specifically, OpenFlow introduces a reduction in the network control timescales which closely approximate flow control timescales. As a result, the implementation of the OpenFlow switch driver, the interface between the control and data plane, becomes a complex and critical performance bottleneck, influencing the control architecture of the network and, effectively, the overall performance.

In this paper we claim that the effective integration of the OpenFlow protocol in production requires a flexible and high-precision open-source measurement platform which will provide a deep understanding of switch capabilities to users. The openness and flexibility is important in order to establish an evolvable community tool. The OpenFlow protocol, currently defining its 1.3 version, has greatly transformed its design over the years, reflecting both the development and deployment experience of the technology on a constantly widening range of network environments. The OpenFlow community *requires* a performance testing platform capable to co-evolve with the protocol and to support rapid prototyping of experimentation scenarios which highlight the impact of new protocol features. Furthermore, as network link speeds increase

exponentially, meaningful packet-level measurements increase their precision requirements. For example, 10GbE has become the de-facto solution for the aggregation layer of modern datacenter networks, a first-class citizen of the SDN ecosystem. For 10GbE links, a measurement precision in the order of hundreds of μsec is insufficient for OpenFlow applications in the security domain, as it translates into a potential unauthorized transmission of 100KB of sensitive information.

In this paper, we present our effort to enhance the measurement capabilities of the OFLOPS [2] switch performance evaluation framework. More specifically, we present the integration we provide with the NetFPGA-10G platform¹ and the design of the Open Source Network Tester (OSNT²) platform [1]. OSNT enhances OFLOPS with sub- μsec precision and 40 Gbps full bi-directional traffic generation and capturing, providing a highly flexible and open platform for OpenFlow experimentation at high data rates.

II. OFLOPS-TURBO DESIGN

A schematic of the design of the OFLOPS-Turbo design is depicted in Figure 1. The OFLOPS-Turbo architecture consists of a software and a hardware subsystem. The software subsystem runs the core OFLOPS functionality, along with the measurement module of the user. A measurement module contains both the control and data plane functionality of the experiment. The hardware subsystem consists of a series of NetFPGA-10G cards running the OSNT design and is responsible to fulfil the data plane requirements of the experiment. The user can interconnect the OFLOPS-Turbo host with one or more switches in arbitrary topologies and measure with high precision specific aspects of the network architecture, both on the data and control plane. For the rest of the section, we describe the design of the OSNT hardware design and the OFLOPS software architecture.

A. OSNT

The OSNT is a fully open-source traffic generation and capturing system. Its architecture is motivated by limitations in existing network testing solutions: closed-source/proprietary, high costs, inflexibility, and lack of important features such as high-precision timestamping

¹<http://www.netfpga.org>

²<http://www.osnt.org>

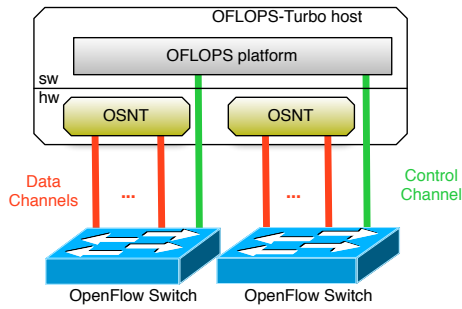


Figure 1. OFLOPS-Turbo design

and packet transmission. Primarily designed for the research and teaching community, its key design goals were low cost, high-precision time-stamping and packet transmission, as well as, scalability. In addition, the open-source nature of the system gives the flexibility to allow new protocol tests to be added to the system. The prototype implementation builds upon the NetFPGA-10G platform – an open-source hardware platform designed to support full line-rate programmable packet forwarding. The combination of Traffic Generator and Traffic Monitor into a single FPGA-equipped device allows a per-flow characterization of a networking system (*i.e.*, OpenFlow-enabled switch) within a single card. Using one or more synchronized OSNT cards, the architecture enables a user to perform measurements throughout the network, characterizing aspects such as end-to-end latency and jitter, packet-loss, congestion events and more.

The OSNT *Traffic Capture* subsystem is intended to provide high-precision inbound time-stamping with a loss-limited path that gets (a subset of) captured packets into the host for further processing. The design associates packets with a 64-bit timestamp on receipt by the MAC module, thus minimizing queueing noise. The timestamp resolution is 6.25nsec with clock drift and phase coordination maintained by a GPS input. The OSNT *Traffic Generation* subsystem consists of several micro-engines, each of which generates traffic according to a given generator function. The traffic generator has an accurate time-stamping mechanism, located just before the transmit 10GbE MAC. The mechanism, identical to the one used in the traffic monitoring unit, is used for timing-related measurements of the network, permitting characterization of measurements such as latency and jitter. When enabled, the timestamp is embedded within the packet at a pre-configured location and can be extracted at the receiver as required.

B. OFLOPS

Measuring OpenFlow switch implementations is a challenging task in terms of characterization accuracy and precision. Predominantly, production-level OpenFlow-switch devices provide proprietary OpenFlow drivers with minimum logging capabilities. Characterising the performance of an OpenFlow-switch requires a black box approach to the problem, where multiple input measurements must

be monitored concurrently in order to characterise the behaviour of the switch.

OFLOPS is an holistic measurement platform which enables the development of custom OpenFlow-based experiments. The platform provides a unified API that allows developers to control and receive information from the data and control channels, as well as, SNMP switch-state information. Experimenters can develop measurement modules on top of OFLOPS, implementing custom OpenFlow applications and measure their performance through the data plane.

Currently OFLOPS provides a wide range of elementary testing modules evaluating the performance of flow actions, flow table management, flow counter extraction, OpenFlow-based packet injection and capturing and detecting potential performance penalties due to OpenFlow operation co-interaction. OFLOPS uses OSNT primarily to control the data plane of the network using a simple interface, provided by the OSNT user-space library. The re-designed DMA engine of the NetFPGA-10G platform improves OFLOPS user-space capturing and transmission capabilities, compared to the earlier OFLOPS hardware design for the NetFPGA-1G platform.

III. CONCLUSIONS

In this paper we have presented OFLOPS-Turbo, an open and flexible OpenFlow testing framework for the next-generation of OpenFlow switches. OFLOPS-Turbo takes advantage of the OSNT hardware design for the NetFPGA-10G platform and provides support for 10GbE traffic generation and capture, coupled with a high precision time-stamping functionality. We believe that OFLOPS-Turbo will provide the tool for the OpenFlow community to better understand the performance impact of OpenFlow implementations and motivate a community of rigorous OpenFlow testing.

Acknowledgements

This work was jointly supported by the EPSRC INTERNET Project EP/H040536/1 and the Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Laboratory (AFRL), under contract FA8750-11-C-0249. The views, opinions, and/or findings contained in this article/presentation are those of the author/ presenter and should not be interpreted as representing the official views or policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the Department of Defense.

REFERENCES

- [1] G. Antichi, M. Shahbaz, Y. Geng, N. Zilberman, A. Covington, M. Bruyere, N. McKeown, N. Feamster, B. Felderman, M. Blott, A. Moore, and P. Owezarski. Osnt: Open source network tester. *IEEE Network*, Special issue on Open Source for Networking: Tools and Applications, 2014.
- [2] C. Rotsos, N. Sarrar, S. Uhlig, R. Sherwood, and A. Moore. Oflops: An open framework for openflow switch evaluation. In *Passive and Active Measurement (PAM)*, 2012.