

Enabling Performance Evaluation Beyond 10 Gbps

Gianni Antichi
University of Cambridge
gianni.antichi@cl.cam.ac.uk

Charalampos Rotsos
Lancaster University
c.rotsos@lancaster.ac.uk

Andrew W. Moore
University of Cambridge
andrew.moore@cl.cam.ac.uk

ABSTRACT

Despite network monitoring and testing being critical for computer networks, current solutions are both extremely expensive and inflexible.

This demo presents OSNT (www.osnt.org), a community-driven, high-performance, open-source traffic generator and capture system built on top of the NetFPGA-10G board which enables flexible network testing. The platform supports full line-rate traffic generation regardless of packet size across the four card ports, packet capture filtering and packet thinning in hardware and sub- μ sec time precision in traffic generation and capture, corrected using an external GPS device. Furthermore, it provides a software APIs to test the dataplane performance of multi-10G switches, providing a starting point for a number of different test cases. OSNT flexibility is further demonstrated through the OFLOPS-turbo platform: an integration of OSNT with the OFLOPS OpenFlow switch performance evaluation platform, enabling control and data plane evaluation of 10G switches. This demo showcases the applicability of the OSNT platform to evaluate the performance of legacy and OpenFlow-enabled networking devices, and demonstrates it using commercial switches.

CCS Concepts

•**Networks** → **Network performance analysis**; *Network measurement*; Network performance modeling; Network experimentation; •**Hardware** → *Hardware accelerators*; Networking hardware;

Keywords

Network Testing; OpenFlow; SDN; High-Performance; OSNT; NetFPGA

1. INTRODUCTION

Computer networks are the hallmark of 21st Century society and underpin virtually all infrastructure in the mod-

ern world. Consequently, society relies on the correct operation of these networks. To make networks more reliable, enormous resources are poured into all phases of the network-equipment lifecycle. The process starts early in the design phase when simulation is used to verify the correctness of a design, and continues through manufacturing and perhaps months of rigorously trials. With a multitude of Internet RFCs and hundreds of IEEE standards, a typical piece of networking equipment undergoes hundreds of conformance tests before being deployed. Finally, when deployed in a production network, the equipment is tested regularly. Throughout the process, a relentless battery of tests and measurement help ensure the correct operation of the equipment. This has led to a multi-billion dollar industry in network test equipment giving rise to companies such as Ixia, Spirent, Fluke, and Emulex/Endace among others.



Figure 1: NetFPGA: the enabling platform of OSNT

However, such equipment has evolved with a number of undesirable characteristics: commonly closed and proprietary systems with limited flexibility well outside the reach of most universities and research laboratories. Researchers and educators would also like to use test equipment to understand current networks and when prototyping new ideas. We therefore built the Open Source Network Tester (OSNT [1]), primarily for the research and teaching community. The first prototype has been conceived on top of NetFPGA-10G platform¹ – an open source hardware platform supporting full line-rate across the four 10G ethernet ports.

OSNT provides two primary distinct functionalities: traffic generation and monitor. The OSNT *traffic monitor* subsystem provides high-precision inbound timestamping 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 minimising queueing noise. The timestamp resolution is 6.25 nsec with clock drift and phase coordination maintained by a GPS input. In addition, the traffic

¹www.netfpga.org

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGCOMM '15 August 17-21, 2015, London, United Kingdom

© 2015 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-3542-3/15/08.

DOI: <http://dx.doi.org/10.1145/2785956.2790036>

capture functionality provides wildcard-enabled packet filters and packet cutting and hashing in hardware. The OSNT *traffic generation* subsystem provides a PCAP replay function with a tuneable per-packet inter-departure time. The traffic generator has an accurate timestamping 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 network measurements, such as latency and jitter. When enabled, the timestamp is embedded within the packet at a preconfigured location and can be extracted at the receiver as required.

The OSNT platform provides a simple and programmer-friendly API to control the traffic generation and monitoring functionality of the OSNT design, enabling the realisation of high precision and throughput measurement tests in software. OFLOPS-turbo is an holistic OpenFlow switch evaluation framework which takes advantage of the OSNT high-precision measurement capabilities. Using OFLOPS-turbo users can develop measurement modules which can access information from multiple measurement channels (data and control plane and SNMP) and measure the impact of the switch OpenFlow implementation in data plane performance with high precision.

OSNT is a highly flexible and function-rich framework, which provides the building blocks for the realisation of a wide range of network test scenarios. By providing an open-source solution we invite everyone from the community to audit (and improve) our implementation as well as adapt it to their needs. We believe that, as an open-source community grows, a low-cost open-source network tester will also prove valuable to the networking industry. We also envisage the enabling of new testing and validation deployments that are simply financially impractical using commercial testers. Such deployments may see the use of hundreds or thousands of testers, offering previously unobtainable insights and understanding.

2. DEMONSTRATION DESCRIPTION

The first part of the demo showcases OSNT and how the original implementation can be used to evaluate the performance of multi-10G legacy switches. The second part of the demo showcases how OSNT can be improved and adapted to different needs. In particular, how the OSNT APIs can be used in order to develop a testing framework to evaluate the performance of multi-10G OpenFlow-enabled switches. Specifically, we present OFLOPS-Turbo [2]: an effort to enhance the measurement capabilities of the OFLOPS [3] switch evaluation framework with support for the emerging protocol requirement.

Part I: Legacy Switch testing.

This part of the demo showcases OSNT main functionalities and how to exploit them to accurately measure the packet-processing latency of a legacy switch under different load conditions. OSNT consists of a software driver supporting command-line and graphic-user interfaces (CLI and GUI), traffic generators and monitors modules. Users can generate traffic at a finely-controlled rate up to 10Gbps per port, and capture high-resolution timestamped packets to evaluate the achievable bandwidth and latency of a network device. Two OSNT ports will be connected to the tested switch. One of the ports will be used to generate traffic at variable rates with the transmission timestamp embedded

in each packet, while the other port will be used to capture packets after they traverse the switch. Packets will be received by a userspace application with transmission and capture timestamps and the application will accurately estimate the switching latency.

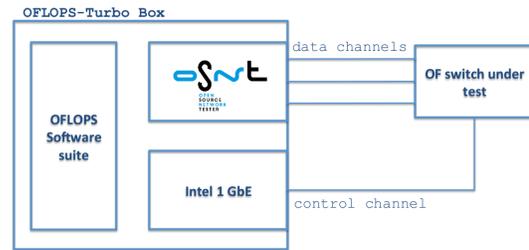


Figure 2: The proposed demo scenario

Part II: OpenFlow Switch testing.

This part of the demo will showcase the ability to leverage the traffic generation and monitoring capabilities of the OSNT platform in software and implement high precision evaluations of a switch control plane functionality. Specifically, we will setup an instance of the OFLOPS-turbo framework on a host and run multiple measurement tests against a production OpenFlow switch. OFLOPS-turbo implements numerous baseline measurement scenarios that target the elementary OpenFlow protocol interactions. During the demo we will focus on a test which measures the latency to modify the entries of the switch flow table through control and data plane measurements, using the topology presented in Figure 2. Through the conducted measurements, we will elaborate on the impact of the control plane implementation on the network performance, as well as on forwarding consistency during large flow table updates.

Acknowledgements

We thank the NetFPGA team and community. This project is jointly supported by the EPSRC INTERNET Project EP/H040536/1, the EPSRC TOUCAN project EP/L020009/1 and by the European Union's Horizon 2020 research and innovation programme under the ENDEAVOUR project (grant agreement 644960).

3. REFERENCES

- [1] ANTICHI, G., SHAHBAZ, M., GENG, Y., ZILBERMAN, N., COVINGTON, A., BRUYERE, M., MCKEOWN, N., FEAMSTER, N., FELDERMAN, B., BLOTT, M., MOORE, A., AND OWEZARSKI, P. OSNT: Open Source Network Tester. *IEEE Network Special issue on Open Source for Networking: Tools and Applications* (2014).
- [2] ROTSO, C., ANTICHI, G., BRUYERE, M., OWEZARSKI, P., AND MOORE, A. W. OFLOPS-Turbo: Testing the Next-Generation OpenFlow Switch. In *ICC* (2015), IEEE.
- [3] ROTSO, C., SARRAR, N., UHLIG, S., SHERWOOD, R., AND MOORE, A. OFLOPS: An open framework for OpenFlow switch evaluation. In *PAM* (2012).