OpenFlow Switch Performance Characterisation in Open Source

Rémi Oudin

ENS Paris-Saclay, Université Paris-Saclay

2017-06-22

Introduction

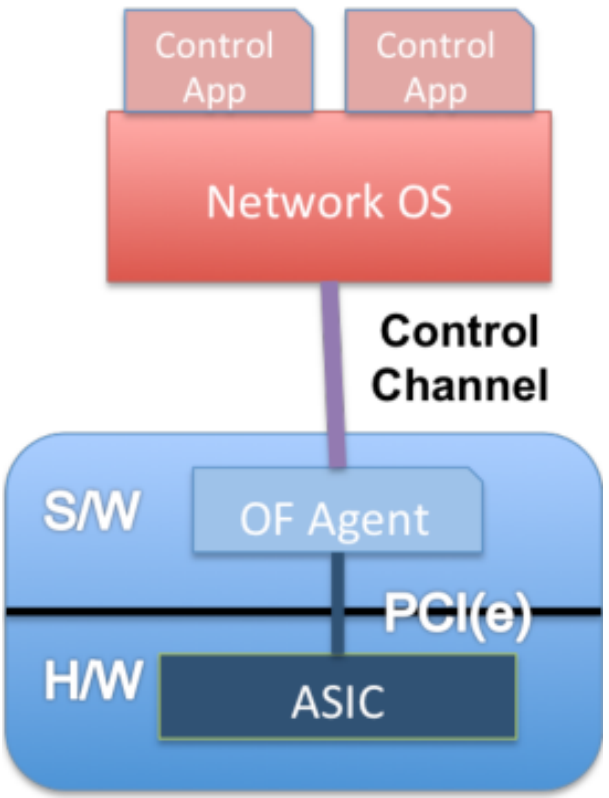

- Control Application complexity
- Control Channel Capacity
- Switch scheduling
- OF Agent implementation
- Driver policy configuration
- ASICs are no more fixed hardware → Programmable ASICs can have an impact on performances.

Introduction

Open-source network measurement

Formerly

- OSNT[1], but no support for the OpenFlow protocol.
- OFLOPS[2], but supported only 1Gb/s datapath, and all measurements done in user space.
- OFLOPS-Turbo[3], limited by hardware resources of old FPGA implementation
- Both were supporting only OpenFlow Protocol 1.0
- We want to go to 10+Gb/s

[1] Gianni Antichi et al. "OSNT: Open Source Network Tester". In: *Network, IEEE* 28.5 (2014), pp. 6–12.

[2] Charalampos Rotsos et al. "OFLOPS: An Open Framework for OpenFlow Switch Evaluation". In: *PAM*. 2012.

[3] Charalampos Rotsos et al. "OFLOPS-Turbo: Testing the next-generation

Let me introduce you to

OMG

==

Open-source Measurement G...

(Do you have any idea for this G ?)

Let me introduce you to

It features:

- An Open-source solution for OpenFlow-enabled switches performance measurement.
- Supports OpenFlow Protocol 1.0 and 1.3.
- Benefits from a "low cost" high capacity hardware support: The NetFPGA SUME board[4].
- User space management of tests.
- Software interface to create test modules.

[4] Noa Zilberman et al. "NetFPGA SUME: Toward 100 Gbps as Research Commodity". In: *IEEE Micro* 34.5 (2014), pp. 32–41. ISSN: 0272-1732. DOI: doi.ieeecomputersociety.org/10.1109/MM.2014.61.

OMG Architecture

- Measurement modules and OMG driver written in C.
- Each test module is compiled as a shared library.
- Event-based system to schedule messages.
- Provide an interface to create new tests.
- Packet generation and capture offloaded to OSNT.
- Hardware tx/rx timestamping.

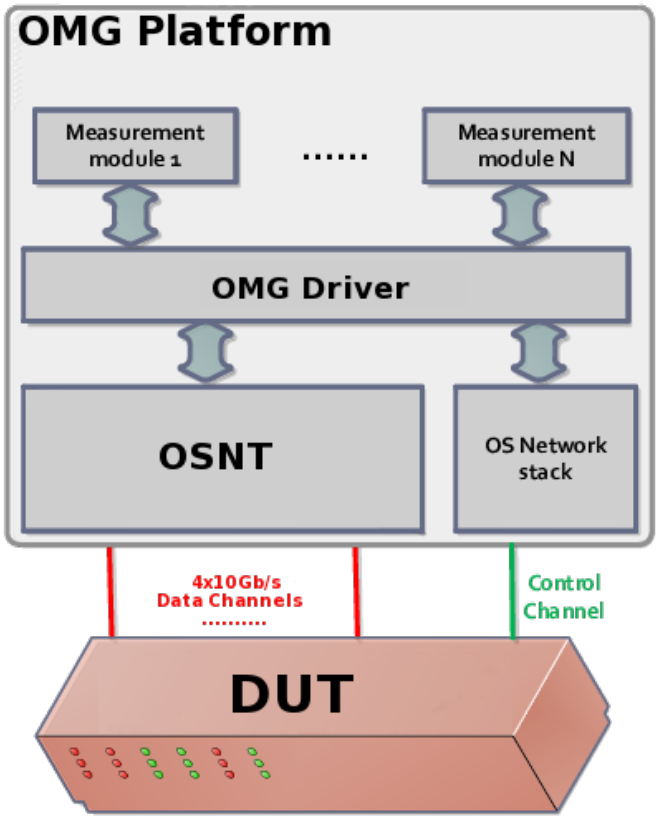

Figure: Global Architecture of OMG

Experiment run

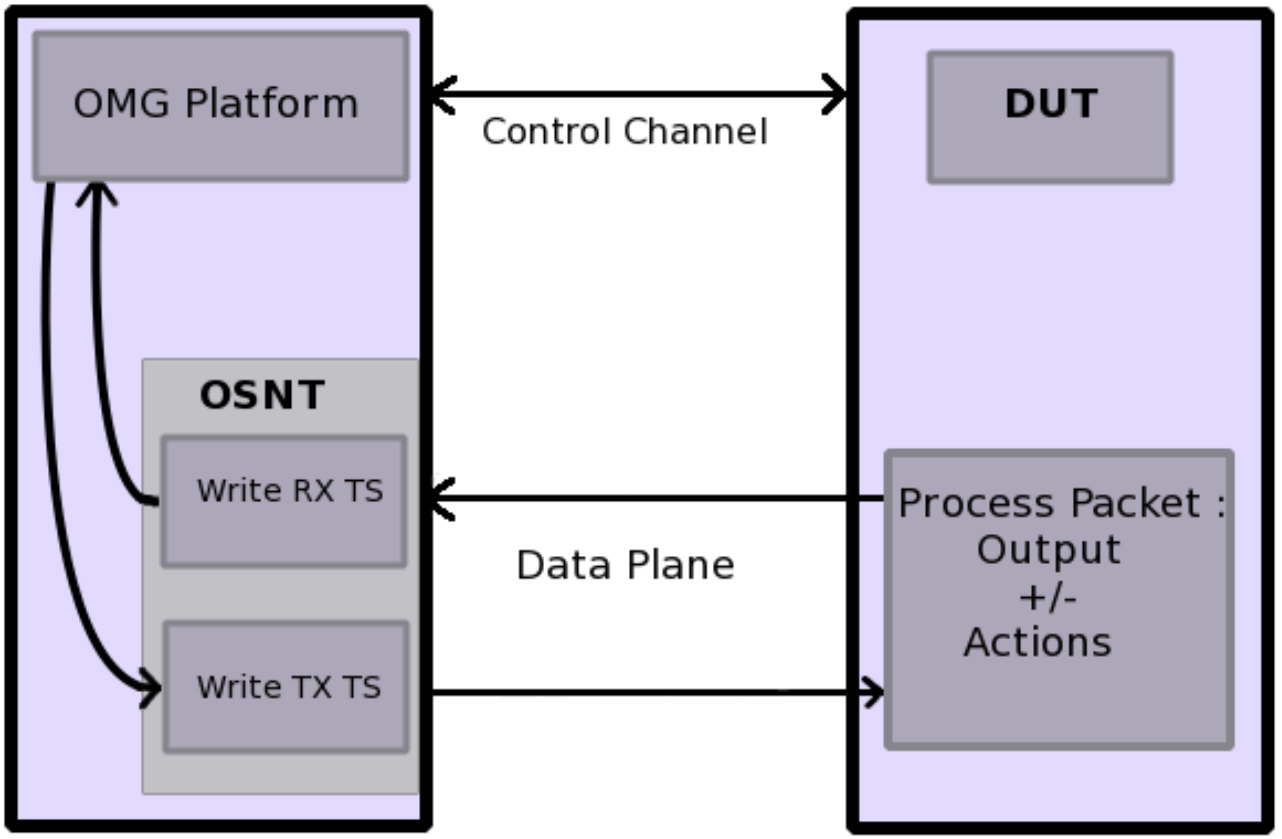

Figure: Graphic of the experiments

First results

Forwarding Latency

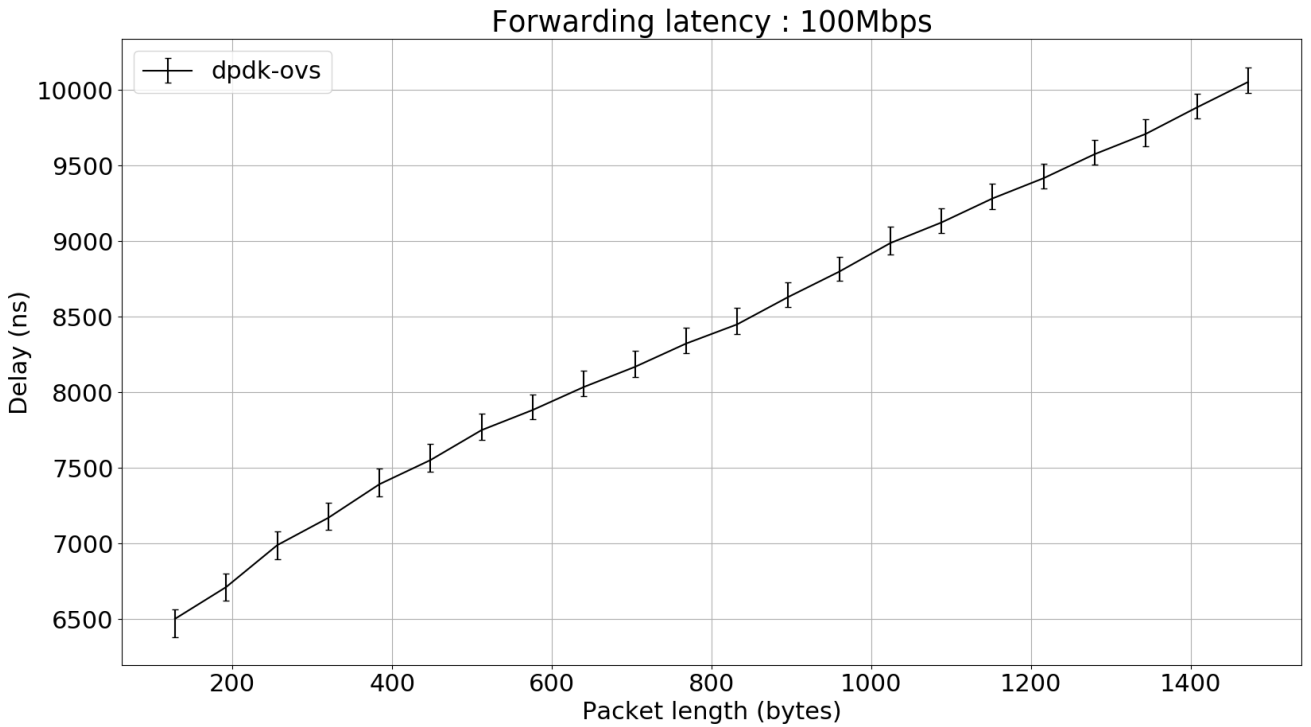

Figure: Forwarding latency at 100Mb/s

First results

Delay created by OpenFlow actions

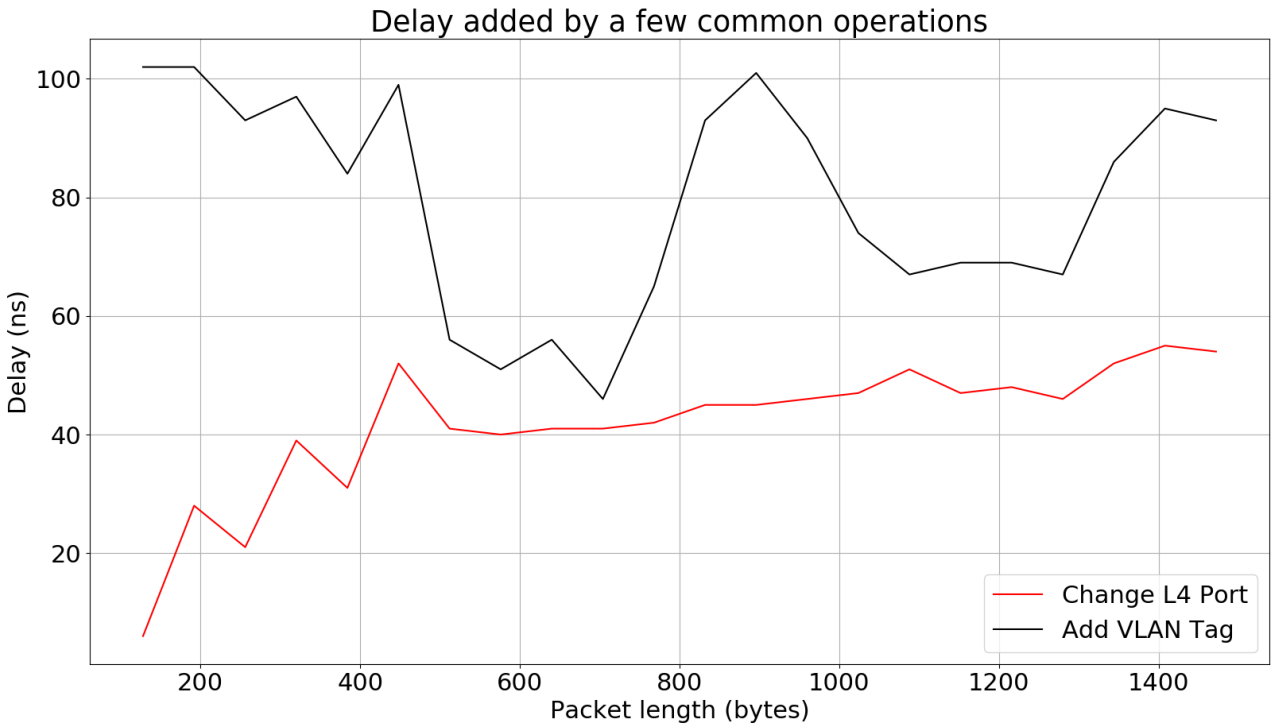

Figure: L4 port modification and adding VLAN id

Why this is nice

- Programmable/Smart NICs are coming for example Netronome[5] and Intel Xeon+FPGA[6]
- Understanding which features could be offloaded in hardware could be awesome.
- Reconciliation between hadware and software ?

[5] Netronome. *Agilio SmartNICs*. https://www.netronome.com/products/smartnic/overview.

[6] P.K Gupta. *Xeon+FPGA Platform for the DataCenter*. https://www.ece.cmu.edu/~calcm/carl/lib/exe/fetch.php?media=carl15-gupta.pdf. 2015.

Conclusion

- The code will be released soon.
- Benefits from Open-source software and hardware solution.
- Feel free to write new test modules !

Thanks!

Gianni Antichi et al. "OSNT: Open Source Network Tester". In: *Network, IEEE* 28.5 (2014), pp. 6–12.

P.K Gupta. *Xeon+FPGA Platform for the DataCenter*. https://www.ece.cmu.edu/~calcm/carl/lib/exe/fetch.php?media=carl15-gupta.pdf. 2015.

Netronome. *Agilio SmartNICs*. https://www.netronome.com/products/smartnic/overview.

Charalampos Rotsos et al. "OFLOPS: An Open Framework for OpenFlow Switch Evaluation". In: *PAM*. 2012.

Charalampos Rotsos et al. "OFLOPS-Turbo: Testing the next-generation OpenFlow switch". In: *2015 IEEE International Conference on Communications (ICC)*. IEEE, June 2015, pp. 1550–3607. DOI: 10.1109/ICC.2015.7249210.

Noa Zilberman et al. "NetFPGA SUME: Toward 100 Gbps as Research Commodity". In: *IEEE Micro* 34.5 (2014), pp. 32–41. ISSN: 0272-1732. DOI: doi.ieeecomputersociety.org/10.1109/MM.2014.61.