

# Experimental Assessment of an SDN-based Control of OPS Switching Nodes for Intra-Data Center Interconnect

F. Agraz<sup>(1)</sup>, W. Miao<sup>(2)</sup>, A. Ferrer<sup>(1),(2)</sup>, G. Bernini<sup>(3)</sup>, H.J.S. Dorren<sup>(2)</sup>, N. Calabretta<sup>(2)</sup>, N. Ciulli<sup>(3)</sup>, J. Perelló<sup>(1)</sup>, S. Peng<sup>(4)</sup>, G. Zervas<sup>(4)</sup>, D. Simeonidou<sup>(4)</sup>, G. Junyent<sup>(1)</sup>, S. Spadaro<sup>(1)</sup>

<sup>(1)</sup>Universitat Politècnica de Catalunya, [spadaro@tsc.upc.edu](mailto:spadaro@tsc.upc.edu), <sup>(2)</sup>Eindhoven University of Technology, <sup>(3)</sup>Nextworks, <sup>(4)</sup>University of Bristol

**Abstract** An SDN-based control framework for OPS switches is experimentally assessed. The OpenFlow messages enabling control communication between SDN controller and OPS switch are validated. Dynamic OPS virtual network slices for intra-data center interconnect are successfully setup and monitored.

## Introduction

Several projects are currently active to re-design today's multi-tier hierarchical architectures for enhancing the scalability, throughput, power efficiency, and latency of the intra-data center network (DCN)<sup>1</sup>. In this respect, the EU FP7 LIGHTNESS project has defined a novel flat intra-DCN architecture integrating Optical Circuit Switching (OCS) and Optical Packet Switching (OPS) technologies<sup>2</sup>. Servers at each rack are interconnected to the hybrid OPS/OCS DCN data plane via the Top of the Rack (ToR) switch, which performs traffic flows aggregation, classification and mapping to either OCS or OPS capabilities (Fig.1). Software Defined Networking (SDN) has been considered as a control framework in support of programmability of network functions and protocols through decoupling data and control planes. In the LIGHTNESS architecture, a centralized SDN-based control plane interacts with the data plane switches through the southbound interface (SB), in order to provide dynamic and on-demand allocation of resources for DC traffic flows. This is performed through a proper configuration and control of the switches via the OpenFlow (OF) protocol<sup>3</sup>.

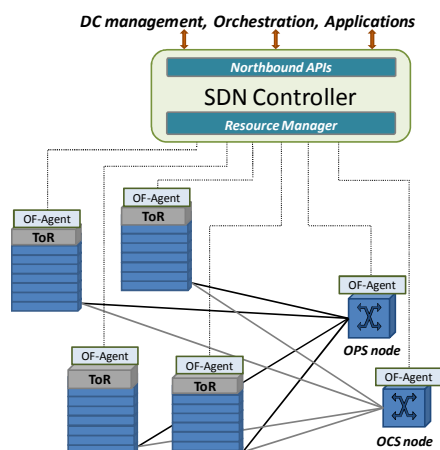


Fig.1: Overall LIGHTNESS architecture

This flexible and highly programmable SDN-based architecture could also enable the

virtualization of the hybrid OCS/OPS DCN and the provisioning of virtual network slices in support of application-to-application connectivity with appropriate QoS guarantee.

The OF protocol extensions for the configuration and control of the OCS switches have been already demonstrated<sup>4</sup>.

In turn, the reconfiguration of virtual DC network slices leveraging statistical multiplexed sub-microsecond flows with QoS guarantees by decoupling the fast OPS switching from the slower control plane operation has also been demonstrated<sup>5</sup>. However, the control employed in that experiment was proprietary. To enable the virtualization of DC network sub-systems by exploiting a centralized SDN control framework, it is crucial to implement the OF protocol also for the OPS switches.

In this paper, we present and experimentally assess the implementation of an SDN-based control framework for the OPS switches. The OpenDayLight<sup>6</sup> (ODL) SDN controller is deployed on top of the OF-enabled OPS switches (with dedicated OF control agents) for switch configuration, monitoring, collection of features/attributes and statistics gathering purposes, performed through the OF protocol.

## SDN-based control for OPS nodes

Fig.2 shows the experimental scenario that consists of one OF-enabled OPS switching node controlled by an ODL SDN controller through a slightly modified version of OF 1.0<sup>3</sup>.

The modular OPS architecture implements a highly distributed FPGA-based switch control for port-count independent reconfiguration time<sup>7</sup>. The RF tones labeling technique allows parallelization of the label bits processing for each input, greatly reducing the label matching time with the Look-Up-Table (LUT) stored in the FPGA-based switch controller. For this experimental assessment, the OPS node has been equipped with an OF agent to let ODL control and monitor the switch. This OF agent,

running in a dedicated server, translates the proprietary interface exposed by the OPS FPGA-based switch controller to the OF protocol, therefore bridging and gluing control and monitoring mechanisms and primitives at the two sides.

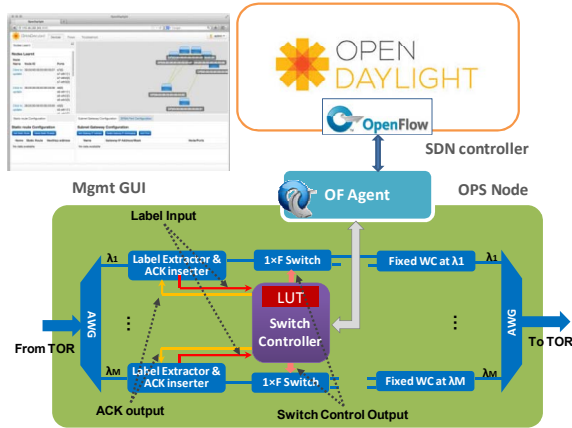


Fig. 2: Control architecture

The SDN controller is implemented by the ODL industry-led collaborative open source platform hosted by the Linux Foundation<sup>6</sup>, in its Hydrogen release (Base edition). ODL has been selected in the context of the FP7 LIGHTNESS project as the reference SDN controller, due to its extensible modular architecture, and the wide set of services available, appliances and northbound primitives for either data center or transport deployments. In this experimental work, the ODL management Graphical User Interface (GUI) has been used to trigger proper actions for the control and monitoring of the OPS switching node, as described in the next section. The actual interaction between ODL and OPS is performed through a slightly

modified version of OF 1.0<sup>3</sup>, implemented at the ODL southbound interface by the dedicated plugin available in the Hydrogen release.

The aim of the experimental work presented in this paper is to assess operation of the developed SDN-based control of OPS within the optical DCN for remote creation and modification of the LUT, and the monitoring of the OPS switch statistics. As a result, this will enable to manage the DCN OPS-based virtual network slices. Due to the specific nature of the OPS technology and operation<sup>7</sup>, the usage of several attributes conveyed in the messages of OF 1.0 has been modified to control the OPS actions through ODL. Such modifications mainly refer to a new identification of OPS switch physical ports and wavelengths in the OF *FEATURES\_REPLY* message<sup>3</sup>, following the OPS hierarchical resource model<sup>7</sup>. A set of mappings of existing OF flow match and action parameters in the *FLOW\_MOD* message structure<sup>3</sup> into specific OPS actions have also been implemented, such as *dl\_vlan* to encode the OPS label and *SET\_NW\_TOS* to encode the virtual slice traffic load, as shown in Fig.3.

### Experimental assessment

According to the control architecture described above, we experimentally assess the communications through the SB interface between the ODL SDN controller and the OPS node, enabled by the OF agent. The latter is used for the configuration of the switch and for collecting statistics of the optical modules (e.g., LUT information, traffic load and retransmission rate). In this way, the SDN controller performs both the remote reconfiguration of the LUT and

**(a)** OpenFlow Protocol details:

- Header: Type: Flow Mod (CSM) (14), Length: 88, Transaction ID: 221426897
- Match: Input Port: 4161 (→ 0x1041), Input VLAN ID: 5 (→ Label)
- Output Action(s): Output to switch port (0), Output port: 4223 (→ 0x107F), IP TOS bits: 50 (→ Load)

**(b)** Flow Entries table:

Flow Name	Node
OPSFlow2	OF 00:00:00:00:00:00:01
OPSFlow1	OF 00:00:00:00:00:00:01
OPSFlow4	OF 00:00:00:00:00:00:01
OPSFlow3	OF 00:00:00:00:00:00:01

**(c)** Flow Detail for OPSFlow1:

Flow Name	Node	Priority
OPSFlow1	OF 00:00:00:00:00:00:01	500

Input Port	Ethernet Type	VLAN ID	VLAN Priority	Source MAC	Dest MAC	Source IP	Dest IP
4161	0x800	5					

**(d)** LUT capacity and active entries:

```
Module with id=1
LUT capacity: 10 entries
Active entries: 4
Label: 5 -> Destination: 1
Label: 6 -> Destination: 2
Label: 1 -> Destination: 1
Label: 2 -> Destination: 2
END OF LUT
Number of InputPorts: 2
```

Fig. 3: Flows allocation process

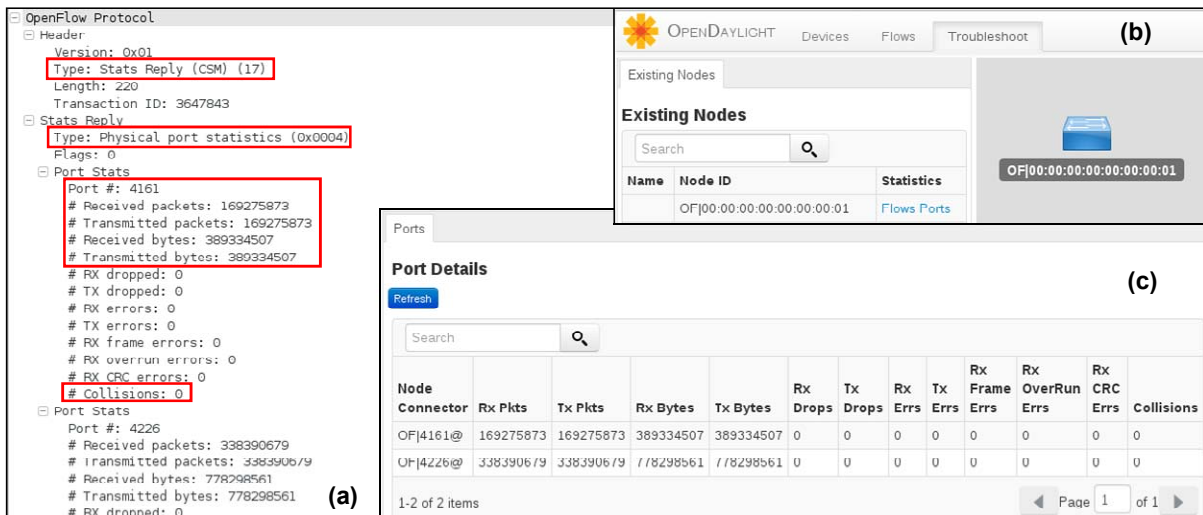


Fig. 4: OPS switch port statistics information

the monitoring of the OPS switch statistics. In particular, the following OF protocol messages have been implemented at the OF agent: 1) *FEATURES\_REPLY* to expose to the controller the capabilities and attributes of the OPS node upon the reception of a *FEATURES\_REQUEST* from ODL; 2) *FLOW\_MOD* to create and delete OPS-based virtual network slice (flow for ODL GUI, as shown in the figures) and, 3) *STATS\_REPLY* to send switch port statistics to ODL when requested by a *STATS\_REQUEST*, for monitoring purposes. A top-down approach has been used; the ODL GUI triggers the request of resource configuration for flows allocation. Fig. 3 illustrates the different steps to control the OPS when a flow has to be created. Once the flow creation command has been executed from the GUI, the SDN controller assigns: 1) the input and output optical switch ports, and 2) the proper label (to be then used by the optical packets). Additionally, in order to prevent excessive packet contention, the allowed load per wavelength can be set. Then, an OF *FLOW\_MOD* message carrying a *New Flow command* (Fig. 3(a)) is sent to the OF agent of the OPS switch that processes the message and configures the OPS switch accordingly. More specifically, the virtual network slice information conveyed by the *FLOW\_MOD* message is used to update the switch LUT (Fig. 3(d)). Fig. 3(b) summarizes the created flows and Fig. 3(c) shows the details of one of the flows, as depicted by the GUI. Moreover, OF messages to communicate statistics per switch port have been implemented at the OF agent. Such statistics are used by the SDN controller for monitoring purposes. As an example, Fig. 4(a) shows the OF *STATS\_REPLY* message; it is sent by the OF agent to the SDN controller upon a request

for gathering the port statistics. In particular, the number of transmitted and received packets (and bytes) is depicted. The optical packets retransmissions due to contentions in the switch are also reported in the collisions field of the OF message. Such statistics are collected by the OF agent from the OPS node which stores them as counters. Finally, Fig. 4(b) and (c), show the port statistics as depicted in the GUI.

**Conclusions**

This experimental work assesses OPS switches dynamically controlled by an SDN controller. This also enables creation/deletion and monitoring of virtual DCN slices through control actions from/to the SDN controller.

**Acknowledgements**

The authors would like to thank the FP7 LIGHTNESS (n°. 318606) and the ELASTIC (TEC2011-27310) projects for supporting this work.

**References**

- [1] C. Kachris et al., "A Survey on Optical Interconnects for Data Centers" Commun. Surveys Tuts., vol.14, (2012).
- [2] J. Perelló et al., "All-Optical Packet/Circuit Switching-based Data Center Network for Enhanced Scalability, Latency and Throughput" IEEE Network, vol. 27, (2013)
- [3] OpenFlow v1.0, <https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow/openflow-spec-v1.0.0.pdf>
- [4] M. Channegowda et al., "Software-Defined Optical Networks Technology and Infrastructure: Enabling SDN Optical Network Operations", JOCN, vol. 5, (2013)
- [5] W. Miao et al., "Demonstration of Reconfigurable Virtual Data Center Networks Enabled by OPS with QoS Guarantees", submitted to ECOC 2014
- [6] OpenDayLight project, <http://www.opendaylight.org/>
- [7] W. Miao et al., "Novel flat datacenter network architecture based on scalable and flow-controlled optical switch system", Optics Express vol. 22, (2014).