

SDN-Enabled Reconfigurable Optical Data Center Networks Based on Nanoseconds WDM Photonics Integrated Switches

Xuwei Xue*, Kristif Prifti, Fu Wang, Fulong Yan, Bitao Pan, Xiaotao Guo, and Nicola Calabretta

IPI-ECO Research Institute, Eindhoven University of Technology, Eindhoven, the Netherlands

e-mail: x.xue.1@tue.nl

ABSTRACT

An SDN-controlled optical DCN based on photonics integrated switch enabling QoS-driven network-slice provisioning and reconfiguration is experimentally assessed. Network-slice can be dynamically reconfigured within 150 ms to maintain specified QoS per network-slice. Controlled by the optical flow control protocol, the deployed photonics integrated switches perform error-free 10 Gb/s switching with < 2dB penalty and zero packet-loss caused by the packet contention.

Keywords: optical data center, software-defined networking, network slicing, photonics integrated switch.

1. INTRODUCTION

With the emerging of cloud computing, Internet of Things (IoT), and the incoming of 5G mobile communications, the traffic communications inside the data centers (DC) are imposed stringent requirements in terms of low latency, high capacity, and high cost and power-efficiency [1]. With the aim to satisfy the scalable growth in both network traffic volume and connected endpoints while decreasing the cost and the energy consumption, transparent optical DC networks (DCNs) based on fast optical switches have been considered, featuring the data rate and format transparency and eliminating the power consuming O/E/O conversions [2,3].

The system-level validations of the fast optical switch node based on developed prototypes have been reported [4, 5], which has demonstrated the promising capabilities for the potential applications in high-capacity and low-latency DCNs. Commercial discrete components are used in the developed prototypes to implement the fast optical switch, whereas the practical implementation would require the integration of hundreds or even more of those optical components, resulting in power-inefficient bulky systems. The realization of photonic integrated circuits (PICs) brings light to this, with the promises of reduced footprint and power consumption [6]. It would be especially beneficial for application scenarios like the DCNs, resulting in diminished networking complexity, easy management and less concerns with the cooling issues.

The optical DCN based on distributed nanoseconds photonic integrated switches featuring with statistical multiplexing and high-throughput allows for low latency, and high capacity and connectivity. However, despite these advantages of photonic integrated switches based DCN, a flexible and reconfigurable network infrastructure that enables the full exploitation of the high network throughput and nanoseconds reconfiguration time is essential. Moreover, with the critical pressure in terms of capacity and interconnectivity from the emerging traffic-boost applications, future DCNs require to allow the deployment of multiple network slices (NS) over the same physical infrastructure to achieve efficient resource utilization with high tenant density, and the maintenance of specified Quality of Service (QoS) per slice needs to be guaranteed for slice consumers [7].

In this work, we propose and experimentally assess a software defined network (SDN) enabled reconfigurable optical DCN based on nanoseconds WDM photonics integrated switches (PISs). Where 4×4 WDM nanoseconds PIS has been designed and fabricated exploiting the modular architecture featured by the fast optical switch node. More than 100 components including the SOAs, AWGs and couplers are integrated in the same chip. An Orchestration plane and a SDN controller is deployed to enable the dynamic provisioning and reconfiguration of

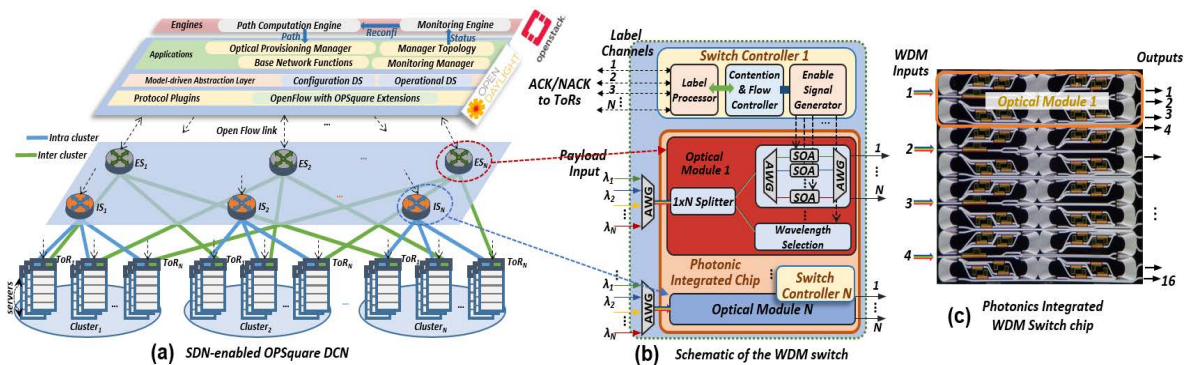


Figure 1: (a) SDN-enabled optical DCN employing photonics integrated switches; (b) Schematic of the photonics integrated switch and switch controller; (c) The fabricated photonics integrated switch chip.

NS services. The real-time network statistics (counts of lost and retransmitted packets) are periodically reported to the SDN control plane which will aggregate them into NS monitoring information and send this aggregated information to the Orchestration layer that is equipped with optimized path calculation and priority class assignment algorithms. Upon decision of the Orchestrator, the SDN controller then re-configures the path to reduce the packet loss and latency, and then achieve a guaranteed QoS.

2. SDN-ENABLED OPTICAL DCN BASED ON PHOTONICS INTEGRATED SWITCHES

The optical DCN with distributed photonics integrated switches and flow control empowered by SDN control plane is depicted in Fig. 1(a). The $N \times M$ intra-cluster photonics integrated switches (ISes) and $M \times N$ inter-cluster photonics integrated switches (ESes) are dedicated for intra-cluster and inter-cluster communication, respectively. The optical flow control between ToRs and IS/ES is implemented to solve possible contention. The flow control signals (ACK/NACK) are generated by the IS/ES controller and sent back to the top of rack (ToR) for releasing the packets or requesting a packet retransmission. FPGAs are used to implement the ToRs, IS/ES switch controllers and the flow control protocol. The OpenDaylight (ODL) platform is set as the base SDN controller connecting the IS/ESes and ToRs by means of integrated OpenFlow (OF) agents implementing an extended OF protocol. These OF agents report network statistics extracted from FPGA registers to the SDN controller, which are in turn aggregated to perform per-NS monitoring. Triggered by the monitored statistics, the SDN controller updates the set of actions of ToRs and photonics integrated switches in real time to provision and reconfigure the NS, depending on the functional characteristics and availability of network resource.

The data plane layout and physical distribution information is stored in the Topology Manager (TM). The Optical Provisioning Manager (OPM) module is developed to configure the underlying devices (i.e., IS, ES and ToR) required to set up the specific network connectivity for NS deployment. The OPM contacts the Path Computation Engine (PCE) of the OpenStack-based orchestrator, which relies on abstracted topological information from the TM to provide the ODL controller with a ToR-to-ToR path computation service that computes the best NS deployment (placement and interconnection). Furthermore, the Monitoring Manager (MM) collects optical data plane statistics and aggregates them into an NS level. Such aggregated information is collected by the Monitoring Engine (ME) of the orchestrator to trigger the needed actions (e.g. reconfiguration) to maintain the expected QoS.

The buffer-less photonics integrated switch is schematically shown in Fig. 1(b). The arrayed waveguide gratings (AWG) groups WDM wavelengths coming from different ToRs and each respective optical module consists of a 1:N splitter to broadcast the WDM channels to the N wavelength selective switches (WSS). The outputs of the N WSSs are connected to the respective N output ports. Each WSS can select one wavelength channel and forward the channel to the output port according to the switching control signals. Turning on/off the N SOAs determines which wavelength channel is forwarded to the output or is blocked. As shown in Fig. 1(b) each optical module forwards the input WDM signals from the ToRs in the cluster by the WSSs to the other ToRs residing in the cluster. Specifically, the $6 \times 4 \text{ mm}^2$ fabricated photonic switch chip in Fig. 1(c) integrates 4 optical modules that can be used to implement 4 ES or 4 IS for 4 ToRs intra-cluster or inter-cluster interconnection. At the input of each module, an $800 \text{ }\mu\text{m}$ booster SOA is employed to compensate the 6 dB losses of the 1:4 splitter and partially the AWGs losses at the WSS. The passive 1:4 splitter is realized by cascading 1×2 multimode interferometer (MMI). Each of the four identical modules processes one of the four WDM inputs and forwards them to the dedicated outputs. Therefore, one photonics integrated switch chips can be used to implement the 2 ES and 2 IS switches to interconnect 8 ToRs grouped in 2 clusters.

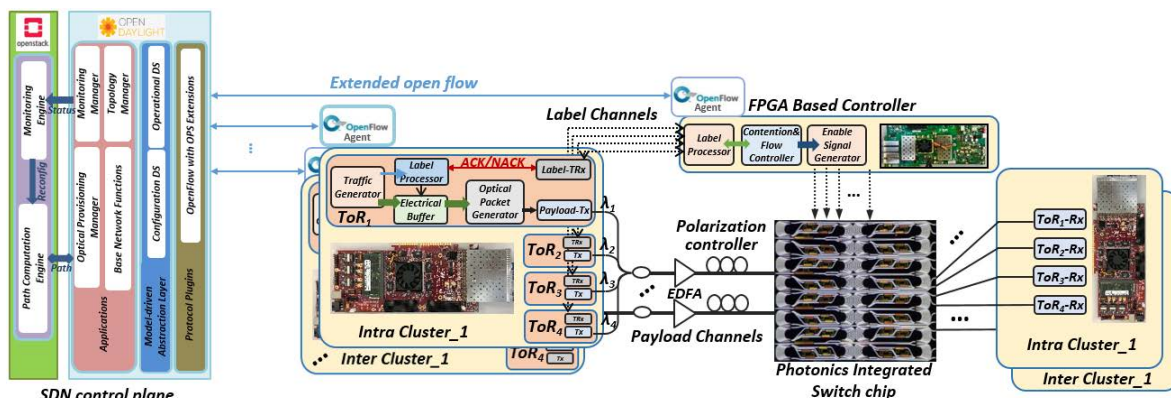


Figure 2. Experimental set-up employed for performance evaluation.

3. EXPERIMENTAL SET-UP AND RESULTS

The experimental set-up to assess the SDN-enabled reconfigurable optical DCN based on the nanoseconds photonics integrated switches is shown in Fig. 2. It consists of 8 FPGA-based ToRs grouped in 2 clusters, connected by 2 ES and 2 IS equipped FPGA-based controller and OF agents. An ODL-based controller is placed over the data plane to control the configuration and updating of the network interconnectivity. The Orchestrator implements the management operations. The first and second modules of PIS are used to implement the intra-cluster switches IS-1, IS-2 that interconnects the 4 ToRs of cluster 1, cluster 2, respectively. The third and fourth modules of PIS are utilized to implement the inter-cluster switches ES-1, ES-2 that interconnects the 2 *i*-th ToRs of these 2 clusters (*i* = 1, 4). Note that given the modularity of the architecture (4 ToRs interconnected by the intra-cluster switch IS, and 2 inter-cluster ToRs connected by the inter-cluster ES), and being all the optical modules of the PIS the same copy, the assessment of one optical module (IS-1) connecting the 4 ToRs in the cluster (cluster 1) is representative for all the other intra-cluster and inter-cluster switch modules operation and performance. The ToRs are implemented by an FPGA that integrates the packets (labels and payloads) generation, the electronic buffers, and the optical Flow Control protocol for each distinct ToR. The FPGA-based ToRs of cluster 1 is equipped with 10 Gb/s SFP transceivers at 1525.0 nm, 1528.9 nm, 1532.9 nm and 1536.8 nm to generate the distinct optical packets for the four ToRs, respectively. The optical packets time slot is 600 ns (540 ns payload time and 60 ns guard time). The four ToRs optical channels are amplified and injected to the first module of PIS via an EDFA, while the optical label signals are sent to the FPGA-based switch controllers. After detecting the optical labels and resolving the packets contention, the switch controller sets the SOA gates to dynamic control the PIS switch.

First, the provisioning and reconfiguration of NS infrastructure are investigated to validate the configurability of optical DCN. Each NS comprises several virtual network function (VNF), and different types of connection (intra-cluster and inter-cluster) among them that can be flexibly reconfigured depending on physical resource distribution and application switchovers. Figure 3(a) shows an example of NS provisioning and reconfiguration, where NS1 consists of VNF1 and VNF2. First, the PCE allocates VNF1 and VNF2 in racks 1 and 8 respectively according to the network resource availability. Next, the OPM coordinates with the PCE to configure the best path among ToR1 and ToR8 to provide connectivity to NFVs. The selected path between ToR1 and ToR8 is IS1<->ToR4<->ES2. After the provisioning, the controller starts the statistics collection to monitor the packets re-transmitted due to collisions in the switch, and the packets lost at the ToR due to buffer overflow. The OF-agent reads the monitored counters from its controlled FPGA devices and then reports the aggregated per-port values to the ODL controller through the extended OF protocol message as shown in Fig. 3(b). Once the ME detects that the packets losses are over a preventive threshold 8000000 (96053079 > 8000000), it triggers the NS reconfiguration through an alternative path, to maintain the requested QoS. The OPM coordinates with the PCE to calculate the new network connectivity among those ToRs (e.g., ES1<->ToR5<->IS2). Figure 3(c) shows the new path configuration information at the flow used to update set of actions inside the ToRs and the optical switches controllers. Then the NS reconfiguration is completed through the new end-to-end connection within 125 ms which is less than the application switchover time. Once the connections are reconfigured, the packet loss will below the threshold and there will be less retransmitted packets as shown in Fig. 3(d). The distributed optical Flow Control protocol utilized to prevent the packets loss is also demonstrated. Figure 3(e) shows the optical labels of ToR 1-4, RequestMessage_NextDestination signals (RM), with the packets destined ports information to the switch controller at every time slot. The FPGA switch controller processes the 4 RM signals and operates the contention resolution protocol to generate the RequestResponse_NextDestination (RR) signals back to the ToRs. Once packet contention happens, as shown in Fig. 3(e), the switch controller sends the ACK

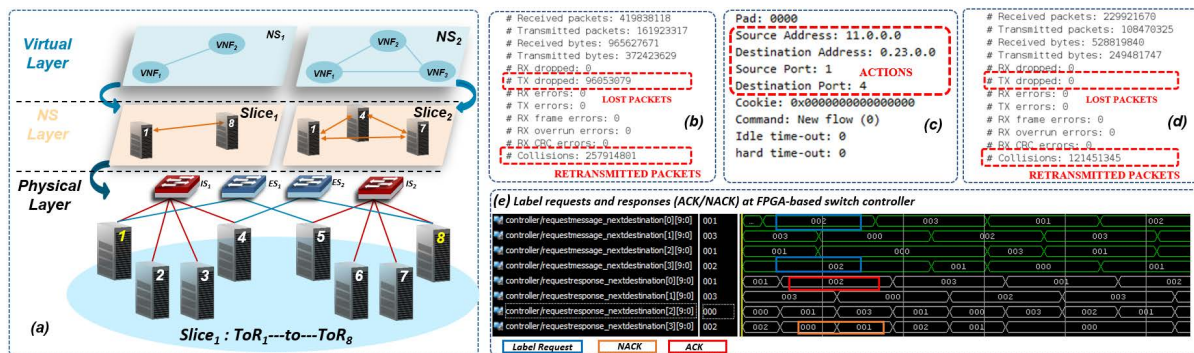


Figure 3: (a) NS provisioning and reconfiguration; (b) Statistics monitoring before NS reconfiguration; (c) OF FlowStats used to update the set of actions; (d) Statistics monitoring after NS reconfiguration; (e) Flow Control signals processed at switch controller.

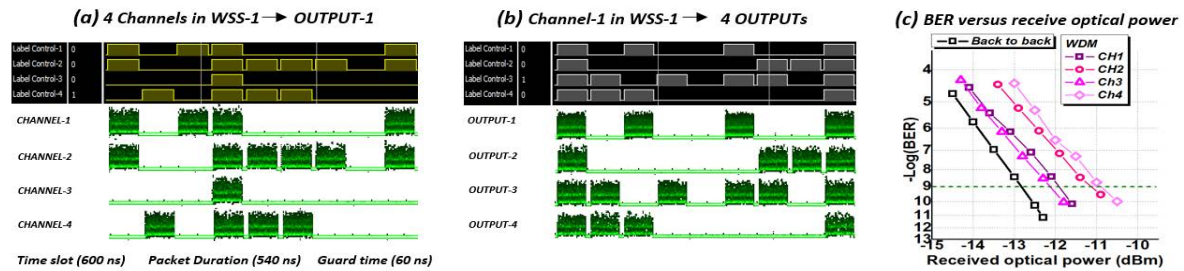


Figure 4: (a) Traces for WSS1; (b) Traces for 4 outputs; (c) BER for WDM channels at 10 Gb/s. CH: channel.

signal (RR = RM) to the ToR1, which means optical packets generated by ToR1 will be forwarded to the destination directly. The NACK signal (RR \neq RM) is send back to the ToR4 to trigger the retransmission of the dropped optical packet.

Second, we assess and validate the dynamic switching operation of the photonics integrated switch. In particular, we assess the switching operation of the WSS1 of the IS-1 module. The switch enabling signals (label control 1-4) generated by the FPGA switch controller after processing the optical labels (RR) of the 4 ToRs, and the channel traces at the WSS1 (output port 1) are reported in Fig. 4(a). The enabling signals are synchronized with the optical packets and a bias current of 40 mA is applied for the “on” state of the SOA gates. The traces indicate that the optical packets are dynamically switched according to the FPGA control signals. Moreover, the dynamic switch control operation of the packets to the four WSSs output ports has also been validated as shown in Fig. 4(b). Packets from ToR 1 has been switched to the 4 output ports by dynamic controlling the SOA gates of the 4 WSSs. Bit error rate (BER) curves for the switched 4 ToRs channel inputs are reported in Fig. 4(c). The back-to-back (B2B) curve is included as reference. Error-free operations with < 1 dB have been measured for ToR 1 (CH 1) and ToR 3 (CH 3) packets, while for ToR 2 (CH 2) and TOR 4 (CH 4) packets the penalty is around 2 dB, but still enough quality to be correctly detected. Those results confirm that the fast dynamic control of the PIS, and validate the switching operation in space, wavelength, and time domain.

4. CONCLUSIONS

We experimentally assess the SDN enabled reconfigurable optical DCN based on photonics integrated switches with dynamic NS provisioning and reconfiguration services. Joint efforts of both Controller and Orchestration layers allow for the best optical path layout for NS provisioning, and the network statistics monitoring dynamically triggers the NS reconfiguration within 125 ms to guarantee the QoS per slice. The photonics integrated switches can be fast and dynamically controlled by the optical flow control protocol implemented between the FPGA-based switch controllers and ToRs. Experimental results indicate that the photonics switch based DCN can switch error-free the 10 Gb/s traffic in space, wavelength, and time domain with no packets loss and < 2 dB penalty.

ACKNOWLEDGEMENTS

The authors would like to thank the H2020 Passion (780326) and H2020 Qameleon (780354) projects for partially supporting this work.

REFERENCES

- [1] Cisco Global Cloud Index: Forecast and Methodology, 2016–2021 White Paper, USA, 2018.
- [2] X. Xue and N. Calabretta: Flexibility assessment of the reconfigurable OPSquare for virtualized data center networks under realistic traffics, in *Proc. ECOC*, pp. 1-3, 2018.
- [3] W. Xia *et al.*: A survey on data center networking: Infrastructure and operations,” *IEEE Commun. Surveys and Tutorials*, 2017.
- [4] F. Yan and N. Calabretta: HiFOST: A scalable and low-latency hybrid data center network architecture based on flow-controlled fast optical switches, *IEEE/OSA JOCN*, vol. 10, no. 7, pp. 1-14, 2018.
- [5] F. Yan and B Pan: FOScube: A scalable data center network architecture based on multiple parallel networks and fast optical switches, in *Proc. ECOC*, pp. 1-3, 2018.
- [6] N. Calabretta and R. Stabile: Photonic integrated WDM cross-connects for optical metro and data center networks, vol. 10946, p. 1094603, *International Society for Optics and Photonics*, 2019.
- [7] X. Xue and F. Wang: Experimental assessment of SDN-enabled reconfigurable OPSquare data center networks with QoS guarantees, in *Proc. OFC*, 2019, paper M3F-4.