SDN Control Framework with Dynamic Resource Assignment for Slotted Optical Datacenter Networks

Giada Landi(1), Ioannis Patronas(2), Kostas Kontodimas(3), Muzzamil Aziz(4), Kostas Christodoulopoulos(3), Angelos Kyriakos(2), Marco Capitani(1), Amirreza Fazely Hamedani(4), Dionysios Reisis(2), Emmanuel Varvarigos(3), Paraskevas Bakopoulos(2) and Hercules Avramopoulos(2)

[1]*Nextworks, Via Livornese, 1027, 56122 Pisa, Italy.*
[2]*National Technical University of Athens, Institute of Communication and Computer Systems, 15780 Zografou, Greece.*
[3]*Communication Networks Laboratory of the Computer Engineering and Informatics Department, University of Patras, Greece*
[4]*Gesellschaft für wissenschaftliche Datenverarbeitung mbH (GWDG), Göttingen, Germany.*
Author e-mail address: g.landi@nextworks.it

Abstract: An SDN control framework is demonstrated enabling slotted operation for dynamic resources assignment in optically-switched datacenters. The demonstration includes the SDN controller with scheduler plugins and north-/southbound interfaces, and the SDN agent communicating to data-plane.
OCIS codes: 060.4250 Networks, 200.4650 Optical interconnects, 060.4251 Networks, assignment and routing algorithms.

1. Overview

Optical switching is gaining momentum as a potential path for gracefully scaling data center networks (DCNs) due to its inherent speed, energy efficiency and transparency to bitrate and protocol. Recent advancements in the data plane involve slotted operation for dynamic allocation of the hardware resources. To reap the benefits of this dynamically reconfigurable data plane, network control within a software-defined networking (SDN) framework remains an outstanding challenge. This is a nontrivial task due to the idiosyncrasy of optical switches, notably concerning lack of buffering as well as limitations in terms of size and speed of current components.

This demonstration showcases an SDN framework capable of controlling an optical DCN with dynamic resource allocation in the optical layer, by means of Time Division Multiple Access (TDMA). As a reference network topology we use the NEPHELE project data plane (www.nepheleproject.eu) that consists of pods interconnected with each other via WDM rings (Fig. 1). Each pod consists of multiple racks interconnected in a star topology through their top-of-rack switch (ToR) with a POD switch being the center of the star. Communication leverages fast tunable lasers at the ToRs and passive wavelength routing (using arrayed waveguide grating – AWG) for intra-pod communication (via the POD switch), whereas space switching is used to add traffic to the ring and active wavelength selective switches (WSS) are used to maintain/drop traffic in/from the ring (involving the intermediate POD switches on the ring).

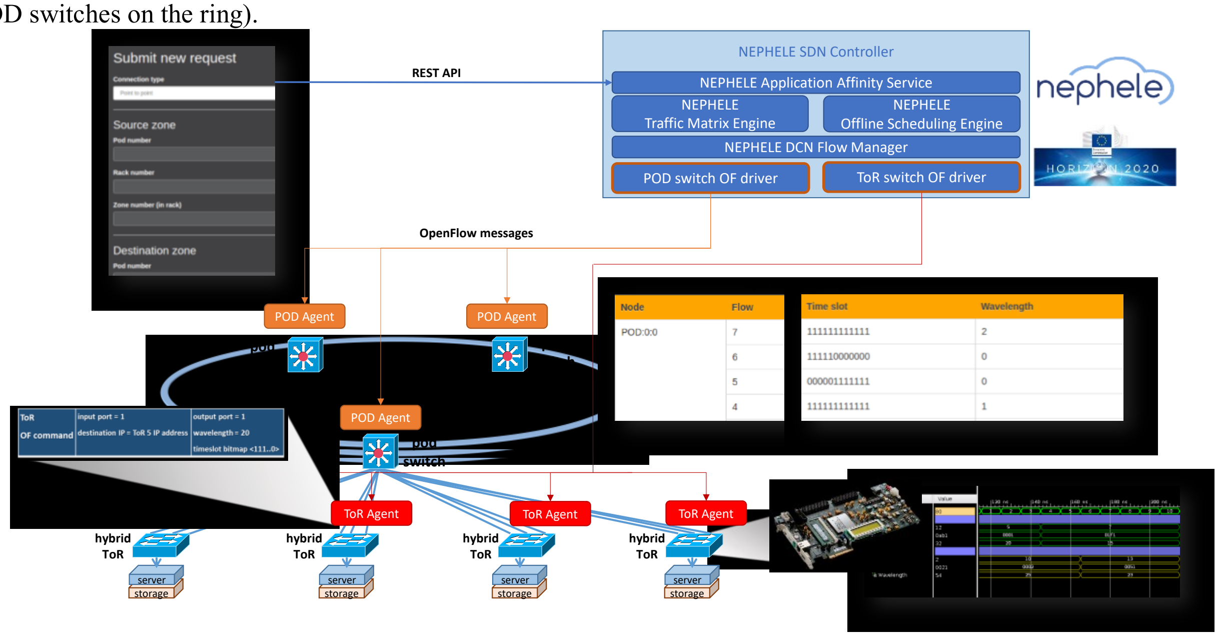

Fig. 1. TDMA-enabled controller with north- and southbound interface, operating over the NEPHELE data plane architecture.

The demo exhibits the NEPHELE SDN controller, an OpenDaylight-based controller extended with an OpenFlow plugin for optical resources, and with SDN applications for dynamic resource allocation in the NEPHELE DCN. An SDN agent is also demonstrated, and its communication with the data plane is implemented.

The NEPHELE controller, shown in Fig. 1, offers a NorthBound Interface (NBI) which allows cloud orchestrators (e.g. OpenStack) to declare the requirements of the virtual connectivity between Virtual Machines (VMs), thus enabling an application-aware DCN configuration. In this demonstration, the controller operates over an emulated optical network with a simplified topology of 2 planes with 3 POD switches/plane organized in a double ring, each of them connected to 4 hybrid ToR switches interconnecting the NEPHELE Innovation Zones. Based on the requests received at the NBI, the controller continuously updates the expected DCN traffic matrix and computes an efficient resource allocation solution exploiting the TDMA paradigm of the NEPHELE DCN.

In the demonstration, the Application Affinity service at the NEPHELE SDN controller is triggered from the DCN management GUI through its REST API, requesting connections with different profiles between source and destination virtual machines residing in the Innovation Zones. The Traffic Matrix Engine aggregates such requests into a Traffic Matrix, representing the global load of traffic that the DCN should carry. The Scheduling Engine then applies a global optimization algorithm to allocate the DCN resources in terms of time slots and planes for that communication (and all traffic represented in the traffic matrix). Finally, the Flow Manager generates the OpenFlow (OF) rules and sends them to configure the devices (PODs and ToRs) according to the resource allocation determined by the Scheduling Engine.

The SDN agent receives the OF commands, translates them and forwards them to the related data plane device. The prototype SDN agent demonstrates this functionality in two different steps. First, it makes a connection with OpenDaylight's southbound driver to receive the NEPHELE specific OF commands (including TDMA slot assignments), parses and translates them. Second, it demonstrates the flow addition process of the agent to the device FPGAs. The flows are sent to the FPGAs over a PCI Express bus interface. The development of the agent is done in Java language, whereas the communication channel between an agent and a device-specific FPGA is established using RIFFA (Reusable Integration Framework for FPGA Accelerators) framework.

2. Innovation

The NEPHELE SDN controller extends OpenDaylight in terms of network resources handled at the south-bound interface and in terms of algorithms, SDN applications, procedures and interfaces to enable application awareness and dynamicity in the configuration of Data Centre Networks. Support for NEPHELE data plane technologies has been implemented through the representation of wavelengths and TDMA timeslots in the OF protocol and in YANG information models used internally in the controller modules. At the SDN applications level, new resource allocation algorithms have been implemented to enable an efficient and fully automated path allocation, enabling QoS guaranteed DCN configuration. Extra effort was put to develop algorithms that would be efficient in large DCN but also perform fast calculations to enable rapid network reconfiguration. In this perspective, the novel NBIs proposed by NEPHELE have a key role since they open the DC network programmability to cloud orchestrators, enabling application-awareness and a finer tuning of resource allocation based on application profiles or support new types of applications that manage and fine-tune their networking requirements to achieve unprecedented performance.

The prototype SDN agent is able to act as a proxy for both legacy Ethernet and novel optical switching devices. This is why it implements the parsing mechanism for both standardized and extended (NEPHELE specific optical extensions) OpenFlow (OF) 1.3 commands. Furthermore, it has the ability to detect out-of-order arrival of control plane commands, and re-order them in a particular schedule before pushing the flows to the device FPGA. Last but not least, the design of the SDN agent is made extensible to accommodate any vendor specific device instructions/extensions in the future i.e., only the translation mechanism needs to be updated in this regard without disturbing the other modules.

3. Relevance to OFC conference

The demo exhibits timely innovations recently attracting the community's attention, extending ubiquitous open-source SDN frameworks with TDMA functionality for dynamic resource allocation. All key aspects are covered, from the SDN controller to the slotted resource assignment algorithms, the north- and southbound interfaces and APIs, the agent and its communication with the data-plane. Thus the demo has strong impact in the topics addressed and poses interest to a broad OFC audience, addressing mainly subcommittees N2, DSN6 and extending to S2.