SDN and Optical Flow Steering for Network Function Virtualization

(Extended Abstract)

Ming Xia, Meral Shirazipour, Ying Zhang, Howard Green, Attila Takacs
Ericsson Research Silicon Valley, San Jose, CA

Introduction

Middlebox traffic steering, the capability of routing traffic through different sequences of middleboxes, has been a popular application in software-define networks (SDN) [5]. In this work, we focus on network function virtualization (NFV) [2][3], especially on core virtualized network functions that process high-volume traffic, such as session border controllers (SBCs) and serving/gateway GPRS support node (SGSN/GGSN). While these core network functions (NFs) run in a virtualized environment in a data center/cloud instead of dedicated hardware, their locations are more dynamic and thus there is a strong need to steer traffic through them dynamically. Similar to [5], traffic will need to be routed across various intermediate NFs before reaching its destination, e.g. at the customer PE. This process is called traffic steering for NFV. Different from traditional middleboxes, the core NFs handle much higher volumes of traffic, and usually don't have the same fine-grained traffic steering requirements that normal middleboxes do. Thus, we argue that existing Layer2/Layer3 packet based traffic steering will have scalability issue for large traffic volume, and may also be associated with high energy cost.

We propose to use software-defined optics to achieve the required flexibility, while reducing the processing load of core packet switches. It also eliminates the need to reinstall multiple packet rules each time a virtualized NF is instantiated or torn down, which can be error prone. We believe the proposal is practical as the use of optics in data centers is becoming prevalent [1]. Assuming high traffic loads, optics can achieve much lower cost per bit than packet switching, meanwhile offering high capacity and energy efficiency. Photonic integration promises further-lower cost per bit of 100GbE and 400GbE interfaces. The dense wavelength division multiplexing (DWDM) and flexible grid technologies allow a single fiber to carry tens of simultaneous non-uniform wavelength channels for ultra-high transmission capacity and spectrum efficiency [4]. The major drawback of optical technology is coarse traffic granularity compared with the packet solutions. However, we argue that not all scenarios require fine-grained traffic steering; aggregated traffic steered in the optical domain may achieve higher throughputs and scalability more efficiently. Moreover, thanks to the increasing agility of optical equipment, the time for establishing wavelength paths is acceptable when considering the time required to instantiate a virtualized NF.

This paper proposes an SDN-based architecture to enable efficient traffic steering in support of NFV. This architecture introduces an optical steering domain joining the packet switching domain, to steer classified packet flows as aggregated flows carried by wavelengths inside the data center connecting to core virtualized NFs. The architecture employs a novel configuration of minimum three inter-connected wavelength selective switches (WSS) and a fiber-loopback scheme. Under the control of a centralized orchestration layer and SDN controller, wavelength paths can be flexibly set up to connect virtualized NFs hosted in different server racks.

Approach Overview

Figure 1 provides an overview of the proposed architecture. The operator-defined policies for traffic steering are configured via the OSS/BSS system, and configurations are pushed down to the network by the SDN controller and the cloud management system. Traffic destined to core NFs can be from an aggregation network which could either directly employ DWDM as its main aggregation/transport mechanism, or be packet based with DWDM transceivers to convert aggregated traffic into wavelengths. The entry point of the optical steering domain is a reconfigurable optical add/drop module (ROADM), which receives DWDM traffic and dispatches wavelengths to either the optical domain, or the conventional packet domain, as predefined by policy. The ROADM avoids the conversion between light and electricity for traffic to/from the optical domain.

The optical domain operates at the DWDM layer to avoid power-intensive operations typically incurred by packet switching. Aggregated traffic carried by a wavelength is treated as a *lambda flow* in the optical domain, which is the basic entity steered based on its central frequency. The right part of Figure 1 shows the optical steering architecture employing three WSSes, which supports flexible wavelength dispatch to its tributary ports with a configuration time around 10s milliseconds. Each NF server is connected to one tributary port of WSS_1 and WSS_2, where WSS_1 switches a lambda flow to the desired NF on its wavelength path, and WSS_2 multiplexes lambda flows on the fiber to WSS_3. One tributary port of WSS_3 (P_5 in this example) is designated as *exit port* to steer traffic out of the optical domain.

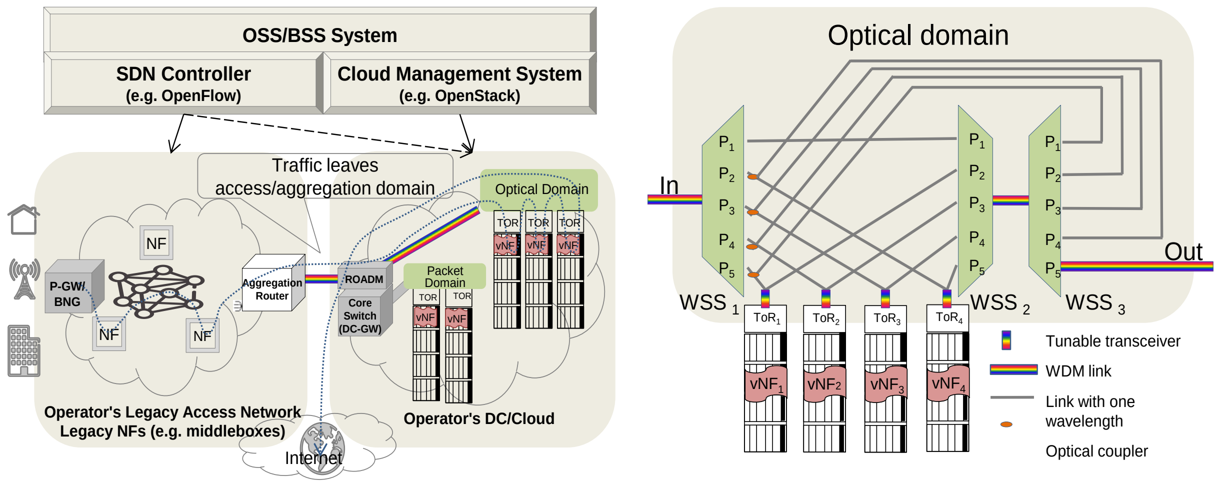

Figure 1 Overall architecture (left), Optical steering domain (right).

The remaining tributary ports loop lambda flows back to other NFs through optical couplers. An optical coupler is a 2x1 passive device, allowing an optical signal to enter the device from either of the two input ports. The three WSSes are controlled by the SDN controller, and coordinated to provide wavelength paths traversing sequences of NFs. The two input ports of an optical coupler cannot simultaneously carry incoming signals, even at different frequencies. After all NF processing in optical domain are complete, lambda flows will be steered back to the ROADM, and will be routed to its destination. Note that it is possible that traffic exits the data center in the form of DWDM signal without going through the core packet switch.

The centralized architecture provides a natural support for SDN. The intelligence of the architecture sits in the controller, which controls all the configurable network elements. The northbound interface above the controller provides an abstraction layer to present the optical infrastructure for resource orchestration. This API also translates operator-specific NF policies into the network for NF mapping and wavelength-path configuration. Coordinated with the aggregation/transport network, the traffic is first classified and grouped. The cloud manager allocates data-center resources for NF instantiation, and the SDN controller allocates the network resources to connect the NFs using a southbound API like OpenFlow. The SDN controller is also responsible for wavelength routing and assignment subject to wavelength availability and contention. The orchestration layer can coordinate a joint optimization between the two for optimal resource utilization and load balancing. Besides flexible traffic steering performed in optical domain, the proposed architecture also supports connectivity for virtual machine (VM) migration. When VM migration is needed, wavelength paths can be dynamically set up between the servers within the same optical domain, or the servers in different domains (packet and optical).

Preliminary Analysis

We study the power-saving benefits of the proposed architecture using a scenario with growing traffic. Flow rates increase from 10Gbps to 100Gbps, and number of NFs per flow increases from 1 to 4. The main benefit is avoiding the repeated power-intensive switching at core switches in the conventional packet-based approach. While packet and optical steering have similar power consumption for the low-demand case in our study, optical steering can save up to 60% power consumption per lambda flow at the highest demand. Given the flexibility enabled by the proposed architecture and the continuous cost reduction promised by silicon photonics, optical traffic steering can effectively complement packet traffic steering for NFV in data centers, especially for bulky or aggregated traffic scenarios.

References

[1] N. Bitar, S. Gringeri, T.J. Xia, "Technologies and protocols for data center and cloud networking", *IEEE Communications Magazine*, vol. 51, no. 9, pp. 24-31, September 2013.

[2] "Network Functions Virtualization, An Introduction, Benefits, Enablers, Challenges & Call for Action", White Paper [Online] http://portal.etsi.org/NFV/NFV_White_Paper.pdf, October 2012.

[3] "Network Functions Virtualisation (NFV) - Network Operator Perspectives on Industry Progress", White Paper [Online]: http://portal.etsi.org/NFV/NFV_White_Paper2.pdf.

[4] ITU-T G694.1, [online] available: http://www.itu.int/rec/T-REC-G.694.1/

[5] Z. A. Qazi, C. Tu, L. Chiang, R. Miao, V. Sekar, and M. Yu, "SIMPLE-fying Middlebox Policy Enforcement using SDN," *Proc. ACM SIGCOMM*, Hong Kong, August 2013.