Load Balancing in SDN Enabled Integrated Packet/Circuit Networks, First Experimental Demonstrations

Raimena Veisllari*, Steinar Bjornstad*†, Kurosh Bozorgebrahimi‡
*Dept. of Telematics, Norwegian University of Science and Technology, Trondheim, Norway,
veisllar@item.ntnu.no
†TransPacket, Drammensveien 126, 0277 Oslo, Norway steinar@transpacket.no
‡UNINETT, Abels gate 5, 7465, Trondheim Norway

***Abstract*—Carriers are under a constant pressure to meet the ever-increasing bandwidth demand. Enhancing the network throughput with reduced total cost is therefore attractive. Integrated packet/circuit networks provide absolute QoS and high throughput efficiency. We demonstrate how load balancing can be performed in the packet part of the network without any impact on the circuit traffic. The circuit paths are reconfigured on wavelength reconfiguration timescale, while packet paths are reconfigured based on feedback of the load on the available paths. This scenario is applicable in an SDN-enabled network with full knowledge of both circuit and packet layers and knowledge on resource utilization of the two layers.**

Keywords—software defined networking (SDN); load balancing; packet; circuit; optical network.

I. Introduction

Carriers struggle to keep pace with the ever increasing bandwidth demand from their customers. While the demand for bandwidth rises, there is not a strong willingness to pay for the additional bandwidth. This motivates for systems and mechanisms that can increase the bandwidth utilization of the network without adding significant costs. The utilization of wavelengths in optical networks is in many cases low due to e.g. temporarily low traffic pressure or bandwidth reserved for resiliency purposes [1], [2]. This bandwidth can be efficiently exploited in Integrated Hybrid Optical Networks (IHON), without any impact to performance of the traffic in the wavelength [3]-[5]. However, the amount of traffic that can be added on a wavelength varies strongly according to the traffic load of the existing traffic in the wavelength. In a network with multiple wavelength paths available through the network, the total amount of available traffic on the different paths will, because of non-correlated statistical variations in the traffic patterns, vary less than for each single wavelength-path. Hence, if a suitable amount of packet traffic can be added on each of the individual wavelength-paths, packet loss on the added traffic can be minimized or avoided while the network utilization can be increased.

Balancing the load across wavelengths in the IHON network in a dynamic fashion motivates the need for monitoring of traffic-load in the wavelengths as well as mechanisms for dynamically directing traffic-flows to wavelength paths not saturated by traffic. In a Software Defined Network (SDN), the controller is able to both monitor and dynamically direct traffic flows in both circuit and packet domains, e.g. [6]-[8]. By adding a load-balancing application communicating with the SDN controller, increased bandwidth utilization in the network is achievable, as demonstrated e.g. in [8] and [9].

In this work we apply SDN to an integrated hybrid optical network demonstrating a dynamic load balancing scheme. The IHON consists of circuit switched wavelength paths used for dedicated transport of customer traffic and a packet switched network applying statistical multiplexing. While traffic in wavelength paths for dedicated customer traffic must receive an absolute QoS with fixed latency and zero packet loss, traffic in the packet switched part of IHON is statistically multiplexed and has less demanding QoS requirements. By combining SDN with IHON we show experimentally how added traffic in the packet switched layer can dynamically be divided on two different wavelength paths, creating a more stable available throughput for the added traffic.

The rest of this paper is organized as follows. An overview of the integrated packet/circuit network architecture is described in Section II. SDN related work and the IHON framework model are presented in section III. The load balancing example as an SDN application is discussed in Section IV. The latest experimental results and their analysis is reported in Section V. Finally future work and conclusions are drawn in Section VI.

II. Integrated Hybrid Optical Networks

Integrated hybrid optical networks, also known as integrated packet/circuit networks, have the distinctive characteristic of transmitting two main traffic classes, namely circuit and packet, on the same physical resource, i.e. wavelength. The wavelengths are provisioned exclusively for the circuit traffic. Hence, this traffic is transported with strict guaranteed quality of service (QoS): no packet jitter, low deterministic delay comparable to the propagation delay and no loss. Provisioning circuits of wavelength granularity leads to the well-known issue of low resource utilization in optical circuit switching and wavelength routed optical networks (WRON). To maximize throughput, IHON detects the idle time-gaps in between the packets of the circuit stream and inserts packets from the best effort traffic class if they fit in the gaps. Thus, the circuit traffic is oblivious to- and not affected by- the additional packet traffic. We refer to the traffic class transported in the

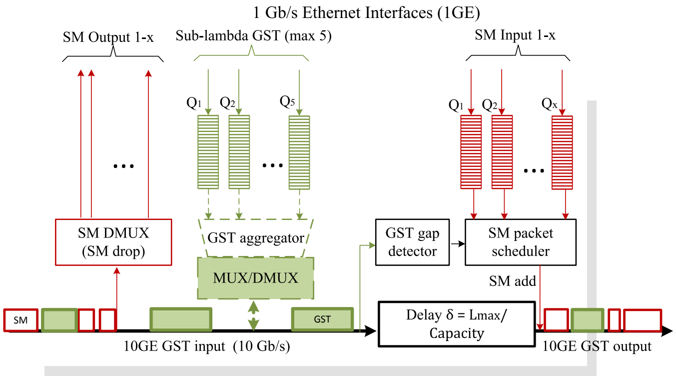

Fig. 1: Schematic diagram of the unidirectional (left to right) transport through the fusion hybrid node with ten 1 Gb/s Ethernet interface for local add/drop and one 10Gb/s Ethernet wavelength. The wavelength can either carry an aggregate GST of five 1GE or one 10 GE GST stream without GST aggregation. Idle time-gaps between GST packets are detected and SM packets are inserted at the output channel.

circuit as guaranteed service transport (GST) and statistically multiplexed (SM) to the low priority added packet traffic.

1) The Fusion solution: Fusion is a technology building on the IHON architecture with the goal of enabling Ethernet packet transport for both types of traffic: GST with circuit QoS and increased resource utilization through SM. Thus, the current solution, the TransPacket H1 node [10], moves from all-optical switching technologies e.g. optical packet switching, towards using standard Ethernet technology over the optical medium.

The schematic diagram of the hybrid node is presented in Fig. 1. The node has two 10 Gb/s Ethernet (10GE) interfaces for the wavelength transport channel. Ten Gigabit Ethernet (GE) interfaces are applied to increase the channel utilization by adding SM traffic or transport GST sub-wavelengths. Packets entering a hybrid node are tagged with a Virtual Local Area Network label (VLAN-ID) indicating the type of service. Any SM packet received at the input 10GE interface is dropped to one of the GE interfaces while the GST packets received at the 10GE interface pass through to the other 10GE interface with absolute priority. At the network level, packets classified as GST are forwarded through the network along an end-to-end dedicated path.

III. SDN-enabled Integrated Packet/Circuit Network

A. Related work on SDN

Software defined networking (SDN) is revolutionizing the networking world once again towards the centralized network control, but most importantly bringing the software actors in central stage on a new networking ecosystem. The last few years have generated an explosive R&D work, involving from network device designers to application layer developers, and related development and deployment by several software and hardware vendors. Recently, optical network equipment vendors are showing interest and joining the network movement towards the SDN controlled network. While SDN controllers come in more than 20 flavors, what is important for openness and standardization is the southbound and northbound application programming interfaces (APIs) [11]-[13]. An illustrative example is how the emergence of the OpenFlow protocol itself, i.e. a standardized southbound API, generated all the available open source controllers.

One of the perspectives of the traditional network vendors on SDN seems to be the ability to allow full programmability. Cisco's recent southbound protocol OpFlex [14] is perceived as a solution for making its southbound semi-proprietary and shifting SDN back in the direction of smart network devices, allowing the equipment vendors as itself to remain an important hub in the networking market. Juniper is offering its own controller and southbound technology, the OpenContrail open source Apache 2.0-licensed technology, while offering open northbound APIs, as well as a router and network analytics engine. Hence, it seems that the network nodes essentially will continue performing the same functions; the next step is to allow for full network programmability allowing the applications to access and modify all the network control information via the API. The application domain can then allow the service providers a more flexible network that can be configured more easily towards their business objectives: e.g. traffic engineering, minimizing CapEx by maximizing resource utilization, minimizing OpEx by routing through the least expensive path, network monitoring, etc.

B. The FUSION scenario

Although OpenFlow is the most well-known protocol for southbound communication and the spark for the SDN evolution, other protocols are being adapted under the SDN umbrella as is the case with the IETF standardized network management protocol, Network Configuration Protocol (NetConf) [15]. It uses the Extensible Markup Language (XML) for communicating with the nodes and making configuration changes. The Fusion H1 nodes are built on NetConf and Yang [16], which is the NetConf-oriented data modeling language, and the combination offers a powerful tool for managing the Fusion nodes. Yang provides a strict XML data model that gives leverage to all layers in the network management systems to use the standard models and to re-use their networking functions code independently of the device model or vendor. In this work is employed the proprietary FusionControl network management system as illustrated in Fig. 2. However, this could be replaced with any SDN controller employing NetConf clients to communicate with the Fusion NetConf agents. The XML schema defined by Yang could then be provided from the Service layer through the northbound API, for example as Java classes through XML-Java binding or, as in FusionControl, a SOAP server API providing a set of Web services, e.g. provisioning the paths, monitoring, alarms, statistics gathering etc.

IV. Load balancing

A load balancing application on the northbound layer can interact with the SDN controller which receives information by the SDN agents on the southbound layer, regarding the resource usage by each traffic type. If SM experiences losses and/or longer end-to-end delays, or the GST traffic instant load increases on the path, the load balancer will compute

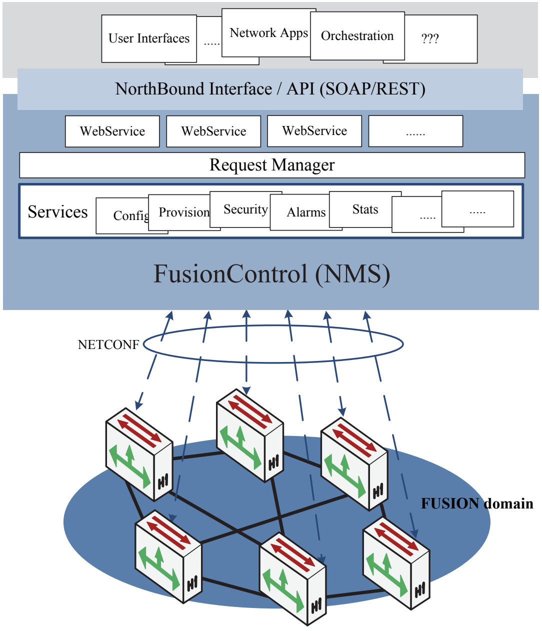

Fig. 2: SDN framework for IHON (the FUSION scenario).

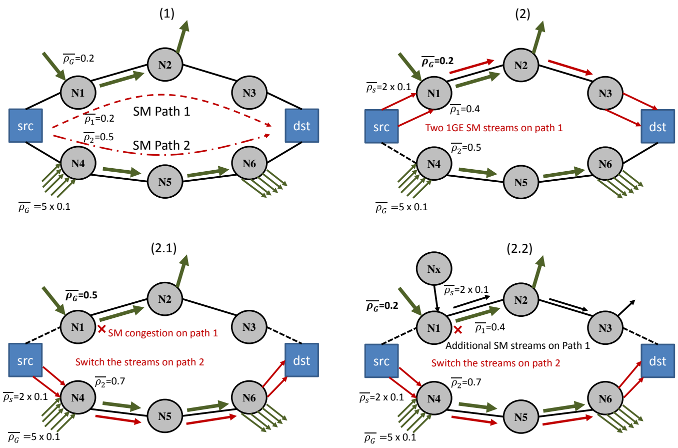

Fig. 3: Load balancing use cases.

a new path or SDN could setup forwarding on an already pre-calculated path.

The integrated hybrid offers a slightly different context on load balancing as the total load is an aggregate of two traffic classes with absolute deterministic guarantees and statistical multiplexed guarantees. In this work, there is neither load balancing nor disruption for GST. The lightpaths are fully provisioned for the circuits, as in a WRON and with the same establishment time-frames. Thus, the increase in the GST load, or SM for that matter, will not affect any of its guaranteed QoS parameters. The SM connections are established and dynamically re-routed based on feedback on the GST utilization of the paths, i.e. effective residual bandwidth. This SM path establishment method allows for network traffic load-balancing, e.g. minimizing processing load on the network elements, or for efficient utilization of the paths, e.g. increasing resource efficiency on paths with low GST average loads. Thus the following benefits are achieved:

- load-balancing in the network with multiple SM routes (paths);
- resource utilization efficiency by increasing the throughput on paths with low GST loads;
- protection of SM traffic when GST peaks, ensuring a minimum QoS;
- SM re-route for additional injected traffic (downstream nodes need to add SM streams);

First, the system needs to be characterized for understanding the performance of GST and SM when the lightpath is dedicated to one 10GE GST connection of 10 Gb/s, and when the lightpath is provisioned for an aggregate of five 1GE GST streams. Next, the SM performance is monitored for load-balancing in the network and SM protection in the lightpath.

V. Experimental Results

A. The testbed setup

In Fig. 4 is illustrated the testbed at the UNINETT lab with four Fusion H1 nodes and one packet generator/analyzer. It emulates the two-path connectivity depicted in Fig. 3, where the packet generator is both source and destination. In the upper branch, equivalent to *path 1*, the GST stream is generated from a 10 Gb/s interface, is forwarded by both nodes N_1, N_2 with absolute priority and minimum processing, i.e. only the VLAN tag, and received back to the tester to be measured. Two additional 1GE SM streams are added on the first node to be inserted in time-gaps between this GST stream on the output channel, i.e. port xe_1 of node N_1. The packet succesfully received at N_2 will be physically looped and sent back to the tester through N_1 which now emulates N_3 in Fig. 3. Hence, the SM streams are virtually passing three nodes N_1, N_2, N_3 before reaching the destination. The lower branch emulates *path 2* in Fig. 3. Five 1GE GST streams are aggregated and transported with absolute priority on the output channel xe_1 of node N_4, looped back at node N_5 and received at the tester through $N_4(N_6)$. Additionally, two 1GE SM traffic streams, identical to the ones on the upper branch, are being transmitted. The nodes are configured and controlled through the FusionControl NMS.

B. System characterization

In Fig. 5 is plotted the average end-to-end delay of all streams on both paths as a function of the offered load normalized for the path capacity of 10Gb/s. To understand the system's limitation on the maximum achievable throughput for different GST traffic patterns, i.e. an aggregate of five 1GE streams as on *path 2* or single packets from a 10GE stream as in *path 1*, we fix the offered load from the two SM streams to its maximum ρ^s=0.2, and vary only the GST load so that $\rho_1^g = \rho_2^g$. Note that ρ_2^g is the sum of the load offered by each GST GE stream on node N_4 normalized on the 10 Gb/s path. The results are plotted for both cases when SM is traversing *path 1* and *path 2*. We observe that the saturation on *path 1* is reached at ρ_1^g=0.45 for $\rho^s = 0.2$, i.e. ρ_1=0.65. At this load, SM starts experiencing losses and high delay in the hundreds of milliseconds, comparable with the buffering delay of the first

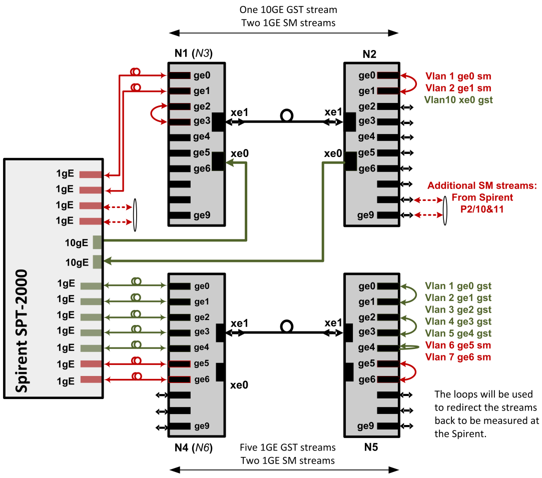

Fig. 4: Network connectivity in the experimental setup emulating the load-balancing scenario with two paths in Fig. 3. The Spirent packet generator is programmed to automatically switch the SM streams between paths if loss is detected.

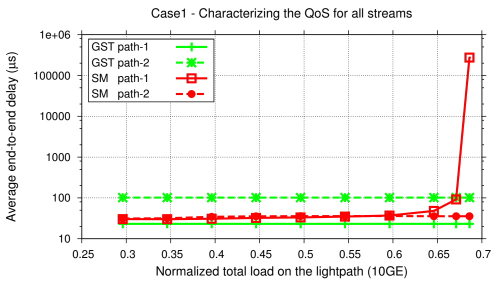

Fig. 5: Characterization of the SM and GST performance with proportional equal contribution to the total load increase.

node; the second link is GST free and SM is not competing with other SM streams. At the same time we see that the GST traffic is transported without being affected by the SM insertion: fixed constant delay, ultra-low jitter in the range of 30ns and no loss. The same is valid for the GST on *path 2*: the delay is constant independently of ρ^g or ρ^s. However, we observe that SM on *path 2* does not saturate at these loads. The reason has been investigated in earlier works and is related with the higher fragmentation of the left-over bandwidth when GST arrives in single packets as compared to bursty arrivals, as is the case for GST aggregation [4], [5].

C. Load balancing

In Fig. 6 and Fig. 7 are illustrated the results on average end-to-end delay for all streams on both paths for the load-balancing use-case examples depicted in Fig. 3, respectively (2.1) and (2.2). For the first case at time t_0 the average GST load on *path 1* is $\rho_1^g = 0.2$. At time t_1 we observe at the destination node that the end-to-end delay of the SM streams is in the range of 200ms and there are losses of 1%. Next statistics from N_1 are gathered and ρ_1^g is reported to be 0.45, thus the system is in saturation. At time t_3 the SM streams are re-routed towards *path 2* where $\rho_2^g = 0.5$ and the SM end-to-end delay is now in the range of 35μs. The GST delay

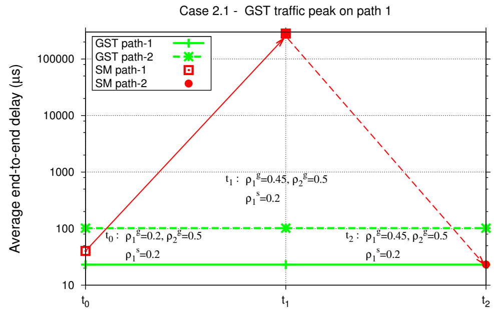

Fig. 6: Average end-to-end delay $[\mu s]$ at time: (t_0) SM is traversing *path 1* and ρ_1^g=0.2; (t_1) GST load on *path 1* increases to ρ_1^g=0.45 as illustrated in case (2.1) and SM experienses packet losses and high delay; (t_3) SM is forwarded on *path 2* where GST aggregate load is ρ_2^g=0.5.

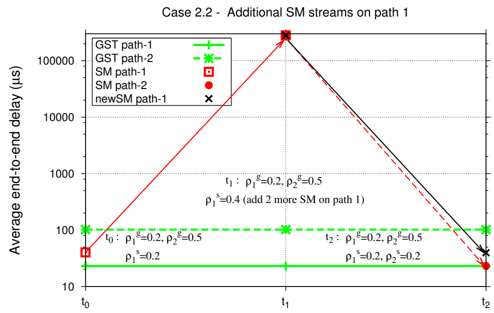

Fig. 7: Average end-to-end delay $[\mu s]$ at time: (t_0) SM is traversing *path 1* and ρ_1^g=0.2; (t_1) two additional streams need to be forwarded on *path 1* as illustrated in case (2.2) and SM experienses packet losses and high delay, ρ_1=0.6; (t_3) SM is forwarded on *path 2* where GST aggregate load is ρ_2^g=0.5 but offers better performance. The network load is balanced out and the network throughput is increased by accepting addiotional traffic on the provisioned GST paths.

on *path 2* however, is not affected by the new SM streams eventhough its delay is higher due to the aggregation process in the H1 nodes and the additional hop. In the (2.2) case at t_1, see Fig.7, ρ_1^g continues to be 0.2 but new SM streams need to be forwarded in *path 1*. Hence these new streams are competing for resources with the source SM streams of interest and node N_1 with a total offered load of $\rho_1 = 0.6$ is in saturation. The source SM streams are re-routed on *path 1* at t_3 and the performance of both source and new SM streams goes back in the normal range, i.e. delay in the range of tens of microseconds.

Scenario (2.1) is emulating load-balancing for the hybrid network and SM protection specifically, while scenario (2.2) is equivalent to re-assigning the network resources as to increase the total network throughput efficiency. For this example in the lab the Spirent traffic generator was programmed to monitor the SM stream performance with 5 min intervals and re-route accordingly, i.e. source re-routing with pre-configured paths. The detection time $\delta_d = t_3 - t_2$ and response time $\delta_r = t_3 - t_2$ of traffic congestion are parameters that can be fine-tuned by the algorithms run at the application layer. However, if strict protection timing is required, i.e. less than 1s, monitoring has to be pro-active at the nodes themselves without involving the controller in real-time.

VI. Conclusions

In this paper we reported the first results on an SDN controlled integrated hybrid optical network, illustrated through load balancing as an application that can successfully be facilitated by SDN. IHON has major benefits from load balancing as SM QoS is sensitive both to the GST load on the path and its traffic pattern. QoS on the SM class can therefore be increased considerably as demonstrated in this paper. Energy saving algorithms, exploiting the daily traffic calendaring (fluctuation) and using intelligent traffic profiling, could be implemented in a network application to profile the daily effective residual capacity leftover from GST and interact with load balancing algorithms for adaptive dynamic SM provisioning and protection, and network traffic load-balancing.

References

[1] A. Saleh, J.M. Simmons, "Technology and architecture to enable the explosive growth of the internet," in *IEEE Comm. Mag.*, vol. 49, no. 1, pp. 126-132, January 2011.

[2] M. Xia, S. Dahlfort, L. Lu, G. Wang, and S. Shen, "CapEx Model and Analysis for Metro Networks: DWDM vs. Packet," in *Proc. of OFC 2014*, paper W2A.12, March 2014.

[3] N. Stol, C. Raffaelli, M. Savi, "3-Level integrated hybrid optical network (3lihon) to meet future qos requirements," in *Proc. of IEEE GLOBECOM 2011*, December 2011.

[4] R. Veisllari, S. Bjornstad, K. Bozorgebrahimi, "Integrated Packet/Circuit Hybrid Network Field-Trial with Production Traffic," in *J. of Opt. Commun. and Netw. (JOCN)*, vol. 5, no 10, Oct. 2013.

[5] S. Bjornstad, R. Veisllari, J. Braute, and K. Bozorgebrahimi, "Integrated Packet/Circuit Hybrid Network Field-Trial Demonstrating Sub-Wavelength Aggregation," in *Proc. of OFC 2014*, paper W4H.4, March 2014.

[6] S. Azodolmolky, R. Nejabati, E. Escalona, R. Jayakumar, N. Efstathiou, D. Simeonidou, "Integrated OpenFlow-GMPLS control plane: An overlay model for software defined packet over optical networks," in *Proc. of ECOC'11*, paper Tu.5.K.5, Sept. 2011.

[7] M. Channegowda, R. Nejabati, M. Rashidi Fard, S. Peng, N. Amaya, G. Zervas, D. Simeonidou, R. Vilalta, R. Casellas, R. Martnez, R. Muoz, L. Liu, T. Tsuritani, I. Morita, A. Autenrieth, J. Elbers, P. Kostecki, and P. Kaczmarek, "Experimental demonstration of an OpenFlow based software-defined optical network employing packet, fixed and flexible DWDM grid technologies on an international multi-domain testbed," in *Opt. Express*, vol. 21, pp. 5487-5498, 2013.

[8] S. Jain et al., "B4: experience with a globally-deployed software defined WAN," in *Proc. of the ACM SIGCOMM 2013*.

[9] R. Doverspike, et al., "Using SDN Technology to Enable Cost-effective Bandwidth-on-Demand for Cloud Services," in *Proc. of OFC 2014*, paper W1E.6, March 2014.

[10] TransPacket whitepaper, "TRANSPACKET H1; a fusion networking add-drop muxponder," [Online]: http://www.transpacket.com/fusion-h1/fusion-ethernet-h1/, last checked May 2014.

[11] K. Hyojoon, N. Feamster, "Improving network management with software defined networking," in *IEEE Comm. Mag.*, vol. 51, no. 2, pp. 114-119, Feb. 2013.

[12] Y. Jarraya, T. Madi, M. Debbabi, "A Survey and a Layered Taxonomy of Software-Defined Networking, in *IEEE Comm. Survey & Tutorials*, Vol PP(99), April 2014.

[13] Enterprise Strategy Group, "Market Landscape Report : The Emergence of Software-defined Network Architectures ", [online], May 2013.

[14] Cisco whitepaper, "OpFlex: Framework for a Broad Partner Ecosystem", [online], April 2014.

[15] R. Enns et. al, "Network Configuration Protocol (NETCONF)," IETF RFC 6241, June 2011.

[16] M. Bjorklund, "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)," IETF RFC 6020, October 2010.