Mayoral, A., Vilalta, R., Muñoz, R., Casellas, R., Martínez, R., Moreolo, M. S., Fàbrega, J. M., Aguado, A., Yan, S., Simeonidou, D., Gran, J. M., López, V., Kaczmarek, P., Szwedowski, R., Szyrkowiec, T., Autenrieth, A., Yoshikane, N., Tsuritani, T., Morita, I., ... Kitayama, K. (2017). Control orchestration protocol: unified transport API for distributed cloud and network orchestration. *IEEE/OSA Journal of Optical Communications and Networking*, 9(2), A216-A222.
<https://doi.org/10.1364/JOCN.9.00A216>

Peer reviewed version

Link to published version (if available):
[10.1364/JOCN.9.00A216](https://doi.org/10.1364/JOCN.9.00A216)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via OSA at <https://www.osapublishing.org/jocn/viewmedia.cfm?uri=jocn-9-2-A216&seq=0&html=true>. Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Control Orchestration Protocol: unified transport API for distributed cloud and network orchestration.

Arturo Mayoral, Ricard Vilalta, Raul Muñoz, Ramon Casellas, Ricardo Martínez, Michela Svaluto Moreolo, Josep M. Fabrega, Alejandro Aguado, Shuangyi Yan, Dimitra Simeonidou, Jose M. Gran, Víctor López, Pavel Kaczmarek, Rafal Szwedowski, Thomas Szyrkowicz, Achim Autenrieth, Norboru Yoshikane, Takehiro Tsuritani, Itsuro Morita, Masaki Shiraiwa, Naoya Wada, Masato Nishihara, T. Tanaka, Tomoo Takahara, Jens C. Rasmussen, Yuki Yoshida, and Ken-ichi Kitayama

Abstract—In the context of the fifth generation of mobile technology (5G), multiple technologies will converge into an unified end-to-end system. For this purpose, SDN is proposed as the control paradigm will integrates all network segments and heterogeneous optical and wireless network technologies together with massive storage and computing infrastructures.

The Control Orchestration Protocol (COP) is presented as a unified Transport API solution for joint cloud/network orchestration, allowing interworking of heterogeneous control planes to provide provisioning and recovery of QoS-aware end-to-end services. E2E QoS is guaranteed by provisioning and restoration schemes, which are proposed for optical circuit/packet switching restoration by means of signal monitoring and adaptive modulation, and adaptive route control, respectively.

The proposed solution is experimentally demonstrated in an international multi-partner testbed which consists of a multi-domain transport network comprising OCS and OPS domains controlled by SDN/OpenFlow and GMPLS control planes and a distributed cloud infrastructure. The results show the dynamic provisioning of IT and network resources and recovery capabilities of the architecture.

I. INTRODUCTION

Software Defined Networking (SDN) is consolidated as the most promising candidate to fulfill the stringent control requirements of the upcoming fifth generation of mobile technology (5G). SDN is aimed to achieve the integration and convergence of all network segments and heterogeneous optical and wireless network technologies together with massive storage and computing infrastructures. SDN promotes the separation between control and data planes through the definition of common abstractions and well defined Application Programming Interfaces (APIs). Many different SDN controllers implementations have emerged in the last few

years, each one with its own proprietary northbound interfaces (NBIs), increasing complexity to design upper control applications which are forced to implement a different protocol library for each SDN controller implementation. Moreover, for complex scenarios where multiple network domains and different transport layers are involved, a single monolithic controller is not a feasible solution. In this context, SDN orchestration has been proposed as a feasible and scalable solution for multi-domain, multi-technology network scenarios to provide end-to-end (E2E) network services, enabling to apply inter-domain coordinated actions based on the intelligent correlation of multiple network information sources.

The need of coordination between different control planes and SDN solutions, has driven the definition of a common transport API, which may serve as a common interface for different SDN controllers (and other control technologies) to enable E2E transport services by reducing complexity of hierarchical orchestration architectures. The OIF and ONF presented the results of their Global Transport SDN Prototype Demonstration [1] last year, where several SDN controllers and hierarchy levels were analyzed, as a result, a Transport API is being standardized at ONF. The Control Orchestration Protocol (COP), proposed in [2], has been developed within the STRAUSS project to address these requirements by offering a common set transport services (i.e., topology service, connectivity provisioning and path computation) shared among multiple control systems.

The COP abstracts a set of control plane functions commonly used by SDN controllers, allowing the interworking of heterogeneous control plane paradigms (e.g., OpenFlow (OF) or GMPLS/PCE). The proposed COP provides a common NorthBound / SouthBound Interface (NBI/SBI) for SDN controllers allowing different SDN controllers can be orchestrated using a single common protocol but also enabling hierarchical orchestration approaches [3]. Moreover, COP enables the joint orchestration of cloud (i.e., computing, storage) and network resources by allowing the provisioning of E2E transport services between Virtual Machines (VM) across multiple network domains.

This paper extends the work in [4]. In summary, it presents a detailed description of the Control Orchestration Protocol and the experimental validation of an integrated orchestration of IT and heterogeneous network resources for E2E transport services' provisioning and recovery with QoS. The E2E recov-

Manuscript received June X, 2016.

Arturo Mayoral, Ricard Vilalta, Raul Muñoz, Ramon Casellas, Ricardo Martínez, Michela Svaluto and Josep M. Fabrega are with CTTC, Spain (e-mail: arturo.mayoral@cttc.es).

Alejandro Aguado, Shuangyi Yan and Dimitra Simeonidou are with University of Bristol, United Kingdom.

Jose Maria Gran and Victor López are with Telefonica I+D, Spain.

Pavel Kaczmarek, Rafal Szwedowski, Thomas Szyrkowicz and Achim Autenrieth are with ADVA Optical Networking.

Norboru Yoshikane, Takehiro Tsuritani and Itsuro Morita are with KDDI R&D Laboratories, Japan.

Masaki Shiraiwa and Naoya Wada are with NICT, Japan.

Masato Nishihara, Toshiaki Tanaka, Tomoo Takahara and Jens C. Rasmussen are with Fujitsu LTD, Japan.

Yuki Yoshida and Ken-ichi Kitayama are with Osaka University, Japan.

Fig. 1: Proposed LIGHTNESS-STRAUSS scenario

ery mechanism has been previously presented in [5], however neither QoS-aware E2E transport was considered nor the COP was employed. In this paper we present the provisioning and QoS recovery results in an experimental multi-partner testbed (Fig.1) with distributed Data Centers (DCs) and multi-domain Optical Circuit/Packet Switching (OCS/OPS) networks, covering Spain, U.K., Germany and Japan. Data plane quality of transmission (QoT) monitoring is integrated into the control plane for the first time to monitor cross-domain connections, which enables both per-domain and E2E QoS recovery.

II. CONTROL ORCHESTRATION PROTOCOL (COP)

The COP provide an abstract and technology neutral definition of a set of control services: provisioning, topology discovery and path computation services. The COP has been designed for the E2E service provisioning across multi-domain, multi-layer and multi-technology network scenarios, to this aim, the modeled objects and operations are not coupled to a specific network technology but they abstract a set of common control functions applicable to different networks. The COP information model is described in YANG modeling language, with RESTconf as transport protocol using JSON encoding for data transmission. Asynchronous notifications of the monitored optical parameters are also modeled in the COP definition models, using Websockets as transport protocol. The open-source implementation of the COP is available for use and contribution in the STRAUSS project Github repository: <https://github.com/ict-strauss/COP>.

In the following subsections we introduce the different services supported by COP which has been defined in different YANG information models.

1) *Call Service*: The COP service call is defined as the E2E provisioning interface. It is based on the concept of Call/Connection separation where a service is supported by a set of effective connections established in the network. The Call object formalize the intent of a service provisioning between two Endpoints, it defines the type of service that is requested or served (e.g., DWDM link, Ethernet E2E transport, MPLS), the Traffic Engineering (TE) parameters requested for the service (e.g., bandwidth, QoS class, latency) and may also include filtering parameters or matches (i.e., MAC, IP or TCP headers) which specifies the granularity of the services to be served.

The Call object also includes the list of effective Connections created in the data plane, to implement the E2E service. A Connection is single network domain scoped, however multiple connections may be established within a single domain. The Connection includes the path across the network topology that the data traverses, which may be fully described or abstracted depending on the orchestration/control schemes used.

The Call service provides the interface for E2E provisioning services but also includes monitoring capabilities for already deployed connections/calls. The COP allow an external entity to subscribe to the notification services of another control entity, through Websockets, to asynchronously receive information about the state of the calls controlled by the control entity.

2) *Topology Service*: The COP allows to retrieve the topological information about the network, by providing a technology agnostic information model for the description of topology elements: Nodes, Edges and EdgeEnds.

The information model is arranged on a tree structure which base node is a Topology object which includes as leafs the list of the nodes and edges included in the related network. A Node must contain a list of ports or edgeEnds and their associated switching capabilities. An Edge object is defined as the connection link between two EdgeEnds. Due to the need of conforming to a common model among different transport network technologies, the definition of the three main objects described (Node, Edge, EdgeEnds) is extensible, featuring polymorphism programming concept, to define specific Node and Edge types. So far, the COP includes the definition of DWDM and Ethernet nodes and edges which defines the technology related switching capabilities for now packets, wavelengths.

3) *Path Computation Service*: The Path Computation service provides an interface to request and return Path objects which contain the information about the route between two Endpoints. These operations are modeled as Remote Procedure Calls (RPCs) in the YANG definition and they are included into the operations subset of the RESTCONF API. Path computation is highly related to the previous group of resources. In the service Call, the Connection object has been designed to contain information about the traversed Path. The Path model is the same in both, the service Call and at the Path Computation. Furthermore each component in the Path object is represented as an Endpoint with TE information associated to it.

Fig. 2: Multidomain SDN Orchestrator ABNO-based architecture

III. INTEGRATED NETWORK AND IT ORCHESTRATION ARCHITECTURE

Fig.1 shows the integrated cloud and heterogeneous network scenario. The proposed scenario includes a distributed datacenter (DC) infrastructure placed in two different locations and managed by different institutions (LIGHTNESS and CTTC). DCs are interconnected by and heterogeneous transport network, consisting on isolated administrative OPS/OCS domains. Each network domain is controlled either by an SDN controller, an Optical Network Hypervisor (ONH) or an Active Stateful PCE.

On top of the control plane the multi-domain SDN Orchestrator (MSO) is responsible of coordinate the different controllers to provide E2E network services [6]. Fig.2 shows

the MSO's architecture based on the IETF ABNO specification [7]. It integrates the COP as southbound interface to communicate with the different controllers. Each domain provides its abstracted topology (node abstraction) through the COP Topology Service. All topological information is gathered by the Topology Manager MSO's component which is responsible to compose the multi-domain abstracted topology. The inter-domain connectivity is pre-loaded into the MSO by static configuration files. Fig.3 shows the topology composed by the SDN orchestrator.

The COP also has been integrated as a northbound interface of the MSO enabling the SDN IT and network orchestrator (SINO) to request E2E connectivity services between the endpoints defined in the abstracted multi-domain topology. The SINO was first introduced in [8] and it is responsible of providing joint orchestration of IT and network resources through the distributed cloud infrastructure and the MSO.

Finally, network service monitoring is enabled at the MSO and the SINO. The MSO architecture includes the OAM Handler component which is responsible of process the COP asynchronous notifications. The MSO subscribes the Notification Servers available on each domain controller through two different services: *updateServiceCall* and *removeServiceCall*. The former, keep the MSO update with changes on the already established calls on every domain, the later notifies the MSO when a call has ceased to be operative. The SINO also subscribes the MSO's *updateServiceCall* notification service to keep up-to-date information about its requested network services.

Fig. 3: Abstracted network and distributed datacenter scenario.

IV. PROOF OF CONCEPT I: DYNAMIC VM AND E2E TRANSPORT SERVICE DEPLOYMENT WITH QoS PROVISIONING

This section presents the experimental validation of the proposed integrated IT and SDN orchestration architecture using COP as unified transport API. The first proof of concept consist on the dynamic creation of two VMs distributed among the two DC locations included in the reference scenario, and the provisioning of a E2E transport service to interconnect the virtual instances. Fig.4a shows the provisioning workflow which involves the orchestration of computing and network resources to offer an E2E system.

E2E QoS capabilities are introduced by the definition of two QoS classes (Table I). Each QoS class defines a certain packet loss rate (PLR) for OPS domains, and a certain OSNR for OCS domains, for a given bandwidth request. The SDN

(a) Workflow

(b) Traffic capture

Fig. 4: Proof of concept I: IT and E2E connectivity provisioning.

orchestrator will convert the high level QoS classes into the corresponding traffic parameters (OSNR, PLR) based on the values (Table I) pre-installed in the SDN orchestrator. Those parameters are included in the *trafficParams* object within the Call request messages to be sent to the domain controllers.

The provisioning of the VMs is requested to each responsible cloud controller implemented by OpenStack. The network service is requested to the MSO through the COP Call service by requesting the creation of a new Call (ID: 1) between the network Endpoints to which the VMs are attached, specifying the QoS class (Fig.5, trafficParams). The MSO computes the path across the different domains and issues a call creation request (IDs: 00001,00002,00003,00004) to the involved domain controllers. Those calls represent the abstract connections provisioned on each domain and, once they are established, they are included in the connections list of the E2E Call. Fig.4b shows the Wireshark captures at the integrated cloud and network orchestrator and at the SDN orchestrator. VM (1 CPU, 20GB disk, 2GB RAM) creation process took around 15 seconds for each instance, the total E2E connectivity service in the current multi-domain scenario is 1.13 seconds.

Fig. 5: E2E QoS transport service request (Call Object).

QoS Class	OCS Domains		OPS Domains	
	OSNR		PLR	
Gold	< 23 dB (PM-16QAM, 224Gb/s)		< 0.1%	
Silver	< 14 dB (PM-QPSK, 112 Gb/s)		< 4%	

TABLE I: QoS classes

V. PROOF OF CONCEPT II: PER-DOMAIN / E2E SERVICE RECOVERY WITH QoS

Fig.6 shows three conducted experiments for QoS recovery: in an OPS domain (scenario A), in an OCS domain (scenarios B, C) and finally E2E QoS recovery (scenario D).

A. Per-domain QoS recovery through adaptive route control in the OPS network

Fig.7 shows the experimental setup of the OPS domain at the NICT premises in Japan. The OPS nodes used are optical packet and circuit integrated nodes [9], including one SOA-based 4 x 4 OPS. In the control plane, an OF-based SDN controller is used to control the OPS nodes. Four OPS nodes

Fig. 6: Per-domain / E2E service recovery with QoS workflow

with optical packet counters are used, including OF agents and OPS transmitters and receivers.

Fig. 7: OPS domain.

The OF agent periodically reads and provides to the SDN controller the optical packet count information that is measured. In this use case (scenario A, Fig.6), two E2E connections are setup involving the OPS domain, flow1 with a packet occupancy rate of 10% and flow2 with a packet occupancy rate of 2%. In this case, the PLR for flow1 measured by a tester is around 4%. When we increase the packet occupancy rate of flow2 from 2% to 6%, the optical packet counter of OPS node 4 reaches the pre-defined threshold indicating packet congestion. Fig.8 shows the measured packet counts of Node 4. With the increase of the packet occupancy rate, the packet count at OPS node 4 is finally smaller than 17000, the OF agent attached to OPS node 4 detects packet congestion and sends an alarm to the SDN controller. The SDN controller receives the alarm and then issues the route adaption for the switching table of node 2, in order to improve the PLR. After the route control, the obtained PLR for flow1 measured by the traffic tester is reduced from around 4% to 0.1%. The route adaptation is announced to SDN orchestrator by an *updateServiceCall* notification to update the call service information with the new *connectionID* which refers to the new intra-domain connection created across the new path.

B. QoS recovery in an OCS domain

For same BER performance, the required OSNR value will relax when a signal with a lower order modulation format is used [10]. Fig.9b shows the tested OSNR vs. BER curve for our 28Gbaud PM-QPSK and PM-16QAM transmitters (Tx's). QPSK requires an OSNR value about 9dB less than 16QAM at HDFEC threshold ($3.8E-3$). Moreover, OSNR monitoring of a circuit flow can detect the OSNR degradation for optical

The figure displays two Wireshark packet capture screenshots. The top screenshot shows a packet list with the following details:

- Frame 1310: 76 bytes on wire (608 bits), 76 bytes captured (608 bits) on interface 0
- Ethernet II, Src: Fujitsu, Dst: 192.168.1.101 (08:00:27:5e:95:25)
- Internet Protocol Version 4, Src: OPCI (192.168.1.25), Dst: OPA (192.168.1.101)
- Transmission Control Protocol, Src Port: telnet (23), Dst Port: 44828 (44828), Seq: 91, Ack: 27, Len: 1
- Telnet
- Data: 029a347b7c7c

The packet count is 43660411, and the packet rate is 48 packets per second.

The bottom screenshot shows a packet list with the following details:

- Frame 1338: 76 bytes on wire (608 bits), 76 bytes captured (608 bits) on interface 0
- Ethernet II, Src: Fujitsu, Dst: 192.168.1.101 (08:00:27:5e:95:25)
- Internet Protocol Version 4, Src: OPCI (192.168.1.25), Dst: OPA (192.168.1.101)
- Transmission Control Protocol, Src Port: telnet (23), Dst Port: 44831 (44831), Seq: 91, Ack: 27, Len: 10
- Telnet
- Data: 02a7f77b7c7c

The packet count is 4434969, and the packet rate is 170000 packets per second.

The figure illustrates the process of route control, which involves selecting the best path for data transmission based on factors such as packet count and packet rate.

Fig. 8: Adaptative Route Control.

links. The receiver-side error-vector-magnitude (EVM) based monitor provides in-band OSNR monitoring without deploying new hardware [11]. With these monitoring information, the COP can orchestrate multi-domain E2E service efficiently and reconfigure the network according to the traffic and link conditions to maintain QoS. Fig.9c shows the performance of OSNR monitor for both QPSK and 16QAM TxS. The monitor notifies the SDN controller, when the OSNR is degraded up to a threshold specified by the multi-domain orchestrator during the provisioning. Then SDN controller reconfigures the link either to use another path (scenario B) or to adopt a lower order modulation format signal with a multi-format transceiver (scenario C). Fig.9a shows the experimental setup for the OCS domain. At node A, 10 carriers with QPSK or 16QAM are launched into the optical network. The 28Gbaud QPSK/16QAM transmitter with central wavelength at 1548.9nm is connected to the OPS domain with an integrated OPS/OCS interface. The signals are transmitted 175km from node A, through node B toward D, and demultiplexed at node D for coherent detection. The coherent receiver deploys an EVM-based OSNR monitor. ASE noise is added between node A and B to emulate the OSNR degradation. By adding more ASE noise to the link, the monitored OSNR will decrease to a threshold of 23dB. Then the monitor instructs the SDN controller to reconfigure the link. In scenario B, the SDN controller reconfigures the optical path from node A to node D to the direct route with a reduced distance link. In this scenario, the total link bitrate will not change, but the SDN orchestrator is notified about the new connection. In scenario C, the QPSK/16QAM Tx is instructed by the SDN controller to reconfigure its modulation format from 16QAM to QPSK, to provide a bitrate of only half of the initial value. Recovered constellation diagrams are shown as insets in Fig.9a for both 16QAM and QPSK. The SDN orchestrator is notified of the new bitrate.

C. E2E QoS recovery

When a transport domain is unable to recover itself from a failure, or QoS cannot be ensured, the SDN orchestrator must

Fig. 9: QoS recovery in an OCS domain: a) OCS domain; b) OSNR vs. BER curve for PM-QPSK and PM-16QAM Tx; c) EVM-based in-band OSNR monitored for PM-QPSK and PM-16QAM.

be notified to perform an end-to-end QoS recovery (Fig.6, scenario D). This end-to-end recovery has a multi-domain scope, it involves the calculation of a new alternative route that satisfies the required QoS constrains, which avoids the network domain on which the failure has occurred.

Fig.10 shows the control messages between the orchestrator and the controller instances during the conducted experiment. First the IT and Network orchestrator subscribes the *updateCallService* and *removeCallService* services of the MSO through the corresponding URL (*/restconf/streams/serviceName*), the same process is performed by the MSO to each domain controller. When the MSO receives the notification from the *removeCallService* of Domain C, it calculate the new E2E path (excluding Domain C) which satisfies the required QoS (A-D-E) and it requests the necessary call deletions (IDs: 00002, 00003), modifications (IDs: 00001, 00004) and the establishment of a new call in Domain D (ID: 00005) to the corresponding SDN controllers. Once the E2E service is restored, it informs the cloud and network orchestrator about the E2E call modification (ID:1). The whole E2E recovery process in the current scenario took 0.624 seconds since the *removeCallService* notification was received by the orchestrator until the last Service Call response is received by the orchestrator too.

VI. CONCLUSION

This paper presents the Control Orchestration Protocol (COP) as a transport API solution for cloud and heterogeneous network orchestration. An integrated IT and network orchestration which implements COP is presented for the dynamic provisioning and restoration of QoS-aware E2E services between IT resources deployed in a distributed cloud computing infrastructure.

The proposed architecture is demonstrated in a multi-partner international control and data plane testbed, by two proof of concepts: the dynamic provisioning of E2E IT and network resources across the aforementioned network; and E2E QoS

Fig. 10: E2E recovery traffic capture.

assurance, including per-domain and end-to-end QoS recovery based on data-plane QoT monitoring.

ACKNOWLEDGMENT

This work was partially supported by EU FP7 projects STRAUSS (608528) and LIGHTNESS (318606) and Spanish MINECO project DESTELLO (TEC2015-69256-R).

REFERENCES

- [1] "White paper: Global transport sdn prototype demonstration," tech. rep., OIF/ONF, October 2014.
- [2] R. Vilalta, V. López, A. Mayoral, N. Yoshikane, M. Ruffini, D. Siracusa, R. Martínez, T. Szyrkowicz, A. Autenrieth, S. Peng, R. Casellas, R. Nejabati, D. Simeonidou, X. Cao, T. Tsuritani, I. Morita, J. P. Fernández-Palacios, and R. Muñoz, "The need for a control orchestration protocol in research projects on optical networking," in *Networks and Communications (EuCNC), 2015 European Conference on*, pp. 340–344, June 2015.
- [3] R. Muñoz, A. Mayoral, R. Vilalta, R. Casellas, R. Martínez, and V. Lopez, "The need for a transport api in 5g networks: the control orchestration protocol," in *Optical Fiber Communication Conference*, p. Th3K.4, Optical Society of America, 2016.

- [4] A. Mayoral, R. Vilalta, R. Muñoz, R. Casellas, R. Martínez, M. S. Moreolo, J. M. Fàbrega, S. Y. Yan, A. Aguado, E. H. Salas, S. Peng, G. S. Zervas, R. Nejabati, D. Simeonidou, J. M. G. Josa, V. Lopez, O. G. de dios, J. P. Fernández-Palacios, P. Kaczmarek, R. Szwedowski, T. Szyrkowiec, A. Autenrieth, N. Yoshikane, X. Cao, T. Tsuritani, I. Morita, M. Shiraiwa, N. Wada, M. Nishihara, T. Tanaka, T. Takahara, J. C. Rasmussen, Y. Yoshida, and K.-I. Kitayama, "First experimental demonstration of distributed cloud and heterogeneous network orchestration with a common transport api for e2e services with qos," in *Optical Fiber Communication Conference*, p. Th1A.2, Optical Society of America, 2016.
- [5] A. Aguado, M. Davis, S. Peng, M. V. Álvarez, V. López, T. Szyrkowiec, A. Autenrieth, R. Vilalta, A. Mayoral, R. Muñoz, R. Casellas, R. Martínez, N. Yoshikane, T. Tsuritani, R. Nejabati, and D. Simeonidou, "Dynamic virtual network reconfiguration over sdn orchestrated multitechnology optical transport domains," *Journal of Lightwave Technology*, vol. 34, pp. 1933–1938, April 2016.
- [6] Y. Yoshida, A. Maruta, K. Kitayama, M. Nishihara, T. Takahara, T. Tanaka, J. Rasmussen, N. Yoshikane, T. Tsuritani, I. Morita, S. Yan, Y. Shu, Y. Yan, R. Nejabati, G. Zervas, D. Simeonidou, R. Vilalta, R. Muñoz, R. Casellas, R. Martínez, V. López, A. Aguado, and J. Beltran, "SDN-based Network Orchestration of Variable-capacity Optical Packet Switching Network over Programmable Flexi-grid Elastic Optical Path Network," *Journal of Lightwave Technology*, vol. 33, pp. 609–617, Feb. 2015.
- [7] D. King and A. Farrel, "A PCE-Based Architecture for Application-Based Network Operations." RFC 7491, Mar. 2015.
- [8] A. Mayoral, R. Vilalta, R. Muñoz, R. Casellas, R. Martínez, and F. J. Vilchez, "Integrated IT and Network Orchestration Using OpenStack, OpenDaylight and Active Stateful PCE for Intra and Inter Data Center Connectivity," in *The European Conference on Optical Communication (ECOC)*, Sept. 2014.
- [9] H. Harai, H. Furukawa, K. Fujikawa, T. Miyazawa, and N. Wada, "Optical packet and circuit integrated networks and software defined networking extension," *J. Lightwave Technol.*, vol. 32, pp. 2751–2759, Aug 2014.
- [10] G. Bosco, P. Poggiolini, A. Carena, V. Curri, and F. Forghieri, "Analytical results on channel capacity in uncompensated optical links with coherent detection," *Opt. Express*, vol. 19, pp. B440–B451, Dec 2011.
- [11] R. Schmogrow, B. Nebendahl, M. Winter, A. Josten, D. Hillerkuss, S. Koenig, J. Meyer, M. Dreschmann, M. Huebner, C. Koos, J. Becker, W. Freude, and J. Leuthold, "Error Vector Magnitude as a Performance Measure for Advanced Modulation Formats," *IEEE Photonics Technology Letters*, vol. 24, pp. 61–63, 2012.