

Control and Orchestration of Multi-domain Optical Networks with GMPLS as Inter-SDN Controller Communication [Invited]

Ramon Casellas, Ricardo Martínez, Raúl Muñoz, Ricard Vilalta, Lei Liu

Abstract—Emerging use cases, such as the orchestration of heterogeneous domains or the interconnection of remote data centers, require the design of new network control, management and orchestration architectures that are adapted to the requirements associated to end-to-end services. This heterogeneity and diversity, stemming from actual deployments, is related to the involved data and control technologies or due to network segmentation. It is acknowledged that, in these scenarios, the deployment of a single Software Defined Networking (SDN) controller may not be practical or may have fundamental limitations, either in terms of scalability, complexity or inter-operability. To address this trend, and as a means to cope with the complexity growth, research in this area is considering the concept of SDN orchestration (the coordinated control of heterogeneous systems), abstraction (the selection of an entity relevant attributes) and the use of multiple controllers, commonly arranged in hierarchical or flat configurations.

In this paper, we consider an hybrid approach, combining elements from distributed control with elements from centralized control, relying on the concepts of abstraction and aggregation, and hierarchical Traffic Engineering (TE). We propose a multi-domain and multi-vendor network control architecture organized in layers, the abstract network layer and the control specific layer, resulting in a mesh of generic SDN controllers that use GMPLS protocols as their East/West interfaces (OSPF-TE and RSVP-TE as Inter-SDN controller communications protocols), forming an abstracted topology. To validate the approach, and to evaluate factors such as end-to-end network service deployment delay or control overhead, the proposed architecture is implemented and validated in a control plane testbed. While emphasis is given to the control of optical (flexi-grid) networks, the proposed approach can be generalized to other transport technologies.

Index Terms—Optical Network Control and Management, Network and Topology Abstraction, Orchestration, Path Computation Element (PCE), Software Defined Networking (SDN), Orchestration, Stateful, Optical OpenFlow.

I. INTRODUCTION

SIMPLE Software Defined Networking (SDN) control deployments are scoped to a single domain, where a (logically) centralized controller has full topology visibility of the network. Even if the network is segmented, e.g., for scalability reasons, a common deployment still considers a single controller. However, new use cases and scenarios call for more robust and scalable solutions, considering the concept of orchestration, abstraction and the use of multiple controllers, commonly arranged in hierarchical or flat configurations. An overview of selected use cases and referred configurations is provided in [1]. Orchestration, roughly defined




Fig. 1. Overview of the hybrid architecture combining elements from distributed and centralized control.

as the coordinated control of heterogeneous systems usually involving multiple interfaces and its related workflow, is becoming a key component of control and management solutions due to common use cases such as: i) multi-domain scenarios, networks segmented in vendor islands or heterogeneous control plane interworking (with SDN/OpenFlow [2] and Generalized Multi-Protocol Label Switching, GMPLS [3] potentially involving multiple controllers), ii) integration of cloud computing IT and networking resources with joint and dynamic provisioning of resources [4] or iii) coordination of network segments such as access, aggregation and core transport. Abstraction, defined as the selection of an entity relevant attributes, is considered since although Traffic Engineering (TE) information is only contained within each domain for scaling and confidentiality, end-to-end services may require a constrained TE path, so it is efficient to expose a limited amount of TE information.

This paper extends our previous work in [5]. In summary, we consider a hybrid solution, sketched in Figure 1 based on a combination of centralized and distributed elements, where a mesh of controllers collaborate running a dedicated control plane instance [6] using GMPLS protocols as their East/West (E/W) interfaces. Within each domain, different control approaches are possible. This approach complements [7], where we proposed a multiple controller model for overarching control and orchestration based on a hierarchy. We adopt the concept of TE reachability as defined in [8], so a subset of TE reachability information should be provided from each domain so that a client can determine whether they can

Manuscript received June 1, 2015.

R. Casellas, R. Muñoz, R. Martínez and R. Vilalta are with CTTC, Spain (e-mail: ramon.casellas@cttc.es).

L. Liu was with University of California, Davis. One Shields Ave, Davis, CA 95616, USA. He is now with Fujitsu Laboratories of America, Inc., Sunnyvale, CA 94085, USA.

establish a multi-domain TE path. The proposed architecture addresses several use cases such as multi-domain overarching control, migration scenarios, or support for GMPLS User Network Interface (UNI) for a SDN controller. It combines the benefits of the GMPLS and Automatically Switched Optical Networks (ASON) routing (robustness, scalability, stability) while locally enabling network operators to deploy SDN features in their scoped and controlled domains (control and data plane separation, logically centralized controllers, open interfaces and modular software designs, etc.) in view, for example, of progressive migrations.

The rest of this paper is organized as follows: in Section II we detail the proposed control plane layered architecture. In Section III we elaborate on the actual procedures and message flows. Section IV includes the experimental performance evaluation, focusing on parameters such as setup delay and control plane overhead and finally, in Section V we conclude the paper.

II. CONTROL PLANE ARCHITECTURE

We consider a *domain* as a set of data plane elements (such as Reconfigurable Optical Add Drop Multiplexers or ROADMs) under the control of a technology specific control plane, assuming a deployment model of a single SDN controller per domain, with full topology visibility. Even if the domain deploys a distributed control plane such as GMPLS, we assume that an Active Stateful Path Computation Element (AS-PCE) [9] acts as a proxy single point of entry. If we consider basic set of functions associated to an SDN controller, related to the separation of control and data planes, its centralized nature, the call/connection provisioning capabilities and aspects such as topology management, let us note that an AS-PCE is an SDN controller. It is able to provision Label Switched Paths (LSPs), delegating to an underlying GMPLS control plane, provided, for completeness, that it exports a certain North Bound Interface (NBI) Application Programming Interface (API) to applications, such as a Network Management System (NMS) or orchestrator that requests the provisioning from the PCE.

The design of the overarching control and orchestration architecture is based on the following underlying principles and initial requirements, involving a trade-off between flexibility, scalability and availability.

First, the main target for deployment of the architecture is a reduced number of domains (the applicability to a large number of domains is not considered), in the scope of one or few number of operators with peering agreements, the most common being a single operator that spans multiple domains.

Second, each domain is constrained to its own control technology, due to several factors such as return on investment of existing deployments or vendor lock-in in which the SDN controller is tightly coupled to the domain elements. Within each domain, intra-domain control is up to the domain, abstracted to the higher layers: either OpenFlow, GMPLS or even based on a traditional NMS.

Third, is the use of domain abstraction to enable scalability and hierarchical TE. Key to this is TE network abstraction, which is the synthesizing of reported TE attribute information for each domain and inter-domain links (abstract topology). This representation does not include all possible connectivity options, but instead provides a connectivity

matrix based on attributes as reported from within each domain. The abstraction makes use of available TE information, but is subject to network policy and management. By selecting relevant TE attributes, a hierarchy of topologies can be constructed. Note the recursive property of the architecture, where the abstract network can be in turn placed behind a SDN controller or AS-PCE and be part of a wider orchestration and overarching control layer (i.e., to jointly provision IT and network resources).

Next, in view of the number, disposition and resulting mesh domains, the distributed GMPLS architecture and framework is selected for the abstract network layer. The choice is motivated by the maturity, robustness, and available feature set, including the ability of resource and neighbor discovery, path provisioning and recovery. A potential drawback of the approach is related to inter-operability. It has been argued that the interworking of GMPLS implementations of different vendors is problematic, due in part to proprietary vendor extensions; the complexity of the protocols – notably when dealing with transport data planes such as Optical Transport Networks (OTN) – and the different interpretation of standards. We believe that these issues can be mitigated to a large extent when considering network topology abstraction, in which the selected relevant TE attributes are mature and well understood. That said, inter-operability between vendors is not likely to disappear, in spite of efforts related to the definition on common hardware models and open and standard protocols.

In view of the aforementioned considerations, the architecture (cfr. Figure 2) is thus based on a mesh of SDN controllers that implement GMPLS routing and signaling as their East/West interfaces. In other words, the GMPLS Open Shortest Path First with Traffic Engineering extensions (OSPF-TE) routing protocol and the Resource Reservation protocol with Traffic Engineering extensions (RSVP-TE) signaling protocol are selected as Inter-SDN controller communication. In the proposed architecture, cooperating SDN controllers are responsible for the abstraction of their domains (in terms of announcing an underlying abstract topology) and for adaptation of the control procedures, which remains technology specific. Such GMPLS enabled SDN controllers (GSC) maintain a domain abstract topology map and execute a separate GMPLS instance in which routing and signaling adjacencies are congruent to the domain connectivity. Each GSC disseminates the abstraction of its domain by means of dedicated OSPF-TE Link State Advertisements (LSAs). An AS-PCE placed on top of the abstracted network provides the functions of multi-domain path computation and multi-domain network service provisioning, operating on the abstracted Traffic Engineering Database (TED). Let us note that there may thus be several independent AS-PCE instances: one for the abstract network layer, and one for each domain that is controlled by a GMPLS control plane.

In view of the choice of this E/W interfaces, each GSC/SDN controller has thus a Path Computation Client (PCC) module that interacts with the AS-PCE (for path computation, LSP establishment and LSP control delegation), a Routing Controller (RC) that uses OSPF-TE and a Connection Controller (CC) that participates in signaling sessions. The part of the GSC corresponding to the specific control layer is responsible for network abstraction and control adaptation, address and resource identifier mappings and process adaptation.

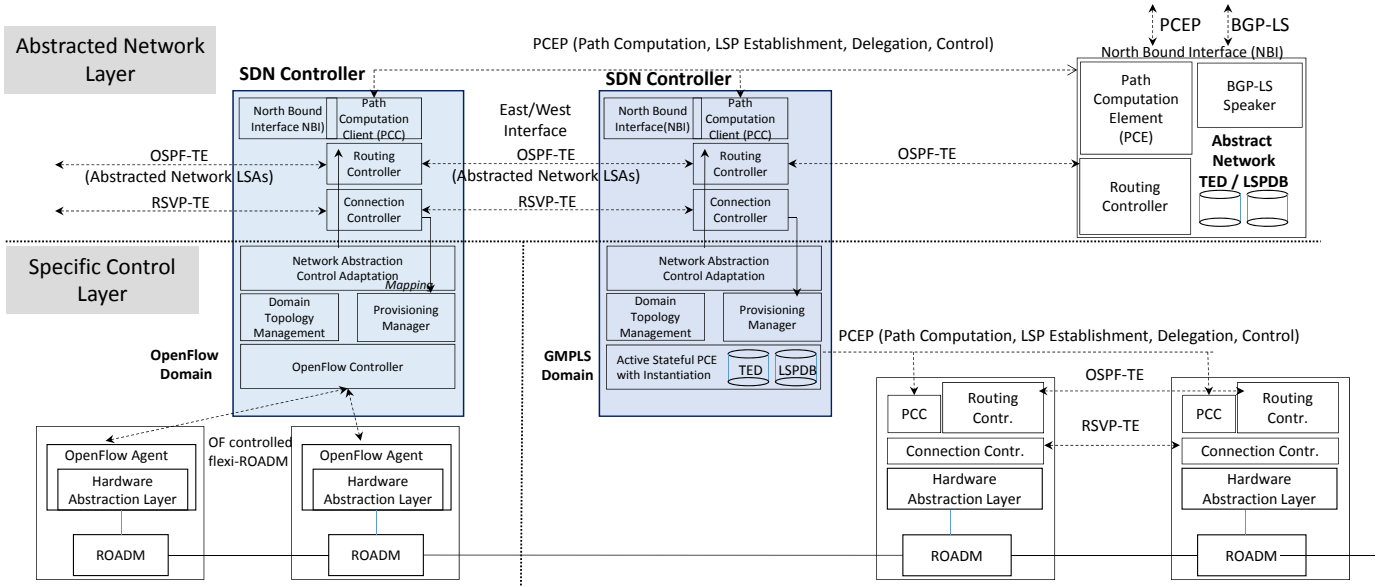


Fig. 2. Hybrid control and orchestration architecture based on a mesh of SDN controllers forming an abstracted network layer

III. CONTROL PLANE PROCEDURES AND MESSAGE FLOWS

The involved GSCs disseminate OSPF-TE LSAs describing the abstract topology, as defined by their domain internal policy. This allows the AS-PCE in the abstract network to obtain a TED of the abstract network layer. The workflow for the provisioning is described within Figure 3. An end-to-end provisioning is triggered by the AS-PCE in the abstract network, e.g., by sending a PCE communications Protocol (PCEP) [10] PCInitiate message to the ingress abstract node, and relying on distributed signaling over the Data Communications Network (DCN) formed by the controller mesh. The ingress GSC uses an end-to-end RSVP-TE [11] signaling session, indicating the Explicit Route Object (ERO) as provided by the AS-PCE, and the SESSION address being set to the destination abstract node. Each GSC processes RSVP-TE signaling messages, mapping resource reservations to actual segment connectivity requests, which are, in turn, delegated to the underlying, specific technology (e.g., a cross-connect at an abstract nodes between input and output ports induces a segment connection in the corresponding domain). Specific examples of mappings are provided in Section IV.

In particular, a receiving GSC processes the RSVP-TE PATH and RESV messages and identifies the involved abstract TE links, from the PATH RSVP_HOP and ERO objects, mapping such abstract link identifiers to its domain client port addresses, as defined in its topology aggregation policy. The sender descriptor traffic parameters (SENDER_TSPEC) conveys the requested frequency slot width and, if available, SUGGESTED_LABEL objects may specify the upstream domain preference, whereas the RESV GENERALIZED_LABEL contains the downstream allocated frequency slot. The underlying segment is established either using GMPLS/PCE control plane or OpenFlow, since the GSC South Bound Interface (SBI) is technology specific.

For an OpenFlow domain, the GSC configures the switching behavior of the network elements using a modified OpenFlow for flexi-grid networks, directly programming cross-connections in the domain network elements (by means of modified OpenFlow messages CFLOW_MOD

and CPORT_STATUS to support flexi-grid extensions [12] shown in Figure 3). For a GMPLS/AS-PCE domain, the approach is applied recursively, thus having an additional RSVP-TE session that is scoped to the domain. To follow standard GMPLS procedures, the RESV message at the abstract layer is sent upstream upon successful completion of the establishment of the segment. The architecture does not preclude additional methods within a domain, as long as the sequential ordering is respected. The end-to-end service is thus the result of *stitching* the segments established within each domain. In case of errors, standard methods apply: a GSC may send a PATH_ERROR message upstream and a PATH_TEAR message downstream to release the resources. GMPLS Recovery applies without change at the abstract layer (e.g. supporting break-before-make rerouting). Finally, the successful provisioning of the end-to-end service is reported via a PCEP Report (PCRpt) message back to the AS-PCE.

IV. EXPERIMENTAL PERFORMANCE EVALUATION

In order to validate the proposed architecture and to obtain performance indicators, we have deployed a control plane testbed with emulated optical hardware. This implies that each GMPLS controller or OpenFlow agent located in a data plane node accepts configuration requests via its Hardware Abstraction Layer (HAL) and assumes the underlying hardware is properly configured without error. Each of the involved AS-PCEs, GMPLS controllers as well as the OpenFlow controllers and OpenFlow agents is running in dedicated (physical) GNU/Linux servers with Core2 Duo processors and 1-4 Gb RAM, a hardware that suffices for the purpose. The servers are running in a LAN, with dedicated VLANs and GRE tunnels emulating the DCN.

A. Testbed description

The testbed has been configured as a multi-domain network with 5 domains. Each domain is controlled by a GSC that performs single TE node network abstraction, using OSPF-TE LSAs to disseminate node identifiers (domain X

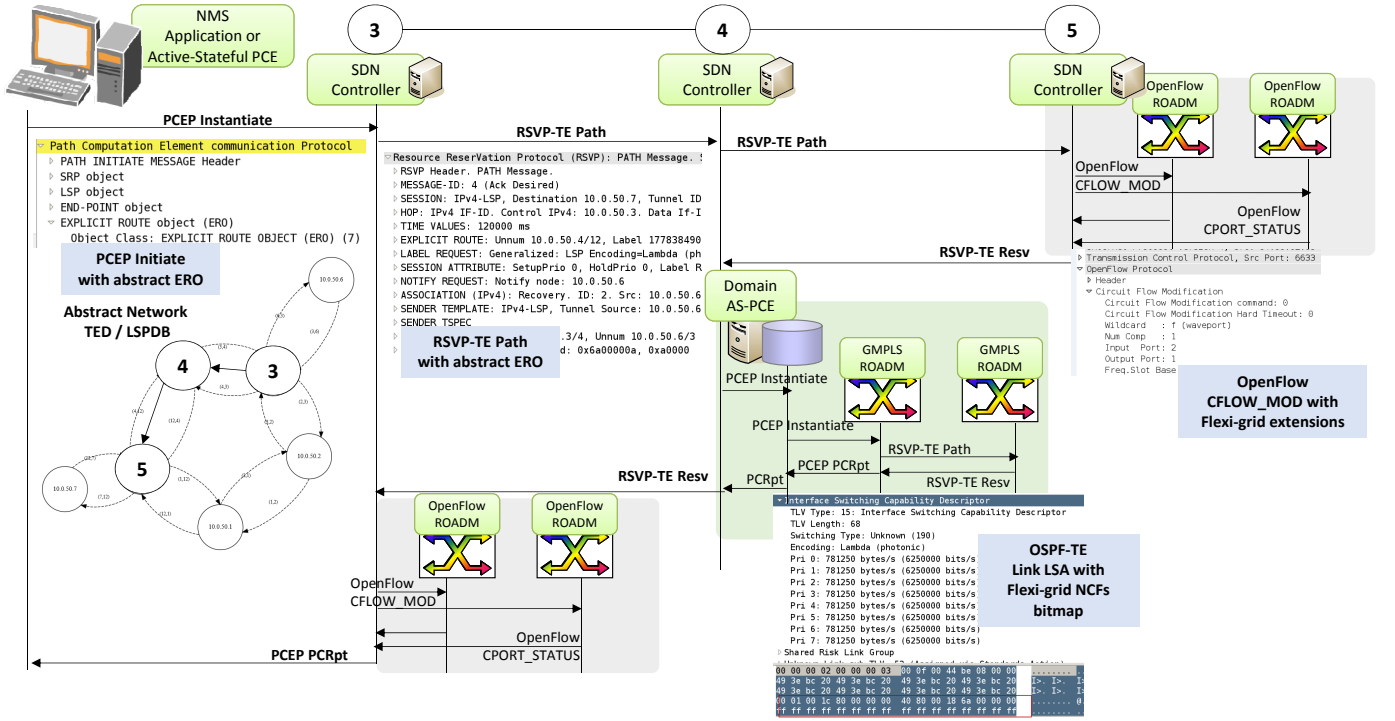


Fig. 3. Example of message flow for the provisioning of end-to-end connections, showing the end-to-end RSVP-TE session and the intra-domain segment establishment (just 3 domains shown for simplicity).

has 32 bit address 10.0.50.X) and link TE attributes. Additional nodes representing GMPLS clients/endpoints - 6 and 7 - have been added, resulting in the abstract topology shown in Figure 4.

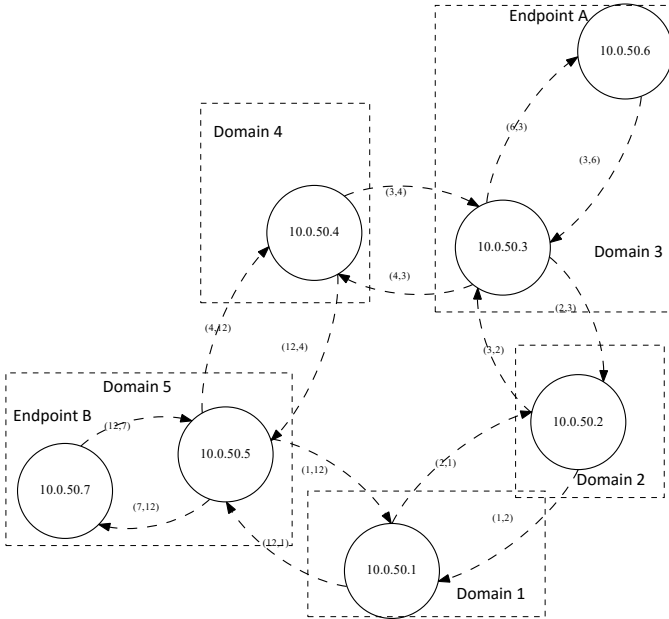


Fig. 4. Considered multi-domain mesh topology seen by the AS-PCE in the abstract network layer

Domains 3 and 5 are optical OpenFlow islands with 4 core nodes each and three additional OpenFlow datapath ids that represent endpoints of the OpenFlow network (Figure 5).

On the other hand, domains 1, 2 and 4 are GMPLS domains with an AS-PCE to instantiate LSPs. The biggest

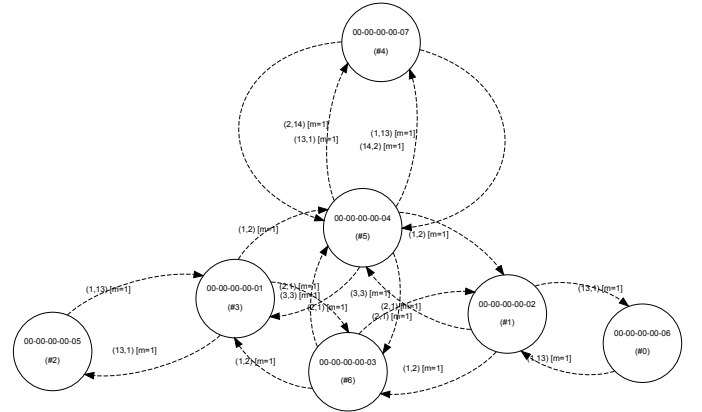


Fig. 5. Deployed OpenFlow flexi-grid network topology for domains 3 (datapath ids 00-0x) and 5 (datapath ids 01-0x, symmetrical).

domain, domain 4, has 14 nodes representing a Spanish-wide optical mesh network, as shown in Figure 6. Let us note that the arrangement of OpenFlow - GMPLS - OpenFlow domains is motivated by a particular use case in which e.g., remote data centers controlled by OpenFlow are interconnected by a GMPLS controlled core transport network. Domains 1 and 2 in this context provide path diversity.

All links are considered basic flexi-grid optical links with finest granularity in the selection of nominal central frequencies and frequency slot widths. All links are homogeneous having 128 usable nominal central frequencies, ranging from with n parameter [13] $n = 0$ to $n = 127$. The BV-ROADMs do not present any asymmetry and are assumed to be able to switch from any incoming port to any outgoing port.

Each GSC is responsible for mapping addresses and identifiers (such as endpoints or unnumbered interfaces) from

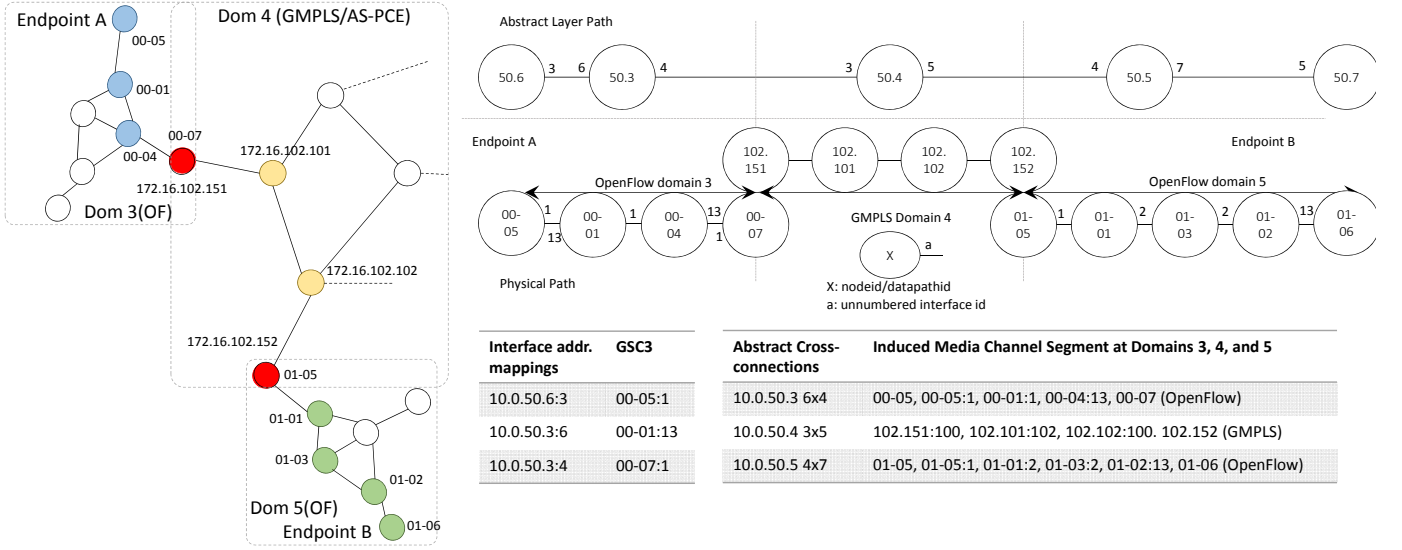


Fig. 7. Example of multi-domain path in the abstract layer, its corresponding path in the physical multi-domain topology and mapping identifiers and cross-connections to underlying segment media channels. Connections from Endpoint A (represented in the abstract network as node 10.0.50.6 and corresponding to OF datapathid 00-05 in domain 3) to Endpoint B (represented in the abstract network as node 10.0.50.7 and corresponding to OF datapathid 01-06 in domain 5).

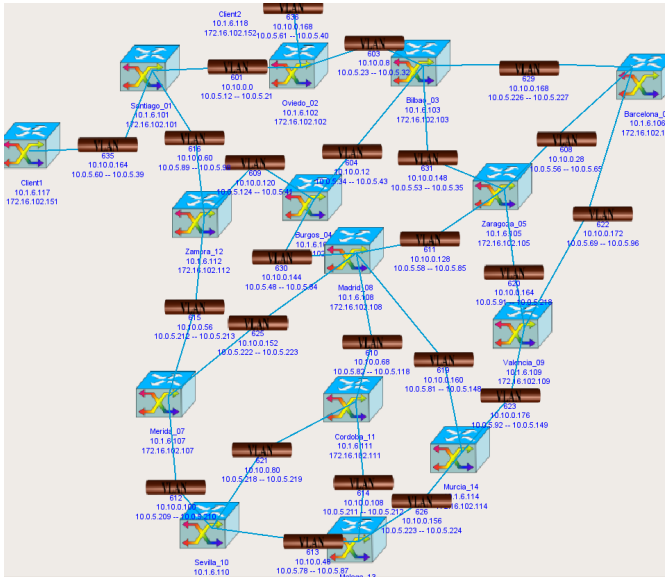


Fig. 6. GMPLS controlled network emulating a Spanish 14-node topology.

the address space of the abstract network to its underlying domain-local address space. To illustrate this, Figure 7 shows a path in the abstract network layer, the corresponding actual physical path with the considered topologies and (selected) mappings configured in domain 3 (table, left) and how the cross-connections at the abstract nodes induce a (segment) media channel at each domain.

B. Experimental Results and analysis

To test the system, connections are set up from Endpoint A (represented in the abstract network as node 10.0.50.6 and corresponding to OF datapathid 00-05 in domain 3) to Endpoint B (represented in the abstract network as node 10.0.50.7 and corresponding to OF datapathid 01-06 in domain 5), for a frequency slot of width $m=1$. In order to

0.000000	Abstract_AS_PCE	Endpoint_06	PCEP	Initiate
0.006812	Endpoint_06	Domain_03	RSVP	PATH Message
0.007039	Domain_03	Endpoint_06	RSVP	ACK Message
0.150216	Domain_03	Endpoint_06	RSVP	RESV Message
0.150388	Endpoint_06	Domain_03	RSVP	ACK Message
0.154733	Endpoint_06	Abstract_AS_PCE	PCEP	Path Comput

Transmission Control Protocol, Src Port: 47759 (47759), Dst Port	
Path Computation Element communication Protocol	
Path Computation LSP State Report (PCRpt) Header	
SRP object	
LSP object	
EXPLICIT ROUTE object (ERO)	
BANDWIDTH object	
LABEL object	

0000	00 04 03 0a 00 00 00 00	00 00 00 00 00 00 08 00
0010	45 00 00 ec be 3f 40 00	40 06 2f e9 0a 00 05 e3	E...?@.
0020	0a 00 32 01 ba 8f 10 5d	37 ea 25 2b ef 0b 9b 23	..2....]

Fig. 8. Wireshark captured at Endpoint_06 with the reception of PCInitiate that triggers the signalling session in the abstract network. Endpoint_06 refers to node 10.0.50.6 in the abstract layer (without prefix 10.0.50.)

decompose the process, let us start with Figure 8. The AS-PCE in the abstract network sends a PCInitiate with the ENDPOINTS object and a ERO that contains the domain sequence to cross, encoded as an ordered list of abstract node identifiers. The endpoint then proceeds with the signaling by sending a RSVP-TE PATH message to the abstract node GSC (Domain 3). After the reception of the RSVP-TE RESV message implying a successful establishment, the endpoint sends a PCEP Report message (PCRpt) to the AS-PCE. From the capture we can see that this provisioning took 154 ms.

At each domain, upon reception of the RSVP-TE RESV message, the GSC needs to actually expand the segment and establish the underlying connection, mapping the procedures to the underlying control plane. In particular, when the GSCs located in Domains 3 and 5 receive the RESV message, it sends OpenFlow CFLOW_MOD messages to the OF agents located in the nodes and gets OpenFlow CPORT_STATUS messages acknowledging the configuration. When the GSC

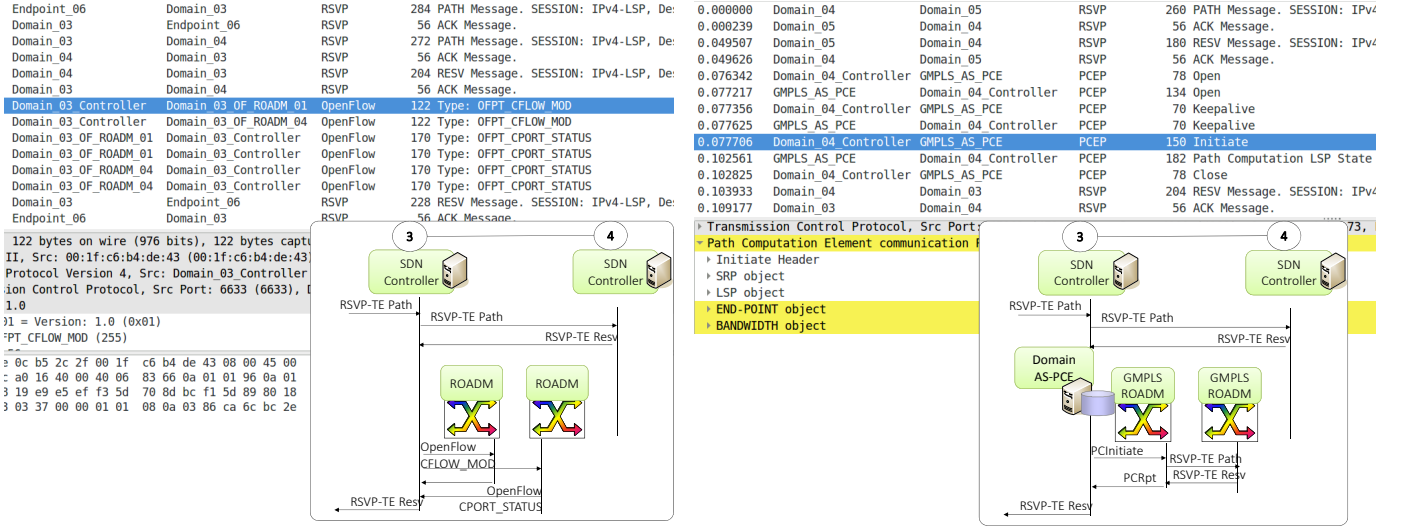


Fig. 9. Wireshark captured Domain 3 (left) and Domain 4 (right) upon reception of a RESV message from a downstream abstract node.

located in Domain 4 receives the RESV message, it applies the initial process recursively: the GSC requests the AS-PCE in that domain to setup the segment, which proceeds by sending a PCInitiate to the head-end node of that domain, and sending a PCRpt upon completion. The processes can be seen in Figure 9.

It is worth noting that the configuration of ROADMs in OpenFlow domains can be parallelized, as opposed to GMPLS in where it is serialized by the order of processing of RESV messages.

C. Network Service Setup Delay

There are several significant values as reported in Table I. Starting from the service delay seen by an application or NMS on top of the AS-PCE of the abstract network layer, averaging 180ms, which includes two main components: the PCEP interface to the ingress node, with the PCInitiate and PCRpt exchange, and the actual signaling delay seen by the ingress node, measured from the sending of the PATH message to the reception of the RESV message, which is, on average 155 ms. This value is, in turn, the result of the sequential provisioning that happens at each domain. In particular, the highest value corresponds to Domain 4 (109 ms) seen by GSC 4, that covers the whole PCEP handshake with the AS-PCE in the domain and the reception of the PCRpt (note that this value could be further reduced if a persistent PCEP connection was used). The setup delay purely from the point of view of the AS PCE in the Domain 4 (from the PCInitiate to the PCRpt from the head end node in the segment) is 30 ms in average, with the actual segment provisioning involving 4 GMPLS nodes (16 ms of PATH/RESV signaling exchange). Segment provisioning in OpenFlow domains is the lowest value, averaging 12 ms and a result of the inherent parallelization in the domain.

The histogram of the path setup delay and the Cumulative Distribution Function (CDF) can be seen in Figure 10; the average setup value seen by an application using the NBI of the PCE in the abstract layer is around 180 ms (of which 155 ms correspond to the actual signaling delay seen by the ingress/head end abstract node). Let us note that flexi-grid optical hardware may require a significant amount of time

Decomposition of Control Plane latencies		(Avg. ms)
Setup delay seen by NMS/App to AS-PCE in Abstract Network		180
Setup delay seen by ingress abstract node		155
Segment delay Domain 4 SDN Controller		109
Segment establishment in Domains 3 and 5 (symmetric topology domains) using OpenFlow parallel configuration and acknowledgements		12
Segment setup in Domain 4 GMPLS from AS-PCE		30
Segment setup in Domain 4 GMPLS from ingress node		16

TABLE I
MEASURED LATENCIES FOR THE NETWORK SERVICE SETUP DELAY.

(can reach the order of seconds) for configuration and optical connections may need to be validated before actual traffic is transported on the provisioned optical channel. Consequently, values are provided as a performance indicator, but real deployments would definitely show higher values.




Fig. 10. Setup delay (as seen by the abstract layer GMPLS instance and by the PCE) and CDF.

D. Control Plane overhead and Processing requirements

Finally, another performance indicator is the control plane overhead, including not only the overhead in terms of required bandwidth in the control plane network due to message exchanges, but also in the processing requirements and resources related to session management notably in centralized elements. Generally speaking, the main factors to determine the overhead are the rate of arrivals of requests, the centralized or distributed nature of the involved elements and aspects such as the intra-domain topology aggregation policies.

Starting with the abstract layer, and for service requests arriving every few seconds, there are no significant issues regarding control plane overhead, but it is worth noting the effect of the topology dissemination delay which may cause the AS-PCE to act upon outdated information if the rate of changes is high. In this case, a conservative policy of aggregating domains as nodes significantly reduces the overhead when compared to policies (e.g. virtual links reflecting domain internal connectivity) [14] that require complex or frequent re-computation of exported TE attributes due to internal changes. The overhead is then due exclusively to the periodic refreshes of LSAs, changes in the inter-domain links and the signaling sessions. For targeted scenarios involving a few domains the system scales reasonably. Generally speaking, assuming dedicated control plane links, the required control plane bandwidth is of a few Mb/s, resulting from the having message sizes rarely bigger than a few hundred bytes. To illustrate this, a single LSP provisioning results in the head end abstract node having to process PCEP initiate and report messages (272, 252 bytes) RSVP-TE PATH (284 bytes) and RESV (228 bytes) with ACK (56 bytes each) in 155 ms, with an average of 40 packets per second resulting in 0.06 MBit/s throughput. In the GMPLS domain 4, during the window of a segment provisioning upon a reception of an abstract layer RESV message from domain 5, abstract node 4 sees 9 OSPF-TE LS Updates (248-444 bytes) resulting from the reception of LSAs generated by downstream abstract nodes and the generation of their own interdomain links totaling 2832 bytes with an aggregated throughput of 0.183 MBit/s.

Within each domain, there is a clear trade-off: for GMPLS domains message overhead and control bandwidth is notably higher than the centralized OpenFlow domains due to the additional synchronization required for OSPF-TE LSAs and the signaling exchanges, but OpenFlow domains impose a higher stress in centralized controllers due to the need of processing all the messages and updating the internal databases. The overhead and scalability limitations are increasing with the domain size due, notably, to the number of managed sessions: for the AS-PCE it has to maintain N PCEP sessions with the edge nodes and for the OpenFlow controller it must maintain M OpenFlow connections with all the network elements. In the considered topologies, OpenFlow domains must typically configure 2 flexi-ROADMs involving 6 messages (2 cross-connection requests and 4 link notifications) which, in the worst case, happens in a 2ms window exchanging 924 bytes and resulting in 3.84 MBit/s peak rate.

In summary, the obtained values are within expected ranges for the considered scenarios and topologies, and dedicated control interfaces commonly implemented in terms of

Ethernet interfaces are largely sufficient. The main bottlenecks remain in the processing capabilities of the nodes and the propagation delay of TED updates and synchronization, although mitigated by the expected traffic patterns.

V. CONCLUSIONS

In this paper, we have proposed and implemented an architecture for overarching control and orchestration of multi-domain optical networks with heterogeneous control planes, illustrated with both GMPLS and OpenFlow deployments. The architecture, based on network abstraction and on a mesh of SDN controllers using GMPLS as East/West interfaces, has been validated on an experimental testbed focusing on aspects such as service provisioning latency and control plane overhead. The main drivers to consider such an approach, based on the hybrid combination of both centralized and distributed control, are scalability issues and robustness, isolating intra-domain control for inter-domain. The use of an abstracted domain mesh scales to a reasonably large number of domains, although we expect that such deployments are confined and scoped to a single or reduced number of operators with peering agreements. Within each domain, the scalability is defined by its own intra-domain control.

ACKNOWLEDGMENTS

The research leading to these results has received funding from the European Community's Seventh Framework Programme under grant agreement n. 317999 IDEALIST, the EU-Japan project STRAUSS with grant agreement n. FP7-ICT-2013-EU-Japan 608528) and from the Spanish MINECO project FARO (TEC2012-38119).

REFERENCES

- [1] R. Casellas, R. Muñoz, R. Martínez, R. Vilalta, A. Mayoral, L. Liu, T. Tsuritani, and I. Morita, "Overarching Control of Flexi Grid Optical Networks: Interworking of GMPLS and OpenFlow Domains," *IEEE/OSA J. Lightw. Technol.*, vol. 33, no. 5, pp. 1054 – 1062, Mar. 2015.
- [2] ONF, "OpenFlow Switch Specification, version 1.4 (Wire protocol 0x5)," Open Networking Foundation, Oct 2013. [Online]. Available: <https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow/openflow-spec-v1.4.0.pdf>
- [3] E. Mannie, "Generalized Multi-Protocol Label Switching (GMPLS) Architecture," RFC 3945 (Proposed Standard), IETF, Oct. 2004. [Online]. Available: <http://www.ietf.org/rfc/rfc3945.txt>
- [4] A. Mayoral López de Lerma, R. Vilalta, R. Muñoz, R. Casellas, and R. Martínez, "Experimental seamless virtual machine migration using an integrated sdn it and network orchestrator," in *Optical Fiber Communication Conference*. Optical Society of America, 2015, pp. Th2A–40.
- [5] R. Casellas, R. Martínez, R. Muñoz, R. Vilalta, and L. Liu, "Multi-domain overarching control of flexi-grid networks with GMPLS as Inter-SDN controller communication," in *Optical Fiber Communication Conference (OFC 2015), paper Tu2B.5.*, Mar. 2015.
- [6] R. Muñoz, R. Casellas, R. Vilalta, and R. Martínez, "Dynamic and Adaptive Control Plane Solutions for Flexi-grid Optical Networks based on Stateful PCE," *IEEE/OSA J. of Lightw. Technol.*, vol. 32, no. 16, pp. 2703–2715, Jun. 2014.
- [7] R. Casellas, R. Muñoz, R. Martínez, R. Vilalta, L. Liu, T. Tsuritani, I. Morita, V. López, O. G. de Dios, and J. P. Fernández-Palacios, "SDN orchestration of OpenFlow and GMPLS flexi-grid networks with a stateful hierarchical PCE [Invited]," *IEEE/OSA Journal of Optical Communications and Networking, Special Issue on OFC 2014*, vol. 7, no. 1, pp. A106–A117, Jan. 2015.

- [8] A. Farrel, N. Bitar, G. Swallow, D. Ceccarelli, and X. Zhang, "Problem Statement and Architecture for Information Exchange Between Interconnected Traffic Engineered Networks," *IETF, internet draft, work in progress*, Mar 2015.
- [9] E. Crabbe, I. Minei, and S. S. R. Varga, "Pcep extensions for pce-initiated lsp setup in a stateful pce model," Internet Engineering Task Force, April 2013.
- [10] J. Vasseur and J. L. Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)," RFC 5440 (Proposed Standard), Internet Engineering Task Force, Mar. 2009. [Online]. Available: <http://www.ietf.org/rfc/rfc5440.txt>
- [11] L. Berger, "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions," RFC 3473 (Proposed Standard), IETF, Jan. 2003, updated by RFCs 4003, 4201, 4420, 4783, 4874, 4873, 4974, 5063, 5151, 5420, 6002, 6003, 6780. [Online]. Available: <http://www.ietf.org/rfc/rfc3473.txt>
- [12] R. Casellas, R. Martínez, R. Muñoz, L. Liu, T. Tsuritani, and I. Morita, "An Integrated Stateful PCE/OpenFlow controller for the Control and Management of Flexi-Grid Optical Networks," in *(OFC 2013), paper OW4G.2, Anaheim, CA, USA, Mar. 2013.*, 2013.
- [13] ITU-T Recommendation G.694.1, "Spectral grids for wdm applications: Dwdm frequency grid," Feb. 2012.
- [14] R. Casellas, R. Martinez, R. Muñoz, L. Liu, T. Tsuritani, I. Morita, and M. Tsurusawa, "Dynamic virtual link mesh topology aggregation in multi-domain translucent wson with hierarchical-pce," *OSA Optics Express Journal*, vol. 19, no. 26, December 2011.

Ramon Casellas (SM'12) graduated in telecommunications engineering in 1999 by both the UPC-BarcelonaTech university and ENST Telecom Paristech, within an Erasmus/Socrates double degree program. After working as an undergraduate researcher at both France Telecom R&D and British Telecom Labs, he completed a Ph.D. degree in 2002. He worked as an associate professor at the networks and computer science department of the ENST (Paris) and joined the CTTC Optical Networking Area in 2006, with a Torres Quevedo research grant.

He currently is a senior research associate and the coordinator of the ADRENALINE testbed. He has been involved in several international R&D and technology transfer projects. His research interests include network control and management, the GMPLS/PCE architecture and protocols, software defined networking and traffic engineering mechanisms. He contributes to IETF standardization within the CCAMP and PCE working groups. He is a member of the IEEE Communications Society and a member of the Internet Society.

Ricardo Martínez graduated in telecommunications engineering 2002 and received a Ph.D. degree in telecommunications in 2007, both from the Technical University of Catalonia, Spain. Since June 2007 he has been a research associate in the CTTC Optical Networking Area. He has participated in several Spanish, EU-funded (EU FP6, FP7 and CELTIC) and Industrial projects in the context of optical networking. His research interests include GMPLS architecture, distributed control schemes, packet and optical transport networks. He has published over 80 scientific conferences and peer reviewed journal papers.

Raül Muñoz (SM'12) graduated in telecommunications engineering in 2001 and received a Ph.D. degree in telecommunications in 2005, both from the Universitat Politècnica de Catalunya (UPC), Spain. After working as undergraduate researcher at Telecom Italia Lab (Turin, Italy) in 2000, and as assistant professor at the UPC in 2001, he joined the Centre Tecnològic de Telecomunicacions de Catalunya (CTTC) in 2002.

Currently, he is senior research associate. Since 2000, he has participated in several R&D projects (FP7 IP STRONGEST and NoE BONE, FP6 IP NOBEL and NOBEL2, FP5 LION, etc.) and technology transfer projects. He leads the Spanish project DORADO, and previously has led the projects RESPLANDOR and GAUDI. He has published over 50 journal and international conference papers in this field. His research interests include GMPLS architecture, protocols and traffic engineering algorithms (provisioning, protection and restoration strategies) for highly dynamic optical transport networks.

Ricard Vilalta graduated in telecommunications engineering in 2007 and received a Ph.D. degree in telecommunications in 2013, both from the Universitat Politècnica de Catalunya (UPC), Spain. He also has studied Audiovisual Communication at UOC (Open University of Catalonia) and is a master degree on Technology-based business innovation and administration at Barcelona University (UB). Since 2010, Ricard Vilalta is a researcher at CTTC, in the Optical Networks and Systems Department. His research is focussed on Optical Network Virtualization and Optical Openflow. He is currently a Research Associate at Open Networking Foundation.

Lei Liu (M'11) received his B.E. and Ph.D. degrees from Beijing University of Posts and Telecommunications (BUPT), Beijing, China, in 2004 and 2009 respectively. From 2009 to 2012, he was a research engineer at KDDI R&D Laboratories Inc., Japan, where he was engaged in research & develop on intelligent optical networks and their control and management technologies. From 2013 to early 2015, he was with University of California at Davis, USA, working on many US national (including NSF, GENI, ARL, DOE, etc.) and international projects regarding optical networking and communication technologies. He is now with Fujitsu Laboratories of America Inc. in Silicon Valley, as a Member of Research Staff. His main research interests include the Future Internet architecture (such as Information-Centric Networking), optical networking and communications, network control and management techniques (such as MPLS/GMPLS, PCE, OpenFlow), software-defined networking (SDN), and network protocols and algorithms. He has coauthored more than 100 peer-reviewed papers related to the previously mentioned research topics in international journals, conferences, and book chapters, and He is the first author of more than 50 of them. He presented several invited papers in international journals and conferences, including many prestigious ones such as IEEE JLT, IEEE JOCN, ECOC'11, ECOC'12 and Globecom'12, etc.