

[Invited] End-to-end Service Orchestration From Access to Backbone

V. Lopez, J. M. Gran Josa, V. Uceda, F. Slyne, M. Ruffini, R. Vilalta, A. Mayoral, R. Muñoz, R. Casellas, R. Martínez

Abstract— The rise of new types of business models with service providers, content providers, and virtual network operators entering in competition with traditional telco operators, is reducing revenue margins, making cost-effective provisioning mechanism a necessity. Costly massive over-provisioning needs thus to be replaced by more intelligent dynamic resources allocation. In addition, it is increasingly recognized that many upcoming 5G applications and services will require assured end-to-end quality of service.

Operators have thus started to look at Network Function Virtualization and Software Defined Networks as means to address the challenges of cost effective and highly dynamic end-to-end provisioning.

In this paper we present a test case of an SDN-driven end-to-end service orchestration using a transport API called Control Orchestration Protocol. Our testbed, interconnecting a core network (within the Telefonica premises in Madrid) and an access network (within the Trinity College of Dublin facilities), demonstrates the possibility to operate sub-second end-to-end capacity reservation, showcasing SDN provisioning across multi-domain networks.

Index Terms— SDN; PCE; ABNO; Network Orchestrator; Network Controller; Transport API; COP; LR-PON.

I. INTRODUCTION

For the past many years, network operators have typically managed their network in a rather static fashion, by analysing network traffic matrices and statically provisioning network resources to meet the traffic growth prediction for the envisaged planning period (typically 3 to 5 years).

Recently, the rapid growth of new applications (e.g., social media with video sharing) and type of services (e.g., mobile fronthaul), have led operators to consider technologies enabling more dynamic service provisioning. This is exemplified by the 5G motto of "reducing service deployment

time from 90 days to 90 minutes"[1].

This quest for fast resource provisioning is being addressed through the concepts of Software Defined Networks, Network Virtualization and Network Function Virtualisation (NFV), which moves services from dedicated hardware into software running in data centres of different scales spread across the metro area.

When considering the exact type of 5G services that we will see by 2020 and beyond, there is indeed still much uncertainty. However whether they might involve autonomous driving, in-car entertainment, augmented and virtual reality with very high resolution video streams, they will require assured end-to-end data delivery by the network in order to fulfil end user expectations.

Due to the increasingly competitive telecommunications market, with many virtual operators and service providers entering in direct competition with the traditional operators, providing quality of service by massively overprovisioning network resources is not a viable option. Therefore, new architectural solutions are needed to deliver the huge expected increase in traffic in a cost-effective way, and ensure low cost broadband Internet access. Long-Reach Passive Optical Networks (LR-PON) [1], for example, addresses the issues by proposing a network design that can substantially reduce the number of network nodes, by transparently interconnecting end users to a small number of Metro-Core (MC) nodes.

In addition to physical network architectures aimed at reducing installation and maintenance costs, it is mandatory to have a control plane that enables the configuration of the resources on demand, in order to direct capacity where and when needed instead of operating ubiquitous overprovisioning. Software Defined Networking (SDN) proposes to decouple the control plane from the data plane to allow this dynamic operation.

This paper defines and demonstrates architecture to support E2E service from access to backbone. The remaining of the article is structured as follows. Section 2 provides literature background on SDN-driven network orchestration. Section 3 defines the SDN architecture to allow end-to-end services. This work uses the Control Orchestration Protocol, which is presented in Section 4. Section 5 introduces the test case that motivate the work. Section 6 presents the

Manuscript received September 30, 2016.

V. Lopez, J. M. Gran Josa and V. Uceda are with Telefonica I+D, Madrid, Spain (e-mail: victor.lopezalvarez@telefonica.com). J. M. Gran Josa is now with Telefonica de España.

F. Slyne and M. Ruffini are with Trinity College Dublin, Ireland.

with Centre Tecnològic de Telecomunicacions de Catalunya (CTTC/CERCA), Castelldefels, Spain.

experimental validation and the results. Finally, Section 7 concludes this article.

II. STATE OF THE ART

The utilization of SDN for access networks is motivated by operators [3], who need highly dynamic control to deploy new services, reduce their provisioning time and ultimately reduce the cost of network ownership.

EU FP7 SPARC (Split Architecture) [4] was one of the first projects to demonstrate the use of SDN in the Access and Aggregation networks as well as prototyping an idea of Network Function Virtualisation, through a Virtual Home Gateway and a Virtual BRAS. SPARC's PCE architecture facilitated the concentration of control capability in a centralised system, as well as the separation of access/aggregation and backbone/core networks. Standard IP/MPLS control protocols such as OSPF, LDP, RSVP-TE and BGP were used to provide the necessary glue between control domains.

Other relevant work was carried out by the EU-FP7 project OFELIA. While not primarily looking at SDN in the access network, it demonstrated the evolving use of SDN and particularly Openflow in the Wide Area Networks and the effect on traditional carrier networks [5]. OFELIA demonstrated Optical Wavelength switching, Optical Flow Switching and Multi-service technology control. Associated projects such as EU-FP7 project ALIEN presented a generic model using abstraction in the data plane to allow a wide range of access devices based on FPGA's and Network processors.

ICT IP STRONGEST [6] targeted instead core network control plane, demonstrating an evolutionary ultra-high capacity multilayer transport network, based on optimized integration of optical and packet nodes, and equipped with a multi-domain, multi-technology control plane. The GMPLS/PCE Unified Control Plane aspects of STRONGEST were further developed by CTTC ADRENALINE[7]. ADRENALINE is an optical circuit-switched WDM testbed that deploys reconfigurable OADMs based on AWG's and tunable lasers. End-to-end lightpaths are set up and torn down dynamically and in real time by means of a GMPLS-based control plane (switched connections) and a distributed management plane (soft-permanent connections). Moreover, the utilization of GMPLS in core optical networks has enabled an SDN-ready solution for the backbone network. The appearance of Path Computation Element (PCE) and Application Based Network Operation (ABNO) architecture allows having a SDN solution for fix and flexi grid networks. These paradigms were demonstrated in recent works [8], [9] showcasing the dynamic provisioning of services in backbone networks, considering multi-layer and multi-technology environments.

From an access network control plane perspective, a dynamic SDN control plane can significantly speed up service provisioning to end users in power-split passive optical networks [11], as well as improved delivery performance by in foxed/mobile converged scenarios [10].

In terms of SDN network controllers, while there are many options available, we briefly mention some of the work

carried out on three main solutions, focusing our attention particularly to the ONOS controller.

OpenDayLight [12] is an OpenSource SDN architectural framework, based on the Cisco Extensible Network Controller (XNC). The Service Provider edition has renderers for IETF's NetConf configuration, BGP and PCEP. OpenContrail is a tactical SDN framework, which has been adopted by Juniper as a control framework (Contrail) for its SDN compatible equipment [13]. Architecturally it is composed of four subsystems. vRouters handle network slicing, traffic steering and MPLS or VXLAN based overlay networks. The configuration subsystem manipulates the high-level service data model into a form for consumption by the devices. The Controller component manages and monitors network state.

The ONOS project [14] is a carrier-grade performance SDN platform purposely designed and supported through ONF and ON.Lab. The ONOS project has defined a number of use cases to demonstrate the carrier capability of the system. These include an SDN IP Peering use case, a Network Function Virtualisation as a Service (NFVaaS) use case which demonstrates a virtual OLT (vOLT) solution for GPON. ONOS does not rely solely on Openflow as its SDN control plane technology, as demonstrated in the Segment Routing use case. A PCE [x] use case addresses the issue of over-dimensioning (typically by a factor of 4 or 5) of current Packet Optical cores typically used to cater for network outages and peak bursts. The ONOS architecture and use cases demonstrates that there is accommodation for SDN protocols other than Openflow, particular for the orchestration of lower layers, as well as the co-ordination of multiple domains.

The ONOS-based Central Office Re-Architected as a Data Center (OpenCord) [1] project represents probably one of the most influential implementation of the SDN control framework for network operators. Its architecture replaces proprietary legacy hardware components with software running on commodity servers and off the shelf white box switches and access devices. XOS is a service orchestration layer built on top of OpenStack and ONOS that manages scalable services running in CORD. The CORD architecture consists of three sub-projects namely residential CORD (R-CORD), Mobile CORD (M-CORD) and Enterprise CORD (E-CORD). Each sub-project is a proof of concept use case for the CORD framework for demonstrating its ability to accommodate a wide range of technologies in a software-defined architecture.

III. SDN ARCHITECTURE ENABLING E2E SERVICE ORCHESTRATION

A. Reference Network Architecture

The reference network architecture for this work was taken from the DISCUS project, based on LR-PONs.

While the main contribution of this paper is end-to-end network orchestration across generic access and backbone networks, our test case is based on a LR-PON scenario.

The main characteristic of the network architecture is the flexibility that can be provided by a power split optical network with wavelength tunable terminations. By tuning network termination cards (Optical Line Terminals - OLTs)

and users side termination equipment (Optical Network Units - ONUs) it is possible to dynamically move end users into different channels, to better manage capacity across all the wavelength channels of a PON.

In addition, as shown in Figure 1, the node architecture is centred around a large port-count optical switch that enables flexible connectivity of multiple technologies to access and core fibre. For example a business customer could be served by a 400G link operating over coherent transmission technology, using the same PON physical infrastructure of other residential and business users. The large port-count in the optical switch is achieved through a 2-stage folded Clos architecture capable of scaling to over 12000 ports using 192-port switching matrices [16].

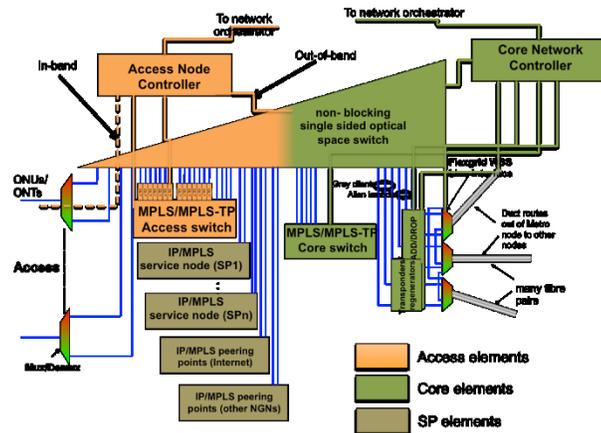


Fig. 1. DISCUS metro/core node architecture

B. SDN Architecture

The SDN control plane architecture is shown in Fig. 2, and it is based on a hierarchical structure of controllers. This architecture is consistent with the Open Network Foundation (ONF) SDN framework. Three main logical components can be identified for the SDN control plane. The access network controller, in charge of controlling the access network elements; the backbone network controller, in charge of controlling the elements carrying out backbone transmission and the network orchestrator, in charge of taking requests from application plane and translating them into high-level commands for the access and backbone network controllers.

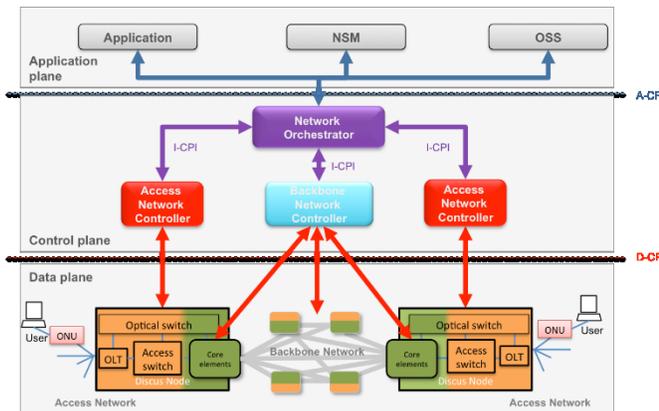


Fig. 2. Proposed SDN Architecture

In terms of interfaces, the overall architecture follows the

ONF architecture using three main Controller plane interfaces (CPI): A-CPI, I-CPI and D-CPI. The A-CPI interface describes the interaction between the control plane and the application plane. Thus is the interaction between the service provider and the network orchestrator. The I-CPI interface operates between the network orchestrator and the access and core network controllers. Lastly, the D-CPI interface operates between the access and backbone controllers and the physical devices. There are different protocols to cope with the functionalities in each interface. The demonstration in this work has selected a set of protocols to fulfil the requirements in each domain aligned with the industry trends.

C. Network Orchestrator

The network orchestrator is defined as a parent controller or a centralized “controller of controllers”, which handles the automation of end-to-end connectivity provisioning, working at a higher, abstracted level and covering inter-domain aspects between the access and the metro/core network. The network orchestrator interfaces with the controllers to get topological information about the resources in each controller’s domain. Each controller may have different interfaces, which requires the orchestrator to have a method to support multiple technologies or interfaces.

When an application, such as Network Management System (NMS) or Operation Support System (OSS) requests a service, the network orchestrator must deal with the end-to-end path computation and service request. This process can be done by the orchestrator or delegated to the access and backbone controllers. Once the services are set-up, the network orchestrator is in charge of updating its status and notify to the application plane.

D. Access Network Controller

The SDN Access Network Controller shown in Fig. 3 provides the required capacity over the access network following a request from the network orchestrator. The orchestrator communicates with the access controller through a Java Script Object Notation (JSON) (REST/API) interface.

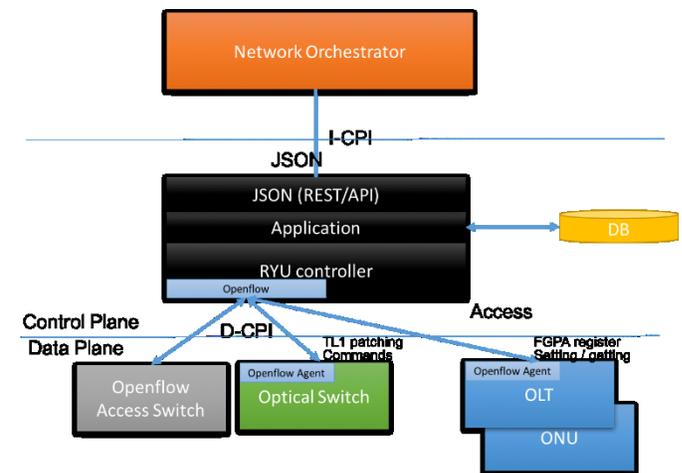


Fig. 3. Access-Metro Network Controller

The application module processes the incoming request, based on the state information present in the database. This

module implements functionalities such as path calculation, path recovery, wavelength selection, bandwidth assignment, PseudoWire (PW) assignment and Link State Protocol (LSP) assignment. Based on the request and the network state, it determines whether the request can be satisfied or should be declined. The application module triggers the appropriate OpenFlow commands using the RYU OpenFlow controller.

Quality of service assurance is achieved through the use of meter tables in the OpenFlow access switch to control Peak Information Rate (PIR) and Committed Information Rate (CIR) for each flow. This is achieved by pipelining OpenFlow tables, so that PIR is applied first to discard data above this rate, and CIR subsequently by marking packets that exceed the CIR bitrate as low priority ($\text{prec_level}=0$). The database module stores all information from the MC node on routing, wavelengths, capacity, MPLS labels, and detail of flows and meters being used.

The Access Network controller uses OpenFlow also to manage the other network components. We have devised an OpenFlow Agent that provides a mediation interface between the upper OpenFlow Control plane and the lower non-native OpenFlow network devices such as the FPGA-based LR-PON OLTs and the optical switch. The ONU is instead controlled by the OLT through an internal management channel, as it typically occurs in today's PON systems.

E. Backbone Network Controller

The backbone controller is in charge of receiving commands from the network orchestrator and transforming them in the D-CPI for the metro/core network. Similarly, it exports the topology to the network orchestrator, so it can have a view of the resources in the backbone network. The network orchestrator can request a path computation to the backbone controller, so it must support path computation within its domain.

A backbone controller is used as an entity, which is in charge of the specifics of the underlying backbone technologies. The technologies controlled by the backbone controller are Optical Transport Network (OTN), Wavelength Switched Optical Network (WSO), Spectrum Switched Optical Network (SSON) networks, which are based on the GMPLS distributed control plane. If the GMPLS is enabled, the best interface with the nodes is Path Computation Element Protocol (PCEP), as demonstrated in [8]. This hybrid approach with GMPLS and a controller allows maintaining the advantages of a reliable distribute control plane, while having a central entity to optimize the overall network.

IV. CONTROL ORCHESTRATION PROTOCOL

The Network Orchestrator is in charge of the control (e.g., E2E transport service provisioning), at a higher, abstracted level, of end-to-end resources across multiple domains with heterogeneous multi-layer multi-vendor transport network technologies regardless of the specific control plane technology employed in each domain. Typically, domain SDN controllers have a northbound interface (NBI) which is technology and vendor dependent, so the Network Orchestrator has to implement different plugins for each of the domain SDN controller's NBI [17]. During the latest

OIF/ONF Transport SDN demonstration [1], vendors tested prototypes of SDN controllers NBI for Service Request and Topology functions in development by the OIF. The framework of the demo was application-based bandwidth-on-demand between data center sites.

The Control Orchestration Protocol (COP) [18] is the first defined Transport API, which abstracts the particular control plane technology of a given transport domain. COP provides a research-oriented multi-layer approach using YANG and RESTconf. COP was the first to define an SDN controller NBI using YANG. Currently, several SDOs are focusing on defining such a standard for transport networks (i.e., ONF Transport API [19] and IETF TEAS topology and tunnel models [20][21]). The authors are actively contributing in these SDOs.

It can be both seen as a solution for the I-CPI, between the Network Orchestrator and domain SDN controllers, as well as a detailed A-CPI for an application (i.e., NMS, OSS) running on top of the Network Orchestrator. Moreover, the COP also enables the integration of heterogeneous radio access networks (5G, mmWave, LTE/LTE-A, Wi-Fi, etc.) with transport networks as well as the orchestration of cloud resources and transport resources for DC interconnection.

In brief, the COP is composed of four main base functions:

- 1) *Topology service*, which provides topological information about the network, including a common and homogeneous definition of the network topologies included in the Traffic Engineering Databases (TED) of the different control instances;
- 2) *Call service*, based on the concept of Call/Connection separation, and providing a common provisioning model which defines an end-to-end connectivity provisioning service.
- 3) *Path computation service*, providing an interface to request and return path objects which contain the information about the route between two endpoints;
- 4) *Notification Framework*, which provides an asynchronous mechanism for providing event information messages (event reports) for the COP.

The COP definition is open for discussion and can be downloaded and contributed at <https://github.com/ict-strauss/COP>.

The **Topology Service** provides abstracted topological information about the network. It must include a common and homogeneous definition of the network topologies included in the TED of the different control instances. The COP Topology provides the interface for the exchange of detailed network topology between SDN controllers. An abstract Topology object may consist of a set of nodes and edges, which form a tree structure. A Node must contain a list of ports/endpoints and their associated switching capabilities (e.g., packet/lambda switch capable). An Edge object is defined as the connection link between two Endpoints, which includes some traffic metrics and characteristics. 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, Endpoint) must be extensible, able to include TE extensions to describe different switching capabilities (i.e., time-slots,

packets, wavelengths, frequency slots).

The **Call service** is defined in the scope of COP as the provisioning interface of each domain SDN controller. The Call is defined as the API's basic data-container and it completely solves the provisioning of an End-to-End connectivity service. A Call object describes the type of service that is requested or served by it (e.g., DWDM, Ethernet, MPLS). It also contains the endpoints among which the service is provided. The Call object also includes the list effective connections made into the data plane, to support the service call. Moreover, a Call also includes a set of traffic parameters that need to be supported, such as Quality of Service, or allocated bandwidth. Finally, a Call shall include the Match parameter, which refers to the traffic description to which the Call refers. A Connection object is used for a single network domain scope. It should include the path object (i.e., the route) across the network topology the data traverses, which may be fully described or abstract depending on the orchestration/control schemes used. Each connection must be associated with a single control plane entity (e.g. a SDN controller) responsible for the configuration of the data path. Finally, the Call also introduces the necessary TE parameters (e.g., bandwidth) that the service requests.

Both the Topology and Call services can be extended in order to not only provide abstracted information, but to provide low level technological dependant details, such as optical spectrum grid or CIR/PIR.

The **Path Computation** service should provide an interface to request and return Path objects, which contain the information about the route between two Endpoints. 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. Furthermore, each component in the Path object is represented as an Endpoint with TE information associated to it.

Finally, the **Notification Framework** should provide Notifications refer to the set of autonomous messages that provide information about events, such as alarms, performance monitoring (PM) threshold crossings, object creation/deletion, attribute value change (AVC), or state changes.

Notifications specifications are generally written around a model of a subscriber and a notification server. The term subscriber is used to name an entity that requests notification subscriptions and receives notification messages. The term notification server is used to designate an entity that recognizes events, turns them into notification messages, and transmits them to pertinent subscribers.

Currently, the COP only provides few notification examples, in order to asynchronously provide an alarm that a Call has been modified (e.g., due to recovery), or to notify that a Call service cannot be fulfilled due to a change in network conditions.

Finally, it is important to describe that COP enables hierarchical control of networks, which allows better

scalability through abstraction. Regarding COP complexity, JSON parsers have demonstrated [22] to be able to handle within the order of milliseconds the received instructions. In view of this, and the connectivity setup times in real hardware, there is not a clear need for a binary-oriented protocol.

V. TEST CASE FOR DEMONSTRATION

The test case shown in Fig. 4 exemplifies how capacity constraints in the Core and Access Metro network of a Telecommunications network may be overcome both in terms of restoration of QoS but also opportunistically for the dynamic provision of high bandwidth services. In the Metro Access network, we implement a DWA mechanism that can facilitate allocation of capacity (an entire wavelength or a portion of a TDM-PON channel) to specific PON users on a dynamic and temporary basis, with applications for on-demand multimedia and big data transfers. In the cases of temporary and dynamic migration of wavelength it is important that the wavelength switching time is minimized to avoid impacting users of other network services. In Step 1, an end user (customer) elects through a portal frontend to the ABNO controller to transmit a fixed bandwidth transport (100Mbps) between two end points aEnd (10.0.50.3) and zEnd (10.0.50.4). This action could also be automated by the specific application, for example for a quality-assured video or video conference streaming. In Step 2, the ABNO orchestrates the provision of the path, according to its knowledge of the full End-to-End topology covering the Core and Metro Access Networks. The Idealist core backbone Network Ccontroller (NC) and the Metro Core (DISCUS) SDN NC, are each instructed to provision an explicit path in steps 3 and 4 respectively. For the access portion of the MC node (the path from 10.0.50.2 to 10.0.50.4), the DISCUS SDN NC is instructed to provision the path according to the route 10.0.50.2 to 10.0.50.1 to 10.0.50.4. The DISCUS SDN Controller provisions the path through the metro (openflow) switch (step 5), and the PON (primary/secondary OLT and ONU – step 6). Using a custom implemented PLOAM message, the primary OLT requests the ONU tune to a wavelength provisioned out of the secondary OLT (step 7). In step 8, the DISCUS SDN NC, acknowledges that it has completed the provisioning of the path through the access portion of the network. In step 9, the video transmission is triggered to start.

VI. EXPERIMENTAL VALIDATION

A. Test-bed

There are two different labs set-up to demonstrate the scenario of this work, one in Telefonica premises in Madrid and another in the Optical Network Architecture (ONA) lab of Trinity College of Dublin. Fig. 4 shows a schema of the lab set-up for this experiment. The data plane interconnection is through a low bandwidth VPN.

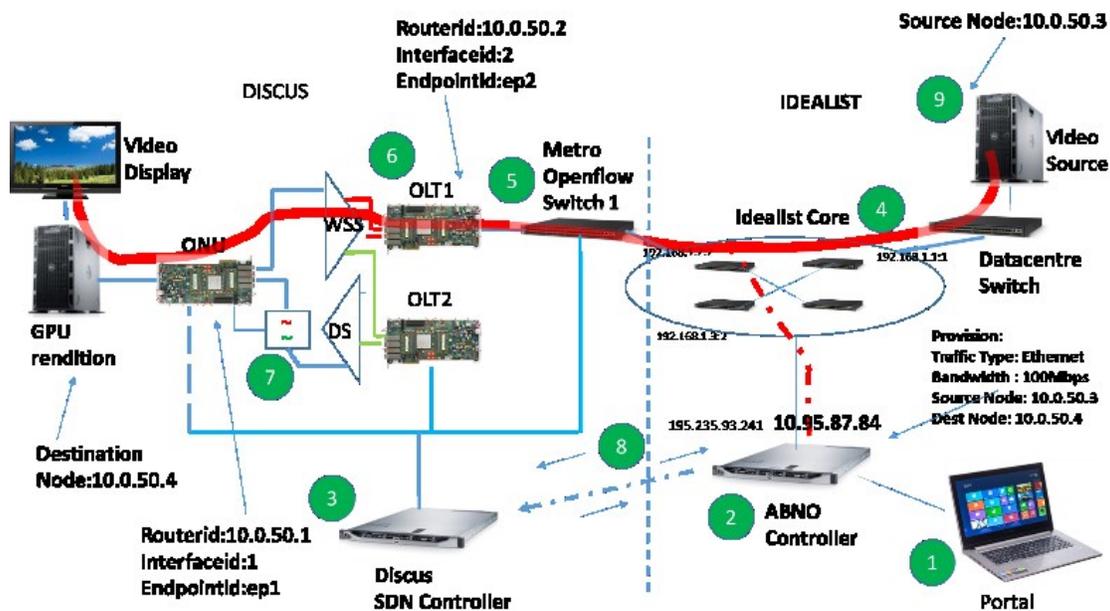


Fig. 4. Experimental laboratory set-up

The network orchestrator is located in Telefonica labs and is based on Netphony ABNO implementation¹. The north and south-bound interface of the orchestrator is implemented using the STRAUSS COP [18].

The backbone controller uses the netphony ABNO, which in addition to netphony PCE controls a GMPLS emulated control plane. The GMPLS nodes uses the protocol suite developed by Telefonica I+D and is released in github². This setup is built with 30 virtual machines, which run in a Linux server distribution. Each emulated node implements a GMPLS stack (including RSVP, OSPFv2 and PCEP) with the extensions to support flexgrid developed in IDEALLIST project.

The LR-PON Protocol is implemented over three Xilinx VC709 FPGA boards acting as primary and secondary OLTs and ONU. The LR-PON protocol is a partial implementation of the XGPON standard, modified to suit the longer logical reach and larger split size of a LR-PON (additional details are available in [19]). The PON backplane connection to the core network contains a 10G Ethernet physical layer and Media Access Control Layer, allowing it to be plugged into any 10G capable network element. In this experiment the PON backplane is connected to a 48-port 10G Openflow switch. A Microblaze soft processor, which is collocated on the Virtex FPGA board, provides a (North Bound) UART management interface to the PON OLT and ONU hardware. Through this interface most PON functionality can be controlled such as resetting the hardware, viewing hardware status, simulating hardware failure, loading bandwidth map and setting XGEM mappings. Dynamic Wavelength Assignment is supported through the addition of laser and filter control to the LR-PON protocol hardware and control mechanisms. The tunable laser is controlled across an i2c bus to the Skylake 10G SFP+ tunable lasers and the tunable filter is controlled through a UART. To implement DWA in

the physical layer, we employ a splitter and filter in the downstream, and a wavelength multiplexer in the upstream direction patched to ITU channels 32.5 and 31 of the primary and secondary OLT respectively. In order to select the OLT and ONU transmission wavelengths, the OLT provides a North Bound interface. Through this Interface, the control plane can tune an OLT's transceivers to a given wavelength. Since the ONU is remote from the control plane, tuning of an ONU's laser and filter is performed through this interface by the invocation of a custom PLOAM message within the LR-PON protocol. The wavelength of the OLT and ONU may be selected by writing to control registers in the OLT. Each individual OLT laser wavelength can be set by writing the ITU channel number to its local register. To select the transmission wavelength for an ONU, the ITU channel number is set by writing the target ITU channel number to register of the OLT to which it is homed. The ONU Id must also be specified so as to distinguish an individual from multiple ONU's homed off a single OLT.

For the experimental LR-PON access network, the configuration is comprised of a Pronto 3780 switch with 48 10G interfaces, running release 2.4 (Openflow v1.4 compatible firmware). The Pronto switch is configured with multiple virtual bridges.

A Video Server (VLC) application is co-located with the ABNO controller interface in the Telefonica premises. This transmits a UDP based video stream across the Tunnel between the two testbeds, traversing the DISCUS PON and being received by the GPU workstation for display by the TV display. It should be noticed that UDP was used in order to assess the raw network performance, undisturbed by the TCP flow control mechanism. However we envisage that video streaming will generally be operated over TCP. The latency between the two testbeds was measured over 100 measurements at between 45 and 48 ms. It was not possible to transmit the video stream through the Idealist network, as

¹ <https://github.com/telefonicaid/netphony-abno/wiki>

² <https://github.com/telefonicaid/netphony-network-protocols>

a suitable physical DataPath was not available.

B. Workflow steps to obtain the topological information

This section explains how the topological information is obtained from the access and backbone controller. The STRAUSS ABNO pulls periodically each of the controllers in its context: for this scenario the backbone (IDEALIST) and access controllers (DISCUS). Another policy to obtain the topology is to request the information when a create service call is received. The problem with this approach is that it increases the delay in the provisioning process. Moreover, we do not expect a very dynamic topology, so we have chosen the periodical pull for this experiment. Let us highlight that each controller is in charge of obtaining the topology in their domain. For instance, the IDEALIST PCE obtains the topological information based on the IGP which is OSPF in the Netphony implementation. The workflow is displayed in Fig. 5.

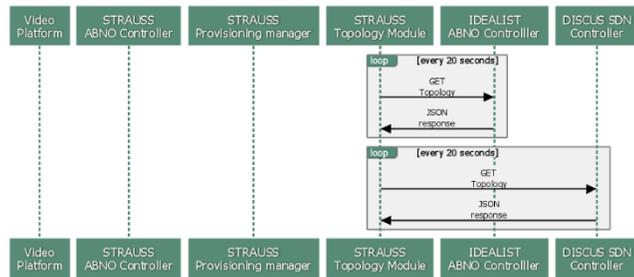


Fig. 5. Workflow steps for the topology service

```
127.0.0.1 127.0.0.1 8087 228 GET /restconf/config/topologies/topology/0.0.0.1/ HTTP/1.1
127.0.0.1 127.0.0.1 48800 75 HTTP/1.1 200 OK (application/json)
127.0.0.1 127.0.0.1 8078 228 GET /restconf/config/topologies/topology/0.0.0.1/ HTTP/1.1
127.0.0.1 127.0.0.1 34622 78 HTTP/1.0 200 OK
```

Fig. 6. Whireshark capture with the topology service

```
topology [1]
  edges [5]
  0 {0}
  1 {0}
  nodes [4]
    0 {5}
    1 {5}
    2 {5}
      domain : 1
      nodetype : transport
      name : ONU
      edgeEnd [2]
        0 {2}
          edgeEndId : 1
          name : 1
        1 {2}
          edgeEndId : 2
          name : 2
      nodeId : ONU
    3 {5}
      domain : 1
      nodetype : transport
      name : OLT1
      edgeId : OLT1_2_ONU_1
      edgeType : eth_edge
      metric : 1.0
      source {5}
        domain : 1
        nodetype : transport
        name : OLT1
        edgeEnd [2]
          0 {2}
            edgeEndId : 1
            nodeId : OLT1
          1 {2}
            edgeEndId : 2
            name : 2
        localIfid {2}
          edgeEndId : 2
          name : 2
        remoteIfid {2}
          edgeEndId : 1
          name : 1
        target {6}
          domain : 1
          nodetype : transport
          name : ONU
          edgeEnd [2]
            nodeId : ONU
            underlayAbstractTopology [0]
          2 {0}
          3 {0}
```

Fig. 7. JSON object for the topology service

Fig. 6 shows the message exchange for the topology acquisition process. This is tested in the lab environment. As displayed in Fig. 6, the STRAUSS ABNO sends HTTP GET queries to the DISCUS controller and the IDEALIST PCE. Fig. 7 shows the JSON object with the response from the DISCUS controller. It shows the abstracted topology, essentially a list of nodes and links: the nodes contain information about their domain, name, id, type and interfaces (edgeEnd). Also the link from OLT1 to ONU is showed as example: the information contained in a link includes name, id, type, metric and the node and interface information for source and destination.

C. Workflow steps to create a service

This section explains step by step how this demonstration is performed (Fig. 4). The topological information is obtained, as previously explained. The portal requests a new video service, which cannot be processed within the access area scope. This means that there is a request of the video platform to provisioning an end to end path between the client and the video server. Therefore, the STRAUSS ABNO receives a COP call service set-up to establish the connection. The STRAUSS ABNO carries out a path computation, which crosses different networks, the backbone (IDEALIST) and access network (DISCUS). Therefore, the STRAUSS ABNO sends a COP calls service set-up to each controller to configure the nodes in their domain. The IDEALIST PCE configures the GMPLS nodes, while the DISCUS controller configures the access elements. The workflow is explained in Fig. 8.

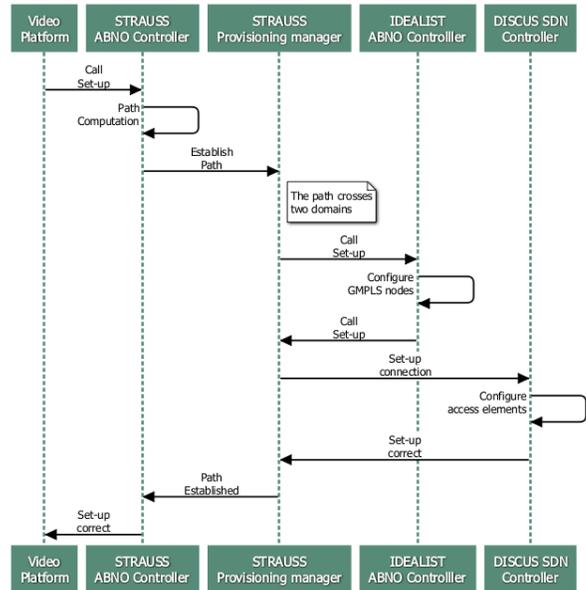


Fig. 8. Workflow steps for the service-call set-up

The previous workflow was demonstrated in the lab set-up. Fig. 9 shows the message exchange between the different elements.

Time	Source	Destination	Length	Protocol	Application	Notes
127.0.0.1	127.0.0.1	127.0.0.1	547	POST	/restconf/config/calls/call/1/ HTTP/1.1	ABNO
127.0.0.1	127.0.0.1	127.0.0.1	132	Path Computation Request (PCReq)		PCE
127.0.0.1	127.0.0.1	127.0.0.1	156	Path Computation Reply (PCRep)		PM
127.0.0.1	127.0.0.1	127.0.0.1	164	Path Computation LSP Initiate (PCInitiate)		Idealist
10.95.164.204	10.95.164.204	10.95.164.204	73	POST	/restconf/config/calls/call/1/ HTTP/1.1	Controller
10.95.164.204	10.95.164.204	10.95.164.204	73	HTTP/1.1	200 OK (application/json)	Controller
127.0.0.1	127.0.0.1	127.0.0.1	820	POST	/restconf/config/calls/call/1/ HTTP/1.1	Discus
127.0.0.1	127.0.0.1	127.0.0.1	107	HTTP/1.0	200 OK	Controller
127.0.0.1	127.0.0.1	127.0.0.1	164	Path Computation LSP State Report (PCRpt)		
127.0.0.1	127.0.0.1	127.0.0.1	73	HTTP/1.1	200 OK (application/json)	

Fig. 9. Wireshark capture with a COP service-call set-up

As it is shown in Fig. 9, the STRAUSS ABNO receives an HTTP POST request with COP syntax. Fig. 10 shows the JSON object with the request parameters. The aEnd and zEnd routerIds identifies the client and the video server. The traffic parameters are set to request a 100Mbps connection and a 100 ms latency. This request is sent as an Ethernet service.

Once the STRAUSS ABNO receives the request, it asks its PCE for a path computation between the two end points. To do so a PCReq-PCResp process is performed. Now, the PCE can calculate the path and response to ABNO with a PCResp, which contains the Explicit Route Object with the path. The ABNO controller with the ERO information call to Provisioning Manager (PM) via a PCInitiate message. The PM splits the route in different domains and forwards a COP message call to each controller to create a path in each domain (IDEALIST and DISCUS). When the path is created each controller sends respective http messages with an OK status. With this information the PM response to the ABNO controller with a PCReport message and finally ABNO reports to the video platform with an HTTP response.

```

object {5}
  callId : call_1
  aEnd {2}
    routerId : 10.0.50.3
    endpointId : ep1
  zEnd {2}
    routerId : 10.0.50.4
    endpointId : ep2
  trafficParams {2}
    latency : 100
    reservedBandwidth : 100000000
  transportLayer {2}
    layer : ethernet
    direction : bidir

```

Fig. 10. JSON object for a COP service-call set-up

Across 10 repetitions of the experiment, the total completion time of the workflow was measured at 275ms. Of this, 35 ms (with associated inter-testbed latency) related to the blocking element of the call to the DISCUS SDN Controller. The non-blocking elements of the DISCUS SDN proceed in parallel with the completion of the return calls by the ABNO controller.

VII. CONCLUSIONS

This paper has demonstrated the use of a SDN architecture to provision end-to-end resource reservation across multiple domain and technologies. The case study and

implementation are based on the network architecture of two major European Projects: DISCUS, for the design of a long-reach optical access and aggregation controlled by the DISCUS SDN controller, and IDEALIST providing the core network design and ABNO node controller and orchestrator. The communication protocol between the ABNO orchestrator and access/metro controller was developed by the Control Orchestration Protocol (COP) developed by the European STRAUSS project.

The demonstration has shown how the COP enables seamless interoperability between the orchestrator, which receives a request for dynamic end-to-end resource allocation, the core controller, which sets up a GMPLS path across the core, and the access-metro controller, operating dynamic wavelength allocation of a LR-PON to provide the requested capacity to the end user. The testbed results demonstrated the possibility to set up an end-to-end capacity reservation across two labs separated by over two thousand kilometres in less than 300 ms, practically showcasing how network operators can deploy end-to-end SDN solutions across multi-domain access and backbone networks.

As future work, the authors would like to explore scenarios where the video content is delivered between different network operators. There is a EU research project 5GEx [24] which is working in these topics.

ACKNOWLEDGMENT

The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n. 318137 in DISCUS project, in. 317999 in IDEALIST research project and no 608528 by the Japanese Ministry of Internal Affairs and Communications (MIC) and National Institute of Information and Communications Technology (NICT) through the EU- Japan coordinated project STRAUSS; from the Spanish MINECO project DESTELLO (TEC2015-69256-R); and from the Science Foundation Ireland grant 14/IA/2527.

REFERENCES

- [1] 5G empowering vertical industries, technical report, www.5g-ppp.eu, Feb. 2016.
- [2] Ruffini, M.; Wosinska, L.; Achouche, M.; Jiajia Chen; Doran, N.; Farjady, F.; Montalvo, J.; Ossieur, P.; O'Sullivan, B.; Parsons, N.; Pfeiffer, T.; Xing-Zhi Qiu; Raack, C.; Rohde, H.; Schiano, M.; Townsend, P.; Wessaly, R.; Xin Yin; Payne, D.B., "DISCUS: an end-to-end solution for ubiquitous broadband optical access," in Communications Magazine, IEEE , vol.52, no.2, pp.S24-S32, February 2014
- [3] N. Bitar, SDN and Potential Applicability to Access Networks, Paper Th1G.4, OFC 2014.
- [4] EU-FP7. "SPARC - Split Architecture Carrier Grade Networks", <http://www.fp7-sparc.eu/>. 2010
- [5] Kpsel, A., et al. . "Design and implementation of the OFELIA FP7 facility: the European OpenFlow testbed." 2013
- [6] EU ICT IP. "STRONGEST - Scalable Tunable and Resilient Optical Networks Guaranteeing Extremely-high Speed Transport" ,<http://www.ict-strongest.eu/>. 2010
- [7] Martinez, R., R. Munoz, R. Casellas, and R. Vilalta. "OpenFlow-based Hybrid Control Plane within the CTTC ADRENALINE Testbed." In Proc. OFELIA Workshop. 2011.
- [8] A. Aguado, V. López, J. Marhuenda, Ó. González de Dios and J. P. Fernández-Palacios: ABNO: a feasible SDN approach for

- multi-vendor IP and optical networks , in *Journal of Optical Communications and Networking*, February 2015, Vol. 7, Iss. 2, pp. A356–A362.
- [9] R. Muñoz, R. Vilalta, R. Casellas, R. Martínez, F. Francois, M. Channegowda, A. Hammad, S. Peng, R. Nejabati, D. Simeonidou, N. Yoshikane, T. Tsuritani, V. López and Achim Autenrieth: Transport Network Orchestration for End-to-End Multi-layer Provisioning Across Heterogeneous SDN/OpenFlow and GMPLS/PCE Control Domains, in *Journal of Lightwave Technology*, April 2015, Vol. 33, No. 8, pp. 1540 - 1548.
- [10] X. Li, et al., Joint Bandwidth Provisioning and Cache Management for Video Distribution in Software-Defined Passive Optical Networks, Paper Tu3F.5, OFC 2014.
- [11] M. Ruffini, F. Slyne, C. Bluemm, N. Kitsuan, S. McGettrick, Software Defined Networking for Next Generation Converged Metro-Access Networks, *Optical Fiber Technology*, Volume 26, Part A, December 2015, Pages 31-41.
- [12] Medved, Jan, Robert Varga, Anton Tkacik, and Ken Gray. "Opendaylight: Towards a model-driven sdn controller architecture." In *A World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, 2014 IEEE 15th International Symposium on, pp. 1-6. IEEE, 2014.
- [13] Guan, Jing, Songwei Ma, Bingli Guo, Juhao Li, Shanguo Huang, Xiaoyuan Cao, Zhangyuan Chen, Zhengbin Li, and Yongqi He. "First Field Demonstration of Network Function Virtualization via Dynamic Optical Networks with OpenContrail and Enhanced NOX Orchestration." In *Asia Communications and Photonics Conference*, pp. AF2C-5. Optical Society of America, 2013
- [14] Berde, Pankaj, Matteo Gerola, Jonathan Hart, Yuta Higuchi, Masayoshi Kobayashi, Toshio Koide, Bob Lantz et al. "ONOS: towards an open, distributed SDN OS." In *Proceedings of the third workshop on Hot topics in software defined networking*, pp. 1-6. ACM, 2014.
- [15] Al-Shabibi A, Peterson L. *CORD: Central Office Re-architected as a Datacenter*. Open Stack Summit. 2015.
- [16] J. Chen , et al., Final report on the specification of the metro/core node architecture. DISCUS deliverable D6.5, July 2015, <http://cordis.europa.eu/docs/projects/cnect/7/318137/080/deliverables/002-318137DISCUSD65renditionDownload.pdf>
- [17] V. López, O. González de Dios, L.M. Contreras, J. Foster, H. Silva, L. Blair, J. Marsella, T. Szyrkowiec, A. Autenrieth, C. Liou, A. Sasdasivarao, S. Syed, J. Sun, B. Rao, F. Zhang, and J.P. Fernández-Palacios, "Demonstration of SDN orchestration in optical multi-vendor scenarios", OFC 2015.
- [18] "Global transport sdn prototype demonstration," OIF / ONF White Paper, 2014.
- [19] R. Vilalta, V. López, A. Mayoral, N. Yoshikane, M. Ruffini, D. Siracusa, R. Martínez, T. Szyrkowiec, A. Autenrieth, S. Peng, R. Casellas, R. Nejabati, D. Simeonidou, X. Cao, T. Tsuritani, I. Morita, J. P. Fernández-Palacios, R. Muñoz, The Need for a Control Orchestration Protocol in Research Projects on Optical Networking , in *Proceedings of European Conference on Networks and Communications (EuCnC 2015)*, 29-2 July 2015, Paris (France).K. Sethuraman (ed), et al., "Functional Requirements for Transport API", ONF TR-527, Jun. 2016.
- [20] X. Liu, et al. "YANG Data Model for TE Topologies.", draft-ietf-teas-yang-te-topo-06, IETF, Oct. 2016.
- [21] T. Saad (ed), et al. "A YANG Data Model for Traffic Engineering Tunnels and Interfaces", draft-ietf-teas-yang-te-05, IETF, Oct. 2016.
- [22] A. Mayoral, R. Vilalta, R. Muñoz, R. Casellas, R. Martínez, M. Svaluto-Moreolo, J. M. Fàbrega, S. Yan, A. Aguado, E. Hugues-Salas, S. Peng, F. Meng, Y. Shu, G. Zervas, R. Nejabati, D. Simeonidou, J. M. Gran, V. López, O. González 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, K. Kitayama, First experimental demonstration of a distributed cloud and heterogeneous network orchestration with a common Transport API for E2E services with QoS, ECOC, 2016.
- [23] S. McGettrick, F. Slyne, N. Kitsuan, D.B. Payne, and M. Ruffini, Experimental End-to-End Demonstration of Shared N:M Dual Homed Protection in SDN-controlled Long Reach PON and Pan-European Core. To appear in *Journal of Lightwave Technology*. Pre-print access available at http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=7517316&filter%3DAND%28p_IS_Number%3A4357488%29
- [24] 5G Exchange (5G Ex) <https://5gex.mit.bme.hu/>
- Victor López** (M.Sc. 2005 - Ph.D. 2009) is a Technology Expert at the Technology and Planning for Transport, IP and Interconnection Networks department in Telefonica Global CTO. He is involved in the research projects as well as in the definition of the Telefonica network guidelines for IP/optical segment. He started his carrier in Telefónica I+D as a researcher (2004). He obtained his Ph.D at Universidad Autonoma de Madrid, where he worked as an Assistant Professor. In 2011, he joined Telefonica I+D as Technology specialist working on funded research project from the Telefonica group and the European Commission. He has co-authored more than 100 publications and contributed to IETF drafts.
- José Manuel Gran** (B.Sc 2015) ia a research engineer at telematics department in Universidad Carlos III de Madrid (UC3M). He obtained his four-year degree (with the highest qualification in the degree thesis) at Universidad de Valladolid, where he started to collaborate with Telefónica I+D, although he continued working at IMDEA headquarters in Global CTO. He has also been working at IMDEA Networks Institute and collaborating with Telefónica S.A. His work has always been related to SDN and NFV technologies focused on PCE, BGP and 5G.
- Victor Uceda** is a senior student in the double B.Sc. degree in Mathematics and Computer Science of Universidad Autónoma de Madrid. In 2013, 2014 and 2015, he was awarded the excellence scholarship granted by Comunidad de Madrid. He has collaborated with the High Performance Computing and Networking (HPCN) research group at UAM where he has developed mechanisms for both traffic- capture storage and analysis. Currently, he collaborates with Telefonica I+D where he has developed functionalities for network orchestration software.
- Frank Slyne** is completing his PhD at Trinity College Dublin in the research area of Metro Nodes as implemented in a flat-core and Long Reach Passive Optical Network (LR-PON) architecture. His interests are the application of software defined networks for the optimisation of capacity, energy and performance. He has worked for a number of years at Eircom where he was responsible for ISP platforms and systems. He has a B.E. (Elect) from UCC (Cork) and an M.Eng from DCU (Dublin).
- Marco Ruffini** received his M.Eng. in telecommunications in 2002 from Polytechnic University of Marche (Italy), and after working for Philips R&D in Germany he joined Trinity College Dublin (TCD) in 2005, where he received his Ph.D. in 2007. Since 2010 he's assistant professor (tenured 2014) at TCD. He is part of the CTVR/CONNECT telecommunications research centre and a technical coordinator of the FP7 DISCUS project and holder of a personal Science Foundation Ireland Principal Investigator award. He has authored more than 80 publications and 10 patents. His research focuses on flexible high-capacity fiber broadband architectures and open access SDN.
- Ricard Vilalta** (M.Sc. 2007 – Ph.D. 2013) is a researcher at CTTC, in the Optical Networks and Systems Department. 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, His research is focused on Transport SDN, NFV and cloud/fog computing. He is currently a Research Associate at Open Networking Foundation. He has been involved in several international, EU,

national and industrial research projects. He has also authored and co-authored more than 100 journals, conference papers and invited talks.

Arturo Mayoral (Madrid, 1987) graduated in Telecommunications Engineering by the Universidad Autónoma de Madrid in 2013, now is a Ph.D. student within the department of Theory of Signals and Communications (TSC) of Universitat Politècnica de Catalunya (UPC). He started his professional career in 2012, as an undergraduate researcher in Telefonica I+D (R&D) where developed his Final Career's Project (PFC), and he was awarded with the Best Final Project on Networks and Telecommunications Services Prize by the COIT and FUNDACIÓN TELEFÓNICA. In September 2013, he joined the Centre Tecnològic de Telecomunicacions de Catalunya (CTTC) as SW engineer and Research Assistant in the Communications Network Division (CND) within the ICT-STRAUSS European Project. His research areas include Software Defined Networking (SDN), Network Function Virtualization (NFV) and Cloud computing.

Raúl Muñoz 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 Researcher, Head of the Optical Network and System Department, and Manager of the Communication Networks Division. His research interests include control and management architectures, protocols and traffic engineering algorithms for future optical transport networks.

Ramon Casellas 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.

Ricardo Martínez graduated and PhD in telecommunications engineering by the UPC-BarcelonaTech university in 2002 and 2007, respectively. He has been actively involved in several public-funded (national and EU) R&D as well as industrial technology transfer projects. Since 2013, he is Senior Researcher of the Communication Networks Division (CND) at CTTC. His research interests include control and network management architectures, protocols and traffic engineering mechanisms for next-generation packet and optical transport networks within aggregation/metro and core segments.