

NFV Orchestration Over Disaggregated Metro Optical Networks With End-to-End Multi-Layer Slicing Enabling Crowdsourced Live Video Streaming

ABUBAKAR SIDDIQUE MUQADDAS^{1*}, RODRIGO STANGE TESSINARI¹, RAMON CASELLAS², MIQUEL GARRICH³, EMILIO HUGUES-SALAS¹, ÓSCAR GONZÁLEZ DE DIOS⁴, LOURDES LUQUE⁴, ALESSIO GIORGETTI^{5,6}, ANDREA SGAMBELLURI⁵, FILIPPO CUGINI⁵, FRANCISCO-JAVIER MORENO-MURO³, ROBERTO MORRO⁷, KRISTAN FARROW⁸, ADRIAN WONFOR⁹, MAYUR CHANNEGOWDA¹⁰, PABLO PAVÓN-MARIÑO³, ANDREW LORD⁸, REZA NEJABATI¹, AND DIMITRA SIMEONIDOU¹

¹High Performance Networks Group, University of Bristol, Merchant Venturers Building, Woodland Road, Bristol, BS8 1UB, United Kingdom

²CTTC/CERCA, Av. Carl Friedrich Gauss n7, 08860 Castelldefels, Barcelona, Spain

³Universidad Politécnica de Cartagena, Cuartel de Antiguones, Plaza del Hospital 1, 30202 Cartagena, Spain

⁴Telefónica Global CTO, Spain

⁵CNIT/SSSUP, Via Moruzzi n. 1, 56124 Pisa, Italy

⁶CNR-IEIT, Pisa, Italy

⁷TIM, Via G. Reiss Romoli n. 274, 10148 Torino, Italy

⁸British Telecommunications Ltd, United Kingdom

⁹University of Cambridge, United Kingdom

¹⁰Zeetta Networks Limited, Bristol, United Kingdom

*Corresponding author: abubakar.muqaddas@bristol.ac.uk

Compiled June 22, 2021

Network infrastructure must support emerging applications, fulfill 5G requirements and respond to the sudden increase of societal need for remote communications. Remarkably, Crowdsourced Live Video Streaming (CLVS) challenges operators' infrastructure with tides of users attending major sport or public events that demand high bandwidth and low latency jointly with computing capabilities at the networks' edge. The Metro-Haul project entered in scene proposing a cost-effective, agile and disaggregated infrastructure for the metro segment encompassing optical and packet resources jointly with computing capabilities. Recently, a major Metro-Haul outcome took the form of a field trial of Network Function Virtualization (NFV) orchestration over multi-layer packet and disaggregated optical network testbed that demonstrated a CLVS use case. We showcased the average service creation time below 5 minutes, which met the Key Performance Indicator (KPI) as defined by 5G PPP. In this paper, we expand our field trial demonstration with a detailed view of the Metro-Haul testbed for the CLVS use case, the employed components and their performance. The throughput of the service is increased from approximately 9.6 Gbps up to 35 Gbps per VLAN with high-performance VNFs based on SR-IOV technology. © 2021 Optical Society of America

<http://dx.doi.org/10.1364/ao.XX.XXXXXX>

1. INTRODUCTION

Telecom operators seek cost-effective technologies for their infrastructure while responding to the pressure coming from at least three fronts. In line with forecasts, operators have to accommodate the ever-expanding cloud services and ever-increasing traffic carried by Content Delivery Networks (CDN) [1]. At the same time and unexpectedly, the recent crisis generated by the COVID-19 pandemic has further increased the usage of re-

ote communications in society both for personal and business needs [2, 3]. On top of that, networks must evolve toward a new generation (5G and beyond) that is aimed at supporting new services with stringent requirements [4]. This three-fold combination poses challenges for services provided on top of carriers' infrastructure.

Crowdsourced Live Video Streaming (CLVS) is a clear example of such emerging applications to be supported by operators' infrastructure. CLVS is a Network Service (NS) in which

thousands of users attending an event (sports, concerts, etc.) stream video from their smartphones to a platform hosted in cloud-computing facilities [5]. The content from all the users is edited in real time, producing an aggregated video, which can be broadcast to a large number of viewers. Such service represents a challenge for the network infrastructure as it combines tides of users, high-bandwidth and low-latency requirements for a valuable experience.

In order to support such new vertical applications, Metro-Haul project [6] has designed and built a cost-effective, agile and disaggregated network infrastructure for the optical and packet resources in the metro segment. Moreover, Metro-Haul has contemplated the key role of compute capabilities at the network edge, addressing the capacity increase and characteristics such as low latency and high bandwidth. Metro-Haul control plane consists of the Control, Orchestration and Management (COM) system [7], based on the principles of ETSI Network Function Virtualization (NFV) framework and a hierarchical SDN control plane, which facilitates the deployment of multi-layer end-to-end network slices, including Virtual Network Functions (VNFs) in multiple datacenters (with multiple Virtual Infrastructure Managers (VIMs)) and simultaneously dedicates packet and optical network resources, including Layer 2 (L2) VPNs and photonic media channels. Within the context of Metro-Haul, we recently reported a field trial of NFV orchestration using COM system over multi-layer packet and disaggregated optical network testbed that demonstrated a CLVS use case with service creation time below 5 minutes for the reported NS [8].

This paper extends [8] with the following contributions. First, our work is contextualized with a literature review of related work. Second, we expand our description of the Metro-Haul testbed for the CLVS use case encompassing all elements in the data plane and all components in the control plane. Third, we report the in-depth services' and per-component performance in terms of setup and configuration times respectively, which are explained in companionship with an extensive workflow. In terms of performance, we report that throughput is increased by means of SR-IOV technology as compared to VirtIO based VNF network interfaces. Note that we utilize the terms: datacenter, site or node interchangeably in this work.

The remainder of the paper is organized as follows. In Sec. 2, we describe the CLVS operation principle and its conception as a network service. Sec. 3 reviews relevant works and highlights the completeness of our field trial reported here. Sec. 4 provides an in-depth description of the Metro-Haul testbed for the CLVS use case. Sec. 5 describes the workflow for deploying the CLVS service, which entails a hierarchy of services from the top of the COM system down to the data plane. In Sec. 6, we report results in terms of setup times, optical connectivity and throughput. Finally, we present our concluding remarks in Sec. 7.

2. CROWDSOURCED LIVE VIDEO STREAMING NETWORK SERVICE

In the past few years, videos streamed on CLVS platforms have generated a staggering amount of traffic: 3.8 million people broadcasted live-streaming videos [9] in Feb. 2020 on Twitch [10], a CLVS platform very popular with the gaming community. The amount of traffic generated from this huge user base in context of a CLVS platform [5] requires a high capacity network in combination with low latency network paths for multiple video feeds. Furthermore, a CLVS-based NS may also require compute resources to process the video traffic [11]. A reference CLVS

NS is shown in Fig. 1. It consists of CLVS traffic as multiple video streams in different VLANs, from access network being processed by a Deep Packet Inspector (DPI) VNF for traffic analysis. Via a metro-optical network, the CLVS traffic is forwarded to a Firewall (FW) VNF which can block any video stream per VLAN based on some criteria. Following this, the CLVS traffic is forwarded to the core network for upload to a CLVS platform.




Fig. 1. ETSI-NFV network service targeting CLVS use case

Metro-Haul project which provides an agile and programmable metro optical network, is targeted for various use cases from vertical industries having varying latency and bandwidth requirements. CLVS use case can directly benefit from Metro-Haul infrastructure, since its COM system can rapidly provision a slice on top of the end-to-end compute/network resources and support the high capacity and low latency requirements.

3. RELATED WORK

This section reviews relevant works in the context of our field trial, which commonly comprise experimental validations and demonstrations in which software modules interact for allocation of compute and network resources, namely implementations that have real-time actions over infrastructure. We divide these works in three categories. First, we cover related initiatives, highlighting the uniqueness of our work reported as it encompasses from the top vertical application down to the underlying optical infrastructure. Second, we cover works which emphasize developments on top of hierarchical architectures regarding planning and orchestration. Finally, we conclude this section by focusing on the bottom part of hierarchical architectures reviewing recent contributions in optical SDN controllers.

A. Related initiatives

Here, we report some contemporary initiatives to the Metro-Haul project, which encompass backbone networks, orchestration, Software Defined Networking (SDN) and NFV in 5G scenarios.

Authors in [12] reported NFV orchestration over wireless networks, where the VNFs hosted in different testbeds are interconnected via LTE, WiFi or mmWave technologies. In contrast, our work focuses on the NFV orchestration over metro optical networks, where the work in [12] may be targeted for the NFV orchestration on the access network side. Another initiative namely 5TONIC [13] is an NFV Management and Network Orchestration (MANO) platform having the capability of incorporating an external site to exploit the resources for multi-site network service experimentation. However, the inter-site communication is facilitated by an external VPN service in 5TONIC which is not part of the NFV orchestrator. Another relevant work [14] uses the combination of Open Source MANO (OSM) NFV Orchestrator (NFVO), OpenStack VIM and OpenDaylight SDN controller with extended networking capabilities and names it as *Open³++*; where the objective is multi-domain orchestration. In *Open³++*, the authors are using a version of OSM which does not allow using an external WAN Infrastructure Manager (WIM)

to interconnect VNFs; therefore they resolved this issue via a proxy which abstracts all the underlying entities (multiple datacenters hosting their own VIMs, and one WIM to interconnect datacenters) as a single datacenter towards the OSM orchestrator. The previous two works differ from our work in which we utilize an extended version of OSM, which has an intrinsic capability to provision end-to-end network connectivity via an external WIM.

We describe several particular projects in the context of the 5GPPP Phase 1. The ORCHESTRA project [15] addresses monitoring solutions in backbone networks. The 5GEx project focuses on cross-domain orchestration of services over heterogeneous technologies and domains [16]. 5G-XHaul project targets dynamic and reconfigurable optical-wireless backhaul/fronthaul networks [17]. 5G-Crosshaul focuses on integrated backhaul and fronthaul transport network enabling flexible/software-defined reconfiguration of networking elements in a multi-tenant and service-oriented management environment [18]. Some other projects include iCirrus, which proposes an intelligent converged 5G network consolidating radio and optical access around user equipment [19], 5GinFIRE involves a 5G-Oriented experimental playground for vertical industries through an extensible 5G NFV-based reference ecosystem [20].

It is worthwhile mentioning initiatives from 5GPPP phase 2 as well. The BLUESPACE project focus on Spatial Division Multiplexing (SDM) in the RAN, using optical beamforming interfaces designed for wireless transmission, seamlessly integrated into optical access network infrastructures [21]. The PASSION project, targets high-capacity flexible photonic technologies [22]. The QAMELEON project proposes new photonic devices [23]. 5G-CLARITY project targets beyond 5G private networks integrating 5G, Wi-Fi and LiFi technologies, managed through AI-based autonomic networking [24]. 5G Transformer project which defines SDN/NFV-based mobile transport and computing platform, brings “slicing” into mobile transport networks [25]. 5G-COMLETE focuses on converged 5G infrastructure including wireless, optical and packet network in a disaggregated RAN approach [26].

To the best of our knowledge, Metro-Haul has been the only initiative to encompass from the top vertical services with two applications (video surveillance [27] and CLVS [8]) down to the underlying optical infrastructure, while covering compute and network resources with NFV and SDN principles. We mention that a similar result shown in this work related to component level setup times has been in shown in the video surveillance use-case demonstration [27]. The testbed in [27] utilizes an optical line system, whereas we utilize a disaggregated optical network, and show the configuration time at the ROADM level as well.

B. Network planning and orchestration over optical networks

Net2Plan is an open-source Java-based tool for planning, optimization, and evaluation of networks that cover compute and multi-layer network resources. Net2Plan was chosen as the Metro-Haul’s planning tool performing joint optimization of compute and network resources. Capacity planning and dimensioning studies leveraged Net2Plan in non-real-time simulation techno-economic analyses, which considered 5G requirements while exploiting control-plane programmability and spectrum management in the optical data plane [28]. On the other hand, Net2Plan participated in experimental demonstrations for the allocation of compute and network resources in real-time [29, 30]. Recently, a Net2Plan-based architecture for network optimization as a service (OaaS) showcased the possibility of reusing

algorithms for both the offline capacity planning and the online resource allocation points of view [31].

C. ONOS SDN Controller

The Open Network Operating System (ONOS) [32] is one of the leading open source SDN controllers, designed for carrier-grade solutions that leverage the economics of white box hardware, while offering the flexibility to create and deploy new dynamic network services with simplified Northbound API (NBI). It includes a base platform and a set of applications, on top of which the MetroHaul COM platform was developed. In this work, the main extensions involve the ability to setup optical channels across partially disaggregated (using a dedicated open line system controller) or fully disaggregated (using OpenROADM devices) systems. While ONOS is a standard SDN controller, the application to OpenROADM-based networks and the use of a Transport API (TAPI) NBI interface was lacking. However, in the recent years, this has been addressed. Recently in [33], a description of ONOS controller features is provided, in link with the establishment of the Open and Disaggregated Transport Network (ODTN) working group, specifically focused on the introduction of required functionality to control and monitor disaggregated transport networks. The paper [33] describes a set of experiments performed on a setup including both emulated and real optical devices controlled with ONOS. In [34], the authors validate a control plane architecture for multi-domain disaggregated transport networks that relies on the deployment of network elements compliant with the OpenROADM Multi Source Agreement device model [35].

4. TESTBED FOR CLVS USE CASE

In this section, we describe the testbed supporting the CLVS use case as presented in Fig. 2. This includes the data plane infrastructure and the COM system. We further describe the particularities of the VNFs used in the demo with their specific types of network interfaces. We assume that CLVS traffic from user handsets originates from an access network, where many users are capturing video streams via their handheld devices. It traverses the metro optical network and is then uploaded to an online CLVS platform.

A. Data plane infrastructure

The data plane for the testbed consists of the Metro-Haul data plane nodes which host both compute and network infrastructure. Moreover, the nodes are connected by an optical infrastructure described as follows.

A.1. Data plane nodes

Metro-Haul data plane nodes are SDN/NFV enabled with compute servers to host applications as VNFs and multi-layer packet and optical networking capabilities. They comprise i) *Access Metro Edge Node (AMEN)* connected to the metro and access network; and ii) *Metro Core Edge Node (MCEN)* connected to the metro and core network [36]. AMENs may host VNFs close to the access network for low latency to end user; whereas MCENs may host computationally intensive VNFs as it is equipped with considerable compute resources. AMEN and MCEN are connected together by a metro optical network. We demonstrate the CLVS use case where the AMEN and MCEN are hosted at High Performance Networks (HPN) Lab, University of Bristol (Bristol, UK) and BT Research Labs (Ipswich, UK), respectively.

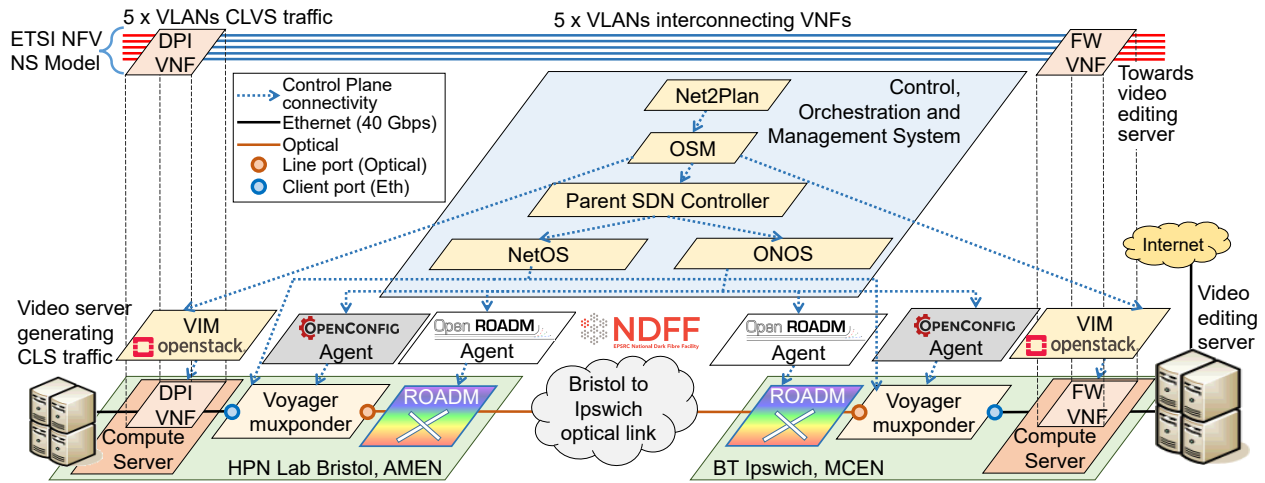


Fig. 2. Testbed for CLVS DPI application

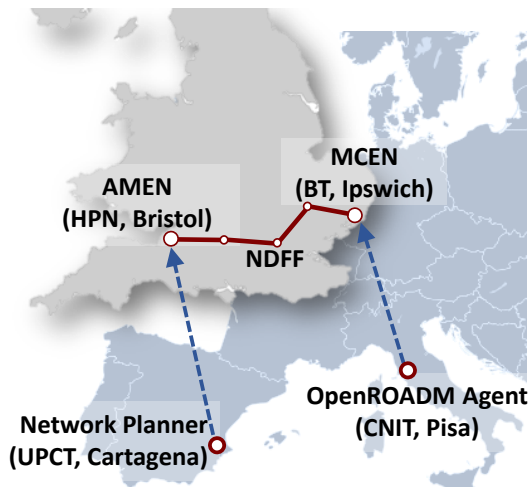


Fig. 3. Pan European control plane; and UK optical connectivity: Bristol to Ipswich¹

At each node, a compute server is hosted with 2×40 Gbps Ethernet interfaces; where one of the interfaces receives CLVS traffic from the access network at the AMEN, and one of the interface transmits the aggregated and edited CLVS traffic to the CLVS platform at the MCEN. The other 40 Gbps interface at each compute server is connected to a multi-layer Voyager muxponder [37]. The Voyager muxponder interfaces the packet domain at each node with the optical domain, and can switch VLAN-tagged Ethernet traffic on a client port towards the optical network on a DWDM coherent transponder (line port). The Voyager muxponder is further connected to a ROADM at each node for optical switching. For the purposes of demonstration, we utilize a simple broadcast-and-select architecture for the ROADMs. In particular, the one hosted at AMEN is a complete multi-degree ROADM including multiple Wavelength Selective Switches (WSSs) while the ROADM at MCEN is comprises of a single degree including one WSS. We adopt vendor neutral OpenConfig [38] and OpenROADM [35] based YANG data models and agents to control the Voyager muxponder and ROADM devices respectively. This enables disaggregated optical networks [33] by allowing equipment assembled by multiple

vendors however having the same consistent software to control them. We expand on the OpenConfig and OpenROADM agents in Sec. A.4.

A.2. Optical connectivity infrastructure

The aforementioned two sites are interconnected via two optical fibre infrastructures. The National Dark Fibre Facility (NDFF) connects Bristol to Cambridge, and a second fibre-link connects Cambridge to Ipswich with a total length of 1060.68 km (530.34 km unidirectional) of standard single mode fibre as shown in Fig. 3. DWDM filtering at Cambridge and Ipswich limits the ASE noise generated by the EDFAs in the links.

A.3. Crowdsourced live video streaming setup

A video server at AMEN generates video streams, emulating CLVS traffic using VLC media player instances [39] on $5 \times$ VLANs towards the AMEN; where it is received by a video editing server at MCEN after traversing the VNFs and the metro network. The video editing server hosts OBS studio software [40] to aggregate the received CLVS video streams as a single video and broadcasts to Twitch via an Internet uplink.

A.4. OpenROADM and OpenConfig agents

The optical devices have been modeled using OpenConfig and OpenROADM models, respectively, for Voyager muxponders and ROADMs. Thus, two different ONOS drivers have been implemented at the controller and two different SDN agents have been deployed in the data plane to operate the devices. The communication between the SDN controller and the SDN agents exploits the NETCONF protocol. The architecture of both SDN agents is decoupled in two main components: the *model processor* and the *device driver* used for actual configuration of the underlying hardware. The device drivers are coded as Linux dynamic libraries associated to the specific component to be configured (e.g. the circuit-packs i.e. the WSSs involved in the connections); they are loaded during the initialization of each specific device. The model processor communicates with the device driver using proprietary APIs. Such a software architecture is very flexible to manage a wide range of devices. For example, it is easy to implement node emulators, since it is possible to create *dummy* drivers performing no actions.

¹Europe image licensed under Creative Commons Attribution-Share Alike 3.0 Unported license, source: <https://w.wiki/34jR>

The OpenROADM agent has been implemented using the transAPI [41] framework within Netopeer [42], an open source implementation of the NETCONF protocol based on the libnetconf library. The transAPI framework allows invoking call-back functions whenever an edit-config RPC operation performs changes on a specific branch of the configuration. The agent is therefore composed by call-back functions managing controller requests for the creation of the interfaces that, according to the OpenROADM model, abstract the optical spectral window used during cross-connections and for the setup of the cross-connection itself.

The OpenConfig agent has been implemented using the ConfD framework. Thus the NETCONF protocol enables both the configuration and the monitoring of the main key transmission parameters of the muxponders (i.e. config: target-output-power, central-frequency and operational-mode; state: pre-FEC-BER, Q factor, ESNR, output-power, input-power). The agent presents an internal database (DB) organized according to the OpenConfig YANG modules to store those parameters. Two ad-hoc sockets (i.e. config socket and monitoring socket) with custom syntax have been designed in order to enable the communication with the device driver modules.

B. Control, Orchestration and Management (COM) system

The Metro-Haul control plane includes the COM system [7] to enable flexible service provisioning. It requires an NS described using ETSI standard Network Service Descriptor (NSD) template, as depicted in Fig. 4 (similar to the NS described in Fig. 1) consisting of VNF Descriptors (VNFDs) and Virtual Links (VLs) interconnecting VNFs [43]. VNFDs include the required resources (CPU, RAM, storage) per VNF. Additionally, there may be VLs to manage the VNFs by the NS administrator. The COM system provisions a slice of resources based on the NSD by mapping each VNF and VL into actual deployed Virtual Machines (VMs) and connectivity respectively over the metro network, using a hierarchy of services. Following sections describe the management network interconnecting the components, the services and the components providing them.

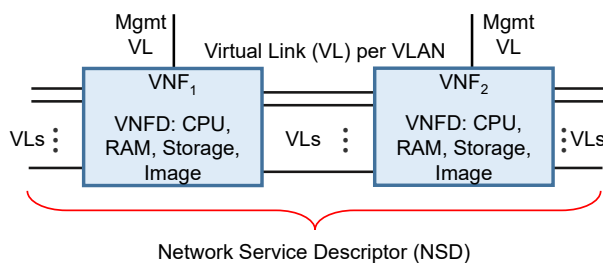


Fig. 4. ETSI-based Network Service Descriptor

B.1. Management Network

Most of the software components in the testbed as part of the COM system were hosted in United Kingdom; whereas some were hosted elsewhere in Europe: this includes the Net2Plan hosted in Universidad Politécnica de Cartagena (UPCT), Spain and one of the OpenROADM agents hosted in CNIT, Pisa, Italy, which controls the MCEN ROADM. To interconnect these various software components in the control plane as well as the data plane, a management network is used. This network is accessed by various partners in the Metro-Haul consortium via an OpenVPN instance hosted in Amazon Web Services. The

components hosted in Spain and Italy used this management network to connect to other components in UK as shown in Fig. 3. This network is not only utilized to access the NBI or Southbound API (SBI) of the control plane components, but is also used to access the VNFs by the administrator of the CLVS use-case once the NS is deployed.

B.2. Network planning service

Net2Plan [44] is in charge of providing the *network planning service*. As reviewed in Sec. 3, Net2Plan leverages Java-based extensions for interfacing with software components, with a variety of objectives encompassing OSM for NFV orchestration [45], the ONOS SDN controller [46] and multiple OpenStack [47] instances simultaneously [29]. Here, we leverage on those experimental validations and demonstrations for the allocation of compute and network resources, namely online implementations that have real-time actions over infrastructure. For providing the network planning service, Net2Plan interfaces with OSM, takes into account the service requirements, computes the allocation of the computing and multi-layer network resources, and generates as outcome the location of the VNFs. Further details on the working principle inside Net2Plan for providing the network planning service are described in Sec. 5.

B.3. End-to-end network service

OSM [48] is used as the reference MANO system, otherwise also known as an NFVO, which offers the *End-to-end network service* to deploy the CLVS NS. OSM is aligned with the ETSI-based NFV architecture [49] to deploy and manage NSs and the lifecycle of its constituent VNFs. It utilizes the ETSI NFV Information Model to model the NSs in terms of NSDs and VNFs; in addition to using an ETSI standardized SOL005 NBI. It deploys the NS using the NSD. This includes deployment of VNFs and Container Network Functions (CNFs) as VMs or containers respectively using underlying VIMs (for e.g. OpenStack or Kubernetes). To interconnect the VNFs/CNFs, OSM utilizes WIMs which may be a hierarchical system of SDN controllers. We use an extended version of OSM [50] to be able to query external WIMs to deploy network connectivity. For the CLVS use-case, we use an OpenStack VIM (to control the compute resources) [47] hosted at each site: AMEN and MCEN, and a Parent SDN Controller (PSC) as the WIM. On OSM, prior to the CLVS NS deployment, we specify the endpoints, which map the connectivity of each site with the WAN network: this mapping is used to deploy the network connectivity between the VNFs.

B.4. End-to-end connectivity service

PSC is a hierarchical SDN controller which is utilized as the WIM in the CLVS use case and provides the *End-to-end connectivity service*. It implements the IETF L2SM model as NBI [51], where OSM specifies the AMEN and MCEN network endpoints and VLAN ID while requesting a point-to-point connectivity. PSC deploys the connectivity between VNFs, composed individually of an optical connectivity service (provided by ONOS SDN controller) and packet connectivity service (provided by NetOS SDN controller).

B.5. Optical connectivity service

ONOS SDN controller provides the *optical connectivity service*. It involves being able to configure, upon request, a network media channel between the muxponders. For this purpose, we have implemented the TAPI v2.1 photonic media layer extensions as ONOS NBI, so the PSC can request a connectivity service (optical channel - OCh) between Voyager muxponder line ports.

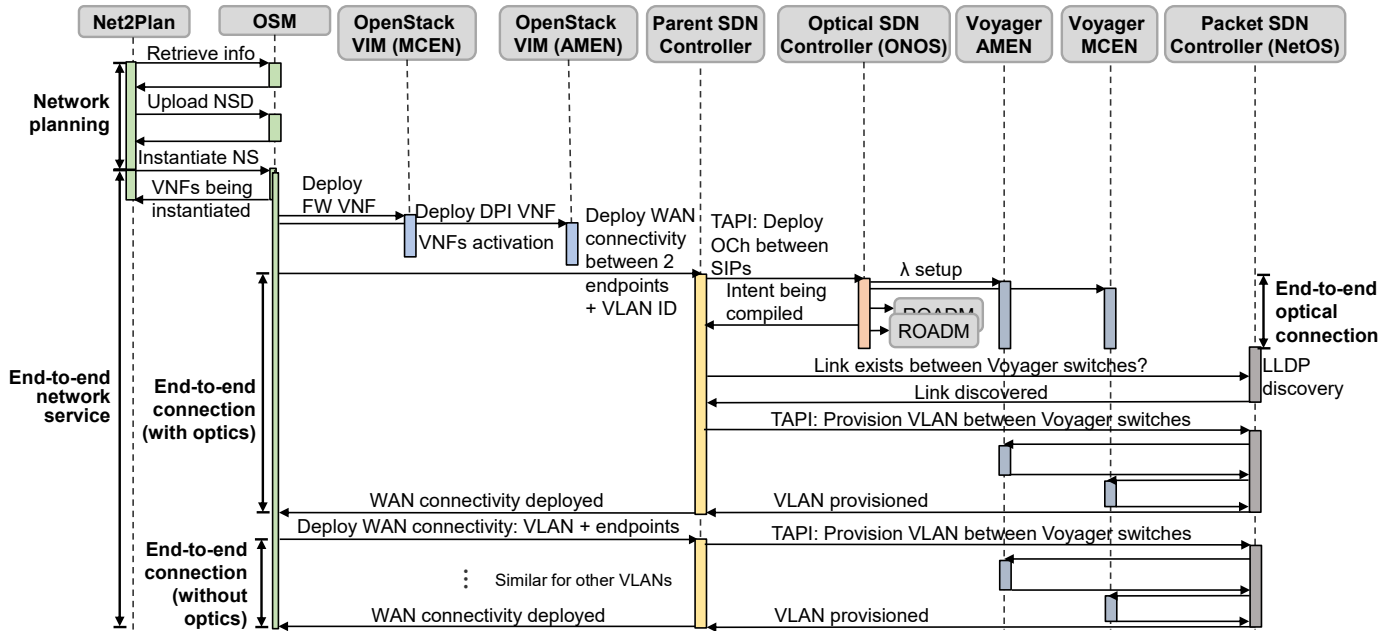


Fig. 5. Crowdsourced live video streaming NS deployment workflow

Moreover, we have also implemented device drivers (using NETCONF/YANG) in ONOS for both the Voyager muxponder (OpenConfig terminal device data model) and for the ROADMs (OpenROADM device data model) [33].

C. VNFs used in CLVS demo

For the demonstration of the CLVS use-case, we utilize proof-of-concept DPI and FW VNFs which are described as part of the overall CLVS NS in Sec. 2. Both VNFs use Ubuntu 18.04.5 LTS as the base operating system, where Open vSwitch (OVS) [52] software switches are used to switch traffic between ingress and egress ports. The DPI VNF consists of an open source library called ntopng [53], which performs protocol classification by analyzing the packet header as well as payload information. The FW VNF utilizes multiple OVS instances between the ingress and egress ports to allow or block traffic based on the port numbers, since the compute server forwards untagged traffic to the FW VNF, which is then identified based on the ingress/egress port. For this purpose, OpenFlow rules are specified on the OVS instances to enforce policy (allow or block) on the traffic. Within the VNFD, we specify 6 vCPUs, 8 GB RAM, and 60 GB storage as the resources for both VNFs.

D. VNF interfaces: VirtIO and SR-IOV

VNFs on a compute server can be connected to L2 networks in multiple ways: the most notable ways include attaching standard VirtIO or Single Root Input/Output Virtualization (SR-IOV) interfaces to the VMs. VirtIO is a standardized software-based interface which allows a VMs to access physical network interfaces on the compute servers [54]. Here, the VM's interface, which is based on VirtIO, is connected to the physical network interface via a standard Linux bridge while using the Linux Bridge mechanism driver, as part of the OpenStack Neutron networking drivers. This Linux bridge has limited performance, as it requires the packets to be processed by the Linux kernel, as will be shown in the results. Whereas an SR-IOV-based interface behaves differently. In this scenario, an SR-IOV

enabled physical network interface can be virtualized as multiple Virtual Functions (VFs). Each VF can be directly assigned to a VM as an interface, bypassing the bottleneck of the Linux kernel; consequently leading to considerable throughput and latency savings. We deploy the VNFs in the CLVS NS with either VirtIO or SR-IOV interfaces by specifying the interface type in the VNFD and compare the network level results in Sec. 6. For the management VL, we use a VirtIO-based interface as it does not require a high performance network.

5. NETWORK SERVICE DEPLOYMENT WORKFLOW

We present the CLVS NS deployment workflow (depicted in Fig. 5) using the Metro-Haul COM services. First, a CLVS application administrator requests a CLVS NS from Net2Plan [55], which provides the *network planning service*. An ad-hoc extension of Net2Plan has been developed for this use case. In particular, the extension has additional fields for indicating the requirements of the CLVS application and also includes the client for interfacing with OSM. Within this particular extension, Net2Plan retrieves from OSM information in the form of a list of NSs. Net2Plan leverages on a set of algorithms for computing resource allocation algorithms that consider both compute and network resources, being the latter multi-layer, i.e. both packet and optical. Based on the requirements introduced by the CLVS application administrator in terms of compute resources, network latency and bandwidth, Net2Plan executes a resource allocation algorithm based on a NFV-based shortest path approach, producing NSD and VNF locations as output. The output of Net2Plan is the NSD with the location of the VNFs, i.e. DPI VNF at the AMEN and FW VNF at the MCEN. Net2Plan uses this output to deploy the network slice by instantiating the CLVS NS using an NFVO while specifying VNF locations.

The OSM NFVO, which offers the *end-to-end network service*, deploys the VNFs specified in the NSD as VMs on the compute servers at both locations using OpenStack VIM. Once the VNFs are active, OSM requests connectivity between VNFs using the PSC as the WIM, while sharing the VLAN ID used for the net-

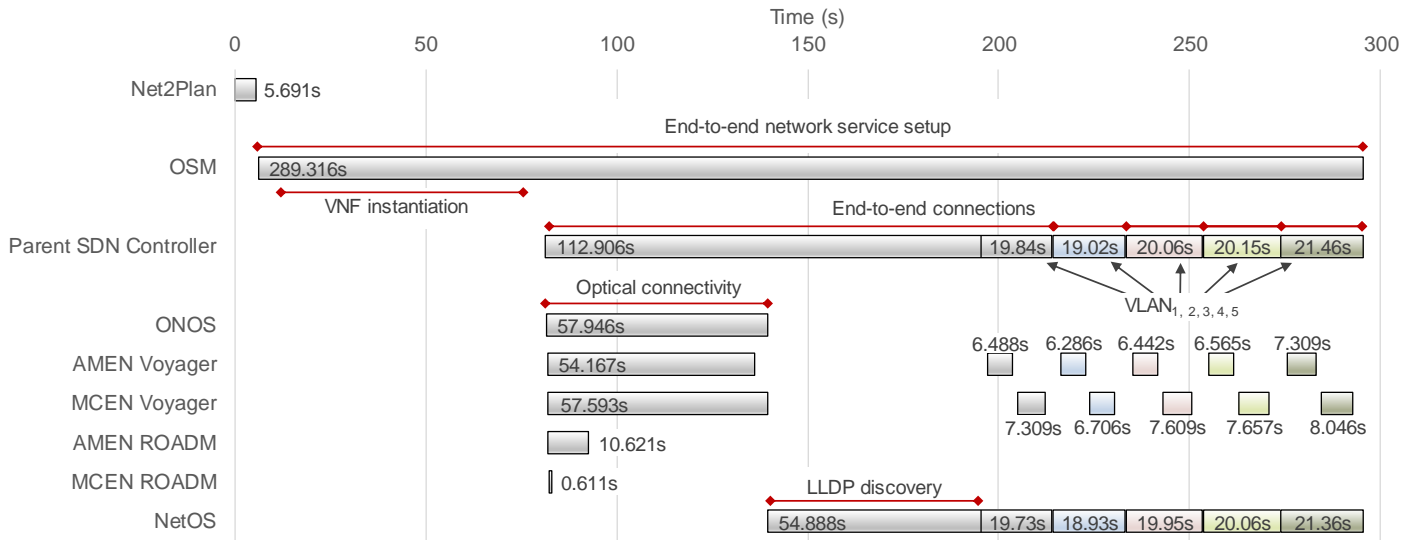


Fig. 6. Detailed setup times per Metro-Haul components. The times reflect one particular experiment run.

Table 1. Metro-Haul service setup times. Averages and 95% confidence interval shown.

Service	Setup time (s)
Network planning	6.059 ± 0.283
End-to-end network service	292.679 ± 6.056
End-to-end connection (with optics)	131.549 ± 5.241
End-to-end optical connection	56.149 ± 0.649
End-to-end connection (no optics)	20.129 ± 0.710

Table 2. End-to-end connection: component time breakdown. Averages and 95% confidence interval shown.

Metro-Haul component	Optics connection (s)	VLANs connection (s)
AMEN Voyager	54.747 ± 0.253	6.597 ± 0.273
MCEN Voyager	54.488 ± 0.912	6.908 ± 0.562
AMEN ROADM	10.6 ± 0.001	n/a
MCEN ROADM	0.870 ± 0.63	n/a

works connected to the VNFs locally at each site (AMEN and MCEN). Moreover, OSM shares the endpoints of each site with the PSC as part of the request.

The endpoints sent by OSM can then be used by the PSC, along with the VLAN ID, to identify the two sites between which the end-to-end connectivity is to be deployed. First, PSC identifies if there is an optical connectivity established between the two sites; otherwise it proceeds to request ONOS to deploy an optical connection between the Voyager line ports to interconnect the AMEN and MCEN. ONOS then proceeds to configure the Voyager line ports as well as the ROADMs at each site to provide optical connectivity.

Once the optical connectivity is deployed, the PSC deploys the L2 VLAN-based packet connectivity between the Ethernet client ports of the Voyager muxponder using NetOS packet SDN controller [56]. NetOS exposes the topology via TAPI context to the PSC, where the optical link between Voyager muxponders is discovered using LLDP. To interconnect the Voyager muxponder client ports, NetOS configures the VLAN on both the client and line ports, which are internally connected to the same bridge on the Voyager muxponder. As a VLAN is deployed, the end-to-end connectivity is established.

Since the NS is composed of multiple VLs (which map to multiple VLANs), OSM proceeds to deploy each VLAN using the PSC. However in each further VLAN deployment, the optical connectivity is already deployed between the sites and PSC only uses NetOS to deploy the packet connectivity for the new VLAN. Once all the VLANs between the VNFs are deployed, the CLVS NS is fully deployed.

6. RESULTS

This section covers the performance results of the Metro-Haul COM system services and data plane. We present the the whole system's setup times, including the planning phase, VNF instantiation, and optical and VLAN connectivity. We also discuss the performance comparison between VirtIO and SR-IOV VNF interfaces. Each experiment was repeated 30 times, and we present the results in the form of averages and 95% confidence intervals. Finally, we show the optical link performance between the two sites.

A. Setup times

Metro-Haul's most relevant Key Performance Indicators (KPIs) to this work relate to the various service establishment times. For this particular experiment, the deployment of the network service consists of several steps, as described in Section 5, Fig. 5.

Table 1 summarises the results from 30 experiments, showing averages with 95% confidence intervals for an NS with 5 VLANs. The data show that the whole network service is deployed within ≈ 5 minutes (even when considering the network planning phase). From those 5 minutes, ≈ 6 s are taken by the network planner, whereas remainder ≈ 292 s are taken by the network service establishment per se. These ≈ 292 s can be further broken down into ≈ 80 s consumed by the VMs instantiation and OSM internal processing (not shown in Table 1) and ≈ 212 s remaining to establish the end-to-end connections. The end-to-end connections includes the VLAN configuration and can include optical connections or not. An end-to-end connection with optics takes ≈ 131 s (including ≈ 56 s to configure the

optical connection, and ≈ 75 s NetOS processing and discovering time), whereas each end-to-end connection without optics takes ≈ 20 s.

To better illustrate the aforementioned results, Fig. 6 shows, in detail, the times of one of the experiment runs. Additionally, Table 2 presents the times taken by the equipment utilised in this work. As expected, the Voyager switches take much more time configuring its optics than the VLANs, and the ROADMs are not active during VLAN configuration since the optical connectivity is already deployed. Moreover, the difference of configuration times between AMEN and MCEN ROADMs is due to its implementation, as discussed in Section 4.

At this point, it is important to expand on the scalability of Metro-Haul architecture. The network service setup time is dependent on the number of VNFs as well as the links interconnecting the VNFs which are deployed as VLANs in the network. Since the configuration of the optical equipment, i.e. ROADMs and muxponders, is performed in parallel, the addition of extra devices does not impact the obtained results in any meaningful way. On the other hand, since the VLANs are configured in series, the total number of VLANs do increase the overall establishment time in a linear fashion. As Table 1 shows, each extra VLAN allocated on top of an already existing optical connection adds ≈ 20 s.

When comparing these results against the 5G PPP KPI service creation time of 90 minutes [4], it is clear that the presented solution can meet the specification if the number of VLANs needed by a given service does not exceed ≈ 270 VLANs (over one single optical channel).

B. Performance comparison VirtIO vs SR-IOV interfaces

During the deployment of the NS, it is possible to specify both VNFs to run with either VirtIO or SR-IOV interfaces within the VNFD at the OSM level. In both cases, we check the round-trip latency, from the video server at AMEN to the video editing server at MCEN by using ICMP ping utility. Moreover, we compare the throughput performance of both interfaces, on a single VLAN, by running multiple parallel² iPerf3 [57] clients and servers on the video server and the video editing server respectively. We show the performance in terms of averages and 95% confidence intervals in Table 3, for the cases where the sites are connected back-to-back (we replicate the setup of MCEN at Bristol) and while using the optical connectivity infrastructure.

The result shows that the SR-IOV-based interface has 1.008 ms lower latency than a VirtIO-based interface in the case of using the optical fiber infrastructure; this is considerable saving if a 5G Ultra-Reliable Low Latency Communication (URLLC) use-case is considered with stringent latency requirements.

We observe that the SR-IOV based interfaces give a throughput of 35.13 Gbps, which is almost 3.67 times faster than the throughput provided using VirtIO interfaces (9.58 Gbps). This is considerable since the industry is moving towards NFV-based 5G network service deployments, and VNF performance needs to be at par with middleboxes and other hardware-based solutions. Furthermore, with technology advancements where we may see CLVS traffic comprised of 4K videos with high frames per second having very high data rates in the near future, such a throughput is highly anticipated.

It is important to mention that the throughput performance does not reach the maximum of 40 Gbps, which is the bandwidth

on all the compute server interfaces in the deployed testbed. This is due to the reason that further optimization such as DPDK-based VNFs need be to used which optimize the packet processing within the VMs and are out of scope of this work. The reader may explore [58, 59] for further explanation.

Table 3. Performance comparison: VirtIO vs SR-IOV interfaces. Averages and 95% confidence interval shown.

Link	Interface type	Latency (ms)	Throughput (Gbps)
back-to-back (0 km)	VirtIO	0.857 ± 0.005	9.58 ± 0.13
	SR-IOV	0.409 ± 0.003	35.13 ± 0.42
AMEN-MCEN (1060 km)	VirtIO	6.939 ± 0.080	9.58 ± 0.13
	SR-IOV	5.931 ± 0.071	35.13 ± 0.42

C. Optical connectivity

The optical connectivity was implemented by the Voyager muxponders, with the channels centred at 1546.9 nm wavelength (193.8 THz), using PM-QPSK modulated signals resulting in a 100 Gbps data rate (Fig. 7). The measured pre-FEC BER was 1.33×10^{-3} for Ipswich-Bristol link and 4.257×10^{-5} for Bristol-Ipswich link. The FEC in the Voyager muxponders was set at 25%, sufficient to correct errors during the transmission in both links.

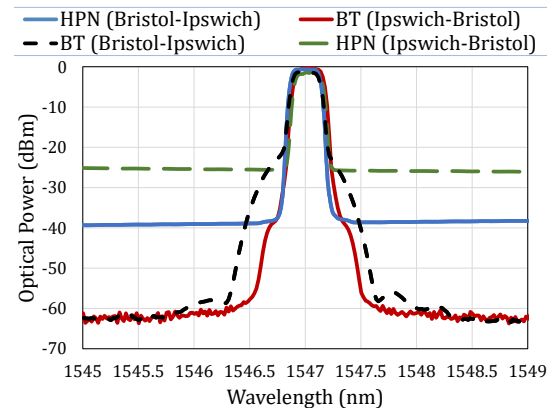


Fig. 7. Optical spectrum Bristol-Ipswich link

7. CONCLUSIONS

In this paper, we presented an SDN and NFV-enabled solution provided by Metro-Haul project, which consists of a multi-layer slice deployed via NFV orchestration over packet and disaggregated metro optical networks. We applied this solution to a well known CLVS use-case which has stringent requirements of high-bandwidth and low latency, relevant to 5G networks. We began by stressing the fact that the benefits of Metro-Haul infrastructure can be directly leveraged by the CLVS use-case. We reported some relevant initiatives and concluded that Metro-Haul is the only project covering all aspects of a multi-layer slice deployment, spanning from network planning down to controlling disaggregated optical network components.

To validate the solution, we demonstrated the CLVS use-case on a test-bed with data plane components hosted in the UK,

²Multiple parallel iPerf3 sessions allows each iPerf client/server instance to run on different CPU threads, consequently avoiding processing bottlenecks.

where the control plane components were deployed throughout Europe. The testbed involved two SDN and NFV enabled data plane nodes, with multi-layer packet and optical networking and compute capabilities, interconnected via an optical network. The data plane is accompanied by a COM system which provides a hierarchy of network service as well as connectivity services, to enable the CLVS use-case. The proof-of-concept VNFs used in the demonstration utilized both VirtIO and SRIOV interfaces.

We show the NS deployment workflow and the messages exchanged between various components to deploy the COM system services. This is followed by comprehensive results, showcasing the setup times of the services as well as the components. Using a reference CLVS NS with 5 VLANs interconnecting its constituent VNFs, we show that the network service setup time meets the 5G PPP KPI service creation time; thanks to the agility and flexible service provisioning enabled by SDN and NFV technologies. Moreover, we compare the performance of VirtIO and SR-IOV based VNF interfaces; where we deployed an end-to-end NS with VNFs having SR-IOV-based interfaces via an NFVO. The results concluded that the SR-IOV based interfaces provide significantly higher throughput and lower latency as compared to VirtIO interfaces. This is very relevant to the increasing softwarization of telecommunication networks, where physical network functions in the form of middleboxes are being converted to VNFs being hosted in commodity compute hardware.

ACKNOWLEDGMENTS

This work is supported by EU METRO-HAUL project (Grant Agreement 761727) and UK EPSRC NDFP project (EP/S028854/1).

REFERENCES

- Cisco, "Cisco Annual Internet Report (2018–2023), White Paper," IEEE Access (2019). <https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.pdf>, Accessed: 2021-03-01.
- Sandvine, "The Global Internet Phenomena Report COVID-19 Spotlight," (2020). <https://www.sandvine.com/phenomena>, Accessed: 2021-03-01.
- A. Feldmann, O. Gasser, F. Lichtblau, E. Pujol, I. Poese, C. Dietzel, D. Wagner, M. Wichtlhuber, J. Tapiador, N. Vallina-Rodriguez *et al.*, "The Lockdown Effect: Implications of the COVID-19 Pandemic on Internet Traffic," in *Proceedings of the ACM Internet Measurement Conference*, (2020), pp. 1–18.
- 5G Initiative Technology Board, "5G PPP Phase-II Projects Performance KPIs," Tech. rep., 5G-PPP (2019). <https://bscw.5g-ppp.eu/pub/bscw.cgi/312793>, Accessed: 2021-03-01.
- F. Chen, C. Zhang, F. Wang, J. Liu, X. Wang, and Y. Liu, "Cloud-Assisted Live Streaming for Crowdsourced Multimedia Content," *IEEE Transactions on Multimed.* **17**, 1471–1483 (2015).
- Metro-Haul 5G Project. <http://metro-haul.eu>, Accessed: 2021-03-01.
- L. Velasco, R. Casellas, S. Llana, L. Gifre, R. Martínez, R. Vilalta, R. Muñoz, and M. Ruiz, "A control and management architecture supporting autonomic NFV services," *Photonic Netw. Commun.* **37**, 24–37 (2019).
- A. S. Muqaddas, R. S. Tessinari, O. G. de Dios, E. Hugues-Salas, R. Casellas, L. Luque, M. Channegowda, A. Giorgetti, A. Sgambelluri, F. Cugini, F. J. Moreno-Muro, M. Garrich, P. Pavón-Mariño, R. Morro, K. Farrow, A. Lord, R. Nejabati, and D. Simeonidou, "Field Trial of Multi-Layer Slicing Over Disaggregated Optical Networks Enabling End-to-End Crowdsourced Video Streaming," in *2020 European Conference on Optical Communications (ECOC)*, (2020), pp. 1–4.
- TwitchTracker, "Statistics & Charts," <https://twitchtracker.com/statistics>, Accessed: 2021-03-01.
- Twitch. www.twitch.com, Accessed: 2021-03-01.
- 5G-PPP Software Network Working Group, "Cloud-Native and Verticals' services - Whitepaper," (2019). <https://5g-ppp.eu/wp-content/uploads/2020/02/5G-PPP-SN-WG-5G-and-Cloud-Native.pdf>, Accessed: 2021-03-01.
- N. Makris, C. Zarafetas, A. Valantasis, and T. Korakis, "Service Orchestration over Wireless Network Slices: Testbed Setup and Integration," *IEEE Transactions on Netw. Serv. Manag.* pp. 1–1 (2020).
- B. Nogales, I. Vidal, D. R. Lopez, J. Rodriguez, J. Garcia-Reinoso, and A. Azcorra, "Design and Deployment of an Open Management and Orchestration Platform for Multi-Site NFV Experimentation," *IEEE Commun. Mag.* **57**, 20–27 (2019).
- P. Karamichailidis, K. Choumas, and T. Korakis, "Enabling Multi-Domain Orchestration using Open Source MANO, OpenStack and OpenDaylight," in *2019 IEEE International Symposium on Local and Metropolitan Area Networks (LANMAN)*, (2019), pp. 1–6.
- EU ORCHESTRA project (H2020). <http://www.orchestraproject.eu>, Accessed: 2021-03-01.
- 5GPPP 5GEx project (H2020). <https://5g-ppp.eu/5gex>, Accessed: 2021-03-01.
- 5GPPP 5G-XHaul project (H2020). <https://5g-ppp.eu/5g-xhaul>, Accessed: 2021-03-01.
- 5GPPP 5G-Crosshaul project (H2020). <https://5g-ppp.eu/xhaul>, Accessed: 2021-03-01.
- iCirrus project (H2020). <https://www.icirrus-5gnet.eu>, Accessed: 2021-03-01.
- 5GinFIRE project (H2020). <https://5ginfire.eu>, Accessed: 2021-03-01.
- BLUESPACE project (H2020). <https://5g-ppp.eu/bluespace>, Accessed: 2021-03-01.
- PASSION project (H2020). <https://www.passion-project.eu>, Accessed: 2021-03-01.
- GAMELEON project (H2020). <https://ict-gameleon.eu>, Accessed: 2021-03-01.
- 5G-CLARITY project (H2020). <https://www.5gclarity.com>, Accessed: 2021-03-01.
- 5G Transformer project (H2020). <https://5g-ppp.eu/5g-transformer>, Accessed: 2021-03-01.
- 5G-Complete project (H2020). <https://5gcomplete.eu>, Accessed: 2021-03-01.
- B. Shariati, J. J. Perdeno-Manresa, A. Dochhan, A. S. Muqaddas, R. Casellas, O. G. de Dios, L. L. Canto, B. Lent, J. E. Lopez de Vergara, S. Lopez-Buedo, F. J. Moreno, P. Pavón, L. Velasco, S. Patri, A. Giorgetti, F. Cugini, A. Sgambelluri, R. Nejabati, D. Simeonidou, R.-P. Braun, A. Autenrieth, J.-P. Elbers, and R. Freund, "A Latency-Aware Real-Time Video Surveillance Demo: Network Slicing for Improving Public Safety," in *2021 Optical Fiber Communications Conference and Exhibition (OFC)*, (2021), pp. 1–3.
- P. Pavon-Marino, F. Moreno-Muro, M. Garrich, M. Quagliotti, E. Riccardi, A. Rafel, and A. Lord, "Techno-economic impact of filterless data plane and agile control plane in the 5g optical metro," *J. Light. Technol.* **38**, 3801–3814 (2020).
- M. Garrich, M. Hernández-Bastida, C. San-Nicolás-Martínez, F. J. Moreno-Muro, and P. Pavon-Marino, "The net2plan-openstack project: It resource manager for metropolitan sdn/nfv ecosystems," in *2019 Optical Fiber Communications Conference and Exhibition (OFC)*, (2019), pp. 1–3.
- F. J. Moreno-Muro, M. Garrich, C. San-Nicolás-Martínez, M. Hernández-Bastida, P. Pavón-Mariño, A. Bravalheri, A. S. Muqaddas, N. Uniyal, R. Nejabati, D. Simeonidou, R. Casellas, and O. G. de Dios, "Joint vnf and multi-layer resource allocation with an open-source optimization-as-a-service integration," in *45th European Conference on Optical Communication (ECOC 2019)*, (2019), pp. 1–4.
- M. Garrich, J. Romero-Gázquez, F. Moreno-Muro, M. Hernández-Bastida, and P. P. Mariño, "Joint optimization of it, ip and wdm layers: From theory to practice," in *2020 International Conference on Optical Network Design and Modeling (ONDM)*, (2020), pp. 1–4.
- ONOS. <https://opennetworking.org/onos>, Accessed: 2021-03-01.
- A. Giorgetti, A. Sgambelluri, R. Casellas, R. Morro, A. Campanella, and

- P. Castoldi, "Control of open and disaggregated transport networks using the Open Network Operating System (ONOS) [Invited]," *IEEE/OSA J. Opt. Commun. Netw.* **12**, A171–A181 (2020).
34. R. Casellas, R. Martínez, R. Vilalta, and R. Muñoz, "Abstraction and Control of Multi-Domain Disaggregated Optical Networks with OpenROADM Device Models," *IEEE J. Light. Technol.* (2020).
 35. OpenROADM MSA. www.openroadm.net, Accessed: 2021-03-01.
 36. A. Dochhan, J. K. Fischer, B. Lent, A. Autenrieth, B. Shariati, P. W. Berenguer, and J. Elbers, "Metro-Haul Project Vertical Service Demo: Video Surveillance Real-Time Low-Latency Object Tracking," in *2020 Optical Fiber Communications Conference and Exhibition (OFC)*, (2020), pp. 1–3.
 37. A. S. Muqaddas, A. Giorgetti, R. S. Tessinari, T. Diallo, A. Sgambelluri, R. Nejabati, and D. Simeonidou, "OpenConfig-Extension for VLAN-Based End-to-End Network Slicing over Optical Networks," in *2020 Optical Fiber Communications Conference and Exhibition (OFC)*, (2020), pp. 1–3.
 38. OpenConfig. www.openconfig.net, Accessed: 2021-03-01.
 39. VideoLAN, "VLC Media Player," www.videolan.org/vlc, Accessed: 2021-03-01.
 40. Open Broadcaster Software® | OBS. www.obsproject.com, Accessed: 2021-03-01.
 41. Libnetconf: Transaction API (transAPI). <https://rawgit.com/CESNET/libnetconf/master/doc/doxygen/html/d9/d25/transapi.html>, Accessed: 2021-03-01.
 42. CESNET – Netopeer, NETCONF Toolset. <https://github.com/CESNET/netopeer>, Accessed: 2021-03-01.
 43. ETSI, "GS NFV-MAN 001 V1.1.1 Network Function Virtualisation (NFV); Management and Orchestration," (2014).
 44. Net2Plan. www.net2plan.com, Accessed: 2021-03-01.
 45. F. J. Moreno-Muro, C. San-Nicolas-Martinez, E. Martin-Seoane, M. Garrich, P. Pavon-Marino, O. G. de Dios, and V. López, "Joint optimal service chain allocation, vnf instantiation and metro network resource management demonstration," in *2018 Optical Fiber Communications Conference and Exposition (OFC)*, (2018), pp. 1–3.
 46. F. J. Moreno-Muro, C. San-Nicolás-Martínez, M. Garrich, P. Pavon-Marino, O. G. De Dios, and R. L. Da Silva, "Latency-aware optimization of service chain allocation with joint vnf instantiation and sdn metro network control," in *2018 European Conference on Optical Communication (ECOC)*, (2018), pp. 1–3.
 47. OpenStack - Open Source Cloud Computing Infrastructure. <https://www.openstack.org>, Accessed: 2021-03-01.
 48. Open Source MANO.
 49. ETSI, "GS NFV 006 V2.1.1 Network Functions Virtualisation (NFV) Release 2; Management and Orchestration; Architectural Framework Specification," (2021).
 50. A. Bravalheri, A. S. Muqaddas, N. Uniyal, R. Casellas, R. Nejabati, and D. Simeonidou, "VNF Chaining Across Multi-PoPs in OSM using Transport API," in *2019 Optical Fiber Communications Conference and Exhibition (OFC)*, (2019), pp. 1–3.
 51. B. Wen, G. Fioccola, C. Xie, and L. Jalil, "A YANG Data Model for Layer 2 Virtual Private Network (L2VPN) Service Delivery," RFC 8466 (2018). <https://rfc-editor.org/rfc/rfc8466.txt>, Accessed: 2021-03-01.
 52. Open vSwitch. <https://www.openvswitch.org>, Accessed: 2021-03-01.
 53. ntopng – High-Speed Web-based Traffic Analysis and Flow Collection. <https://github.com/ntop/ntopng>, Accessed: 2021-03-01.
 54. Mellanox – SR-IOV vs. VirtIO; Comparison of two OVS/vRouter Acceleration Techniques. <https://www.mellanox.com/download/white-paper/comparison-of-ovs-vrouter-acceleration-techniques-sri-iov-and-virtio>, Accessed: 2021-03-01.
 55. M. Garrich, J. Romero-Gázquez, F. Moreno-Muro, M. Hernández-Bastida, M. B. Delgado, A. Bravalheri, N. Uniyal, A. S. Muqaddas, R. Nejabati, R. Casellas, O. G. de Dios, and P. P. Mariño, "IT and Multi-layer Online Resource Allocation and Offline Planning in Metropolitan Networks," *J. Light. Technol.* **38**, 3190–3199 (2020).
 56. Zeetta Networks, "High Level Overview of NetOS architecture," <https://zeetta.com/technology/netos>, Accessed: 2021-03-01.
 57. iPerf3 - The TCP, UDP and SCTP network bandwidth measurement tool. <https://iperf.fr>, Accessed: 2021-03-01.
 58. M. Kourtis, G. Xilouris, V. Riccobene, M. J. McGrath, G. Petralia, H. Koumaras, G. Gardikis, and F. Liberal, "Enhancing VNF performance by exploiting SR-IOV and DPDK packet processing acceleration," in *2015 IEEE Conference on Network Function Virtualization and Software Defined Network (NFV-SDN)*, (2015), pp. 74–78.
 59. T. Zhang, H. Qiu, L. Linguaglossa, W. Cerroni, and P. Giaccone, "NFV Platforms: Taxonomy, Design Choices and Future Challenges," *IEEE Transactions on Netw. Serv. Manag.* pp. 1–1 (2020).