

Orchestration of Network Services Across Multiple Operators: The 5G Exchange Prototype

A. Sgambelluri¹, F. Tusa², M. Gharbaoui¹, E. Maini², L. Toka³, J. M. Perez⁴, F. Paolucci¹, B. Martini⁵, W. Y. Poe⁶, J. Melian Hernandez⁷, A. Muhammed⁸, A. Ramos⁷, O. G. de Dios⁹, B. Sonkoly³, P. Monti⁸, I. Vaishnavi⁶, C. J. Bernardos⁴, R. Szabo¹⁰

Abstract— Future 5G networks will rely on the coordinated allocation of compute, storage, and networking resources in order to meet the functional requirements of 5G services as well as guaranteeing efficient usage of the network infrastructure. However, the 5G service provisioning paradigm will also require a unified infrastructure service market that integrates multiple operators and technologies. The 5G Exchange (5GEx) project, building heavily on the Software-Defined Network (SDN) and the Network Function Virtualization (NFV) functionalities, tries to overcome this market and technology fragmentation by designing, implementing, and testing a multi-domain orchestrator (Mdo) prototype for fast and automated Network Service (NS) provisioning over multiple-technologies and spanning across multiple operators. This paper presents a first implementation of the 5GEx Mdo prototype obtained by extending existing open source software tools at the disposal of the 5GEx partners. The main functions of the 5GEx Mdo prototype are showcased by demonstrating how it is possible to create and deploy NSs in the context of a Slice as a Service (SaaS) use-case, based on a multi-operator scenario. The 5GEx Mdo prototype performance is experimentally evaluated running validation tests within the 5GEx sandbox. The overall time required for the NS deployment has been evaluated considering NSs deployed across two operators.

Keywords— 5G, Multi Domain Orchestrator, Multi Operator, Slice as a Service (SaaS), Network Service Orchestration.

I. INTRODUCTION

Market and technology fragmentation results in a multitude of telecommunication networks and operators each focused on different countries, regions, and technologies. This makes it difficult and costly to create an infrastructure for a flexible service provisioning that should not be hindered by technology/vendor specific solutions and administrative domain boundaries.

One of the key concepts of 5G networks is the coordinated allocation of IT (e.g., compute, storage) and networking

resources for service provisioning. In other words, a 5G service will require a unified infrastructure service market, integrating multiple operators and technologies in order to guarantee that, on one hand, service performance and functional requirements are met, and on the other, that the network infrastructure is used efficiently. In this context, Software-Defined Network (SDN) and Network Function Virtualization (NFV) represent very promising tools able to provide high programmability and great elasticity while deploying network services (NSs) [1].

The European 5G Exchange (5GEx) project [2] tries to overcome the aforementioned market fragmentation providing an innovative solution to the challenge of orchestrating services across multi-domain, multi-technology, 5G networks. Using the functionalities offered by SDN and NFV, we have designed, implemented, and tested a multi-domain orchestrator (Mdo) prototype, which represents the first step for a fast and automated NS provisioning over multiple-technologies spanning across multiple-operators. The three main functionalities of the 5GEx Mdo prototype are the following: (i) ability to discover the 5GEx Mdo of the other operators and related capabilities/resources, (ii) ability to receive and deploy NS or Resource requests from other 5GEx Mdos, (iii) ability to monitor the performance of the deployed requests and report the performance to the requesting 5GEx Mdo.

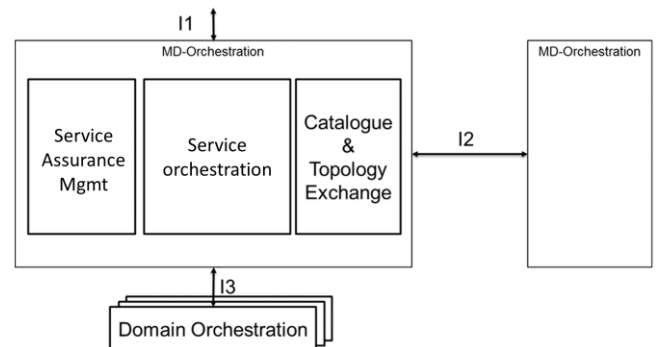


Figure 1 Functional architecture of the 5GEx Mdo

Based on a multi-operator scenario, the paper describes how it is possible to create and deploy a NS in the context of a Slice as a Service (SaaS) use-case.

II. RELATED WORKS

This section provides an overview of the state of the art solutions and activities relevant to this work (i.e., the 5G Exchange prototype). This information has been considered during the 5G Exchange prototype design and implementation

¹ Scuola Superiore Sant'Anna, Pisa, Italy, emails {a.sgambelluri, m.gharbaoui, fr.paolucci}@santannapisa.it

² University College London, UK, emails: {francesco.tusa, e.maini}@ucl.ac.uk

³ Budapest Univ. of Technology and Economics, emails: {sonkoly, toka}@tmit.bme.hu

⁴ Universidad Carlos III de Madrid, emails: {jmbc@it.uc3m.es, j.mperez@alumnos.uc3m.es}

⁵ CNIT, Pisa, Italy email: barbara.martini@cnit.it

⁶ Huawei Technologies Duesseldorf GmbH, Munich, Germany, emails: {wint.yi.poe, ishan.vaishnavi}@huawei.com

⁷ ATOS Spain SA, emails: {javier.melian, aurora.ramos}@atos.net

⁸ KTH Royal Institute of Technology, Stockholm, Sweden, email: {ajalmu, pmonti}@kth.se

⁹ TID Telefonica Investigaci3n y Desarrollo (TID), Madrid, Spain email: oscar.gonzalezdedios@telefonica.com

¹⁰ Ericsson Research, Budapest, Hungary, email: robert.szabo@ericsson.com

process. It is interesting to notice that none of the listed activities and solutions provide at this stage the capabilities for the full orchestration of NS over multiple administrative domains (multi-operator). First, we will look at the results for standardization bodies regarding multi-domain orchestration in software controlled networks are presented. Then, we will analyze the main single-domain and multi-domain orchestration frameworks currently available.

With respect to standardization bodies, the ETSI MANO group [3] defines the network functions virtualization orchestrator (NFVO) as the entity responsible for resource and service orchestration and life cycle management across virtualized infrastructure managers (VIMs). Currently, the open platform of the NFV group is looking into the realization of an ETSI NFV architectural framework by integrating and extending existing open source projects, such as OpenStack, OpenDaylight, KVM and Open vSwitch, among others. The ITU-T study group 13 (SG13) [4] is standardizing Future Networks (FNs) with the objectives of service (network virtualization), data, environmental (energy saving for FNs), and socio-economic awareness. In 2015, they published a recommendation on resource management in (multi) cloud computing. The recommendation presents general concepts of end-to-end resource management for cloud services and a vision for adopting cloud resource management in a telecommunication-rich environment [5]. The IEEE next-generation service overlay network (NGSON) [6] is a reference framework for the control and delivery of composite services over diverse IP-based networks with QoS support. NGSON supports the creation and deployment of composite services across different SPs and over heterogeneous network domains. It also promotes the provisioning of services to end users with guaranteed quality of experience. In particular, it allows the adaptation of composite services to the dynamic context of users, devices, services, and networks while optimizing network and computing resources consumption.

With respect to existing frameworks, a set of single domain and multi-domain orchestration implementation are already available. Ericsson Harmonizer [7] provides data center interconnections crossing heterogeneous networking domains that differ in terms of switching technology, control system, and vendors. It provides dynamic and carrier grade end-to-end transport connectivity combining heterogeneity, elasticity, and traffic engineering capabilities in each domain. The work in [8] presents a prototype of a SDN orchestrator which provides elasticity at the level of the data service delivery path. The framework performs the orchestration of network connectivity resources allowing it to take appropriate actions whenever data delivery degradations (e.g., congestions) are detected in the network paths connecting the whole or part of service chain. Further adaptability of the framework at the VNFs level including their lifecycle management via a cloud management platform is considered as future work.

There are also a number of research works that tackle orchestration over multiple technological domains, and address some limited orchestration features across several administrative domains. T-NOVA [9] introduces a novel framework allowing operators not only to deploy virtualized

Network Functions (NFs) for their own needs but also to offer them to their customers, as value-added services by means of an NFV Marketplace. Below the Marketplace layer, T-NOVA implements a management/orchestration platform (TeNOR) for the automated provision, configuration, monitoring, and optimization of Network Functions-as-a-Service (NFaaS) over virtualized Network/IT infrastructures in compliance with the ETSI NFV MANO architecture. T-NOVA leverages on and enhances cloud management architectures for the elastic provision and (re-) allocation of IT resources assigned to host the NFs. ESCAPE [10] is a proof of concept prototype developed in the UNIFY FP7 project [11] which supports the development of the main functionalities of a service chaining architecture (i.e., VNF implementation, traffic steering, virtual network embedding and so on). ESCAPE can operate as domain orchestrator for different technological domains. It can also act as a multi-domain orchestrator, which is capable of instantiating service elements in several separate domains. It mainly: (i) provides a Service layer that receives service requests in the format Network Function Forwarding Graph NF-FG; (ii) maps the received NFFG to the available resources; (iii) configures properly the NF; and then (iv) performs the Resource control for the resource domains. Another relevant framework is FROG [12], a Flexible and pROGrammable edge router that offers (to different actors including end users) the possibility to deploy network functions on (a portions of) the network traffic flowing through the edge router. It supports multi-tenancy, guarantees traffic isolation among NFs of different tenants, and executes thousands of NFs on the same server. Finally, Nestor [13] is a network service embedding (NSE) orchestrator which addresses the main challenges faced by the assignment of service chains across multiple NF providers (NFPs). It provides a holistic solution including network service request partitioning, generation of NF-subgraphs with DC gateway bindings, and NF-subgraph mapping.

III. THE 5GEX MDO PROTOTYPE ARCHITECTURE

The 5GEx Mdo prototype consists of the three functional blocks described in Figure 1. The Catalogue and Topology Exchange function is responsible for exchanging (abstracted) topology information with other orchestrators as well as the 5GEx Mdo service level capabilities. The Service Orchestration component performs the actual deployment of the multi-operator slices or the multi-operator services based on the information that has been discovered. After services are deployed the service assurance management system kicks in ensuring the performance of the service across the different domains within the agreed operational parameters. The first release of the prototype in the 5GEx project focuses on the service orchestration functional block and the basic information needed from the catalogue and topology exchange block to perform the orchestration.

The implemented architecture for the 5GEx Mdo prototype is shown in Figure 2. The architecture has been obtained starting from the study proposed in [14]. The service orchestration component gathers the abstracted view of the topology of other orchestrators via the topology abstraction component. In addition, it is aware of its own domain's

topology. The GUI can be used to make slice deployment requests where services are constructed from a catalogue management system. Future versions of the prototype will be able to exchange catalogue information with other orchestrators (this feature is currently under implementation). Based on the chosen constructed service, a network service description (NSD) is created and handed over to a preliminary implementation of a Network Service Orchestrator (NSO). The NSO translates the NSD to a Virtual network function forwarding graph (VNFFG) based on the information contained in the VNF store and as described in ETSI NFV [3]. A resource orchestrator (RO-ESCAPE [10]) breaks down the VNFFG and forwards parts of it to other domains for deployment. This way a multi-domain service is realized by this architecture.

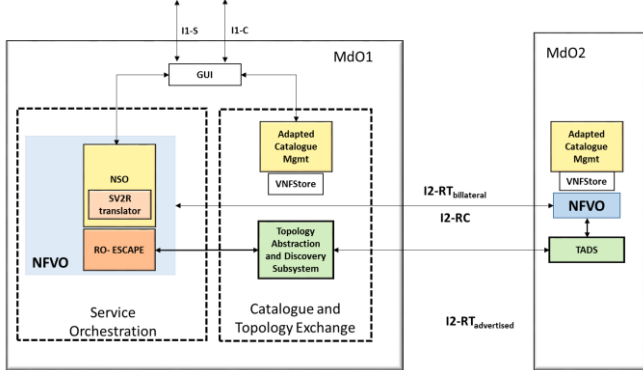


Figure 2 Implementation architecture of prototype 1

IV. EXPERIMENTAL VALIDATION

This section focuses on the description of the experiment carried out in order to assess the performance of the 5GEx MdO prototype. In particular, the testbed organization, the reference scenario and the experiment are highlighted in the following subsections.

Testbed

The validation experiments have been performed using the 5GEx Sandbox multi-domain testbed (see Figure 3). The 5GEx Sandbox connects up to 13 project partners' premises/local testbeds, hereafter referred to *islands*, by means of dedicated tunnels. The Sandbox includes a control plane network and a data plane network. The control plane network connects controllers, orchestrators, and 5GEx MdOs of different islands and is realized by means of a full-mesh p2p IPsec tunnels configuration. The data plane connects the configurable IT resources of each island through GRE tunnels by means of hierarchical domains (i.e., Tier2/3, and Tier1 domains) offering IP/MPLS connectivity services. Data plane multi-domain routing is enabled by running Border Gateway Protocol (BGP) instances at each domain border area router (currently each island has a Quagga router instance running BGP as a gateway). In the experiment hereafter described, a number of islands have been involved. In particular, the UCL island (London), the KTH island (Stockholm), the BME island (Budapest), and the UC3M island (Madrid) act as Tier3 islands offering virtual computational resources (i.e., VMs) to be used for instantiating network services formed of different VNFs. At

the control plane of each island, a 5GEx MdO prototype orchestrates the inter-operator VNF requests with the 5GEx MdOs of different islands and the intra-operator resources with a local Domain Orchestrator (DO). Finally, SSSA island (Pisa), besides virtual computational resources, acts as Tier1 connectivity domain, deploying an IP/MPLS/Segment Routing network domain of 10 Juniper routers [14] connecting all the other Tier3 islands.

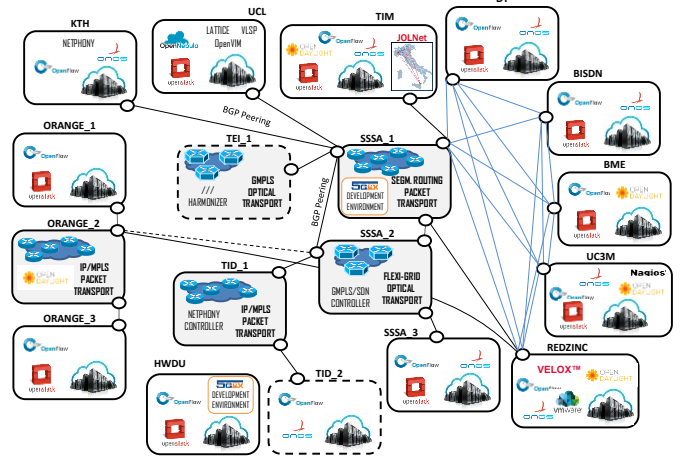


Figure 3 5GEx Sandbox islands topology

Reference scenario

For the experimental work, we defined a reference scenario (Figure 4) which captures all the relevant features and capabilities of our 5GEx MdO prototype. More specifically, it refers to a VNFaaS use-case, involving two 5GEx providers (5GEx Provider #1 and #2). The different functional components of the 5GEx MdO prototype are instantiated as Docker containers (DoCs). The control & orchestration plane communication between 5GEx MdOs and Domain Orchestrators (DOs) is performed using a dedicated network and each containerized 5GEx MdO runs in a dedicated machine (i.e., VM#1 and VM#2) including:

- NF Store (DoC#1)
- T-NOVA Marketplace [9] (adapted to 5GEx), composed of several containers (DoC#2-13)
- Network Service Orchestration (DoC#10-14)
- ESCAPE RO (DoC#15)
- TADS (DoC#16)

In the current version of the 5GEx MdO, the communication with the underlying DOs is performed via the Unify Virtualizer interface (I3) [11]. The data plane is composed by two Docker-Host (DH) machines: Docker-Host1 (DH1) for 5GEx Provider1 and Docker-Host2 (DH2) for 5GEx Provider2. Each DH in addition to the functionalities needed for instantiating Docker images (i.e., a Docker engine), also includes the following elements:

- a Domain Orchestrator that activates the Unify-based communication with the 5GEx MdO, create slices, and perform the slices orchestration

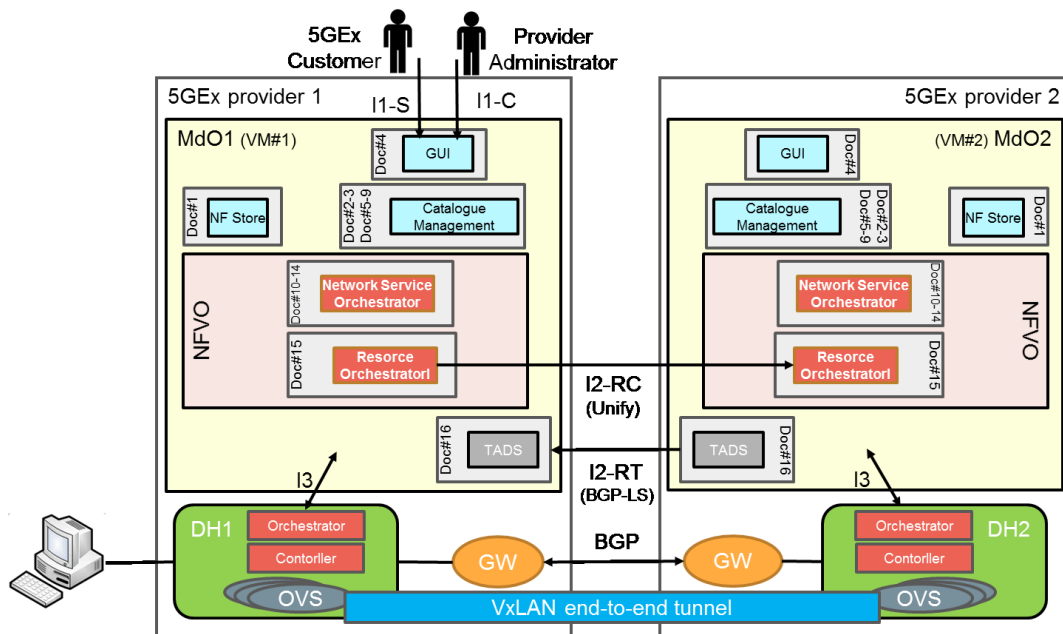


Figure 4 Proposed reference scenario

- an OpenFlow controller (OF version 1.0) to setup connectivity between containers and the input/output interfaces
- one or more instances of OVS switch in order to enable the data plane inside the created slices within the DH.

The end-to-end connectivity between the DHs of the two providers is performed by means of Quagga routers (GWs) with an active BGP session. Over this L3 connectivity, a VxLAN tunnel (i.e., overlay solution) is created in order to extend the L2 data plane to the network slice hosted by other operators. The workstation connected to the DH1 acts as end-user of the system. Since it is connected to the data plane of the slice, it is able to access the deployed NSs and is used to verify the overall behavior of the system from the end-user perspective.

Experiment

The experiment discussed in this section demonstrates and validates the successful instantiation of Network Services (NSs) across different 5GEx testbed sites using the current 5GEx MdO prototype. Considering the reference scenario, two different Sandbox islands (i.e., acting as two providers) are involved in a single experiment run. In each site, a 5GEx MdO instance is running in order to enable 5GEx functionalities. Moreover, a separate data plane node is considered (the previous mentioned DH), including its own domain orchestrator above the virtual infrastructure manager (i.e., the Docker engine). The Docker nodes support the deployment of containerized virtual network functions (VNFs). More specifically, the VNFFG in this particular experiment consists of two VNF images, i.e., a Web Server and a deep packet inspection (DPI) instance, linked together and accessible by the customer through a Service Access Point. We configured the VNF capabilities of each testbed site such that the two images of the service could not be deployed at the same place. We

considered the experiment successful when the chain is deployed involving both sites and the customer could access a test web page available on the Web Server VNF running at DH1 and inspected by the DPI VNF (some of the web pages contained words that were prohibited and thus blocked by the DPI) deployed at DH2, or vice versa. The experiment was reproduced under the same conditions considering different couple of remote sites (KTH-SSSA, KTH-UCL, etc.). Moreover, a live demonstration about a similar experiment on a multi-domain NS provisioning has been carried out in [15].

In order to test the different components/functionalities of the 5GEx MdO prototype we identified different experiment steps as follows:

- A resource slice including a (sub)set of the whole infrastructure has to be created on the involved DHs through the API provided by the local Domain Orchestrators.
- The two slices created in the DHs have to be linked through a layer 2 tunnel (connecting the OVS in each slice segment).
- The 5GEx MdO has to be activated at each site in order to: (i) initialize the NF Store and perform the VNF onboarding; (ii) start the 5GEx GUI; (iii) activate the NFVO in order to enable the orchestration of the NS requests; (iv) connect the local slice segment, and (v) discover the 5GEx MdO in the other site through TADS, in order to extend its domain with the resource slice provided.
- Initiating a connection to the Marketplace as Service Provider (I1-C) to perform the on-boarding and publishing of the reference NS (i.e., NS1, composed by a Webserver VNF and a DPI VNF).
- The 5GEx customer connects to the Marketplace and requests the deployment of NS1 as a Service instance

on the available slice. This operation requires user interaction with the Marketplace by means of the provided GUI interface (I1-S).

- A NVFFG resulting from the previous service instantiation is generated and sent to the Resource Orchestrator of the request receiving 5GEx MdO (i.e., MdO1) that will perform the adaptation and orchestration algorithms, involving the other 5GEx MdO (i.e., MdO2).
- Resources will be allocated on the slice by the 5GEx MdOs (according to the outcome of the orchestration process) by interacting with the underlying Docker domain orchestrators.

In order to validate the experiment, an HTML web page is downloaded by the end-user connected to the service access point from the Web Server VNF created during the service instantiation. In case such a web page is requested that contains banned words, the DPI filters out the content. We measured the elapsed time since a customer requests the instantiation of the NS (Web Server and a DPI), and the instant when each VNF is instantiated by the two DOs. The measurements were repeated 100 times. In Figure 5 the distribution of the obtained samples is shown. In particular, the average overall deployment time is $\mu=1852\text{ms}$, with a standard deviation of $\sigma=183\text{ms}$, and the 90th percentile was 2112ms. Moreover, we evaluated how long it

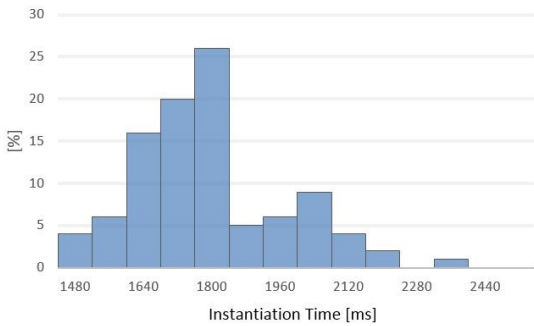


Figure 5: Instantiation time of the 100 NS requests.

takes to deploy the two 5GEx MdOs that are required for the orchestration of the NS. Considering 50 5GEx MdOs deployments, Figure 6 shows the distribution of the MdO




Figure 6: Distribution of the MdO deployment time.

deployment time. We computed an average of $\mu=82\text{s}$ to deploy each of them, with a standard deviation of $\sigma=45\text{s}$, and a 90th percentile of 86s.

V. CONCLUSIONS

The 5GEx MdO prototype presented in this paper is the result of the integration process of properly extended open source software tools at the disposal of the 5GEx partners. Considering the Slice as a Service (SlaaaS) use-case in the context of a multi-operator scenario, both the creation and the deployment of NSs have been considered in order to evaluate the performance of the prototype. The results obtained during the experiments, carried out in the 5GEx sandbox, show the effectiveness of the system, with few seconds needed for the deployment of the NS involving two operators.

ACKNOWLEDGMENT

This work has been performed in the framework of the H2020-ICT-2014 project 5GEx (Grant Agreement no. 671636), which is partially funded by the European Commission. This information reflects the consortium's view, but neither the consortium nor the European Commission are liable for any use that may be done of the information contained therein.

REFERENCES

- [1] M. Fiorani, B. Skubic, J. Mårtensson, L. Valcarenghi, P. Castoldi, L. Wosinska, P. Monti, "On the Design of 5G Transport Networks," *Springer Photonic Network Communications (PNET) Journal*, Vol. 30, No. 3, pp. 403-415, December 2015.
- [2] 5GEx webpage <http://www.5gex.eu/>.
- [3] ETSI NFV ISG, Network Functions Virtualisation: An Introduction, Benefits, Enablers, Challenges & Call for Action. In *SDN and OpenFlow World Congress: Darmstadt-Germany*, 2012.
- [4] ITU-T. Study Group 13 2016.
- [5] ITU-T. *Cloud Computing Framework for End to End Resource Management ITU-T Y*, Vol. 3520. Global Information Infrastructure, Internet Protocol Aspects And Next-Generation Networks, 2015.
- [6] IEEE Std 1903–2011. Standard for the Functional Architecture of Next Generation Service Overlay Networks, October 2011.
- [7] P. Iovanna "Effective elasticity for data centers interconnection in multi-domain WAN: Information modelling and routing". *European Conference on Optical Communication (ECOC)*, Valencia, 2015; 1–3.
- [8] A.A. Mohammed, M. Gharbaoui, B. Martini, F. Paganelli, P. Castoldi, SDN controller for network-aware adaptive orchestration in dynamic service chaining, NetSoft Conference and Workshops (NetSoft), 2016 IEEE. 126-130.
- [9] FP7 Project T-NOVA. <http://www.t-nova.eu/>.
- [10] Balazs Sonkoly. et al. Multi-domain service orchestration over networks and clouds: a unified approach. *ACM Conference on Special Interest Group on Data Communication*, 2015, ACM.
- [11] FP7 Project Unify. <https://www.fp7-unify.eu/>.
- [12] I. Cerrato, F. Risso, A Scalable and Massively Multi-tenant Platform for User-oriented Network Services, draft paper, 2012.
- [13] D. Dietrich, A. Abujoda, P. Papadimitriou, Network service embedding across multiple providers with Nestor, IFIP Networking Conference (IFIP Networking), 2015.
- [14] R. Guerzoni et al., "Multi-domain Orchestration and Management of Software Defined Infrastructures: a Bottom-Up Approach", European Conference on Networks and Communications (EuCNC2016), 2016.
- [15] A. Sgambelluri, F. Paolucci, A. Giorgetti, F. Cugini and P. Castoldi, "Experimental Demonstration of Segment Routing," in *Journal of Lightwave Technology*, vol. 34, no. 1, pp. 205-212, Jan.1, 1 2016. doi: 10.1109/JLT.2015.2473656.
- [16] A Sgambelluri, et al. "A Multi-Operator Network Service Orchestration Prototype: The 5G Exchange" Optical Fiber Communication Conference, 2017, Los Angeles, CA, USA