This work has been accepted for publication in European Conference on Networks and Communications (EuCNC 2018).
DOI: 10.1109/EuCNC.2018.8443210.

5G OS: Control and Orchestration of Services on Multi-Domain Heterogeneous 5G Infrastructures

Sevil Dräxler*, Holger Karl*, Hadi Razzaghi Kouchaksaraei*, Azahar Machwe†,
Crispin Dent-Young†, Kostas Katsalis‡ and Konstantinos Samdanis‡
*Paderborn University, Paderborn, Germany
†Zeetta Networks Limited, Bristol, UK
‡Huawei Technologies, Co., Ltd, Germany

***Abstract*—A heterogeneous 5G infrastructure includes physical and virtual computation, storage, and networking resources, exploiting different technologies and spanning across several administrative domains. These resources can be combined into end-to-end slices that can host different services. Services consist of virtual or physical elements and have different requirements, e.g., in terms of resource demands, latency, performance guarantees. The underlying resources need to be efficiently allocated to the slices and their corresponding services. To abstract away the complexities of the underlying 5G infrastructure and to provide the common functionalities required for efficient and flexible service and slice management and orchestration, we propose 5G Operating System (5G OS). We describe the requirements of such a system, present the high-level architecture of 5G OS, and describe the design challenges for different interfaces within 5G OS. Our proposed architecture includes scalable hierarchies of multi-domain and domain-specific slice and service orchestrators, software-defined network controllers, and network function virtualization management and orchestration frameworks.**

I. INTRODUCTION

A 5G infrastructure consists of heterogeneous physical and virtual resources. Resources include computation, storage, and network resources, using different technologies, e.g., optical network, MPLS, etc. We refer to a group of resources having the same technological specifications as a *technological domain*. Moreover, different resources may be governed by different usage and operational rules and policies. We refer to a group of resources under the same administrative policies as an *administrative domain* (e.g., the wireless access administrative domain).

This infrastructure needs to support a diverse set of use cases with different requirements, e.g., in terms of resource demands, latency, end-user mobility, and reliability. Examples of these use cases include industrial automation and control, video streaming, online gaming, Internet of things, and cloud applications [1].

Use cases are realized by deploying *services* on the 5G infrastructure. Services may be composed of multiple *service elements*. For example, a video streaming service consists of stream selection and stream aggregation elements, which can be deployed as virtual network functions (VNFs), content distributor elements, which can be deployed as cloud-based content servers, as well as access points and gateway elements, deployed as physical network functions (PNFs). Of course, a service can also be realized by a single element like a Home Subscriber Server (HSS), which is deployed as a single application server.

Main stakeholders in deploying and using services on top of the underlying infrastructure include: (1) *infrastructure owners*, (2) *infrastructure providers* (who may be owners or resellers of the infrastructure), (3) *service providers* (who may be owners or resellers of services), (4) *service users* (who may be end users or may be service providers using a service for providing another service).

Service users and providers benefit from large-scale heterogeneous hardware and software resources but should not have to deal with the complexities of managing and operating them. Infrastructure providers and owners also need appropriate tools for efficient management of their resources and handling the requests from their customers. Moreover, well-defined interfaces for dealing with different types of resources are required, for unifying the distributed, heterogeneous resources into a common 5G infrastructure that can host different services. For this, an abstraction layer is required, analogous to a computer operating system that acts as an intermediary between programs and the computer hardware, manages the hardware and software resources, and provides common services for computer programs and users [2], [3].

Different resources of a heterogeneous 5G infrastructure are suitable for service requirements. For example, a function deployed on as virtual machine on a general-purpose hardware has a different performance compared to the same function deployed using special-purpose programmable hardware. A service element that can be deployed on different resource types can react efficiently to dynamic load situations. This requires different (software) versions of the service element to be implemented.

An interesting use case for 5G OS is controlling and orchestrating such *multi-version services*. Based on the predefined objectives for service performance and resource utilization, the best composition of multi-version service elements can be selected and mapped to the resources. Mapping service requirements onto such a diverse set of resources requires awareness of the existing supply of compute, storage, and network resources, as well as information about the elasticity of these resource, i.e., possibility and costs of modifying, creating, re-assigning such resources. This allows optimal decisions to be made regarding the level of application split, allocation of compute and storage resources, and network link allocations. This is not a trivial problem given that resource availability is dynamic and new services must not disrupt existing ones.

We introduce the architecture of **5G Operating System (5G OS)** that

- provides control and management for services that run on top a multi-domain, heterogeneous 5G infrastructure,
- implements common functionalities between virtual and physical network functions (VNFs and PNFs) as service elements on one hand and the programs that manipulate or use them (e.g., NFV orchestrators, service managers, BSS, OSS) on the other hand,
- combines control and management of services with resource management by handling the underlying physical and virtual resources in the form of end-to-end (E2E) *slices*. E2E slices consist of compute, storage, network resources and management functions, which can be allocated and modified as needed.

Aligned with the network softwarization trend that is paving the way towards 5G systems [1], software-defined networking (SDN), network function virtualization (NFV), and end-to-end slice orchestration over multi-domain infrastructures can be used as underlying concepts for realizing such an operating system. 5G OS can be realized as a framework consisting of multiple components. For flexible and programmable provisioning of network, compute, and storage resources as well as efficient management of services on top of them, SDN controllers and NFV management and orchestration (MANO) systems need to collaborate. Additionally, E2E slice orchestrators need to be integrated with these systems to abstract and isolate the disaggregated 5G resources shared by a variety of services. These components become functional in combination with information management systems, infrastructure and service monitoring, and decision algorithms for automatic management and orchestrations of services and resources.

There have been several developments (research, implementation, and standardization activities) regarding SDN controllers, NFV MANO frameworks, and slice orchestrators. The relationship among these components is, however, not very well-defined yet. In this paper, we present the initial design of the 5G OS architecture as a proposal for the relationship and interplay of these components. This proposal will be refined in the context of the 5G-PICTURE project [4] by investigating questions including the following:

- How are the functionalities of 5G OS distributed among these components? E.g., is there a need for a full-fledged slice orchestrator or can we handle *network* slicing as a part of SDN control and *compute/storage* slicing as a part of NFV MANO?
- How are the areas of responsibility defined for each instance of a 5G OS component? E.g., does each slice require dedicated NFV MANO frameworks or can NFV MANO frameworks be responsible for several slices?
- How and when should the components of the 5G OS be scaled?
- How can monitoring information from different resources and different services be propagated and used for better control and management decisions?

In the rest of this paper, we first give an overview of the state of the art (Section II). Afterwards, we present the 5G OS architecture and the challenges and implementation options for realizing such an operating system (Section III) before concluding the paper.

II. State of the Art

Currently, there are no comparable frameworks that can handle E2E slice orchestration, SDN control, and NFV management and orchestration to support complex services on top of the multi-technology and multi-administrator heterogeneous 5G infrastructure. Therefore, in this section, we give an overview of existing work in three major fields relevant to the functionality of the 5G OS, namely, NFV management and orchestration, SDN control, and network and compute slicing.

A. NFV Management and Orchestration

Regarding NFV orchestration, the European Telecommunications Standards Institute (ETSI) Management and Orchestration (MANO) architectural framework describes the functional blocks, data repositories and interfaces for orchestrating virtualized infrastructures, network functions, and services [5]. This reference architecture is the base for the design of the NFV MANO component in 5G OS. This architecture, however, cannot fulfill all requirements of 5G OS as it is. For example, no clear guidelines regarding integration of SDN and NFV is given in this reference architecture. Moreover, the proposed information models that accompany this reference architecture cannot be used for uniformly specifying the heterogeneous resources, the slices composed out of resources, and the resource demands of services running on top of this.

The MANO architecture is realized by several open-source solutions. For example, Open Source MANO (OSM) [6], Tacker [7], ONAP [8], JOX [9], and SONATA [10]. Most of these solutions focus on orchestration of services consisting of virtual network functions on top of limited resource types. There are several practical issues to solve for the integration of them into the rest of the 5G OS framework, for example, supporting multi-domain infrastructures. In this regard, the 5G Exchange project [11] has been dealing with hierarchical multi-domain NFV orchestration. Their solution is based on a peer-to-peer model for the communication of multi-domain orchestrators, which limits the possibilities for supporting different administrative domains as well as multiple technological domains as the same time.

B. SDN Control

For controlling network resource within a domain, popular SDN controllers like OpenDaylight [12], ONOS [13], Floodlight [14], and Ryu [15] can be used as a starting point. However, these solutions do not provide standard virtualisation features which allow, for example, exposing PNFs (e.g., a physical OpenFlow-enabled switch) as VNFs to the layers above.

For managing networking resources over multiple domains, there are options like NetOS® [16], a network operating system based on OpenDaylight, which currently only provides a limited support for non-OpenFlow devices. Tzanakaki et al. [17] have also presented a multi-domain network architecture considering both wireless and optical domains. Katsalis et al. [18] have described some implementation experience for multi-domain SDN solutions for integrating wireless and optical networks. 5G OS can build upon these solutions to provide support for controlling heterogeneous network resource over multiple domains.

C. Network and Compute Resource Slicing

Network slicing has been defined by various standards bodies. Network slicing in the context of 5G is a concept introduced by NGMN (Next Generation Mobile Network) in the NGMN 5G White paper [19], in Feb. 2015. For the mobile network, 3GPP has addressed network slicing from the architectural perspective [20], studied the related management and orchestration issues [21], and highlighted the related RAN slicing aspects [22]. Additionally, 3GPP has presented the enhance Décor (eDecor) approach for slicing the core network [23]. According to ITU-T [24], network slices are perceived as Logical Isolated Network Partitions (LINP) composed of multiple virtual resources, isolated and equipped with a programmable control and data plane.

Inspired by the existing specifications and recommendations, we consider four abstract building blocks for defining a slice: compute resources, storage resources, network resources, and control/management functions. We define a slice as an isolated, self-contained, virtualized set of resources that can be used for hosting services.

While there are several definitions and specifications for slices and slicing systems, no practical implementation of network and compute resource slicing, integrated into SDN control and NFV orchestration frameworks, as described in Section III, is available.

NFV orchestration frameworks like OSM [6] provide compute slicing functionality as a part of the management and orchestration of virtual network functions running over the virtualized infrastructure. Along with these frameworks, SDN controllers are required to provide the underlay and overlay network connectivity and enable compute slicing.

III. 5G Operating System

Fig. 1 shows an overview of 5G OS, consisting of different boxes that represent functional components of 5G OS, the high-level interfaces among them, as well as the underlying 5G infrastructure.

As shown in the example network in Fig. 2, the infrastructure is organized into pre-defined technological and administrative domains. These domains may change over time, e.g., when a new technology is installed in a part of the network, but 5G OS cannot make changes to these domains. Instead, 5G OS receives the domain specifications as input and can combine compute, storage, and network resources from different domains into E2E *slices* that can be used for service instantiation. An administrative domain may span across multiple technological domains. For example, in Fig. 2, the Core administrative domain consists of the technology domains Synchronous Digital Hierarchy (SDH), Dense Wavelength Division Multiplexing (DWDM), Multiprotocol Label Switching (MPLS), and Ethernet WAN. Similarly, different parts of one technological domain may be under different

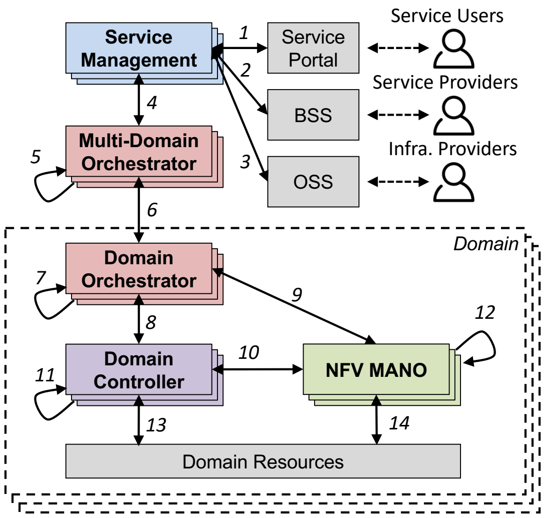

Fig. 1. 5G Operating System architecture

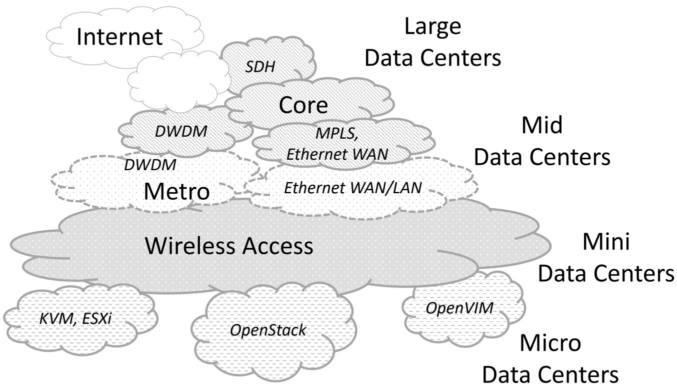

Fig. 2. Examples of different technological and administrative domains in the 5G infrastructure

administrative domains. For example, DWDM in Fig. 2 spans across Core and Metro administrative domains. The E2E slices composed by 5G OS may span across several technological and administrative domains.

The components of 5G OS are organized in a hierarchical way, receiving service and resource management requests and propagating them down towards the actual infrastructure. Different layers of this hierarchy, as we describe in this section, are responsible for end-to-end resource and service orchestration, SDN control, and NFV MANO operations. Describing all interfaces among all components is not possible due to space restrictions. Therefore, we give a brief overview of the important interfaces and present some open questions regarding realization of these interfaces.

A request for instantiating, managing, or modifying a slice or a service can enter 5G OS via the following three components:

- *Service Portal*, used by service users for requesting and managing services offered by service providers.
- *Business Support System (BSS)*, used by service providers for requesting and managing resources and services on top of the infrastructure offered by infrastructure providers.
- *Operations Support System (OSS)*, used by infrastructure

providers for requesting and managing slices they provide to their customers, on top of the underlying resources offered by infrastructure owners.

These components may be under different administrative domains with different business and management policies. They offer different interfaces to the corresponding stakeholders. The interactions between these components and the rest of 5G OS happen via the *Service Management* component using interfaces marked with numbers 1–3 in Fig. 1.

The Service Management component has the important task of converting high-level requests of the aforementioned stakeholders into requirements in terms of resources and slices [1], [21]. Different instances of the Service Management component can be deployed, e.g., to serve different service users, service providers, and infrastructure providers in different administrative domains. No direct interaction is required between different instances of this component. The actual communication required for end-to-end management and control of resource and services over multiple domains happens in the lower levels of the hierarchy within the architecture shown in Fig. 1.

Each instance of the Service Management component is responsible for FCAPS (Fault, Configuration, Accounting, Performance and Security) management of slices and maintaining Service-Level Agreements in its corresponding responsibility area. Additionally, the Service Management component is responsible for making the initial service placement decisions by interacting with and selecting from the instances of *Multi-Domain Orchestrators (MDOs)* it is integrated with. The placement calculation takes the current resource availability into account and, ideally, does not degrade the performance and quality of existing services. The information required for solving such a multi-objective optimization problem is provided to the Service Management components via interface 4 from MDO components.

MDOs are responsible for lifecycle management decisions of compute, storage, and network resource slices that span across multiple technological and/or administrative domains. Each MDO has information about the set of existing slices and can request instantiation, modification, or termination of slices within the domains it is responsible for. The way a request from a Service Management component is delivered to MDO instances is related to the realization of the interaction among MDOs, represented by interface 5 in Fig. 1. This interface can be realized in different ways. For example, MDO instances may have a hierarchical relationship, in which the highest MDO in the hierarchy receives the requests and distributes it to other instances. Alternatively, the MDOs may have a peer-to-peer relationship, where different MDOs can communication and cooperate directly, in case a certain request cannot be fulfilled by then individually. Within the 5G-PICTURE project, we are investigating the best option for realizing interface 5.

Via interface 6, each MDO instance can receive information from a group of *Domain Orchestrators (DOs)* it is associated with and can request lifecycle management operations from each DO, for the corresponding parts of end-to-end slices that are under the responsibility of that DO. Each DO is responsible for compute, storage, and network slice lifecycle management decisions within a certain technological or administrative domain. DOs abstract away the differences in these domains and expose a uniform interface to MDOs for requesting and managing the resources from different domains. Similar to the inter-MDO interactions, inter-DO interactions via interface 7 can also be realized in different ways (e.g., hierarchical, peer-to-peer, etc.), which also affects the data and control flow among DOs as well as between DOs and MDOs.

DOs break the slice specification into concrete actions to be performed on top of networking resources and compute and storage resources. Via interface 8, DOs interact with *Domain Controller (DC)* component(s), responsible for controlling the network resources. Each DC can provide paths and network resources across its specific domain. Different DCs might be required inside one domain, e.g., when there are multiple technologies, each requiring their own network controller within an administrative domain. Interface 11 among different DCs can be implemented in different ways, as discussed for interface 5 and 7. DCs use interface 13 to directly control and manipulate the underlying infrastructure.

Via interface 9, DOs can interact with *NFV MANO* component(s), responsible for managing and orchestrating compute and storage resources. NFV MANO manages and orchestrate the virtualized service elements. For example, each NFV MANO instance can calculate the placement for a specific VNF or scale it if necessary within the resources that are available to it. or simplicity, we do not show the details of the MANO system but we assume a model compatible to the ETSI NFV reference architecture [5]. Multiple instances of the NFV MANO component might be required within one domain, e.g., to cover multiple geographically distributed data centers. Interface 12 is used for the interactions among different NFV MANO instances, which may follow different models. Via interface 14, NFV MANO instances do the actual resource allocation to service elements.

Interface 10 stands for the interactions among DC and NFV MANO instances. Defining this interface and the relationship among these components is a challenging task we are tackling, towards implementing 5G OS. ETSI NFV has provided recommendations for roles of SDN controllers in an NFV ecosystem [25]. In this document, the resources that can be controlled by an SDN controller are defined as physical/virtual switches and routers, softswitches, as well as switches and routers implemented as VNFs. For example, depending on the scenario, 5G OS may need to deal with one of the following cases, which require different interfaces and control flows between DC and NFV MANO:

1) SDN controller as an underlay network controller [ETSI case 5], controlling the physical or virtual infrastructure that serves the components on top of the virtual infrastructure. This implies that for managing a slice or a service, the DC decides about network resource allocation and instructs the NFV MANO to use the compute and storage resources corresponding to this allocation, for instantiating the service entities.
2) SDN controller as a Virtualized Infrastructure Manager (VIM) [ETSI case 1 and 3], controlling the internal network of an infrastructure node (e.g., a data center). This means the NFV MANO instructs the DC to create the required network paths after the placement of service entities is calculated.

In the 5G OS architecture, Service Management, Multi-Domain Orchestrator, and Domain Orchestrator components are organized in a functional hierarchy. Domain Controller and NFV MANO components can interact with one another if necessary and perform the tasks assigned to them when triggered by the higher level of the hierarchy. This model is consistent with the 3GPP recommendations for management and orchestration of network slicing [21]. For example, the Communication Service Management Function (CSMF) role defined by 3GPP is represented by the Service Management component in Fig. 1. Our Multi-Domain and Domain Orchestrator components include the functionalities of the Network Slice Management Function (NSMF) and Network Slice Subnet Management Function (NSSMF) from 3GPP definitions.

An important step towards realizing 5G OS is to provide a common description model for slices and services to support different use cases, such as enabling external control of network slices, allowing slice pinning based on physical infrastructure (e.g., ports) and requesting specific resources from the infrastructure provider. For example, the ETSI OSM Information Model [26] is an attempt to standardize network service descriptors to include virtual links, virtual network functions and service forwarding graphs. But this does not include physical network functions. ETSI OSM (Release Three) also demonstrates only limited multi-data center orchestration capabilities.

Once the interfaces are clearly defined, existing implementations can be used as a starting point for different components of 5G OS. For example, the MANO frameworks of OSM [6] or SONATA [27] can be used as the base for NFV MANO component implementation. For the relationship of multi-domain and domain-specific components, 5GEx [11] multi-domain NFV orchestration interfaces can be used as a working example. For multi-domain SDN control, NetOS® [16] can be used as a network operating system based on OpenDaylight [12].

IV. Conclusion

5G Operating System (5G OS) is a framework for orchestration of complex services and end-to-end slices on top of a heterogeneous, multi-domain 5G infrastructure. We present the high-level architecture of 5G OS in this paper that is based on a scalable hierarchy of different components. In this design, the resources in different technological and administrative domains are controlled and manipulated by domain-specific SDN controller and NFV management and orchestration systems. Domain-specific orchestrators build slices of the network, compute, and storage resources, while multi-domain orchestrators handle the lifecycle management of end-to-end slices that can span across multiple administrative domains and include several technological domains to be able to provide the required service levels to different stakeholders of the 5G OS.

Acknowledgment

This work has been partially supported by the 5G-PICTURE project, funded by the European Commission under grant number 762057 through the Horizon 2020 and 5G-PPP programs. The authors would like to thank all colleagues from the 5G-PICTURE project for the discussions and contributions that have led to this work.

References

[1] 5GPPP Architecture Working Group, "5G Architecture White Paper Version 2.0," Dec. 2017.

[2] W. Stallings and M. M. Manna, *Operating systems: internals and design principles*. Pearson, 2015.

[3] I. A. Dhotre, *Operating systems*. Technical Publications, 2009.

[4] "5G-PICTURE Project Website," accessed on 08.02.2018. [Online]. Available: http://www.5g-picture-project.eu/

[5] ETSI NFV ISG, "GS NFV-MAN 001 V1.1.1 Network Function Virtualisation (NFV); Management and Orchestration," Dec. 2014.

[6] "Open Source MANO Project Website," accessed on 08.02.2018. [Online]. Available: https://osm.etsi.org/

[7] "Tacker Website," accessed on 08.02.2018. [Online]. Available: https://wiki.openstack.org/wiki/Tacker

[8] "Tacker Website," accessed on 08.02.2018. [Online]. Available: https://www.onap.org/

[9] K. Katsalis, N. Nikaein, and A. Huang, "Jox: An event-driven orchestrator for 5g network slicing," in *Accepted in IEEE Network Operations and Management Symposium (NOMS)*, 2018.

[10] S. Dräxler, H. Karl, M. Peuster, H. R. Kouchaksaraei, M. Bredel, J. Lessmann, T. Soenen, W. Tavernier, S. Mendel-Brin, and G. Xilouris, "SONATA: Service programming and orchestration for virtualized software networks," in *2017 IEEE International Conference on Communications Workshops (ICC Workshops)*, May 2017, pp. 973–978.

[11] "5GEx Project Website," accessed on 08.02.2018. [Online]. Available: https://www.5gex.eu/

[12] "OpenDaylight Website," accessed on 08.02.2018. [Online]. Available: https://www.opendaylight.org/

[13] "ONOS Website," accessed on 08.02.2018. [Online]. Available: https://onosproject.org/

[14] "Flooglight Website," accessed on 08.02.2018. [Online]. Available: http://www.projectfloodlight.org/floodlight/

[15] "Ryu Website," accessed on 08.02.2018. [Online]. Available: https://osrg.github.io/ryu/

[16] "NetOS Website," accessed on 08.02.2018. [Online]. Available: http://www.zeetta.com/netos/

[17] A. Tzanakaki, M. P. Anastasopoulos, S. Peng, B. Rofoee, Y. Yan, D. Simeonidou, G. Landi, G. Bernini, N. Ciulli, J. F. Riera *et al.*, "A converged network architecture for energy efficient mobile cloud computing," in *Optical Network Design and Modeling, 2014 International Conference on*. IEEE, 2014, pp. 120–125.

[18] K. Katsalis, B. Rofoee, G. Landi, J. Riera, K. Kousias, M. Anastasopoulos, L. Kiraly, A. Tzanakaki, and T. Korakis, "Implementation experience in multi-domain sdn: Challenges, consolidation and future directions," *Computer Networks*, vol. 129, pp. 142–158, 2017.

[19] NGMN Alliance, "NGMN 5G White Paper," Feb. 2015.

[20] 3GPP, "3GPP TS 23.501 V15.0.0 System Architecture for the 5G System," Dec. 2017.

[21] ——, "3GPP TR 28.801 V15.0.0 Study on Management and Orchestration of Network Slicing for Next Generation Network," Sep. 2017.

[22] ——, "3GPP TR 38.801 V14.0.0 Study on new radio access technology: Radio access architecture and interfaces," Mar. 2017.

[23] ——, "3GPP TR 23.711 V14.0.0 Enhancements of Dedicated Core Networks Selection Mechanism," Mar. 2017.

[24] ITU-T, "Y.3011 Framework of Network Virtualization for Future Networks," Jan. 2012.

[25] ETSI NFV ISG, "GS NFV-EVE 005 v1.1.1 Network Function Virtualisation (NFV); Ecosystem; Report on SDN Usage in NFV Architectural Framework," Dec. 2015.

[26] "OSM Information Model," accessed on 08.02.2018. [Online]. Available: https://osm.etsi.org/wikipub/index.php/OSM_Information_Model

[27] "SONATA Project Website," accessed on 08.02.2018. [Online]. Available: http://sonata-nfv.eu/