

“© © 2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.”

Chirivella-Perez, E., Salva-Garcia, P., Alcaraz Calero, J. M., Wang, Q., & Neves, P. M. (2019). Inter-business orchestration for resource and service provisioning in 5G network slicing. In 2019 European Conference on Networks and Communications (EuCNC) (IEEE Conference Proceedings). IEEE. <https://doi.org/10.1109/EuCNC.2019.8801947>

Inter-Business Orchestration for Resource and Service Provisioning in 5G Network Slicing

Enrique Chirivella-Perez, Pablo Salva-Garcia, Jose M. Alcaraz Calero, Qi Wang

The University of the West of Scotland, UK

Pedro M. Neves

Altice Labs, S.A. Aveiro

ABSTRACT

The Fifth Generation (5G) networks leverage softwarisation and virtualisation as cornerstones to build cost-effective and flexible infrastructures and services. Meanwhile, the advanced orchestrator is entailed to manage and orchestrate the resources and services at various levels of infrastructures, including the physical layer, virtualisation layer, service layer etc. The challenges become more evident when the multiple business roles, enabled by the various levels of operations in this new business model paradigm, are taken into account in an integrated operational environment. This paper proposes a new layered, inter-business vertical orchestration architecture to address the highlighted challenges in softwarised and virtualised 5G networks. The architecture distinguishes not only the infrastructure levels but also the business roles, both in an integrated framework with interactions among them specified. It is noted that the proposed vertical orchestration architecture is compatible with network slice orchestration in 5G networks. The proposed orchestration framework is prototyped in a testbed with realistic 5G infrastructures, and its performance is empirically evaluated focusing on the service creation time 5G Key Performance Indicator (KPI).

***Index Terms*—5G networks, 5G KPI, service creation time, infrastructure orchestrator, vertical orchestrator, automatic provisioning, 5G testbed, performance evaluation**

I. INTRODUCTION

The forthcoming 5G ecosystem of infrastructures and services in the telecommunication industries is fostering the emergence of different business roles enabled by key technologies such as virtualization and softwarization. Thus, an Infrastructure Service Provider (ISP) controls physical infrastructures to rent Infrastructure-as-a-Service achieving multi-tenancy isolation among the different tenants of the ISP. These infrastructures can be consumed by Cloud Service Providers (CSPs), which are in turn used to offer virtualized Infrastructure-as-a-service to the different CSP tenants through virtualization [1]. Thus, a Digital Service Provider (DSP) a business that consumes either physical or virtual infrastructures to deploy a software service that

is exposed as a service to its Digital Service Consumers (DSCs). If the DSC is a vertical business, it acts as platform-as-a-service to be used by the vertical businesses to provide a service to the final user. If the DSC is the final user, it offers software-as-a-service. Depending on the business strategy, different organisations are exploiting one or more roles in order to compete in the market for their niche. Network slicing can be defined as a way to warrant that the resources and services requested by a consumer are met by the provider of such resources and services. Therefore, slicing a service will warranty such expected service level agreement.

The orchestration is defined as a workflow for the execution of an ordered set of steps to achieve an operational goal. In this context, current orchestration architectures have traditionally been designed to perform horizontal orchestration strategies defined as a set of steps for the control and management of resources and services that belong only to one of these business roles previous described. However, there is a clear need to explore vertical orchestration strategies where there is support for the execution of steps involving the control and management of services and resources belonging to different business roles. This vertical orchestration is especially relevant when network slicing needs to be enforced along the value chain created through the interactions among all the business roles involved.

This vertical orchestration is key for the understanding of how virtualisation, softwarisation and network slicing can be utilised for commercial usage among various business roles in a coordinated way. The main contributions of this research are triple-fold. Firstly, the paper presents a new layered, vertical orchestration architecture with various business roles differentiated and the vertical interactions among business roles explicitly addressed. Secondly, the proposed orchestrator is designed and prototyped with support for vertical orchestration strategies that allow the provisioning of services and resources across the business value chain while allowing enforcing network slicing to provide warranties along with the value chain. Thirdly, empirical results are generated through experiments in a testbed with realistic 5G infrastructures and analysed to show that it meets the challenging requirements of the 5G service creation time Key Performance Indicator (KPI).

The rest of this paper is structured as follows. Section II reviews related work. Section III describes the architecture of the proposed orchestration framework. Section IV presents and analyses the experimental results based on a proof-of-concept implementation in a realistic 5G testbed. Section V provides additional discussions on business roles. Finally, Section VI concludes the paper.

II. RELATED WORK

In light of the importance of orchestration in 5G networks [2], a number of projects are addressing this topic especially in the context of network slicing. For instance, DASMO [3] presents a proof-of-concept orchestrator that uses a distributed and autonomous slice management architecture to solve the problem of scalability in the management plane. DASMO uses the ETSI MANO orchestration combined with multiple in-slice deployed autonomous management platforms. Chang et al. [4] devise an architecture over ultra-dense networks to handle a large amount of data in 5G and the slicing paradigm. The work optimises the network resources in terms of orchestration and management systems. The orchestrator interacts with the RAN modules to enable supporting slice-based multi-service chain creation and chain placement. Raza et al. [5] create a proof of concept of SDN/NFV-based orchestration to share infrastructure resources among different tenants. The orchestrator takes profit of the utilisation of the infrastructure provider using an efficient utilisation of its resources through dynamical network slicing. Farrel [6] describes a service function chaining and network slicing to provide support for 5G services across backhaul and metro networks. This allows NFV support for services to enable 5G connectivity. Zanzi et al. present an Overbooking Network Slices (OVNES) solution [7] where network orchestration across different network domains is proposed to install and inter-connect network slices through cloud data centres and the core network segment. Montero et al. [8] propose optical network control and resource orchestration into software-defined networks to fulfil the requirements imposed. The work defines different types of services defined for 5G verticals and presents an architecture for E2E slicing provisioning and monitoring 5G services over the optical segment. Li et al describe [9] an orchestration enabler for slicing with efficient load distribution and arbitration among network slices to meet the requirements of the E2E services for different verticals. Moreover, this orchestrator enables multi-domain service orchestration. Alvizu et al. [10] proposes Network Orchestration for Dynamic Network Slicing with DNS over MPLS tunnels. The work provides customisable slices with different degrees of service level agreement (SLA) in terms of bandwidth, delay and availability requirements.

Table I provides a comparison of all these different orchestrators in terms of the types of orchestration supported. Both horizontal and vertical orchestrations have been analysed. As can be seen in the table, the vast majority of the solutions either do not cover any type of vertical orchestration between different business roles or have mainly focused just on the

interactions between DSP and ISP. In particular, no existing solutions available in the literature has provided a complete integration among the three business roles ISP, CSP and DSP. At maximum, the vertical solutions proposed in other work only contain two of the three. The complexity of our research work presented in this paper is the integration of all these infrastructures and network slice provisioning controllers to provide effective service provisioning for all the business roles involved. Through this comparison, it is noted that our orchestration solution is a more complete and integrated approach in addressing the multiple business roles in both vertical and horizontal directions.

TABLE I
COMPARISON ORCHESTRATION SOLUTIONS

	Vertical Orchestration Support			Horizontal Orchestration Support			Technologies
	CSP ISP	DSP CSP	DSP ISP	ISP	DSP	CSP	
[3]	✗	✗	✗	✗	✓	✗	[11] [12]
[4]	✗	✗	✗	✓	✗	✗	[13]
[5]	✗	✗	✗	✓	✗	✗	[14]
[6]	✗	✗	✓	✓	✓	✗	[14] [15]
[7]	✓	✗	✗	✓	✗	✓	[16]
[8]	✗	✗	✓	✓	✓	✗	[11] [17]
[9]	✗	✗	✓	✓	✓	✗	[11] [17]
[10]	✗	✗	✓	✓	✗	✗	[18] [14]
Ours	✓	✓	✓	✓	✓	✓	[11] [17] [19] [20]

III. ARCHITECTURE FOR VERTICAL ORCHESTRATION FOR SLICE PROVISIONING OF 5G INFRASTRUCTURES

Figure 1 shows an overview of the different architectural elements needed for each of the different business roles to offer their services. Each business role has its own data, control and management planes. The ISP and The CSP are compliant with the ETSI MANO model. There is an infrastructure manager, Physical Infrastructure Manager (PIM) and Virtual Infrastructure Manager (VIM) (*see 6*) respectively, to control the life cycle of resources, labelled as NFPI and NFVI (*see 4*) respectively. There are network function managers, Physical Network Function (PNF) Manager and Virtual Network Function (VNF) Manager (*see 9*) respectively, to control the life cycle of the services deployed over such resources. These services deployed over the resources are represented by both PNFs and VNFs for the ISP and the CSP respectively (*see 10*). The management of such services is conducted by an Element Management System (EMS), named as P-EMS and V-EMS (*see 7*) for the ISP and the DSP respectively. The DSP also has a Service Manager (*see 9*), which shares the same responsibility of the VNF and the PNF Managers but has been labelled with a different name to make it explicit that it is under the control of the DSP. Similarly, the Service EMS (*see 7*) under the DSP has the same responsibility that both P-EMS and V-EMS have but has been explicitly labelled as S-EMS to indicate that it is controlled by the DSP. This is, in fact, a vertical orchestration between the Service

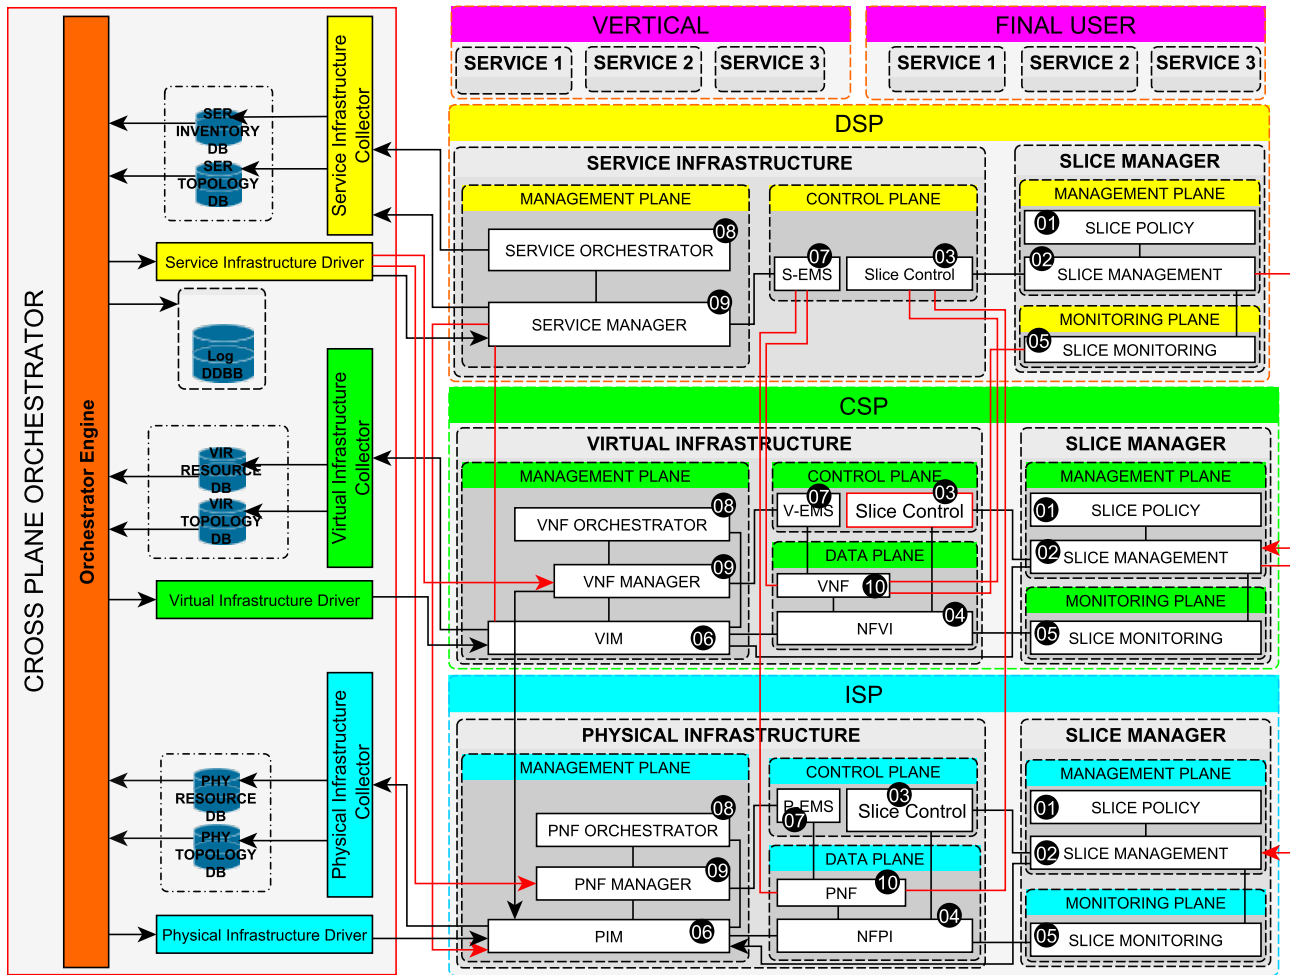


Fig. 1. Proposed architecture

Manager and the management of the resources running in both the ISP and the DSP. Each business role has an orchestrator that will allow automatic management of operations within the resources and services associated with the same business plane (horizontal orchestration).

Analogously, the architecture to allow network slicing is similar among the different business roles. They are composed of the following components. First, a Slice Policy module (*see 1*) to allow the configuration of network slices in the infrastructure. Second, a Slice Manager decides the coverage of the application of the network slices to allow the enforcement of the slices dynamically by interacting with the third component, the Slice Control component, which enforces the slices in the locations indicated by the Slice Manager. The Slice Control (*see 3*) acts as an actuator over the resources and services. It is noted that there is a Slice Management functionality residing not only in the ISP and the CSP but also in the DSP. Network slicing takes place in all the infrastructure layers including a slice of the service that is provided to the verticals and final users. In the case of a telecommunication provider acting as a DSP, this is a communication service that needs to be sliced, for example, among all the users of this

provider. Each of the different business roles contains as well a Slice Monitoring (*see 5*) module, which gathers metrics from resources and services to ensure each of the slices enforced meets the guaranteed service requested. i.e. ISP meets CSP's requirements, CSP meets DSP's requirements, and so on.

Furthermore, there is another vertical invocation between the Slice Management module in the DSP and the counterpart in the CSP. And the same happens between the one in the DSP and the one in the ISP. This vertical interaction is required to enforce a new configuration of a network slice in a cascade until it reaches the ultimate hardware and thus true enforcement of the slices is carried out. It is noted that the Slice Control component of the DSP (*see 3*) interacts directly with the services deployed in the CSP and the ISP in order to enforce the slicing in the resources. This is yet another vertical interaction between different business roles. Over such a complex infrastructure with several business models, a vertical orchestrator is proposed to allow the interactions of the management planes of the different business roles in an integrated way in order to achieve a significant level of automation in the provisioning of services and slices, thereby reducing operational costs. This vertical orchestrator

is depicted in the left part of Figure 1. The vertical orchestrator has two different interfaces to interact with each of the different infrastructures exposed by the business roles. There is no information exposed by the third party, these interfaces only send a request when resources are not enough. They share a similar logical architecture even though they are completely different in implementation details. Thus, each business role has a driver to allow the orchestrator to interact with the management interfaces and it has as well a collector to gather both resources/services and topological information. Thus, the orchestration engine will make use of this key information about resources and services to perform automatic tasks involving all the different management interfaces of the various business roles.

IV. EXPERIMENTAL RESULTS

A. Proof-of-Concept Implementation

The proposed vertical orchestration framework has been prototyped as a proof of concept using key technologies such as MaaS for the ISP management plane, and OpenStack, Neutron and OpenDayLight with slicing extensions for the slice management for the CSP management plane. Juju is employed for the PNF/VNF and Service manager with its integration with both MaaS and OpenStack. OSM is adopted for the CSP and the DSP.

B. Execution Testbed

A realistic virtualised 5G testbed has been established to empirically investigate the orchestration for resources and service provisioning for network slices. This testbed has been instrumented to measure the time consumed when the provisioning of new resources is made in each of the layered infrastructures during the vertical orchestration. The orchestration and provisioning of all the elements involved in the architecture have been started from bare-metal, i.e. from a wipe hard disk. The testbed has been created using 8 physical machines provides by the ISP and used by the CSP as managed computers, and each one has 8 cores, 32 Gbytes of RAM, and 4x1Gbps Ethernet NICs + IPMI Ethernet. Each physical machine hosts up to 8 virtual machines provided in the CSP infrastructure and used by the DSP. Therefore, the managed infrastructure consists of 50 virtual machines with 4Gbytes of RAM and 1 core. These machines are managed by another Intel Xeon Processor E5-2630 v4 with 32GBytes and 3x10Gbit Ethernet NIC acting as a management plane and where the orchestrator has been deployed as depicted in Figure 1.

C. Empirical Results

To empirically validate the proposed architecture, we have designed a set of experiments where the number of resources is increased. Experiments are executed 10 times. In Figure 2, x-axis shows number of resources in the CSP provisioned by the orchestrator. The number of physical machines associated to ISP is exactly the number of the resources in CSP divided by 8. Y-axis records the results on resource provisioning times




Fig. 2. Resources Provisioning Time

in minutes. Each of the stacked bars plotted in Figure 2 shows an experiment executed from bare-metal until the number of resources (5G NFV Services) are completely created in DSP infrastructure. The trends of the average and deviations are also displayed. Due to the number of resources that are involved in the experiment, it is shown the average and standard deviations for ISP, CSP and DSP respectively. STD differs due to increasing in the number of resources deployed at the same time. The time to deploy resources from bare-metal in the ISP is less than 2 minutes on average. The times for the CSP are between 8 to 10 minutes on average spent in spawning and allocating the VM image in the physical machine, and launching it. Times for provisioning the resources for the DSP have the same linear trend. Results show that the system is highly scalable and performs well, even in the most stressed scenario with 8 machines in the ISP, 50 Virtual Machines in the CSP and 50 Services in the DSP (one service on each VM), 45 EnB and 5 SPGW.

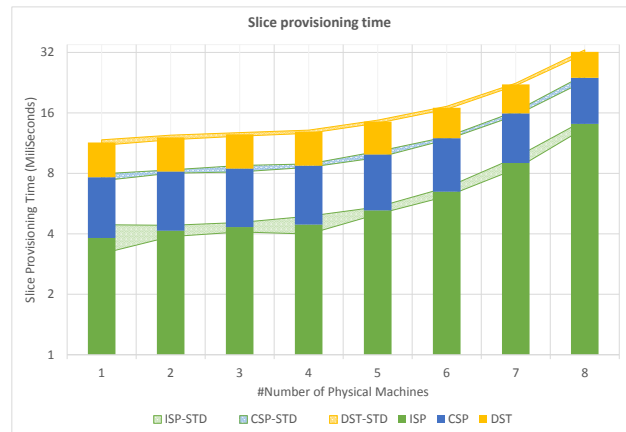


Fig. 3. Slices Provisioning Time

Figure 3 shows time consumption when the slice is provisioned in each of the infrastructures. X-axis shows the number of physical machines in each experiment, and y-axis shows average, together with the standard deviations, of the provisioning times in the experiments. As shown, times in the ISP are higher than those in the other two business roles (CSP and DSP). This is due to complexity and the number of rules required to enforce slicing in the data path for the ISP. Notice that when OpenStack creates VXLAN tunnels, it creates a full-mesh network interconnecting the corresponding number of physical machines (as shown in the x-axis) to all computers available to establish connections point to point with $N(N - 1)/2$ mesh links. It is worth highlighting that even in the most demanding case when we summarise all the times with 8 machines in the ISP, 50 Virtual Machines in CSP and 50 Services in DSP, the maximum slice provisioning time is only 31 ms, which is negligible. These superior results demonstrate the high scalability in service provisioning for network slicing.

V. DISCUSSIONS ON BUSINESS ROLES

The proposed integration among business roles is a significant enabler in automating the deployment of services covering all the different steps of the value chain across business roles. This inter-business architectural consideration enables a number of business roles. For example, the orchestrator can be owned by the vertical business to perform the complete automation of its businesses over the services and resources exposed by third parties. It can be also owned by another new business role who can provide such service to the vertical businesses. Alternatively, it can be hidden under each of the existing business roles (DSP, CSP and ISP) to enhance their functionality exposed to the final user. In addition, it can be used in scenarios where two or more of these existing business roles belong to the same administrative domain, thereby allowing even tighter integration of the information models across the different infrastructures.

VI. CONCLUSION

5G networking has fostered the emergence of new business roles and the value chains among them. This paper has proposed a novel inter-business orchestration framework to address the resource and service orchestration in this new paradigm. Architecturally, the required functional components across the various business roles' domains have been proposed with cross-roles' interactions specified in deploying resources and services in an integrated architecture. Experimental results have demonstrated the performance and the scalability of the proposed solution, which meets the requirement of 5G service creation time KPI in proof-of-concept implementation.

ACKNOWLEDGMENT

This work is funded in part by the European Commission Horizon 2020 5G-PPP Program under Grant Agreement Number H2020-ICT-2016-2/761913 (SliceNet: End-to-End Cognitive Network Slicing and Slice Management Framework in Visualized Multi-Domain, Multi-Tenant 5G Networks)

REFERENCES

- [1] E. Chirivella-Perez, J. Gutiérrez-Aguado, J. M. Claver, and J. M. A. Calero, "Hybrid and extensible architecture for cloud infrastructure deployment," in *Computer and Information Technology; Ubiquitous Computing and Communications; Dependable, Autonomic and Secure Computing; Pervasive Intelligence and Computing (CIT/IUCC/DASC/PICOM), 2015 IEEE International Conference on*. IEEE, 2015, pp. 611–617.
- [2] E. Chirivella-Perez, J. M. A. Calero, Q. Wang, and J. Gutiérrez-Aguado, "Orchestration architecture for automatic deployment of 5g services from bare metal in mobile edge computing infrastructure," *Wireless Communications and Mobile Computing*, vol. 2018, 2018.
- [3] S. Kukliński and L. Tomaszewski, "Dasmo: A scalable approach to network slices management and orchestration," in *NOMS 2018-2018 IEEE/IFIP Network Operations and Management Symposium*. IEEE, 2018, pp. 1–6.
- [4] C.-Y. Chang, N. Nikaiein, O. Arouk, K. Katsalis, A. Ksentini, T. Turletti, and K. Samdanis, "Slice orchestration for multi-service disaggregated ultra dense rans," *IEEE Communications Magazine*, 2018.
- [5] M. R. Raza, C. Natalino, A. Vidal, M. Santos, P. Öhlen, L. Wosinska, and P. Monti, "Demonstration of resource orchestration using big data analytics for dynamic slicing in 5g networks," in *2018 European Conference on Optical Communication (ECOC)*. IEEE, 2018, pp. 1–3.
- [6] A. Farrel, "Recent developments in service function chaining (sfc) and network slicing in backhaul and metro networks in support of 5g," in *2018 20th International Conference on Transparent Optical Networks (ICTON)*. IEEE, 2018, pp. 1–4.
- [7] L. Zanzi, V. Sciancalepore, A. Garcia-Saavedra, and X. Costa-Pérez, "Ovnes: Demonstrating 5g network slicing overbooking on real deployments," in *IEEE INFOCOM 2018-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 2018, pp. 1–2.
- [8] R. Montero, F. Agraz, A. Pagès, and S. Spadaro, "End-to-end 5g service deployment and orchestration in optical networks with qoe guarantees," in *2018 20th International Conference on Transparent Optical Networks (ICTON)*. IEEE, 2018, pp. 1–4.
- [9] X. Li, J. Mangues-Bafalluy, I. Pascual, G. Landi, F. Moscatelli, K. Antevski, C. J. Bernardos, L. Valcarengi, B. Martini, C. F. Chiasserini *et al.*, "Service orchestration and federation for verticals," in *Wireless Communications and Networking Conference Workshops (WCNCW), 2018 IEEE*. IEEE, 2018, pp. 260–265.
- [10] R. Alvizu, S. Troia, V. M. Nguyen, G. Maier, and A. Pattavina, "Network orchestration for dynamic network slicing for fixed and mobile vertical services," in *2018 Optical Fiber Communications Conference and Exposition (OFC)*. IEEE, 2018, pp. 1–3.
- [11] J. Quittek, P. Bauskar, T. BenMeriem, A. Bennett, M. Besson, and A. Et, "Network functions virtualisation (nfv)-management and orchestration," *ETSI NFV ISG, White Paper*, 2014.
- [12] A. Pras, B. J. F. Beijnum, and R. Sprenkels, *Introduction to TMN*. Centre for Telematics and Information Technology, 1999.
- [13] A. Ksentini and N. Nikaiein, "Toward enforcing network slicing on ran: Flexibility and resources abstraction," *IEEE Communications Magazine*, vol. 55, no. 6, pp. 102–108, 2017.
- [14] E. Haleplidis, K. Pentikousis, S. Denazis, J. H. Salim, D. Meyer, and O. Koufopavlou, "Software-defined networking (sdn): Layers and architecture terminology," Tech. Rep., 2015.
- [15] J. Halpern and C. Pignataro, "Service function chaining (sfc) architecture," Tech. Rep., 2015.
- [16] M. Corici, F. Gouveia, T. Magedanz, and D. Vingarzan, "Openepc: A technical infrastructure for early prototyping of ngmn testbeds," in *International Conference on Testbeds and Research Infrastructures*. Springer, 2010, pp. 166–175.
- [17] O. Foundation, "Openstack," 2015. [Online]. Available: <https://www.openstack.org/>
- [18] E. Rosen, D. Tappan, G. Fedorkow, Y. Rekhter, D. Farinacci, T. Li, and A. Conta, "Mpls label stack encoding," Tech. Rep., 2000.
- [19] Juju. [Online]. Available: <https://jujucharms.com/>
- [20] A. Sirbu, C. Pop, and F. Pop, "MaaS advanced provisioning and reservation system," in *Proceedings of the 1st International Workshop on Automated Incident Management in Cloud*, ser. AIMC '15. New York, NY, USA: ACM, 2015, pp. 13–18. [Online]. Available: <http://doi.acm.org/10.1145/2747470.2747473>