A reference architecture for network slicing

Sławomir Kukliński[*◇]
Lechosław Tomaszewski[*]
and Tomasz Osiński[*]
[*]Orange Polska, [◇]Warsaw University of Technology
Warsaw, Poland

Adlen Ksentini
and Pantelis A. Frangoudis
Eurecom
Sophia Antipolis, France

Eleonora Cau
and Marius Corici
Fraunhofer FOKUS
Berlin, Germany

***Abstract*—Network slicing is an important technology that influences the way in which new networking solutions will be designed and operated. Currently, there are several approaches to network slicing, but so far there is no a systematic approach that includes all aspects of the technology. In this paper, we present a generic network slicing framework that includes embedded in each slice some management and operations related mechanisms in order to cope with slice management scalability and to address the multi-tenancy issue. The presented solution uses NFV MANO for slice orchestration and supports multi-domain slicing.**

I. Introduction

Network slicing enables creation of many, fully fledged virtual networks over a common infrastructure while maintaining isolation between the slices. A network slice can be glued with applications. The key enabler of network slicing is ETSI Network Functions Virtualization (NFV) technology [1]. Network slicing was introduced by PlanetLab in 2002 [2] and more recently reinvented by NGMN [3]. It brings three main benefits: (i) dynamic deployment of networking solutions with short time to market and low CAPEX; (ii) ability to create networks that are tightly coupled with their service(s); (iii) delegation of almost complete network slice management to a slice tenant (a vertical).

In this paper we propose a reference framework for network slicing, addressing such issues as slice related operations, management and orchestration in a single and multi-domain environment. Our work is based on an exhaustive analysis of different network slicing approaches and relevant standardisation efforts (§ II). We have tried to identify key commonalities of different approaches as well as gaps. The features, requirements and design assumptions that drive our vision are described in § III. § IV presents the overall architecture, § V focuses on operations, management and orchestration issues in single domain sliced networks and § VI depicts extensions related to multi-domain slicing. § VII concludes the paper.

II. Related work

In the NGMN approach [4], a *Network Slice Instance (NSI)* is built over physical or logical resources (computation, storage and transport) that are fully or partially isolated from other resources. The NSI is built using *Network Functions* and is a complete, instantiated logical network that meets certain characteristics as required by a *Service Instance(s)*. The NSI can be composed of many *Sub-network Instances (SNIs)* that are similar to NSI, but do not have to form a complete logical network. Multiple NSIs may share the SNIs. In this concept, the service instances are separate to network slice instances.

NFV has no special support for network slicing, but this is the key technology commonly used in most network slicing solutions. In [5] it has been proposed to add to the NFV architecture a slice specific entity (Slice Controller), as a part of the OSS/BSS domain. It communicates with MANO via the Os-Ma-nfvo reference point, beyond which everything relies on standard NFV concepts and procedures. ETSI NFV group in Release 3 plans to address some network slicing requirements like scalability and multi-tenancy.

The IETF has initiated works on network slicing [6], but so far no Working Group on the topic has been formed. The IETF works concern the requirements, architecture and the adaptation of IETF protocols to network slicing needs.

The 3GPP concepts DÉCOR [7] and eDÉCOR [8], applicable to 4G networks, allow for the deployment of several Evolved Packet Cores (EPC) sharing the same RAN, but only one network slice can be used by the end-user. The 5G network slicing by 3GPP splits slices into planes (user/control) and divides the control plane functions into common and slice-specific [9]. Common control plane functions include slice selection, authentication and mobility management. A benefit of the common functions is the terminal's ability to be attached to several slices simultaneously. The 5G System [10], [11] incorporates the New Radio (NR), which enables RAN slicing [12]. The 3GPP has started working on several standards documents which deal with 5G network slice management [13], provisioning [14], selection [15] and security [16]. Recently finalized 3GPP studies, [17], [18] and [19], address network slice management and orchestration.

In the 5G PPP research program, several projects focus on network slicing. They commonly use NFV and SDN technologies. The 5G PPP project 5G NORMA II [20] proposes a mobile network architecture for the multi-tenant environment. NORMA assumes that some control plane functions should be common for multiple slices. The 5GEx project [21] focuses on transport networks only. It targets multi-domain slicing by extending the ETSI NFV MANO framework towards cross-domain orchestration and management.

III. Requirements and design assumptions

To make the network slicing framework generic as well as compatible with other approaches we have selected the following requirements that it has to fulfill:

This work has been realized in the framework of the EU-Japan project 5G!Pagoda under grant agreement 723172.

978-1-5386-3416-5/18/$31.00 ©

- *Slice isolation*. Slice isolation may be full or partial, logical or physical [4]. Limited isolation capabilities of some solutions cause the relaxed isolation requirements. Isolation concerns both slice resources and slice oriented operations (including management and orchestration).
- *Slice creation rights*. The orchestrator operator, 3rd parties or end-users should be able to create a slice.
- *Support for per-plane slicing*. Per plane slicing should be supported. In some cases, it is an implementation constraint, whereas keeping some functions as common (e.g. authentication, handover) is reasonable and can make slices lightweight and faster to provision.
- *Mapping of services to slices*. Slices should support single or multiple services. In the latter case, the service instance lifetime management should be separated from slice lifetime management.
- *Sub-slices*. A slice can be created as a combination of existing or on-demand created sub-slices.
- *Multi-domain slicing*. Slices should cross technological or administrative domain borders. For that purpose, the slice should be composed of sub-slices.
- *Slice provisioning time*. Some on-demand created slices may require short provisioning time.
- *3GPP compatibility*. Compliance with 3GPP slicing approaches (dedicated slices selection mechanisms, not sliced RAN of (e)DÉCOR, common control plane functions of the 5G Core) is required.
- *Legacy solutions inclusion*. Legacy mobile networks will coexist with 5G networks. Incorporation of migrated legacy systems in the overall picture is highly desirable.
- *Grouping of slices*. A group of service-oriented slices can be, e.g., used by the MVNO. The grouping should impact the access rights to slices (no additional authentication, etc.).
- *Slice management rights*. The slice operator/owner (3rd party) should have management capabilities that include: policy-based management, configuration, security operations, accounting and performance monitoring.
- *Management scalability*. As the number of possible slices can be high, management operations should be automated as much as possible, and slice management should be scalable.

IV. The overall architecture

The proposed approach for a single orchestration domain is presented in Fig. 1. The overall management and orchestration of slices are distributed into several functional blocks. The *Global OSS/BSS* is a logically centralised entity that drives the behaviour of the entire system. The NFV MANO compliant orchestration is used in the concept without modifications, following the arguments presented in [5].

A slice (or a sub-slice) in our concept is composed of three groups of functions. The first group, *Sliced Network (SN)*, is the same set of functions of the non-sliced network. The second group, *Slice Operations Support (SOS)*, supports operations related to slice selection, subscription, user authentication, and stitching of sub-slices of different domains to obtain the end-to-end slice. The third group, *Slice Manager (SM)* performs slice management. The SM allows for slice management by its tenant.

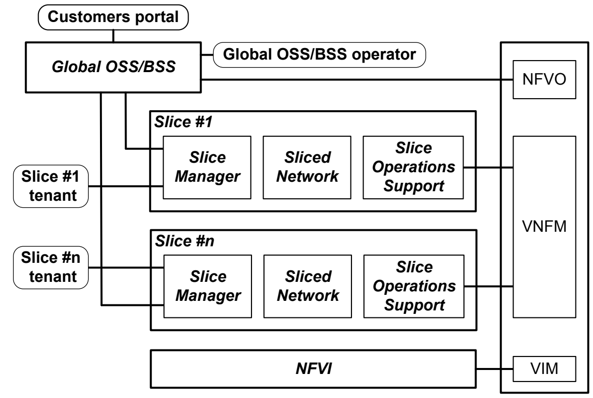

Fig. 1. System architecture in a single orchestration domain

A. The Sliced Network

Each Sliced Network is composed of data, control and application planes (see Fig. 2) that realise the same functions as in the non-sliced case (e.g. EPC). Such split provides high flexibility of slicing by enabling independent slicing at each plane level. Moreover, it provides compatibility with existing approaches, like (e)DÉCOR.

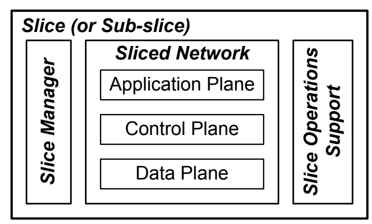

Fig. 2. Slice or sub-slice generic structure

The approach allows for sharing of functions of a specific plane by several or all slices of the domain, whereas other planes can be fully based on slice-dedicated functions. We have decided to group the shared/common functions into a single domain sub-slice called *Common Sub-Slice (CSS)* and group the functions that are dedicated for a specific goal into the *Dedicated Sub-Slice(s) (DSS)*. The DSSes can access CSSes services through the CSS APIs. Both sub-slices "stitched vertically" form a single Dedicated Slice. It is expected that in a single domain there will be a single CSS and multiple DSSes. An example of a slice composed of CSS and a single DSS is presented in Fig. 3.

The CSS is generally optional, but in some cases its existence may be enforced by system limitations, whereas in other cases it is a solution of choice. For example, there is no way to avoid CSS in legacy solutions, and the use of the CSS to handle mobility and authentication of users attached to several DSSes is reasonable. The life-cycles of CSS and DSSes are separated – from the DSS point of view, the CSS is a permanent slice. As CSS functions are used, the DSSes may have a smaller footprint and can be provisioned faster, what is of premium importance for on-demand created slices. The

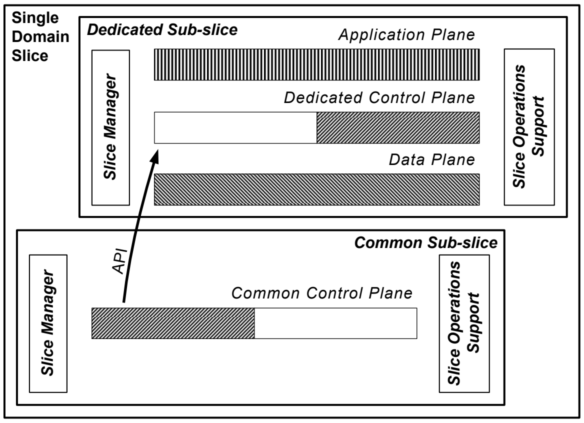

Fig. 3. Common and Dedicated Sub-slices concatenation (example)

CSS can be used for the definition of the common, lightweight control plane with functions can be exploited by the dedicated slices. The set of CSS functions should not be fixed and it should be possible to add new functions to the running CSS. A combination of CSS and DSS can also be seen as a new CSS sub-slice. That way the concept of CSS-DSS combination can become recursive. It has to be noted that in contrast to 3GPP [17] we do not allow for the existence of individual network functions that are shared between several DSSes. The reason is an additional complexity of the management of such individual functions. The proposed solution is to merge such individual, new functions with the existing CSS.

There is no doubt that the existence of common functions (i.e. CSS) also raises some problems. The limitations of CSS services can negatively impact the overall functionality and flexibility of a CSS/DSS slice. Moreover, the separation between DSSes is also reduced due to the use of CSS.

B. Slice Operations Support

A multi-slice environment requires a set of operations not present in non-sliced networks. From the user perspective, these operations should include network slice discovery, selection, subscription and authentication. Slice chaining (creation of the end-to-end slice from "horizontally" stitched sub-slices) also requires mechanisms like exposing the abstracted view of a single domain slice to other domains and supporting inter-slice operations. These functionalities are grouped within the *Slice Operations Support (SOS)* block, deployed together with the Sliced Network by enhancing its Blueprint by SOS functions. These enhancements should be tailored to the needs of the horizontally or vertically stitched sub-slices.

The detailed definition of the slice selection mechanism that has to be combined with the slice description scheme (a part of SOS) is out of the scope of this paper. However, several options should be supported. For example, slice selection can be hard-coded (e.g. in case of IoT devices), done in a way completely agnostic to the end-user by the network, or based on a negotiation process between the network and the end-user. The latter may include end-user preferences (price, performance, etc.). Moreover, there should exist a "default" slice to be used if a proper match of a dedicated slice cannot be found. It should be also possible to create the on demand slice. In the latter case, the default slice should provide the basic communication and services until the requested slice is created. Existing legacy, non-sliced systems, e.g. the LTE network, can be nicely used as a default slice. Their use in the described way allows for graceful migration from non-sliced to sliced solutions.

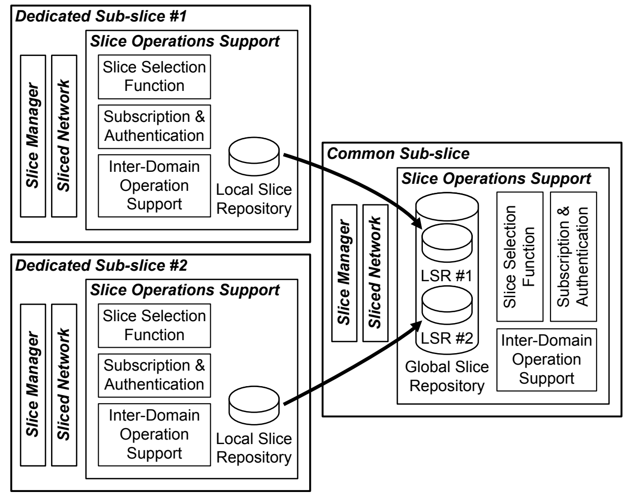

Fig. 4. Functional entities of the Slice Operations Support and relationships between Slice Repositories of CSS and DSSes (example)

Fig. 4 presents an example of slice selection functionality of SOS split between CSS and DSSes. Initial slice selection is performed by the CSS *Slice Selection Function (SSF)* that is combined with a *Global Slice Repository (GSR)*. Each DSS may have a local, optionally modified copy of GSR, the *Local Slice Repository (LSR)*. Slice selection scalability and reliability justify the existence of SSF and LSR as a part of DSS. Moreover, the LSR may have a subset of GSR slices only, e.g. only the slices that are operated by the same tenant (cf. MVNO).

The multi-domain slicing requires a set of operations that support the exchange of information between the slices of neighbouring domains to provide the end-to-end activity. Such functionality is placed in the *Inter-Domain Operations Support (IDOS)* block of SOS (Fig. 4). It provides exposure of a slice to the neighbouring domains (e.g. topology abstraction) and supports protocols that enable information exchange between the domains. More details about SOS functionality in the multi-domain environment will be provided in Section VI.

C. Slice Manager

The *Slice Manager (SM)* is an element of a slice or sub-slice acting as its management entity (see Fig. 2). It is only a part of the overall management system, cooperating with the main management platform – the Global OSS/BSS, which plays the master role in the overall management and orchestration picture (cf. Fig. 1).

Two reasons justify the inclusion of the SM into a slice. The first one is scalability–the management of multiple, functionally isolated slices by a single management system is not scalable and raises problems related to the separation of management operations between slices. The second is business oriented. It is commonly assumed that slices will be tailored to the needs of specific services that can be offered by 3rd parties (verticals), being slice tenants. Each tenant should have some management capabilities to operate its slice(s) efficiently.

The SM plays the role of the management plane of the Sliced Network, but it is slightly different from the management plane of the non-sliced networking solution. The slice tenant (typically not a professional operator) obtains with the SM a simple and comfortable management of the slice via a dedicated interface. Moreover, the SM handles slice (or sub-slice) faults and performance. Tenants request creation of a slice via the *Global OSS Tenants Portal* (see Fig. 6), but then they use SM for most slice oriented operations, except those that are related to slice life-cycle management and accounting.

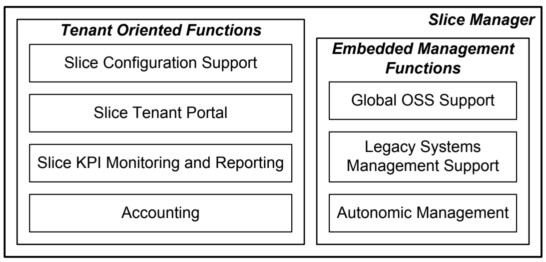

Fig. 5. Functional entities of Slice Manager

The internal architecture of the SM is shown in Fig. 5. The SM is composed of the *Tenant Oriented Functions* and the *Embedded Management Functions (EMF)*. Tenant operations are performed via the *Slice Tenant Portal* with the involvement of the *Slice Configuration Support* entity and enable modification of slice runtime policies, configurations of slice services and subscriptions of users. The *Slice KPI Monitoring and Reporting* entity provides the tenant information about its slice health and usage, and with a combination of the *Accounting* component allows for charging of slice users. Accounting information is additionally stored in the Global OSS/BSS, because the SM of a slice disappears on slice termination.

To make slice management lightweight and comfortable, most of the tenant's management operations (i.e. EMF) should be automated, providing the overall management system scalability and short response time for events, alerts, etc. For that purpose, the use of the cognitive [22] or autonomic [23] network management paradigms is expected. If legacy systems with their own management are used, the EMF should include support for legacy management functions.

V. Operations, Management and Orchestration of Slices

We first describe management and orchestration for the single orchestration domain case. We exclusively assume NFV MANO compliant orchestration. Thus, orchestration is slice-agnostic, while slice management is slice (or sub-slice) specific. Some details about the scalability of the proposed approach can be found in [24].

A. Single domain slice related operations

For the single orchestration domain, the operations provided by SOS have already been described in § IV. In such case, the IDOS entity of SOS does not need to be implemented.

B. Single domain management and orchestration

Single domain slice management and orchestration is shown in Fig. 6. The management is split between SMs and the Global OSS/BSS. The role of the SM in case of a single domain has been already described. The Global OSS/BSS acts as a master and also drives the MANO compliant orchestration. The first group of Global OSS/BSS functions includes

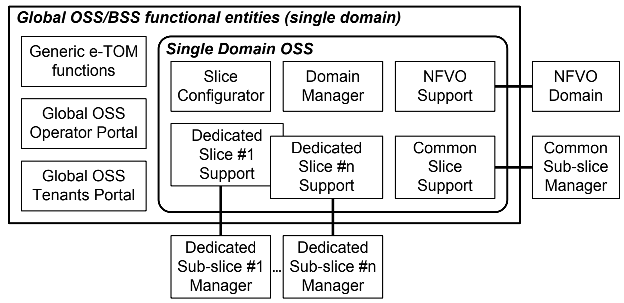

Fig. 6. Management and orchestration architecture in case of single orchestration domain

generic eTOM functions and portals for orchestrator operator and tenants. The tenants use the Global OSS Tenants Portal (see Fig. 6) for requesting slice creation and termination, and for accessing current and historical data related to the slice.

The second group of Global OSS/BSS functions (*Single Domain OSS*) drives MANO orchestration of slices and provides a single domain slice management and orchestration. The *Slice Configurator* analyses slice requests, blueprints as well as users' and operators' policies. In cooperation with the *NFVO Support* entity, which keeps the catalogue of network slices, it creates the Network Slice Description that will be used by the NFVO for slice deployment. At the same time, the *Slice Support* entity is created, which handles the cooperation with this slice's SM. The *Domain Manager* keeps information about all resources available in the domain, their usage, their allocation to slices, as well as overall information about alerts, etc. It can also do the arbitrage of resource allocation to slices based on their priorities. The *Common Slice Support* may include legacy (i.e. hardware-based) management systems if they are not implemented as an SM part of the CSS.

As already stated, we use the ETSI MANO as it is. However, MANO orchestration scalability can be also a problem. This can be partly solved by using multiple VNFMs and NFVOs. The latter case requires the existence of multiple VIMs sharing the same infrastructure and dynamic allocation of resources between them. This necessitates a modification of ETSI NFV recommendations, which is outside the scope of this paper.

VI. Multi-domain slicing

A single orchestration domain's slices, treated as sub-slices, can be horizontally stitched to form an end-to-end slice. For example, RAN and EPC in LTE networks can be stitched to form a complete mobile network. Additionally, the transport domain (able also to serve as a backhaul) could be integrated. The creation of multi-domain slices raises new issues related to slice operations, management and orchestration.

A. Multi-domain operations

The horizontal stitching of several slices impacts the implementation of SOSes. It specifically concerns the slice selection oriented operations and IDOS functional entities (both functions are a part of SOS). Slice selection capabilities should be implemented in the "edge sub-slices" only. The IDOS entities act as inter-slice gateways that implement an exchange of the information between neighbouring domains to provide their efficient cooperation. In this context, IDOS should expose an abstracted view of its domain as well as enable inter-domain communication via supporting necessary protocols. An example of such functionality, however not labelled as slice level operations, can be found in [21]. The SOS configuration of each sub-slice (or single domain slice) should be done during the multi-domain slice provisioning by each domain Slice Configurator (part of the Single Domain OSS).

B. Multi-domain management and orchestration

The multi-domain management and orchestration is shown in Fig. 7. In comparison to the single-domain case, the Global OSS/BSS has now several Single Domain OSSes and two new entities, namely *Multi-Domain Slice Configurator (MDSC)* and *Multi-Domain Manager (MDM)*.

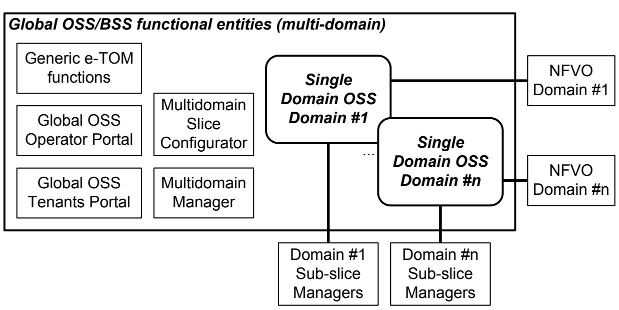

Fig. 7. Multi-domain management and orchestration architecture

The MDSC translates the business requirements (described in the sub-slice or slice blueprint) into technology-specific ones. Then, in cooperation with Slice Configurator of each domain, it configures each domain SOS for proper inter-domain operations as described in the previous section.

The MDM interacts with each Domain Manager to solve multi-domain management issues. Such information exchange could be beneficial for end-to-end optimisation. According to the "domain oriented philosophy", though, exchanging management information between domains should be minimised.

In our concept, we have avoided a direct orchestration of multi-domain slices. Instead, we propose to create the end-to-end slices as a chain of single domain slices that are orchestrated by a single orchestrator. Therefore there are no multi-domain orchestration issues.

VII. Conclusions

In this paper, a top-down approach to network slicing has been presented. We have defined our framework in a systematic way, according to the requirements presented in § III. As per our design, each slice has functions responsible for slice operations and management. The slice management component provides a dedicated management interface to slice tenants. The slice operations support component supports slice selection and multi-domain slicing. The proposed slice structure is generic, allows for flexible, per plane slicing as well as for horizontal or vertical slice stitching. Multi-domain slices are defined as a concatenation of single domain slices, thus avoiding multi-domain orchestration issues. However, even in a single domain, the scalability of MANO orchestration can be an issue, and this problem is yet to be solved.

References

[1] "Network Functions Virtualisation (NFV); Architectural Framework", ETSI GS NFV 002, V1.2.1, Dec. 2014.
[2] "PlanetLab", http://www.planet-lab.org/.
[3] "NGMN 5G White Paper", NGMN Alliance, white paper, Feb. 2015.
[4] "Description of Network Slicing Concept", NGMN Alliance, Sep. 2016.
[5] B. Chatras, S. Tsang Kwong U, N. Bihannic, "NFV Enabling Network Slicing for 5G", in Proc. ICIN, 2017.
[6] A. Galis (ed.), S. Kukliński et al., Network Slicing - Revised Problem Statement (draft-galis-netslices-revised-problem-statement-03), Internet-Draft, informational, Mar. 2018.
[7] "Architecture enhancements for dedicated core networks; Stage 2", 3GPP TR 23.709, v13.0.0, Dec. 2014.
[8] "Enhancement of Dedicated Core Networks selection mechanism", 3GPP TR 23.711, v14.0.0, Sep. 2016.
[9] "Study on architecture for next generation systems", 3GPP TR 23.799, v14.0.0, Dec. 2016.
[10] "System Architecture for the 5G System", 3GPP TS 23.501, v15.0.0, Dec. 2017.
[11] "Procedures for the 5G System", 3GPP TS 23.502, v15.0.0, Dec. 2017.
[12] "NR; Overall description; Stage-2", 3GPP TS 38.300, v15.0.0, Jan. 2018.
[13] "Management of network slicing in mobile networks; Concepts, use cases and requirements", 3GPP TS 28.530, v0.5.0, Feb. 2018.
[14] "Provisioning of network slicing for 5G networks and services", 3GPP TS 28.531, v0.3.0, Feb. 2018.
[15] "5G System; Network Slice Selection Services; Stage 3", 3GPP TS 29.531, v1.0.0, Mar. 2018.
[16] "Study on security aspect of 5G network management slicing", 3GPP TR 33.811, v0.4.0, Jan. 2018.
[17] "Study on management and orchestration of network slicing for next generation network", 3GPP TR 28.801, v15.0.0, Jan. 2018.
[18] "Study on management and orchestration architecture of next generation network and service", 3GPP TR 28.800, v15.0.0, Jan. 2018.
[19] "Study on management aspects of next generation network architecture and features", 3GPP TR 28.802, v15.0.0, Jan. 2018.
[20] B. Sayadi et al., "SDN for 5G Mobile Networks: NORMA perspective", in Proc. CrownCom, 2016.
[21] R. Guerzoni et al., "Multi-domain Orchestration and Management of Software Defined Infrastructures: a Bottom-Up Approach", in Proc. EuCNC, 2016.
[22] http://www.cognet.5g-ppp.eu/.
[23] "Generic Autonomic Network Architecture (An Architectural Reference Model for Autonomic Networking, Cognitive Networking and Self-Management)", ETSI GS AFI 002, v1.1.1, Apr. 2013.
[24] S. Kukliński, L. Tomaszewski, "DASMO: A Scalable Approach to Network Slices Management and Orchestration", in Proc. 3rd IEEE/IFIP 5GMan (in conjunction with IEEE/IFIP NOMS), 2018.