

Date of publication xxxx 00, 0000, date of current version xxxx 00, 0000.

Digital Object Identifier XX.XXXX/ACCESS.XXXX.DOI

A Future-proof Architecture for Management and Orchestration of Multi-domain NextGen Networks

VINCENZO SCIANCALEPORE¹ (Member, IEEE), CHRISTIAN MANNWEILER², FAQIR ZARRAR YOUSSEF¹, PABLO SERRANO³ (Senior Member, IEEE), MARCO GRAMAGLIA³, JULIE BRADFORD⁴ and IGNACIO LABRADOR PAVÓN⁵

¹NEC Laboratories Europe, Heidelberg, Germany (e-mail: {vincenzo.sciancalepore,zarrar.youssaf}@nec-lab.eu)

²Nokia Bell Labs, Munich, Germany (e-mail: christian.mannweiler@nokia-bell-labs.com)

³University Carlos III of Madrid, Leganés, Spain (e-mail: {pablo.mgramagl}@it.uc3m.es)

⁴Real Wireless, London, United Kingdom (e-mail: julie.bradford@real-wireless.com)

⁵ATOS, Madrid, Spain (e-mail: ignacio.labrador@atos.net)

Corresponding author: Marco Gramaglia (e-mail: mgramagl@it.uc3m.es).

This work has been performed within the 5G-MoNArch project, part of the Phase II of the 5th Generation Public Private Partnership (5G-PPP) program partially funded by the European Commission within the Horizon 2020 Framework Program. The 5G-MoNArch project builds on the results of the 5G-PPP Phase I project 5G-NORMA.

ABSTRACT The novel network slicing paradigm represents an effective turning point to operate future wireless networks. Available networking and computational resources may be shared across different (instantiations of) services tailored onto specific vertical needs, envisioned as the main infrastructure tenants. While such customization enables meeting advanced Key Performance Indicators (KPIs) introduced by upcoming 5G networks, advanced multi-tenancy approaches help to abate the cost of deploying and operating the network. However, the network slicing implementation requires a number of non-trivial practical considerations, including e.g. (i) how resource sharing operations are actually implemented, (ii) how involved parties establish the corresponding agreement to instantiate, operate and terminate such a sharing or, (iii) the design of functional modules and interfaces supporting these operations.

In this paper, we present a novel framework that unveils proper answers to the above design challenges. While existing initiatives are typically limited to single-domain and single-owner scenarios, our framework overcomes these limitations by enlarging the administrative scope of the network deployments fostering different providers to collaborate so as to facilitate a larger set of resources even spread across multiple domains. Numerical evaluations confirm the effectiveness and efficiency of the presented solution.

INDEX TERMS 5G mobile communication, Computer network management, Network architecture, Network function virtualization

I. INTRODUCTION

THE fifth generation (5G) of wireless and mobile communication networks needs to natively support multiple advanced services. This requirement significantly deviates from the legacy approach, which is built on the one-fits-all paradigm: the same telecommunication services (both voice and data) are provided to any kind of customer, regardless of the mobile applications that are actually carried over the network. This monolithic view is no longer sufficient when dealing with a large ecosystem of use cases as the one currently envisioned for 5G (that could be extended with novel services).

5G networks should fulfill a set of requirements that are

identified by Key Performance Indicators (KPIs) such as, e.g., high data rates, extremely low latency, robustness, and reliability. To efficiently address such a stringent set of KPIs, a key-technology has been proposed, namely *network slicing*.

A network slice is a set of physical and logical resources (e.g., radio, transport and cloud/edge resources) gathered from a common pool and assigned to an infrastructure tenant to deliver some specific services [1]. Such services could be, e.g., enhanced/extreme Mobile Broadband (e/xMBB) for “regular” end-users, or high-reliable and ultra-low latency services for autonomous driving [2]. The potential advantages brought by network slicing are evident [3], given its

ability to provide highly customizable services over the same shared infrastructure and, in turn, creating new revenue opportunities for the involved stakeholders.

This ability to customize services is attracting a lot of attention from researchers (both academic and industrial). In particular, one key challenge is the design of mechanisms to efficiently chain different network functions at different locations while fulfilling the target KPIs [4], [5]. But there is (at least) another key aspect that has been overlooked so far: *the analysis of diverse agreements among multiple stakeholders that need to collaborate to bring network slicing into practice*. Some of these stakeholders are network operators (from one or multiple domains), infrastructure providers, service providers, and tenants, which shall agree in terms of control and management of the resources involved to instantiate a service.

In this paper, we present a novel network management and orchestration architecture that fully supports the customization of network slices across services and tenants. While our framework is *backward compatible* with current standardization efforts (namely, ETSI NFV MANO [6]), it further introduces a number of key features that enable the development of a series of use cases of economic relevance, as we discuss next. These main features our architecture are:

- A formal definition of the different actors of the ecosystem, along with the various roles they can play as well as different agreements that can be reached to support and deliver intended services.
- Efficient support for multi-domain environments. Instead of complex peer-to-peer interactions, our architecture enables a clean coordination between systems, thus providing a versatile framework to e.g. deploy new services over multiple domains, or extend traditional services using the novel infrastructure.
- Design of a novel element, namely Inter-Slice Resource Broker, to enable an efficient sharing of resources across slices while guaranteeing the reservation of resources as specified in the Service Agreements across actors.

The remainder of the article is organized as follows. Section II motivates the need for network slicing from an economic standpoint. Section III formally introduces the different actors and roles of the ecosystem while illustrating the flexibility of the envisioned solution via two representative and timely use cases. Section IV provides a brief overview of the current solution for network slicing and identifies its limitations for the case of multi-domain scenario. Section V presents the proposed architecture, which is composed of both novel modules and extensions over the ETSI NFV MANO framework thereby describing the lifecycle of a network slice. Section VI empirically assesses the improvements introduced by our proposal, and Section VIII provides concluding remarks.

II. ECONOMIC MOTIVATION FOR NETWORK SLICING

As discussed before, there are clear technical advantages that network slicing will bring into next generation networks.

Here we discuss both qualitative and quantitatively another dimension in which network slicing can provide significant benefits, namely, the economic axis. Indeed, network slicing can contribute to address some of the revenue/cost issues that current network deployments face, which are discussed next.


We first summarize the main economic motivations for the development of the proposed architecture and associated technologies. Fig. 1 presents an analysis of historical mobile data volumes and revenues for the UK (based on [7]). The figure shows a slight downward trend in prices paid by mobile subscribers, driven by competition in the UK mobile market and limits on consumers' willingness to pay. This is despite subscribers receiving much higher data volumes through their subscriptions over time, and has created a steep downward trend in the mobile revenue received per Gigabyte (GB) of data delivered. Alongside this, subscriber volumes have grown only marginally with headline retail mobile revenues in the UK falling by 8%, i.e., from GBP 17 billion to GBP 15.6 billion between 2012 and 2017.

This difficulty in growing or maintaining revenues from consumer-focused services puts pressure on reducing network costs to maintain margins. However, with increasing volumes of data being consumed on mobile networks, the opportunity to reduce costs has been limited. In fact, according to GSMA [8], for the period 2010–2017 the Capital Expenditure (CAPEX) investment from European MNOs has been at a rate between 12% and 18% of revenues.


The adoption of the architecture proposed in this paper has the potential to both improve margins and de-risk the long-term business case for providers:

- **Increasing the revenue per GB**, by developing relationships with end users and tenants who place a higher value than regular consumers on tailored mobile services with a guaranteed quality of service level.
- **Reducing the cost per GB**, compared with many disparate and dedicated private networks, due to economies of scale and scope of combining multiple services on the same network.

The two above effects have been investigated in [9], where we consider a multi-service virtualized network delivering a range of services in a central London study area. This analysis considers initially the costs and revenues associated with a “business as usual” scenario over a 2020 to 2030 time period. The network costs are based on making use of existing sites in the area (representative of a typical existing MNO), and then either (i) adding antennas and/or frequency bands to these, or (ii) choosing to build new sites, depending on which approach is most cost-effective to serve the increasing. Tables 1 and 2 present an extract of the reported revenues, costs and resulting discounted cash flow for the years 2020, 2025 and 2030 with the accumulated total over the 2020–2030 time period given in the final column. Comparing the total income against costs for the entire 2020 to 2030 period gives the Net Present Value (NPV), which is reported in Table 3 and can subsequently be translated into a Return on Investment (ROI) for the period.



(a) UK monthly mobile data volume and revenue per subscriber.



(b) UK MNO revenue per GB trends

FIGURE 1. UK data, cost and revenue trends [9]

TABLE 1. V2I cost and revenue results combined (2020-2030), middle scenario, £ million, Central London

	2020	2025	2030	Total
Full costs: eMBB plus service as indicated				
eMBB only	11.85	15.74	31.19	205.1
Semi-automated driving	11.85	15.85	33.15	204.2
Assisted driving	11.85	16.79	34.16	210.0
Vehicle infotainment	12.18	17.97	35.68	221.7
Full revenues: eMBB plus service as indicated				
eMBB only	17.04	19.60	21.88	215.2
Semi-automated driving	17.04	20.46	27.98	234.1
Assisted driving	17.18	20.59	26.10	232.3
Vehicle infotainment	17.29	21.93	28.95	247.8

TABLE 2. V2I cumulative discounted cash flow, middle scenario, £ million, Central London

	2020	2025	2030
eMBB only	5.19	23.29	14.81
eMBB & Semi-automated driving	5.19	26.05	26.84
eMBB & Assisted driving	5.33	24.40	22.16
eMBB & Vehicle infotainment	5.11	24.05	24.20

The analysis in [9] also aims at understanding the business impact of delivering more bespoke services alongside existing consumer mobile broadband services, using the same network infrastructure set (via network slicing). To this aim, the analysis considers the provision of vehicle to infrastructure (V2I) services in combination with eMBB services over a virtualized multi-service 5G network. This analysis includes:

- Deriving revenue forecasts for the V2I services considered, based on the benefits derived by end users and hence their willingness to pay. For example, the impact on insurance premiums were considered for some of the V2I services in this context.
- Deriving the additional network CAPEX and opera-

TABLE 3. V2I summary of financial measures, middle scenario, £ million, Central London

	NPV £~million	Indicative ROI
eMBB only	14.81	5%
eMBB & Semi-automated driving	26.84	15%
eMBB & Assisted driving	22.16	11%
eMBB & Vehicle infotainment	24.20	12%

tional expenditure (OPEX) beyond the planned “business as usual” eMBB network, to accommodate the envisaged V2I services in terms of their capacity and coverage requirements. This included developing service definitions and demand forecasts for the three candidate V2I services considered: (i) infotainment, (ii) assisted driving (non-critical traffic, navigation, maintenance etc. information updates to drivers), and (iii) semi-automated driving (critical real time updates on upcoming hazards).

The resulting impact on revenue and costs of offering one of the three V2I services considered in combination with the eMBB baseline case is shown on Table 1, with the resulting discounted cash flows on Table 2. Finally, the impact on total revenues against costs over the 2020 to 2030 is reported in the NPV and ROI values given in Table 3. These show up to a 10% improvement in ROI compared to the baseline eMBB case. While only a limited set of services were considered in this example, this is indicative that extending existing mobile networks to deliver a wider range of more bespoke mobile services via network slicing could help to improve the long-term business case challenges of reducing margins that the mobile industry currently faces. This is a benefit that will likely increase with the number of services combined on the same network. The same study also showed significant socio-economic benefits from delivering a wider range of services than currently, such as smart metering and V2I services, from mobile networks.

III. ROLE, BUSINESS MODELS, AND USE CASES OF THE 5G ECOSYSTEM

In this section we describe the potential use cases that are enabled by the architecture proposed in this paper. First, we describe the actors in our architecture and their main interactions.

A. ACTORS AND ROLES

An **actor** is an individual or institution playing a role in our ecosystem. Actors could therefore be e.g. a telecommunications company, a (medium or large scale) HW and/or SW provider or an enterprise from a vertical industry. The **role** is the function assumed by the actor in the corresponding

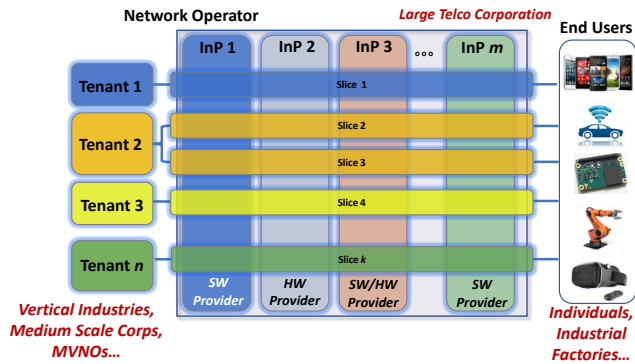


FIGURE 2. Actors and roles in a network slicing scenario.

scenario, e.g., providing the infrastructure, operating the network slice or consuming the service. As we discuss next, an actor could play different roles.

We illustrate the above concepts in Fig. 2, which represents a scenario where a network operator is providing different communication services to a number of tenants via a number of slices. We use red/italics to identify the different actors, while black/bold is used to identify the different roles (we formally introduce the different roles next). The “telecommunications company” acting as “Network Operator” could be, in the current ecosystem of enterprises, a large telco corporation (e.g., the one operating at national scale). This operator would run a number of slices over a heterogeneous infrastructure of HW and SW elements, which could be provided by HW/SW vendors. Individuals could play as “end users” (e.g., using a mobile phone), but also certain companies could act as the final consumers of a service (e.g., to monitor a fleet of cars or robots). Mobile Virtual Network Operators (MVNOs) or companies from vertical industries could leverage on slices provided by the network operator to create their services, where each slice would contain HW/SW infrastructure from the different available infrastructure providers, depending on the needs for each one.

Based on the above description, we can identify the following roles:

- **Infrastructure Providers (InPs):** They own and manage part or all of the network infrastructure (physical and/or logical). It can be further distinguished between Radio Access Network (RAN) infrastructure provider (owning the physical infrastructure such as the antenna sites and the HW equipment for the antennas), and datacenter/cloud infrastructure provider (managing local and central datacenters). The former provides access to radio resources (i.e., spectrum), while the latter provides virtualized resources such as virtualized computing, storage, and networking, by deploying a virtualization environment to logically abstract the physical infrastructure.
- **Network Operators (NOs):** They operate the physical and virtualized network functions and the communications links to realize a mobile network, and provide mo-

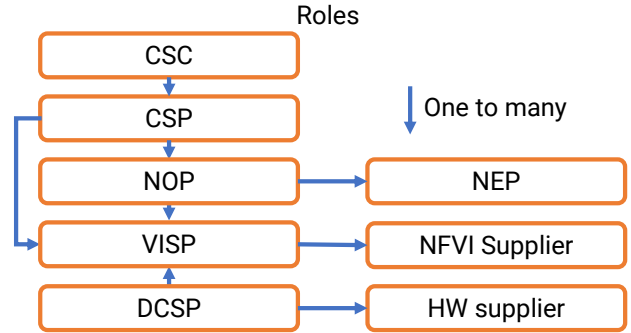


FIGURE 3. The high level roles as defined by 3GPP in [10]

bile connectivity towards mobile end-users. They also sell dedicated mobile network instances (*network slices as a service*), each one realizing specified telecommunication services such as, e.g., massive machine-type communication (mMTC), to tenants. A network operator could lease the needed physical or logical resources from one or more specialized InPs; also, an InP could additionally act as a Mobile Service Provider (MSP) using its own infrastructure, and probably, leasing certain resources from other specialized InPs.

- **Tenants:** They are business entities that rent and leverage a network slice provided by the Network Operator. They could be MVNO, other enterprises (e.g. vertical industries), or any other organization requiring a telecommunication service for their business operations.
- **End Users:** They are the ultimate consumers of the service provided by the tenant, which could be individuals subscribed to a “classical” communication service (e.g., voice, video), employees accessing a company-private VPN service, or sensors requiring some means to transport for their gathered data (e.g. for an IoT scenario).

One important aspect of the actor/role distinction, already hinted above, is that a specific actor can play different roles. For instance, the telecommunications company acting as Network Operator can also act as Infrastructure Provider (which is what usually happens today as well: they provide their own infrastructure to build-up services, and they also can hire certain resources to third parties). Also, a vertical industry company could play the role of Tenant, but also the role of Infrastructure Provider. Another important feature is that network slices could be sold to the companies acting as tenants, but also the telecommunication company acting as Operator could have its own slices to provide certain services (as operators do now).

These roles are fully aligned with the ones currently studied in 3GPP [10], which are depicted in Fig. 3. The utmost level in the 3GPP management architecture is represented by Communication Service Customers (CSC), who are enjoying a service (e.g., industry 4.0) provided by Communica-

tion Service Providers (CSPs). Therefore, CSCs are totally aligned with the End Users concept defined above, while Tenants, instead, perfectly match with the definition of CSP provided by 3GPP. CSPs organize and structure a communication service on top of the Network Operator (NOP) by, for instance, requesting for a network slice template from the NOP portfolio (e.g., the Industry 4.0 service may require mMTC and eMBB slice instances). NOPs, thus, have a direct mapping with the roles described above (i.e. the Network Operators). To build the network slice instances they use VNFs from one or more Network Equipment Providers (NEPs), which are finally executed in the underlying cloud infrastructure. In 3GPP terminology, the roles that agglomerate the lower layers of the architecture are the Virtualization Infrastructure Service Provider (VISPs), who provide virtualized connectivity by using the hardware from one or more NFVI supplier (offering e.g., transport network functionality) and the Data Centre Service Provider (DCSP), who finally offers the bare metal (acquired from an HW supplier) to run the service requested by CSC. From the above discussion, it is clear that our proposed actors and roles setup is fully compliant with the one standardized by 3GPP.

B. NETWORK SLICE SERVICE PORTFOLIO

There are different ways, from the managing and control perspective, in which a Network Operator can provide a network slice to a given tenant. These range from the provision of a “bearer service” for data traffic, with no associated control, to the provision of a set of control primitives that enable tailoring the composition of the functions to transport this data. Of course, these offer types have different requirements in terms of the potential impact on efficiency, the required interfaces and security considerations [11]. We envision the following portfolio of slice services:

- **Offer type N** (No control): in this case, the Network Operator operates the slice and provides communication services on behalf of the tenant. A tenant requests the commissioning of a network slice by providing the high-level requirements of the telecommunication service to be provided. Operation of the network slice is completely handled by the NO, and the tenant might only receive coarse-grained performance reports. This allows for a flexible Network Slice market that is ultimately enabled by the algorithms running in a module dedicated to this aim.
- **Offer type L** (Limited control): here the Network Operator allows the tenant to have some partial configuration and control options over the slice. Apart from the coarse-grained specification of the previous case, such as e.g. bandwidth, in this case, the tenant can specify more fine-grained configuration options for the requested network slice. Moreover, selected network operations (e.g., subscriber data management, QoS control) are performed by the tenant. Still, the major part of network operation is handled by the NO, although the network slice could integrate a set of tenant-owned

functions that are customized (and certified) for its needs. This interaction of these *onboarded* functions with the NO’s systems would be strictly monitored and controlled by the NO.

- **Offer type F** (Full control): finally, with this case, the Network Operator allows extended slice configuration and control options for the tenant. In addition to the control provided in the previous case, here the tenant has rather wide control over the deployed network functions. This can go as far as, e.g., the tenant onboarding its own network functions for selected areas (mobility, session management, etc.), contributing with its own infrastructure, or operating a part of the network slice independent of the NO.

The above represents a service portfolio that a Network Operator could define to provide different types of slices to tenants. We illustrate the advantages of this flexibility in the next section, describing two different use cases of relevance for future networking scenarios.

One major difference between our proposal and the ones currently defined in the state of the art [12] is the flexibility. Indeed, [12] envisions the type N operation, but type L and type F are out of the scope. As a matter of fact, [12] explicitly states in clause 5.1.1.2 that the Network Operator may create network slices on top of the Infrastructure Provider resources, but the management system shall be extended to provide more or different models, that necessarily rely on novel management system such as the one presented here. Analogously, the proposal in [13] distinguishes between two types of slices, Internal or External, depending on their intended use. For the External slices, only two offers are supported, Provider managed or Tenant managed, which would correspond to Offers N and F, respectively –no limited control is envisioned.

C. USE CASE: ON-DEMAND VIDEO SLICES

Nowadays, content providers, such as e.g. Youtube or Netflix use the underlying mobile network as a ‘data pipe’ to the end user, relying on the standard best effort configuration. This has several drawbacks, including (i) no KPI is guaranteed, which could result in a poor service, (ii) the content provider cannot tune the network configuration, e.g., asking for prioritized handling of traffic; and (iii) the NO cannot monetize this traffic, as it is not offering any added value to the over the top service (OTT). The possibility of dynamically instantiating network slices can solve these problems, allowing thus for a flexible Network Slice Market.

While much mobile content is delivered today as OTT services based on best effort quality, the provision of content over mobile networks continues to be an evolving area. Some mobile operators, in an effort to differentiate themselves, have offered data plans which include video from particular applications as so-called “zero-rated services” (i.e., they do not count towards a user’s data allowance [14]). One well-known example of this service in the US is T-Mobile’s Binge On. Furthermore, AT&T and Verizon have also launched

unlimited data plans with managed video traffic, which is offered at different video quality depending on the price of the data plan. Additionally, venues are looking for ways to differentiate their offered experience, with novel services such as 360 degree or virtual reality content, developed by groups such as Intel [15]. Other examples of the on-demand multimedia service delivery are vehicular environments [16], mobile edge computing (MEC) [16] and Smart Cities [18]. Finally, marketing and advertising are also evolving to make use of mobile technologies, mixed reality, and geo-location data [19].

The introduction of dynamic network slicing will help to support quickly trialing these new products, and potentially drive the need for a flexible Network Slice Market. Following these lines, we envision the future network slicing landscape as follows. We assume a scenario in which the InP is a NO, owning the RAN infrastructure along with datacenters at the edge, while it leases cloud resources from another InP for the central cloud. The NO supports deploying customized network slices to sustain dynamic requests from various tenants, for instance:

- **Event-tailored:** The tenant is an event organizer who requires a “video upload” slice for a localized area (e.g. a stadium and its surroundings) during a given period of time corresponding to the event (e.g. a football match). The requested service requires a high throughput in the uplink, an relatively average latency, and support for relatively little mobility in a localized coverage.
- **Nationwide delivery-tailored:** The tenant requires a “congestion free” video slice for wide (national) coverage, supporting HD video. To fulfill this demand, the MSP needs to set up a slice providing high throughput and reliability, and medium latency and mobility, with nationwide coverage. These are requirements different from a “common” eMMB slice, which might also include optimizations to better support transmission protocols like e.g. DASH.

If the tenant has no networking expertise, or does not want to manage the network slicing mechanics, the NO will provision and operate the corresponding network slice(s) tailored for this communication service (i.e., Offer type N). In this case, the KPIs of interest to the tenant could be the desired coverage area, target throughput, the number of devices, etc. The use of this offer enables better optimization of the resources required to implement the slices, as the MSP keeps the full control on those resources, and therefore algorithms can leverage multiplexing gains of traffic among slices.

On the other hand, in case the tenant wants to partially or completely control the actual resources implementing the slice (i.e., offer types L or F), the potential gains due to multiplexing are reduced. In this case, the MSP needs to carefully review the allocation of resources before accepting new slice requests, as the tenant might not accept any change in the agreed resource allocations (cf. Sec. III-B regarding

resource commitment models) in case a new slice instance shall be commissioned, or another tenant is experiencing a lack of resources.

D. USE CASE: INDUSTRY 4.0

Here we consider the case where the tenant is a vertical enterprise that owns some Industry 4.0 factory sites.¹ More specifically, we assume that the tenant wants a secure network for highly sensitive traffic from monitoring sensors in the product line of the factory floor. The requirements for this indoor-only service are: low latency, average throughput, and very-high reliability, with no support for mobility.

The driver for this scenario comes from the vertical’s requirement for a secure private network within the factories fully isolated for its critical monitoring sensors network. Full isolation of the vertical’s traffic against any network provider or network user is only possible by running a private network on the infrastructure owned by the vertical. The vertical relies on the resources provided by its own private network infrastructure inside the factory premises. In that case, the vertical combines the role of tenant, NO, and infrastructure provider. More precisely, the various organizations of the vertical (production line, delivery line entities) could be considered as “inner tenants” of the vertical. In this ecosystem, the MSP becomes a possible business partner, selling to the vertical its expertise into designing and rolling out IoT networks onto the vertical’s private infrastructure, or some software assets for realizing the non-critical slices.

In addition to critical IoT communications, the vertical uses network slicing for optimizing its own network for its various organization entities (e.g., product line, delivery line, commercial service). An example of three slices (critical IoT, non-critical IoT, and corporate eMBB) deployed by the Industry 4.0 vertical on its own private infrastructure is depicted in Fig. 4. These slices are:

- **Private slice for critical IoT, indoor coverage.** This slice may require customization down to the PHY layer, with only transmission (and reception) functionality shared across other network slices. It will implement its own specific radio scheduler, which will result in higher complexity when managing radio resources for multiples slices, and will simplify Network Access Stratum (NAS) signaling, as sensors inside the factory do not need any mobility management.
- **Private slice for non-critical IoT, factory campus coverage.** This slice serves the forklift sensors inside the factory campus. The non-critical IoT slice may not need a customized PHY or MAC layer, which increases the deployment flexibility and reduces their costs, given that no proprietary technology is needed.
- **Private slice for eMBB, factory campus coverage.** This slice provides employees with access to its private cor-

¹Industry 4.0 refers to a fourth industrial revolution combining production methods with state-of-the-art information and communication technology [20]

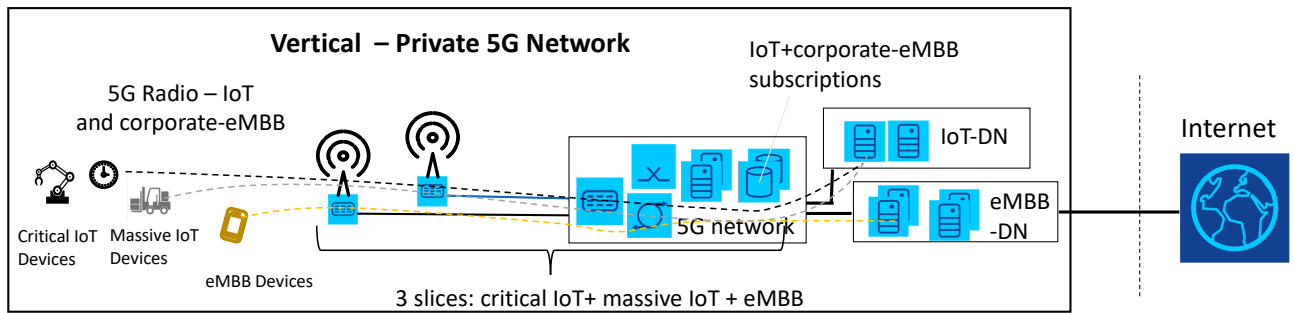


FIGURE 4. Architecture for deploying industry 4.0 slices onto a vertical private network

porate data network, which may be connected to the Internet and thus allow also Internet access. The vertical's employees have mobile devices and subscriptions to access the corporate network, with both of them (devices and subscriptions) being managed by the vertical.

IV. ETSI NFV MANO AND ITS LIMITATIONS

In this section, we first provide a quick overview of the ETSI NFV MANO Framework, describing its most relevant elements (some of which will be extended by our architecture), and summarizing the different strategies that could be used to assign resources to network slices. Then, we discuss the limitations of the framework when dealing with multi-domain operation.

A. SUMMARY OF THE ETSI NFV MANO FRAMEWORK

Fig. 5 depicts the NFV MANO system architecture proposed by the ETSI ISG NFV [6], which is composed of the following three main functional blocks:

- *Virtualized Infrastructure Manager (VIM)* for the management of NFV Infrastructure (NFVI) resources like computing, networking, and storage.
- *Virtualized Network Function Manager (VNFM)* for the lifecycle management (LCM) of Virtualized Network Functions (VNFs) that are deployed and instantiated over the NFVI.
- *NFV Orchestrator (NFVO)* for the service and resource management of the network services (NS) that are formed by chaining multiple VNFs over virtual links (VL) and characterized by the VNF Forwarding Graph (VNF-FG).

These three functional blocks interact with each other using standard interfaces, which are specified for the relevant reference points and serve to provide lifecycle management of virtualized resources belonging to different realms.

In addition to these functional blocks, there are various catalogs that contain relevant descriptor files such as, e.g., the VNF Descriptor (VNFD) file and the NS Descriptor (NSD) file, specifying the operational, functional, resource, performance, and policy requirements of the VNFs and NSs, respectively. The NFV MANO system allocates resources,

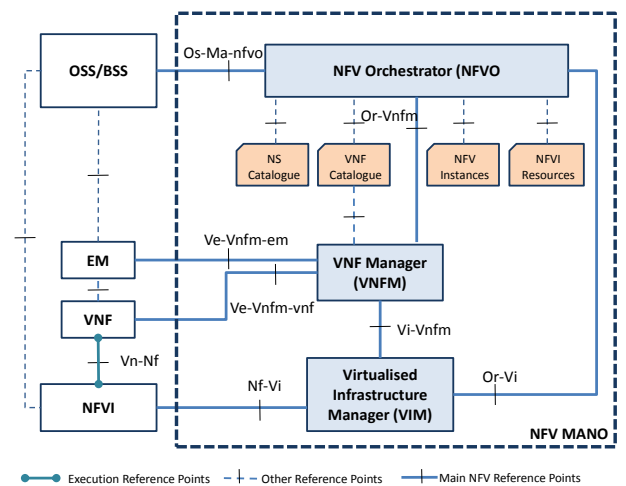


FIGURE 5. The NFV Management and Orchestration (MANO) framework as specified by ETSI (cf. [6]).

deploys, and instantiates VNFs/NSs over the NFVI according to the request and requirements specified in the respective VNFD/NSD files.

The multiple NS instances (that could belong to different tenants or to the Network Operator itself) are managed by the NFV MANO system. As part of the LCM tasks, the system has the capability to instantiate, migrate, scale-in/out/down/up, update/upgrade, and terminate VNF/NS instances. Also, the system can orchestrate network resources and NSs based on a set of policy rules.

The architecture presented in this section is the current state of the art for NFV Orchestration, and it has several real-life implementations coming from different working groups such as OSM [21].

B. RESOURCE COMMITMENT MODELS

ETSI directives on resource management procedures [22] define three so-called “resource commitment models”: reservation model, quota model, and on-demand model: with the **quota model**, the NFVI limits the resources that a slice can obtain from a particular NFVI-PoP (Point of Presence); with the **reservation model**, a specified amount of resources is

statically allocated to a particular tenant or slice, even if this results in the resources remaining idle for most of the time; finally, with the **on-demand** resource commitment model, no reservation or preemptive allocation of resources is made: NFVI resources are assigned only once they are requested.

Regarding the co-existence of the quota model and the reservation model, a VIM will, as the default behavior, also apply the slice quota to the slice reservation being made. However, further rules will determine the behavior of the VIM if a reservation exceeds the specified slice quota. In our architecture, these rules are determined from the policies implemented in a novel element, the Inter-slice Resource Broker.

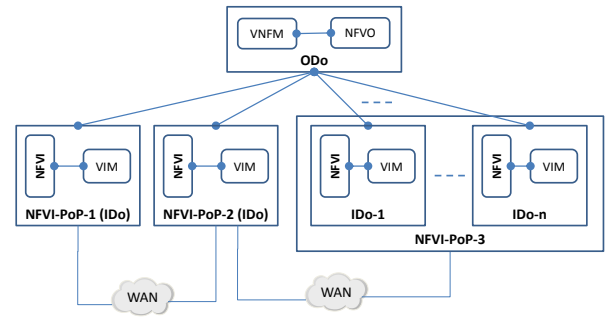
C. ISSUES IN MULTI-DOMAIN SCENARIOS

The multi-domain issue is being duly addressed in ETSI ISG NFV that has published two reports. Report [23] deals with managing the connectivity of an NS deployed over multiple NFVI-PoPs, and assumes a single MANO system that manages interconnectivity issues over the WAN links connecting these NFVI-PoPs. The second report [24] highlights the different architectural options and recommendations to support MANO operations in multiple administrative domains. According to [24], two distinct administrative domains are specified: one domain is characterized by the NFVIaaS provider and the other by the NFVIaaS consumer. These two domains may either be owned by the same or different actors, and the NFVIaaS provider domain can support multiple NFVIaaS consumers. For the sake of clarity and explanation, this work re-uses a few ETSI concepts and therefore refers to the NFVIaaS provider’s administrative domain as “Infrastructure Domain” (IDo) and the NFVIaaS consumer’s administrative domain as “Orchestration Domain” (ODo).

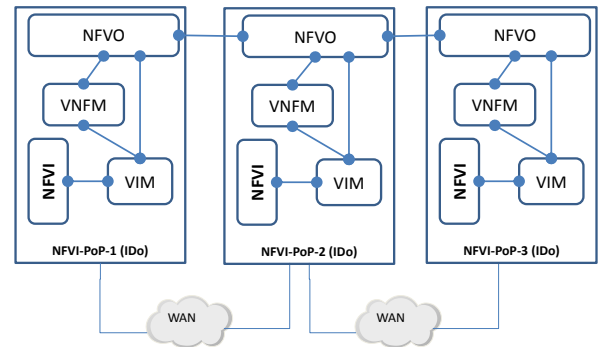
The IDo consists of NFVI resources and has at the minimum one VIM. The ODo, on the other hand, consists of the VNFM and the NFVO functional blocks of the NFV MANO system. It may also be that the IDo consists of the VIM and VNFM in which case the ODo will be composed of NFVO only. A single IDo may support one or more isolated ODo(s), and/or a single ODo may consume the resources of multiple IDos. The latter is the case when the ODo has to take care of the LCM of NS(s) deployed across multiple IDos, and this also is one of the focus issues of this article.

There are several deployment options for IDo/ODo, and each one will have performance implications on the ODo performance when providing MANO services. We focus on the scenario where a NS is deployed across multiple IDos and analyze the implications on the MANO performance. It should be noted that [25] and [26] introduce the notion and definition of Quality of Decision (QoD), which can be used to quantify the performance of a MANO system. QoD is measured in terms of the following criteria:

- How resource efficient the management action is. The resource efficiency is in turn measured in terms of:
 - Whether both the long-term and short-term resource requirements of the managed VNF will be



(a) Centralized Management of Infrastructure domains



(b) Distributed Management of Infrastructure domains

FIGURE 6. MANO Deployment options in multi-domain scenario.

fulfilled in the selected compute node.

- How non-intrusive a management action has been for other VNFs that are already provisioned in the selected compute node. That is, to what extent will the managed VNF VM affect the performance of other VNFs in the selected compute node in terms of resource availability.
- Number of times the management action has to be executed before the most-suitable compute node is determined to migrate/scale the managed VNF to.
- The timeliness of the computation and execution of MANO LCM actions. The latter criterion is more relevant in the management of a multi-site NS scenario as described below.

Fig. 6 shows two main deployment options of the MANO in a multi-site environment. Fig. 6(a) shows a scenario where a central ODo is used to manage multiple IDos in same and different NFVI-PoPs. In case the IDos and the ODo are collocated within the same NFVI-PoP (e.g., NFVI-PoP-3 in the figure), then the LCM operations on the NS instance(s) can be imparted without much impact on MANO execution time.

However, such collocated management of NS instances is no longer suitable when the IDos are deployed in geographically dispersed NFVI-PoPs and interconnected via WAN

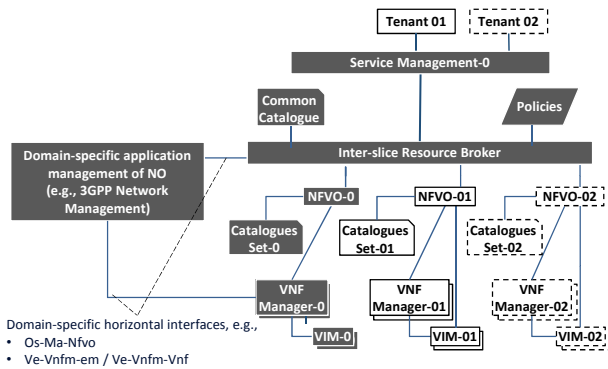


FIGURE 7. ISRB and tenant-specific ETSI NFV MANO stack instances

infrastructure. In such a scenario, i.e., when managing an NS that is deployed across multiple NFVI-PoPs, issues in the WAN infrastructure can impact the performance and hence the QoS of the NFVO/VNFM in the ODo. For instance, WAN delays will impact the timely delivery of performance monitored data/KPIs from the different IDos towards the centralized ODo. This, in turn, will delay the NFVO/VNFM to analyze and derive appropriate LCM decision and will also result in the delay of the application of the LCM actions. It could also render the monitored data, and hence the corresponding LCM decision, as stale by the time it gets processed. Moreover, a central ODo also introduces a single-point-of-failure.

To overcome the above issues of managing a multi-site NS from a central ODo, the [24] proposes to distribute the ODo such that each NFVI-PoP has its own MANO stack. In other words, each IDo domain will have its own ODo (i.e., NFVO and VNFM) as illustrated in Fig. 6(b). The MANO operations on the multi-site NS instance(s) are then coordinated in a peer-to-peer manner between the respective NFVO instances over a newly proposed Or-Or reference point [24]. However, this approach not only brings more complexity, but is sub-optimal in view of the delay-sensitive nature of MANO operations. Considering the above challenges, we propose an architectural option that is based on a distributed MANO system (as in Fig.6(b)), but instead of a peer-to-peer interaction, our architecture proposes that the coordination between the different MANO systems is carried out by an over-arching entity. Our proposed architecture extends the standard ETSI NFV MANO system (see Fig.5) and offers a versatile solution. The details of our proposed multi-domain MANO system architecture is presented in the next section.

V. A NOVEL MULTI-DOMAIN MANAGEMENT AND ORCHESTRATION ARCHITECTURE

Multi-domain orchestration allows network operators (NOs) to deploy network slices in multiple administrative domains to support, e.g., some specific performance requirements or particular tenant requests. In general, ETSI assumes that in a typical NFV scenario, there will not be a single organization controlling and maintaining a whole NFV system [27]. It




FIGURE 8. Tenant-specific ETSI NFV MANO stack instances

therefore provides multiple options for the specific mapping of NFV MANO functions to administrative domains. While a subset of options assumes a rather horizontal split into IDos and ODo (cf. Sec. IV-C), this section first depicts a solution where a single IDo-0 hosts multiple ODo. In a second step, IDo-0 integrates NFVI resources of another IDo-1 but maintains the single IDo perspective as seen by the ODo.

If the NFV Infrastructure (NFVI) spans across several locations (i.e., multiple NFVI-PoPs), the network providing connectivity between these locations can be regarded to be part of the NFVI [27]. NFVI resource management across an NO's IDos can be performed using one or more VIMs as needed. Usually, however, VIMs are intended to manage the NFVI resources within one NO's IDo. For the multi-domain orchestration framework proposed in this article, required resources from these different domains are integrated into a single NFVI block that is orchestrated from just one administrative domain. From the business perspective, specific agreement among the different administrative domains should be reached to make this possible. From the technical perspective, the domain designated for orchestration is granted access to the resources in the "external" domain(s).

The proposed Management and Orchestration architectural framework, therefore, includes the novel function of an Inter-Slice Resource Broker (ISRB) which enables orchestration per ODo by managing a set of complete ETSI NFV MANO stacks, i.e., with their corresponding NFVO, VNFM, VIM and catalogue sets as depicted in Fig 7. The NO as the owner of the IDo operates a dedicated, so-called Operator-MANO (o-MANO) stack (this and further management functions of the NO are depicted as solid rectangles). Both the NO (or specific departments thereof) and tenants can define their own ODo which is associated with the dedicated so-called tenants-MANO (t-MANO) stack. An ODo-specific t-MANO stack orchestrates a tenant's network slice(s). In particular, the VIM(s) of each ODo manage(s) the set of NFVI resources assigned to the tenant (as outlined next, NFVI resources can come from multiple infrastructure domains). This architecture has the additional flexibility for enabling per tenant t-MANO stacks based on the paradigm proposed in [28].

As a second major extension beyond ETSI NFV MANO, Fig. 8 shows the framework for cross-IDo orchestration: NFVI-0, owned by InP-0 (i.e., the NO), virtually extends its IDo by integrating resources from NFVI-1, owned by InP-1. Each InP could have instantiated different t-MANO

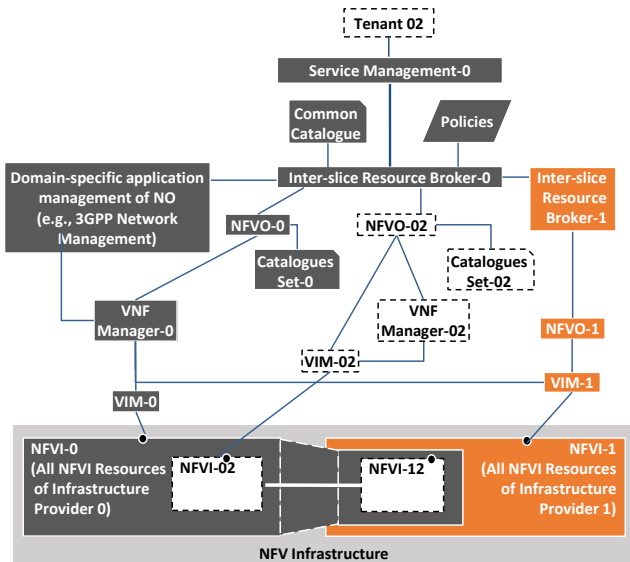


FIGURE 9. Multi-domain networking in the infrastructure domain.

stacks for different tenants, beyond the InP’s o-MANO stack. The framework for combined multi-ODo and multi-IDo orchestration is depicted in Fig. 9. For the sake of simplicity, no t-MANO stacks are depicted in the NFVI-1 domain. The following subsections describe the conceptual details and the new components of the proposed framework. For a detailed description of a first implementation and operation of the multi-domain management and orchestration system, the reader is referred to [29], which describes a 5G testbed setup in a commercial seaport environment. The setup leverages on the presented technology to orchestrate several network slice instances across the administrative domains of the port authority and the mobile network operator. Another implementation is provided in [30], which describes an open source implementation of a multi-slice radio access network and the orchestration framework.

A. NEW COMPONENTS AND CONCEPTS

1) Inter-slice Resource Broker (ISRB)

Sharing the same infrastructure across tenants and network slices entails the following trade-off: resource reservation versus flexible resource sharing. Specific service-level agreements (SLAs) define the concrete embodiment of reservation and on-demand allocation rules. For example, a tenant may request a fixed amount of NFVI resources (in such a case, the t-MANO stack assigned to a tenant could exclusively manage the allocated quota of resources). In another case, a tenant may agree that a percentage of the associated, but unused resources, may be dynamically allocated to other tenants or slices, thus realizing cost savings and allowing multiplexing gains. The role of the ISRB is to have a general view of the whole infrastructure that can be offered within a single administrative domain, as well as monitoring the usage of the resource subsets allocated to a tenant’s t-MANO stack. It controls the dimensioning of resources that are assigned

to each tenant and their status, including those resources not yet assigned. Upon instantiation of a tenant’s network slice, initial quotas are assigned to the responsible t-MANO stack instance. Despite this initial assignment, resource allocation can be reshaped at runtime if tenants request for that. This also implies that in case a tenant’s slices do not utilize all allocated resources, the reserved but idle resources will not automatically be made available to slices of other tenants. Moreover, special terms in SLAs can allow a tenant to exceed its assigned quota for a certain time and at a certain cost. In case a t-MANO stack is permanently decommissioned, the associated resources are released to NFVI-0 again.

Except for the NO (or InP, respectively), tenants are neither aware of the existence nor the resource utilization level of other tenants. They only have an SLA specifying their right to use certain resources in a certain manner.

The rules for the resource commitment model utilized by individual tenants as well as for resource sharing and prioritization, are kept in a policy catalog maintained by the ISRB. This catalog also contains an inventory map of currently available infrastructure resources and their allocation to each tenant. The tenant-specific NFVO reports resource utilization to the ISRB that may use this information to reshape the current association patterns according to new external triggers, such as a new slice creation request or a re-orchestration request from an already hosted slice.

In the case that NFVI-0 resources are not sufficient to accommodate all slice instances and their requirements, NFVO-0 triggers a request for additional resources with ISRB-0. As one option, ISRB-0 can subsequently request resources from another IDo, e.g., ISRB-1, cf. Fig. 9. If ISRB-1 can accommodate the request, the o-MANO stack of InP-1 will then partition an according subset of NFVI-1 resources and expose them to InP-0’s o-MANO stack as one or multiple separate NFVI-PoP(s). For this purpose, VIM-1 of InP-1 will integrate towards VNFM-0 and NFVO-0 using the respective reference points. After the successful integration of the NFVI-1 resource subset into the NFVO-0 IDo, ISRB-0 can henceforth re-allocate them to other t-MANO stacks within InP-0.

2) NFV Infrastructure Concept

The IDo NFVI-0 constitutes a conceptual re-design of the ETSI NFV IDo. This comprises two major characteristics: (1) Disposition of NFVI-0 resources into tenant-specific subsets according to the resource commitment model requested by the tenant and (2) integration of NFVI resources from other domains into NFVI-0 domain forming a virtually integrated, single domain, cf. Fig. 9. For the depicted setup, InP-0 and InP-1 have to reach an agreement on:

- The amount and type of InP-1 NFVI resources to be allocated into the InP-0 Infrastructure Domain; this includes the typical NFVI resources (compute, storage and networking).
- The commitment model to allocate them (static quota, dynamically scalable quota, allowed/supported proto-

cols, etc.).

- The agreement could also include the VIM(s) in case the rented resources would require a specific VIM (VIM-1 in Fig. 9). From the technical point of view, the only condition for integration is that such “external” VIMs shall provide the ETSI NFV Or-Vi and Vi-Vnfm reference points in order to connect with the VNFM(s) and the NFVO.
- Additional operational policies (including scaling rules, security requirements, redundancy, and over-provisioning levels).

Regarding the fundamental realization of multi-domain orchestration, this article considers two scenarios:

- 1) InP-0 extends the NFVI-0 infrastructure in order to provide a better service to the hosted tenants. E.g., Tenant-02 could request more infrastructure resources from InP-0 than NFVI-0 can satisfy. Then, the InP-0 signs a business contract with another InP (InP-1) to include resources from that provider into his infrastructure domain. The request of Tenant-02 can now be satisfied and the new resources can be made available to Tenant-02’s t-MANO stack, no additional business contract between InP-1 and Tenant-02 is required. Technically, the fact that some NFVI resources used by Tenant-02’s t-MANO stack are actually located in the InP-1 domain is transparent from the tenant perspective. In other words, for Tenant-02 it looks like resources from the InP-0 domain are used and only a single business relationship is maintained.
- 2) The tenant explicitly wants to extend its slice using specific infrastructure from another InP (InP-1). This could happen when a tenant already has certain infrastructure up and running on a different InP and does not want to take the effort of migrating that infrastructure into the InP-0 domain. I.e., the tenant already has a business contract with both, InP-0 and InP-1. In this case, besides the corresponding update of these both contracts, a new contract involving both InPs needs to be agreed.

B. EXTENDED ETSI NFV MANO COMPONENTS

In order to support multi-domain orchestration, the NFV Orchestrator NFVO-0 (cf. Fig. 9) plays a special role since it oversees and executes, respectively, the management of NFVI-0 resources based on the constraints received by ISRB. Specifically, it needs to support and/or implement the following technical aspects:

- InP-1 shall isolate and assign the requested NFVI resources and provide the necessary interfaces/reference points to InP-0. In particular, this applies to the Nf-Vi reference point when the InP-0 deploys its own VIMs, or the Or-Vi and Vi-Vnfm reference points when the VIMs are supplied by the InP-1.
- InP-0 shall take the newly integrated NFVI resources and associate those to a o-MANO or one (or multiple) t-



FIGURE 10. Lifecycle of a network slice [12].

MANO stacks, thus forming a “merged”, single-domain set of resources for these stacks.

- While Fig. 9 does not assume any deployment constraints of VIMs or the entire o-/t-MANO stack instances, for performance reasons (e.g., latency) it might be necessary to add such constraints. However, this does not change the functional perspective of the architecture.
- A tenant’s t-MANO stack should have the full information about logical node topology and resources, address space, etc. within the tenant’s dedicated shares of the two IDos (NFVI-02 and NFVI-12 in Fig. 9). Hence, NFVI-02 and NFVI-12 can be used, together with the associated VIMs, to set up the interconnections between VNFs in the tenant’s Odo accordingly.
- Security issues: A tenant may not want to rely on trust in arbitrary InPs –hence an SLA between tenant and NO may restrict the choice of InPs the NO can use to host tenant functions. It can be further noted that in a multi-domain scenario, inherently the number of involved parties, data centers, interfaces, and software routines will be higher, thus increasing the attack surface. There is no specific remedy. Against this, the general security rules apply, in particular designing a sound security architecture and implementing it carefully in order to minimize the threat of exploitable vulnerabilities.
- Adjustment of resources assigned to a specific NFV MANO stack instance relies on an ‘offline’ procedure (as for the single-domain use case). I.e., t-MANO stacks could be assigned with a specific quota of resources (regardless of whether those resources come from a single domain or multiple InP domains) that should be properly dimensioned to dynamically scale during the slice operation. In both cases (single-domain or multi-domain), an offline re-negotiation is necessary if the assigned quota of resources have not been sufficient to meet the tenant’s necessities.

C. NETWORK SLICE LIFECYCLE

When a network slice shall be commissioned, a network slice descriptor is provided to the ISRB. Such a descriptor does not only contain information on control and data layer network functions of a network slice, but also on the associated MANO stack instance for managing these network functions (NFs). Hence, the network slice descriptor is comprised of two major parts that specify the functions, resources, and policies that are required to, respectively, (i) perform life-cycle management for a network slice and (ii) realize the network service requested by the tenant. While the former

TABLE 4. Scope of LCM operations

	Monitoring	Global scope	Local scope	Realtime
Instantiation/Deletion			✓	
Scale-up	✓	✓		✓
Scale-down	✓	(✓)	✓	✓
Scale-out	✓	✓		✓
Scale-in	✓	(✓)	✓	✓
Migration	✓	(✓)	✓	✓
SW Upgrade			✓	
Configuration			✓	

TABLE 5. Empirical evaluation parameters

Parameters	Values
Number of VNFs	100
Traffic distribution	Pareto
Monitoring frequency	100ms
Threshold	20% of available resources
Simulation time	20s
Simulation seeds	1000
Traffic load	50MBit/s

point comprises a specification of the NFV MANO stack instance (NFVO, VNFM, VIM, NFVI instances, catalogs for network services and functions, etc.) that is dedicated to the lifecycle management of the network slice, the latter includes the network service descriptor(s), i.e., the collection of VNFs and PNFs that, as a whole, form the control and data layer architecture of the particular network slice instance.

According to [12], the network slices lifecycle management is composed of four distinct phases (depicted in Fig. 10): (i) preparation, (ii) instantiation, configuration and activation, (iii) run-time, and (iv) decommissioning. The network slice descriptor contains the necessary information to carry out phases (ii) to (iv) appropriately.

In a first step, the ISRB uses the network slice descriptor to commission a new NFV MANO stack. In the second step, the same network slice descriptor is utilized to generate the necessary objects and models which the ETSI NFV MANO stack instance operates on, i.e., NFV service catalogue, VNF/PNF catalogues, NFV instances, and NFVI resources. For the allocation of the NFVI resources that are under control of this MANO stack instance, the ISRB uses a combination of the resource commitment models as outlined in Section IV-B. Commissioning of the network slice control and data layer functions is triggered by the Inter-slice Resource Broker via the Os-Ma-Nfvo reference point of the NFVO by providing or referring to the set of network service descriptors to be instantiated.

The network slice lifecycle management is now delegated to the NFV MANO instance and the according domain-specific application management functions (see Fig. 9). This includes several tasks, including the instantiation and configuration of the network services and associated network functions, the activation of the network slice, the runtime supervision and reporting as well as upgrading, reconfiguration, scaling, and finally, the deactivation and termination of the network slice.

VI. PERFORMANCE EVALUATION

Practicability and implementability issues are the main aspects that must be taken into account when considering an advanced multi-domain orchestration. Hereafter, we highlight the main advantages along our novel solution with respect to traditional and legacy approaches, as well as the experienced limitations using off-the-shelf equipment. Last, we show how our solution can be easily integrated to existing standard interfaces without requiring significant architectural changes.

A. LCM OPERATIONS TAXONOMY

For our validation analysis, we consider a legacy multi-domain orchestration wherein a single MANO stack is deployed across different infrastructure providers. In particular, each infrastructure domain is provided with its own Virtualized Infrastructure Manager (VIM) component and (might be provided) with the VNF Manager (VNFM). The NFV Orchestrator (NFVO), as the entity in charge of taking Lifecycle Management (LCM) decisions, is shared among different infrastructure domains.

We detail in Table 4 the relevant LCM operations considered within our simulation campaign. When an LCM operation is executed, the VNF Forwarding Graph (VNFFG) is adjusted accordingly. However, NFVO can trigger LCM operations locally, i.e., without affecting external domains' VNF Forwarding Graphs (VNFFGs), or globally, i.e., LCM operations on other infrastructure domains are required. Note that in some cases, operations can be executed within local and global scope. This is described in the network service descriptor (NSD) following an event-threshold definition: if the resource increase request exceeds such a threshold, i.e., it may impact on VNFs chained within the same network service but running on different infrastructures, the operation will be executed globally. While some of those operations might also be monitored to avoid unhandled service degradations, this might further incur in the additional overhead on the communication means, as shown in the remainder of the section.

B. VALIDATION RESULTS

We provide our preliminary validation results by studying the communication overhead as well as the end-to-end service delay while deploying three relevant multi-domain network slices.

Our proposal relies on the multi-MANO deployment, i.e., different independent classical MANO stacks are deployed on each single infrastructure domain, namely NFVI-PoP, and all of them are connected to the Inter Slice Resource Broker (ISRB) component—which can be envisioned as a stand-alone software running our algorithms, as the one described in [31]—through a dedicated interface.² The MANO deployment is realized through OpenStack heat template

²Note that, such a new interface can be readily mapped onto the Or-Or interface as per the standard report [24].




FIGURE 11. Validation reference scenarios.




FIGURE 12. Validation reference scenarios.

based on a pre-configured ONAP deployment.³ To clarify the concept, we depict the baseline architectural solution and our novel framework in Fig. 11. This brings a two-fold advantage: (i) *trustworthiness*, as the ISRB is recognized by all connected infrastructure domains as the only trusted entity so as to avoid any security threat, and (ii) a clean *master-slave relationship*, with the ISRB constantly keeping track of the status changes of each infrastructure domains, and taking decisions on LCM operations, playing as a centralized entity only when unexpected LCM operations might affect the overall performance.

We consider three different network services that automatically chain distinct virtualized network functions distributed over different infrastructure domains, as shown in Fig. 12. The first considered multi-domain slice delivers enhanced Mobile BroadBand (eMBB) services by deploying a virtualized Evolved Packet Core (vEPC), virtualized switch and transport nodes on the first infrastructure domain (NFVI-1). A virtualized scheduler and traffic shaper is then deployed

³Advanced changes to existing MANO orchestrator solutions are out of the scope of this paper. However, they will be addressed in future work.

on the second infrastructure domain (NFVI-2) whereas the virtualized firewall function is executed on the third infrastructure domain (NFVI-3). The second network slice provides streaming services with a virtualized video optimizer function and a virtualized load balancing function deployed on different infrastructures. We assume such two network services as delay-tolerant services. Last, we consider an ultra-reliable low latency communication (URLLC) multi-domain slice running public safety network services. Specifically, the network service deploys a virtualized imaging processing function on the second infrastructure domain while the target-matching virtualized function (and its associated database) on the third NFVI-PoP.

We carry out an exhaustive simulation campaign to evaluate the complexity of our solution, namely *ISRB*, against a *Legacy case* in terms of the overhead of different interfaces as well as the end-to-end service delay that may play a fundamental role in case of low-latency applications (for e.g. URLLC services). Simulation parameters are listed in Table 5. We implement and emulate the communication between each deployed virtualized function and we generate synthetic traffic traces based on a Pareto statistical distribution.⁴ In the legacy scenario, the communication between the NFV orchestrator and VNF Managers is stable as the NFVO continuously collects information on running NFVs (monitoring) while quickly reacts in case of unexpected changes. In our novel multi-domain orchestration solution, our novel ISRB continuously retrieves the network service descriptor for each network slice. Once resources have been set up on different NFV infrastructures, the NFVO locally gathers monitoring information while transmitting, with a fixed frequency, few packets on the status of the NFVI to the ISRB. Some LCM operations are taken locally without requiring the intervention of the ISRB, as specified in Table 4. However, some of those operations, such as scale-down, scale-in and migration, might trigger the ISRB reconfiguration only if the event requires an number of resources that exceeds a pre-defined threshold, as previously explained.

⁴We consider this distribution as it exhibits a long-tail in the density function that helps while evaluating queues of packets and, in turn, experienced delay.

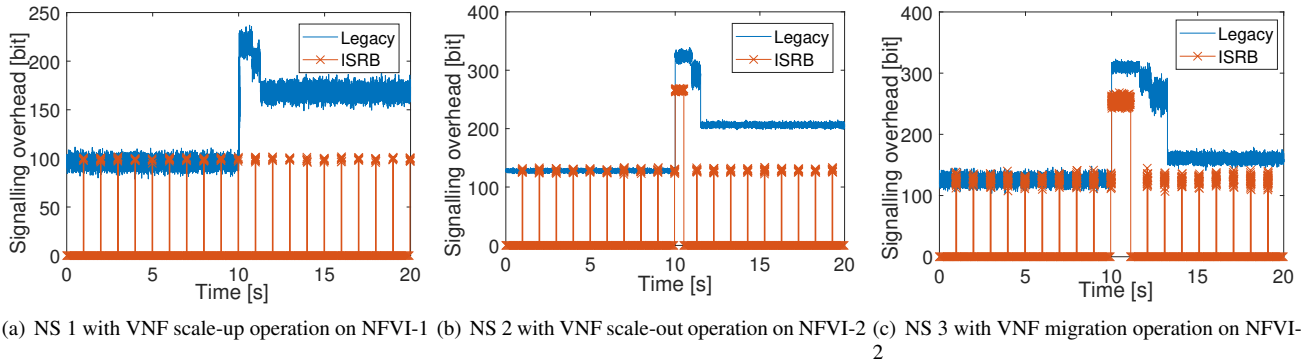


FIGURE 13. Complexity analysis with increasing traffic at time $t = 10$ seconds.




FIGURE 14. Whiskers plot for the overall end-to-end service delay.

In Fig. 13, we evaluate the communication overhead between the centralized entity (NFVO in case of legacy and ISRB in case of our novel solution) and the distributed NFV infrastructure. We run our simulations for 20 seconds with a plot granularity of 1 ms. In the first 10 seconds, the average amount of offered load is constant. For the first network service (please refer to Fig. 12), we assume networking traffic suddenly increasing at time $t = 10$ seconds. This automatically triggers a scale-up operation of the vEPC function in NFVI-1, as shown in Fig. 13(a). Several messages are exchanged to deal with the resource increasing resulting in a huge overhead that stabilizes after few seconds. However, the number of messages exchanged is higher due to the larger number of used NFV resources. Conversely, ISRB solution does not show any criticism while dealing with the scale-up operations. The number of required resources is within the fixed threshold thereby preventing other NFVIs to apply local reconfigurations. In Fig. 13(b), when traffic flows suddenly increases a scale-out operation is triggered on NFVI-2. In particular, multiple instances of the virtualized Video Optimizer functions are instantiated to deal with the unexpected traffic boost. This automatically affects functions running on NFVI-3 so that both legacy and ISRB exhibit signaling message exchange. However, due to the limited number of messages required by the ISRB, even in that case ISRB results in a short-time and low overhead that turns into regular frequent messages after less than 1 second. Last, net-

work service 3 has a significant impact on other infrastructure domains when the traffic increases at time $t = 10$ seconds, as shown in Fig. 13(c). In particular, a function migration is required, and both solutions deal with an increasing signaling overhead that lasts roughly 1 and 3 seconds for ISRB and legacy, respectively.

While the relative gain of ISRB could appear limited, this token example unveils the difference in terms of overhead between the two considered solutions for a short-time period (20 seconds) and three different NFVI-PoPs. Real NFV deployments may support up to hundreds of NFVI-PoPs with a time window that could incur in a huge signaling burden.

Finally, we depict in Fig. 14 the whiskers plot for experienced delay values over 20 seconds of simulation. Network service 1 and 2 provide high end-to-end delay as they are delay-tolerant services. However, ISRB significantly outperforms the legacy scheme when the network service 2 is in place. In case of network service 3, although a centralized approach (legacy) could slightly benefit the system in terms of manageable delay, the complexity required does not pay off. Note that in such case, ISRB still provides affordable delay performance (below 20 ms).

VII. RELATED WORK

The architecture presented in this paper extends the one proposed by ETSI NFV to take into account specific characteristics of network slicing, multi tenancy, and service personalization, as already described above. In this section, we describe the main difference between our proposed architecture and one of the most prevalent alternatives to the ETSI NFV MANO Architecture, namely, the one proposed by the Open Network Automation Platform (ONAP) [32]. We also compare our proposal with another relevant architecture for 5G Networks Management and Orchestration, which is proposed by the European 5G-Exchange (5G-Ex) project.

A. ONAP

The ONAP initiative was launched in 2017 with the goal of providing a common platform to deliver differentiated network services on a shared infrastructure. As the main objective of ONAP is generality, in the latest version of their

architecture, the ONAP consortium proposes a clear split between the general, abstract models that tackle the problem of service design and the specific modules that control the lifecycle management of such services. More specifically, they define the *Service Design and Creation (SDC)* and the *Runtime Framework* realms. In a nutshell, they perform tasks that are commonly categorized under Network Management (SDC) and Orchestration (Run Time).

Therefore, all the tasks related to the abstraction of resources and the high-level deployment of network services are performed within the Design Time Framework, while all the tasks related to the lifecycle management and the actual representation of those resources are performed by the run time execution framework. The full specification of the architecture modules is depicted in the Fig. 15.

Within ONAP, a network service is thus defined as a collection of *recipes* that specify the behavior of a specific service which are deployed in the ONAP Operation Manager Portal. Recipes detail, among other things, factors such as the VNF deployment, the metric that have to be analyzed and the self-healing of the network.

The ONAP and the architecture presented in this paper share the same field of operation (i.e., the management, orchestration, and operation of a multi-service network) although from a very different standpoint. ONAP is very much code-oriented and module-driven, while our proposal builds on top of the ETSI NFV framework and tackles the same problem with a top-down approach. In the following, we describe how the different modules of the two architectures relate among them.

Management and Service Orchestration: our architecture perform such tasks at the Service Management and (partially) at the Inter Slice Resource Broker, that defines the interface towards the NFV-O for the subsequent resource orchestration procedure. Within ONAP, this functionality is performed by the SDC framework that then interfaces towards the run-time modules for lifecycle management.

Resource Orchestration and Lifecycle Management: in this work we specify procedures for the network slice lifecycle management by leveraging on the ETSI NFV Architecture modules. ONAP also adopts a similar approach, being the Virtualized Function Controller a replacement of the ETSI ENI Orchestration Stack.

From the above discussion, we can recognize one major difference between the two proposals. While in the ONAP architecture the concepts of network slicing and multi-tenancy are left open and possibly enforced through the ONAP Operations Manager, in our architecture we clearly define specific roles for the involved stakeholders. We believe that this tighter definition of the interaction between the roles of the tenants / service providers and infrastructure providers as done within our proposed architecture will eventually lead to a better and clearer interaction of concurrent services provided on the same infrastructure.

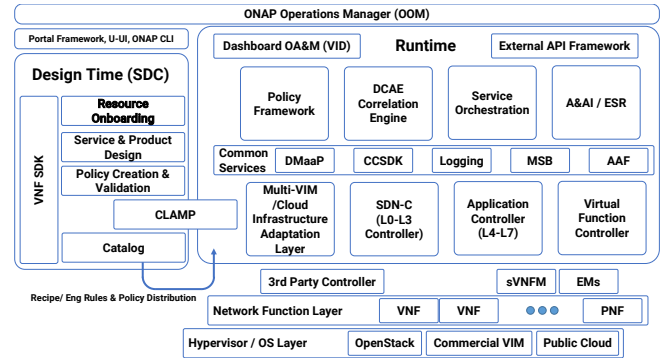


FIGURE 15. The ONAP architecture

B. 5G-EX

The 5G-Ex project targeted exactly the same problem as ours: how to provide multi-domain orchestration in a multi-slice, multi-tenant network [33], [34]. Their approach, analogously to ours, define a hierarchy of orchestrators to solve the multi-domain problem, with a multi-domain orchestrator (that belong to different network operators) linked to specific domain orchestrators.

However, the main difference of 5G EX compared to our approach is the limited flexibility in the type of offers. That is, the interaction between the tenants and the network operators, just happens through Business to Customer (B2C) interfaces that allow only no or limited control to the tenant. Instead, we believe that a more flexible management API will enable new and more efficient business models such as the ones described in Section III.

VIII. CONCLUSIONS

In our work, we have presented a novel 5G management and orchestration architecture that overcomes the main limitations of the current state-of-the-art frameworks. Namely, we have designed our architecture with the goals of being backward compatible while natively taking into account the novel concepts of multi-tenancy and network slicing across multiple infrastructure domains. A core-contribution is also the economic analysis presented in the article that further motivates the need for such a flexible architecture. In particular, it comprises (i) a novel Inter-slice Resource Broker entity and the (ii) NFV infrastructure concept. Validations results over realistic deployments show how the proposed modules outperform legacy solutions at affordable costs while supporting fundamental operations in 5G multi-domain networks.

REFERENCES

- [1] NGMN Alliance, "Description of network slicing concept," *White Paper*, 2016.
- [2] 3GPP, "Study on new radio (NR) access technology physical layer aspects," Tech. Rep. TR 38.802 V15.0.0, 2018.
- [3] P. Rost, C. Mannweiler, D. S. Michalopoulos, C. Sartori, V. Sciancalepore, N. Sastry, O. Holland, S. Tayade, B. Han, D. Bega, D. Aziz, and H. Bakker, "Network Slicing to Enable Scalability and Flexibility in 5g Mobile Networks," *IEEE Communications Magazine*, vol. 55, no. 5, pp. 72–79, May 2017.

- [4] Y. Xie, Z. Liu, S. Wang, and Y. Wang, "Service Function Chaining Resource Allocation: A Survey," *arXiv:1608.00095 [cs]*, Jul. 2016, arXiv: 1608.00095.
- [5] W. Hahn, B. Gajic, and C. Mannweiler, "Compound implementation of chained network functions and virtual resource management performance evaluation," in *NOMS 2016 - 2016 IEEE/IFIP Network Operations and Management Symposium*. Istanbul, Turkey: IEEE, Apr. 2016, pp. 1301–1304.
- [6] ETSI GS NFV-MAN 001, "Network Functions Virtualisation (NFV); Management and Orchestration," Tech. Rep. v1.1.1, Dec. 2014.
- [7] OFCOM, "Communications Market Report 2018," Aug. 2018. [Online]. Available: http://bit.ly/ofcom_2018
- [8] GSMA, "The mobile economy - Europe 2018," 2018. [Online]. Available: http://bit.ly/gsama_mobeco
- [9] 5G NORMA, "Deliverable D2.3 - Evaluation architecture design and socio-economic analysis - final report," Dec. 2017. [Online]. Available: <http://bit.ly/5GNORMA-D23>
- [10] 3GPP SA5 WG, "Management and orchestration; Concepts, use cases and requirements," Tech. Rep. TS 28.530 V15.0.0, 2018.
- [11] 5G NORMA, "Deliverable D3.3 - 5G NORMA network architecture - final report," Tech. Rep., Dec. 2017. [Online]. Available: <http://bit.ly/5GNORMA-D33>
- [12] 3GPP SA5 WG, "Study on management and orchestration of network slicing for next generation network," Tech. Rep. TR 28.801 V15.1.0, 2018.
- [13] L. M. Contreras and D. R. López, "A Network Service Provider Perspective on Network Slicing," *IEEE Softwarization*, Jan. 2018. [Online]. Available: <https://sdn.ieee.org/newsletter/january-2018/a-network-service-provider-perspective-on-network-slicing>
- [14] G. Yigit, "MNOs can improve mobile video monetisation with new pricing models and intelligent optimisation." [Online]. Available: http://bit.ly/video_monetisation
- [15] Intel, "NFL + Intel® True View." [Online]. Available: http://bit.ly/NFL_Intel
- [16] I. Al Ridhawi, Y. Kotb and Y. Al Ridhawi, "Workflow-Net Based Service Composition Using Mobile Edge Nodes." *IEEE Access*, vol. 5, pp. 23719–23735, Oct. 2017.
- [17] I. Al Ridhawi, M. Aloqaily, Y. Kotb, Y. Al Ridhawi, and Y. Jararweh, "A collaborative mobile edge computing and user solution for service composition in 5G systems." *Transactions on Emerging Telecommunication Technologies*, vol. 29(11), Jun. 2018.
- [18] I. Al Ridhawi, M. Aloqaily, B. Kantarci, Y. Jararweh, and H.T. Mouftah, "A Continuous Diversified Vehicular Cloud Service Availability Framework for Smart Cities." *Computer Networks*, vol. 145, pp. 207-218, Nov. 2018.
- [19] P. Salter, "How Mixed Reality Is Revolutionizing Marketing." [Online]. Available: http://bit.ly/mixreality_revol
- [20] 5G-PPP, "5g and the Factories of the Future," Tech. Rep., 2015. [Online]. Available: http://bit.ly/5GPP_FOF
- [21] OSM, "Open Source MANO." [Online]. Available: <https://osm.etsi.org/>
- [22] ETSI GS NFV IFA 010, "Network Functions Virtualisation (NFV); Management and Orchestration; Functional requirements specification," Tech. Rep. V2.1.1, 2016.
- [23] ETSI GS NFV IFA 022, "Network Function Virtualisation (NFV); Management and Orchestration; Report on Management and Connectivity for Multi-Site Services," Tech. Rep. V3.1.1, 2018.
- [24] ETSI GS NFV IFA 028, "Network Function Virtualisation (NFV); Management and Orchestration; Report on architecture options to support multiple administrative domains," Tech. Rep. V3.1.1, 2018.
- [25] F. Z. Yousaf, C. Goncalves, L. Moreira-Matias, and X. C. Perez, "RAVA — Resource aware VNF agnostic NFV orchestration method for virtualized networks," in *2016 IEEE PIMRC*. Valencia, Spain: IEEE, Sep. 2016, pp. 1–6.
- [26] F. Z. Yousaf and T. Taleb, "Fine-grained resource-aware virtual network function management for 5G carrier cloud," *IEEE Network*, vol. 30, no. 2, pp. 110–115, Mar. 2016.
- [27] ETSI GS NFV 003, "Network Functions Virtualisation (NFV); Terminology for Main Concepts in NFV," Tech. Rep. V1.2.1, 2014.
- [28] F. Zarrar Yousaf, V. Sciancalepore, M. Liebsch and X. Costa, "MANOaaS: A Multi-tenant NFV MANO for 5G Network Slices," *IEEE Communication Magazine* vol. 57, no. 5, pp. 103–109, May 2019.
- [29] P. Rost, M. Breitbach, H. Roreger, B. Erman, C. Mannweiler, R. Miller, and I. Viering, "Customized Industrial Networks: Network Slicing Trial at Hamburg Seaport," *IEEE Wireless Communications*, vol. 25, no. 5, pp. 48–55, Oct. 2018.
- [30] G. Garcia-Aviles ; M. Gramaglia, P. Serrano and A. Banchs, "POSENS: a practical open source solution for end-to-end network slicing," *IEEE Wireless Communications*, vol. 25, no. 5, pp. 30–37, Oct. 2018.
- [31] J.X. Salvat, L. Zanzi, A. Garcia-Saavedra, V. Sciancalepore and X. Costa-Perez, "Overbooking Network Slices Through Yield-driven End-to-end Orchestration," *ACM CoNEXT 2018*, pp. 353–365, Dec. 2018.
- [32] The Linux Foundation, "Open Networking Automation Platform." [Online]. Available: <https://onap.org/>
- [33] C. J. Bernardos, B. P. Gerö, M. Di Girolamo, A. Kern, B. Martini, and I. Vaishnavi, "5GEX: realising a Europe-wide multi-domain framework for software-defined infrastructures" *Transactions on Emerging Telecommunications Technologies*, vol. 27, no. 9, pp. 1271–1280, Sep. 2016.
- [34] G. Biczok, M. Dramitinos, L. Toka, P. E. Heegaard, and H. Lonsethagen, "Manufactured by Software: SDN-Enabled Multi-Operator Composite Services with the 5G Exchange," *IEEE Communications Magazine*, vol. 55, no. 4, pp. 80–86, Apr. 2017.



technologies (Wireless and Networking) issued by GTTI in 2015.

VINCENZO SCIANCALEPORE (S'11–M'15) received his M.Sc. degree in Telecommunications Engineering and Telematics Engineering in 2011 and 2012, respectively, whereas in 2015, he received a double Ph.D. degree. Currently, he is a 5G researcher at NEC Laboratories Europe in Heidelberg, focusing his activity on network virtualization and network slicing challenges. He was also the recipient of the national award for the best Ph.D. thesis in the area of communication



EU projects and authored more than 20 papers appeared in international conference and journals.

MARCO GRAMAGLIA is a post-doc researcher at University Carlos III of Madrid (UC3M), where he received M.Sc (2009) and Ph.D (2012) degrees in Telematics Engineering. He held post-doctoral research positions at Istituto Superiore Mario Boella (Italy), the Institute of Electronics, Computer, and Telecommunications Engineering (IEIIT) of the National Research Council of Italy (CNR, Torino, Italy) CNR-IEIIT (Italy) and IMDEA Networks (Spain). He was involved in



architectures for future mobile networks. Christian has worked in several nationally and EU-funded projects covering the development of cellular and industrial communication systems.

CHRISTIAN MANNWEILER received his M.Sc. (Dipl.-Wirtsch.-Ing.) and Ph.D. (Dr.-Ing.) degrees from the Technische Universität Kaiserslautern (Germany) in 2008 and 2014, respectively. Since 2015, he is a member of the Network Automation research group at Nokia Bell Labs, where he has been working in the area of cognitive network management and SON for 5G systems. He is (co-)author of numerous articles and papers on wireless communication technologies and



economic assessment of 5G network architectures within the European Commission 5G NORMA and 5G MoNArch projects. Previously, she was Communications Engineer at QinetiQ UK, consultant at PA Consulting UK, and Senior Systems Engineer for Airvana UK.

JULIE BRADFORD received her 1st class MEng degree in Electronic Engineering with Business Management from the University of York UK, 2002, and a Postgraduate Certificate in Business Management with a distinction from Bath University UK. She works at Real Wireless in techno-economic analysis to quantify coverage, capacity and user experience in wireless networks. Most recently Julie has been part of the Real Wireless team providing the commercial and socio-




protocols for optimizing overall network performance. He has one granted patent and 15 filed patents while his research work has been published in several peer reviewed journals, conferences and book chapters

FAQIR ZARRAR YOUSAF is a Senior Researcher at NEC Laboratories Europe, Germany. He completed his PhD from TU Dortmund, Germany. His current research interest is on NFV/SDN in the context of 5G networks. He is also an active contributor to ETSI ISG NFV standards organization where he holds Rapporteurship for four work-items. He has extensive experience in the design, development, modeling, simulation and prototyping of communication systems and



he contributed as Work Package Leader in different 5G related research projects. His main contributions in this area have been in 5G-Norma (a 5G novel radio multiservice adaptive network architecture) and NGPaaS (a 5G cloud-based PaaS).

IGNACIO LABRADOR PAVÓN has been in the computing and electronics industry for over 30 years, most of them in the mobile telecommunications industry developing value-added services for different mobile network operators. During this time Ignacio worked in different technical areas and with different responsibilities, from software developer up to projects leader. Over the last several years he works as a member of the Research and Innovation department in Atos Spain, where



IEEE INFOCOM and IEEE WoWMoM.

PABLO SERRANO (M'09–SM'16) received his degree in telecommunication engineering and his Ph.D. from UC3M in 2002 and 2006, respectively. He has been with the Telematics Department of UC3M since 2002, where he currently holds the position of associate professor. He has over 80 scientific papers in peer-reviewed international journals and conferences. He has served as a Guest Editor for Computer Networks, and on the TPCs of a number of conferences and workshops including

...