

Research Article

Expanding GÉANT Testbeds Service to Support Pan-European 5G Network Slices for Research in the EuWireless Project

Álvaro Rios ¹, Barbara Valera-Muros ¹, Pedro Merino-Gomez ¹ and Jerry Sobieski²

¹Universidad de Málaga, Andalucía Tech, Edificio de Investigación Ada Byron, Málaga 29071, Spain

²NORDUnet, Nordic Gateway for Research and Education NORDUnet, Copenhagen, Denmark

Correspondence should be addressed to Pedro Merino-Gomez; pedro@lcc.uma.es

Received 21 December 2018; Accepted 4 April 2019; Published 23 April 2019

Guest Editor: Pinar Kirci

Copyright © 2019 Álvaro Rios et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

This paper presents the design options for creating a Pan-European mobile network for research in the context of the European Horizon 2020 EuWireless project. The most likely direction is a platform that makes it easier to create network slices for research. In this context, we identify one promising technology to implement network slicing in 5G networks: the framework GÉANT Testbeds Service (GTS). GTS is currently a production service by GÉANT that offers remote construction and use of virtual testbeds for wired networks mapped to the real GÉANT infrastructure. These GTS-virtualized testbed environments conform to Software Define Networks (SDNs) principles and offer compute, storage, and switching resources, at scale and with line rate performance. In this paper, we explain how the current (wired oriented) GTS can be extended with the 5G components, such as radio access nodes (gNBs), transport networks, user devices, etc., in order to implement 5G network slices. Our first conclusion is that using GTS for EuWireless implementation is feasible, dramatically increasing the potential impact of this service in the research community.

1. Introduction

Traditionally, experimentation in computer networks has a low barrier entry. The equipment needed to set up complex architectures is inexpensive, and the algorithm to control communication between peers is usually open and flexible to test new communication protocols and processes. Moreover, with the adoption of virtualization and Software Defined Networks (SDNs) techniques, researchers can create realistic large-scale networks with hosts behaving identically to physical ones [1]. When we move to wireless, Wi-Fi access technology devices offer the same Application Programming Interface (API) to the upper layers of the network as the wired ones; so the interface that controls access to the lower layers can also be easily modified to test new configurations, maintaining this barrier entry at its lower level [2].

In contrast, the barrier entry in experimentation with cellular mobile networks is high, mainly due to two reasons. Firstly, obtaining the required equipment and software to set up a mobile network for experimentation is arduous because

of the costs of acquaintance and the hesitance of manufacturers to sell the components to anyone other than commercial operators, as it can disrupt their business model. Secondly, and more important, research in mobile networks requires access to state-regulated radio spectrum in order to recreate realistic environments and to use commercial equipment as part of the new developments. Governments worldwide use public auction schemes for the frequencies needed to operate, and commercial operators are normally the only ones capable of bidding for their use. Even when the experimenters have access to the right equipment, they can only set up a complete network by isolating it in radio-frequency shields to prevent radiation outside the boundaries of the laboratory, greatly diminishing the value of the tests [3].

Regarding the lack of access to the basic infrastructure needed to perform experiments, it was recently lessened by the appearance of the Software Defined Radio (SDR) technology. SDR allows recreating base stations and other radio elements of the network with commercial off-the-shelf

(COTS) components placed inside a computer. In addition, several open-source projects [4, 5] were created to provide the elements of the core network by implementing the published standards directly in software. Nevertheless, SDR technology is not able to achieve the radio performance of the real equipment, and the different software solutions are still unstable or do not provide all the characteristic of the commercial ones.

Recognizing that mobile technology research is a key component of scientific prowess, different initiatives try to provide researchers with access to commercial equipment outside the umbrella of telecommunications operators. The European Commission, as part of research programs like FIRE, has set up different technology laboratories, i.e., NITOS in Greece or PerformNetworks in Spain [6], equipped with state-of-the-art components that are offered to researchers and businesses interested in this kind of experimentation. Although these projects allow basic research in the mobile networks field, they still have a major caveat. The experiments are linked to fixed laboratories, and the mobility part of wireless technology research is limited or has to be simulated by other means [7, 8].

In this context, the EuWireless project [9] is envisioned with the purpose of helping researchers and experimenters to test new developments in a live, realistic mobile network across Europe. The main objective of the project is to create a research environment capable of providing access to the usually hidden internal mechanism of the networks and also coexisting with commercial operators, without interfering with commercial exploitation. In that direction, this paper addresses how to implement 5G network slices for research at a large scale reusing a mature technology for automatic creation of network testbeds, namely, the GÉANT Testbed Service (GTS) [10, 11]. GTS is already being used by the GÉANT research community to seamlessly create virtual networks between different locations across Europe, and, with some modifications, it could be able to integrate radio resources, spectrum sharing, and mobile network entities into a single homogeneous user interface. In the paper, we introduce the design options for EuWireless and analyze how to map EuWireless concepts to GTS technology. This analysis confirms the technical feasibility of implementing EuWireless 5G slices in the midterm.

The rest of the paper is organized as follows. Section 2 lists the design principles and goals that the EuWireless architecture must fulfil to be differentiated from other projects. Section 3 compares different high-level architectures to be used as the foundation of the EuWireless infrastructure, briefly summarizing pros and cons of each one. A technology capable of coordinating all the components to be used in the initial architecture proposed, the GÉANT Testbeds Service (GTS), is detailed in Section 4. In Section 5, we present the most relevant related work that is currently ongoing, and, finally, conclusions are addressed in Section 6.

2. EuWireless Design Principles

The main goal of the EuWireless project is to design and build a comprehensive architecture capable of providing the

researcher community, experimenters, and end-users access to low-level 5G resources, network capabilities, and configuration options by combining a pool of heterogeneous resources into a homogeneous interface, without the constraints of a laboratory setting or the rigidity of commercial environments. This section introduces the design principles of EuWireless and the main concepts to be considered as inputs to design the architecture, as discussed in the next section.

2.1. Design Principles. EuWireless has identified the following design principles that the final architecture must meet to differentiate itself from other existing projects:

- (i) Support for concurrent researchers in their own isolated network: The framework that controls the infrastructure must be able to accommodate several experiments in parallel, sharing resources from the physical network, but without interfering with each other.
- (ii) Cross-country deployment: Although the EU has invested in the spread of the European research infrastructure, many sites are concentrated around certain universities or technological areas. EuWireless aims to be able to offer experimentation capabilities near researchers across Europe.
- (iii) Scalable Point of Deployment as the main core object: Instead of relying on a centralized control of all the infrastructures deployed in Europe, it is desirable to build a distributed architecture that provides local high-performance and great scalability.
- (iv) Adaptable to integrate new mobile technologies: the design must be flexible enough to accommodate new standards and paradigms in the future, without relying too much on current, fixed implementations.
- (v) Fully automated: researchers must be able to define, reserve, and use the resources needed for any given experiment in an autonomous way, whereas EuWireless must maintain the separation between experiments and a fair user policy to prevent resource exhaustion from greedy experiments.

2.2. Network Slicing. The key enabler to meet the project objectives and design principles is the ongoing convergence between traditional wired networks and mobile ones, which is already a goal in the recently published Fifth Generation (5G) standard. This convergence is known as Network Slicing (NS) and proposes the adoption of virtualized network paradigms and new schemas of resource sharing between the different stakeholders that provide network services to the end user [12]. NS technology combines a set of diverse resources such as radio spectrum, network capacity, or processing power and offers them for a specific purpose or service [13]. This concept can go further, allowing the aggregation of virtual resources with physical ones in a

transparent way from the point of view of the service being offered [14]. Thus, a telecommunications operator is able to create core network instances and reserve them for the exclusive use of a group of users or kind of traffic, guaranteeing certain Key Performance Indicators (KPIs) regardless of the network load. These KPIs are defined for each use case based on the needs of the different vertical industries, as analyzed by the GSMA in [15].

The key concept in NS is the complete separation of resources to be used by different services, in the understanding that this separation is a logical one and the physical resources are used concurrently by all the users but with different priorities. An example of this kind of resource sharing is depicted in Figure 1.

Thanks to the slicing concept, EuWireless can offer each experimenter a network composed of real resources, by using an abstraction layer that distributes those resources in a differentiated way, that is, completely isolated from networks used in other experiments. A remarkable advantage of network slicing is that the same principles can be applied to resources that are not owned by the same company or operator providing the service. This offers the network operator the opportunity of expanding the coverage area and the range of services that could be included in the networks provided for experimentation by entering into agreements with commercial operators with key resources, such as spectrum, link capacity, or raw deployments in a wide area.

2.3. 5G Slices for Research as a Service. Expanding on the design principles stated previously, the prime goal is to deliver a platform that would provide the research community with an end-to-end solution to perform experiments in a realistic way, with commercial equipment but without the constraints of maintaining a highly reliable commercial operation.

In order to do this, the idea is to allow the experimenters to reserve the specific resources needed, run the experiment and free the resources afterwards. In each experiment, the researchers would see a complete 5G operator, from the radio hardware communicating with the UE to the network level, including the core network and any Multiaccess Edge Computing (MEC) services they may need to perform all kinds of experiments [16]. The project should be able to create a private network for each of the experiments in a way that would not interfere with one another. From the low-level radio access to the components that route the user's traffic through the operator network to other users and networks, each experimenter will perceive a complete "telecommunication infrastructure" at his or her disposal, regardless of the networks deployed by other experimenters. Figure 2 shows the entities and logical relations of a complete 5G network, from the User Equipment (i.e., the smartphone or modem), to authentication and network access.

A traditional approach would involve replication of these components for each researcher and compartmentalization to ensure isolation. However, network slicing enables this isolation in a seamless way without the need of duplicating the resources for each experiment. Moreover,

the combination of the network components in a virtual environment and the proficient allocation of the resources needed to perform each task of the experiment will result in the possibility of creating an ample number of "virtual" networks, always maintaining a strict separation between them (Figure 3).

2.4. Distributed Deployment of EuWireless. Another goal of the project is to have it distributed, instead of being physically concentrated in a laboratory or a testbed, with the intention of providing equal opportunities for access and performance to researchers located across Europe. If the resources are located in a single Point of Presence (PoP), this goal would be impossible due to unavoidable physical constraints that will decrease data flow performance. Thus, the main components of EuWireless must be decentralized and decoupled in nature, making it possible to expand the coverage of the service just by deploying one set of hardware and software at each desired location. Additionally, it is desirable for these services to be easily connectable with already deployed installations, so as to compose a mesh network of nodes capable of providing the same kind of features to the experimenters.

This leads to the concept of PoP as the core object in the EuWireless infrastructure. As mentioned above, there is a need to provide services geographically close to where the experiment is running. A PoP would consist in the set of hardware and software that is capable of configuring and managing the slices created for each of the experiments, with the capacity of running as a single node or as a part of a network of PoPs. PoPs must be able to provide services isolated in a certain geographical or logical area, but also to interconnect in a seamless and decentralized way. All the software running in the PoP must be able to interact with the entities present in the current deployment, hardware or software, as well as with the ones present on other PoPs. All of this will naturally provide scalability, since an overloaded deployment can be scaled up just by adding processing power or network links, or even installing new PoPs in areas with a great demand of services to increase performance. As a result, there will be a homogeneous interface with all the resources available to be used, even if the PoP located close to the researcher does not have the resources needed by the experimenter. Furthermore, the project can grow organically from the initial pilots to large deployments running an ample range of experiments concurrently by just connecting new PoPs, as shown in Figure 4, avoiding the reconfiguration of the underlying infrastructure.

The EuWireless project aims to build a future-proof framework where new technologies are tested in conjunction with the research on the current one. Any architecture chosen to provide the services of a project like EuWireless should be sufficiently abstract or expandable to accommodate new paradigms to be tested without fixating in the rigid structure of a particular generation of mobile networks. The following section discusses some design options for this architecture.

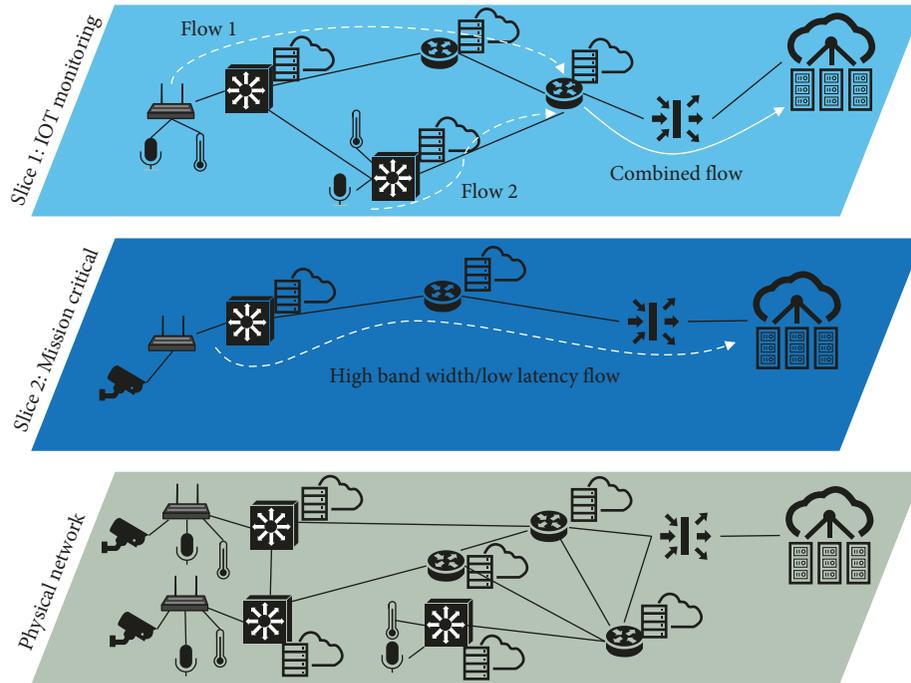


FIGURE 1: Different slices over the same physical topology.

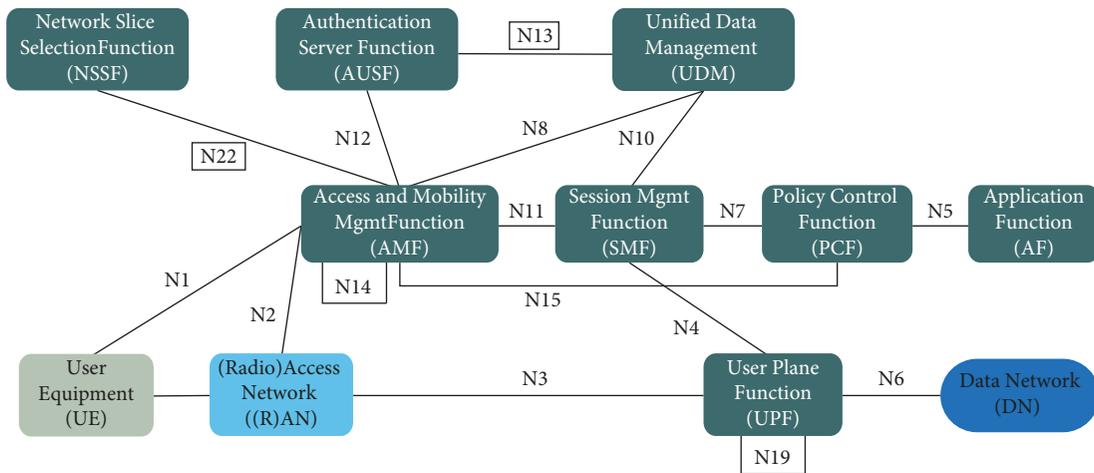


FIGURE 2: Entities in a 5G network.

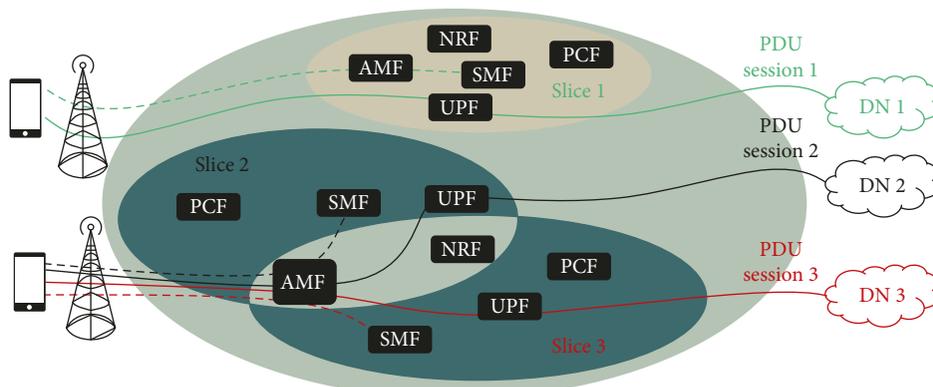


FIGURE 3: Network slicing in 5G networks.

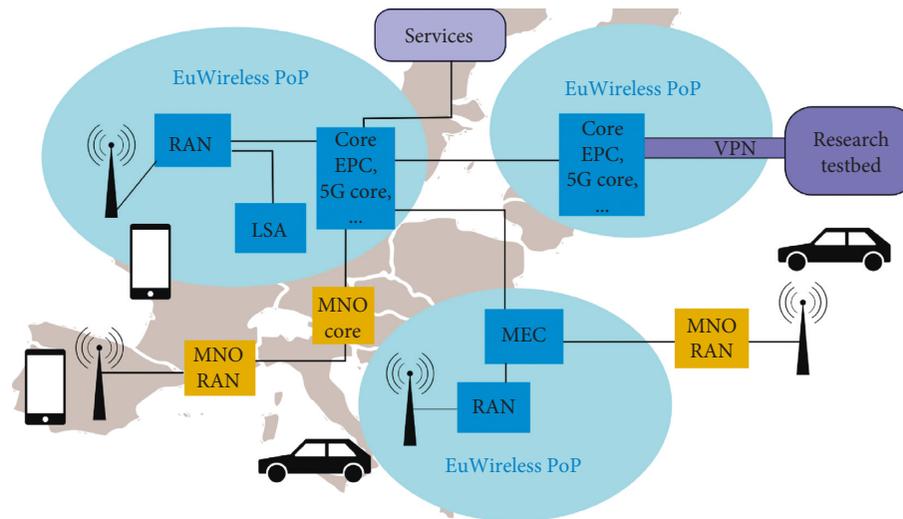


FIGURE 4: EuWireless high-level deployment overview [9].

3. Design Options for EuWireless

Several reference architectures are considered in order to implement the functionality to be offered by EuWireless, combining private and public infrastructures, which will also drive the kind of access and resources to be offered to the researchers.

3.1. Creation of a Full 5G Operator. The first option is to create a full 5G operator to provide the kind of network for experimentation with all the resources owned by EuWireless. This would give the project complete freedom in the configuration and selection of the elements that compose a particular slice. Any kind of end-to-end experiment can be created focusing on any layer of the stack, and the researchers will have direct access to monitor the internals of the architecture to gather the raw performance data from within the network.

The main drawback of this approach is the cost, both in economic and in human resources, needed to complete the deployment of such an operator across Europe. Moreover, additional regulatory constraints might hamper completeness of the service provided: to be able to operate in a multicountry environment EuWireless would need to comply with local legal requirements in each territory in which it operates [17], as well with European regulations regarding the security and the privacy of the data passing through its network, even when these data are generated automatically by other systems [18]. In addition, even if all the infrastructures are owned by the project, there are still limitations on the kind of access that can be provided to low-level network capabilities in order to maintain the operational needs of the whole network.

A related, more realistic approach to this architecture would be one in which EuWireless owns the core of the network where experiments take place but relies on external operators to ensure network capacity and last-mile access in areas where it cannot provide radio coverage. Figure 5

presents the schematic architecture for this kind of infrastructure.

3.2. Management of External Resources. The second approach is the opposite of the first one. In this case, EuWireless would not own any physical infrastructure apart from the minimum required to provision, configure, and manage the slices and could be considered a Virtual Mobile Network Operator (vMNO) that provides services on top of the physical resources of a single commercial operator, or several of them. EuWireless would lease/rent 5G slices of the commercial operators and offer them to the researchers to test new services. By entering into an agreement with operators of different countries, it would be possible to extend EuWireless coverage as far as desired, using PoPs near the point of connection with the carriers to provide access to the EuWireless network and the slice-monitoring facilities.

The main problem of this architecture is that it completely depends on the resources of external operators, which might not be available for the entire duration of the experiment or for which the Quality of Service (QoS) offered may suddenly change depending on the commercial operation of the physical infrastructure. Another constraint would be, most probably, that slices would be limited to 5G resources, as it is the only technology that supports the concept of slicing, and research in the interconnection with other technologies would require external capabilities outside the scope of the project.

Nevertheless, the openness of the configuration is attractive since it enables creating proprietary slices with advanced features and minimal costs, as the network elements are already being used in a commercial operation, even when this flexibility depends on the agreements with the commercial provider, which is ultimately responsible for limitations on the usage of its network, and may or may not be eager to offer the concrete resources experimenters would like to use in some deployments. Figure 6 shows the

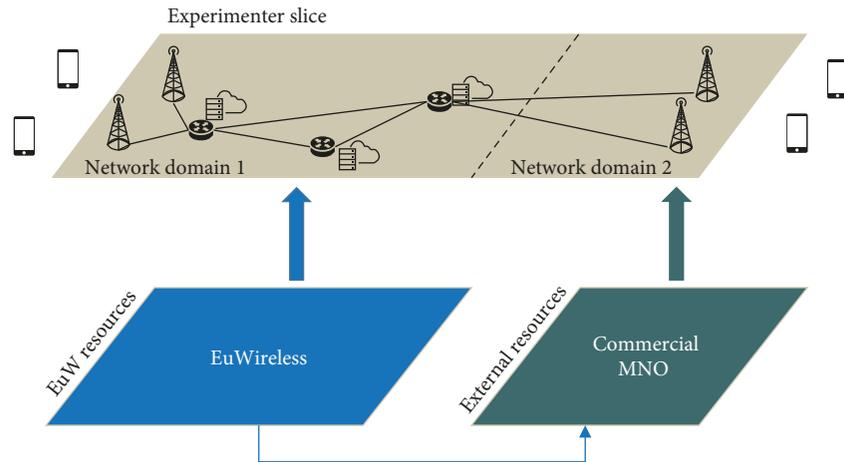


FIGURE 5: First architecture considered for the project. EuWireless has resources with extended coverage provided by commercial MNOs.

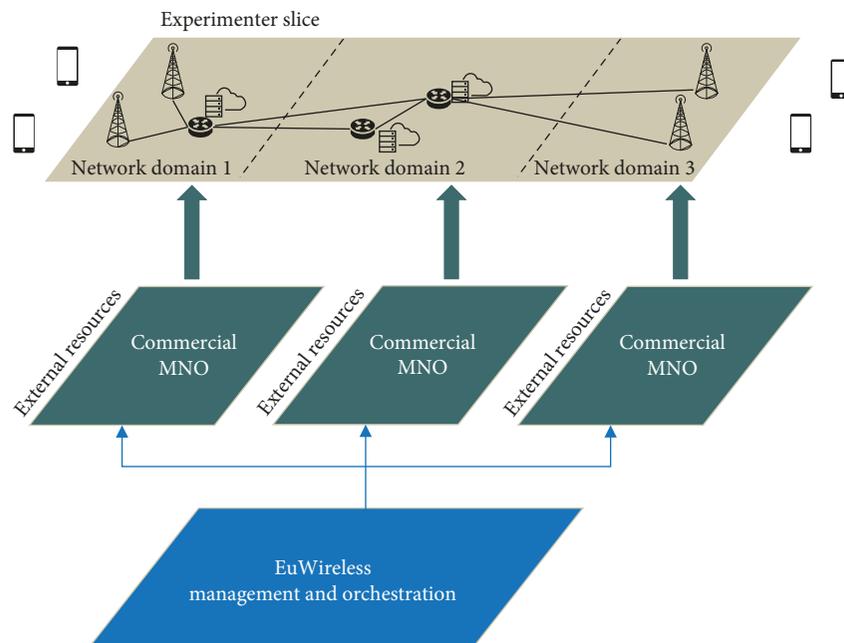


FIGURE 6: Second architecture considered for the project. EuWireless manages resources of commercial MNOs.

only component required, Management and Orchestration, coordinating resources of different operators to create a slice for a particular experiment.

3.3. Network Slices Virtualization and Management. The previous architectures have the issue of only being able to provide 5G slices and, even in the case when the infrastructure is completely owned and controlled by EuWireless, the inability to access low-level components of the network. One solution would be to use the slices created as the foundation of a complete network. Ultimately, a network slice defines the links between entities with different characteristics and the performance expected when they are interconnected. However, the state of the art of real networks is moving towards virtualization, and these entities could be perceived as virtual objects; so, in fact, any software can be

deployed and any network infrastructure can be defined in the slices to interconnect them. By adding a new level of abstraction to the resources, it is possible to present different elements aggregated as “raw slices” on top of which new services can be deployed, treating different resources of different operators in a homogeneous way.

Moreover, with this concept of raw connectivity slices, researchers are not strictly bound to experiment in 5G technology because additional network elements, even entities from previous generations, can be instantiated as Virtualized Network Functions (VNF) [19]. With this architecture (Figure 7), some of the constraints of the two previous proposals are reduced: for example, the limitation of only experimenting in a 5G environment is lowered because the researcher can deploy any kind of “virtualized” environment, previous generation entities can be deployed as virtual machines using the computer processing power of

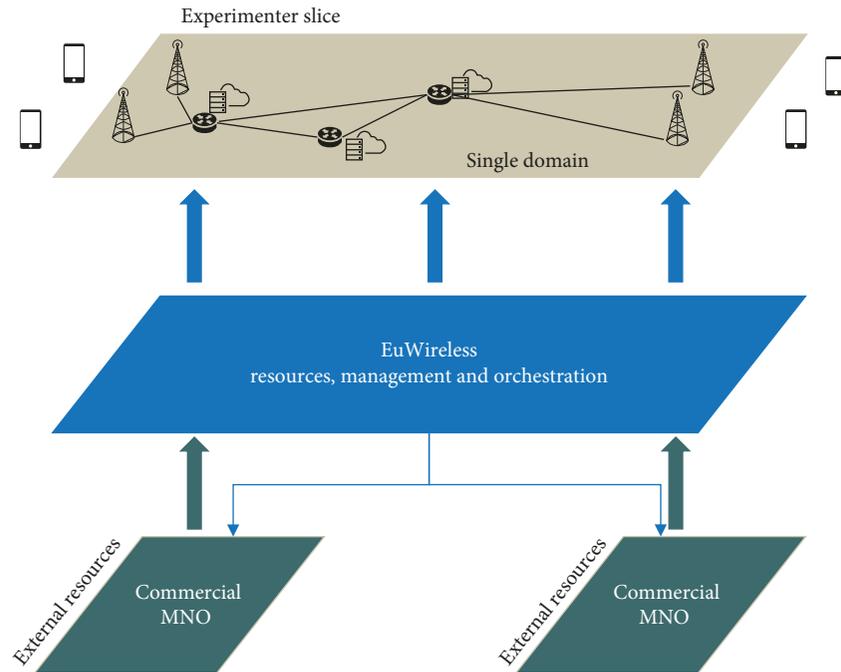


FIGURE 7: Third architecture considered for the project. EuWireless is a “virtual slice” provider.

the slice, and connected with 5G entities, Wi-Fi access points can be attached in any link of the network as an IP device [20], and so on. In addition, the architecture “seen” by the experimenter will not depend directly on the network status of the external operators, as the EuWireless coordination layer will manage the resources during the experiment lifecycle, being capable, for example, of selecting different links if the conditions of the underlying physical network worsen.

In this approach, the only thing outside the scope of the experiment will be the infrastructure that allows creating the slice, such as the MANO, or the slice manager used to coordinate the resources provided by the different actors. But even in this case, if the researcher needs those components for a specific experiment, they can actually be installed as virtual appliances to control another virtual layer on top of the raw one, although this kind of virtual-on-virtual slices imposes a nonzero performance penalty that must be analyzed before undertaking the complex task of deploying a scenario of this type.

Another benefit of this approach is hinted in Figure 7: in the two previous architectures, the traffic had to pass through the different resource providers’ network domains, whereas in this one, a single domain can be created for each slice and experiment.

4. Role of GTS in the Implementation of 5G Slicing

The third architecture, chosen as the basis of the EuWireless infrastructure, allows great flexibility in the selection and configuration of the network elements, but a key component is still missing: the coordination layer that manages and orchestrates all the different resources into a single unified

slice to be offered to the experimenter, ensuring link quality and network isolation between different experiments. During the evaluation of available standards and technologies for this purpose, a new architecture developed by the European research network GÉANT emerged, namely, the GÉANT Testbeds Service (GTS) [10, 11]. This section explains how to use GTS to implement the selected architecture.

4.1. GÉANT Testbed Service. The GÉANT Testbeds Service (GTS) is an innovative network virtualization architecture that offers experimenters access to wide area network infrastructure integrated within the GÉANT network footprint. Through GTS, the research community can deploy and refine novel and/or experimental networking concepts “at scale,” using real network components, with real users, under real-world conditions with the addition of Software Defined Networks (SDNs) paradigms and network-based computing, storage, and switching of resources for research networking and distributed applications. This notion of allocating network infrastructure components to particular projects or applications, under the direct control of the research user, is quite similar to the “Network Slicing” in 5G but applied only to wired networks. From the user perspective, GTS offers virtual testbeds, as represented in Figure 8.

In the original GÉANT implementation, these various components were allocated and stitched together by human-mediated operational processes or using different applications and protocols on a project-by-project basis. GTS offers enhancements to this original work with dynamic and automated slice provisioning and employs advanced abstraction techniques to enable a broad range of new network

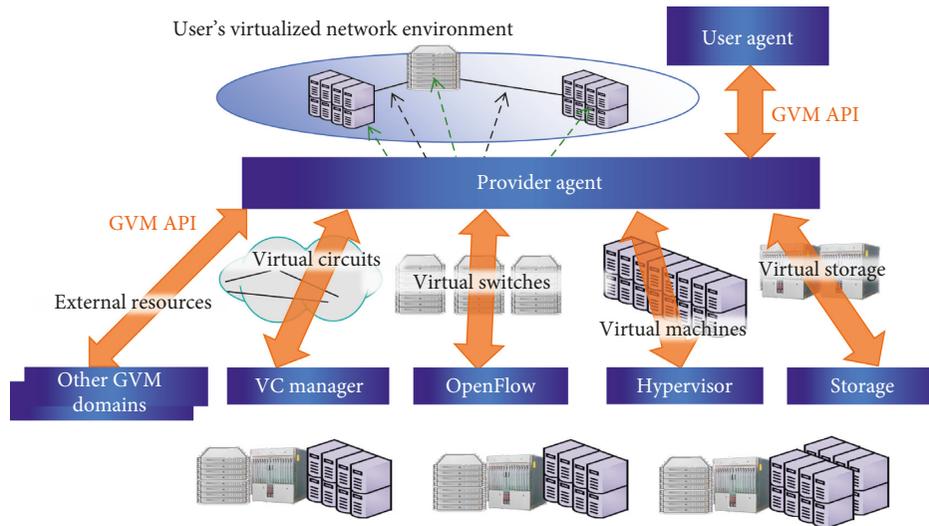


FIGURE 8: Vision of a virtualized network environment provided by GÉANT.

resources to be offered as fully virtualized service objects—such as VMs, virtual circuits, virtual switching/forwarding elements, storage, instruments, sensors, wireless/cellular resources, and other functional network objects.

4.2. Generic Virtualization Model architecture. Since its initial implementation in 2013, GTS has evolved into a sophisticated open-source Software Suite and a Generic Virtualization Model (GVM) architecture. The service is able to dynamically allocate a wide range of SDN-capable network infrastructure elements and provide these resources as insulated and isolated networks, according to the user’s direction and under user control. GTS is also able to dynamically allocate both virtual machines (VMs), in a range of increasingly powerful configurations, and dedicated hardware platforms, or bare metal servers, where the user has access to the entire physical server interconnected with a fully virtualized SDN. To connect all of these virtual “nodal” objects, there are virtual circuits provisioned with hard QoS guarantees. All of these resources can be scheduled or reserved in the time domain and placed spatially within the geographic footprint of the GTS service reachability.

The GVM is conceptually simple: it defines an opaque service domain that offers virtualized resources at user request, meaning that the underlying infrastructure and processes that produce the resources are not exposed (and are not relevant) to the user. Within the GVM architecture, a resource can be anything, and resources that share common characteristics can be grouped into types, or classes. Each resource class defines the attributes that govern the behavior of the object, which users can tune to their needs when requesting such an object. A key aspect of defining resources is specifying how the resource communicates with other resources. This is done by specifying I/O ports and networking capabilities. Ports enumerate data paths from/to each resource instance. These ports will be connected to other ports to create the network slice topology. Using the GTS approach, when requesting a slice,

the user must specify two aspects: the set of resources the experiment requires and the interconnection of those resources relative to one another by specifying port adjacencies. This set of links defines the topology of the network slice, or the data flow graph when the resources are functional objects rather than hardware infrastructure analogs. In this model, all objects that make up the virtual network slice are realized as resources—including virtual circuits. Thus, the network slice is reduced to a set of individual resource objects derived from the original network diagram, and a set of adjacencies that indicate which ports are interfaced to or flow into other ports. Figure 9 shows the real architecture of the virtualized network using the GVM and hints what is the work to be done in EuWireless to expand GTS.

4.3. GTS Experiment Lifecycle. The user interacts with the service through a set of web service primitives that move the resources through a basic lifecycle. The user experiences the resources as virtually the same as the real objects they are intended to model. By hiding these internal allocation and provisioning processes, the service providers are accorded maximum flexibility to use whatever products, external resources, or technologies they prefer. This opaque principle is also the basis of the Network Service Interface (NSI) [21]. This technology-agnostic service model allows the service architecture to scale by allowing providers to retain control of their local infrastructure engineering and design; as long as the service provider delivers the “virtualized” service objects as defined, they are not otherwise constrained to use particular products or technologies.

Each virtual object class requires a means of realizing instances of that object. This “runtime” environment, which supervises the sharing of the underlying infrastructure, is essential for each GVM resource class. In addition to these class-specific runtime services, there are processes that manage the lifecycle of the service object and that each virtual object requires in order to become real. Since the

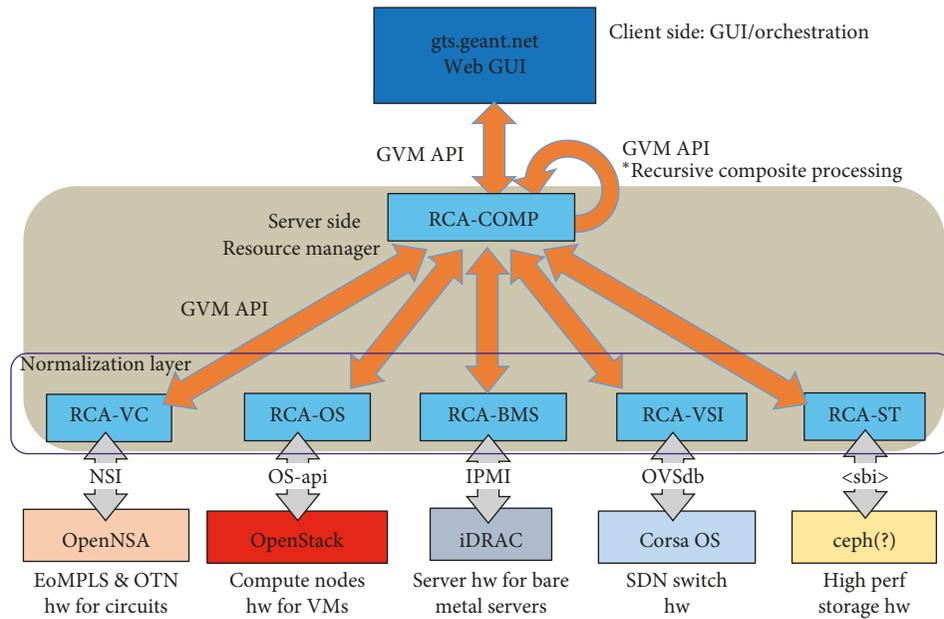


FIGURE 9: GVM.

internal mechanisms required to find and reserve infrastructure for each resource class may be quite different, GTS defines Resource Control Agents (RCAs) as the software modules that implement these processes for each class.

Every resource class must implement the following five basic control primitives, which move the resource, whether physical or virtual, through the states of its lifecycle, shown in Figure 10:

- Reserve(): Instantiate the resource, allocate physical infrastructure components
- Activate(): Configure the hardware to realize the resource, configure the VNF with the required information, move from RESERVED state to ACTIVE state
- Deactivate(): Move from ACTIVE state to RESERVED state, deconfigure VNFs
- Release(): Destroy the instance, shutdown the VNF
- Query(): Get information about the actual state of the resource

4.4. Expanding GTS with 5G components. The EuWireless effort will be focused on the definition and implementation of the lifecycle of 5G network components to be used within the GTS environment. Once their GVM template is complete, we can leverage the already existing traditional network models to create slices with 5G access capabilities, as the additional core network entities of 5G can be deployed as virtualized functions. Figure 11 reflects this virtualization proposed to address network research in the context of GTS and EuWireless. Compared with Figure 9, the radio access node (gNB) and the AMF (one of the key components of the 5G core) are now virtualized to be offered to the researchers in their testbeds.

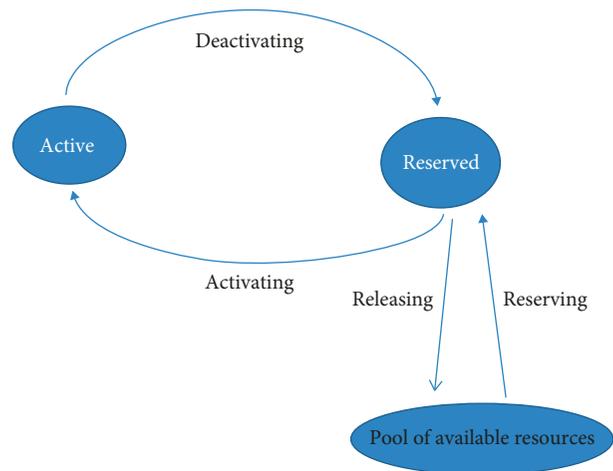


FIGURE 10: GTS objects lifecycle.

Each 5G component will be considered a virtual object in the GTS architecture. As mentioned above, virtual object instances interact with each other by exchanging information across their virtual I/O ports. Thus, the definition of the ports required for each 5G entity is fundamental in the implementation and corresponds with the interfaces of each component. The definition of a virtual object includes the name of the object, the description of its behavior, the ports required, and the set of attributes, where the latter are the parameters the user can tweak in defining the object instance. These attributes characterize the reservation of resources to be taken into account by the system.

Following the previous example, in the case of the virtual object that abstracts the Access and Mobility Management Function (AMF) entity, it will have, at least, eight ports that represent the logical relations with the entities of the 5G network, i.e., N1 with the User Equipment (UE), N2 with the

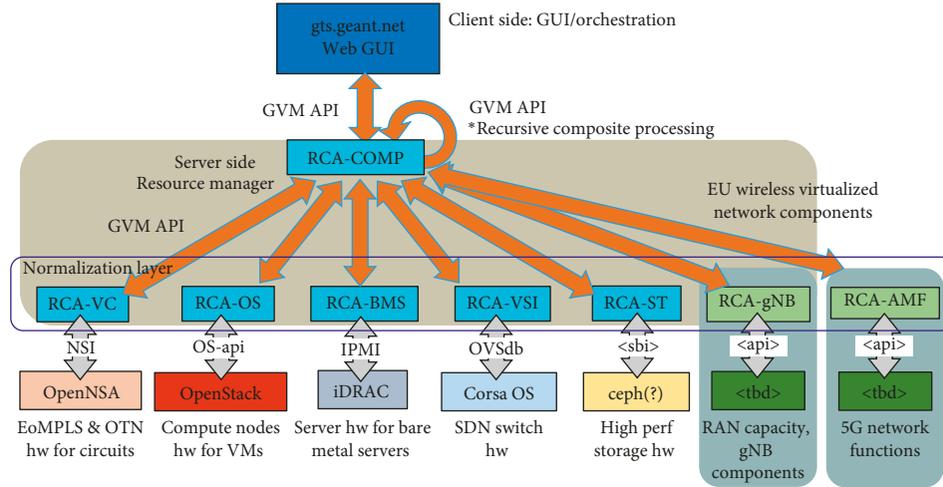


FIGURE 11: GVM incorporating EuWireless RCAs.

(Radio) Access Network ((R)AN), N8 with the Unified Data Management (UDM), N11 with the Session Management Function (SMF), N12 with the Authentication Server Function (AUSF), N14 with other AMF entities in the network, N15 with the Policy Control Function (PCF), and N22 with the Network Slice Selection Function (NSSF). Ports in GTS are linked with virtual circuits that might or might not be distributed and will establish the connection between the virtual object, in this case the AMF entity, and the destination entity or external component. These virtual circuits are able to transport any protocol the entities might use, such as Stream Control Transmission Protocol (SCTP) or User Datagram Protocol (UDP), that are the most used between mobile network components. In addition, there might be optional ports available such as the port connecting the AMF with the control plane through the N_{amf} in case the architecture is based on services, as shown in Figure 12. In this case, only three ports are required in the AMF, two if the UE connects via 3GPP technologies since the N1 will be in this case a logical interface through the physical N2 interface.

In accordance with the control primitives of the object lifecycle introduced in the previous section, after the reservation of resources and the initialization of the element, each entity is contextualized by the software that will configure the element as a network function in the 5G slice. In the case of the RCA that manages the virtual object representing a gNB, the reservation implies to find a compute node sufficient to run the virtual gNB software stack, with port capacity and spectrum capacity. Thus, steps will be reserve the port capacity towards IP network, then ensure that spectrum capacity is available, then reserve compute resources such as memory, number of cores, disk capacity, and finally any other facilities to/from/among physical gNB agents. The activation will spin up the VMs and establish bridges to the phy-gNB interfaces, that is, spin up the v-gNB VM(s) with context, set bridges to physical gNB ports, and configure them for spectrum sharing.

The virtual objects in GTS might be considered atomic, which hide all lower layer constructs, or composite, that enable the specification of complex virtual objects as a

grouping of others. The latter may be parameterized during reservation, contain other composites allowing the existence of hierarchically constructed object instances, have external ports that can be set adjacent to ports of internal children, and be used to further parameterize children object instances. Following this reasoning, the network slice is understood as the root level composite object.

Table 1 shows a summary of the actions to be performed by the main components of the 5G stack in response to the execution of each GVM primitive.

5. Related Work

Experimentation in 5G networks is the main scope of the EuWireless project. However, there are several projects that already offer experimenters different state-of-the-art components to allow basic research in mobile networks. One of them is the European Commission's FIRE program. This program provides technology laboratories equipped with components to perform mobile networks research. However, the experiments must be executed in the laboratory, and thus wireless technology is usually simulated by other means. Similarly, in the US, the GENI program offers large-scale Internet testbeds and in 2013 funded the project SciWiNet [22] to support wireless networking systems research based on the vMNO model. SciWiNet offers cellular data services via Sprint's 3G, WiMAX, and LTE networks. Thus, a wireless testbed for the academic research community is provided, enabling research in areas such as eHealth, intelligent transportation systems, smart buildings and structures, homeland security, and Internet of Things, although the underlying infrastructure is not under investigation.

Recently, the European 5G PPP Initiative launched three projects as part of Horizon 2020 that will address the challenge of creating 5G end-to-end facilities. One of these projects, 5GENESIS [23], will deploy five large-scale platforms to compose a pan-European test platform addressing multiple vertical use cases. Thus, the architecture will be owned by each project, with an approach completely opposed to the SciWiNet vMNO.

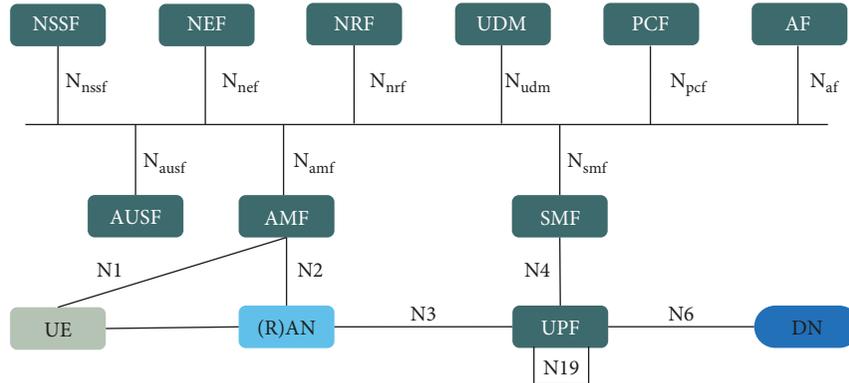


FIGURE 12: Entities in a 5G network service-based architecture.

TABLE 1: Actions in response to primitives.

GTS primitive/5G component	User equipment (UE)	gNB	AMF	UPF	AUSF/UDM	SMF
Reserve	Locate one specific UE and connect	Ensure that the gNB is available on the testbed	Instantiate AMF VNF	Instantiate UPF VNF	Instantiate AUSF VNF	Instantiate SMF VNF
Activate	Attach to 5G slice	Activate N2, N3 interfaces to the correct AMF, UPF and start PLMN broadcast	Connect interfaces N8, N11, N12 with the rest of the components	Connect interfaces N4, N6 with SMF and DN	Register user credentials and keys, connect to AMF	Configure UE IP address pool, traffic steering, connect to AMF/UPF
Deactivate	Detach from 5G slice	Close logical association with AMF/UPF	Disconnect from other entities	Disconnect from other entities	Disconnect from other entities	Disconnect from other entities
Release	Free from testbed	Return resource to the pool	Shutdown VNF	Shutdown VNF	Shutdown VNFs	Shutdown VNF
Query	Get connectivity status and measurements	Get logical association status, number of UE connected, measurements	Get connection status	Get connection status	Get connection status, registered users	DL data notification, traffic routing

In comparison with the above-mentioned projects, EuWireless proposes an architecture that combines both solutions by using virtualized network slices that will reserve the resources required to perform the different experiments from a pool of external and internal resources available, which will be transparent to the researchers due to the implementation of an abstraction layer.

There are also less theoretical projects in which slices are not only described but also created. In this context, the main goal of the MATILDA project [24] is to provide clear interfaces for the multisite management of cloud/edge computing and Internet of Things (IoT) resources. For this purpose, the project proposes a three-layer architecture: the development environment layer, the 5G application orchestrator, and the programmable 5G infrastructure. Thus, MATILDA includes the design and development of 5G-ready applications on top of network slices that are application-aware and can lead to optimal application execution [25].

The 5Gtango project aims at improving the flexibility and programmability of 5G networks. Thus, 5Gtango proposes the creation of a platform with validation and verification mechanisms for VNFs and Network Services, which

is vendor-independent and owns an orchestrator compatible with existing Virtual Infrastructure Managers (VIMs) and SDN controllers [26].

The SLICENET project [4] focuses on the deployment of real end-to-end slicing in virtualised multidomain, multi-tenant 5G networks. The project targets three main use cases: the smart grid oriented to the energy vertical, the eHealth with connected ambulances, and the smart city use case. In [5], SLICENET presents results for a disaggregated Ultra-Dense Network deployment with slicing capabilities based on the RAN runtime.

The 5G-EmPOWER open platform [27] targets an open ecosystem where new 5G services can be tested in realistic conditions. Currently, the platform supports both Wi-Fi and LTE RAN, and developers are able to fully visualize the state of the network and deploy and orchestrate network services dynamically.

6. Conclusions

This paper presents different infrastructure choices that can be used in the final design of the research network proposed

- [17] B. Custers, F. Dechesne, A. M. Sears, T. Tani, and S. van der Hof, "A comparison of data protection legislation and policies across the EU," *Computer Law & Security Review*, vol. 34, no. 2, pp. 234–243, 2018.
- [18] S. Wachter, "Normative challenges of identification in the internet of things: privacy, profiling, discrimination, and the GDPR," *Computer Law & Security Review*, vol. 34, no. 3, pp. 436–449, 2018.
- [19] ETSI, *Network Functions Virtualisation, An Introduction, Benefits, Enablers, Challenges & Call for Action*, ETSI, Sophia Antipolis, France, 2012.
- [20] M. Richart, J. Baliosian, J. Serrat, J.-L. Gorricho, and R. Agüero, "Slicing in WiFi networks through airtime-based resource allocation," *Journal of Network and Systems Management*, 2018.
- [21] T. Kudoh, G. Roberts, and I. Monga, "Network Services Interface: an interface for requesting dynamic inter-datacenter networks," in *Proceedings of the Optical Fiber Communication Conference and Exposition and the National Fiber Optic Engineers Conference (OFC/NFOEC)*, Anaheim, CA, USA, March 2013.
- [22] M. Berman, J. S. Chase, L. Landweber et al., "GENI: a federated testbed for innovative network experiments," *Computer Networks*, vol. 61, pp. 5–23, 2014.
- [23] H. Koumaras, D. Tsolkas, G. Gardikis et al., "5GENESIS: the genesis of a flexible 5G facility," in *Proceedings of the 2018 IEEE 23rd International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD)*, pp. 1–6, Barcelona, Spain, September 2018.
- [24] P. Gouvas, A. Zafeiropoulos, C. Vassilakis et al., "Design, development and orchestration of 5g-ready applications over sliced programmable infrastructure," in *Proceedings of the First International Workshop on Softwarized Infrastructures for 5G and Fog Computing (Soft5 2017), Co-Located with the 2017 29th International Teletraffic Congress (ITC 29)*, pp. 13–18, Genoa, Italy, September 2017.
- [25] P. Gouvas, A. Zafeiropoulos, E. Fotopoulou et al., "Separation of concerns among application and network services orchestration in a 5G ecosystem," in *Proceedings of the 2018 European Conference on Networks and Communications (EuCNC) Workshop "From Cloud Ready to Cloud Native Transformation: What It Means and Why It Matters"*, Ljubljana, Slovenia, June 2018.
- [26] C. Parada, J. Bonnet, E. Fotopoulou et al., "5GTANGO: A beyond-MANO service platform," in *Proceedings of the IEEE European Conference on Networks and Communications (EuCNC)*, Ljubljana, Slovenia, June 2018.
- [27] R. Riggio, M. K. Marina, J. Schulz-Zander, S. Kuklinski, and T. Rasheed, "Programming abstractions for software-defined wireless networks," *IEEE Transactions on Network and Service Management*, vol. 12, no. 2, pp. 146–162, 2015.



Hindawi

Submit your manuscripts at
www.hindawi.com

