THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Network Slicing in 5G: Survey and Challenges

Citation for published version:

Foukas, X, Elmokashfi, A, Patounas, G & Marina, MK 2017, 'Network Slicing in 5G: Survey and Challenges' *IEEE Communications Magazine*, vol. 55, no. 5, pp. 94-100. <https://doi.org/10.1109/MCOM.2017.1600951>

Digital Object Identifier (DOI):

[10.1109/MCOM.2017.1600951](https://doi.org/10.1109/MCOM.2017.1600951)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

IEEE Communications Magazine

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Network Slicing in 5G: Survey and Challenges

Xenofon Foukas*, Georgios Patounas†, Ahmed Elmokashfi†, Mahesh K. Marina*

*The University of Edinburgh, Edinburgh, United Kingdom

†Simula Research Laboratory, Oslo, Norway

*{x.foukas, mahesh}@ed.ac.uk, †{gpatounas, ahmed}@simula.no

Abstract—5G is envisioned to be a multi-service network supporting a wide range of verticals with a diverse set of performance and service requirements. Slicing a single physical network into multiple isolated logical networks has emerged as a key to realizing this vision. This article is meant to act as a survey, the first to the authors’ knowledge, on this topic of prime interest. We begin by reviewing the state of the art on 5G network slicing and we present a framework for bringing together and discussing existing work in a holistic manner. Using this framework, we evaluate the maturity of current proposals and identify a number of open research questions.

I. INTRODUCTION

Mobile devices have become an essential part of our daily lives and as such the mobile network infrastructure that connects them has become critical. It is set to take on an even bigger role with the next-generation 5G mobile systems envisioned to support a wide array of services and devices. In this article, we consider the architectural aspect of mobile networks looking towards 5G. Examining the evolution of mobile networks until now suggests that the changes across generations have been driven largely by the need to support faster data oriented services. For instance, spectral efficiency in the Radio Access Network (RAN) has increased by a factor of 30 from 2G to 4G. On the Core Network (CN) front, the packet switched (IP) component introduced initially in the 2.5G (GPRS) system eventually supplanted the legacy circuit switched component in 4G systems.

What 5G systems are going to be is yet to be determined. However, it is conceivable that the eventual 5G system would be a convergence of two complementary views that are currently driving the research and industrial activity on 5G. One is an *evolutionary* view focusing on significantly scaling up and improving the efficiency of mobile networks (e.g., 1000x traffic volume, 100x devices, and 100x throughput). Much of the research focus around this view is on the radio access side looking at novel technologies and spectrum bands (e.g., massive MIMO, millimeter waves).

The other *service-oriented* view envisions 5G systems to cater to a wide range of services differing in their requirements and types of devices, and going beyond the traditional human-type communications to include various types of machine-type communications. This requires the network to take different forms depending on the service in question, leading naturally to the notion of *slicing* the network on a per-service basis, the focus of this article. Realizing this service-oriented view requires a radical rethink of the mobile network architecture to turn it into a more flexible and programmable fabric,

leveraging technologies like SDN and NFV, that can be used to simultaneously provide a multitude of diverse services over a common underlying physical infrastructure. We take this view in this article as it is intertwined with the 5G mobile network architecture, although the evolutionary view also has architectural implications.

We aim to survey the existing work on network slicing in the 5G context and identify challenges remaining to be addressed to make the service-oriented 5G vision a reality. This article is, to the authors’ knowledge, the first survey on this important topic. We start by outlining representative 5G architectural proposals that highlight the crucial role network slicing is expected to play to meet the widely different requirements of various use cases. We then present a generic 5G architectural framework made up of infrastructure, network function, and service layers as well as the cross-cutting aspect of service management and orchestration (MANO). With respect to this framework, we discuss the state of the art on network slicing in the mobile/5G context. This leads us to identifying some key outstanding challenges to realize a slice-able, softwarized 5G mobile network architecture.

II. NETWORK SLICING IN THE 5G ARCHITECTURE

This section highlights the need for a flexible 5G architecture to accommodate diverse use cases; outlines representative architectural proposals that indicate the crucial role network slicing is expected to play; and presents a generic framework that broadly represents various proposals and will be used as the reference for our network slicing literature review in following sections.

A. Use Cases and Requirements

The 5G network is expected to be the basis for a range of verticals and use cases. For example, the ITU and 5G-PPP have identified three broad use case families: enhanced mobile broadband, massive machine type communications and critical communications. Within those, it is possible to define several specific use cases [1] ranging from general broadband access with global coverage to specialized networks for sensors or extreme mobility. The stark differences between these use cases translates to a set of heterogeneous requirements (Fig. 1) that cannot be satisfied by a one-size-fits-all architecture. With this in mind, alternative architectural proposals for 5G have emerged recently to accommodate use cases with diverse requirements; we outline two such proposals next.

Fig. 1. Key 5G use cases and their requirements. In this illustration, further the distance of a requirement from the center, more important it is to the corresponding use case. It is inspired by a similar illustration from ITU [2]. Diverse use cases need to be mapped to suitably tailored network structures. It is therefore vital for a 5G architecture to be flexible to realize different structures as needed.

B. Architecture

NGMN's architectural vision [1] advocates a flexible softwarized network approach. This views network slicing as a necessary means for allowing the coexistence of different verticals over the same physical infrastructure. Initial proposals were limited to slicing the CN, but NGMN has argued for an end-to-end (E2E) scope that encompasses both the RAN and CN. To realize this and provide the needed context awareness, both parts need to be flexibly sliced into several overlaid instances serving different types of users, devices and use cases. This whole process needs to be orchestrated by an E2E MANO entity that has a central role in the architecture.

The overall NGMN architecture is split into three layers: infrastructure resource, business enablement and business application. Realizing a service in this proposal follows a top-down approach via a network slice blueprint that describes the structure, configuration and work-flows for instantiating and controlling the network slice instance for the service during its life cycle. The service/slice instance created based on the blueprint may be composed of several sub-network instances, each in turn comprising of set of network functions and resources to meet the requirements stipulated by the service in question.

5G-PPP's architectural vision [3] offers a more elaborate examination of the roles and relationships between different parts of the 5G network. Overall, 5G-PPP shares the NGMN view that a potential 5G architecture must support softwarization natively and leverage slicing for supporting diverse use cases. 5G-PPP's architectural proposal is divided into five layers: infrastructure, network function, orchestration, business function and service layers. Relating this to the NGMN proposal, while both are built on infrastructure and network function (business enablement) layers, there are a couple of differences: the orchestration/MANO is viewed as a separate layer in the 5G-PPP proposal; and the business application layer in the NGMN proposal is divided into two

Fig. 2. Generic framework representing various 5G architectural proposals. We review and appraise the 5G network slicing literature with respect to this framework.

layers (business function and service) in the 5G-PPP case.

More generally, there seems to be a broad consensus on the need for native support for softwarization and network slicing as a means to realize widely different services in 5G. Moreover, various 5G architectural proposals can be broadly mapped to a generic framework shown in Fig. 2. This framework is composed of three main layers: the infrastructure layer, the network function layer and the service (or business) layer. It also consists of a MANO entity that translates use cases and service models into network slices by chaining network functions, mapping them to infrastructure resources, configuring and monitoring each slice during its life-cycle.

III. STATE OF THE ART ON 5G NETWORK SLICING

In this section, we discuss the existing work on network slicing for 5G using the generic framework presented in the previous section (Fig. 2) as a reference. Table I summarizes this discussion. Fig. 3 gives an overview of the various research issues related to network slicing, and indicates where future research should be focused.

A. Infrastructure Layer

Scope: The infrastructure layer broadly refers to the physical network infrastructure spanning both the RAN and the CN. It also includes deployment, control and management of the infrastructure; the allocation of resources (computing, storage, network, radio) to slices; and the way that these resources are revealed to and can be managed by the higher layers.

Existing work: The related work focuses on two main subjects; the composition of the network infrastructure and its virtualization.

1) *Composition of Network Infrastructure:* It has been advocated that in order to realize slicing we need to move towards the Infrastructure-as-a-Service (IaaS) paradigm [4], where different infrastructural elements covering different requirements can be leased to accommodate the needs of the various slices. This paradigm is well-known in the context of cloud computing, but it needs to be further adapted for the 5G context.

TABLE I
APPROACHES FOR ADDRESSING DIFFERENT ASPECTS OF NETWORK SLICING AND THEIR (DIS)ADVANTAGES

		Advantages	Disadvantages
Virtualization of Radio Resources	Dedicated Resources	Naturally ensures resource isolation, ease of realizing virtual base station stacks and supporting multiple RATs	Inefficient utilization of radio resources
	Shared Resources	More efficient use of radio resources	Requires more sophisticated techniques to ensure isolation of radio resources
Granularity of Network Functions	Coarse-grained	Easier deployment and management of network functions	Less flexible and adaptive to changes in underlying network conditions
	Fine-grained	More flexible and easier to conform to SLAs	Service chaining and interoperability of functions challenging
Service Description	Human-readable format	Easier to express service requirements	MANO role challenging in mapping requirements to network components
	Set of functions and network components	Non-intuitive way to express service requirements	Simpler realization of network slices

Considering the *core network (CN)*, there is a broad consensus in favor of using generic hardware infrastructure for the deployment of virtual network functions [5] [6] [7]. However, due to the differing constraints between various services deployed over the network, a simple centrally positioned cloud infrastructure might not be suitable for all the slices. For example, a sub-millisecond latency requirement of a tactile service such as remote surgery deployed over a dedicated network slice cannot be accommodated if the cloud infrastructure is located far away from the RAN. Therefore, some network slicing architectures [5] [7] propose a mix of central and edge cloud computing infrastructures where resources can be allocated to either of them, depending on the slice requirements.

On the other hand, the *radio access network (RAN)* comprising of multiple base stations is expected to span diverse radio access technologies (RATs) including LTE and Wi-Fi. Moreover, since the slices are expected to be created dynamically with their service requirements not known a priori, the RAN infrastructure needs to be flexible enough to provide support for various RATs on-the-fly, following a RAN as a Service (RaaS) paradigm. This is why a large number of architectural proposals [4] [7] [6] for network slicing expect the deployment of generic software-defined base stations composed of centralized baseband processing units and remote radio heads as the logical next step.

2) *Infrastructure Virtualization*: The ability to virtualize the underlying infrastructure and provide isolation among services is essential for network slicing. This not only means virtualization and full isolation of the underlying resources (processing, storage, network and radio) among slices but also the capability to support different types of control operations over the resources in a virtualized manner based on the service requirements. This characteristic of providing a virtualized end-to-end environment, that can be potentially opened up and fully controlled by third parties, is one of the key features that separates network slicing from the already existing network sharing solutions [5] [4].

Considering the virtualization of the CN infrastructure, research done in the context of cloud computing can be leveraged. Specifically, technologies like Kernel-based Virtual Machines (KVM) and Linux Containers (LXC) can provide isolation guarantees in terms of processing, storage and network resources at the OS or process level. These isolation

guarantees combined with the capabilities offered by platforms like OpenStack for the pooling of resources can greatly simplify the on-the-fly creation of virtualized CNs. Due to the high maturity level of the aforementioned technologies, concrete prototype implementations of slicing frameworks are already available, enabling the deployment of virtual core network functions (virtual MME, virtual SGW, etc.) over cloud infrastructures (e.g., [7]).

On the other hand, virtualization approaches for the RAN are at an early stage. Applying VM and container-based solutions in this domain does not fully address the problem as they do not deal with the additional dimension of virtualizing and isolating radio resources (spectrum and radio hardware). Existing RAN virtualization approaches that account for this dimension fall into one of two categories: (i) providing a dedicated chunk of spectrum for each virtual base station (slice) to deploy a full virtual network stack on top of it [7]; (ii) dynamically sharing the spectrum between different virtual base station instances (slices) by employing common underlying physical and lower MAC layers [8]. The dedicated spectrum approach is easier to implement, especially with dedicated radio hardware per slice, since isolation of radio resources is guaranteed through the static fragmentation of the spectrum but it can result in inefficient use of radio resources. The other approach of fine-grained and dynamic spectrum sharing has the opposite problem of making isolation between slices challenging.

B. Network Function Layer

Scope: The network function layer encapsulates all the operations that are related to the configuration and life-cycle management of the network functions that, after being optimally placed over the (virtual) infrastructure and chained together, offer an end to end service that meets certain constraints and requirements described in the service design of the network slice.

Existing work: The research interest in this layer mainly revolves around the technologies that can act as enablers for the deployment and the management of network functions, as well as around issues regarding the granularity and the type of the deployed functions.

1) *Enabling Technologies*: There already seems to be a consensus among researchers and the industry about the role of SDN and NFV [7] [9] [6] [5]. NFV is an ideal technology for the life-cycle management and orchestration of the network functions, while SDN can inherently act as an enabler of NFV by allowing the configuration and control of the routing and forwarding planes of the underlying infrastructure through standardized protocols (e.g., Openflow).

2) *Granularity of Network Functions*: One particularly interesting aspect of this layer that is thoroughly discussed in various relevant works is the granularity (scope) of the available virtual network functions [5] [10]. On one end, we have coarse grained functions, where each one is responsible for a large portion of the network's operations (e.g., individual functions for LTE eNodeBs, MMEs, S-GWs). On the other end, we have functions with a very fine granularity, where each of the coarse grained functions mentioned above are divided further into many sub-functions. For example, in [11] the LTE EPC is broken into functions responsible for mobility and forwarding traffic (MME, S-GW, P-GW), which in turn are further decomposed into sub-functional entities including signaling load balancers, mobility managers and functions dedicated to the forwarding of either control or data plane traffic.

The coarse grained approach offers a more simplified way of placing and managing the network functions of a slice. However, this comes at the expense of a slice that is less flexible and adaptive to the changes of the underlying network conditions, something that can be critical when the slice needs to conform to specific service level agreements (SLAs) [10]. For example, the radio resource scheduler in a slice might need to be swapped for another one with a different scheduling policy when a large number of mobile devices appear concentrated in a specific location in order to avoid violating the slice's SLA. If the scheduler is tightly coupled and packed as a single function with the rest of the eNodeB, performing the swapping operation can become a challenge. However, it has also been argued that, despite its benefits on the adaptation of a slice to the network conditions, a fine granularity can be problematic for the interfacing and chaining of the network functions since the more network functions that exist, the more interfaces need to be defined for their inter-communication [5]. This is particularly an issue when virtual network functions are made available by third-parties through some kind of a network function store [7], since without common interfaces, their interoperability is not guaranteed. As a workaround, the use of a container-based protocol that will wrap the interface of the contained functions has been proposed [5], however there is no concrete description as to how to achieve this.

C. Service Layer & MANO

Scope: Perhaps the most important element that distinguishes network slicing in the context of 5G from other forms of slicing that have been considered in the past (e.g., cloud computing) is its end-to-end nature and the requirement to

express a service through a high-level description and to flexibly map it to the appropriate infrastructural elements and network functions. This observation regarding the operation of slicing in the context of 5G naturally leads to two new high-level concepts: (1) a service layer that is directly linked to the business model behind the creation of a network slice; and (2) network slice orchestration for the hypervision of a slice's life-cycle.

Existing work: Due to the novelty that this layer introduces in terms of concepts and ideas, the related research in this domain naturally focuses on answering fundamental questions regarding network slicing architectures. More specifically, the topics considered are related to the way that services should be described and how they should be mapped to the underlying network components, and the architecture of network slicing managers and orchestrators.

1) *Service description*: Regarding the service layer and the way that the business model of a service should be described in high-level terms, there are two different proposals. In one approach, the service level description (manifest) is just simply a set of traffic characteristics, SLA requirements (e.g., for performance related aspects like throughput and latency) and additional services (e.g., localization service) [6]. In the second approach, the service description is more detailed in the sense that it can identify specific functions or RATs that are bundled together and should be used for the creation of the slice [7] that provides a specific service (slice as an application). The main difference lies in the way that the network slice will be generated. In the first case, the slice orchestrator will be assigned the more complex task of identifying the appropriate functions and technologies that will guarantee the fulfillment of the requirements described in the slice's manifest, while in the second case things are more simplified since the required building blocks of the slice are already identified in its description. However, the second approach can be less efficient as it leaves less flexibility to the slice orchestrator to tune the components of the slice.

2) *MANO architecture*: The exact form that the network slicing management and orchestration entity should have is still unclear with different works presenting different ideas. Some proposals envision that network slicing will come through an evolution of the current 3GPP standards and therefore propose enhancements in terms of interfaces and functionalities for the existing mobile architecture [12]. Others envision a more radical clean-slate approach where the slice management and orchestration will be implemented as an application over an SDN controller which will oversee both the wired and the wireless domain [4] [5] [9]. Concrete implementations of MANO reference frameworks such as Open Source MANO¹ (OSM) are already making their appearance, which enable experimental studies on end-to-end 5G network slicing.

3) *Mapping of services to network components*: Another very important issue is how to map and stitch together the

¹<https://osm.etsi.org/>

components that are available to the various layers of the architecture in order to compose an end-to-end slice. Two types of mapping have been considered: (1) the functional/SLA mapping of the service requirements to network functions and infrastructure types; and (2) the mapping of network functions and infrastructure type to vendor implementations [6] [7].

The first type of mapping refers to the way that MANO chooses appropriate high-level network elements that are required to create a slice for a given service in order to meet its functional requirements and SLA. For example, if a slice has a need to cover devices over a wide area without any capacity concerns, choosing an LTE deployment with macro cells might be a good option. For this mapping, it has been proposed that the available infrastructural elements and network functions should reveal their capabilities to the MANO in a form of meta-data, describing the types of services that they can support [6].

Once the type of functions and infrastructural elements required for the slice have been identified, there is a need for a further mapping of these elements to concrete vendor implementations. Depending on the implementation of a function by a vendor, different levels of services can be offered. For example, alternative software implementations of an LTE eNodeB could provide support for a different number of users, with different performance guarantees or even with different capabilities (e.g., flexible modification of the MAC scheduler). Here too, the high level solution to this problem seems to be the use of meta-data in the elements provided by the vendors of the functions and the infrastructure. Such meta-data could describe both the capabilities of the vendor specific functions and hardware [6] [5] as well as their deployment and operational requirements (connectivity, supported interfaces and infrastructural KPI requirements) [6] [7], providing the MANO with sufficient information to perform the best possible configuration for the slice.

IV. CHALLENGES

From the last section, it is apparent that 5G network slicing has already received fair amount of attention from the research community and industry. At the same time, there are several aspects key to end-to-end network slicing that are not well understood as captured by the illustration in Fig. 3. With this in mind, we now elaborate on several significant outstanding challenges that need to be addressed to fully realize the vision of network slicing based multi-service softwareized 5G mobile network architecture.

A. RAN Virtualization

As already discussed in Section III-A2, the main challenges for infrastructure virtualization lie in the RAN. Solutions that pre-allocate distinct spectrum chunks to virtual base station instances (slices) are straightforward to realize and provide radio resource isolation but have the downside of inefficient use of radio resources. The alternative dynamic and fine-grained spectrum sharing based RAN virtualization approach

does not have this limitation and therefore is desirable. However ensuring radio resource isolation is a challenge for this approach. This can potentially be addressed by adapting SD-RAN controllers like [13]. As 5G networks are expected to span multiple RATs (including emerging technologies like 5G new radio and NB-IoT), it is vital for RAN virtualization solutions to be able to accommodate multiple RATs. This presents an additional outstanding challenge as it is unclear whether multiple RATs can be multiplexed over the same possibly specialized hardware or each needs its own dedicated hardware; the answer to this question might depend on the set of RATs under consideration.

From the RAN virtualization viewpoint, realizing the RaaS paradigm is another major challenge over and above the ones outlined above. This is a significant step over the notion of RAN sharing that involves sharing of radio resources among tenants (e.g., MVNOs) via a physical mobile network operator; various solutions for RAN sharing exist (e.g., [8], [13]). RaaS paradigm requires going beyond radio resource and physical infrastructure sharing [4] [9] to have the capability to create virtual RAN instances on-the-fly with tailored set of virtualized control functions (e.g., scheduling, mobility management) to suit individual slice/service requirements while at the same time ensuring isolation between different slices (virtual RAN instances).

B. Service Composition with Fine-Grained Network Functions

Ease of composing a service out of the available network functions is as discussed in Sec. III-B2 directly dependent on the granularity of these functions. Coarse-grained functions are easy to compose as fewer number of interfaces need to be defined to chain them together but this comes at the cost of reduced flexibility for the slices to be adaptable and meet their service requirements. Fine-grained network functions do not have this limitation and are more desirable. However we lack a scalable and interoperable means for service composition with fine-grained functions that could be implemented by different vendors. The straightforward approach of defining new standardized interfaces for each new function is not scalable as the functions increase in number and the granularity becomes finer.

C. End-to-End Slice Orchestration and Management

A significant challenge for realization of a network slice is how to go from high-level description of the service to the concrete slice in terms of infrastructure and network functions. The problem of describing services has already been identified in the literature (Sec III-C1) but without satisfactory resolution. A good approach to address this void is to develop domain-specific description languages that allow the expression of service characteristics, KPIs and network element capabilities and requirements in a comprehensive manner while retaining a simple and intuitive syntax (e.g., in the philosophy of [14]). Two important features that such languages should inherently provide are the flexibility/extensibility to accommodate new network elements that may appear in the future (e.g., new network functions, new RATs) and the applicability

Fig. 3. Maturity level of various aspects of 5G network slicing research.

to be used in multi-vendor environments. A desirable feature would also be the capability to compose complex rules and expressions out of simpler ones, introducing abstraction layers in the expression of service requirements.

As noted in Sec III-C2, concrete MANO frameworks like OSM have emerged in recent years. While such platforms are essential to flexibly realize network slices end-to-end and as needed, there's a more significant challenge that is only starting to be addressed. This concerns holistic orchestration of different slices so that each meet their service/SLA requirements while at the same time efficiently utilizing underlying resources. This calls for a sophisticated end-to-end orchestration and management plane. Such a plane should not be limited to trivial slice generation that does mapping of slices to network components and statically allocates them resources. Instead it should be adaptive, ensuring that the performance and resiliency requirements of the deployed services are met. To achieve this, it should efficiently and holistically manage resources by making decisions based on the current state of slices as well as their predicted state/demands in the near future [4].

Such issues have been thoroughly investigated in the context of cloud computing and data centers, where many concrete

solutions have already been proposed (e.g., [15]). While underlying principles from these other contexts can be leveraged, mechanisms targeting 5G network slicing should be suitably adapted and extended considering additional types of resources. Specifically, not just the resources found in cloud environments (memory, storage, network), but also radio resources need to be included, considering their correlation and how adjusting one resource type could have a direct effect on the efficiency of some other and therefore on the overall service quality. The problem of meeting requirements of different services while efficiently managing underlying network resources in the 5G network slicing context is also somewhat analogous to the quality of service (QoS) provisioning in the Internet.

V. SUMMARY

We have presented what we believe to be the first survey of the state of the art on 5G network slicing. To this end, we present a common framework for bringing together and discussing existing work in a holistic and concise manner. This framework essentially groups existing slicing proposals according to the architectural layer they target; namely the infrastructure, network function, and service layers along with

the MANO entity. With respect to this framework, we evaluate the maturity of current proposals and identify remaining gaps. While several aspects of network slicing at the infrastructure and network functions layers are quickly maturing, issues like virtualization in the RAN are unresolved. Also, approaches for realizing, orchestrating, and managing slices are still in their infancy with many open research questions.

REFERENCES

- [1] NGMN Alliance. 5G White Paper, Feb 2015.
- [2] ITU-R. IMT Vision – Framework and overall objectives of the future development of IMT for 2020 and beyond, Sept 2015.
- [3] 5G PPP Architecture Working Group. View on 5G Architecture, Jul 2016.
- [4] I. F. Akyildiz, P. Wang, and S. Lin. SoftAir: A software defined networking architecture for 5G wireless systems. *Computer Networks*, 85(C), 2015.
- [5] P. Rost et al. Mobile network architecture evolution toward 5G. *IEEE Communications*, 54(5), 2016.
- [6] X. Zhou, R. Li, T. Chen, and H. Zhang. Network slicing as a service: enabling enterprises’ own software-defined cellular networks. *IEEE Communications*, 54(7), 2016.
- [7] N. Nikaein et al. Network store: Exploring slicing in future 5G networks. In *Proc. 10th ACM International Workshop on Mobility in the Evolving Internet Architecture (MobiArch’15)*, Sep 2015.
- [8] X. Costa-Pérez, J. Swetina, T. Guo, R. Mahindra, and S. Rangarajan. Radio access network virtualization for future mobile carrier networks. *IEEE Communications*, 51(7), 2013.
- [9] A. Banchs, M. Breitbach, X. Costa, U. Doetsch, S. Redana, C. Sartori, and H. Schotten. A novel radio multiservice adaptive network architecture for 5G networks. In *Proc. 81st IEEE Vehicular Technology Conference (VTC Spring’15)*, May 2015.
- [10] M. Sama, X. An, Q. Wei, and S. Beker. Reshaping the mobile core network via function decomposition and network slicing for the 5G era. In *Proc. IEEE WCNC*, Apr 2016.
- [11] F. Z. Yousaf, P. Loureiro, F. Zdarsky, T. Taleb, and M. Liebsch. Cost analysis of initial deployment strategies for virtualized mobile core network functions. *IEEE Communications*, 53(12), 2015.
- [12] K. Samdanis, X. Costa-Perez, and V. Sciancalepore. From network sharing to multi-tenancy: The 5G network slice broker. *IEEE Communications*, 54(7), 2016.
- [13] X. Foukas, N. Nikaein, M. M. Kassem, M. K. Marina, and K. Kontovasilis. FlexRAN: A flexible and programmable platform for software-defined radio access networks. In *Proc. ACM CoNEXT*, Dec 2016.
- [14] Q. Zhou, C. Wang, S. McLaughlin, and X. Zhou. Network virtualization and resource description in software-defined wireless networks. *IEEE Communications*, 53(11), 2015.
- [15] D. Shue, M. J. Freedman, and A. Shaikh. Performance isolation and fairness for multi-tenant cloud storage. In *Proc. 10th USENIX Symposium on Operating Systems Design and Implementation (OSDI’12)*, Oct 2012.