WILEY | Hindawi

Research Article

Network-Cloud Slicing Definitions for Wi-Fi Sharing Systems to Enhance 5G Ultra Dense Network Capabilities

Maxweel Carmo,[1,2] Felipe S. Dantas Silva [ID],[2,3] Augusto Venâncio Neto,[2,4] Daniel Corujo,[4] and Rui Aguiar [ID][4]

[1]*Federal University of Mato Grosso (UFMT), Barra do Garças/MT, 78698-000, Brazil*

[2]*Department of Informatics and Applied Mathematics (DIMAp), Federal University of Rio Grande do Norte (UFRN), Natal/RN, 59078-970, Brazil*

[3]*LaTARC Research Lab, Federal Institute of Education, Science and Technology of Rio Grande do Norte (IFRN), Natal/RN, 59015-000, Brazil*

[4]*Instituto de Telecomunicações, Aveiro, 3810-193, Portugal*

Correspondence should be addressed to Felipe S. Dantas Silva; felipedantas@gmail.com

Ultradense Networks (UDNs) seek to scale the 5th-Generation mobile network systems at unforeseen amounts of networks, users, and mobile traffic. We believe that the Wi-Fi sharing service is an asset in expanding 5G UDN capacity requirements for higher coverage and ubiquitous wireless broadband connectivity. However, the limitations of the Wi-Fi sharing pioneer deployment, along with other related works, has led our team to carry out further research. As a result, it was found that FOg CloUd Slicing for Wi-Fi sharing (FOCUS) is a suitable means of expanding 5G UDN capacities. FOCUS applies end-to-end Network-Cloud slice definitions on top of the Wi-Fi sharing technology, with the aim of offering multitenancy and multiservice support for a wide range of services, while meeting carrier-grade requirements and resource control at runtime and making full use of a "softwarized" approach. The feasibility of the FOCUS system is assessed in a real testbed deployment prototype, which allows an accurate view to be obtained of the basic functional principles and system-level proof-of-concept alongside the FON *de facto* Wi-Fi sharing service. The results suggest that FOCUS offers much greater benefits than FON, owing to its capacity to provide end-to-end Network-Cloud Slices while ensuring independent/isolated service delivery with resource adaptation at runtime.

1. Introduction

The increasing integration of Information and Communication Technologies in a broad spectrum of devices has led to an expected massification of information being exchanged in telecommunications networks. As a result, there has been an unforeseen number of connected devices in Ultradense Networks (UDNs) that require the deployment of smaller cells to encompass the immense amount of traffic [1]. In addition, this growth is not confined to a particular use case or type of communication technology but entails a widespread deployment over different scenarios with disparate requirements, ranging from low-latency to high broadband.

The 5th Generation (5G) of Telecommunication Networks [2] extends beyond improvements in isolated performance but enables the network to tackle these different requirements in a flexible and dynamic way. Among the key features of 5G networks, "Softwarization" [3] is highlighted as an outstanding core mechanism, which addresses the needs of different traffic demands through cognition and programmability at run-time. Virtualization is a technological enabler of this "Softwarization," through its ability to orchestrate Cloud Computing (and its extensions), thus resulting in Software-Defined Networking (SDN) [4] and Network Function Virtualization (NFV) [5] as key enablers in a 5G "softwarized" system. The Network Slicing concept [6] has been progressively incorporated in the 5G architecture to provide these flexible services while meeting the requirements of targeted services and applications. Slicing allows the network resources to be partitioned in logical representations, where each slice is regarded as an isolated substrate. This enhancement enables multiple tenants to share

a common physical infrastructure, while allowing high-level service isolation, customization, and monetization for service providers.

There are two ways to implement the slicing concept in 5G systems, namely, slicing at the cloud portion of the network (i.e., the Cloud Slicing (CS) concept) and slicing at the radio access network (RAN) (i.e., the Network Slicing (NS) concept). The former focuses on the slicing of the computing infrastructure of the network service provider, where its network functions are virtualized, whereas the latter is concerned with spectrum frequency, and space or time multiplexing mechanisms. When combined, both slicing concepts provide end-to-end slicing capabilities.

The 5G hybrid system foresees network capacities reaching unprecedented levels to cope with the explosive growth of mobile data traffic and connected devices. In addition to slicing, Fog computing [7] allows the application of cloud principles to network provisioning operations, but, in the end, it is the devices themselves that allow the network to deal with the huge amounts of data generated at the edge. Moreover, 5G is expected to support the rapid wireless network densification and meet the different requirements of emerging service technologies and vertical markets, such as Smart Cities, Internet of Things, eHealth, the automotive industry, energy, food and agriculture, etc. By covering such a wide scope, UDNs represent a new paradigm that promises to deploy short-range, low-power, low-cost, and high-density wireless access. In fact, 5G UDNs will take different forms depending on service demands and thus have to face extreme challenges to meet the huge demand for ultrahigh connection density and high utilization of resources.

The WLAN-sharing facility has emerged as an asset for 5G UDN scenarios, by allowing greater opportunities for external access to wireless broadband through an agreement with Wi-Fi network owners to share a part of their spectrum. WLAN-sharing makes the Wi-Fi slicing concept a reality and provides a single access point for different virtual Wi-Fi networks, as well as ubiquitous wireless connectivity to authorized users. However, although the current deployments (such as the widely used FON (https://fon.com)) make Wi-Fi more widely available for shared facilities, it is not capable of operating under the above mentioned 5G UDN flexible and dynamic provisioning capabilities. As a result, the shared networking fabric provided by these market-oriented solutions only allows traffic isolation at the Wi-Fi medium level. Moreover, it only provides differentiated services for network-layer facilities, while fully sharing the backhaul and carrier cloud resources. By default, the Wi-Fi sharing networking system allows a close to "best-effort" delivery service, with added capabilities that require manual configuration (which is not expected by most home equipment owners), or even not be supported by the system at all. Moreover, the different types of emerging applications require high-level service differentiation that goes beyond the Wi-Fi domain.

In this paper we propose the FOg CloUd Slicing for Wi-Fi sharing (FOCUS) approach, which aims to make a significant improvement in the current Wi-Fi sharing technology through provisioning an end-to-end fully slice-defined environment (from the home network to the network carrier cloud). Moreover, FOCUS provides the network operator with capabilities for high-level networking services as a service, including full isolation, customization, and control at runtime. To achieve this, FOCUS transforms Wi-Fi off-the-shelf Consumer-Premise Equipment (CPE) hardware into a fog node (fogfication) and materializes the network-cloud slicing concept by orchestrating both cloud- and network-slicing resources to operate in a collaborative way. In addition, FOCUS controls them in a carrier-grade manner at runtime by leveraging a fully softwarized system, which makes it unique. Progressing from the study carried out in [6], FOCUS was evaluated through the FON system (the Wi-Fi sharing deployment that is most widely used) in a real testbed, to obtain an accurate perspective. The outcome confirms that FOCUS can provide greater benefits than FON, by defining the end-to-end network-cloud slicing concept and highlighting high-level isolation, as well as offering customization in the form of different VNFs and application services use cases, all controlled in a carrier-grade manner at runtime through softwarized substrates.

The remainder of the paper is structured as follows. Section 2 summarizes the background of Wi-Fi sharing technology, together with related studies. This is followed by Section 3, which describes our FOCUS framework. A proof-of-concept is established and evaluated in Section 4, along with a discussion of the results. Finally, the paper is concluded in Section 6, together with suggestions for future studies.

2. Background and Related Work

According to the ITU-R Report M.2320 [8], an Ultradense Network (UDN) is a current technology that has advanced capabilities to meet the requirements of the huge volume of data traffic and high data rates of the future 5G mobile communications. Improved coverage and cellular network capacity are a critical requirement raised by 5G UDN. According to [9], UDN scenarios involve environments such as offices, apartments, open-air gatherings, stadiums, subways, and railway stations. Given this range, these environments raise the following challenges in terms of density: (*i*) very high user density, with a 25% increase in the number of people per square meter; (*ii*) very high traffic density, with a forecast aggregated bitrate close to 10 Mbps/square meter; and (*iii*) very high AP density, where the number of small cells is expected to increase, to ensure enough throughput (e.g., APs in an office can be as little as 10m apart).

In addition to low power small cells (such as femtocell and picocell), which are overlaid on top of the macrocells to exploit spatial reuse of the spectrum, wireless data traffic can be offloaded to dense small cells belonging to indoor wireless systems, such as Wi-Fi, in which more than 80% of the data traffic effectively occurs. We claim that Wi-Fi sharing has emerged as an ideal solution for improving both the coverage and broadband data rate capacity of 5G UDN. In general terms, Wi-Fi sharing arises when users agree to share a part of their Wi-Fi broadband spectrum to have access to external user devices.

FON is a pioneer in deploying on-premise Wi-Fi-sharing technology with the aim of providing global Wi-Fi connectivity to its user community. Currently, FON implements wireless social networking by aggregating both residential and public Wi-Fi footprints and making an interconnection between Wi-Fi networks for seamless access to over 21 million hotspots worldwide. Fonero (identified here as FON-enabled CPE) corresponds to a node that runs a particular Operating System [10] that features capabilities allowing it to seamlessly connect FON-registered users into nearby FON hotspots, for ubiquitous Wi-Fi connectivity. Fonero materializes the Wi-Fi-slicing concept, whereby virtualization is implemented to accommodate two virtual networks within the common Wi-Fi spectrum for shared connectivity.

Fonero creates two Virtual Access Points (VAPs) over a single physical Wi-Fi AP and provides two SSIDs, for each Wi-Fi-shared slice: (*i*) a "public" one, devoted to FON-registered users, and (*ii*) a "private" one, for exclusive devices attached to the Wi-Fi owner's network. Apart from its coverage and potential increase in broadband data-rate capacity, our analysis of the Fonero SO highlighted the following critical issues in the domain of 5G UDN:

(i) The wireless traffic coming from the different Wi-Fi-shared slices has to be indiscriminately subjected to the same ecosystem services network (i.e., switching, routing, firewall, NAT, etc.) behind the underlying CPE. This hampers Fonero's ability to meet the different and rigorous requirements imposed by futuristic 5G applications and use cases (i.e., Enhanced Mobile Broadband (eMBB), Ultrareliable and Low Latency Communications (URLLC), and Massive Machine Type Communications (mMTC))

(ii) The Fonero Wi-Fi slicing solution provides local-level CPE QoS facilities to offer a degree of service differentiation up to the network-layer (queuing, bandwidth reservation, packet classification, and the like). However, the correct configuration of this service provisioning requires considerable networking expertise (which most owners lack) and is thus an impracticable assumption

(iii) Finally, but no less important, all the wireless traffic from the "public" and "private" Wi-Fi-shared slices only benefits from isolation at the wireless medium. All the wireless traffic will be fully transmitted through the backhaul link towards the carrier cloud, along with all concurrent traffic coming from other Fonero nodes. At both the network operator backhaul and cloud tiers, Fonero is thus capable of providing a minimum level of isolation (at the Wi-Fi medium) and lacks any service customization for a differentiated treatment of traffic handling.

FONTech (https://fontech.com) is a service in the FON catalog especially designed for operators and businesses willing to offer a FON Wi-Fi sharing service to their subscribers. FONTech adds a set of complementary software-based tools within a fully integrated architecture to carry out carrier-grade management and monitoring of Wi-Fi-shared FON hotspots for Wi-Fi seamless connectivity. Although it allows both FON client operators and businesses to deliver carrier-grade Wi-Fi-shared services (with the promise of secure, scalable, and flexible capabilities), FONTech lacks the ability to provide provisioning at both high-level isolation and customization, since it inherits the Wi-Fi medium slicing scheme from the traditional FON service.

The limitations of FON-like Wi-Fi sharing deployments in the domain of 5G UDN have driven our team to carry out further investigations. The ETSI standardization group [11] defines a set of use cases that illustrate how NFV is applicable in essential areas. Among them, [12] stands out because of its virtualization solutions that are tailored to CPEs in the network domain of both homes and small businesses. By adopting this approach, the classical CPE functions (e.g., DNS, Firewall, NAT, DHCP, mobile client authentication, and set-top boxes) [13] that are usually implemented in home CPEs are executed in the cloud data center premises instead. As a result, the CPEs (that generally encompass off-the-shelf hardware appliances) play a simpler role, by running lightweight function sets (usually L1 and L2 level services). The virtual CPE architecture has the inherent advantage of reducing the complexity of local network management [14], such as frequent *in loco* technical visits, since management services (monitoring, updating, maintenance, etc.) can be carried out remotely by the ISP operators. On the one hand, the decoupling between network functions and the local hardware platform provides the prospect of greater agility and flexibility, as well as integrating new network services into the CPE (along with a reduction in CAPEX and OPEX). On the other hand, the virtualization of services and their execution at centralized data centers may cause scalability problems, as well as adding to the complexity of management in the carrier domain.

Our research has involved integrating Wi-Fi sharing technology into the 5G network to offer enhanced UDN capabilities (in terms of a higher coverage and broadband wireless data-rate). While satisfying the heterogeneous demands of future scenarios (e.g., IoT), it seeks to extend current CPE capacities through the following key factors: (*i*) supporting dynamic multitenancy and, hence, multiservice evolutionary concepts, allowing the CPE's resources to be shared in regards to the demands of particular tenantsparticular tentant's demands (including its data, configuration, management, functionality, etc.); (*ii*) decoupling between the CPE data and control planes. The purpose of this is to simplify multiservice customized provisioning, and provide a manageable device through a network programmable approach at runtime; and (*iii*) hosting virtualized versions of both application services and network functions. These can be assumed to be key technology enablers that can allow both (i) and (ii) to provide a consistent service. In the following, we discuss a set of works that have an important bearing on the fields of research underlying this paper, and attention is drawn to their main benefits and drawbacks.

The instantiation of network services is performed on the operators side by employing cloud data center infrastructures to offload the CPE functions. However, by adopting this approach, some strategies [15] encounter a set of constraints,

due to their application being restricted to the fulfillment of a fixed set of L1 and L2 layer functions (e.g., packet forwarding for remote virtual functions). In other cases, the resulting platforms [16–18] support neither multitenancy nor traffic isolation capabilities.

Instead of using virtual structures such as VMs or containers, other systems [19] make use of software sets for implementing network functions that are generally available in Internet access CPEs (such as classical CPE functions). However, the introduction of an additional software layer raises a number of issues that are not addressed by the proposed solutions: (*i*) the lack of hardware resources isolation between functions (e.g., CPU isolation); (*ii*) the lack of security mechanisms (i.e. functions supported by the OS can be installed remotely, and nothing prevents them from behaving differently from what was expected); and (*iii*) the lack of flexibility (strict dependence on host OS configurations to execute the functions).

Centralized management of WLANs in complex environments (such as businesses and university campuses) poses considerable challenges (e.g.,the multiplicity and heterogeneity of the equipment, which involve a great configuration task on the part of the administrators) [20]. In addition, the configuration of network services and equipment (not just in wireless environments) incurs the highest costs for the service provider networks [21]. These challenges have led to new approaches being adopted, such as the CPE WAN Management Protocol (CWMP) [22], which is an open specification for WLAN remote management. Known as TR-069 [23], the protocol is a well-established standard and is also focused on tasks such as autoconfiguration, remote manual configuration, updating firmware images, obtaining performance information, and diagnosing CPE problems. Although they play an essential role in the management of CPEs, several solutions based on these protocols [24–27] show a limited set of functionalities (mainly regarding scalability and flexibility), which prevents the implementation of more complex tasks such as the management of virtual APs [28] or slicing WLAN resources [29].

In the same context, the Wireless Access Point Control and Provisioning Protocol (CAPWAP) [30] provides a centralized controller architecture for the management of APs in WLAN scenarios. This architecture standardizes communications between APs and a WLAN controller, by providing a uniform interface for performing control and management operations such as initial procedures for controller discovery, updating firmware images, and monitoring and the configuration of the AP behavior. However, schemes that are only based on this approach (e.g., [31–33]) raise serious concerns about flexibility, mainly owing to the inability of CAPWAP to provide network programming capabilities [34].

The lack of a fine granularity control to coordinate data transmissions both in the wireless and in the backhaul links, and the need for flexibility in the planning of WLANs, has led to the development of SDN-based solutions for this domain. The schemes outlined in [35–44] reflect the importance of expanding the netprogrammability capacities until the last hop of the WLANs.

As highlighted in [45], cloud computing makes it feasible the realization of IoT in an industrial level by providing advantages such as computing elasticity and cost-efficient storage for a massive amount of monitoring data originated from sensors/things. Nonetheless, the inadequacy of the cloud to meet specific IoT requirements such as high capacity and availability of links that connect them to the cloud, lower latency, and support for localized services motivated the conception of the fog computing, an architecture that extends the cloud to the edge of the network by placing computing, storage, and network resources closer to devices and sensors. Fog computing emerges as a common platform for execution of applications from a range of domains (smart metering, wind farms, urban transportation, eHealth, etc.). As so, it requires orchestration solutions that rely on "softwarization" of network and computing resources [46] to support the coexistence of applications from different tenants [47] and satisfy nonfunctional requirements such as QoS, reliability, and security [48]. Nonetheless, fog orchestrators solutions are yet at its infancy when compared to the orchestration of pure cloud environment. In spite of initial efforts such as the ones concentrating on conceptual generic frameworks [48–51] or targeted on specific issues [52–54], a comprehensive orchestration and management solution that tackle key issues such as scalability, reliability and robustness [55], location awareness, and the heterogeneity and distribution of fog nodes is still missing.

Our previous study argued that WLAN slIcing as a SErvice (WISE) [6] addresses the challenges of effectively coping with the increase of massive mobile data demands in UDN 5G use cases, by deploying the current Wi-Fi-sharing technology. WISE expands the computational capabilities of Wi-Fi sharing CPEs through the application of Fog computing principles, which allow the implementation of slice-based capabilities. The main principle behind this approach is to efficiently connect Wi-Fi slices into fog slices that feature differentiated services on top of the same infrastructure, through customized, isolated, and independent digital building blocks. Finally, it is assumed that fog slices can virtualize application services in addition to network functions, and are thus capable of offering ultra-low latency rates by directly linking to data producer services and applications. The feasibility of this strategy is assessed via experiments in a real testbed, that allow an insight into its proof-of-concept. Although it improves current Wi-Fi sharing deployments to cope with 5G UDN, WISE lacks any carrier-grade control, which means that all the (re)configurations must be set at the fog node premise manually. Moreover, the way from the fog node to the operators cloud, lacks any kind of traffic isolation.

Our related work studies clearly show that none of the illustrated technologies is capable of providing the required support to enhance 5G UDN services. In general, the currently available Wi-Fi sharing solutions lack the following: multitenancy or multiservice dynamic support; carrier-grade netprogrammability at run-time; and end-to-end dedicated/customized resources, as well as independent/isolated service delivery. In light of this, the next section describes our proposed solution, which raises as unique and advances the state-of-the-art through a new Wi-Fi sharing

service architecture that enables an efficient framework to be established for 5G UDN.

3. Description of the FOCUS Approach

The FOCUS approach aims to create and manage Network-Cloud slices inside the service provider, providing an end-to-end view all the way from a Wi-Fi slice to the carrier cloud system. A Network-Cloud slice is understood as a partition of both computing and network resources, featuring high-isolation by definition, upon which network functions and application services can be executed in the form of virtualized structures. At one extreme of the Network-Cloud slice, Wi-Fi sharing is applied to leverages the densely available Wi-Fi networks at urban centers so as to promote ubiquitous wireless networking connectivity. FOCUS manages this pool of geographically distributed WLAN CPEs to dynamically offer services closer to end nodes (e.g., by selecting and configuring a subset of CPEs to provide network connectivity wirelessly and host local services to a sort of mobile nodes). On the basis that a sort of applications or services requires particular resources to run in an effective way, FOCUS also manages carrier's cloud computing resources and realizes the integration between the WLAN CPEs and the carrier's cloud.

The slicing technique, along with the collaboration of both "softwarization" and "cloudification" technologies, has emerged as a viable ecosystem that can suit the functional requirements of 5G UDN future scenarios. In general terms, Slicing deploys a view of segmented physical resources in the form of independent virtual elements, each capable of allowing virtualized instances of network functions and application services to run on the basis of preallocated resources. In enabling it to cope with the current 5G UDN challenges, while suiting its functional requirements, the FOCUS approach has been designed (though the integrated native orchestration and management of both cloud- and network-Slicing techniques), to offer multiple independent virtual service structures that feature high-level isolation all the way between the Wi-Fi slices and carrier cloud slices. The Network-Cloud slicing approach that FOCUS deploys aims to provide an end-to-end view, by deploying a set of physical components (i.e., cloud servers, network nodes, and other suitable devices) that are capable of supporting both application services and network functions to efficiently run in virtual instances (i.e., virtual machines or containers).

In view of the need to exploit Wi-Fi sharing systems in both Home Office and Small Office Home Office (SOHO) environments, solutions must be tailored with a low-cost off-the-shelf WLAN CPE platform (a.k.a. legacy CPE), which offers resource-constrained hardware capabilities. Legacy CPEs traditionally support Internet access network functions (e.g., DNS, Firewall, NAT, DHCP, Wi-Fi AP, etc.), and, depending on the OS distribution, advanced tools can be available (e.g., QoS control). In this kind of resource-restricted device platform, virtualization is at risk of negatively impacting the system due to its resource-consumption approach. To overcome this danger, the Virtual CPE (vCPE) concept [13] can be applied, whereby data are offloaded to the cloud so that it can be applied to corresponding network functions. On the one hand, vCPE is promising because of its ability (a) to reduce CAPEX/OPEX, (b) to break free from specific hardware dependencies and upgrade time constraints, and other factors. On the other hand, the service agility is questionable since there is a need to face delays in seeking to reach the cloud by applying the target VNF, which can become essential (depending on the network conditions) to suit the rigorous requirements of Ultrareliable and Low Latency (URLL) 5G usecases.

Focus is able to overcome the performance penalties discussed above by expanding the computational capabilities of the WLAN-sharing CPEs. This is carried out through applying Fog computing technology to support cloud-slice definitions at the edge. Thus, FOCUS allows ultralow latency rates to be offered by making a direct link with IoT devices. However, it is beyond the scope of this paper to provide technical guidelines on how computational resources of low-cost off-the-shelf WLAN CPEs can be expanded to become a fog node. There are different ways available for achieving this target deployment (Section 5 describes FOCUS prototyping and outlines our strategy).

FOCUS adopts a fully softwarized design to carry out Network-Cloud life-cycle operations (creation, configuration, activation, elasticity, and tear-down), both dynamically and flexibly, and provide a remotely manageable platform driven by carrier-grade systems. During the Network-Cloud slice creation, FOCUS is driven by requirements to allocate resources (i.e., connectivity, computing, and storage), and all other control operations (elasticity and tear-down) are carried out at run-time. In the Network-slicing life-cycle, FOCUS relies on SDN northbound and southbound APIs to dynamically enforce the configuration of network resources (i.e., classification, interface chaining, bandwidth reservation, etc.). In what concerns the Cloud-slicing life-cycle, FOCUS relies on the assistance of the Virtualized Infrastructure Manager (VIM) to dynamically request the enforcement of cloud resources (i.e., computing, memory, and mass storage) for target systems. It should be noted that cloud resources refer to the computing resources for allocation at both the fog and cloud parts of the ecosystem. The fully softwarized Network-Cloud approach of FOCUS anticipates that the Wi-Fi sharing experience will be of great value in improving 5G UDN technology capacities at unprecedented levels. Figure 1 provides (from a high-level view) the basic differences between the Wi-Fi sharing system offered by both classical and FOCUS Network-Cloud approaches.

According to Figure 1, FOCUS allows the CPE service architecture to be significantly simplified by only keeping the essential networking services and softwarization and virtualization substrates. All the other networking functions are offloaded to the carrier cloud infrastructure, which makes use of high computational capabilities to achieve an outstanding performance. FOCUS demonstrates its benefits through the orchestration of specific Network-Cloud resources (e.g., CPU cycles, memory, storage, and network) at the fog networking (Wi-Fi and backhaul) and cloud infrastructures. The FOCUS platform allows the execution of multiple logical systems, along with respective services, within a common end-to-end shared infrastructure that embodies the wireless part
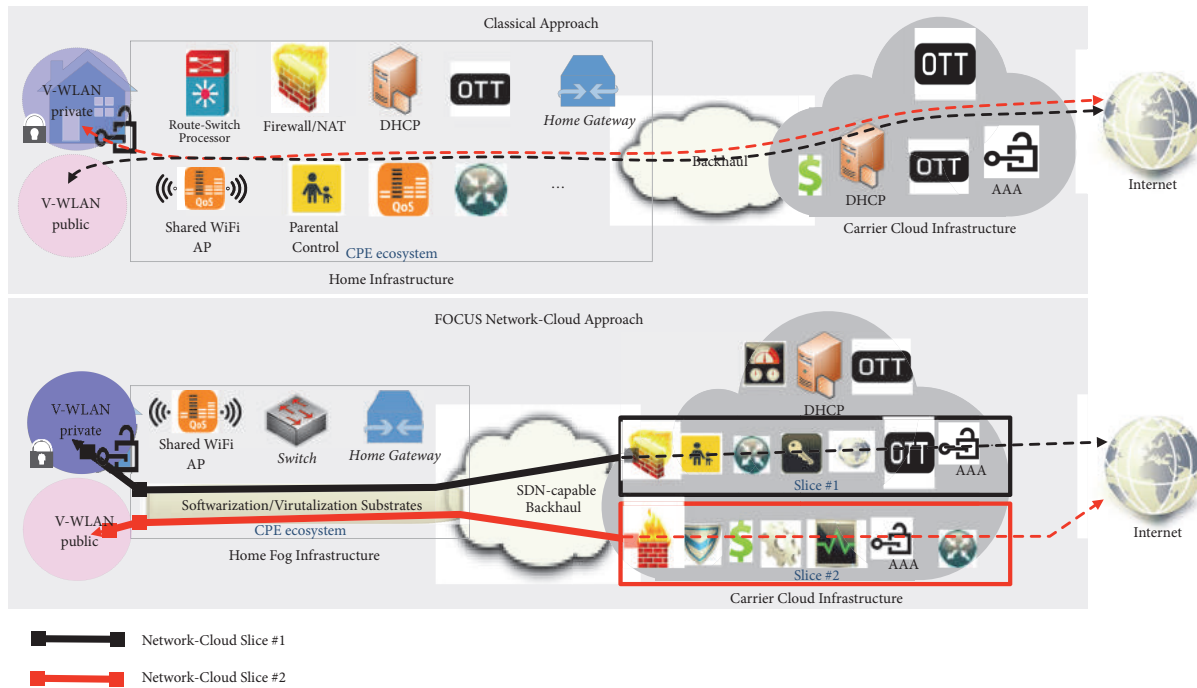
FIGURE 1: Wi-Fi sharing classical and FOCUS Network-Cloud approaches.

all the way to the carrier data-center infrastructure. By definition, resources allocated to a slice are isolated and used independently, since the FOCUS is responsible for fulfilling the appropriate guarantees.

3.1. Architectural Overview of the FOCUS Approach. After providing an overview of the main capabilities featured in FOCUS, this section introduces its functional architecture. FOCUS is designed in the form of a modular architecture that comprises a set of functional blocks and internal/external interfaces. Figure 2 depicts the FOCUS functional architecture, together with the end-to-end Network-Cloud Slice structure.

As shown in Figure 2, the functional architectural design of FOCUS consists of two different tiers operating at the same functional level, each featuring different functional blocks and internal subsystems that are interoperable via well-defined interfaces so that they can play different roles according to their location. The Wi-Fi Sharing tier depends on the participation of the Fog-Slicing Control Plane, in which corresponding functional blocks are interoperable to enforce and monitor the fog node resources. The Carrier Cloud tier embedding functional blocks play, in turn, a broader role in the FOCUS system. This involves cooperating in an attempt to carry out all the complex Network-Cloud slicing management and control decisions at both the Carrier Cloud and Wi-Fi Sharing tiers. Figure 3 shows a high-level view of the end-to-end Network-Cloud Slice that FOCUS structures.

As outlined in Figure 3, FOCUS instantiates at the Fog node part a Fog Slice #1 (FS#1) that includes containerized

VNFs (VNF#1, VNF#2, and VNF#3 in this case), each configured with individual resources (CPU and RAM). Moreover, FOCUS creates three Virtual Network Interfaces (VNetIf), and each containerized VNF is accordingly attached to a corresponding VNetIf. A fourth VNetIf will be used to transfer the packets belonging to the FS#1. Finally, FOCUS sets the Virtual Switching tool with the forwarding rule that must be applied to all the packets linked to the NS#1, and in this case, follow the following sequence: V-WLANprivate@CPE#1 $\longrightarrow$ VNF#1 $\longrightarrow$ VNF#2 $\longrightarrow$ VNF#3 $\longrightarrow$ NS#1. At the Carrier cloud infrastructure, FOCUS instantiates: a cloud slice (CS#1) assigned to the indicated resources (CPU, RAM and disks); two containers that run VNF#4 and the application service (APP); and a virtual network interface attached to each container. Assuming that there is a peer-to-peer link that physically connects the Fog node and the carrier cloud systems, FOCUS sets corresponding forwarding rules at the could part VSwitch, and finally the end-to-end Network-Cloud is activated.

3.2. Network-Cloud Control Plane. The Cloud-Slice Control Plane is an essential block of the FOCUS architecture that can be found in the network operator cloud infrastructure. The Cloud-Slice Control Plane consists of a set of functional blocks that are interoperable and can orchestrate the elementary network-cloud slice resources in both the Carrier Cloud and Wi-Fi Sharing tiers. The Cloud-Slice Control Plane architecture is described in the next sections.

3.2.1. Slice Descriptor. The Slice Descriptor plays a crucial role in the Network-Cloud Control Plane as the receiving
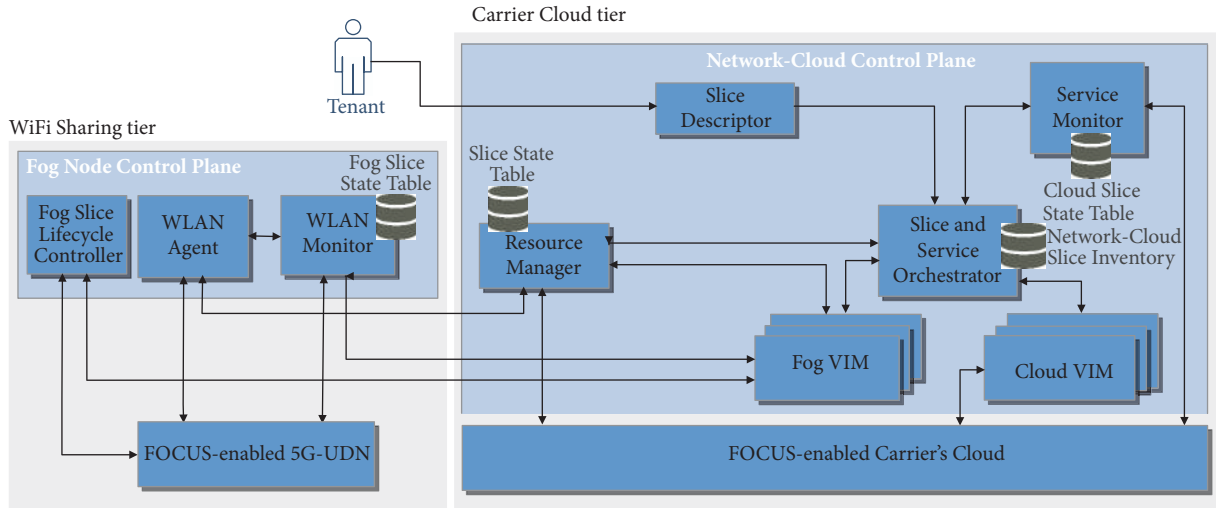
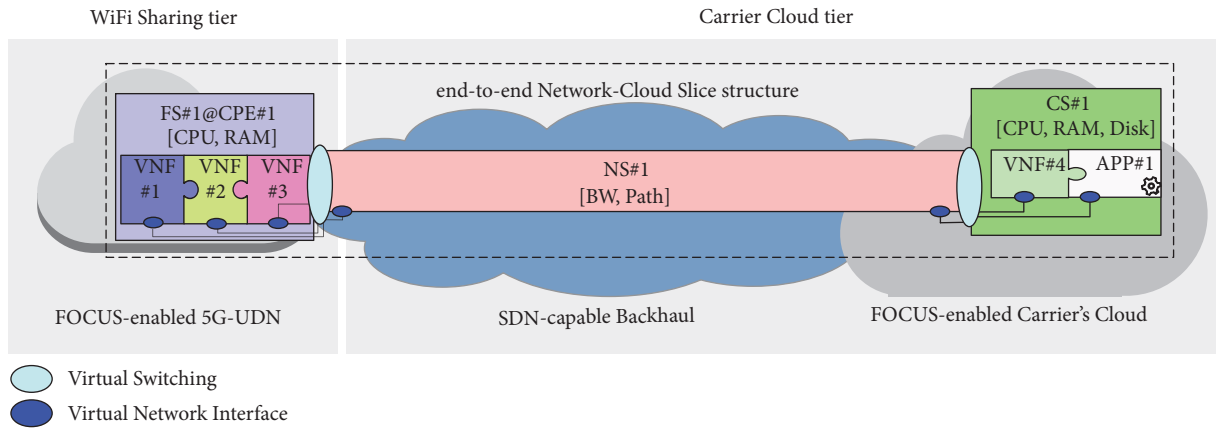FIGURE 2: Functional architecture of the FOCUS approach.

FIGURE 3: FOCUS End-to-end network cloud slice structure.

point of the tenant's slice creation requests, via a well-defined external interface. A tenants slice creation request includes a description of a slice framework that must comply with the TOSCA template language [56] and provides both mandatory and optional information. The information about Network-Slice resources, infrastructural requirements, and geographical location is mandatory. In the Network-Cloud slice resources, there are both computational (CPU, memory, and storage) and networking (bandwidth) minimum capacities. The infrastructure requirements record both VM flavors and network QoS tolerance (delay, loss, jitter, etc.). The geographical location indicates the fog node positioning that will root the target Network-Cloud slice. With regard to the WLAN coverage, the tenant can specify either a precise geospatial location or a geographical area. In the case of the former, FOCUS attempts to instantiate the fog slice that matches the given geospatial position. In the latter, FOCUS will select a group of fog nodes located within the limits of the geographical area. Optional information is also allowed in the slice description. For instance, the tenant can make a decision to specify a general component scope (e.g., a web

server) or a particular component implementation (e.g., the Nginx HTTP Server [57]). On the basis of the incoming slice description, the Slice Descriptor functional block forms a FOCUS-compliant Slice Specification framework.

3.2.2. Resource Manager. By definition, a Network-Cloud slice requires resources to support its running services. The Resource Manager is the functional block inside the FOCUS architecture that provides knowledge about both the cloud-level and network-level resource availability in the entire FOCUS-enabled ecosystem. The cloud-level knowledge provides the Resource Manager with two operational modes, namely, passive and reactive. In the former, the Resource Manager passively waits for the occurrence of incoming Fog node announcements and Fog slice feedback. Whereas the Fog node announcements are processed at the system boot-strap (carrying both current resource status and geospatial positioning input), the Fog slice feedback is dynamical events that document QoS conditions. In the latter, the Resource Manager reacts by responding to an explicit request from the Slice and Service Orchestrator functional block (such as the

list of fog nodes participating in a given citys geographical area) as well as resources available in the carrier cloud infrastructure. In both operational modes, the Resource Manager's responses are supported by the Slice State Table local subsystem, which registers stateful knowledge of all of the end-to-end Network-Cloud Slices activated in the FOCUS ecosystem. In the domain of network-type resources, the Resource Manager collaborates with an existing network control-enabling tool (e.g., SDN Controller), which is capable of providing the network map of the entire carrier system infrastructure, including nodes, links, and current capacities.

3.2.3. Slice and Service Orchestrator. The Slice and Service Orchestrator functional block plays a key role in the FOCUS architecture by finding a way to combine the fog, network, and cloud slice parts into a single end-to-end Network-Cloud slice. In seeking to achieve this, the Slice and Service Orchestrator functional block processes the slice specification supplied by the Slice Descriptor to support the following: (*i*) selecting the required artifacts that the Network-Cloud slice will serve; (*ii*) identifying the tiers to deploy each of the Network-Cloud slice parts; (*iii*) instantiating suitable VIMs to handle each of the Network-Cloud slice parts; and (*iv*) invoking the VIMs to enforce the Network-Cloud slice resources, which will finally connect the Network-Cloud slice parts. The artifacts refer to the software component(s) that will run in the Network-Cloud slice, in the form of VNF or a service application. The Network-Cloud Slice Inventory is a subsystem of the Slice and Service Orchestrator functional block that keeps all the available artifacts.

3.2.4. Cloud-VIM and Fog-VIM. In the FOCUS architecture, the main goal of both the Cloud-VIM and the Fog-VIM functional blocks is to set up an end-to-end Network-Cloud slice, together with the services that can be offered. In light of this, they support the necessary capacities for dynamically controlling and managing the computing, storage, and network resources within the cloud and fog physical infrastructures respectively. Technologically speaking, both Cloud-VIM and Fog-VIM refer to particular tools that are capable of converting generic calls, specific commands, or methods in accordance with a particular VIM implementation. Examples of this can be OpenStack (http://www.openstack.org) and Kubernetes (http://kubernetes.io) to handle cloud-slice resources, and Opendaylight (http://www.opendaylight.org) SDN Controller to deal with Network-Slice resources. In the FOCUS approach, tenants are allowed to indicate a particular VIM implementation as part of the Slice Definition. On the basis of this clear indication, the Slice and Service Orchestrator functional block derives the VIM that will be supported during the end-to-end Network-Cloud Slice provisioning task. The stateful knowledge of slices are kept in the Slice State Table (subsystem) of the Resource Manager's functional block.

3.2.5. Service Monitor. End-to-end Network-Cloud slice creation is supported by the enforcement of an elementary resources setup, which in turn is driven by a basic knowledge of the underlying systems. The Service Monitor functional block is designed to collect up-to-date Key Performance Indicators (KPIs) of Cloud Slice, and book them in the Slice State Table. Monitoring data includes information about the cloud slice topology, resources of particular slice parts, and the KPIs of the running service applications. The Service Monitor functional block provides the necessary knowledge to support the management operations of the slice life-cycle.

3.3. Fog Node Control Plane. Highly dense Wi-Fi-sharing enables fog nodes to cooperate in the FOCUS ecosystem with the aim of providing wireless broadband connectivity with enhanced capacities for 5G UDN. The Fog Node Control Plane resides at each of the the Fog nodes (at the Wi-Fi sharing tier of the FOCUS architecture) and is devoted to responding to all the fog node level operations. The subsystems colocated in the Fog Node Control Plane architecture are examined in the next subsections.

3.3.1. Fog Slice Life-Cycle Controller. The Fog Slice Life-cycle Controller functional block takes part in the Fog Node Control Plane to respond to Network-cloud slice life-cycle management request operations (i.e., creation, update, and tear-down) from the Fog VIM functional block. In seeking to provide resources at a given Fog node (as requested by the Fog VIM) the Fog Slice Life-cycle Controller operates together with the colocated third-party enabling tools, as follows. With regard to container-level installation, execution, and isolation tasks (as a part of the fundamental basis for a virtual structure provisioning) a container manager enabling tool (e.g., Docker (http://www.docker.com)) is needed to run a service application or network function. A virtual switch enabling tool (e.g., Open vSwitch (http://www.openvswitch.org)) is provided to dynamically set up forwarding rules so that packets coming from a Wi-Fi slice can be delivered to an offering virtual service in a netprogrammable way. Furthermore, the virtual switch enabling tool is also a significant means of instantiating virtual network interfaces so as to allow the provisioning of a virtual service to receive/transmit packets. The virtual switch enabling tool also provides basic support to provision network slices that connect slice instances at both the Fog node and the carrier cloud.

3.3.2. WLAN Agent. As was pointed out in Section 3.2.2, the Fog nodes implement a resource announcement mechanism to support the Resource Manager's functional block in the task of making Network-Cloud level resource knowledge available. At the Fog node bootstrap system, the WLAN Agent functional block announces the current capabilities and geospatial location that are within the scope of the local Fog node. Furthermore, the WLAN Agent functional block works together with the WLAN Monitor functional block to gather Fog slice monitoring KPIs as feedback for the Fog-VIM about Fog slice QoS (e.g., compute/network load, utilization rates, temperature, among other factors).

3.3.3. WLAN Monitor. The status of the physical resources that the Fog slices are offering at a given time is vital for efficient slice management. The WLAN Monitor represents the functional block inside the Fog Node Control Plane and is
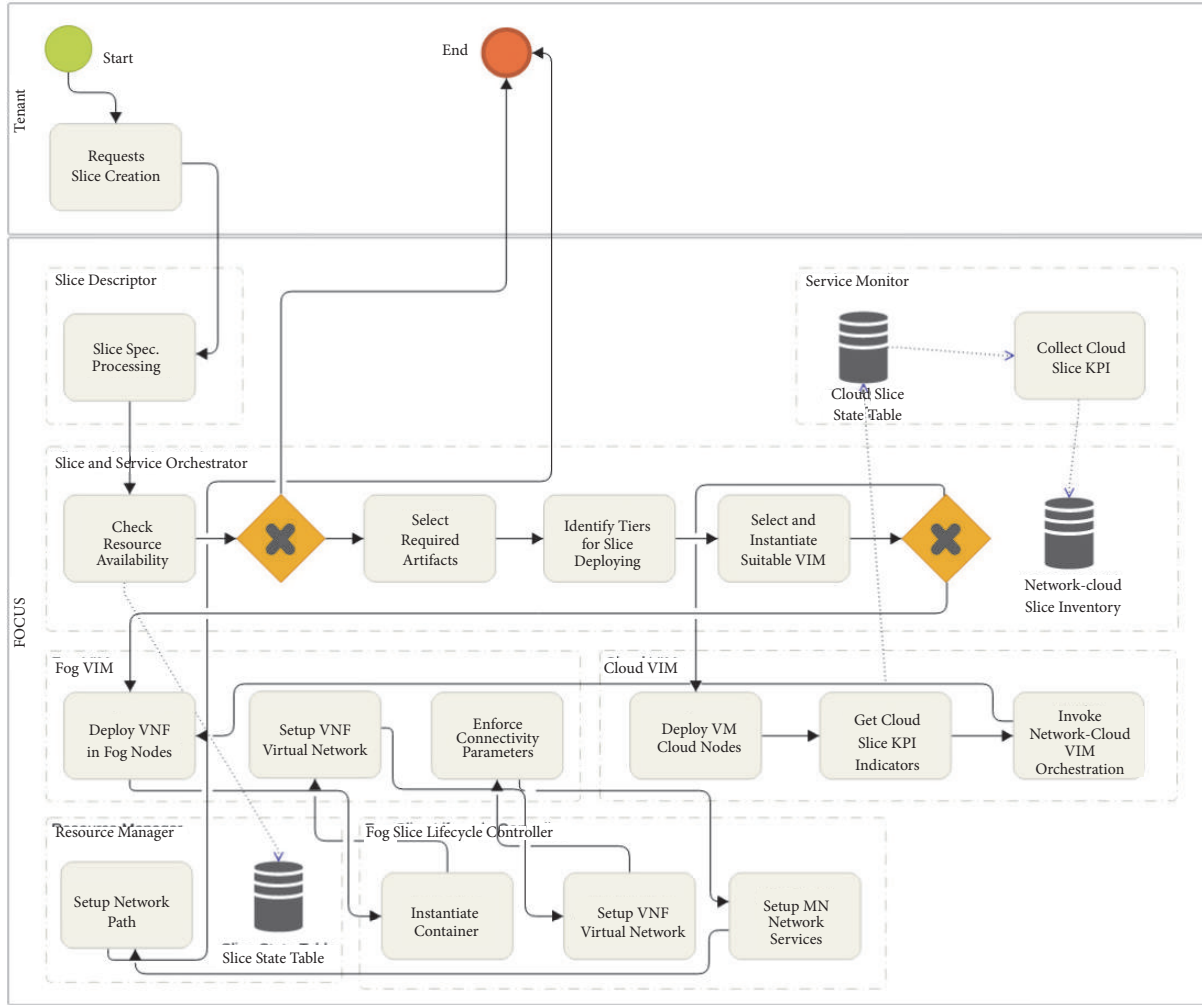
Figure 4: Slice Creation Workflow.

responsible for gathering the monitoring data at the Fog node physical infrastructure, which ranges from per slice computational to networking KPIs. The Fog slice knowledge that the WLAN Monitor functional block maintains in the Fog Slice State Table is shared with both the Resource Manager and the Fog VIM functional blocks of the Network-Cloud Control Plane to support different kinds of decision-making. For instance, the Resource Manager gathers the current resource status to decide whether the Fog node can admit a new slice. The Fog VIM, in turn, uses Fog Slice knowledge to drive scalability (either increasing or decreasing) decision-making, which seeks to meet workload dynamics at runtime. The WLAN Monitor functional block is also prepared to report heartbeat messages to the Fog-VIM at short intervals, as well as to support the detection of Fog node network dynamics (e.g., failure).

4. FOCUS Life-Cycle Workflow

The Network-Cloud Slice life-cycle requires operators to create, configure, activate, update, and tear down, to cope with its dynamic features. The next subsections show how the BPMN (Business Process Model and Notation) methodology can be employed to document the function flow sequence that is invoked to carry out the main Network-Cloud life-cycle operations, namely, creation and tear-down. The descriptions are not exhaustively documented, and this strategy is adopted to provide an easier understanding.

4.1. Slice Creation. Figure 4 shows the slice creation workflow in terms of the architectural features of the FOCUS framework and illustrates the interaction between the modules required to create a new slice.

The slice creation method described in this workflow starts from the creation request made, e.g., by a tenant seated at the Service Provider. On the basis of a set of information that forms the slice description (as explained in Section 3.2.1), the request of the tenant is forwarded to the Slice Descriptor, which interacts with the Slice and Service Orchestrator to determine the availability of resources. On detecting that the resources available are not enough to satisfy the slice request (e.g., through a lack of fog nodes at the desired geographical locations), the Slice and Service Orchestrator exits the slice
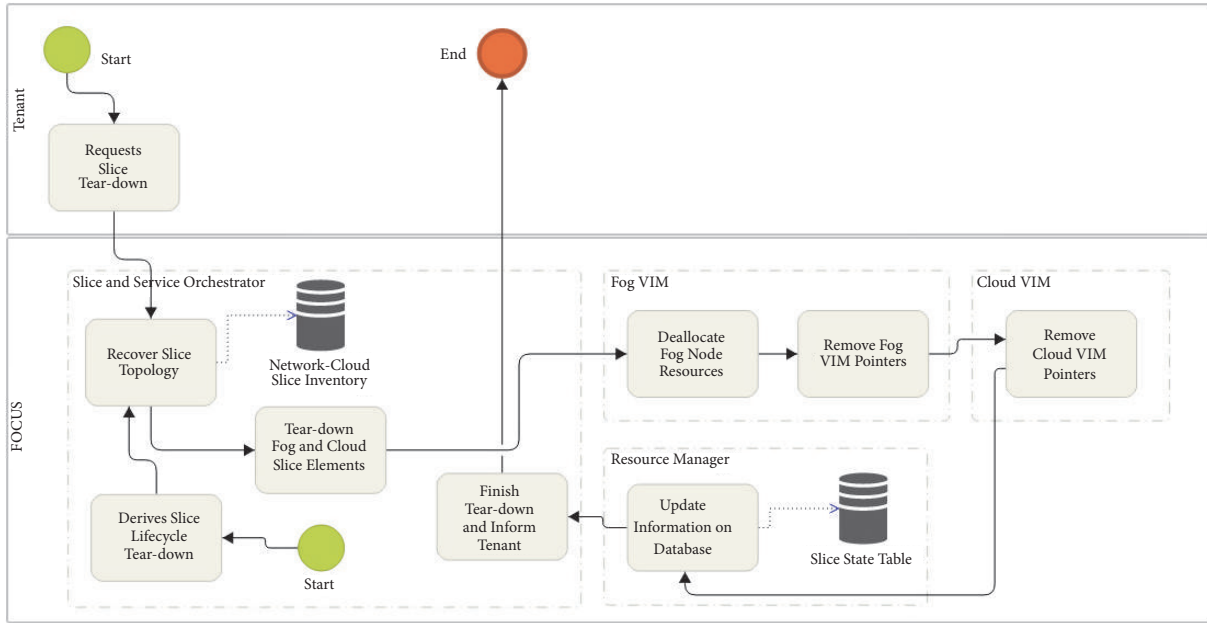
FIGURE 5: Tear-down Workflow.

instantiation and provides feedback to the tenant accordingly. If there are enough resources, the Slice and Service Orchestrator is able to select and instantiate the appropriate VIM through the Fog VIM or Cloud VIM components. The latter will invoke the proper functionalities for deploying and instantiating the fog and cloud nodes. After that, there are subsequent interactions between the Fog VIM and the Fog Slice Lifecycle Controller to provide the necessary network configurations needed to ensure proper connectivity to the requested new slice. The process ends with the Resource Manager delivering all the required network paths to the new slice.

4.2. Slice Tear-Down. Figure 5 illustrates the workflow and interactions behind the Slice tear-down operation.

The tear-down process can be started in two different ways: (*i*) from an explicit request made by the tenant and (*ii*) detecting that an end-to-end Network-Cloud slice, or part of it, is no longer necessary. In the case of the latter, the Slice and Service Orchestrator proceeds to finish its life-cycle, which refers to tear-down of all accompanying physical resources that offer services. In both cases, after being triggered, the Slice and Service Orchestrator processes the request and retrieve information about the slice topology, its elements, and linked resources. Then, the Slice and Service Orchestrator forward the tear-down request to the Fog VIM and Cloud VIM in order to deallocate all the slice parts, including the fog node resources and VIM pointers. After the Tear-down has been completed successfully, the information is updated in the Resource Manager, through the Slice State Table, and the tenant is properly notified.

5. Testing and Evaluation

This section illustrates the evaluation of a proof-of-concept concerning our architecture. The next subsections describe the testbed environment and several outcomes concerning the achieved throughput and delay, when compared with other solutions.

5.1. Testbed Deployment. A real testbed prototyping in the Ubiquitous and Pervasive Systems Lab (UPLab) at the Federal University of Rio Grande do Norte (UFRN), Brazil, is used to demonstrate the feasibility and sustainability of the FOCUS proposal, along with a comparison with the most relevant related work, notably the FON Wi-Fi sharing deployment. The FOCUS prototyping aims at confirming the basic functional principles by bringing together interoperable components that include the functional architecture displayed in Figure 2. Moreover, we carried out a system-level analysis of our prototype, in order to assess how well the FOCUS system addresses the functional requirements (see the discussions in Section 2) to cope with 5G UDN. The testbed configuration progresses from [58], which includes both Wi-Fi sharing and Carrier's cloud tiers connected through an SDN network that connects four OpenFlow-enabled Mikrotik 951G-2HnD (CPU of 600 MHz, and RAM of 128 MB) switch nodes. The testbed prototyping configuration is shown in Figure 6.

A set of enabling technologies are employed in the testbed configuration to allow using elementary resources for the dynamic provision of end-to-end Network-Cloud slices. In the Network-Cloud Control Plane, the enabling technologies colocated in the testbed setup are elicited in the following: (*i*) Open vSwitch (OVS), as the virtual switch enabling tool; (*ii*)
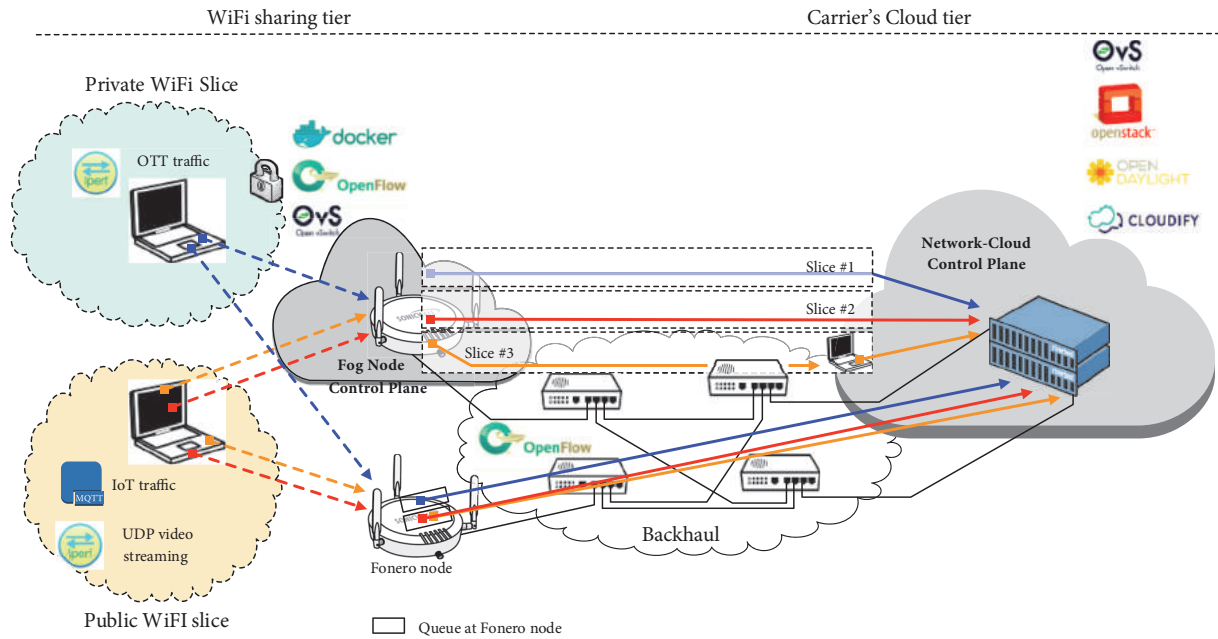
FIGURE 6: Representation of the testbed configuration for the FOCUS and FON set of experiments.

Cloudify (https://cloudify.co/), which provides service and resource orchestration support; (*iii*) OpenFlow, a southbound API to setup networking resources and packet forwarding at both Open vSwitch and SDN-enabled network nodes; (*iv*) OpenStack, for VIM implementation; and (*v*) OpenDayLight, a SDN controller that provides northbound API to trigger and control OpenFlow features. FOCUS Network-Cloud Control Plane facilities are implemented within a cloud that runs into a two clustering rack servers PowerEdge R7425 (2 AMD 32-core EPYC™ processors, 64GB DIMM DDR4 RAM, 4 HDDs of 2TB, and 4 GbEth network cards).

In the FOCUS Fog Node Control Plane side, Docker is responsible for container management, along with OVS and OpenFlow v1.3 support. An off-the-shelf Wi-Fi router TP-LINK TL-WR1043ND v3 (CPU of 720 MHz, and RAM of 64 MB), running the OpenWRT v18.069 and the Fog Node Control Plane implementation, is adopted to provision the Wi-Fi-sharing technology.

When making the FON testbed pilot comparison, a Fonero node running in the off-the-shelf Wi-Fi routers TP-LINK TL- WR1043ND v3 provides the FON-based Wi-Fi sharing system. As a means of allowing a comparative approach between the FOCUS testbed pilots, the Fonero configuration relies on two forwarding queues, one for each Wi-Fi network to handle the corresponding incoming wireless packets. The offering network capacity in the link that connects the Wi-Fi sharing tier and the carrier cloud infrastructure is set to 20 Mbps, for both tests, a value below the average broadband connectivity at South Korea, which, according to the Akamai study [59], boasts the highest Internet connection speed by averaging 29 Mbps (the global average is 7.2 Mbps).

Two laptop clients took part in the testbed pilots for traffic generation: one attached to the Private Wi-Fi Slice and the other to the Public Wi-Fi slice. The IPerf (https://iperf.fr) tool allows UDP traffic to be provided in both Private and Public Wi-Fi slices. The Private Wi-Fi slice offers UDP flows at a maximum rate of 6.8 Mbps (representing OTT multimedia content, mostly consumed by the Internet home users), and 1.4 Mbps at Constant Bit Rate (CBR) for the Public Wi-Fi slicing (thus simulating a surveillance camera video stream flow). A Mosquitto plugin for the Gatling tool (https://gatling.io) generates MQTT v3.1.1 traffic as a suitable means of representing Machine-To-Machine (M2M)/IoT connectivity (primary service offering in Public Wi-Fi slices). The traffic flow is distributed along the prototyping testbed at runtime, which lasts 70 seconds. In the case of the FOCUS testbed pilot, the incoming traffic is subject to the execution of different NFV service chaining instances that run at each Network-Cloud slice, thus denoting customization and independent service provisioning. At the Network-Cloud slice #1, Single Packet Authorization (SPA), Smart Queue Management (SQM), and Hierarchical Token Bucket (HTB) packet scheduling discipline are executed in the form of three NFV chaining instances. At the Network-Cloud slice #2, two NFV chaining instances run the Class-Based Queuing (CBQ) packet scheduling discipline and traffic shaping mechanisms are used. The Network-Cloud slice #3 runs a general-purpose software application, which is assumed to consume the UDP video-like traffic (such as in video surveillance analytics for Smart City use cases). In the case of the FON testbed prototyping pilot, the traffic coming from the Public and that from the Private Wi-Fi Slices are each served by the QoS feature that Fonero OS provisions. In this way, two queues are set with the HTB packet scheduling discipline to allow traffic isolation at the Fonero node. Table 1 summarizes the parameters of the traffic employed for the experiments.

It is worth stressing that the testbed prototyping pilots are not meant to provide strict guidelines or define any enabling technology for use. We assume that there are many

TABLE 1: Traffic Parameters for the testbed experiments pilots.

WiFi Network	Traffic type	Num. of clients	Packet size	Max. transmission rate
Private	UDP (OTT)	1	1024 bytes	6.8 Mbps
Public	TCP (IoT)	300	200 bytes	10.9 Mbps
Public	UDP (video)	1	1000 bytes	1.4 Mbps

mature technologies, of different types and technologies, that are available, and the final decision must be made by the evaluation designers.

5.2. Results and Discussion. In view of the functional requirements for coping with 5G UDNs (see Section 2), the prototyping pilot tests are carried out with a view to answering the following questions: (*i*) which level of service differentiation FOCUS and Fonero can offer for each of the types of traffic coming from the Wi-Fi slices?; (*ii*) how should FOCUS and Fonero respond to dynamic network changes during the test runtime?; and (*iii*) which isolation level FOCUS and Fonero are capable for providing end-to-end Network-Cloud slices? Our analysis exploits the throughput and delay impact on the offering traffic, as the main KPIs that can address these questions.

In order to represent dynamic access behavior in the Wi-Fi sharing slices, we introduce different network resource demanding procedures as in the following. At the beginning of the tests, UDP traffic clients (Public and Private Wi-Fi slices) start to run by scaling flows up to their corresponding maximum transmission rates. In the case of the TCP traffic (Public Wi-Fi slice), all the 300 clients progressively start transmitting during the first seconds of the test runtime (to avoid packets being sent simultaneously by all the clients). After 16 seconds of the running test, the traffic rate from the Private Wi-Fi slice reduces to around 1 Mbps, to represent a real-life situation (e.g., when users leave home to go to work, school, etc.). After 45 seconds of the test running time, the Private Wi-Fi slice's traffic returns to 6.8 Mbps, to denote that the users have returned to their homes, and hence to their home use pattern of behavior.

In response to these dynamics, the FOCUS scheme behaves in our testbed prototyping pilot as follows. The Network-Cloud slices (#1, #2, and #3) are established before the test runtime in both the Fog Node and Network-Cloud Control Plane systems, which is in compliance with the Wi-Fi-sharing scenario. The Network-Cloud slices (#1, #2, and #3) start with a bandwidth capacity of 7 Mbps, 10 Mbps, and 1.5 Mbps, respectively. The reduction of traffic patterns in Network-Cloud slice #1 triggers FOCUS elasticity operations that, at 23 seconds of testbed runtime, allows FOCUS to reduce the Network-Cloud slice #1 bandwidth reservation from 7 Mbps to 2 Mbps and increase the Network-Cloud slice #2 (associated with Public Wi-Fi Slice) bandwidth from 10 Mbps to 15 Mbps. The purpose of this is to adapt to the new traffic demands and provide an enhanced QoS. At 52 seconds of the testbed runtime, FOCUS returns Network-Cloud slices (#1 and #2) to their initial bandwidth reservation patterns (i.e., 7 Mbps and 10 Mbps, respectively), driven by detecting increased traffic offering behavior. The dashed lines


FIGURE 7: FOCUS testbed prototyping pilot results: averaging per Network-Cloud slice throughput impact in the entire data path featuring end-to-end high-level isolation and flexibility to dynamically adapt to network changes at the runtime.

in Figure 7 illustrate those dynamic events in bandwidth reservations.

In the FON testbed prototyping pilot test, each Wi-Fi slice traffic is directed to a particular HTB queue, one for Private-alike Wi-Fi slice traffic and another for Public-alike Wi-Fi slice traffic. The Private-alike Wi-Fi slice queue is set with static bandwidth reservations of 7 Mbps, while the Public-alike queue leverages 11.5 Mbps (the sum of bandwidth reservations set to both slices #2 and #3 of the FOCUS experiment). The static reservation bandwidth is set before the tests and is maintained during the entire test runtime, and thus complies with the QoS feature supported by Fonero that requires manual (re)configuration.

The throughput impact outcomes on a per Wi-Fi slice basis over the test runtime is depicted in Figures 7 and 8, that include the average throughput at all the links entailing the data path. The main goal of the analysis of throughput is to study how both FOCUS and Fonero prototyping pilots behave in response to the network dynamics introduced at runtime. Moreover, an analysis of per Wi-Fi slice throughput along the end-to-end data path is an efficient way to study the isolation level that both FOCUS and Fonero can provide for the testbed prototyping pilots. As dynamic events in the network are introduced in the testbed at different time intervals, a different impact can be expected from both FOCUS and Fonero, with regard to traffic behavior along the

FIGURE 8: FON testbed prototyping pilot results: averaging per Wi-Fi slice throughput impact in the entire data path featuring static bandwidth reservation only at the Fonero node.

FIGURE 9: The impact of bandwidth changes at the MQTT mean delay.

entire data path, from the Wi-Fi slice to the carrier cloud system.

The way the MQTT clients interact with the MQTT broker (located at the carrier cloud) can have an impact on the bandwidth use behavior of Network-Cloud slice #2 (the total bit rate reached by the MQTT clients is around 70% of the corresponding available bandwidth of the slice). The MQTT clients QoS Level 1 marking is an MQTT protocol service that forces the MQTT client to retain the sent message until it receives an acknowledgment from the broker (an MQTT PUBACK message). Additionally, the MQTT clients only send a new message after receiving the acknowledgment from the previous one.

At the FOCUS experiment (as depicted in Figure 7), the increase in the bandwidth reservation of the Network-Cloud slice #2 (from 10 Mbps to 15 Mbps), and corresponding decrease of Network-Cloud slice #1 (from 7 Mbps to 2 Mbps), occurs around 8 seconds after the decrease of UDP traffic at Network-Cloud slice #1 (from 6.8 Mbps to 1 Mbps). During this time interval, FOCUS detects the change in the networking behavior and enforces bandwidth reservations setup in the Network-Cloud slices #1 and #2, and in the on-path selected nodes (not shown in Figure 6) accordingly. As expected, the increase of Network-cloud slice #2 bandwidth reservation triggers the MQTT throughput to scale up (from 7.2 Mbps to 10.9Mbps), and hence leads to a decrease in average delay (from 66 ms to 45 ms), as Figure 9 confirms. The UDP traffic of Network-Cloud slice #1 increases at 45 seconds of the test runtime, although the system only detects the change around 7 seconds later (denoting a processing delay that can entails such process). Naturally, this interval risks jeopardizing the traffic of the Network-Cloud slice #1. However, it is beyond the scope of this paper to recommend and evaluate a particular elasticity mechanism/scheme. Following the readjustment of bandwidth reservations for

both Network-Cloud slices #1 and #2 to the initial amounts (i.e., 7 Mbps and 10 Mbps, respectively), the IoT traffic (Network-Cloud slice #2) falls back to around 7.2 Mbps. The UDP CBR traffic (Network-Cloud slice #3) remains constant (at 1.4Mbps) during the test runtime since the bandwidth amount for the Network-Cloud slice #3 does not change, and the Network-Cloud slice is also isolated from the two others in the entire data path.

All the behavioral patterns outlined above prove that FOCUS is able to deploy independent Network-Cloud slice approaches on top of the same shared system. In these, Network-Cloud slices #1 and #2 are served with different VNFs, and the Network-Cloud slice #3 runs a software application. Furthermore, the throughput outcomes demonstrate FOCUS is able to dynamically adapt bandwidth reservations in the Network-Cloud slice in a carrier-grade way, so that it can respond to network changes at runtime. Finally, FOCUS provides high-level isolation capabilities along the entire offering Network-Cloud Slice data paths, by avoiding traffic flows to interfere with each other (while maintaining the same bit rate patterns) at all the on-path links.

In the FON experiment, the traffic coming from the Private Wi-Fi slice faces the same events as for the Private Wi-Fi slice (#1) of the FOCUS pilot tests (i.e., the reduction of the sender's offered bit rate followed by an increase in the previous value 30 seconds later). In spite of the greater bandwidth availability at the link, the MQTT clients (Slice #2) are unable to consume the entire resources, keeping a bit rate of around 7.2 Mbps along all the test runtime. The reason for this constant bit rate pattern is that the FON system lacks features to reconfigure the queue's bandwidth reservations dynamically (i.e., Fonero only supports static bandwidth reservations, whereas FOCUS provides a dynamic SDN-based resource allocation approach). Furthermore, as the bandwidth isolation is only enforced between the Wi-Fi slices (Public and Private), the UDP traffic (sent at a Constant Bit Rate service model) competes with the MQTT

TCP traffic for the bandwidth resources, as evidenced by the slight variations on the UDP CBR traffic throughput of the Slice #2 depicted in Figure 8 (here the disturbance is more perceptible when examining the measured traffic's one way delay (OWD): for the FOCUS tests the average OWD is 13 milliseconds while for the FON tests it is 57 milliseconds). If there are more aggressive flows (such as UDP at higher rate), the CBR traffic (slice #3) consequently faces greater fluctuations in throughput. This undesirable outcome occurs as the result of a lack of traffic isolation that affects both traffic flows coming from the Public Wi-Fi slice (MQTT and UDP CBR). This happens because the flows compete for the same resources in the data path at the backhaul part that connects the Fonero node with the Carrier's cloud. The lack of service differentiation features tailored for the Public Wi-Fi slice is harmful for 5G UDN scenarios, where a wide range of applications require sophisticated (and specific) services that need to go far beyond providing poor best-effort capabilities.

The analysis of the outcome proves that the FOCUS scheme significantly outperforms the FON Wi-Fi sharing system in different ways, by providing high-level Network-Cloud slice end-to-end provisioning, while giving them independence from each other and offering multiservice customization in the form of chaining VNFs and/or software applications. Moreover, the FOCUS full "softwarization" approach enables Network-Cloud Slice resources to be dynamically enforced/reinforced in a carrier-grade way at running time, thus offering great flexibility and adaptability perspectives. Finally, FOCUS is capable of providing highly-isolated end-to-end (all the way from the Wi-Fi slice to the carrier cloud system) Network-Cloud slice instances. Hence, the ability to cope with the functional requirements foreseen by the UDN scenarios suggests that the FOCUS scheme is potentially able to expand 5G capacities.

6. Conclusion and Future Work

This study has investigated FOCUS, a solution that seeks to enhance 5G UDN techniques by making an efficient use of Wi-Fi-sharing technology for expanding resource perspectives. FOCUS achieves this goal by provisioning end-to-end Network-Cloud slices that cover the whole distance between Wi-Fi -sharing slice and the carrier cloud system. FOCUS allows an efficient sharing of Wi-Fi network resources through the coexistence of multiple slices tailored to the particular needs of corresponding users or things. By turning CPEs and other computing features belonging to the Wi-Fi-sharing network domain into programmable fog nodes, FOCUS was able to adopt a fully softwarized approach that can effectively meet the heterogeneous requirements of a myriad of scenarios on an end-to-end basis, by offering an alternative to traditional ossified Wi-Fi sharing systems such as the FON. FOCUS allows end-to-end high-level isolation and supports flexible carrier-grade resource allocation dynamic functions at the runtime. The prototyping evaluation demonstrates the basic functional principles that the FOCUS interoperable architecture allows. Moreover, the system-level analysis revealed that FOCUS outperforms the FON system to a remarkable degree by addressing the 5G

UDN functional requirements. FOCUS provisions high-level isolated end-to-end Network-Cloud slices, and features multiservice customization in the form of VNFs and software applications. This is because it is controllable in a carrier-grade manner at runtime through a complete end-to-end "softwarization."

In the following, we present key research directions that we intend to take as the next steps to improve FOCUS proposal. Such directions may also apply to the design of fog computing orchestration and management solutions in general, as they involve common concerns. In face of the highly heterogeneous fog infrastructure, (*i*) it is necessary to study means to estimate how specific (resource-constrained in a lot of cases) nodes will perform when executing a given VNF (fog nodes profile). Similarly, fog nodes not often reachable by clients should have a lower probability of being selected. The same happens with nodes located at places of high wireless interference (to explore machine learning techniques upon historical usage of fog nodes may be a possible direction). We intend (*ii*) to automate the decisions about where to deploy the VNFs/applications of a slice: at fog nodes, cloud, or both (since the chaining of the VNFs/applications can be deployed in a distributed way along these domains). It will be necessary to take into account different criteria such as QoS enforcement (e.g., latency limits and bandwidth guarantees) and optimization issues (e.g., VNFs/applications may be required to be deployed at the fog to reduce access network traffic towards the cloud. In addition, in case the slice offers mobility support to end nodes, the instantiation of the VNFs/applications at the cloud could be beneficial, as it could avoid the migration of such components among the WLAN CPE elements, a procedure that has a non-negligible cost). We also plan (*iii*) to conceive effective slice maintenance solutions that take into account the specificities of the fog layer such as the heterogeneity and location of fog nodes elements. For instance, in the case of WLAN CPE failure recovery, nodes compatibility must be taken into account (e.g. location and hardware requirements) when replacing a malfunctioning node. Due to the hardware heterogeneity, it may be the case that the malfunctioning node needs to be replaced by a group of others nodes able to meet the previously allocated resources. The same difficulties apply to services migration. Finally, (*iv*) a slice offering WLAN connectivity throughout a geographical area requires the employment of multiple WLAN VCPEs, instead of just one (e.g., a virtual operator offering data offloading at strategic zones such as the downtown). Such an arrangement increases the fog devices coordination complexity and requires further investigation. For instance, the effective distribution of resources among the selected WLAN CPEs may be a challenging task. Whether the required fog resources (node CPU, backhaul link bandwidth, etc.) will be equally distributed or not, depends upon the end nodes demand and may require readjustments/migration along the time. Also, the elasticity of the geographically spread fog slice (the capacity to increase/decrease the number of WLAN CPEs that form the fog slice at runtime according to the end nodes demands) is an important feature to be incorporated, although its effective realization is not trivial. For instance, the perception of a single mobile node at a

WLAN may not be sufficient to extend the slice to such a place, as the node can sooner move to be out of WLAN's range.

Data Availability

The network performance testing log data used to support the findings of this study have been deposited in the REGINA research group repository at https://github.com/reginagroup/focus/tree/master/results.

Conflicts of Interest

The authors declare that they have no conflicts of interest.

Acknowledgments

References

[1] M. Kamel, W. Hamouda, and A. Youssef, "Ultra-Dense Networks: A Survey," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 4, pp. 2522–2545, 2016.

[2] L. Evariste, D. Corujo, S. Jeon, J. Rodriguez, and L. Rui, *The 5G Internet*, chapter 2, John Wiley & Sons, 2015.

[3] F.-Y. Wang, L. Yang, X. Cheng, S. Han, and J. Yang, "Network softwarization and parallel networks: Beyond software-defined networks," *IEEE Network*, vol. 30, no. 4, pp. 60–65, 2016.

[4] B. A. A. Nunes, M. Mendonca, X.-N. Nguyen, K. Obraczka, and T. Turletti, "A survey of software-defined networking: past, present, and future of programmable networks," *IEEE Communications Surveys & Tutorials*, vol. 16, no. 3, pp. 1617–1634, 2014.

[5] B. Han, V. Gopalakrishnan, L. Ji, and S. Lee, "Network function virtualization: challenges and opportunities for innovations," *IEEE Communications Magazine*, vol. 53, no. 2, pp. 90–97, 2015.

[6] M. Carmo, S. Jardim, A. Neto, R. Aguiar, D. Corujo, and J. J. Rodrigues, "Slicing WiFi WLAN-Sharing Access Infrastructures to Enhance Ultra-Dense 5G Networking," in *Proceedings of the 2018 IEEE International Conference on Communications (ICC 2018)*, pp. 1–6, Kansas City, Mo, USA, May 2018.

[7] Y. Xiao and M. Krunz, "Dynamic network slicing for scalable fog computing systems with energy harvesting," *IEEE Journal on Selected Areas in Communications*, vol. 36, no. 12, pp. 2640–2654, 2018.

[8] ITU-R M.2320, "Future Technology Trends of Terrestrial IMT Systems," Tech. Rep., ITU-R, Nov 2014.

[9] S. Chen, F. Qin, B. Hu, X. Li, and Z. Chen, "User-centric ultra-dense networks for 5G: Challenges, methodologies, and directions," *IEEE Wireless Communications Magazine*, vol. 23, no. 2, pp. 78–85, 2016.

[10] A. Asheralieva, T. J. Erke, and K. Kilkki, "Traffic characterization and service performance in FON network," in *Proceedings of the 2009 1st International Conference on Future Information Networks, ICFIN 2009*, pp. 285–291, China, October 2009.

[11] ETSI Group, *Network Functions Virtualisation (NFV)*, Use Cases, 2013.

[12] S. Tonks, *The Benefits of Virtual CPE - Business User*, White Paper, February 2016, https://builders.intel.com/docs/networkbuilders/The-benefits-of-virtual-CPE-business-user.pdf.

[13] M. Gao, B. Addis, M. Bouet, and S. Secci, "Optimal orchestration of virtual network functions," *Computer Networks*, vol. 142, pp. 108–127, 2018.

[14] E. Hernandez-Valencia, S. Izzo, and B. Polonsky, "How will NFV/SDN transform service provider opex?" *IEEE Network*, vol. 29, no. 3, pp. 60–67, 2015.

[15] T. Cruz, P. Simões, N. Reis, E. Monteiro, F. Bastos, and A. Laranjeira, "An architecture for virtualized home gateways," in *Proceedings of the IFIP/IEEE International Symposium on Integrated Network Management (IM 2013)*, pp. 520–526, IEEE, Ghent, Belgium, May 2013.

[16] J. Lin, H. Bannazadeh, P. Spachos, and A. Leon-Garcia, "SAVI vCPE and internet of things," in *Future Access Enablers for Ubiquitous and Intelligent Infrastructures*, V. Atanasovski and A. Leon-Garcia, Eds., vol. 159, pp. 18–25, Springer International Publishing, Cham, Switzerland, 2015.

[17] H. Zhu and C. Huang, "IoT-B&B: Edge-Based NFV for IoT Devices with CPE Crowdsourcing," *Wireless Communications and Mobile Computing*, vol. 2018, Article ID 3027269, 15 pages, 2018.

[18] A. Lombardo, A. Manzalini, G. Schembra, G. Faraci, C. Rametta, and V. Riccobene, "An open framework to enable NetFATE (Network Functions at the edge)," in *Proceedings of the 1st IEEE Conference on Network Softwarization (NETSOFT '15)*, pp. 1–6, London, UK, April 2015.

[19] R. Bonafiglia, S. Miano, S. Nuccio, F. Risso, and A. Sapio, "Enabling NFV Services on Resource-Constrained CPEs," in *Proceedings of the 5th IEEE International Conference on Cloud Networking, CloudNet 2016*, pp. 83–88, Italy, October 2016.

[20] B. O'Hara, P. Calhoun, and J. Kempf, "Configuration and Provisioning for Wireless Access Points (CAPWAP) Problem Statement," *RFC 3990*, February 2005.

[21] W. Enck, T. Moyer, P. McDaniel et al., "Configuration management at massive scale: System design and experience," *IEEE Journal on Selected Areas in Communications*, vol. 27, no. 3, pp. 323–335, 2009.

[22] A. E. Nikolaidis, S. S. Papastefanos, G. I. Stassinopoulos, M.-P. K. Drakos, and G. A. Doumenis, "Automating remote configuration mechanisms for home devices," *IEEE Transactions on Consumer Electronics*, vol. 52, no. 2, pp. 407–413, 2006.

[23] TR-069 Amendment 5, "TR-069 CPE WAN Management Protocol," Tech. Rep., Broadband Home Technical Working Group, Nov 2013.

[24] K.-M. Lee, W.-G. Teng, and T.-W. Hou, "DRASE: A dynamic rescheduling and self-adaptive estimation technique to enhance ACS throughputs in CWMP," *IEEE Communications Letters*, vol. 20, no. 11, pp. 2161–2164, 2016.

[25] A. E. Nikolaidis, S. Papastefanos, G. A. Doumenis, G. I. Stassinopoulos, and M. P. K. Drakos, "Local and remote management integration for flexible service provisioning to the

home," *IEEE Communications Magazine*, vol. 45, no. 10, pp. 130–138, 2007.

[26] T. Cruz, P. Simões, E. Monteiro, F. Bastos, and A. Laranjeira, "A framework for internet media services delivery to the home environment," *Journal of Network and Systems Management*, vol. 21, no. 1, pp. 99–127, 2013.

[27] L. Zheng, Y. Hu, and S. Chen, "Research and application of CWMP in distributed network management system," in *Proceedings of 2012 International Conference on Computer Science and Service System, CSSS 2012*, pp. 647–650, China, August 2012.

[28] C. Xu, W. Jin, G. Zhao, H. Tianfield, S. Yu, and Y. Qu, "A novel multipath-transmission supported software defined wireless network architecture," *IEEE Access*, vol. 5, pp. 2111–2125, 2017.

[29] O. Sallent, J. Perez-Romero, R. Ferrus, and R. Agusti, "On radio access network slicing from a radio resource management perspective," *IEEE Wireless Communications Magazine*, vol. 24, no. 5, pp. 166–174, 2017.

[30] Z. Yao, W. Zhou, L. Yang, S. Govindan, and H. Cheng, "Objectives for Control and Provisioning of Wireless Access Points (CAPWAP)," *RFC 4564*, July 2006.

[31] M. Bernaschi, F. Cacace, G. Iannello, M. Vellucci, and L. Vollero, "OpenCAPWAP: An open source CAPWAP implementation for the management and configuration of WiFi hot-spots," *Computer Networks*, vol. 53, no. 2, pp. 217–230, 2009.

[32] M. Bernaschi, F. Cacace, A. Davoli, D. Guerri, M. Latini, and L. Vollero, "A CAPWAP-based solution for frequency planning in large scale networks of WiFi Hot-Spots," *Computer Communications*, vol. 34, no. 11, pp. 1283–1293, 2011.

[33] A. Levanti, F. Giordano, and I. Tinnirello, "A CAPWAP architecture for automatic frequency planning in WLAN," in *Proceedings of the 12th IEEE International Symposium on Computers and Communications, ISCC '07*, pp. MW51–MW56, Portugal, July 2007.

[34] A. Sen and K. M. Sivalingam, "An SDN framework for seamless mobility in enterprise WLANs," in *Proceedings of the 26th IEEE Annual International Symposium on Personal, Indoor, and Mobile Radio Communications, PIMRC 2015*, pp. 1985–1990, China, September 2015.

[35] K.-K. Yap, M. Kobayashi, R. Sherwood et al., "OpenRoads: Empowering research in mobile networks," *SIGCOMM Computer Communication Review*, vol. 40, no. 1, pp. 125-126, 2010.

[36] T. Lei, X. Wen, Z. Lu, and Y. Li, "A semi-matching based load balancing scheme for dense IEEE 802.11 WLANs," *IEEE Access*, vol. 5, pp. 15332–15339, 2017.

[37] A. Patro and S. Banerjee, "Coap: A software-defined approach for home wlan management through an open api," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 18, no. 3, pp. 32–40, 2015.

[38] C.-W. Ahn and S.-H. Chung, "SDN-based mobile data offloading scheme using a femtocell and wifi networks," *Mobile Information Systems*, vol. 2017, Article ID 5308949, 15 pages, 2017.

[39] L. Sequeira, J. L. De La Cruz, J. Ruiz-Mas, J. Saldana, J. Fernandez-Navajas, and J. Almodovar, "Building an SDN enterprise WLAN based on virtual APs," *IEEE Communications Letters*, vol. 21, no. 2, pp. 374–377, 2017.

[40] A. C. Baktir, A. Ozgovde, and C. Ersoy, "How can edge computing benefit from software-defined networking: a survey, use cases, and future directions," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 4, pp. 2359–2391, 2017.

[41] J. Saldana, R. Munilla, S. Eryigit et al., "Unsticking the wifi client: smarter decisions using a software defined wireless solution," *IEEE Access*, vol. 6, pp. 30917–30931, 2018.

[42] E. Coronado, R. Riggio, J. Villalon, and A. Garrido, "Joint mobility management and multicast rate adaptation in software–defined enterprise WLANs," *IEEE Transactions on Network and Service Management*, vol. 15, no. 2, pp. 625–637, 2018.

[43] S. M. M. Gilani, T. Hong, W. Jin, G. Zhao, H. M. Heang, and C. Xu, "Mobility management in IEEE 802.11 WLAN using SDN/NFV technologies," *EURASIP Journal on Wireless Communications and Networking*, vol. 1, no. 67, 2017.

[44] J. Chen, B. Liu, H. Zhou, Q. Yu, L. Gui, and X. S. Shen, "QoS-driven efficient client association in high-density software-defined WLAN," *IEEE Transactions on Vehicular Technology*, vol. 66, no. 8, pp. 7372–7383, 2017.

[45] A. V. Dastjerdi and R. Buyya, "Fog computing: helping the internet of things realize its potential," *The Computer Journal*, vol. 49, no. 8, pp. 112–116, 2016.

[46] L. M. Vaquero and L. Rodero-Merino, "Finding your way in the fog: Towards a comprehensive definition of fog computing," *SIGCOMM Computer Communication Review*, vol. 44, no. 5, pp. 27–32, 2014.

[47] C. C. Byers, "Architectural Imperatives for Fog Computing: Use Cases, Requirements, and Architectural Techniques for Fog-Enabled IoT Networks," *IEEE Communications Magazine*, vol. 55, no. 8, pp. 14–20, 2017.

[48] Z. Wen, R. Yang, P. Garraghan, T. Lin, J. Xu, and M. Rovatsos, "Fog orchestration for internet of things services," *IEEE Internet Computing*, vol. 21, no. 2, pp. 16–24, 2017.

[49] P. Bellavista, A. Corradi, and A. Zanni, "Integrating mobile internet of things and cloud computing towards scalability: Lessons learned from existing fog computing architectures and solutions," *International Journal of Cloud Computing*, vol. 6, no. 4, pp. 393–406, 2017.

[50] F. Bonomi, R. Milito, P. Natarajan, and J. Zhu, "Fog computing: A platform for internet of things and analytics," in *Big Data and Internet of Things: A Roadmap for Smart Environments*, vol. 546 of *Studies in Computational Intelligence*, pp. 169–186, Springer International Publishing, Cham, Switzerland, 2014.

[51] P. Hu, S. Dhelim, H. Ning, and T. Qiu, "Survey on fog computing: architecture, key technologies, applications and open issues," *Journal of Network and Computer Applications*, vol. 98, pp. 27–42, 2017.

[52] S. Agarwal, S. Yadav, and A. K. Yadav, "An efficient architecture and algorithm for resource provisioning in fog computing," *International Journal of Information Engineering and Electronic Business*, vol. 8, no. 1, pp. 48–61, 2016.

[53] L. Ni, J. Zhang, C. Jiang, C. Yan, and K. Yu, "Resource allocation strategy in fog computing based on priced timed petri nets," *IEEE Internet of Things Journal*, vol. 4, no. 5, pp. 1216–1228, 2017.

[54] H. R. Arkian, A. Diyanat, and A. Pourkhalili, "MIST: Fog-based data analytics scheme with cost-efficient resource provisioning for IoT crowdsensing applications," *Journal of Network and Computer Applications*, vol. 82, pp. 152–165, 2017.

[55] M. Mukherjee, R. Matam, L. Shu et al., "Security and Privacy in Fog Computing: Challenges," *IEEE Access*, vol. 5, pp. 19293–19304, 2017.

[56] TOSCA, "TOSCA Simple Profile for Network Functions Virtualization (NFV)," 2013.

[57] C. Nedelcu, *Nginx HTTP Server*, Packt Publishing, 3rd edition, 2015.

[58] C. Maxweel, S. Thalyson, A. Medeiros, N. Augusto, and S. Jardim, "Carrier-Grade SDN-Controlled WLAN-Sharing: a Performance Evaluation of OpenFlow-Enabled Commodity-based Hardware Networking Nodes," in *Proceedings of the Workshop on Management and Operation of Networks and Services (WGRS)*, pp. 44–54, Brazilian Symposium on Computer Networks and Distributed Systems (SBRC), May 2017.

[59] D. Belson, "Akamai state of the internet/connectivity report: Q1 2017," Tech. Rep., Akamai Technologies, 2017.