

Micro and Macro Network Slicing: an experimental assessment of the impact of increasing numbers of slices

Flavio Meneses* | Rui Silva | Daniel Corujo | Rui L. Aguiar

Instituto de Telecomunicações and Universidade de Aveiro, Aveiro, Portugal

*Correspondence: flaviomeneses@av.it.pt

Abstract

The Fifth Generation (5G) telecommunications network aims not only to enhance traffic performance and allow efficient management, but also to enable it to dynamically and flexibly adapt to the traffic demands of different vertical scenarios. In order to support that enablement, the underlying network procedures (i.e., network functions) are being virtualized and deployed in cloud-based environments, allowing for a more optimized usage of the infra-structure resources. In addition, such resources can be sliced, allowing isolated provisioning to specific network functions allocated to disparate vertical deployments. As network slices are envisaged by network operators to fulfill a small number of slices, able to cater towards essential 5G scenario demands (i.e., enhanced mobile broadband, massive machine-type communications and ultra reliable low-latency communications), the total amount of slices existing in a system is currently dictated by the underlying operational overhead placed over the cloud infra-structure. This paper explores the challenges associated to a vision where the network slicing concept is applied with a much greater level of granularity, ultimately allowing it to become a core mechanism of the network's operation, with large numbers of co-existing slices. In that respect, this paper proposes an architecture framework for instantiation of network slices among network providers, which in turn are able to instantiate sub-slices tailored to use cases and vertical tenants. The evaluation of this concept is done following a two-pronged approach: firstly, different slice dimensions (i.e., from micro to macro) are proposed and discussed, pointing out the benefits and challenges of each proposed slice; secondly, we deployed a mobile network provider (MNO), using OpenAirInterface and FlexRAN frameworks, and experimentally evaluated the its slicing mechanisms. The objective is to provide insight on the challenges and impact associated with the deployment of an increasing amount of slices, using the same available infra-structural resources.

Keywords: Network Slicing, SDN, NFV.

Personal use of this material is permitted. Permission from Springer must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. This is the final submitted copy of a paper accepted for publication in the Wireless Personal Communications journal.

I. INTRODUCTION

With the dawn of the fifth network generation (5G)[1], research and industry forums have

been contributing towards a heterogeneous network capable of supporting a wide range of use cases requirements, while optimising the

network to fulfil specific vertical demands [2]. Such networks requirements range from the enhance mobile broadband (eMBB) high peak data rates, to the ultra-reliable low latency (uRLL) communications and to the massive machine type communications (mMTC) [2]. In this context, mobile network operators (MNOs) are looking for ways to efficiently orchestrate and manage their network, allowing to dynamically (re)configure the network to meet the traffic and vertical demands.

In this line, network slicing has been pointed out as a key enabler for this network reconfigurability. Network slicing proposes the partition of the infrastructure in logical slices, where each slice is viewed as an isolated network. For this, Network Function Virtualisation (NFV) and Software Defined Networking (SDN) have been the enablers for such dynamic network [2]. On one hand, NFV [3] allows the decoupling of the software functions from the hardware characteristics, enabling network functions (NFs) to be deployed in generic hardware in data-centres. On the other hand, SDN [4] decouples the data and control planes and centralises the network intelligence in a high-level entity, namely the SDN controller. Thus, the SDN controller has a broader view of the network, and uses southbound (SB) and northbound (NB) application programming interfaces (APIs) to dynamically (re)configure the data-plane (i.e., forwarding devices) and to communicate with high-level applications (i.e., SDN applications), respectively. In this context, the Open Network Foundation (ONF) has been pushing the standardisation of OpenFlow (OF) [5] as an SB API. Conversely, Representational State Transfer (REST) has been widely used as a NB API. Currently, network slicing is presented as the chaining and dynamic (re)configuration of physical and virtual network functions (PNFs and VNFs, respectively). As such, while NFV enables NFs to be moved across the network, SDN allows the dynamic (re)configuration of the network in order to chain the necessary NFs to fulfil demands of the verticals.

Despite the deployment flexibility and resources isolation allowed by the utilization of

these concepts for the enablement of network slicing, the introduction of these procedures adds an added layer of complexity towards the provisioning of connectivity to end-users. As operators [6] and manufacturers [7] have been pointing out, the increase in number of slices meets scalability issues associated to today's technical limitations of the underlying infrastructure, despite the high business monetization potential that such increase might bring.

This paper embraces a vision where service providers are able to use the network slicing concept as the key building block of dynamically tailored-based connectivity solution towards verticals. In this way, the paper proposes a network architecture where the infrastructure can be sliced in logically isolated networks, allowing the infrastructure sharing among network providers. Also, verticals are able to specify their requirements allowing network providers to reconfigure their access networks and partition them in sub-slices. In this context, the network provider (e.g., MNO) can dynamically instantiate a sub-slice for fulfilling the sporadic necessity of a sport event, or rather deploying microslices for uRLLC scenarios. The evaluation of this concept is done following a two-pronged approach: firstly, different slice dimensions (i.e., from micro to macro) are proposed and discussed, pointing out the benefits and challenges of each proposed slice; secondly, we deployed a mobile network provider (MNO), using OpenAirInterface and FlexRAN frameworks, and experimentally evaluated its slicing mechanisms. The objective is to provide insight on the challenges and impact associated with the deployment of an increasing amount of slices, using the same available infrastructural resources.

The remainder of this paper is organized as follows: Section II explores the background while presenting network sling initiatives. Our architecture is proposed in Section III presenting the main building blocks and involved entities. Section IV illustrates different slice dimensions and respective benefits for verticals, with possible implementation being presented in Section V. In Section VI a proof-of-concept

framework is deployed and evaluated in Section VII. Finally, the paper concludes in Section VIII.

II. RELATED WORK

Network slicing is the logical partitioning of physical network resources and the combination/chaining of such resources in logical slices, proving to the user a view of an unique network [8]. In this line, end-to-end (E2E) slicing implies the slicing of both radio (or spectrum) and infrastructure [8]. The former is related to the time, space or frequency multiplexing of the spectrum, and the latter to the slicing of physical network resources. Nevertheless, these slice layers should be stitched together in order to compose and E2E slice transparent to users [2]. In this context, projects such as [9, 10] aim to allow the means towards the provision and management of slices tailored to vertical industries, by mapping not only network functions and infrastructure, but also service layer agreements (SLA) of the services and vertical requirements.

The slicing of wireless networks adds complexity to the network slicing concept, due to the variability of wireless links' capacity (which in turn depends on available bandwidth) and the limited resources (limited spectrum) [8]. In this context, wireless slicing is usually technology dependant, since 3GPP and non-3GPP access technologies implement different medium access techniques [8]. In this line, in LTE, medium access control (MAC) is performed by Orthogonal Frequency Division Multiple Access (OFDMA) on downlink and by Single Carrier FDMA on the uplink. Thus, slicing proposals for LTE usually consider the enhancement of scheduler algorithms to decide the amount of radio blocks assignment, taking in account QoS requirements, as proposed in works such as [11, 12, 13, 14]. Additionally, in [15], the authors analyse different RAN slicing approaches comparing them in terms of granularity in the assignment of radio resources and the degrees of isolation and customization. Otherwise, MAC in Wi-Fi for devices in infrastructure

mode is usually set in Distributed Coordination Function (DCF) or Enhanced Distributed Channel Access (EDCA) [8]. Thus, works such as [16, 17, 18, 19] propose enhanced algorithms for packet scheduling. Notwithstanding, network-layer approaches such as [20, 21] try to avoid such low level strategies and propose the resource allocation at higher layers. In this context, such approaches try to slice the Wi-Fi access point (AP), using multiple virtual machines (VMs) over the physical AP, scheduling the use of transmission resources between slices. Also, traffic shaping approaches ([22, 23]) aim to control the traffic sent to the scheduler, modulating its behavior.

Besides the radio access component, Infrastructure slicing has been being pushed by network virtualisation [8]. In this line, works such as [24, 25] propose the softwarization and virtualisation of the mobile network, advocating that the use of SDN and NFV technologies adds a greater degree of flexibility. In this context, OpenAirInterface (OAI) [26] is an opensource software-based implementation of the LTE complaint with 3GPP standards. Providing enhancements to the OAI, Mosaic5G [27] consortium, presents LL-MEC [28] and FlexRAN [29]. On one hand, LL-MEC [28] enhances OAI by following SDN and MEC principles, resulting in a low latency multi-access edge computing (MEC) platform. On the other hand, FlexRAN [29], provides a flexible control over the RAN infrastructure allowing to dynamically instantiate radio slices. Finally, works such as [30] and [31], provide insight on the E2E orchestration of slicing mechanisms for verticals.

Nonetheless, to the best of the authors' knowledge, there is no analysis or experimental assessment of the impact and possibilities that the increasing amount of slices over a system have. This is where our paper contributes, by defining different types of network slice granularity and their deployment feasibility, and by experimentally assessing the dynamic instantiation of sub-slices by network providers, in a dynamic reconfiguration scenario.

III. FRAMEWORK OVERVIEW

In our architecture proposal, network providers (such as, Mobile Network Operators (MNO)) request a network slice to the infrastructure provider, in order to deploy their access network. For this, infrastructure providers offer an service orchestrator that can be interfaced, allowing network providers to announce their requirements. This allows infrastructure sharing among multiple network providers, since network providers request a set of physical and virtual network functions (PNF and VNF, respectively) from the infrastructure provider, which in turn slices the infrastructure (i.e., Base Stations, Access Points, forwarding devices, data-centres) and chains the requested PNFs and VNFs. This is portrayed as Macro Slices, in our architecture. Additionally, network providers, when requesting resources to the infrastructure provider, tailor such requests towards the construction of a suitable network slice, as required by the vertical and traffic demands of each use case. As such, network providers offer a vertical manager, allowing each vertical to specify their requirements. These compose the Micro Slices, in our architecture. Fig. 1 illustrates the architecture proposal, with the different involved layers being described as follows.

i. Infrastructure provider

In our proposal the infrastructure provider is responsible for instantiating network slices

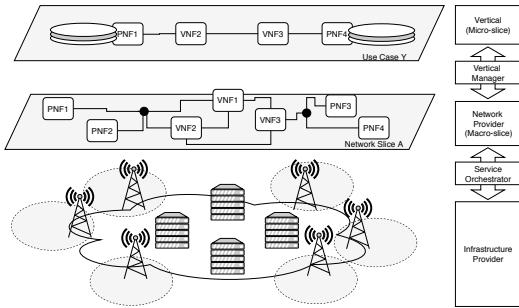


Figure 1: Architecture overview.

when requested by network providers. For this, the infrastructure providers offer a service orchestrator where network providers can choose the necessary PNFs and VNFs from a catalogue, to meet a vertical’s demands. Thus, the infrastructure is responsible for chaining the requested NFs in a logical isolated network.

ii. Network provider

In our proposal, network providers do not own the infrastructure (i.e., physical equipment). Instead, network providers (such as, MNOs) deploy their network over a shared infrastructure by requesting a slice to the infrastructure provider. Here, a network slice is defined as a set of isolated PNFs and VNFs chained together, resulting a logical isolated network. Also, network providers are able to instantiate sub-slices by reconfiguring their NFs, adapting the virtual access network to each vertical and/or use case.

iii. Verticals and Use cases

Verticals are defined as specific industrial use cases such as the automotive and eHealth. In this context, the network provider offers a vertical manager, allowing verticals to specify the necessary traffic requirements. Such information is then used by network operators to (re)configure their slices, in order to adapt the network to the specific traffic demands of each vertical (different types and possibilities of vertical slices are presented in section IV).

iv. Service Orchestrator

The service orchestrator exposes the infrastructure layer to network providers, while coordinating the multiple requests from such providers. As such, the service orchestrator creates a high-level business service to automate network slices instantiation, providing isolation and security among slices and network providers.

v. Vertical Manager

The vertical manager exposes the network provider features and available resources to the vertical tenant. Here, verticals announce the requirements of the use case, and the SLAs are defined. The manger then (re)configures the network slice (e.g., chains of PNFs and VNFs) in order to tailor the slice to the traffic demands.

IV. VERTICAL SLICES AND USE CASES

As discussed previously, network providers, such as MNOs, request to the infrastructure provider PNFs and VNFs (forming a network slice) to deploy their network. Nevertheless, MNOs are able to (re)configure their virtual networks (i.e., network slices), in order to meet traffic demands of each vertical. In this section we propose and analyze a classification of network slices in terms of granularity, going beyond the currently accepted general notion of network slice. Fig. 2 illustrates the proposed slices dimensions, while Table 1 summarizes them by providing use case examples.

In this line, the MNO requests an operator slice to the infrastructure provider, which in turn grants PNFs and VNFs. The MNO is responsible for the orchestration of its slice, allowing it to (re)configure the requested NFs. Thus, the MNO is able to dynamically instantiate sub-slices with different sizes and requirements, according to the use cases' and/or verticals' demands.

Fig. 2a exemplifies macroslices instantiated for covering bigger areas of the network. For example, to meet the sporadic or seasonal traffic demands of a sport event, the MNO may




Figure 2: MNO's sub-slice dimensions.

instantiate a geographical slice, in order to enforce better coverage and enhance QoS and QoE for its users during the event. Also, the MNO may create different classes of users (e.g., premium, regular and low cost) and instantiate a slice with different requirements and QoS for each class.

Current smartphones are able to simultaneously use multiple wireless access technologies (e.g., LTE and Wi-Fi), thus the MNO has the possibility of instantiating a slice per access technology. This allows the traffic redirection to the less crowded access network for QoE enhancement and/or the traffic offloading from the licensed to the unlicensed spectrum (i.e., from LTE to Wi-Fi). Notwithstanding, alternatively, for UE's redundancy and throughput enhancement, the MNO may instantiate a slice covering multiples access technologies of the UE. Also, despite UEs having multiple interfaces of the same access technology, a slice can be instantiated per interface, allowing a UE to be attached to different MNOs simultaneously.

Nevertheless, use cases such as eHealth and automotive require high reliability, with the presence of microslices becoming important (Fig. 2b). In this line, scenarios where network interfaces (NIC) of the nodes (or UEs) attach to multiple slices simultaneously, allow the UE to prioritize certain types of flows. Also, in more specific scenarios, packets of the same flow are transmitted/received in parallel, allowing a greater degree of redundancy, reliability and/or throughput.

V. SLICING INITIATIVES FOR DIFFERENT SLICE DIMENSIONS AND IMPLEMENTATION EFFORTS

As reviewed in section II, different slicing initiatives and proof-of-concept frameworks have been proposed in recent years. However, the majority of such proposals often focus in specific use cases, without presenting their integration with the global architecture, which in turn demands great flexibility in order to tailor a slice for the verticals. In addition, such existing works also address the subject assuming a sin-

Table 1: *Types of slices and use cases.*

Type of slice	Description	Vertical / Use case
Operator	Network slice instantiated by the infrastructure provider with multiple PNFs and VNFs.	Mobile network operator (MNO).
Geographic area	Slice instantiated in geographical area for sporadic gathering.	Sport event or concert in a stadium.
Type of client	Instantiated by the tenant MNO over the operators slice by (re)configuring the requested PNFs and VNFs.	Definition of classes of users for QoS and QoE.
Access technology	The MNO offers network access through different wireless access networks (e.g., LTE and Wi-Fi), instantiating a slice per access technology.	Mobile video offloading.
Node/UE	Independently of the number of UE's network interfaces, the MNO instantiates one slice for the UE.	eMBB (high throughput, since the UE uses multiple interfaces simultaneously).
Interface	For UEs with multiple network interfaces for the same wireless access technology (e.g., dual-sim smartphones) a slice is instantiated per interface.	Resilience and redundancy
Flow	The MNO instantiates slices per types of flow and/or service, allowing UEs to attach to multiple slices simultaneously.	QoS and QoE per flow and/or service (traffic shaping).
Packet	The MNO instantiates multiple slices to which UEs attach allowing the parallel transmission of packet from the flow.	uRLLC and eHealth.

gle type of network slice (or at most, just distinguish between radio access or network services slices at the core). In this line, next we review recent slicing initiatives and proof-of-concept frameworks, viewing them under the lens of our network slicing granularity classification and map them according to our architecture proposal.

Operator: Despite existing works (such as, [9, 10]) proposing infrastructure sharing among network providers, as far as we know, few actually experimentally implement such approach. For example, OAI [26] along with FlexRAN [29] allows to slice the MNOs infrastructure, by virtualising the EPC in a data-centre and instantiate multiple radio slices, respectively. However, such frameworks do not allow the RAN sharing, requiring each MNO

to deploy its own RAN infrastructure (more specifically, eNBs). In this line, it is necessary to develop mechanisms allowing eNBs to be shared among MNOs, in order to allow full infrastructure sharing. Regarding to non-3GPP infrastructure sharing initiatives, works such as [32, 33] present proof-of-concept frameworks to dynamically instantiate Wi-Fi slices, which redirect traffic to the respective core networks, allowing APs deployed in public places to be shared among MNOs.

Geographic area: Frameworks such as FlexRAN [29] and OAI [26] enhanced the deployment flexibility of mobile networks. MEC and fog computing are enablers for this type of slices, as the closer proximity of the network function deployment infrastructure to the end users, allow for geographical-based

performance gains and resource usage optimization. In this line, by enhancing such frameworks, it is possible to dynamically manage the control and data planes in order to instantiate a geographical slice, and use MEC to reduce latencies.

Type of client: Frameworks such as FlexRAN [29], allow the RAN slicing for MNOs, with the instantiated RAN slices being attached to the MNO's EPC. Similarly, in [32], multiple Wi-Fi slices were instantiated with different QoS. Such approaches allow MNOs to instantiate 3GPP and non-3GPP slices for different classes of users, providing different QoS for each slice.

Access technology: Currently, proposed architectures in the literature provide mechanisms to create network slices for different access technologies, disregarding the interoperability between slices of different access technologies. In this line, [29] proposes the slicing of 3GPP networks, while works such as [33] slice non-3GPP points of attachment (PoA) for enhanced QoS and/or infrastructure sharing among operators. However, providing interoperation capabilities between both slices requires both the development of new mechanisms as well as the enhancement of existing procedures. In [32], the authors took the first steps towards this direction and developed a framework that dynamically instantiated a non-3GPP slice for the UE, allowing mobile video offloading while maintaining the QoE.

Node/UE: Aligned with the interoperability between slices of different access technologies, we propose the instantiation of slices per UE. This allows the access technology currently in use by the UE to become transparent, enabling seamless handover scenarios. In this line, the use of a virtualized representation of the UE inside the network provider's virtualized network functions (as in [32]) anchors the UE to the network, allowing the network provider to flexibly move the slice among the different access technologies. Alternatively, the use of

bond interfaces allows the unification of access technologies in the UE, which also allows the transparency of multiple slices to the end-user.

Interface: Here, a different slice is instantiated depending if interfaces are from same access technology, or not. For interfaces the same access technology (e.g., dual-SIM smartphone), the MNO instantiates a slice per interface. For interfaces of different access technologies, the slice is similar to the access technology slice type. This type of slicing allows for greater UE connectivity resilience in case of network failure, and despite being supported by our current slicing implementation, mechanisms for slice handover management in inter-domain scenarios need to be further developed.

Flow: This is the most frequent type of radio slicing. Works such as [19, 14] propose algorithms for the allocation of resource blocks depending on the traffic characteristics. Conversely, [23] noted that such approaches do not guarantee the uplink traffic demands in scenarios with "greedy" flows, and propose instead the use of SDN technologies within the UE, for dynamic queuing of flows.

Packet: This type of slicing aims to offer to the UE the capability to attach the same interface to multiple slices, allowing the parallel transmission of packet of the flow, with the benefit of enhanced throughput, redundancy and resilience. However, as far as we know, currently there is no slicing mechanisms in the literature that allows the UE to dynamically send packets of the same flow through different slices.

VI. USE CASE AND IMPLEMENTATION

Fig. 3a presents the architecture of the instantiated MNO's slice, where the evolved packet core (EPC) was instantiated in a in-house data-centre running OpenStack (Ocata). The deployed EPC was based in the OAI project, enhancing its flexibility by introducing a SDN controller capable of receiving information

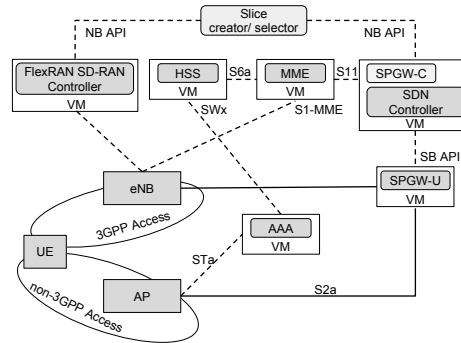
from NB SDN application helping in the network (re)configuration. In this line, 3GPP entities, such as the mobility management entity (MME), the home subscriber server (HSS) and the service and packet gateway (S/P-GW), were instantiated in VMs with a 1 CPU core and 2 GB of RAM, running the Ubuntu 16.04 LTS operating system (OS). Regarding to the RAN, its flexibility was enhanced by deploying the FlexRAN [34] framework. In this context, our framework is composed by both the FlexRAN and the OAI, allowing the development of SDN NB applications to acquire context from both core and ran networks, and assist the SDN controller in the core and access networks management via NB APIs. The eNodeB (eNB) was deployed in a physical machine running Ubuntu desktop 16.04 LTS OS with an Intel Core i7-7700K CPU and 32 GB of RAM, and an USRP B210 software-defined radio (SDR) board with a LP0965 antenna.

i. Scenario


For simplification, due to the complexity of infrastructure sharing among MNOs, in our scenario the slice for the MNO is already instantiated. In this context, the MNO will reconfigure its building blocks (PNFs and VNFs) according to requests from verticals. Fig. 3b illustrates the messages sequence for the reconfigurability of the framework upon the attachment of an UE. Additionally, for our scenario, we developed an SDN NB application (i.e., the *slice creator/selector*) to assist the SDN controller in the management of UEs' requirements. As such, upon an UE attachment, the *slice creator/selector* verifies the requirements of the UE, and if necessary requests a new slice to the FlexRAN SD-RAN controller with the necessary parameters and/or a slice handover.

In this context, when a UE attaches to the 3GPP network, the MME informs the SDN controller via NB APIs, through REST messages we developed (Fig. 3b:1). The SDN controller then reconfigures the SPGW-U (Fig. 3b:2) via SB APIs (i.e., OF for flow based rules and OVSDB for tunnel creation), enabling the com-

munication of the UE. Initially, the UE is attached to a default slice. Nevertheless, after attachment completion, the SDN controller informs the *slice creator/selector*, which in turn verifies the UE's characteristics (e.g., type of client, current flow in use, etc.) and (if necessary) requests to the FlexRAN SD-RAN controller to move the UE to the correct slice (or asks for a new one, moving the UE afterwards).



(a) Deployed framework architecture.



(b) High-level sequence of messages for the scenario evaluation.

Figure 3: Deployed framework and sequence of messages for its reconfigurability.

VII. EVALUATION AND DISCUSSION

In this section we evaluate the proposed framework in terms of UE attachment, radio slice instantiation and slice handover delays, and the impact caused by the simultaneous existence of requests and/or increasing number of instantiated slices. Experiments were run 50 times, with the results being shown in Fig. 4 and presented with a confidence interval of 95%.

i. Slice radio instantiation delay

As previously discussed, our deployment uses both OAI and FlexRAN. In this context, FlexRAN only allows to instantiate a maximum of 10 radio slices. Fig. 4a) presents the instantiation delay of upon a request of multiple radio slices. The results show that the slice instantiation does not depend on the number of requested slices, as of the current version of the FlexRAN software, for the amount of simultaneous slices used in our experiments.

In Fig. 3b we note that the request is sent by our SDN NB application via REST (more specifically, HTTP POST message). However, the FlexRAN SD-RAN controller sends a acknowledgement before the correct implementation of the slice, misleading the slice requester to move UEs to a slice before its correct implementation. To overcome this, we set a delay between the slice instantiation acknowledgement and the UE's slice handover.

ii. Slice handover delay

From Fig. 4b) it is noted that the slice handover of the UE is independent of the number of instantiated slices. As such, in Fig. 3b, when the requester asks for a slice handover of the UE, the network takes about 2s to switch the radio slice that the UE is attached to. In this context, we argue that in scenarios where the network uses handovers from one slice to another as the means to fulfil UE requirements (e.g., better latency), the slice handover delay is required to be much lower in order not to incur any impact to the experienced QoE (especially when considering uRLL scenarios). As such, enhanced mechanisms should be developed in order to increase handover performance.

iii. S/P-GW update delay

Fig. 4c) presents a scalability study of our SDN mechanism for S/P-GW flow control. For this, our MME emulated the attachment of multiple UE simultaneously and we measured the delay for correct flow implementation for the requested number of UEs. The experiment

uses a worst case scenario, since requests are sent consecutively to the SDN controller. Results show that our mechanism escalates well, outperforming a linear scenario.

iv. UE attachment delay

This experiment evaluates the total delay of the UE attachment in a dynamically instantiated slice. Thus, this delay results in the sum of the delay for slice radio instantiation, update of the S/P-GW and slice handover (HO). Fig. 4d) presents the cumulative distribution function (eCDF) of the delay for each stage completion (as such, each stage accounts the delay of previous stages). In Fig. 3b, when an UE attaches to the 3GPP network, the MME informs the SDN controller (Fig. 3b:1) which in turns updates the S/P-GW-U (Fig. 3b:2). Such update takes about 100ms and is related to the installation of two flow-based rules (for uplink and downlink UE's communication) in the OpenFlow switch. The *slice creator / selector* then instantiates a new slice for UE (Fig. 3b:4, 600ms), moving the UE to it afterwards (Fig. 3b:4, 2). In this context, the whole process takes almost 3s to be completed.

v. Scenario messages impact

Fig. 4e) illustrates the message sizes for the different scenario messages. For the UE attachment and detachment (Fig. 3b:1), we used the REST API, more specifically, HTTP POST messages. As such, the former as an impact of 370bytes and the latter 310bytes. When the SDN controller is informed of the UE's attachment, the SDN controller updates S/P-GW-U via OpenFlow messages (i.e., *OF flow_modification*). Thus two messages are sent (Fig. 3b:2): first for uplink (UL) communication (170bytes), and second for downlink (DL) (178bytes). Initially, the UE is attached to a default slice. In case, of specific UE's requirements, the *slice creator/ selector*, uses HTTP POST messages (Fig. 3b:4) to instantiate a slice (558bytes) and to handover the UE for the new slice (66bytes). Finally, for a slice delete and

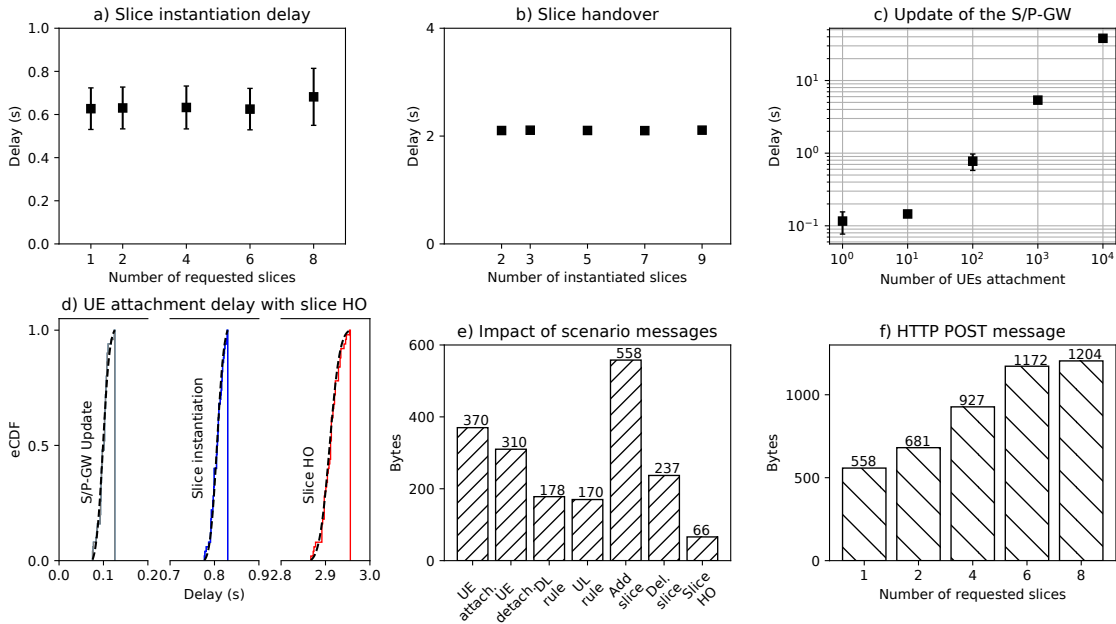


Figure 4: Experimental results for network reconfiguration.

HTTP POST message (237bytes) is sent towards the FlexRAN SD-RAN controller.

Regarding to request slice messages, Fig. 4f) shows how the HTTP POST message increases depending on the number of requested slices. Also, in scenarios where there are already 5 slices, and 3 more slices are needed, the HTTP POST contains the resulted 8 slices with the correspondent slice parameters (resulting in 1204bytes).

VIII. CONCLUSIONS

Network slicing is presented as the game changer for 5G networks. Leveraged by SDN and NFV technologies, network slicing promises to enable the logical partition of the network in slices tailored to vertical' requirements. This paper presented a framework capable of abstracting the partitioning of the network to verticals, while providing the necessary requirements by dynamically reconfiguring the network. Different slice dimensions where proposed and its contribution to verticals and use cases was presented. Additionally, current slicing mechanisms and frame-

works where mapped to the proposed slice types. Finally, a proof-of-concept framework based on OAI and FlexRAN, was deployed and experimentally evaluated in terms of delay and scalability. Results showcased that, despite some of the mechanisms exposed to an increasing amount of slices were able to cope with the experimented scenario's demands, existing technical approached still need to be further enhanced in order to allow not only a larger number of multiple isolated slices, but also to efficiently handover UEs among slices .

ACKNOWLEDGEMENT

This work is funded by FCT/MEC through national funds and when applicable co-funded by FEDER    PT2020 partnership agreement under the project PTDC/EEI-TEL/30685/2017, and by the Integrated Programme of SR&TD    IJSOCA   (Ref. CENTRO-01-0145-FEDER-000010), co-funded by Centro 2020 program, Portugal 2020, European Union, through the European Regional Development Fund. The experimental validation of this work included usage of the FlexRAN software from "Mosaic-

5G", a EURECOM registered trademark (INPI num. 17 4 357 904).

REFERENCES

- [1] Evariste Logota, Daniel Corujo, Seil Jeon, Jonathan Rodriguez, and Rui L Aguiar. The 5g internet. *Fundamentals of 5G Mobile Networks*, pages 29–62, 2015.
- [2] Xenofon Foukas, Georgios Patounas, Ahmed Elmokashfi, and Mahesh K Marina. Network slicing in 5g: Survey and challenges. *IEEE Communications Magazine*, 55(5):94–100, 2017.
- [3] European Telecommunications Standards Institute (ETSI). Network functions virtualisation, 2012.
- [4] Open Network Foundation (ONF). SDN architecture, 2014.
- [5] Nick McKeown, Tom Anderson, Hari Balakrishnan, Guru Parulkar, Larry Peterson, Jennifer Rexford, Scott Shenker, and Jonathan Turner. Openflow: enabling innovation in campus networks. *ACM SIGCOMM Computer Communication Review*, 38(2):69–74, 2008.
- [6] Telefonica’s new open source edge project cuts 5g network slices into strings. <https://www.sdxcentral.com/articles/news/telefonicas-open-source-edge-project-cuts-5g-network-slices-into-strings/> 2018/04/?utm_source=SDxCentral.com+Mailing+List&utm_campaign=61702cf34b-SDxCentral+Newsletter+2018-04-27&utm_medium=email&utm_term=0_c2b6e504a2-61702cf34b-82025313. Accessed: 2018-10-06.
- [7] How many network slices are needed? <https://cloudblog.ericsson.com/digital-services/how-many-network-slices-are-needed>. Accessed: 2018-10-06.
- [8] Matias Richart, Javier Baliosian, Joan Serfat, and Juan-Luis Gorricho. Resource slicing in virtual wireless networks: A survey. *IEEE Transactions on Network and Service Management*, 13(3):462–476, 2016.
- [9] Qi Wang, Jose Alcaraz-Calero, Maria Barros Weiss, Anastasius Gavras, Pedro Miguel Neves, Rui Cale, Giacomo Bernini, Gino Carrozzo, Nicola Ciulli, Giuseppe Celozzi, et al. Slicenet: End-to-end cognitive network slicing and slice management framework in virtualised multi-domain, multi-tenant 5g networks. In *2018 IEEE International Symposium on Broadband Multimedia Systems and Broadcasting (BMSB)*, pages 1–5. IEEE, 2018.
- [10] Antonio de la Oliva, Xi Li, Xavier Costa-Perez, Carlos J Bernardos, Philippe Bertin, Paola Iovanna, Thomas Deiss, Josep Mangues, Alain Mourad, Claudio Casetti, et al. 5g-transformer: Slicing and orchestrating transport networks for industry verticals. *IEEE Communications Magazine*, 56(8):78–84, 2018.
- [11] Yasir Zaki, Liang Zhao, Carmelita Goerg, and Andreas Timm-Giel. Lte wireless virtualization and spectrum management. In *Wireless and Mobile Networking Conference (WMNC), 2010 Third Joint IFIP*, pages 1–6. IEEE, 2010.
- [12] Jonathan Van De Belt, Hamed Ahmadi, and Linda E Doyle. A dynamic embedding algorithm for wireless network virtualization. In *Vehicular Technology Conference (VTC Fall), 2014 IEEE 80th*, pages 1–6. IEEE, 2014.
- [13] Mahmoud I Kamel, Long Bao Le, and André Girard. Lte wireless network virtualization: Dynamic slicing via flexible scheduling. In *Vehicular Technology Conference (VTC Fall), 2014 IEEE 80th*, pages 1–5. IEEE, 2014.
- [14] Jordi Pérez-Romero, Oriol Sallent, Ramon Ferrús, and Ramón Agustí. On the

-
- configuration of radio resource management in a sliced ran. In *NOMS 2018-2018 IEEE/IFIP Network Operations and Management Symposium*, pages 1–6. IEEE, 2018.
- [15] Oriol Sallent, Jordi Perez-Romero, Ramon Ferrus, and Ramon Agusti. On radio access network slicing from a radio resource management perspective. *IEEE Wireless Communications*, 24(5):166–174, 2017.
- [16] Kiyohide Nakauchi, Yozo Shoji, and Nozomu Nishinaga. Airtime-based resource control in wireless lans for wireless network virtualization. In *Ubiquitous and Future Networks (ICUFN), 2012 Fourth International Conference on*, pages 166–169. IEEE, 2012.
- [17] Mahsa Derakhshani, Xiaowei Wang, Tho Le-Ngoc, and Alberto Leon-Garcia. Virtualization of multi-cell 802.11 networks: Association and airtime control. *arXiv preprint arXiv:1508.03554*, 2015.
- [18] Katherine Guo, Shruti Sanadhya, and Thomas Woo. Vifi: virtualizing wlan using commodity hardware. *ACM SIGMOBILE Mobile Computing and Communications Review*, 18(3):41–48, 2015.
- [19] Katerina Koutlia, Anna Umbert, Sergio Garcia, and Ferran Casadevall. Ran slicing for multi-tenancy support in a wlan scenario. In *Network Softwarization (NetSoft), 2017 IEEE Conference on*, pages 1–2. IEEE, 2017.
- [20] Gregory Smith, Anmol Chaturvedi, Arunesh Mishra, and Suman Banerjee. Wireless virtualization on commodity 802.11 hardware. In *Proceedings of the second ACM international workshop on Wireless network testbeds, experimental evaluation and characterization*, pages 75–82. ACM, 2007.
- [21] Lei Xia, Sanjay Kumar, Xue Yang, Praveen Gopalakrishnan, York Liu, Sebastian Schoenberg, and Xingang Guo. Virtual wifi: bring virtualization from wired to wireless. In *Acm sigplan notices*, volume 46, pages 181–192. ACM, 2011.
- [22] Rajesh Mahindra, GD Bhanage, George Hadjichristofi, Ivan Seskar, Dipankar Raychaudhuri, and YY Zhang. Space versus time separation for wireless virtualization on an indoor grid. In *Next Generation Internet Networks, 2008. NGI 2008*, pages 215–222. IEEE, 2008.
- [23] Flavio Meneses, Daniel Corujo, Augusto Neto, and Rui L. Aguiar. Sdn-based end-to-end flow control in mobile slice environments. In *Workshop on Mobility Support in Slice-based network control for heterogeneous environments (MOBISLICE) (IEEE NFV-SDN workshops)*. IEEE, 2018.
- [24] Aditya Gudipati, Daniel Perry, Li Erran Li, and Sachin Katti. Softran: Software defined radio access network. In *Proceedings of the second ACM SIGCOMM workshop on Hot topics in software defined networking*, pages 25–30. ACM, 2013.
- [25] Ian F Akyildiz, Pu Wang, and Shih-Chun Lin. Softair: A software defined networking architecture for 5g wireless systems. *Computer Networks*, 85:1–18, 2015.
- [26] Navid Nikaein, Raymond Knopp, Florian Kaltenberger, Lionel Gauthier, Christian Bonnet, Dominique Nussbaum, and Riadh Ghaddab. Openairinterface: an open lte network in a pc. In *Proceedings of the 20th annual international conference on Mobile computing and networking*, pages 305–308. ACM, 2014.
- [27] Navid Nikaein, Chia-Yu Chang, and Konstantinos Alexandris. Mosaic5g: agile and flexible service platforms for 5g research. *ACM SIGCOMM Computer Communication Review*, 48(3):29–34, 2018.
- [28] Anta Huang, Navid Nikaein, Tore Stenbock, Adlen Ksentini, and Christian Bonnet. Low latency mec framework for sdn-based lte/lte-a networks. In *Communications (ICC), 2017 IEEE International Conference on*, pages 1–6. IEEE, 2017.

-
- [29] Xenofon Foukas, Navid Nikaein, Mohamed M Kassem, Mahesh K Marina, and Kimon Kontovasilis. Flexran: A flexible and programmable platform for software-defined radio access networks. In *Proceedings of the 12th International on Conference on emerging Networking EXperiments and Technologies*, pages 427–441. ACM, 2016.
- [30] A. de la Oliva, X. Li, X. Costa-Perez, C. J. Bernardos, P. Bertin, P. Iovanna, T. Deiss, J. Mangues, A. Mourad, C. Casetti, J. E. Gonzalez, and A. Azcorra. 5g-transformer: Slicing and orchestrating transport networks for industry verticals. *IEEE Communications Magazine*, 56(8):78–84, August 2018.
- [31] Q. Wang, J. Alcaraz-Calero, M. B. Weiss, A. Gavras, P. M. Neves, R. Cale, G. Bernini, G. Carrozzo, N. Ciulli, G. Celozzi, A. Ciriaco, A. Levin, D. Lorenz, K. Barabash, N. Nikaein, S. Spadaro, D. Morris, J. Chochliouros, Y. Agapiou, C. Patachia, M. Iordache, E. Oproiu, C. Lomba, A. C. Aleixo, A. Ro-Drigues, G. Hallissey, Z. Bozakov, K. Koutsopoulos, and P. Walsh. Slicenet: End-to-end cognitive network slicing and slice management framework in virtualised multi-domain, multi-tenant 5g networks. In *2018 IEEE International Symposium on Broadband Multimedia Systems and Broadcasting (BMSB)*, pages 1–5, June 2018.
- [32] Flavio Meneses, Rui Silva, David Santos, Daniel Corujo, and Rui L. Aguiar. Using sdn and slicing for data offloading over heterogeneous networks supporting non-3gpp access. In *29th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*. IEEE, 2018.
- [33] Shamik Sarkar, Christopher Becker, Josh Kunz, Aarushi Sarbhai, Gurupragaash Anasamymani, Sneha Kumar Kasera, and Jacobus Van der Merwe. Enabling wifi in open access networks. In *Proceedings of the 4th ACM Workshop on Hot Topics in Wireless*, pages 13–17. ACM, 2017.
- [34] FlexRAN. <http://mosaic-5g.io/flexran/>. Accessed: 2018-09-30.